University of Liège
Faculty of Sciences
Department of Astrophysics, Geophysics and Oceanography

Determining astrophysical parameters of quasars within the Gaia mission

Author:
Ludovic Delchambre

Supervisor:
Prof. Jean Surdej

A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy in Space Science
in the
Space sciences, Technologies and Astrophysics Research group
Department of Astrophysics, Geophysics and Oceanography

December 2017

Members of the jury:

Prof. Serge Habraken	President, University of Liège
Dr. Yassine Damerdji	Centre de Recherche en Astronomie, Astrophysique et Géophysique, Bouzereah
Dr. Yves Frémat	Royal Observatory of Belgium
Dr. Pierre Geurts	University of Liège
Dr. Éric Gosset	Co-promotor, University of Liège
Prof. Jean Surdej	Supervisor, University of Liège

"Only two things are infinite, the universe and human stupidity, and I'm not sure about the former."

Albert Einstein, 1936

Thesis abstract

Quasars are one of the most peculiar types of objects in astronomy. The supermassive black hole these harbour effectively makes the surrounding matter radiates an enormous amount of energy before getting in the vicinity of the black hole horizon out of which it will never escape. This ironically leads to the most luminous phenomenon in the Universe while being non-transient. It is hence quite natural to rely on these cosmic headlights, visible up to ages when the Universe was still very young, so as to achieve some of the currently most important cosmological applications, notably regarding the determination of the cosmological parameters H_0, Ω_Λ and Ω_m. The Gaia mission, on its side, is one of a kind given the one billion of celestial objects it is intended to observe, among which more than half a million quasars are expected. Furthermore, owing to its exceptional astrometric precision, Gaia stands out to be extremely well suited for the detection of gravitational lens (GL) systems. In the latter, light rays coming from a distant background quasar are deflected by the presence of a massive galaxy being in the line-of-sight that leads to the production of multiple images of this background quasar upon a favourable alignment between the quasar, the galaxy and the observer. Supplemental constraints on the aforementioned cosmological parameters being then gained based on these GLs. Gaia hence provides an unprecedented opportunity to detect and characterize quasars as well as to identify GLs which ultimately bring a better understanding of the Universe we live in. This thesis is accordingly concerned with the development of software solutions dedicated to the determination of the astrophysical parameters (APs) of the quasars that Gaia will observe, on one hand, and to the recognition of the GLs among the billion of sources it will uncover, on the other hand.

Although Gaia provides state-of-the-art astrometric and photometric observations, its capability in characterizing these celestial objects remains however restricted by the relatively low spectral resolution of the blue and red spectrophotometers upon which it is based as well as by the limited signal-to-noise ratio that is associated with faint objects, including quasars. In addition, the overwhelming amount of data that Gaia has to process translates into a stringent need for algorithms having both low numerical complexities as well as low memory usages. These restrictions and shortcomings along with the requirement for reliable APs were at the heart of this research that led to the development of two specifically designed methods that are the weighted principal components analysis and the weighted phase correlation method. The former of these methods allowed us to extract the most significant patterns out of quasars with a view of using these in the production of a spectral library of quasars as observed by Gaia. These were subsequently used in a fast and automated procedure designed to guess the redshift of the quasars within the Gaia mission through the latter mentioned method. Other APs that are the slope of the quasar continua, the total equivalent width of their emission lines and whether these encompass broad

absorption lines or not, being then concurrently derived based on the results of these methods.

Finally, the identification of GL candidates relies on the recognition of the structures and symmetries that are observed within lensed images through supervised learning methods. The specific method we choose to use, based on extremely randomized trees, was shown to yield a low contamination rate on simulated configurations composed of three images as well as a very high probability of detection in cases of four image configurations. Real observations out of the first Gaia data release were processed and resulted in the identification of candidates having three potentially lensed images which are currently waiting for confirmation using ground-based facilities.

Résumé

Les quasars constituent l'un des types d'objets les plus singuliers en astronomie. Le trou noir supermassif qu'ils abritent fait en effet rayonner une énorme quantité d'énergie à la matière environnante avant que celle-ci ne s'approche de l'horizon du trou noir, dont elle ne pourra s'échapper. Ceci produit ironiquement le phénomène le plus lumineux de l'Univers qui ne soit pas transitoire. Il est donc assez naturel d'utiliser ces phares cosmiques, qui sont visibles jusqu'à des époques où l'Univers était encore très jeune, afin d'effectuer certaines des applications cosmologiques les plus importantes de l'époque actuelle, notamment concernant la détermination des paramètres cosmologiques H_0, Ω_Λ et Ω_m. La mission Gaia, de son côté, est unique en son genre étant donné le milliard d'objets célestes qu'elle observera, parmi lesquels plus d'un demi million de quasars sont attendus. De plus, grâce à sa précision astrométrique exceptionnelle, Gaia se trouve être particulièrement bien adapté à la détection de lentilles gravitationnelles (LGs). Dans ces dernières, des rayons lumineux provenant d'un quasar d'arrière plan distant sont défléchis par la présence d'une galaxie massive se trouvant le long de la ligne de visée, ce qui mène à la production d'images multiples de ce quasar pour un alignement favorable entre le quasar, la galaxie et l'observateur. Des contraintes supplémentaires sur les paramètres cosmologiques pré-cités étant alors obtenus sur base de ces LGs. Gaia fournit dès lors une opportunité sans précédent afin de détecter et de caractériser les quasars ainsi que pour identifier des LGs, ce qui mènera, en fin de compte, à une meilleure compréhension de l'Univers dans lequel nous vivons. Cette thèse se concentre donc sur le développement de solutions logicielles dédiées, d'une part, à la détermination des paramètres astrophysiques (PAs) des quasars que Gaia observera et à la reconnaissance des LGs au sein du milliard de sources qu'il découvrira d'autre part.

Bien que Gaia fournisse des observations astrométriques et photométriques de pointe, sa capacité à caractériser ces objets célestes reste cependant limitée par la résolution spectrale relativement faible des spectrophotomètres bleu et rouge sur lesquels elle repose ainsi que par le rapport signal sur bruit limité qui est associé aux objets faibles, comprenant les quasars. Par ailleurs, la quantité importante de données que Gaia doit traiter se traduit par la rigoureuse nécessité d'avoir des algorithmes ayant à la fois de faibles complexités numériques et de faibles utilisations mémoire. Ces restrictions et inconvénients, ainsi que l'exigence d'avoir des PAs fiables étaient au cœur de cette recherche qui a conduit au développement de deux méthodes spécifiques que sont l'analyse en composantes principales pondérées et la corrélation de phase pondérée. La première de ces méthodes nous a permis d'extraire les modèles les plus significatifs des quasars en vue de leur utilisation pour la production d'une bibliothèque spectrale de quasars tels qu'ils seront observés par Gaia. Ceux-ci ont ensuite été utilisés dans une procédure rapide et automatisée conçue afin d'estimer le redshift des quasars au sein de la mission Gaia au travers de la dernière méthode mentionnée. Les autres PAs

que sont la pente du continu des quasars, la largeur totale de leur raies en émission et si ceux-ci présentent des raies d'absorption larges sont ensuite dérivées sur base des résultats précédents.

Enfin, l'identification des LGs repose sur la reconnaissance des structures et des symétries qui sont observées au sein des images lentillées via des méthodes d'apprentissage supervisé. La méthode spécifique que nous avons choisi d'utiliser, basée sur les arbres extrêmement randomisés, a produit un faible taux de contamination pour les configurations composées de trois images ainsi qu'une très forte probabilité de détection dans le cas des configurations à quatre images. Des observations réelles issues de la première publication des données Gaia ont été traitées et ont permis d'identifier des candidats à trois images potentiellement lentillées qui sont actuellement en attente de confirmation à l'aide d'instruments au sol.

Dedicated to my son, Arthur, and to his little sister Clémentine.

Contents

Abstract **v**

Résumé **vii**

Contents **xi**

1 Introduction **1**
1.1 The Gaia mission . . . 2
1.1.1 Scientific objectives . . . 2
1.1.2 Overall design . . . 4
1.1.3 The focal plane assembly . . . 5
1.1.4 Gaia data processing and analysis . . . 10
1.2 Active galactic nuclei and quasars . . . 13
1.2.1 Historical background . . . 13
1.2.2 A multi-wavelength view of AGN . . . 14
1.2.3 A unified model of AGN . . . 26
1.2.4 Cosmological applications . . . 34
1.3 Strong gravitational lensing . . . 40
1.3.1 Introduction to gravitational lens theory . . . 41
1.3.2 The non-singular isothermal ellipsoid in presence of an external shear . . . 45
1.3.3 Applications . . . 48

2 Methods for the characterization of quasars **51**
2.1 The quasar classifier work package . . . 51
2.1.1 Objectives and software requirements . . . 52
2.1.2 Envisaged solution and associated issues . . . 57
2.2 Weighted principal component analysis . . . 61
2.2.1 Paper I . . . 63
2.3 Weighted phase correlation . . . 74
2.3.1 Paper II . . . 75
2.4 Extremely randomized trees . . . 86

3 Determination of the astrophysical parameters of quasars **91**
3.1 Paper III . . . 93

4 Gravitational lenses identification **109**

5 Conclusions **117**

A Numerical complexities **123**
A.1 Weighted least-squares solution to a system of linear equations 123
A.2 Fast Fourier transform 125
A.3 WPC through factorized QR decomposition and lookup tables 126

List of Figures **129**

List of Tables **131**

List of Abbreviations **133**

List of symbols and notations **137**

List of constants and units **139**

1 Introduction

Since the discovery of the two first quasars, 3C273 and 3C48, in the mid-1960s (Schmidt 1963; Greenstein & Schmidt 1964), their interest has grown steadily along with their applications. Indeed, from their extreme luminosities ($L > 10^9 L_\odot$) and large redshift ($0.1 < z < 7$), these are privileged tracers of the history, structure and composition of the Universe. They also play a key-role in fixing the celestial reference frames and constitute ideal candidates for the formation of gravitational lens system upon which some of the most important cosmological parameters can be straightly derived (Refsdal 1964; Suyu et al. 2013).

At the present time, the advent of very large surveys like the Sloan Digital Sky Survey (Alam et al. 2015, hereafter SDSS) has allowed to pass from the few thousands of quasars known in the 90s (see Véron-Cetty & Véron 2010, Table 1) to $\sim 3 \times 10^5$ quasars in most recent catalogues like in Pâris et al. (2017). This large amount of quasars yields the possibility to carry out statistically significant studies on the intrinsic properties of these peculiar objects, on their environment and on their related cosmological applications. The Gaia satellite, launched in December 2013, will also greatly contribute to this number with $5{\times}10^5$ quasars that are expected to be detected.

As we can notice, the study of quasars already entered in the 'big data' era where the human expertise, although invaluable, is rendered impractical because of time issues. Rather, like in almost all fields of astronomy, skilful astronomers have to be supplemented by smart computer algorithms to get their data efficiently and quickly reduced. The detection and characterization of quasars within the Gaia mission similarly rely on such algorithms with the supplemental constraint for these to come along with particularly low execution times as well as very limited memory consumptions. These additional constraints arise from the one billion of objects that Gaia will observe and will have to subsequently process. As an illustrative purpose, let us assume that the time needed for one object to go through the whole Gaia processing chain is one second (including, amongst other things, the pre-processing time, astrometric, photometric and spectroscopic reduction, classification, characterization, variability analysis, ...). It would then take about 32 years of computation in order for the whole set of observations to be reduced on a single CPU. This problematic is explored in the present work along with the identification and determination of the parameters of the gravitational lenses (GL) in the field of the Gaia mission.

The first chapter describes the Gaia mission from a general point of view, quasars

and their applications as well as the gravitational lensing phenomenon. It does not aim at being exhaustive but rather provides the reader with the minimal insights that are required in order for this work to be fully understood. Interested readers willing to go one step further are invited to read Gaia Collaboration et al. (2016) and references therein for a more complete description of the Gaia mission, Beckmann & Shrader (2012) or Peterson (1997) for a comprehensive review of Active Galactic Nuclei and quasars and Schneider (1992) for a rigorous explanation of the gravitational lensing phenomenon.

In the following, Chapter 2 contains the methods that we specifically developed in order to determine the astrophysical parameters (APs) of quasars within the Gaia mission as well as a supervised learning method we used in order to identify and characterize gravitational lens systems (GLs). The precise procedure we implemented for gathering these APs are described within Chapter 3 while the GLs identification and characterization are covered in Chapter 4.

1.1 The Gaia mission

Gaia is one of the cornerstone space mission of the Horizon 2000+ science program of the European Space Agency (ESA) that aims to bring a consensus on the history and evolution of our Galaxy through the survey of a billion celestial objects (Perryman et al. 2001). Initially thought as an interferometric mission (Lindegren & Perryman 1996), it was later re-designed such as to allow fainter objects to be detected. From this early period remains the original mission name: 'GAIA'[1], standing for 'Global Astrometric Interferometer for Astrophysics' though it was no longer an interferometric mission. The satellite was launched from French Guyana on 19 December 2013 at 09:12:19.6 UTC by a Soyuz-STB launcher for a nominal mission of five years around the Lagrange L2 point and cost about one billion euros.

This section briefly describes the mission objectives, the overall satellite design, while particularly focusing on its focal plane, and finally details the computational structures that are put in place so as to process and analyze the huge amount of data coming from the satellite.

1.1.1 Scientific objectives

The composition, formation and evolution of our Galaxy

The objectives of Gaia are multiple but the fundamental scientific goal is the understanding of the formation and of the evolution of our Galaxy: the Milky Way. This objective is fulfilled through the repeated measurements of the astrometric positions and of the photometric fluxes of a billion celestial objects as well as from the spectroscopic measurements of a subset of these objects. These will allow to determine the parallaxes, proper motions, radial velocities and chemical compositions of a significant fraction of the stars composing our galaxy. This snapshot being then used in order to trace the evolution of the Milky Way back to its formation.

By doing a full sky coverage, Gaia will also be able to measure the interstellar extinction in each direction it points to and to create a map of the interstellar medium in the Milky Way. From the knowledge of its structure, kinematics and mass distribution, Gaia will allow to answer one of the most intriguing questions of the present days, that is the distribution of the dark matter in our Galaxy.

[1] The original mission name, 'GAIA', was then set to lower cases in order to become 'Gaia'.

Finally, the combination of the absolute magnitudes of stars with their observed metallicities will enable their individual ages to be determined such as to have a spatial distribution of the latter amongst the Galaxy, to subsequently deduce their birth places and tell the history of its star formation. These observations being then complemented and assessed with those of other numerous distant galaxies.

Stellar physics, evolution and distances

The main strength of Gaia is the huge number of stars it will observe and the thorough statistical analysis that can be conducted based on these. The $\sim 10^8$ stellar parallaxes that will be obtained by Gaia along with its similar number of accurate photometric records will accordingly allow to derive high-quality colour-magnitude diagrams (i.e. Hertzsprung-Russell diagrams, hereafter HR diagrams) based on which precise researches can be undertaken regarding stellar physics, evolution and distances. Indeed, these numerous observations will enable the great diversity of stars to be studied in all their evolution phases, from the pre-main sequence to their final fate, from the most massive O stars down to brown dwarfs. This is even more crucial given that we know some of these evolutionary phases to be extremely fast such as the helium core flash in solar-type stars. Once precisely re-calibrated, the main sequence of the HR diagram will then enable distances from stars to be finely determined even in the case where no trigonometric parallax is available.

Stellar variability, supernovae and the cosmic-distance ladder

The fact that each object will be observed 70 times on average will also allow the detection of $\sim 10^7$ variable stars. Beyond the obvious interest of variability in understanding the associated stellar physics, the parallaxes of some $80,000$ pulsating stars (RR Lyrae and Cepheids) will be used to improve their period-luminosity function and hence their distances. Similarly, this systematic variability survey will be used in order to discover many thousands of supernovae, amongst which type Ia supernovae which can be equivalently used as a fair indicator of their distances. In the end, Gaia will allow a full and independent re-calibration of the whole cosmic-distance ladder through parallaxes, HR diagram, Cepheids & RR Lyrae stars and type Ia supernovae, ...

Stellar multiplicity and exoplanets

The extremely high spatial resolution of Gaia combined with variability analysis and high resolution spectroscopy will enable the detection of many multiple systems of stars either by resolving some of their components, by noting some astrometric wobbles superimposed to their parallaxes and proper motions, from their apparent periodic changes in luminosity in the case of eclipsing stars, from microlensing events or from the back and forth motion of their emission/absorption line(s) in spectroscopy. Similarly, these techniques can easily detect exoplanets of mass equal or higher than that of Jupiter. The collected sample will then be used by other surveys as a basis for the search of habitable earth-size exoplanets protected by a giant planet standing further out.

Asteroids

Although the main targets of Gaia are stars, any point-like sources with relatively low proper motions will be detected, amongst which 10^5 to 10^6 asteroids. These

are unaltered witnesses of the origin and formation of our solar system. The study of their size, mass, shape, albedo and of their taxonomy reveals a vital importance in understanding the formation of planetary systems in general. Gaia was further shown to be able to measure the ephemeris of known asteroids with an unprecedented accuracy while his full sky coverage will enable the discovery of many asteroids in regions of the sky that are usually not covered for their search like those standing far away from the ecliptic plane.

The Local Group

The resolving power offered by Gaia is sufficient for the brightest stars from our local group to be observed. As an example, thousands of stars should be detected within the Andromeda galaxy while a few millions of these should be detected within the Large Magellanic Cloud. In parallel to what is done from stars to stars regarding the Milky Way, the goal would be to shed light on the interactions existing amongst the members of the Local Group through the computation of the three dimensional positions and velocities of the stars composing these (dwarf) galaxies.

Unresolved extragalactic objects and the reference frame

Besides the intrinsic interest of the unresolved extragalactic sources that is further explored in section 1.2.4, these objects have a particular importance in fixing the Gaia Celestial Reference Frame (GCRF) which aims to be the extension of the International Celestial Reference Frame (ICRF) from radio interferometry to visible light. The apparent proper motions of quasars will then principally come from the angular displacement of the Sun around the Galactic center. Another interesting application stands in the comparison of the positions of quasars from the ICRF to their corresponding positions in the GCRF such as to check whether these coincide in both visual and radio wavelengths and thus correspond to similar physical phenomena.

Fundamental physics

From its extremely high spatial resolution, Gaia will have to routinely deal with the correction of relativistic effects like the light bending occurring in the vicinity of strong gravitational fields. Again, the large set of measurements will allow to quantify any bias that would come from the imperfections of the relativistic model we used by simply looking at its residuals. Another experiment would consist in the measurement of the bending of light of stars being close to the limb of Jupiter in order to assess these models.

1.1.2 Overall design

Gaia is mainly composed of two telescopes pointing in directions separated by a basic angle of $\Gamma = 106.5°$. These two telescopes consist each of six mirrors numbered from M1 to M6 (resp. from M1' to M6') and share a common focal plane composed of a patchwork of 106 CCDs. The data acquisition is then carried out by rotating the satellite on its spin axis at an angular velocity of $60''\text{s}^{-1}$ while reading each CCD column at the same rate as the objects cross the focal plane, in the so-called Time Delay Integration (TDI) mode. The validity of this design was already successfully demonstrated by ESA as it is inherited from its predecessor, the HIPPARCOS mission(Perryman et al. 1989). The fact that the objects coming from each field of view (FOV) are mapped onto the same focal plane allows their relative positions to be

determined in an extremely precise manner though these have large angular separations. The combination of many such relationships along with the large focal length of both telescopes and the extremely good sampling of the CCD in their along-scan (AL) direction, which is comparable to the one of the Hubble space telescope, then allow to reach the astrometric precision claimed by Gaia.

From simple considerations we can show that the sensitivity of Gaia regarding the absolute parallaxes in the AL direction is proportional to $\sin\xi \sin\Gamma$, where ξ is the angle between the spin axis of the satellite and the sun (Gaia Collaboration et al. 2016). Intuitively, stars coming from a single FOV will have parallax displacements (i.e. parallax factors) that will be oriented in very similar directions (i.e. towards the Sun). The determination of their absolute parallaxes will hence have to rely on their sole relative position along the (approximately unique) great circle joining these stars to the Sun. A more interesting solution would then be to calibrate these absolute parallaxes with stars having their parallax factors oriented in very different directions, that is at $\Gamma = 90°$ in an ideal case. Nevertheless for the sake of the optimization of the accuracy in the AL direction during short period of time (Makarov 1998), acceptable values were restricted to stand either within $99.4° \pm 0.1°$ or within $106.5° \pm 0.1°$. Manufacturing constraints then favoured the latter solution. Also, given that we would like to optimize the detection of the parallax factors in the AL direction, we would then prefer to have these apparent displacements to be always oriented in this specific direction, which is the case if $\xi = 90°$. However in this configuration the sun light would enter the aperture of the satellite instead of alimenting the solar panels and the thermal stability would not be further ensured. A constant value of $\xi = 45°$ was then chosen as a trade-off between the efficiency and the permanent protection of the satellite by its sun shield and solar panels. The angular velocity of the spin axis ($1\ ''s^{-1}$) was then chosen as a complex compromise between the observational frequency of the sources, their reachable magnitudes, their signal-to-noise ratio (SNR), the limited telemetry of the satellite and the blurring of the point spread function (PSF) during the TDI integration. At last, a slow precession of the spin axis allows to spread the sampling of the variability curves of the observed sources in a more uniform way. The chosen precession period of 63 days still allows an overlapping of the observations over a full revolution of the spin axis.

The two telescopes composing Gaia principally encompass the primary mirrors (M1 and M1'), each of size 1.45×0.5 m and secondary mirrors (M2 and M2'), each having five degrees of freedom such as to ensure a correct alignment of the whole optical bench and a precise focusing of both telescopes. The corrections applied at the M2/M2' mirrors account for the irregularities introduced by the launch vibrations as well as by the routine monitoring of the focus through the Wave-front Sensor system (hereafter WFS, see Section 1.1.3). The fourth mirrors (M4 and M4') are flat folding mirrors designed to gather the beams of light coming from both telescopes onto a common optical path consisting of the M5 and M6 mirrors. The M5' and M6' mirrors being actually alternative names for M5 and M6, respectively. This complex arrangement eventually yield a (required) focal length of 35 m for each of the telescopes. A sketch of the disposition of the mirrors onto the optical bench of Gaia is provided in figure 1.1.

1.1.3 The focal plane assembly

Gaia aims to produce the most complete astrometric and photometric catalogue nowadays. In order to accomplish this tremendous objective, its CCDs are required to be accurate enough so as to produce reliable measurements, to be sensitive to the faintest

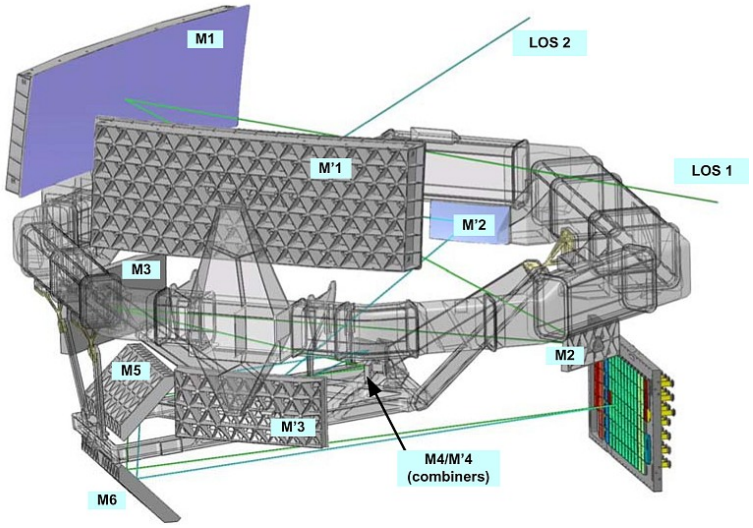

Figure 1.1: The hexagonal optical bench supporting the ten mirrors composing the two telescopes of Gaia. Each telescope is composed of six reflectors numbered from M1 to M6 (M1' to M6', respectively, for the second telescope), two being shared amongst both telescopes (M5 and M6) while the beam recombination occurs at the level of the M4/M4' mirrors. This configuration allows simultaneous observations along two line-of-sights (LOS1 and LOS2), separated by a basic angle of 106.5° along with an effective focal length of 35 m. Image credit: Airbus Defence and Space.

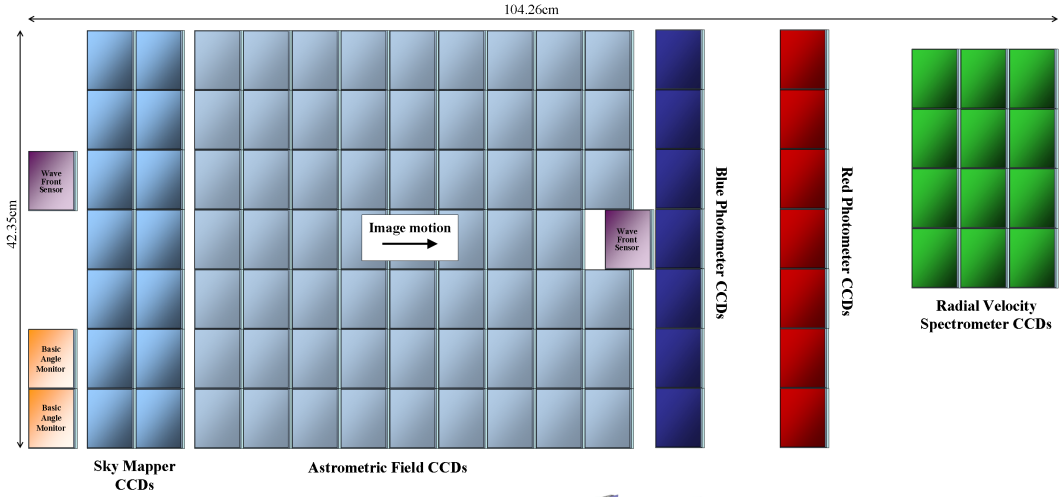

Figure 1.2: The Gaia focal plane composed of 106 CCDs and their associated instrumentations. Image from ESA – A.D. Short, July 2009.

sources as well as to come along with a high dynamic range that allows the flux from the brightest sources to be recorded as well, therefore producing an unbiased, complete catalogue down to a G magnitude of 20. The focal plane is accordingly equipped with 106 CCDs that were specifically designed in order to fulfil the previously mentioned requirements and that form the largest focal plane ever sent in space as illustrated in figure 1.2.

Each of these constituting CCDs consist in a grid of 4500×1966 pixels having a size of 45×60 mm along with an individual pixel size of 10×30 μm (58.9×176.8 mas on the sky). While their accuracy and their sensitivity are mainly and respectively driven by well-studied characteristics of CCDs and by the angular velocity of the spin axis of the satellite, the novelties introduced in the latter allow to outperform by far the dynamic of 'classical' CCDs. Indeed, from its design, Gaia will be able to have some insights on the flux that will be received from each source (see 'Sky mappers' below) and to determine whether these sources will saturate the potential well of the pixels or not. Based on this knowledge, some electrodes standing at the junction of predefined CCDs columns will be activated such as to flush the photo-electrons that were collected during the TDI integration of the previous columns. These bypasses, called TDI gates, then allow the flux coming from each object to be integrated through a suited number of CCD columns and hence to have a well adapted exposure time. Twelve of these TDI gates are present on each CCD composing the focal plane. These stand at CCD column: 2900, 2048, 1024, 512, 256, 128, 64, 32, 16, 8, 4 and 2. This improvement constitutes one of the major strengths of Gaia for its journey through the exploration of one billion stars.

In the following, we describe the instrumentation that is associated with each set of CCDs (see figure 1.2) which will finally provide us with what we though to be a satisfactory review of the Gaia satellite and of its payload.

Wave-front sensors

Each of the telescopes constituting Gaia is equipped with a Wave-front Sensor (WFS) of Shack-Hartmann type composed of an array of 3×11 lenses. This WFS is used in order to monitor the focus and the alignment of the telescopes while actuating the M2/M2' mirrors along their five degrees of freedom in case of defects. To this aim, the spots obtained from the lens focusing of bright stars are compared to those of calibration spots and spots from other stars that were recorded once the best focusing was achieved. This information along with the spot locations then allow the incident wave-front to be reconstructed and further corrected.

Basic angle monitor

As already pointed out, Gaia heavily relies on the accurate measurements of sources standing at high angular distances through its two FOV separated by an angle of $\Gamma = 106.5°$. The sought astrometric precision then requiring this angle to be monitored at the μas level. To this aim, the Basic Angle Monitor (BAM) makes use of two optical benches fed by a common laser source that is split into two parallel beams of collimated light for each telescope that, once re-combined, will produce Young type interference patterns based on which the basic angle, Γ, can be monitored with an accuracy of 0.5 μas. The latter being actuated every 23 s.

Sky mappers

Sky mappers (SM) are the first set of CCDs that are encountered by the scientific sources (i.e. that are not related to the metrology of the satellite). These are composed of two strips of 7 CCDs (SM1 and SM2), each being exposed to the light coming from a single FOV thanks to physical masks placed at the level of the M4/M4' mirrors. Their objectives are first to identify the (point-like) sources that will cross the focal plane and to match each of these to their corresponding FOV and secondly to compute their rough apparent magnitude such as to decide whether or not we should activate some TDI gates in the subsequent CCDs. These are the sole CCDs that are read in full by the TDI mode though a binning of 2×2 is applied. Indeed, the informations gathered from the SM are used in order to produce a windowed reading of the rightmost 'scientific' CCDs, that is: only pixels being in the vicinity of the detected source will be read such as to spare the readout time. The size of the allocated windows then depends on the instrumentation we use such as described in the remaining of this section.

Astrometric field

The astrometric field (AF) consists of 62 CCDs and is therefore the most widespread instrument of the focal plane. Its purpose is to acquire the photometric measurements of the objects detected by the SM in the unfiltered G band[2] covering the wavelength range from 330 nm to 1050 nm. Depending on the magnitudes derived by the SM, each source will transit the AF in allocated widows of 18×12 pixels in cases where these have magnitudes brighter than 16 mag and in windows of 12×12 pixels for fainter magnitudes. CCD readout will then be carried out by binning the 12 pixels in the across-scan (AC) direction in case of objects fainter than 13 mag while no binning will be performed for brighter sources such as to prevent saturation. A complete transit will hence correspond to 9 such photometric records, one per CCD encountered in the AF, at the exception of the fourth row of the AF where only 8 strips are present. The summing of these records then allows a higher SNR to be gained.

The first strip of CCDs among the AF, AF1, has a particular role since it is used in order to confirm detections from the SM as well as to detect moving objects for which the transition window will be extended. Also it can detect inadequacies in the propagation of the objects through the focal plane and send the appropriate corrections to the attitude and control subsystem, which in-turn will produce the impulsions that are necessary in order for the satellite to keep its stable rotation. These additional functionalities consequently require the observations from AF1 to be read in full resolution (i.e. without binning) and to be later combined by the on-board software.

BP/RP spectrophotometers

Gaia comes along with two low-resolution spectro-photometers, namely the blue photometer (BP), observing in the range 330–680 nm, and the red photometer (RP), observing in the range 640–1050 nm. These measure the spectral energy distribution (SED) of the objects transiting the focal plane such as to: (i) apply the chromatic corrections that would otherwise shift the objects centroids because of the optical aberrations of the telescopes (ii) allow the classification of these objects concurrently with the determination of their APs. More specifically, this last point will, amongst

[2] G standing for Gaia

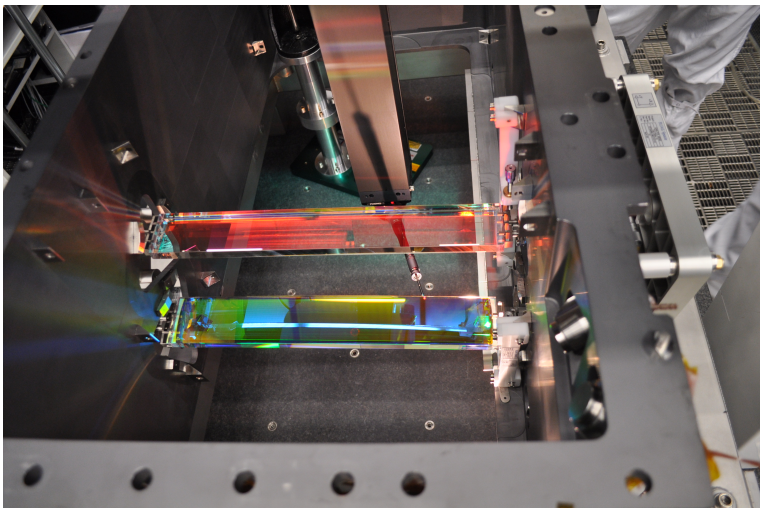

Figure 1.3: Fused-silica prisms of the blue and red spectro-photometers of Gaia. The red photometer (up) disperses the incoming light having wavelengths between 640 nm and 1050 nm with a dispersion curve varying between 7 and 15 nm pixel^{-1} while the blue photometer (down) operates in the range 330–680 nm with a spectral dispersion of 3–27 nm pixel^{-1}. Image credit: Airbus Defence and Space.

other, encompass the characterization of the objects that are getting classified as quasars which is the main topic of the present work.

Both the BP and RP are made of fused-silica prisms which disperse the beams of light over ~ 45 pixels in the AL direction (see figure 1.3). The associated dispersion curves varies from 3 to 27 nm pixel^{-1} for the BP and from 7 to 15 nm pixel^{-1} for the RP. A transition window of 60×12 pixels for both photometers was consequently used in order for the background to be fairly estimated as well and for the errors in the window position and displacement to be taken into account. Similarly to the AF, spectra are binned over the 12 pixels in the AC direction such as to produce one dimensional SED, while spectra having a G magnitude brighter than 11.5 are read in full resolution such as to prevent saturation.

Radial velocity spectrometer

Gaia is further equipped with a high-resolution spectrometer ($R = \lambda/\Delta\lambda \approx 11700$) centred around the Ca II triplet (845–872 nm). Its objectives are manifolds but mainly consist in the determination of the radial velocities of stars for which the magnitude associated with the integrated flux in the concerned band, G_{RVS}, is brighter than 16, hence its name of radial velocity spectrometer (RVS). Some additional objectives being the retrieval of the stellar parameters of objects having $G_{\mathrm{RVS}} < 14.5$ mag or the derivation of the interstellar extinction for stars brighter than $G_{\mathrm{RVS}} \approx 12.5$ mag.

The RVS instrument is an integral field spectrograph consisting in an array of 3×4 CCDs and in an optical module situated between the M6 mirror and the focal plane. This optical module is a complex aggregation of a blazed-transmission grating, of four fused-silica prismatic lenses and of a bandpass filter such as to spread the incoming light over 1100 pixels in the AL direction with a steady wavelength dispersion of 0.0245 nm pixel^{-1} (see figure 1.4). Again RVS spectra were allocated to windows of size 1296×10 pixels, in order to permit background subtraction and errors on windows

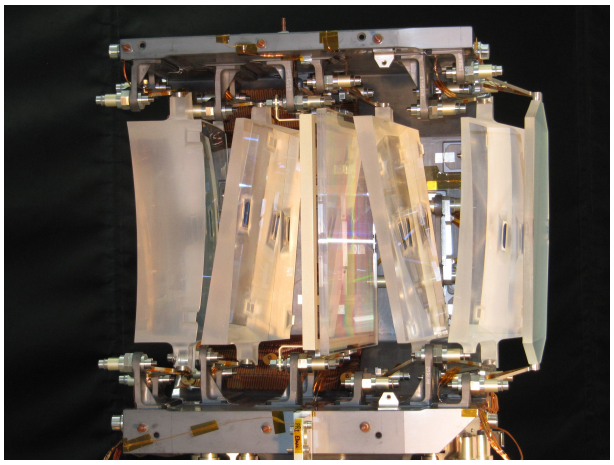

Figure 1.4: The optical module of the radial velocity spectrometers of Gaia. From left to right: two prismatic lenses, a transmission grating plate, two other prismatic lenses and a bandpass filter that together allow high resolution (0.0245 nm pixel^{-1}) spectroscopy over the wavelength range 845–872 nm. Image credit: Airbus Defence and Space.

Table 1.1: Coordination Units from DPAC.

Label	**Name**
CU1	System architecture
CU2	Data simulations
CU3	Core processing
CU4	Object processing
CU5	Photometric processing
CU6	Spectroscopic processing
CU7	Variability processing
CU8	Astrophysical parameters
CU9	Catalogue access and scientific exploration

(dis)placement, and where spectra having $G_{\mathrm{RVS}} > 7$ are binned over the 10 pixels of the AC direction.

1.1.4 Gaia data processing and analysis

The Gaia data processing is entrusted by the Data Processing and Analysis Consortium (DPAC) which is an academic consortium composed of some 450 astronomers distributed over 25 countries, mainly from Europe. This consortium is structured around nine Coordination Units (CUs) which are specialized sub-structures with clearly-defined responsibilities (see Table 1.1). These are coordinated through the DPAC Executive (DPACE) which is a top-level structure serving as an interface between CUs and the Gaia science team, responsible for the final Gaia scientific performances.

Amongst these CUs, CU1 is endorsed with the overall system processing, philosophy, architecture and strategy. It is, amongst other, responsible for the management and design of the data storage that is shared over the whole DPAC, the Main Database (MDB), and of the software architecture based on which all the DPAC work packages

will be able to exchange the results of their calculation (i.e. their software interfaces). CU2 is in charge of the simulation of the Gaia data in a broad sense, that is from its telemetry to its end-of-mission data passing through the simulation of the spacecraft & instrument models, intermediate data products or pixel-level records. The goal of these simulations being then to prepare at best the whole DPAC softwares for their run on real observations. CU3 deals with the astrometric processing of the sources that will be observed, from the pre-processing of their telemetry to the relativistic and colour corrections that must be applied to the latter. CU4 is concerned with the processing of the sources that were flagged within CU3, CU5 or CU6 as having uncommon behaviours these encompass non-single stars (NSS), solar system objects (SSO) and extended objects (EO). CU5 is responsible for the treatment, calibration and validation of the photometric data from SM, AF and BP/RP spectro-photometers, while CU6 has a similar role regarding the RVS spectra. CU7 aims to identify, classify and characterize the photometric and spectroscopic variability of the observed objects. The classification and the determination of the astrophysical parameters of each detected source will be carried out within CU8 while, finally, the content of the final and intermediate catalogues as well as their access will be ensured by CU9. All these efforts being supported by large computing centres, the Data Processing Centres (DPCs), situated in Madrid, Toulouse, Cambridge, Genève, Torino and Barcelona.

The previously described tasks of each CU can be alternatively classified into: *daily* processing regarding the nearly real-time treatment of the Gaia data, *cyclic* processing regarding the iterative process-and-improve data treatment upon which Gaia is based and the data simulation and publication.

Daily processing

The DPAC daily processing consists, on one hand, in the initial data treatment (IDT) which aims to provide a preliminary reduction of the raw scientific data such as the approximated source positions based on a coarse and instantaneous version of the satellite attitude (i.e. the pointing of the telescopes/instruments over the GCRF as a function of time). These approximated positions have a typical accuracy of 100 mas which is sufficient in order for the epoch observations to be cross-matched. Other resources coming from the IDT encompass: preliminary fluxes from AF, BP, RP and RVS instruments, CCD window position and size as well as potential TDI gate activation. On the other hand, the First Look system will use the monitoring informations of the satellite such as to ensure that its whole payload is working as expected through the production of diagnostics logs and plots.

The IDT output is then primarily used as a basis for the science alerts pipeline (http://gsaweb.ast.cam.ac.uk/alerts). Indeed, based on the cross-match performed by the IDT, variability curves can be readily obtained and transient phenomena (e.g. supernovae, cataclysmic variable stars, ...) can be straightly identified and followed-up through ground-based telescopes. If these transients also come along with a fast motion across the focal plane (e.g. 10 mas s^{-1}), they are considered as solar system objects and consequently published as solar system alerts (https://gaiafunsso.imcce.fr/) for a similar follow-up. This distinction is made because such fast moving objects may not get re-observed by Gaia and that their single epoch displacements are insufficient in order to completely reconstruct their orbit, hence necessitating an immediate ground-based follow-up.

Cyclic processing

The self-calibrating nature of Gaia, where a fair calibration of the whole instrumentation is obtained from the observations which in turn rely on a good calibration of the instrumentation, naturally translates into a cyclic processing of its data. This cyclic processing can be viewed as two distinct but complementary parts: *low-level* (or *core*) processing that is straightly concerned with the Gaia instrumentation and *high-level* processing that has to deal with the intrinsic nature of the observed objects. Both processing strongly depend on each other since the low-level processing will be based on products from the high-level processing and conversely. As an example, the control of the attitude of the satellite will be better performed if it is based on fixed points of the celestial reference frame, like quasars, although the identification of quasars is a complex task that has to be carried out separately. Likewise, the precise knowledge of the satellite attitude will improve the precision of the astrometry leading to better predicted positions on the focal plane, to a better modelling of the BP/RP line spread function (LSF) and, finally, to better integrated BP/RP fluxes upon which the quasars identification is based.

Low-level cyclic processing mainly consists in:

- The Intermediate Data Update (IDU), which aims to provide updated images locations and fluxes through the updates of the core astrometric and photometric calibrations (mainly PSF model, CCD bias and dark current estimation) and through the use of the latest attitude and geometric calibrations available. This ultimately leads to an improved cross-matching between the epoch observations. It is worth to mention that this task can be viewed as a more complete version of the IDT while being based on a global approach to the problem of the astrometric and photometric calibration (i.e. using all relevant parameters from the MDB).

- The Astrometric Global Iterative Solution (AGIS), which is the core astrometric solution of Gaia. It consists in a simultaneous least-squares solution to a problem encompassing the astrometric positions of 10^8 well-behaved (primary) sources, the satellite attitude, geometric calibration of the SM and AF as well as other numerous (time-dependent) global calibration parameters with respect to the observed CCD positions and fluxes as received from the satellite.

- The photometric pipeline, which derives the calibrated photometry from the SM, AF, BP and RP instrument. G fluxes coming from the IDU are accordingly converted into calibrated G magnitudes while the BP/RP fluxes are aggregated in order to provide the integrated BP/RP magnitudes, G_{BP} and G_{RP}. BP/RP spectra are similarly calibrated in flux and wavelength using standard stars for which high quality ground-based spectra are available.

- The RVS pipeline, which takes care of the RVS spectra reduction and calibration like the calibration of the wavelength scale, the treatment of stray-lights or the counter-balancing of the charge transfer inefficiency effects.

High-level cyclic processing consists in the already mentioned NSS treatment, SSO treatment and EO analysis from CU4, in the variable star analysis from CU7 and in the astrophysical parameter inference system (APSIS) from CU8. The products of these high-level tasks are aimed to be published as parts of the Gaia data releases.

Data simulation and publication

Apart from the previously described tasks, the DPAC is also in charge of

- The simulation of the data by the CU2 as they will be observed by Gaia. These simulations can be classified in three categories roughly corresponding to the acquisition of these data at the pixel level in the focal plane by the Gaia Instrument and Basic Images Simulator (GIBIS), to its telemetry by the Gaia System-level Simulator (GASS) and to its scientific products through the Gaia Object Generator (GOG).
- The relativistic modelling of the observations, mainly the bending of light owing to the presence of massive objects being in the vicinity of the line-of-sight towards the observed objects. The latter point requiring a precise and up-to-date ephemeris to be used.
- The observation, filtering and compilation of standard calibration objects regarding fluxes and wavelength scales. Accordingly, numerous observational campaigns have seen the day in order to gather photometrically stable stars based on which apparent G magnitudes can be derived, RVS calibration stars as well as a set of FGK benchmark stars that will equivalently serve as references within APSIS.
- The publication of the results, their documentation and validation. Even if apparently straightforward, the publication of the results coming from the iterative processing that is used within the DPAC reveals a particular complexity given that the coherence between each output must be ensured (what to think about an object being classified as a star by the CU8 while it is considered as a quasar by AGIS...). Also the coherence between each publication of the catalogue should be kept as much as possible, even if this implies that only a limited fraction of the observations will be made public. These coherences are being achieved through strong scientific validation of the objects to be published.

1.2 Active galactic nuclei and quasars

Active Galactic Nuclei (hereafter AGN) are the most energetic objects of the Universe while being non-transient. Although their origin was discussed for a while (see Section 1.2.1), these are actually and most probably thought to come from the accretion of matter onto supermassive black holes (hereafter SMBH). Indeed, we expect these SMBHs to be quite common in the Universe as their existence is theoretically well predicted. Furthermore, observations yield evidences that these might stand in the center of each galaxy. AGN thus occur in a specific period of the lifetime of their host galaxy during which a significant amount of matter was still available in the vicinity of the SMBH such as to maintain this powerful accretion process. The characteristics noticed amongst these AGN (see Section 1.2.2) then allow to guess their generating physical processes and to further derive models of the structure, dynamic, composition and evolution of their various constituting regions (see Section 1.2.3). The extreme luminosities resulting from the high accretion rate of some of the most energetic AGN, called quasars, along with the fact that these were preferably occurring in a period of time when the Universe was still young yield a wealth of unique cosmological applications as we will see in section 1.2.4.

1.2.1 Historical background

The first documented observation of an AGN goes back to 1907, where a photographic plate consisting in a 13 hours exposure of the M77 'spiral nebula' taken through a

prism spectrograph installed at the Lick 36-inch reflector telescope then allowed to obtain the spectrum of what would be recognized now as the [O II] λ373 nm, Hγ λ434 nm, Hβ λ486 nm and [O III] λ496, 500 nm emission lines (Fath 1909). The precise nature of M77 (a type II Seyfert galaxy at $z = 0.0038$) was not recognized at that time given that even its extragalactic nature was not yet known. The Great Debate (or Shapley–Curtis Debate) that took place in 1920 shed light on the extragalactic nature of these 'spiral nebulae', presently referred to as galaxies. This debate was however definitely closed in 1929 thanks to the work of Edwin Hubble who studied the period-luminosity relation of forty Cepheid stars amongst the Andromeda galaxy through observations with the 2.5 m telescope of Mount Wilson Observatory such as to provide a final distance of 275 kpc for Andromeda[3], thus proving its extragalactic nature (Hubble 1929b).

In the early 1930s, the advent of the first instruments that are sensitive enough such as to collect the radio fluxes from cosmic sources (Jansky 1933) gives birth to radio astronomy. The decades following world war II then subsequently see the publication of the first maps of astronomical radio sources (Reber 1944; Ryle & Smith 1948). Nevertheless, the limited astrometric precision of these instruments solely allowed the identification of the optical counterparts of a very few number of sources, amongst which Cygnus A, Centaurus A and M87. Two significant milestones were then gained in 1959 with the publication of the Third Cambridge Catalogue of Radio Sources (Edge et al. 1959, hereafter 3C) and in 1962 with the revision of the latter catalogue (Bennett 1962, hereafter 3CR). These encompassed respectively 471 and 328 sources. Many of these 3C and 3CR sources were rather easily cross-matched with relatively nearby extended galaxies but the available angular resolution rendered the identification of point-like sources more uncertain.

The search for such point-like (or stellar-like) optical counterparts of radio sources from the 3C and 3CR catalogues yielded to the discovery of the 3C48 and 3C273 objects (Greenstein & Schmidt 1964) that were accordingly termed quasi-stellar radio sources, quasi-stellar objects (QSO) or quasars. Even though the nature of their redshifts ($z = 0.367$ and $z = 0.156$, respectively) was discussed for a while: once consisting in a gravitational redshift where photons lost their energy as they escape the potential well of the emitting source or being due to the Doppler shift coming from peculiar stars ejected from our Galaxy, the cosmological nature of this redshift was however soon accepted as it was the most straightforward and the least questionable explanation.

Though quasars were first identified thanks to their radio emission, it was later noticed that not all QSOs emit in the radio domain (Sandage 1965) and that 90% of the latter are actually radio-quiet while their optical spectra show features that are similar to those of radio-loud QSOs. Accordingly, the currently used terminology does not actually require these objects to exhibit a radio emission or to be resolved. These are rather part of the more general class of AGN given that both originate from the same physical process as described in section 1.2.3. In the following, we consequently concentrate on the characteristics and on the nature of AGN though all conclusions can be straightly extrapolated to the case of QSOs.

1.2.2 A multi-wavelength view of AGN

One of the major characteristics of AGN is that they emit over the whole electromagnetic spectrum as illustrated in figure 1.5 for the case of the 3C273 quasar. A

[3] the actual distance being around 778 kpc

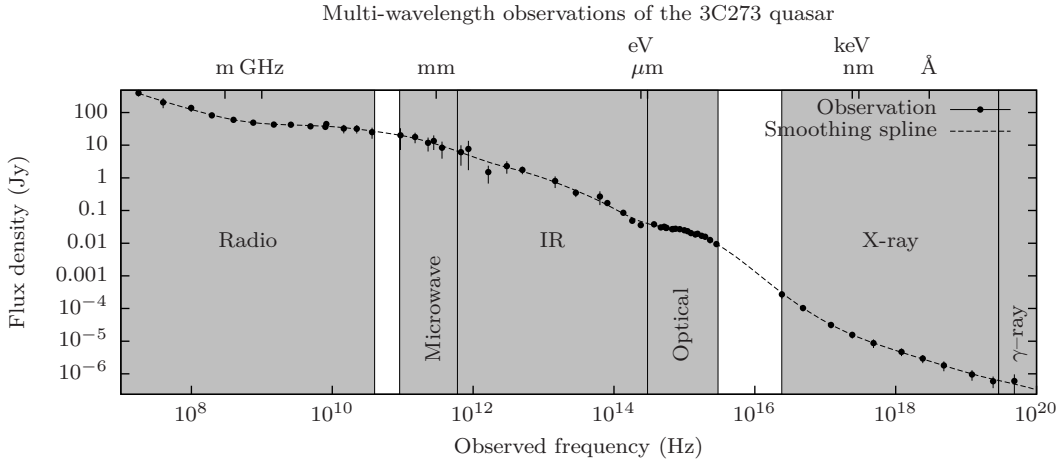

Figure 1.5: Multi-wavelength observations of the 3C273 quasar (Türler et al. 1999).

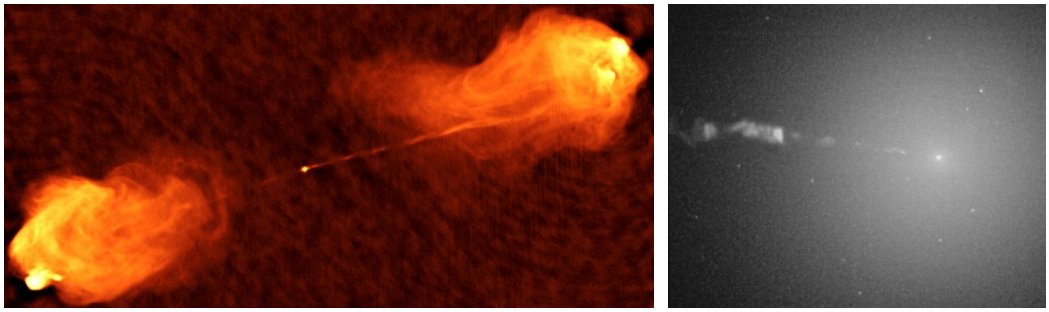

Figure 1.6: (left) The Cygnus A Seyfert II galaxy as seen by the VLA at 5 GHz. Image courtesy of NRAO/AUI. (right) The M87 galaxy as seen by the HST in the optical band. Image credit: Tod R. Lauer, Sandra M. Faber/NASA.

pluridisciplinary analysis is hence necessary in order to fully understand what these objects are and to subsequently try to draw the general frame unifying the various classes of AGN. A recent seminar paper from Padovani et al. (2017) provides a comprehensive review of the various aspects of AGN over a wide range of wavelengths. We summarize here the main AGN characteristics and further try to isolate classes of AGN based on the presence or the absence of some discriminant characteristics.

Radio wavelengths

The radio domain cover the frequency range from 10 MHz to 40 GHz, or equivalently the wavelength range from 0.75 cm to 30 m. As already seen, not all AGN emit in the radio domain and a first distinction should be made between radio-loud AGN and radio-quiet AGN. Observations of sources belonging to the first class predominantly show continuous fluxes in the form of a continuum power-law,

$$f_\nu \propto \nu^{\alpha_\nu}, \tag{1.1}$$

where f_ν is the flux density observed at frequency ν and α_ν is the slope of the continuum that is equivalently termed the 'spectral index'[4]. This continuum part of the

[4]Note that some authors prefer defining the spectral index through $F_\nu \propto \nu^{-\alpha_\nu}$. The reader should hence pay a careful attention to the used convention once using α_ν.

spectrum is particularly noticeable in figure 1.5 for wavelength having $\lambda \geq 1$ m.

This continuum emission is reasonably well interpreted as coming from synchrotron radiation where charged particles are radially accelerated by the presence of a magnetic field which made them whirling in this magnetic field while emitting photons. The strong radio luminosities that are observed require these charged particles to have a very low mass, comparable to the one of the electron. The latter being accordingly the most plausible constituent. The rest-frame frequency of the emitted photons then depends on the relativistic speed of the electrons as well as on the strength of the magnetic field. This hypothesis is further supported by the fact that this radio emission is mostly seen as two strong and highly collimated jets originating from the innermost part of the AGN as depicted in figure 1.6. The observed radio lobes being then produced by the interaction of these particles with the inter-galactic medium at distances up to a few hundreds of kiloparsec from the source. Also, we can notice that such jets are seen in the optical band as well (see figure 1.6, right) thus requiring electrons to move at highly relativistic speed (e.g. with a Lorentz factor of $\gamma \approx 10^6$ if we suppose a magnetic field of strength $B = 10^{-4}$ G). Such a speed being a first hint towards the presence of a SMBH in the center of such systems. For completeness, we have to note that only one jet is frequently visible because of the relativistic Doppler effect that produces a magnification of the jet moving towards us while the counter-jet is conversely attenuated by the same effect (i.e. that is relativistic beaming or Doppler boosting). An example of such a 'single-jet' configuration stands in the image of the M87 galaxy (see figure 1.6). One can further notice in figure 1.6 (left), a compact radio source in the center of the image. The distinction between this compact component and the extended one is easily carried out given that the compact source is expected to be denser and thus optically thick which leads to a flatter spectral index (i.e. $\alpha_\nu \approx 0$) while the extended component is optically thin causing $\alpha_\nu \leq -0.5$.

At the present time, the physical process producing these bipolar jets is still obscure and stands beyond the scope of this work. Furthermore, there is still active discussions about knowing whether the strong synchrotron radiation that is observed in radio-loud AGN comes from the very high energy of the electrons or from the extremely strong magnetic field enclosing them. Whereas we would like to have insights about the physical phenomena producing the various characteristics we will encounter in the optical spectra of Gaia, we have to note that radio emission of AGN can be considered as an additional and apparently independent component regarding higher frequency wavelengths (i.e. at $\nu > 100$ GHz). Accordingly, radio emission will be ignored in the unified model of AGN that will be described in section 1.2.3. Interested readers may find in Urry & Padovani (1995) a detailed approach of such a unification scheme while a more recent review of the subject is provided by Tadhunter (2016). Let us finally note that because of the empirical independence of the radio emission in regards to higher frequencies, these unification schemes remain compatible with the unified model we will later describe.

Infrared wavelengths

The infrared domain covers the wavelength range from 1 μm to 500 μm. It can be further subdivided into: the near-infrared (NIR) band from 1 μm to 3 μm, the mid-infrared (MIR) band from 3 μm to 50 μm and the far infrared (FIR) band from 50 μm to 500 μm. The MIR band is of particular interest regarding AGN since some spectra seem to encompass broad and strong absorption features at $\lambda = 9.7$ μm as well as moderate absorption at $\lambda = 18$ μm (see figure 1.7, the mean spectrum of Seyfert II galaxies for example). Conversely, a potentially different family of AGN (i.e.

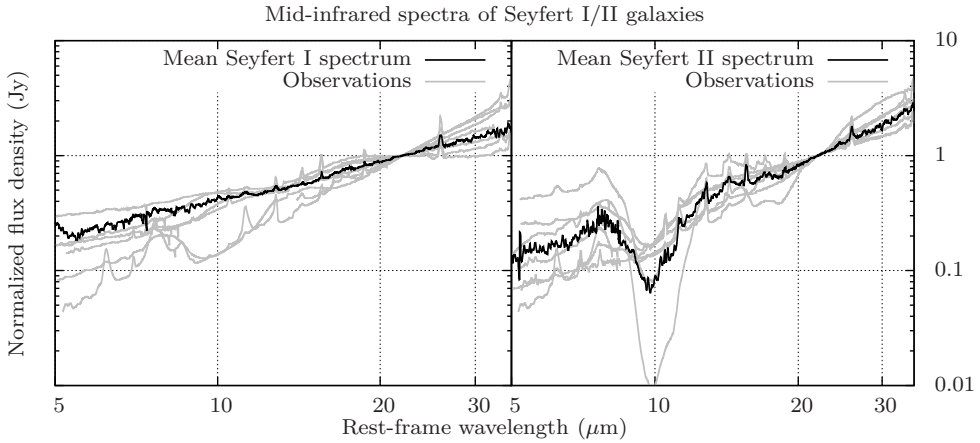

Figure 1.7: Mid-infrared observations of Seyfert I/II galaxies and associated mean spectra. Flux densities were normalized to 1 Jy at $\lambda = 22\mu$m. Spectra were taken from the Combined Atlas of Sources with Spitzer IRS Spectra (CASSIS) based on the classification performed in Wu et al. (2009).

Seyfert I galaxies in figure 1.7) shows emission features–or at least weak absorption features–in those specific regions that correspond to the spectral characteristics that would be produced by dusty silicates. The major difficulty arising in the identification of these features comes from the contamination of the MIR spectra by the emission occurring in the host galaxy of the AGN, mainly from polycyclic aromatic hydrocarbon (PAH) features at $\lambda \approx 6.2, 7.7$ and 11.3 μm as well from dust emission in the 10–30 μm wavelength range. This difficulty is due, amongst other, to the limited angular resolution of the spectroscopic instruments observing in the MIR domain (several arcsec) which do not allow the AGN signal to be disentangled from the one of the host galaxy. It is particularly noticeable in figure 1.7 where some Seyfert I galaxies seem to exhibit strong absorption features while some Seyfert II galaxies exhibit only weak absorption or even no absorption at all.

The modelling of the SED of these silicates features are compatible with the existence of an obscuring torus surrounding the extremely luminous core of the AGN whose inner boundary is delimited by the sublimation of the grains of dust. Various dust repartitions were studied, either consisting of gas, clumpy regions of agglomerated dust or of a mixture between both. However, the actual geometry of the obscuring dust is not already known in detail as it has a limited impact on the SED. It usually varies from an irregular shape to an ideal toroid. The presence of such a torus would first indicate that the spectral appearance of AGN strongly depends on the inclination of this torus with respect to the line-of-sight towards the AGN. Any realistic unification model should hence have to deal with this inclination.

Some other remarkable characteristics of the AGN in the IR domain encompass a minimum at $\lambda \approx$ 1–2 μm corresponding to the sublimation of the dust grains as well as a local maximum between 10 and 100 μm owing to the thermal emission of dust that is heated by the AGN central engine which produces the overall continuum shape of the MIR spectra. Still, the synchrotron radiation can also play a significant role in the continuum emission if the obscuring torus is observed from a favourable angle. IR observations have the advantage over optical observations that the former are less sensitive to obscuration and are consequently particularly well suited for the study of the obscuring torus.

Table 1.2: Dominant emission lines observed in the optical spectra of AGN.

Rest-frame wavelength (nm)	**Element(s)**	**Rest-frame wavelength (nm)**	**Element(s)**
103.382	Lyβ + O VI	372.709	[O II]
121.524	Lyα	388.90	He I
124.081	N V	407.23	[S II]
130.553	O I	410.289	Hδ
133.531	C II	434.168	Hγ
139.761	Si IV	436.444	[O III]
139.98	Si IV + O IV]	486.268	Hβ
148.65	N IV]	496.029	[O III]
154.948	C IV	500.824	[O III]
164.04	He II	630.205	[O I]
166.585	O III]	636.554	[O I]
185.74	Al III	652.903	[N I]
190.873	C III]	654.986	[N II]
232.6	C II]	656.461	Hα
243.95	[Ne IV]	658.527	[N II]
279.912	Mg II	671.829	[S II]
334.679	[Ne V]	673.267	[S II]
342.685	[Ne V]		

Optical wavelengths

The optical domain covers the wavelength range from 100 nm to 1 μm and encompasses the UV domain at $\lambda < 400$ nm. As depicted in figure 1.8, optical spectra of AGN exhibit very complex morphologies in this domain which makes it extremely pertinent for the discrimination between the various types of AGN. Furthermore, the transparency of the Earth atmosphere at these wavelengths renders the acquisition of good quality spectra much easier than in other domains where space observations are often required[5]. This explains, for example, the higher density of observations at these wavelengths in figure 1.5. These advantages ultimately make it the de-facto domain for the identification, classification and characterization of AGN. Nevertheless, this domain will still suffer from some inherent selection effects and observational biases like its low sensitivity to obscured AGN or its high contamination rate by massive stars. From figure 1.8, we notice some dominant spectral characteristics of AGN, namely: (i) a continuum spectrum approximately following a power law, (ii) broad emission lines with FWHM[6] ≥ 1000 km s^{-1} coming from known permitted or semi-forbidden atomic transitions (e.g. the hydrogen Balmer series, Lyα, C IV or C III]), (iii) narrow emission lines with FWHM ≤ 1000 km s^{-1} coming from forbidden atomic transitions (e.g. [Ne III] or [O III]) (iv) narrow absorption lines at $\lambda < 121$ nm coming from neutral hydrogen being along the line-of-sight towards the observed QSO (v) broad absorption lines with FWHM $\geq 5,000$ km s^{-1} presumably coming from the outflow of absorbing matter in the near vicinity of the SMBH. The spectral emission lines that are commonly observed in AGN are summarized in Table 1.2. These are

[5] The earth atmosphere being opaque to all wavelengths except in the high frequency radio domain (30 MHz $< \nu <$ 10 GHz), in visible light (400 nm $< \lambda <$ 700 nm) and in two narrow windows in the NIR (1 μm $< \lambda <$ 2.4 μm) and in the MIR (8 μm $< \lambda <$ 14 μm).

[6] Full-width at half maximum

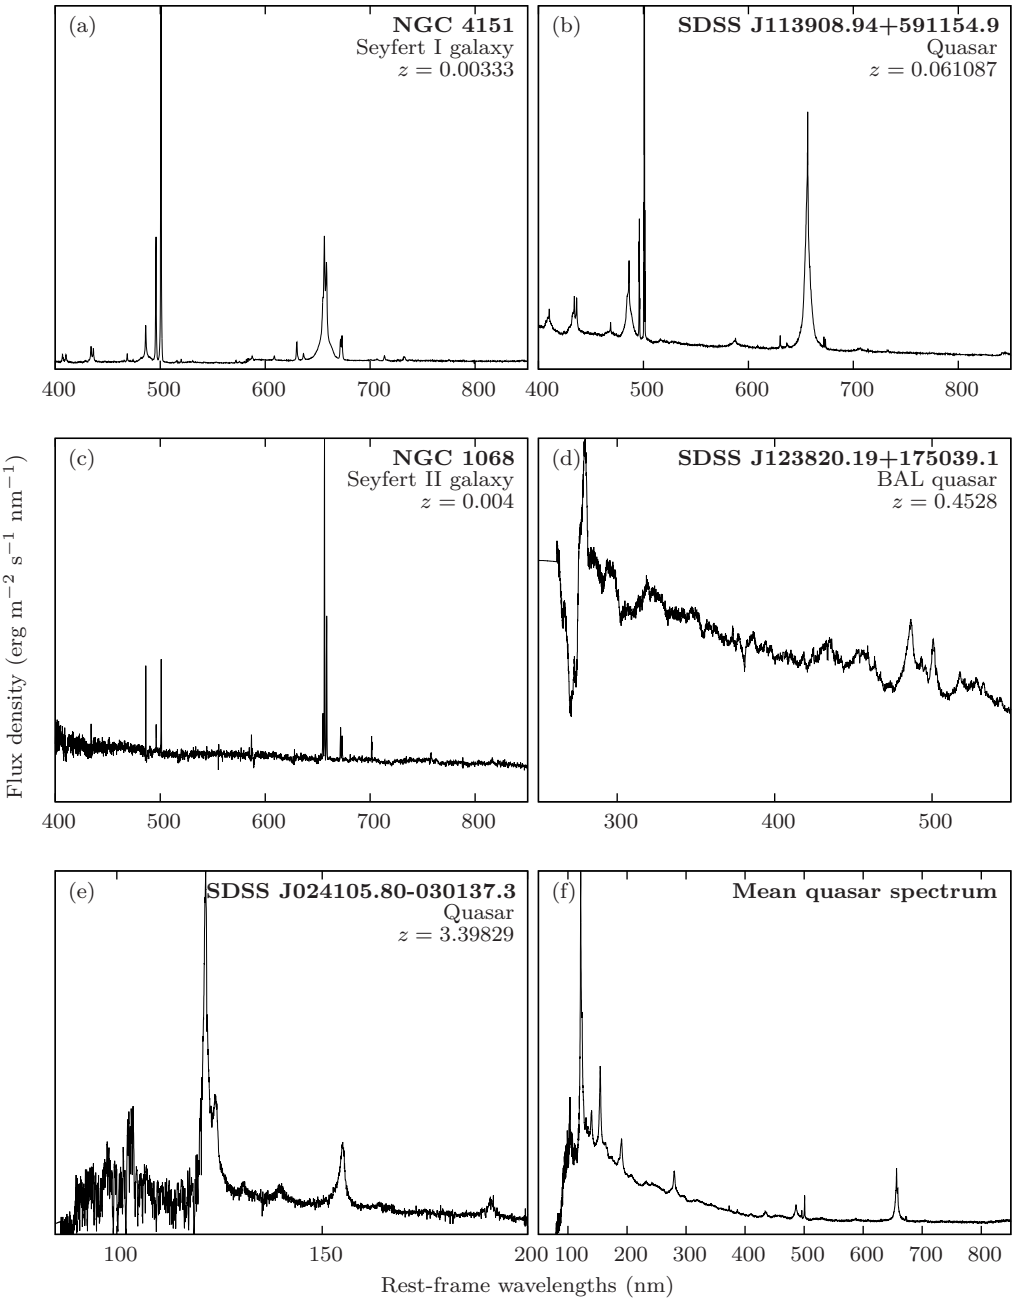

Figure 1.8: Different types of AGN as seen in the optical band: (a) a type I Seyfert galaxy showing a combination of broad and narrow emission lines coming mainly from the Balmer series and from the [O III] doublet, respectively, (b) a quasar spectrum showing patterns similar to those of the Seyfert I galaxy while having a higher Hβ/Hα ratio, a steeper continuum slope and a globally higher absolute luminosity, (c) a Seyfert II galaxy showing only narrow emission lines, (d) a BAL QSO showing a strong absorption feature bluewards the Mg II emission line as well as a very steep continuum slope, (e) a high redshift quasar showing numerous absorption lines bluewards of the Lyα emission line and (f) a mean rest-frame quasar spectrum from Vanden Berk et al. (2001). Spectra (a) and (c) are taken from the Large Sky Area Multi-Object Fiber Spectroscopic Telescope (Cui et al. 2012) while spectra (b), (d) and (e) are taken from the twelfth data release of the SDSS (Alam et al. 2015).

Figure 1.9: Host galaxies of AGN viewed by the HST, from left to right: NGC 6814, a Seyfert galaxy of type I, M77, a type II Seyfert galaxy and 3C273, a type I quasar. AGN correspond to the bright spots seen at the center of each galaxy. In the case of 3C273, the central region was however masked in order to reveal the host galaxy. Also note that the obscuring torus of the AGN is not necessary aligned with the galactic planes as both NGC 6814 and M77 are seen from a rather similar angle while being of different Seyfert type. Images credit: ESA/Hubble & NASA.

the basic tools for the identification, classification and characterization of the AGN and will be used extensively through the present document.

Other noticeable features of AGN in the optical band encompass a flux excess in the UV that is superimposed to the power law continuum. This flux excess, called the 'big blue bump', can be seen in figure 1.5, for example. As already mentioned in the beginning of this section, the power-law component of the spectrum can be fairly well modelled as a synchrotron emission extending up to the optical domain while the origin of the big blue bump rather resides in the thermal emission of the accretion disk which spans over the optical and UV wavelength ranges. It is also worth to mention that narrow emission lines show very little variability when compared to the continuum and broad emission lines, which is consistent with the fact that the region producing these narrow emission line lies at much larger distances than the region producing the broad emission lines and the continuum.

Going deeper in the details of these characteristics, it was soon recognized that these can be straightforwardly explained as a same physical phenomenon that is observed according to different viewing angles (Antonucci 1993). This unified model of AGN is further described in section 1.2.3 but as we already have sufficient insights in order to understand–at least partially–this model, we will still mention some of its basic components. As already suggested by the IR observations, the presence of a clumpy obscuring torus will have as an effect to hide the inner part of the AGN in the optical domain. This inner region consists of a SMBH surrounded by an accretion disk and clouds of gas moving at relativistic speeds which produce on one hand the steep continuum through both its thermal radiation and the synchrotron radiation owing to its magnetic field, and the broad emission lines observed in Seyfert I galaxies and quasars on the other hand. Also, we saw that the clouds of gas producing the narrow emission line seem to be situated further away from the plane containing this torus. Consequently, in cases where this torus is observed nearly edged-on, only narrow emission lines will be visible leading to the appearance of Seyfert II galaxies. The narrow and broad absorption lines seen is some AGN spectra being already explained by the presence of matter being along the line-of-sight towards the observed AGN and from the outflow of gas out of the central region.

The presence or absence of some of these characteristics then allow a straight, though not unique, classification of these AGN into:

- **Seyfert I galaxies** which have broad and narrow emission lines superposed to their continua, have low absolute luminosities, relatively low redshifts and whose host galaxies are often resolved (see figure 1.9, for examples of Seyfert I/II host galaxies).

- **Seyfet II galaxies** which show a strongly damped continuum and no broad emission lines, the latter being replaced by their narrow counterparts (e.g. Hβ emission lines with FWHM ≤ 1000 km s^{-1}), have low redshift, luminosity and frequently resolved host galaxies.

- **Blazars** which show a dominant continuum with no emission lines. This type can be further classified into Optically Violent Variable (OVV) if the AGN show evidence of rapid variability or BL Lac objects otherwise. This peculiar type of AGN would correspond to the case where the obscuring torus is seen nearly face-on, giving access to the innermost part of the AGN. This central region then outshining the spectral features that are otherwise commonly observed.

- **Quasars** which are intrinsically luminous Seyfert I/II galaxies having absolute magnitude of $M_{\rm B}[z=2] \lesssim -23$. Their host galaxies are commonly unresolved owing to their large redshift and the strong luminosities of their cores but let us still mention that this point is not a distinguishing characteristics since some quasars still have their host galaxy resolved (see figure 1.9, for example). Like Seyfert galaxies, quasars can be subdivided into type I and II QSOs according to the FWHM of their emission lines. The observation of type II quasars in the optical domain is however strongly biased by the occultation of most of the associated electromagnetic spectrum by the obscuring torus [7].

- **BAL QSOs** which have spectra that are similar to those of type I quasars while having strong absorption features that are blueshifted when compared to some of their broad emission lines.

Now, if we assume that the major differences between type I and type II AGN is a matter of inclination, we will have that such a binary classification will not describe the underlying distribution of the genuine classes correctly, the latter being accordingly rather smooth and continuous. This gave birth to some intermediate classes like type 1.2, 1.5, 1.8 and 1.9 Seyfert galaxies depending on the existence of some broad emission lines in the spectra of Seyfert II galaxies (Osterbrock 1981). Nevertheless, a straight binary classification is sufficient for the purpose of the present study. Let us finally mention that while quasars represent the bright end of the AGN classification, the faint end is occupied by the Low Ionization Narrow-Line Emission Radio galaxies (LINER). These represent the larger fraction of observed AGN but will not be mentioned anymore in the following given they fit into the type I/II unification scheme.

X-ray wavelengths

The X-ray domain covers the wavelength range from 100 eV to 120 keV. Spectra of this domain are emitted by the innermost part of the AGN where complex interactions take place. Synchrotron emission from the jet as well as thermal emission from the

[7] Nevertheless, a type II QSO was still observed at $z = 3.288$ by Stern et al. (2002).

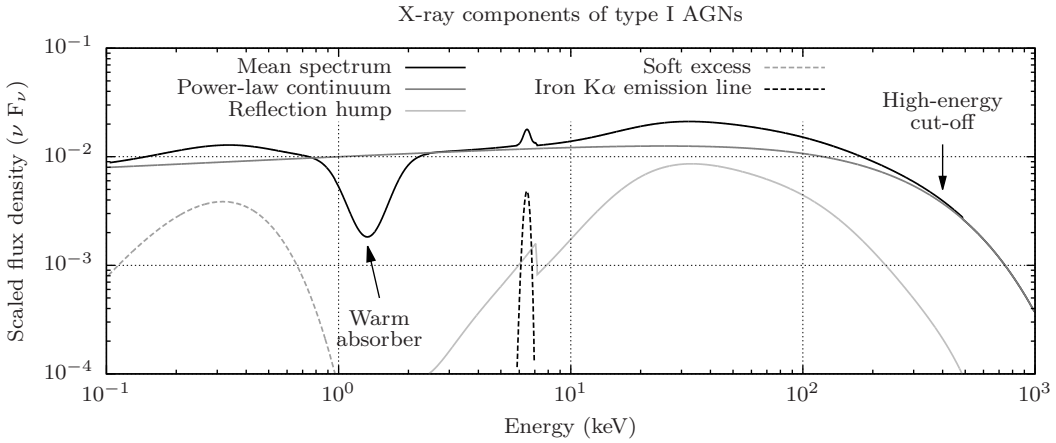

Figure 1.10: Mean X-ray spectrum of a type I AGN. The main contributors are a power-law continuum declining at high energies, a reflection hump around 10–50 keV from Compton scattering in the accretion disk's corona, a thermal component at low energy (the soft excess) and the Kα emission line. Note that the flux density, F_ν, was scaled according to the photons frequency, ν, such as to highlight the deviation of the mean spectrum from a perfect power-law continuum.

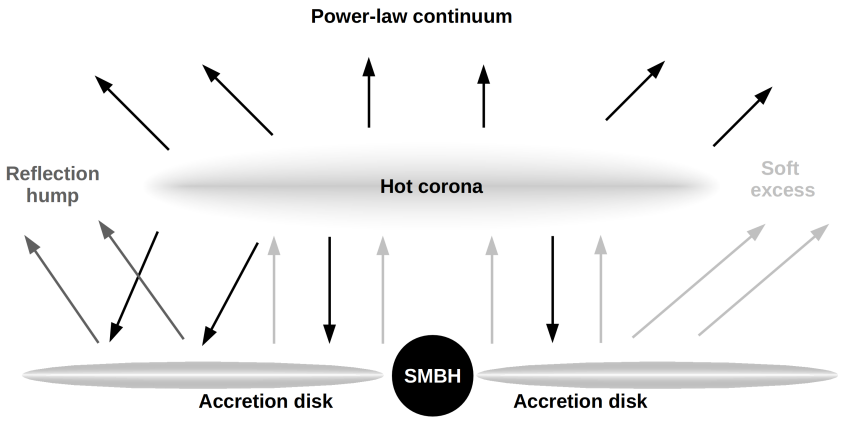

Figure 1.11: Illustration of the accretion disk and hot corona producing the main spectral features observed in the X-ray spectrum of AGN. The thermal emission from the accretion disk (at $k_\mathrm{B}T \approx 100$ eV) radiates in the faint end of the X-ray spectrum where it gives rise to the soft X-ray excess if observed directly. This spectrum can nevertheless be re-processed in a hot corona (at $k_\mathrm{B}T \approx 30$ keV) through inverse Compton scattering, producing higher energy photons in the form of a power-law continuum. Given that this continuum is emitted isotropically by the corona, the latter can be further reflected on the accretion disk yielding the reflection hump observed at 20–30 keV.

accretion disk ($T \lesssim 10^6$ K) are indeed insufficient to explain the wide range of X-ray energies that are observed in some AGN (see figure 1.10, for example). Rather, models have to rely on the presence of a thin layer of hot electrons, a 'hot corona', surrounding the accretion disk. The existence of a similar medium was already assessed in stellar outer atmospheres thus giving even more credibility to this hypothesis. Still, the mechanisms allowing this corona to gain temperatures of the order of 10^8 K are not yet fully understood though magnetic fields reconnection is supposed to play a significant role (Cao 2009). If we assume the existence a such a hot corona, it will have as an effect to boost the low-energy photons coming from the thermal emission of the accretion disk into X-ray photons through the inverse Compton effect. At the opposite of the classical Compton effect, where photons loose their energy as they interact with low energy particles, the inverse Compton effect corresponds to the case where high energy particles transfer their energy to photons during a similar interaction. The resulting continuum will hence follow a power-law that is truncated at very high energies ($E \gtrsim 100$ keV) because of the limited temperature of the corona as well as from the limited energy of the photons emitted by the accretion disk. Owing to the isotropic emission from the hot corona, some of the high energy photons will be reverberated on the accretion disk through simple Compton scattering yielding the reflection hump (or Compton hump) observed in figure 1.10 while lower energy photons get absorbed by the accretion disk giving rise to thermal radiation at lower energies, the so-called 'soft excess', and to the Kα emission line from the de-excitation of the low-ionization stages of iron. The presence of a warm absorber around $E \approx 1.5$ keV in some AGN corresponds to the outflow of matter out of the central engine as later described in section 1.2.3. A sketch of the interactions occurring between the accretion disk and the hot corona regarding the X-ray emission of AGN is given in figure 1.11. The widening of the Kα emission line as well as its asymmetry are part of ongoing debates, once originating from complex re-processing of the spectrum coming from the accretion disk or being due to relativistic effects occurring in the event horizon of the SMBH. Let us mention that the selection of AGN based on their X-ray emission is actually the fairest one regarding its completeness and its very low contamination rate coming from its nearly universal presence in AGN, its ability to go through large columns densities of gas and dust at high energies and from the fact that galactic X-ray emissions are typically weak when compared to those of AGN.

γ-ray wavelengths

The γ-ray domain covers the wavelength range from 120 keV and above (up to ~ 50 TeV, according to current technologies). Spectra covering this band are thought to originate from the inverse Compton scattering of the synchrotron radiation with their emitting electrons (synchrotron self-Compton effect). Blazars, from the privileged inclination of their jets, are accordingly the most probable γ-ray emitters amongst AGN. Still, the involved energies imply that these photons can be converted into electron/positron pairs upon interaction with a particle or another photon once the energy carried out by these photons exceeds the one of two rest-frame electrons (2×511 keV). This yields an opacity in the high energy regime that is however tempered by the previously mentioned Doppler boosting that tends to increase the observed photon frequency while dimming the photon-photon opacity (i.e. the rest-frame frequency of the photon being then $E_{\rm rest} < 1$ MeV). The fluxes measured from extragalactic γ-ray sources might hence not correspond to the intrinsic fluxes because of the photon-photon opacity. This opacity is mainly due to IR photons coming from the first stars and proto-galaxies in the early universe, that is the extragalactic background light

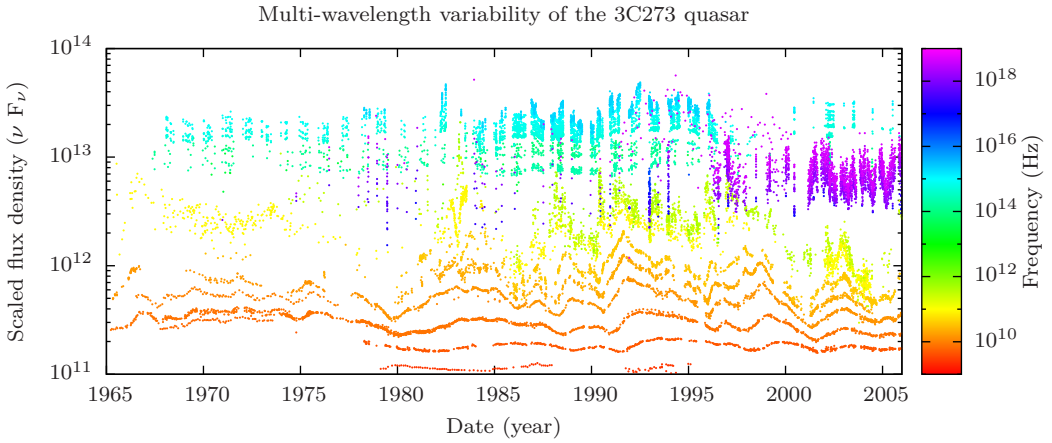

Figure 1.12: Time series of the 3C273 quasar taken between 1965 and 2006 over the wavelength range 10^9–10^{20} Hz. Faster variability is often observed within the high frequency domains that correspond to the innermost part of the AGN. Events occurring in this domain propagate into the low frequency domains along with a slower variability and lags. Observations were taken from the 3C 273 's Database (Soldi et al. 2008).

(EBL). Supposing that the intrinsic SED of AGN in the TeV regime can be guessed, or is universal, then the EBL can be indirectly reconstructed through the impact it has on the continuum slope of the AGN in the γ-ray domain. Let us still mention that AGN are rather dim sources in the range 200 keV $< E <$ 1 GeV if compared to galactic sources while most AGN emitting in this range are found at low redshift because of the photon-photon opacity. Consequently, this domain often yields a biased selection of sources coming with a high contamination rate.

Variability

Variability is one of the characterizing components of AGN. The latter consists in stochastic flux variations over the entire electromagnetic spectrum in time scales ranging between a few decades down to a few minutes depending on the wavelength range in which this variability is observed. The multi-wavelength study of this variability can be used as a way of selection but it principally yields strong constraints on the sizes, interactions and disposition of each sub-structure composing the AGN. Indeed, from the minimal time scale variation in a given band, one can set an upper size on the region producing this variability (Terrell 1967) while the lags that are observed between these various bands can be used such as to guess the distances separating the various emitting components. Figure 1.12 presents the variability curves of the 3C273 quasar taken between 1965 and 2006 over the wavelength range 10^9–10^{20} Hz. Globally, one can notice that the characteristic time scales associated with the high frequency ranges are by far shorter than those at lower frequencies (i.e. the former exhibiting faster variability). Also, long wavelength observations generally tend to lag behind those of shorter wavelengths. This view is consistent with the image of an AGN where the most energetic photons are produced within the innermost part of the system while their propagation induces feedbacks in more distant regions that will reactively produce lower energy components with delays. We can finally note

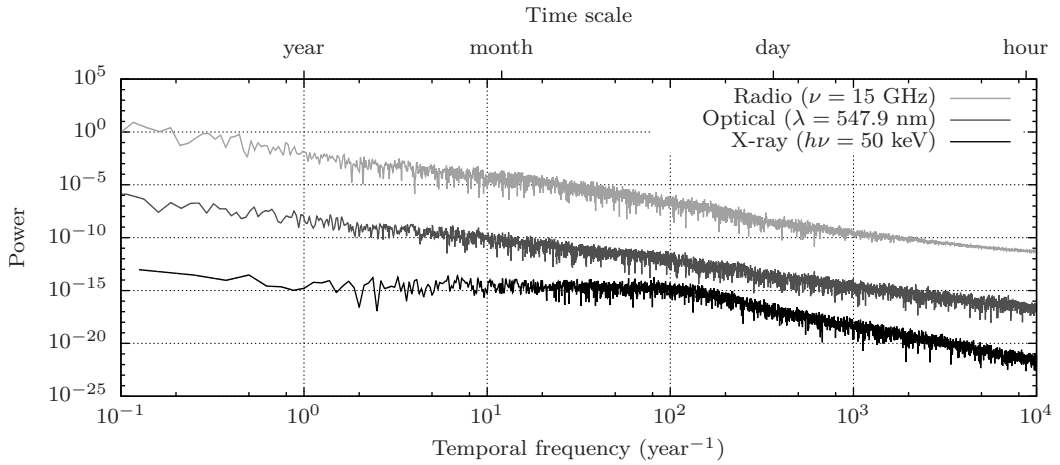

Figure 1.13: Power density spectra of the 3C273 quasar over various observational wavelength ranges. The optical and radio power spectrum density can be fairly represented as two power-law with spectral index $\alpha_\nu = -2.1053$ and -2.4542, respectively while the X-ray PSD exhibits a break around $\nu_b = 100$ year^{-1} from which useful quantities can be derived. Observations were taken from the 3C 273 's Database (Soldi et al. 2008).

that large amplitude variations are mainly driven over long time scale, potentially corresponding to large clouds of gas being accreted.

The optical band, from its complexity regarding AGN, is particularly rich in term of insights we can draw out of variability studies. In radio-quiet AGN, perturbations from the accretion disk will first induce a higher thermal emission that will be converted into X-ray radiations through its reprocessing in the corona. This will in turn produce a rise in the (optical) continuum slope of the AGN followed by a similar rise in the broad emission lines and subsequent reactions in the IR domain. Such highly correlated observations allow together to probe the existence and distances of the accretion disk, of the hot corona, of the broad emission line region and of the obscuring torus, for example. More specifically, a change in the shape and amplitude of the broad emission lines in response to a change in the continuum slope, that is 'reverberation mapping', allows to accurately measure the size of the broad emission line region as well as the mass of the SMBH (Peterson & Horne 2004). In radio-loud AGN, the relativistic jet is considered to be the principal source of emission, the latter often outshining the emission from the accretion disk and corona. The interpretation of this correlated variability being hence less straightforward than in the case of radio-quiet AGN given that the modelling of the variability in radio-loud AGN is rendered difficult as the mechanisms and structure of the jet are currently not well understood.

A useful tool for variability analysis stands in the power spectral density (PSD) defined as the squared amplitude of the time series taken in the Fourier domain. More precisely, if we consider the flux density recorded from a source at time t as $f(t)$,[8] then the PSD of the time series at frequency ν is given by $|F(\nu)|^2$ where $F(\nu) = \int f(t)e^{-2\pi i t\nu}dt$ is the Fourier transform of $f(t)$ evaluated at ν. The resulting PSD can then be used to straightly quantify the various frequencies that are present within the time series and to efficiently constrain the sizes of the emitting components.

[8]We supposed here that time series are continuous which is obviously unrealistic since observations are often unevenly sampled in time but this still prevents unnecessary intricacies to stem the pedagogical goal of the present introduction to the PSD.

From figure 1.13, for example, one can easily notice that the radio and optical PSD seem to follow a power-law in first approximation along with slightly different slopes ($\alpha_\nu = -2.4542$ in the radio domain while $\alpha_\nu = -2.1053$ in the optical domain). This is again consistent with the fact that higher frequency wavelengths tend to produce faster variability in regions that are smaller in size. The X-ray PSD further shows a break around $\nu_b = 100$ year^{-1}. The position of this break can be used as a way to set an upper limit on the radius of the SMBH. Indeed from Terrell (1967) we have that

$$R \lesssim \frac{1}{z+1}\frac{c}{\pi\nu_b} \tag{1.2}$$

yielding in the case of 3C273 $R \lesssim 8.44 \times 10^{-4}$ pc. Secondly it can be used so as to set a limit on on the SMBH mass, $M_{\rm BH}$. Indeed there seems to be a connection between the mass of the SMBHs in AGN and the mass of galactic black hole in X-ray binaries as both show similar shapes that might hence originate from common effects (González-Martín & Vaughan 2012). Differently stated, matter accretion on black holes whether these are supermassive or not tends to produce a rather universal shape of the PSD. Supposing that this relation holds, then one can assume that the relation $\log \nu_b \propto \log M_{\rm BH}$ coming from binary X-ray studies also applies to the SMBH mass giving $M_{\rm BH} \approx 1.19 \times 10^8 M_\odot$ for the case of 3C273. Finally, we can note as expected that no periodic signal can be found in time series of AGN that would otherwise translate into a peak in the PSD.

1.2.3 A unified model of AGN

We will now put all the pieces together in order to provide a concise model of AGN. As highlighted in Netzer (2015), the model we describe here is known to be incomplete as it can not account for the whole variety of observations that are made out of AGN. A more complete model would indeed take into account the covering factor of the obscuring torus, its clumpiness, the accretion rate of the SMBH and its mass, the bolometric luminosity of the AGN, the strength of the jet as well as the host galaxy properties. The model we consider here is based on ideas first introduced by Antonucci (1993) and Elvis (2000). It allows to describe the vast majority of the characteristics observed within the optical spectral of AGN and is accordingly considered as the simplest model that fits our needs. A sketch of the components intervening in this model is given in figure 1.14. Though most of these components were already outlined in section 1.2.2, we detail here the main constituents of this model, their characteristics emission/absorption features as well as their interactions.

The supermassive black hole

The presence of a SMBH as the central engine of AGN was suspected a long time ago (e.g. Pringle et al. 1973). Their existence in the center of AGN can be probed using some simple considerations: (i) As already mentioned, the time delays, Δt, between a change in the continuum slope of the AGN and the associated change in a given broad emission line can be used so as to infer the distance of the region emitting this broad emission through $r = c\Delta t$. Furthermore, the FWHM of this broad emission line, $\Delta\lambda$, is related to the velocity dispersion of the gas around the central massive source through $\sigma = c\Delta\lambda/\lambda$ where λ is the theoretical wavelength of the emission line[9]. Assuming that the clouds of gas are approximately moving around Keplerian orbits

[9]Note that λ and $\Delta\lambda$ should be taken within the same reference frame, that is either in the rest frame of the AGN or in the reference frame of the observer.

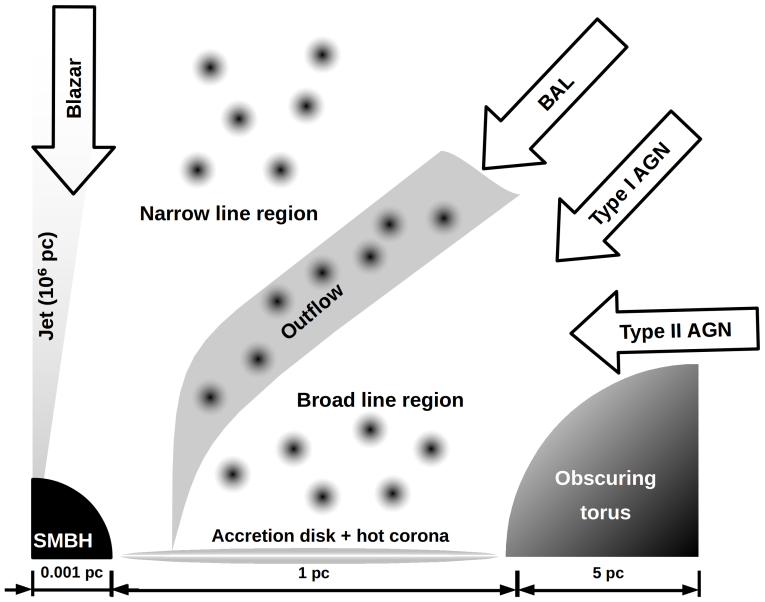

Figure 1.14: Illustration of the components intervening in the unified model of AGN described in section 1.2.3. Differences between the various kinds of AGN (i.e. type I/II AGN, BAL QSOs and blazars) are explained by the viewing angles with respect to the plane containing the obscuring torus, here depicted as labelled arrows. Different viewing angle give access to different emitting components, that together allow to explain most of the diversity encountered within the spectra of AGN. Note that the jet and the outflow of matter are virtually never seen together and might hence constitute a unique structure.

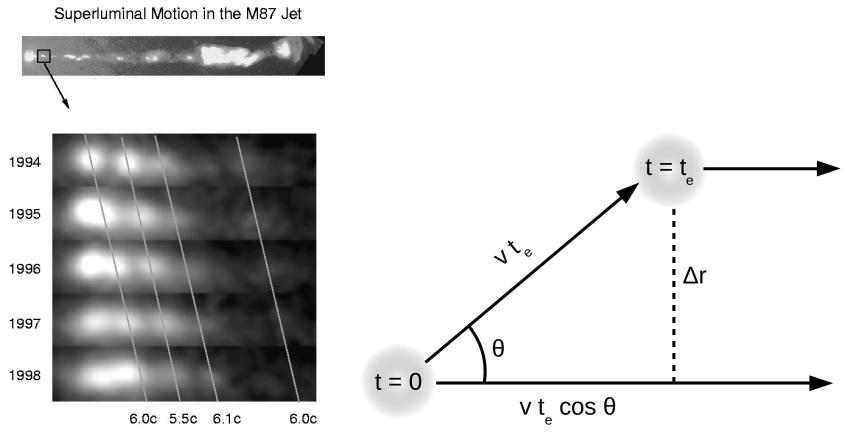

Figure 1.15: (left) Apparent superluminal motion in the jet of the M87 galaxy. Image courtesy of John Biretta, Space Telescope Science Institute (Biretta et al. 1999). (right) Sketch explaining the apparent superluminal motion observed in the jet component. Suppose that the jet component moves at a velocity v towards the observer with an angle of θ with respect to the line of sight, suppose also that this component emits at two distinct times, $t = 0$ and $t = t_e$. From basic geometrical considerations, we find that the transverse motion of the component is given by $\Delta r = vt_e \sin\theta$ while the time delay between both observed emissions is given by $\Delta t = t_e(1 - \beta\cos\theta)$, with $\beta = v/c$. The apparent speed is hence given by $v' = \Delta r/\Delta t = v\sin\theta(1 - \beta\cos\theta)^{-1}$ where nothing prevent v' to exceed c (e.g. $v' \approx 2.4c$ if $v = c$ and $\theta = \pi/4$). See Blandford & Königl (1979) for a thorough analysis of these apparent superluminal motions.

and that their masses are negligible, then we can use the relation (Schneider 2006, Section 5.1.1)

$$M \propto r\sigma^2/G \tag{1.3}$$

so as to have an estimation of the mass that is encompassed within the radius r. The proportionality constant from the latter equation is related to the geometry and viewing angle of the orbit of the broad line region (BLR) and is equal to 1 in the case of a circular orbit viewed from the edge. The knowledge of this proportionality constant allows us to get a precise estimation of the enclosed mass. This technique called 'reverberation mapping' yields times delays that are in the order of $\Delta t(\mathrm{H}\beta) \approx 100$ days along with $\mathrm{FWHM}(\mathrm{H}\beta) \approx 5$ nm (Kaspi et al. 2000), yielding $\sigma \approx 3000$ km s^{-1} and $M \approx 2 \times 10^8 M_\odot$ if we assume that the proportionality constant is 1. (ii) Hard X-ray observations show variability over a period of a few days (as already seen in section 1.2.2) down to a few hours (Mushotzky et al. 1993), yielding typical radius in the order of $R \lesssim 10^{-4}$ pc. Now the Schwarzschild radius that is associated with a mass of $M = 2 \times 10^8 M_\odot$ (i.e. the radius for which the escape velocity against M is equal to the speed of light), is given by

$$r_s = 2\frac{GM}{c^2}, \tag{1.4}$$

in our case $r_s = 2 \times 10^{-5}$ pc, which is of the same order of magnitude as the radius found from X-ray variability. This is a first suggestion for the presence of a SMBH in the center of some AGN.

Another evidence for the presence of a SMBH as the central engine of AGN stands in the apparent superluminal velocities observed in the radio and optical jets of some AGN as illustrated in figure 1.15. In the same figure, we derived the apparent velocity of the jet component as

$$v' = \frac{v \sin\theta}{1 - \beta\cos\theta} \tag{1.5}$$

where v is the effective velocity of the jet component, θ is the angle of the component's motion with respect to the line-of-sight and where $\beta = v/c$. Considering a fixed value of v, we find that the maximal apparent velocity is reached when $\cos\theta = \beta$ yielding

$$v'_{\max} = \gamma v \tag{1.6}$$

where $\gamma = \left(1 - \beta^2\right)^{-0.5}$ is the Lorentz factor associated with the speed v. We hence have that $v' \geq c$ if and only if $v \geq c/\sqrt{2} = 0.70711c$. Such highly relativistic speeds being attained only in black holes and neutron stars, the former being the privileged solution given the previously derived masses.

The accretion disk

The driving mechanism of AGN is the accretion of matter onto a SMBH where the potential energy of the gas is first converted into kinetic energy. Owing to the finite angular momentum of the particles and the friction occurring amongst those particles, the infalling gas takes the form an accretion disk where the kinetic energy is converted into heat due to the internal friction coming from the differential rotational velocities that are present in the disk (i.e. the angular velocity of the particles depending on the distance to the center of the accretion disk). At a time, thermonuclear energy was considered to be the driving mechanism of AGN. Nevertheless, the most efficient nuclear reaction, that is the nuclear fusion of hydrogen into iron, yields an efficiency of

$\epsilon \lesssim 0.008$. Considering the previously derived SMBH masses, we have that the maximal energy that can be produced through the burning of $M \approx 10^8 M_\odot$ of hydrogen into iron over the whole lifetime of the AGN is given by

$$E = \epsilon M c^2, \tag{1.7}$$

yielding $E = 1.4 \times 10^{53}$ J. Now, given that some AGN show radio jet extending up to a few Mpc from their core, their age should then be as large as the time that is needed for the light to reach such distances, that is $t \gtrsim 10^7$ year. The observed bolometric luminosities of AGN being in the order of $L_{\text{bol}} \approx 10^{40}$ W, we have that the total energy produced by an AGN over its lifetime considering that its luminosity remains roughly constant over the ages is given by $E = tL_{\text{bol}} \approx 3 \times 10^{54}$ J. We can hence reject the hypothesis according to which the energy production in AGN is due to thermonuclear activity. This coarse derivation assesses that the conversion of the gravitational potential into heat through the accretion of matter onto black holes must provide a far more efficient way of producing energy, that is $\epsilon = 0.06$ in the case of a non-rotating black hole and $\epsilon = 0.29$ in the case of a black hole with an optimal rotation period, according to current theory (Schneider 2006, Section 5.3.1). This leads to released energies that are compatible with observations.

The hot corona

The radiation emitted out of the accretion disk is limited by the accretion rate of the SMBH as well as by the fact that the last stable orbit around a (SM)BH stands at $r \approx 3r_s$ if we consider a SMBH with an averaged angular momentum (i.e. spin). Consequently, most of the electromagnetic radiation coming from the accretion disk is emitted in the optical/UV domain. Higher energies should hence be produced through another mechanism that relies on the existence of a hot corona, as already described in section 1.2.2 and through figures 1.10 and 1.11. As a remainder, photons from the accretion disk gain energy thanks to the inverse Compton scattering occurring in the hot corona, the latter are then either seen directly as a X-ray continuum or are reprocessed within the accretion disk giving rise to the X-ray soft excess and to the reflection hump depending on whether these are absorbed by the accretion disk or reflected, respectively.

The obscuring torus

The obscuring torus is the key component of the unified model that is presently described. From MIR observations one can deduce the presence of such a disk as already demonstrated in section 1.2.2. Whereas its presence was soon assessed, its dimension remained an open question for a long time owing to the fact that a thick and distant torus shows a similar signature in the observables as a thin and nearby torus. Indeed, it was long believed that the torus may extend up to ~ 100 pc and more according to early theoretical works. From statistical studies of Seyfert I/II galaxies in the IR, Schmitt et al. (2001) first showed that the ratio of the height of the disk to the size of its inner radius is approximately given by $h/r \approx 1$. Combining interferometry and spectroscopy, Jaffe et al. (2004) finally provides a radius of a few parsec for the NGC1068 galaxy along with $h/r \gtrsim 0.6$. Subsequent observation then allowed to resolve such a torus around the core of the Circinus galaxy with $r \leq 2$ pc (Tristram et al. 2007). This size was however a hurdle for the composition of the torus that was commonly considered at that time as being smooth and continuous. The dust emission at 9.7μm effectively requires temperatures in the order of 300K,

and hence large distances from the core that are not compatible with those observed. This issue can nevertheless be solved if we consider the torus as being composed of clouds contained in a toroidal structure (Nenkova et al. 2008). The clumpy nature of the torus yields some refinement of the present model where some AGN viewed edge-on might hence appear as type I AGN given that, in some cases, no cloud will hide our view toward the central source. With this anisotropic obscuration, the type I/II classification is hence based on a probabilistic approach rather than on a straight and unique parameter that is the viewing angle.

The broad and narrow line regions

Emission lines from AGN provide a wealth of informations about the physical processes occurring in these fascinating objects. From the shift they encounter when compared to laboratory calibrators, one can see a straight consequence of the expansion of the Universe through their redshift. To date, the largest redshift that was observed for an AGN stands at $z = 7.085$ (Mortlock et al. 2011), meaning that the SMBH it encompasses (with $M \approx 2 \times 10^9 M_\odot$) was formed only 770 million years after the Big Bang, according to current cosmological models[10]. The broadening of these emission lines, which is a straightly measurable quantity, allows us to determine the angular velocities of the emitting regions around the core of the AGN. The time delay noticed between a change in the continuum slope and the corresponding change in the emission line then allows us to determine the distance of the BLR to the core of the AGN which together with the angular velocities of the BLR provide an estimate for the mass of the SMBH (see the 'reverberation mapping' described above). Absolute fluxes of some of these emission lines (e.g. [O III]) also seem to be correlated with the bolometric luminosities of the AGN (Dicken et al. 2009) while high resolution ($R \gtrsim 5000$) spectra enable us to have an even finer description of the AGN phenomenon. Nevertheless, we here concentrate on medium to low resolution spectral resolutions like those covered by Gaia.

Interpreting the broadening of these lines as coming from the Doppler effect owing to their angular velocities, one might intuitively expect that the BLR (with velocities $v \gtrsim 0.01c$) is situated in the near vicinity of the SMBH, while the narrow line region (NLR) is situated further away. Assuming the latter assumption is true, we then deduce that most of the BLR might be hidden to our view upon a favourable alignment of the obscuring torus with respect to the line-of-sight towards the core of the AGN. This elegant formulation, depicted in figure 1.14, yields an efficient (though incomplete) explanation of the various features noticed amongst the different types of AGN while using a minimal set of parameters (i.e. the viewing angle). A major support to this theory comes from the spectropolarimetry of type II AGN. Indeed, while the BLR is hidden to our view by the obscuring torus, its light can get reflected by charged particles standing at relatively large distances from the core, these hence act as mirrors towards the central source. In such a case, charged particles will capture photons that they will re-emit with a significant degree of polarization through Thompson scattering. Accordingly, type II AGN should then appear with broad emission lines if observed in polarimetry. Using this technique, Antonucci & Miller (1985) effectively showed that the type II AGN NGC1068 consists in a type I AGN whose BLR is hidden to our view. On the other hand, 60% of type II AGN do not show such broad emission lines in polarized light (Wu et al. 2011) whence the potential existence of 'true' type II AGN lacking the BLR. According to the recent work of Trump et al.

[10] Supposing $\Omega_m = 0.26$, $\Omega_\lambda = 0.74$ and $H_0 = 72$ km s^{-1} Mpc^{-1}

(2011), the presence of the BLR might depend on the accretion rate of the AGN given that only high luminosity AGN exhibit broad emission lines in polarized light.

The relativistic jet

As already pinpointed in section 1.2.2, the physical mechanisms powering the relativistic jet and the underlying magnetic fields are presently poorly understood and are still subject to intense debates. Still, the preferred scenario implies the Blandford-Znajeck mechanism (Blandford & Znajek 1977) in which energy can be extracted out of rotating black holes (i.e. Kerr black holes) in the form of magnetic fields having strengths that are compatible with observations. Nevertheless, this scenario comes along with some unresolved issues and questions like the potential intervention of the accretion disk in the jet formation, the observation of disordered magnetic fields while geometrically regular shapes are expected or knowing whether the jet is accelerated in the innermost part of the SMBH or is gradually accelerated up to larger distances. At the opposite to the radio emission, where charged particles can travel up to a few Mpc while emitting in this domain, the optical to X-ray radiation are restricted to a few kpc to the central source owing to the loss of energy of their emitting electrons. Given that optical and X-ray radiations are commonly observed at larger distances, other sources of radiation are envisaged like those standing in the shocks produced upon some flares occurring in the jet which together with the relativistic beaming allow us to explain the presence a such short wavelength radiations at such high distances. Finally, extended jets occur in about 10% of the observed AGN. The ignition of such jets may hence depends either on the intrinsic spin of the SMBH (i.e. its angular momentum) or on evolutionary effects like those observed in the transition between the high/soft state of X-ray binary black holes over much longer time scales. A better understanding of the physics at work in the jet of AGN will mainly pass through the improvement of the computer facilities as well as through dedicated magnetohydrodynamic simulations (Netzer 2015).

The outflow of matter

Some 50% of AGN exhibit narrow absorption lines (NAL) with FWHM of $\sim 10^3$ km s^{-1} in their optical spectra[11] as well as an absorption feature in their X-ray spectra corresponding to a warm ($T \approx 10^6$ K) absorber as depicted in figure 1.10. Furthermore 10% of these AGN also show BAL with FWHM of $\sim 3 \times 10^4$ km s^{-1} in their optical spectra as illustrated in figure 1.8. An empirical model explaining the appearance of these absorption features in the most synthesized way was provided through the work of Elvis (2000). In this model, warm gas rises vertically in a small range of radii from the accretion disk. This flow is then bended outwards by the radiation pressure coming from the continuum emitting region and is accelerated from velocities comparable to those of the BLR up to velocities corresponding to those observed in BAL QSOs. Given the far different velocities observed in BAL and NAL, one has to assume that the vertical rise of the flow has a finite extension such that if the flow is seen from below, NAL will be produced from this vertical extension whereas if the flow is seen from the edge it will produce BAL features. No absorption occurs if this flow is seen from above. This ultimately yields a structure having the shape of a double funnel joined by the tip as sketched in figure 1.14. From empirical considerations, the opening angle of the cone formed this way should stand around 60° so as to have an

[11] These NAL should however not be confused with the Lyα forest already described in section 1.2.2

even ratio of absorbed to unabsorbed AGN. Also, owing to the fact that the radiation pressure accelerating the flow is not parallel to the latter, it will have a non negligible extent of 6°, compatible with the 10% of observed BAL QSOs. Finally, from X-ray observations, the medium constituting this flow must be hot ($T \approx 10^6$ K), dense ($n_e \approx 10^9$ cm^{-3}) and highly ionized. These properties allow the warm highly ionized medium (WHIM) to confine clouds from the BLR into the flow. If we further suppose this outflow to be Thompson thick then it can straightly explain the polarization that is observed in the optical continuum of BALs. Interestingly, no BAL QSOs are found to exhibit extended radio emission while very few show compact radio emission (e.g. Becker et al. 2000). An elegant solution to this issue advanced by Elvis (2000) is that the relativistic jet might be the re-collimation of the outflow of matter and both might hence constitute a unique structure.

Additional components

It would be presumptuous to say that all observations that would be made out of AGN can be explained by the sole aforementioned components. Still assuming that these components provide a fair description of the innermost part of the AGN, we will here mention some external components that may affect the view we have of an AGN. First, the extinction by dust in our Galaxy plays a significant role as it can redden the observed SED by an amount of $E_{B-V} \approx 0.5$ mag leading to a dimming of the observed magnitude in the V band of $A_V \approx 1.5$ mag at $|b| \approx 20°$. The latter can however be partially rectified thanks to the use of dedicated maps like the one of Schlafly & Finkbeiner (2011). Similarly, the extinction by dust in the host galaxy of the AGN should also play a significant role though the latter is hardly constrained and is getting even more complicated given that most of these galaxies are unresolved. While these extinctions dim the recorded flux, stars in the bulge of the host galaxy conversely add a parasitic signal to the observed SED, this contamination being even more severe upon star formation episodes. As briefly suggested here, strong connections exist between AGN and their host galaxies, the former being presumably supposed to ignite star formation episodes while the mass of the SMBH seems to correlate with the mass of the hosting galaxy. All these connections are extensively reviewed in Beckmann & Shrader (2012, chapter 7).

A last point we would like to address is the effect of evolution on AGN. From our local universe, for example, we know that merging of galaxies are frequent and might consequently impact the view we have of AGN. This point is even more critical given that we are looking back in time when the Universe was denser and where most of the gas was not yet converted into stars whence the following open question: can the relatively regular structure of AGN described so far be applied to such mergers? On the other hand, a redshift-dependent analysis provides us with insights about the evolution of AGN (e.g. Kochanek et al. 2012). The luminosity function (LF) of the AGN stands to be the best-known tool of evolutionary tracking as it compares the density of AGN of a given range of bolometric luminosities in terms of the redshift, or equivalently in terms of the lookback time. The general trends, already outlined in the early work of Schmidt (1968), is that very luminous AGN (i.e. quasars) are found to be more common in the early universe than in the present days with a peak of presence around $2 < z < 3$. Low luminosity quasars however, were found to peak at $z < 1$ (Croom et al. 2009), meaning that the most massive and luminous quasars were accreting in the early universe while the least powerful ones are more active in the present time, that is *cosmic downsizing*. Beside these shifted activity periods, X-ray spectra of AGN taken over a wide range of redshift harbour similar shapes(Saez et al.

2011), meaning that the way their central engine works is much the same irrespectively of this redshift.

Validity of the model

The model that is presently described accounts for most of the diversity encountered amongst AGN, however some observations do not straightly fit in this scheme. We will mention here some of these counter-examples while formulating the basis of the possible extensions that can be brought to the model in order to account for these incompletenesses. Let us still mention, that besides these very few annoying observations, this unified model of AGN is empirically well supported as detailed along this section. Though, we have to admit that it lacks of physical explanations about the source of matter feeding the SMBH, about the stable presence of an obscuring torus over a long time scale, about the mechanism giving rise to the relativistic jet or about the origin of the vertical outflows producing the NAL and BAL features, for examples. Similarly, the structure and geometry of most of its constituting components are currently not well settled (e.g. is the obscuring torus a simple clumpy toroid? Does the hot corona have a spherical or disk shape? ...) and the barrier between some components is still blurry so that these may be merged into single constituents (e.g. the aforementioned relation between the relativistic jet and the outflow of matter).

One of the major pitfalls of the current model ironically comes as the counterpart of one of its most successful predictions, namely the already mentioned fact that 40% of type II AGN encompass broad emission lines in their polarized light. These being mainly present in highly luminous AGN, one might wonder whether the 60% of AGN having no trace of the BLR in their polarized light are true type II AGN or if these simply fail to sufficiently illuminate the BLR in order for the latter to be reflected. Both hypotheses explain equivalently well type II AGN for which no absorption is observed in the soft X-ray domain (Pappa et al. 2001). Another incompatibility rising from our model stands from the fact that most luminous AGN in the hard X-ray domain are less frequently seen with absorption in the soft X-ray than the low luminosity ones (Beckmann et al. 2009). A potential explanation of this phenomenon is the scenario of a 'receding torus' first proposed by Lawrence (1991). In this scenario, the distance of the obscuring torus to the center of the system depends on the power of the AGN as the former will be expelled by the radiation pressure of the core. As a consequence, luminous AGN should have more distant obscuring torus and will hence appear as less frequently obscured. Through these examples, we might expect that a more complete model should at least encompass the bolometric luminosity (i.e. the power) of the AGN as one of its parameters as well as the spin of the SMBH that might explain the production of the relativistic jet.

1.2.4 Cosmological applications

Cosmological models

With the advent of modern physics, our view of the cosmos and in a broader sense of the world we live in has dramatically changed. While Newtonian mechanics considered time and space as an absolute framework in which the law of motions are taking place, special relativity taught us that space and time are intimately related as, for example, a clock that is moving with respect to an observer will tick more slowly when compared to the latter one. Similarly, an object moving relatively to this observer will see its shape contracted in the direction of motion. Though very counter-intuitive, these examples have been verified by an overwhelming number of experiments, none of

them contradicting the theory. This theory of special relativity is based on two simple postulates:

> **The principle of relativity**: The laws of physics must be the same in all inertial reference frames[12].
>
> **The constancy of the speed of light**: The speed of light has the same value, c, in all inertial frames, regardless of the velocity of the observer or the velocity of the source emitting the light.

Interested readers may find in Jewett & Serway (2008), a very didactic and complete review of the consequences of the two upper postulates, amongst which the famous equivalence between rest frame energy and mass, $E = mc^2$.

Through the postulated equivalence between the gravitational and inertial mass, or in other words the stated fact that the gravitational 'force' cannot be differentiated from the factitious force that would be experienced in an accelerated frame, Albert Einstein derived the basis of general relativity through its fields equations. This set of ten differential equations is a mathematically complex theory that tells spacetime (i.e. the combination of space and time into a four dimensional continuum) how to curve in presence of massive objects while the curvature of this spacetime tells objects (massive or not) how to move.

A straightforward solution to these fields equations stands in the Minkowski space where an infinitesimal change in spacetime, $\mathrm{d}s$, can be expressed as

$$\mathrm{d}s^2 = -c^2\mathrm{d}t^2 + \mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2 \tag{1.8}$$

where $\mathrm{d}x$, $\mathrm{d}y$, $\mathrm{d}z$ are the classical infinitesimal changes in three dimensional Cartesian coordinates and $\mathrm{d}t$ the infinitesimal change in time. Equation 1.8 immediately results from Special Relativity, in case of the absence of (strong) gravitational fields. Note that based on the constancy of the speed of light, $c^2\mathrm{d}t^2 = \mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2$, we will then have that photons travel along null geodesics (i.e. straight lines in spacetime) meaning that $\mathrm{d}s = 0$. Minkowski space has a very limited impact on cosmology but it still provides a pedagogical transition from the Euclidian space having a metrics given by $\mathrm{d}s^2 = \mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2$ to the expanding spacetime we will next consider.

Based on the observation that the Universe looks identical on large scales in all viewing directions, meaning that the Universe is isotropic and homogeneous on large distance scales (that is the *cosmological principle*), Alexander Friedmann derived the equations of motion of an expanding universe whose metrics is given by[13]

$$\mathrm{d}s^2 = -c^2\mathrm{d}t^2 + a^2\left(\mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2\right), \tag{1.9}$$

where a is the unitless time-dependent scale factor of the Universe, conventionally $a(t_0) = 1$ at present time. The derivation of these equations is rather lengthy and complicated as these rely on the Einstein's fields equations. Still the interested readers may find in Schneider (2006) a very pleasant and comprehensive way of deriving these equations whose starting point is the description of an isotropic and homogeneous universe in the framework of Newtonian gravity along with a very few assumptions from

[12] a frame that is either at rest or moving at a constant velocity along a straight line

[13] For the sake of clarity, we solely consider here the case of a flat universe (i.e. $\kappa = 0$) that is also well supported by observations

General Relativity. Accordingly, we simply provide here the results of this derivation,

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho + \frac{\Lambda}{3} \tag{1.10}$$

$$\left(\frac{\ddot{a}}{a}\right)^2 = -\frac{4\pi G}{3}\left(\rho + \frac{3P}{c^2}\right) + \frac{\Lambda}{3} \tag{1.11}$$

where ρ and P are the mass-energy density and pressure of the Universe respectively and where Λ is the cosmological constant we will soon describe. Basically, both these equations tell us how fast the expansion of the Universe is taking place when compared to the scale factor, a, at a given time t as well as how this expansion accelerates/decelerates. The left hand side of equation 1.10 is hence naturally related to the Hubble parameter, H, through

$$H = \frac{\dot{a}}{a} \tag{1.12}$$

where $H_0 \equiv H(t_0)$ and where H is hence not necessarily constant. From the individual components that are presently known in the Universe, one can relate their respective densities to the pressure they exert. Accordingly, equations 1.10 and 1.11 can be written in terms of the density of matter-energy ρ_m in the Universe, of the radiation density ρ_r and of the density of vacuum energy from quantum mechanics ρ_v. Furthermore, we can consider the cosmological constant Λ as emanating from a hypothetical dark energy whose density is given by $\rho_\Lambda = \Lambda/(8\pi G)$. This cosmological constant was first inserted by Einstein as a way to enable its fields equations to describe a static universe (i.e. with $\dot{a} = 0$) which existence was predominantly postulated at that time. The discovery of the expansion of the Universe by Lemaître (1927); Hubble (1929a) dismissed the use of the cosmological constant though the recent possibility that the Universe may encounter an accelerated expansion (Riess et al. 1998) can give it a second birth. By defining the critical density as

$$\rho_c = \frac{3H^2}{8\pi G}, \tag{1.13}$$

that is the total density of the Universe at time t, we get from equation 1.10 that

$$\Omega_0 \equiv \Omega_m + \Omega_r + \Omega_\Lambda = 1. \tag{1.14}$$

with $\Omega_m = \rho_m/\rho_c$, $\Omega_r = \rho_r/\rho_c$, $\Omega_\Lambda = (\rho_\Lambda + \rho_v)/\rho_c$. Note that equation 1.11 can be re-written according to these normalized densities though we do not present it here as we voluntarily skip the lengthy description of the equivalence between the density and pressure of the various components of the Universe. A still convincing way to understand this equivalence is to consider the first law of thermodynamics stating that the change in the internal energy of a system, $\dot{E}$, is compensated by a change in the volume of the system, $\dot{V}$, according to pressure, P, through $\dot{E} + P\dot{V} = 0$. From special relativity, we have that $E = mc^2 = \rho V c^2$ such that $P = -\dot{E}/\dot{V} = -c^2 \mathrm{d}(\rho V)/\mathrm{d}V$. Given that $V \propto a^3$, we hence have that $P = -c^2 \mathrm{d}(\rho a^3)/\mathrm{d}a^3$. The claimed relation then comes from the derivative standing on the right hand side of the latter equation which obviously depends on the considered component.

Although we only considered flat models of universe, we can already draw some general conclusions out of the previous discussion:

> **The big bang**: given that the right hand side of equation 1.10 is always greater than zero, we have that a is a monotonically increasing function of time. In other words, at some point in the past we should have $a \ll 1$. That is, the Universe used to be denser and hotter in the past (i.e. $a \to 0$ as $t \to 0$).

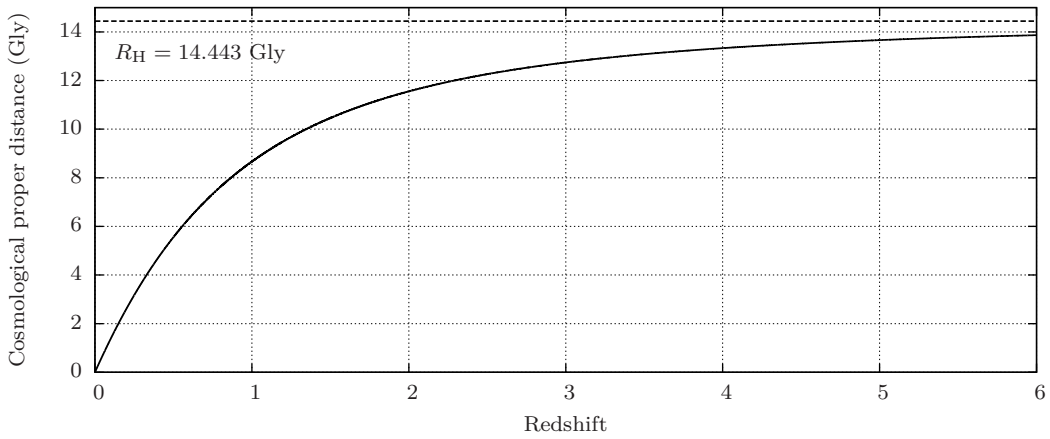

Figure 1.16: Proper distance as a function of redshift in the case of a constant Hubble parameter $H(t) = H_0 = 67.7$ km s^{-1} Mpc^{-1}. Under this assumption the age of the Universe is given by the Hubble time, $H_t = 1/H_0$, that consequently also does match with the largest possible cosmological proper distance we can observe if we expressed it in ly. This corresponds to the Hubble radius, R_{H}, as depicted on the graph by the dotted line.

The cosmological parameters: the past and future of the expansion of the Universe can be determined through the sole knowledge of H_0, $\Omega_{m,0} \equiv \Omega_m(t_0)$, $\Omega_{r,0} \equiv \Omega_r(t_0)$ and $\Omega_{\Lambda,0} \equiv \Omega_\Lambda(t_0)$. Practically, since $\Omega_{r,0} \ll 1$, the latter term is often neglected in the literature.

Distances and redshift in cosmology

AGN, and more particularly quasars, due to their large distances and high luminosities play a crucial role in the parametrization and validation of the cosmological model we just described. From the discovery of the expansion of the Universe (Lemaître 1927; Hubble 1929a), a linear relation was established between the distance of galaxies, d, and their radial velocity, v, such that

$$v = H_0 d. \tag{1.15}$$

Supposing, for the moment, that these velocities are due to the relativistic Doppler shift of the source, we then have that the distance of these objects can be inferred from their observed redshift, z, through equation 1.15 while using the fact that

$$\beta = \frac{v}{c} = \frac{(z+1)^2 - 1}{(z+1)^2 + 1}, \tag{1.16}$$

or in case where $v \ll c$, $v = cz$. From our previous discussion, and more particularly from equation 1.9, we know however that these apparent velocities come from the spacetime expansion of the Universe rather than from an intrinsic velocity of the sources whence the need to reconsider the notion of distance and velocity. The cosmological proper distance[14] is defined as the length of the geodesic connecting two points of the spacetime at a given time t (i.e. we freeze t and then measure the euclidean

[14] Cosmological proper distance should not be confused with the proper distance and proper length from special relativity.

distance). From equation 1.9 and because both measurements are made at a fixed time t, we have that $-c^2\mathrm{d}t^2 = 0$ such that $\mathrm{d}s = a\,\mathrm{d}r$ with $\mathrm{d}r^2 = \mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2$. The *cosmological proper distance* being hence given by

$$d_p = ar \tag{1.17}$$

where r is defined as the *comoving distance* between both points. Because the cosmological proper distance depends on time, we can define a cosmological proper velocity that would be due to the expansion of the Universe as

$$v_p = \frac{\mathrm{d}d_p}{\mathrm{d}t} = \dot{a}r = \frac{\dot{a}}{a}d_p = Hd_p. \tag{1.18}$$

By comparing the latter equation with equation 1.15, we straightly see that the distance that would be inferred from the latter equation roughly corresponds to the cosmological proper distance. One should however pay attention to the fact that the cosmological proper distance assume that both measurements are done in a synchronous way, which is obviously not the case in equation 1.15 as the light emitted from the source takes a finite amount of time to reach the observer. Equation 1.15 then only strictly holds in the case where $H_0 = H(t)$ is constant and provides a good approximation to d_p in case where $z \ll 1$ (assuming that H is smoothly varying). The assumption of the constancy of the Hubble parameter still provides us with a rather good approximation of the cosmological proper distances that are used in figure 1.16 such as to stress the relation existing between the redshift and the cosmological proper distance without having to rely on complex integrations implying equations 1.10 and 1.11. The Hubble time that is derived there ($H_t = 14.4 \times 10^9$ years) corresponds to the age of the Universe as if the latter had encountered a constant expansion and is effectively of the same order of magnitude as the age that would be derived by more complex considerations ($\sim 13.8 \times 10^9$ years).

Now, if we consider an emitter and an observer separated by an infinitesimal comoving distance dr, then the shift that is recorded by the observer due to the expansion of the Universe is given by

$$\mathrm{d}z = \frac{\mathrm{d}\lambda}{\lambda} = \frac{\mathrm{d}v}{c} = \frac{H}{c}\mathrm{d}r = H\mathrm{d}t = \frac{\mathrm{d}a}{a}, \tag{1.19}$$

where we used the fact that $z = v/c$ as $v \ll c$ in this case. From the latter equation, we get $\mathrm{d}\lambda/\lambda = \mathrm{d}a/a$. Supposing an infinite number of such observers along the path joining the source and the final observer and integrating over this path yield $\lambda = Ca$ where C is an integration constant related to the wavelength that is presently observed λ_{obs} (at $a = 1$). We hence have that $\lambda = a\lambda_{\mathrm{obs}}$ such that, from the definition of the redshift,

$$z + 1 = \frac{\lambda_{\mathrm{obs}}}{\lambda} = \frac{1}{a}. \tag{1.20}$$

The measured redshift is hence directly related to the cosmological scale factor at the epoch of emission through the latter equation. This relation is of immense importance as the redshift is often the sole quantity we are able to measure.

Quasars in cosmology

The major role of quasars regarding cosmology is to trace the large scale distribution of matter in the Universe through deep statistical analyses. Indeed, as we suppose the Universe to be isotropic and homogeneous, any noticed over-densities of an idealistic kind of objects that would have been present from the very early phase of the Universe

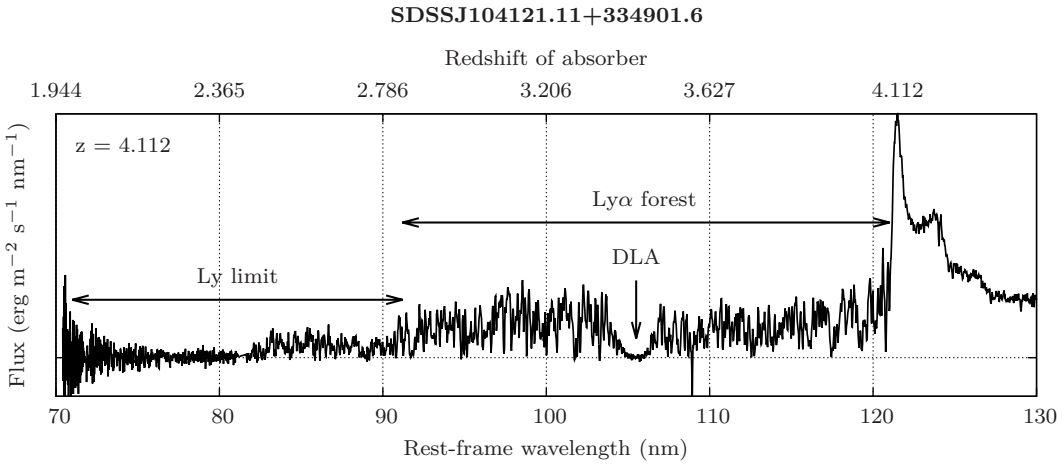

Figure 1.17: Lyman absorbers encountered along the line-of-sight towards the high redshift quasar SDSSJ104121.11+334901.6. The observed absorption lines from the Lyα forest are due to neutral hydrogen having column densities of $N_{\mathrm{H\,I}} < 10^{17}$ cm^{-2}. If this column density exceeds 10^{17} cm^{-2}, then this neutral hydrogen becomes optically thick to photon having wavelengths shorter than the Lyman limit at 91.2 nm while if this density exceeds 10^{20} cm^{-2}, these constitute the so-called Damped Lyα Absorbers (DLA) that are presumably associated with forming galaxies.

while not evolving, would be directly related to the Hubble parameter, H, as well as to all the aforementioned cosmological parameters at the time when these objects emitted their lights. Given that quasars are not these hypothetical objects (e.g. the local universe at $z \lesssim 0.04$ does not contain any quasar), this stresses the need to have a thorough understanding of what quasars are, on how they form and on how they evolve. Accordingly, the cosmological parameters undoubtedly have their imprints on the luminosity function presented at the end of section 1.2.3. Also, given its statistical nature, a very large amount of observations will be preferred in order to confidently constrain these cosmological parameters. One such well-known study is the Baryon Oscillation Spectroscopic Survey (BOSS) of the SDSS that recently allowed to highlight the inhomogeneities that were present in the very early phase of the Universe (Delubac et al. 2015).

From their large redshifts, quasars also have a high probability to see their light being attenuated by intervening gas and galaxies. The spectral signature of these absorbers being hence imprinted to the quasar SED at the various redshifts these absorbers are encountered, that is bluewards to the Lyα emission line. The intergalactic medium being mainly composed of hydrogen, we expect the neutral hydrogen to be the principal cause of such absorptions. Figure 1.17 depicts the three most common sources of external absorption in the SED of (high redshift) quasars. First the Lyα forest consists of sharp absorption lines that can be attributed to diffuse neutral hydrogen having column densities $N_{\mathrm{H\,I}} < 10^{17}$ cm^{-2}. The structures encompassing such column densities are supposed to have sizes that are by far larger than the typical size of galaxies, that is larger than a few hundreds kpc. Regions having column densities higher than 10^{17} cm^{-2} will be furthermore optically thick (i.e. opaque) to photons having energies $E \geq 13.6$ eV corresponding to the Lyman limit standing at 91.2 nm. Finally, regions having column densities of 10^{20} cm^{-2} are supposed to be associated with proto-galaxies and forming galaxies. These typically produce deep and broad absorptions that are similar in shape to the already encountered BAL absorptions

(i.e. with a Lorentz profile) while being unequivocally identified by their detachment from any emission lines. Note that these damped Lyα absorbers (DLAs) further allow the kinematics of the early galaxy type to be studied in detail. A last point we would like to address is the absorption encountered at very high redshift, where cosmological models predict that the intergalactic medium (IGM) was mainly composed of neutral hydrogen, that is hence opaque to wavelengths of 121.5 nm. Quasars situated at redshifts higher than this re-ionization epoch hence encountered a continuous absorption of the Lyα wavelength from the redshift at which they emitted their light up to the redshift where this opacity vanished through the re-ionization of the IGM, that is the so-called Gunn-Peterson trough (Gunn & Peterson 1965). Quasars exhibiting this trough are hence the witnesses of the late stages of the re-ionization process that fenced the 'dark ages' of the Universe. The work of Becker et al. (2001) first identified this Gunn-Peterson trough in a $z = 6.28$ quasar suggesting that the re-ionization epoch ended around $z \approx 6$.

Based on ideas already described in the reverberation mapping techniques (see section 1.2.3), Watson et al. (2011) studied the time lag, τ, between a rise of the optical/UV continuum and a rise in the broad emission lines. He postulated that the size of the BLR, $r = \tau/c$, can be attributed to the ionizing flux coming from the central engine which acts as the inverse squared law of the distance. Accordingly, the squared size of the BLR must be proportional to the intrinsic luminosity of the AGN, $L \propto r^2$, such that the observed flux, F, is related to the cosmological luminosity distance through $D_{\rm L} \propto \tau/\sqrt{F}$. Besides these straight applications, AGN stand at the edge of our knowledge as these combine quantum mechanics as well as general relativity in a single object that is the SMBH. The raise of a new kind of physics might hence emerge from the thorough study of these objects and through the potential observable effects that this new physics might have on AGN.

1.3 Strong gravitational lensing

The previous section discussed the nature of AGN while particularly focusing on quasars. We saw that their extreme luminosities allow these to be seen at very high redshift corresponding to very large distances and hence to epochs when the Universe was still very young. We concluded this section by assessing that these large redshifts also imply a greater chance to encounter some intervening object(s) along the line-of-sight. Finally, we briefly discussed how the spacetime is curved by the presence of massive objects and how this curvature acts on objects through gravity irrespectively of whether these are massive or not. Gravitational lensing (GL) is the place where all these concepts meet. More particularly, strong GL describes the production of multiple images of a single luminous background source whose light rays are deflected owing to the presence of a massive object standing in the near vicinity of the line joining the observer and the source. Note that because of the very small angular separation between the lensed images ($\Delta\theta \ll 1''$) resulting from the GL of a background source by stars, these will often appear as unresolved using current technologies. The recorded flux being however amplified thanks to the presence of these multiple (unresolved) images as well as to the inherent amplification occurring within GL, as we will see. The study of such magnifications/amplifications due to GL is part of the weak and micro GL that will not be considered here. Accordingly, only GL whose deflector is a very massive object (e.g. a galaxy) and for which some lensed images can be unambiguously separated will be discussed here. In the following, GL then implicitly refers to strong GL. This section and the mathematical derivations

it contain are mostly based on the book of Schneider et al. (2006) as well as on the reference didactic report of Refsdal & Surdej (1994).

Short before the final formulation of the general relativity, Albert Einstein derived the deflection angle that would be encountered by a light ray coming from a distant source and passing in the vicinity of a massive compact object as

$$\hat{\alpha}(\xi) = \frac{4GM}{c^2\xi} \tag{1.21}$$

where ξ is the impact parameter of the incoming light ray corresponding to the perpendicular distance between the line joining the observer to the center of mass and the path of the incoming light ray and where a massive compact object means $r_s = 2GM/c^2 \ll \xi$, thus implying that $\hat{\alpha}(\xi)$ is very small. The value from equation 1.21 appears to be exactly twice the value that was previously calculated by Johann von Soldner in 1804 using the Newtonian theory of gravitation. Benefiting from a total solar eclipse on May 29, 1919, Sir Arthur Eddington noticed a deflection angle of the stars close to the limb of the sun of $1.75''$ which is exactly the angle predicted based on general relativity through equation 1.21. This was the first out of many predictions that yield a wide acceptance of Einstein's theory of gravitation. Later, in 1936, Einstein published a note [15] on the angular size of the ring that would be produced by the lensing of a star by another star upon a perfect alignment of the two stars as seen by an observer situated at a distance D from the nearest star. He found that this angular size is given by $\theta_e \propto 1/\sqrt{D}$. He consequently stated: *"Of course, there is no much hope of observing this phenomenon directly."*. One year later Fritz Zwicky pointed out that nebulae (i.e. galaxies) would produce a much more efficient deflector and stated in turn: *"The probability that nebulae which act as gravitational lenses will be found becomes practically a certainty."*. It was no more than forty years later that the first GL, a double image of a background quasar at $z = 1.405$ with image separation of $\Delta\theta = 5.7\ ''$, was discovered. To date, more than 100 GLs are known, most of them being listed in the CfA-Arizona Space Telescope LEns Survey of gravitational lenses (CASTLES) database as well as in the Sloan Digital Sky Survey Quasar Lens Search (SQLS) catalogue (Inada et al. 2012). Gaia from its impressive angular resolution is expected to discover $\sim$ 1600 GLs amongst which 80 have more than two lensed images (Finet & Surdej 2016).

The present section will outline the theoretical basis of GL in section 1.3.1 while the specific model that is being used in the field of the present study is further described in section 1.3.2. Finally, some applications of GLs are reviewed in section 1.3.3.

1.3.1 Introduction to gravitational lens theory

Theory of GL, though being based on general relativity, can be easily interpreted in terms of simple geometrical concepts as well as from equation 1.21. In the following, the reader is referred to figure 1.18 for a graphical view of the lensing phenomenon designed to ease the understanding of the presented concepts. Let us consider now two dimensional coordinates in the lens plane $\boldsymbol{\xi} = (\xi_1, \xi_2)$ whose origin is the center of mass of the deflector, D. We will first start by enunciating some working hypotheses that allow us to efficiently derive the lens equation, namely:

The weak fields approximation Considering galaxies as deflectors, we will have that the distances between the incoming light rays and the center of mass

[15] This investigation was done on the request of a skilful amateur scientist, Rudi Mandl, who was at that time dishwasher in a restaurant

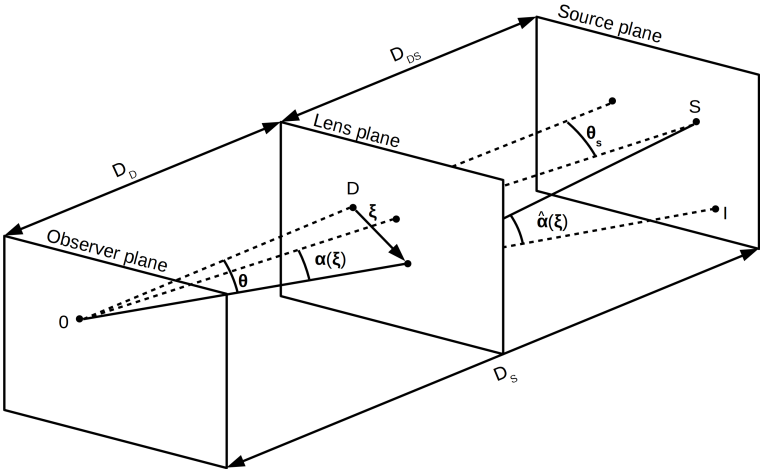

Figure 1.18: Geometry of the gravitational lensing phenomenon. A light ray coming from a distant source, S, and passing in the vicinity of a massive deflector, D, is deflected by an angle $\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi})$ where $\boldsymbol{\xi}$ is the impact parameter of the light ray taken in the lens plane. An observer, O, standing in the path of the deflected light ray hence sees an angular separation $\boldsymbol{\theta}$ between the image of the source, I, and D. The lens equation being then given by $\boldsymbol{\theta_s} = \boldsymbol{\theta} - \boldsymbol{\alpha}(\boldsymbol{\xi})$ where $\boldsymbol{\theta_s}$ is the angular separation between D and S (i.e. the position of the source if no deflector were present).

of the deflecting galaxies will be in the order of a few dozen kpc. In this regime of weak fields, the Einstein fields equations can be linearised such that the deflection angle of an ensemble of massive points will be equal to the vector sum of their individual contribution.

The geometrically thin lens approximation Given that the distances involved in GL are in the order of a few Gpc while the deflection of the light rays only occurs in the near vicinity of the deflecting galaxy, roughly corresponding to its size, we can assume that this deflection occurs instantaneously in the lens plane. From this assumption, we have that the three dimensional mass density of the deflector, $\rho(\boldsymbol{\xi}, x)$ where x runs along the line joining the observer to the center of mass of the deflector, can be squeezed onto the lens plane in order to provide the *projected mass density* defined as

$$\Sigma(\boldsymbol{\xi}) = \int \rho(\boldsymbol{\xi}, x) dx \tag{1.22}$$

The small angle approximation Equation 1.21 considers the incoming light ray as parallel to the line-of-sight towards the deflecting mass. However, most astronomical objects emit their light radially which is in contradiction with the previous assumption. Hopefully, given the cosmological distances at which these objects stand we may still consider the incoming wavefront of photons as flat. Similarly, the various angles that will be discussed in the following are typically very small (e.g. $1.75\ ''= 8.484 \times 10^{-6}$ rad) such that we can assume that $\tan\alpha \approx \sin\alpha \approx \alpha$ are valid approximations.

Along with the previous assumptions and the generalization of equation 1.21 so as to account for a vectorial impact parameter (i.e. $\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = 4\pi G M c^{-2} \boldsymbol{\xi}/|\boldsymbol{\xi}|^2$), we can

immediately write the equation of the *deflection angle* as

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = \frac{4G}{c^2}\int_{\mathbb{R}^2}\Sigma(\boldsymbol{\xi'})\frac{\boldsymbol{\xi}-\boldsymbol{\xi'}}{|\boldsymbol{\xi}-\boldsymbol{\xi'}|^2}\mathrm{d}\boldsymbol{\xi'}, \tag{1.23}$$

that is the deflection angle that would be obtained from an arbitrary mass distribution of the deflector projected in the lens plane at a given impact parameter $\boldsymbol{\xi}$. Note that the wavelengths of the incoming photons do not appear in the latter equation such that GL is achromatic. From the assumption of small angles, we have that $D_{\mathrm{S}}\boldsymbol{\alpha}(\boldsymbol{\theta}) \approx D_{\mathrm{DS}}\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi})$ with $\boldsymbol{\xi} \approx D_{\mathrm{D}}\boldsymbol{\theta}$ such that we can express the *scaled deflection angle* in terms of the classical deflection angle as

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \frac{D_{\mathrm{DS}}}{D_{\mathrm{S}}}\hat{\boldsymbol{\alpha}}(D_{\mathrm{D}}\boldsymbol{\theta}) \tag{1.24}$$

where D_{DS}, D_{S} and D_{D} are the distances between the observer plane, the lens plane and the source plane as sketched in figure 1.18. Further developing the right hand side of the latter equation allows us to describe the scaled deflection angle in terms of dimensionless quantities:

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \frac{1}{\pi}\int_{\mathbb{R}^2}\kappa(\boldsymbol{\theta'})\frac{\boldsymbol{\theta}-\boldsymbol{\theta'}}{|\boldsymbol{\theta}-\boldsymbol{\theta'}|^2}\mathrm{d}\boldsymbol{\theta'}, \tag{1.25}$$

where we have defined the *dimensionless surface mass density*

$$\kappa(\boldsymbol{\theta}) = \frac{\Sigma(D_{\mathrm{D}}\boldsymbol{\theta})}{\Sigma_{\mathrm{cr}}} \tag{1.26}$$

with Σ_{cr}, the *critical mass density*, defined as

$$\Sigma_{\mathrm{cr}} = \frac{c^2}{4\pi G}\frac{D_{\mathrm{S}}}{D_{\mathrm{D}}\,D_{\mathrm{DS}}}. \tag{1.27}$$

From trivial geometrical considerations, the *lens equation* can then be expressed as

$$\boldsymbol{\theta_s} = \boldsymbol{\theta} - \boldsymbol{\alpha}(\boldsymbol{\theta}) \tag{1.28}$$

where $\boldsymbol{\theta}$ is the vectorial angular separation between the center of mass of the deflector and the lensed image(s) and where $\boldsymbol{\theta_s}$ is the vectorial angular separation between the deflector and the source. Consequently, if the latter equation has multiple solutions, $\boldsymbol{\theta}$, for a given source position, $\boldsymbol{\theta_s}$, this means that multiple lensed images of the source will be produced. Though not developed here, we have to note that a mass distribution having $\kappa > 1$ at some point (i.e. $\kappa > \Sigma_{\mathrm{cr}}$) produces multiple images of some given source position whereas these will be visible only if a luminous object stands at this specific source position.

In first order approximation, the lens equation can be linearised so as to provide

$$\boldsymbol{\theta_s} \approx \boldsymbol{\theta_s}(\boldsymbol{\theta_0}) + \mathbf{A}(\boldsymbol{\theta}-\boldsymbol{\theta_0}) \tag{1.29}$$

where the Jacobian matrix $\mathbf{A} = \partial\boldsymbol{\theta_s}/\partial\boldsymbol{\theta}$ evaluated at $\boldsymbol{\theta_0}$ is called the *amplification or magnification matrix* for reasons that will shortly become evident. From equation 1.29, we have that any small extent from the source, $\Delta\boldsymbol{\theta_s} \equiv \boldsymbol{\theta_s} - \boldsymbol{\theta_s}(\boldsymbol{\theta_0})$, can be expressed as a linear transformation of the corresponding extent in the lensed image, $\Delta\boldsymbol{\theta} \equiv \boldsymbol{\theta} - \boldsymbol{\theta_0}$, through

$$\Delta\boldsymbol{\theta_s} = \mathbf{A}\,\Delta\boldsymbol{\theta}. \tag{1.30}$$

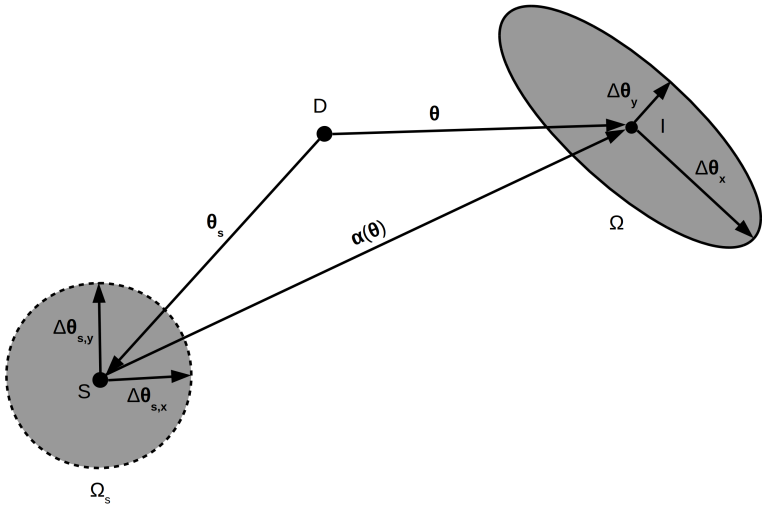

Figure 1.19: Illustration of a gravitationally lensed image, I, of a background source, S, by a massive deflector, D, as seen on the sky. In first order approximation, any small extent of the source, $\Delta\boldsymbol{\theta_s}$, can be mapped in the resulting image by a linear transformation of the form $\Delta\boldsymbol{\theta} = \mathbf{A}^{-1}\Delta\boldsymbol{\theta_s}$ where $\mathbf{A} = \partial\boldsymbol{\theta_s}/\partial\boldsymbol{\theta}$. The solid angle subtended by the lensed image is then given by $\Omega = \pi|\Delta\boldsymbol{\theta_x}||\Delta\boldsymbol{\theta_x}| = \pi|\mathbf{A}^{-1}\Delta\boldsymbol{\theta_{s,x}}||\mathbf{A}^{-1}\Delta\boldsymbol{\theta_{s,y}}| = \pi|\Delta\boldsymbol{\theta_{s,x}}||\Delta\boldsymbol{\theta_{s,y}}|\det\mathbf{A}^{-1} = \Omega_s\det\mathbf{A}^{-1}$ where Ω_s is the corresponding solid angle subtended by the source image. Note that the source itself will not be seen by the observer whence its dotted contour.

We hence have that a small circular extent from the lensed image can be mapped in the source image as an ellipse and conversely as illustrated in figure 1.19. Accordingly, we have that the solid angle subtended by an infinitesimally small circular extent of the source, $\mathrm{d}\Omega_s$, can be related to the solid angle subtended by the ellipse in the corresponding image, $\mathrm{d}\Omega$, through

$$\mathrm{d}\Omega_s = \det\mathbf{A}\,\mathrm{d}\Omega \tag{1.31}$$

as derived in figure 1.19. Given that the surface brightness of the source is conserved upon light bending by a massive object (Etherington 1933), we have that the *amplification or magnification factor* of the lensed image is equal to the ratio of the solid angle subtended by the lensed image to the corresponding solid angle from the source, that is

$$\mu = \frac{\mathrm{d}\Omega}{\mathrm{d}\Omega_s} = \det\mathbf{A}^{-1}. \tag{1.32}$$

Beware that $\det\mathbf{A}$ can become negative, this situation translating in the production of a mirror-symmetric image of the source. The absolute value of μ must hence be considered.

Let us finally note that from simple vector calculus, we have that $\nabla\ln|\boldsymbol{x}| = \boldsymbol{x}/|\boldsymbol{x}|^2$ where $\ln|\boldsymbol{x}|$ is the natural logarithm of the absolute value of $\boldsymbol{x}$, such that we can rewrite the scaled deflection angle as

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \nabla\psi(\boldsymbol{\theta}) \tag{1.33}$$

where

$$\psi(\boldsymbol{\theta}) = \frac{1}{\pi}\int_{\mathbb{R}^2}\kappa(\boldsymbol{\theta'})\ln|\boldsymbol{\theta}-\boldsymbol{\theta'}|\mathrm{d}\boldsymbol{\theta'} \tag{1.34}$$

is the *dimensionless deflection potential* at angular coordinates $\boldsymbol{\theta}$. For convenience, the mapping from $\boldsymbol{\theta}$ to $\boldsymbol{\theta_s}$ can hence be seen as a gradient mapping caused by the deflection potential. Furthermore, using the Laplacian $\nabla^2 \ln|\boldsymbol{x}| = 2\pi\boldsymbol{\delta}(\boldsymbol{x})$ where $\boldsymbol{\delta}(\boldsymbol{x})$ is the two dimensional Dirac delta function[16] yields twice the dimensionless surface mass density in two dimensions, that is henceforth also called the *convergence* in this context,

$$\nabla^2\boldsymbol{\psi}(\boldsymbol{\theta}) = 2\boldsymbol{\kappa}(\boldsymbol{\theta}). \tag{1.35}$$

Along with the previous considerations, we have that the *amplification matrix* can be written as

$$\mathbf{A}_{ij}(\boldsymbol{\theta}) = \delta_{ij} - \frac{\partial\alpha_i(\boldsymbol{\theta})}{\partial\theta_j} = \delta_{ij} - \frac{\partial^2\boldsymbol{\psi}(\boldsymbol{\theta})}{\partial\theta_i\partial\theta_j} = \delta_{ij} - \Psi_{ij} \tag{1.36}$$

where in the last step we defined $\Psi_{ij} \equiv \frac{\partial^2\boldsymbol{\psi}(\boldsymbol{\theta})}{\partial\theta_i\partial\theta_j}$. Further developing equation 1.36 yields

$$\mathbf{A} = \begin{pmatrix} 1-\kappa-\gamma_1 & -\gamma_2 \\ -\gamma_2 & 1-\kappa+\gamma_1 \end{pmatrix} \tag{1.37}$$

where

$$\gamma_1 \equiv \frac{1}{2}\left(\Psi_{11} - \Psi_{22}\right), \qquad \gamma_2 \equiv \Psi_{12} = \Psi_{21}. \tag{1.38}$$

The *determinant of the amplification matrix* being hence given by

$$\det\mathbf{A} = (1-\kappa^2) - (\gamma_1^2 + \gamma_2^2). \tag{1.39}$$

1.3.2 The non-singular isothermal ellipsoid in presence of an external shear

In general, we have that the deflection angle resulting from an arbitrary mass distribution can be computed through the numerical integration of the projected mass density over the lens plane according to equation 1.23. Analytical expressions of the deflection angle (or equivalently of the deflection potential), when available, are however preferred as these often yield to straightly derivable quantities like, for example, an instantaneous calculation of the amplification factors. Regarding this point, an interesting class of GL models are the axially symmetric lens models that are characterized by $\Sigma(\boldsymbol{\xi}) = \Sigma(|\boldsymbol{\xi}|)$. Indeed, for such models, we have that the deflection angle is simply given by

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = \frac{4GM(|\boldsymbol{\xi}|)}{c^2|\boldsymbol{\xi}|^2}\boldsymbol{\xi} \tag{1.40}$$

where $M(\xi) = 2\pi\int_0^\xi \Sigma(\xi')\xi'\mathrm{d}\xi'$ is the projected mass enclosed in a circle of radius ξ centred around the axis of symmetry. Consequently, equation 1.40 allows us to easily derive the analytical solutions of the deflection angle for all models satisfying $\Sigma(\boldsymbol{\xi}) = \Sigma(|\boldsymbol{\xi}|)$. A simple model that fulfils the latter condition while enabling the basic properties of galaxies to be modelled is the singular isothermal sphere (SIS) lens model. This model supposes that the stars constituting the lensing galaxies behave like particles in an ideal gas of constant temperature (i.e. it is isothermal) so as to account for the flat rotation curves that are observed in those galaxies. From the equation of state of an ideal gas, we have that the pressure exerted by a spherical cloud of gas of radius r at constant temperature T is given by $P(r) = \rho(r)k_\mathrm{B}T/m_p$ where k_B is the Boltzmann constant and where m_p is the mean mass of the gas particles (i.e.

[16] $\boldsymbol{\delta}(\boldsymbol{x}) = \begin{cases} \mathbf{1}, & \text{if } \boldsymbol{x} = \mathbf{0} \\ \mathbf{0}, & \text{otherwise} \end{cases}$

stars in our case). The velocity dispersion of the stars along the line-of-sight towards the galaxy being related to T using the Virial theorem as $\sigma_v^2 \approx k_\mathrm{B}\,T/m_p$. Since we suppose the cloud of gas and the galaxies to be in hydrostatic equilibrium, we further have that $\mathrm{d}P(r)/\mathrm{d}r = -GM(r)\rho(r)/r^2$ such that

$$\rho(r) = \frac{\sigma_v^2}{2\pi G r^2}, \qquad \Sigma(\xi) = \frac{\sigma_v^2}{2G\xi}, \qquad M(\xi) = \frac{\pi\sigma_v^2}{G}\xi. \tag{1.41}$$

The attentive reader will notice here that we made a distinction between $M(r)$, the mass enclosed in a sphere of radius r and $M(\xi)$, the projected mass enclosed within the impact parameter of radius ξ. Substituting the latter value within equation 1.40 yields a deflection angle for the SIS model of

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = \frac{4\pi\sigma_v^2}{c^2}\frac{\boldsymbol{\xi}}{|\boldsymbol{\xi}|}, \tag{1.42}$$

that is a constant deflection angle pointing towards the center of mass of the deflector.

The SIS lens model we just described, despite being based on a physical model of the kinematics of the galaxies, suffers from two main drawbacks, namely the fact that its projected mass density diverges as $\xi \to 0$, whence the reason why this model is termed 'singular', and its inability to model asymmetric deflectors that are otherwise commonly observed amongst galaxies. The SIS model would effectively correspond to cases where the deflector is a spiral galaxy that is viewed nearly face-on which hence only covers a limited fraction of the observations. More complex models are hence required in order to accommodate for the observed asymmetry as well as to break the singularity at $\Sigma(\boldsymbol{\xi} = \mathbf{0})$ or equivalently at $\kappa(\boldsymbol{\theta} = \mathbf{0})$. Kormann et al. (1994) concurrently solved both issues by introducing the non-singular isothermal ellipsoid (NSIE) lens model through a straight generalization of the dimensionless surface mass density of the SIS model as given by

$$\kappa(\boldsymbol{\theta}) = \frac{\theta_e\sqrt{f}}{\sqrt{\theta_1^2 + f^2\theta_2^2 + \theta_c^2}} = \frac{\theta_e\sqrt{f}}{\zeta} \tag{1.43}$$

where f is the ratio of the minor axis to the major axis of the isodensity contours of the mass distribution ($0 < f \leq 1$), θ_c is the core radius of this mass distribution and where θ_e is a scaling factor termed the *angular Einstein radius*. The limiting case where $f = 1$, $\theta_c = 0$ corresponding then to the dimensionless mass distribution of the SIS model while $\kappa(\mathbf{0})$ has now a finite value if $\theta_c > 0$, whence the non-singularity of this class of model. The non-singular isothermal sphere (NSIS) lens model corresponding then to the case where $f = 1$, $\theta_c > 0$ while the singular isothermal ellipsoid (SIE) corresponds to the case in which we have $f < 1$, $\theta_c = 0$. The derivation of the normalized deflection angle associated to the NSIE model can be found in Keeton & Kochanek (1998) and results in

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) = \frac{\theta_e\sqrt{f}}{2f'}\left[\operatorname{arctanh}\left(\frac{f'\theta_1}{\zeta + f\theta_c}\right)\boldsymbol{e_1} + \arctan\left(\frac{f'\theta_2}{\zeta + \theta_c/f}\right)\boldsymbol{e_2}\right] \tag{1.44}$$

where $f' = \sqrt{1 - f^2}$ is defined as the *eccentricity* of the isodensity contours of the mass distribution and where $\boldsymbol{e_1} \equiv (1, 0)$, $\boldsymbol{e_2} \equiv (0, 1)$.

A last degree of realism can be added by considering the stretching of the lensed images owing to the presence of a massive object in the vicinity of the deflector (Kovner 1987). This *external shear* being then straightly given by

$$\boldsymbol{\alpha}_{\boldsymbol{\gamma}}(\boldsymbol{\theta}) = \gamma\mathbf{R}(\omega)\boldsymbol{\alpha}(\boldsymbol{\theta}) = \gamma\begin{pmatrix}\cos\omega & -\sin\omega \\ \sin\omega & \cos\omega\end{pmatrix}\boldsymbol{\alpha}(\boldsymbol{\theta}), \tag{1.45}$$

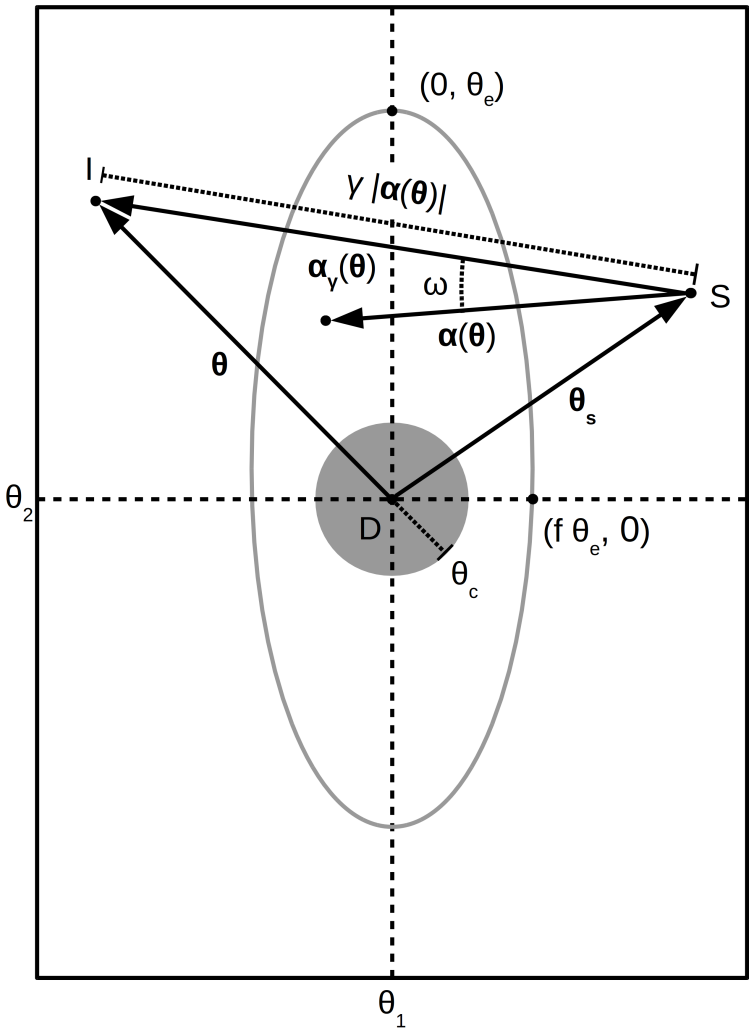

Figure 1.20: Illustration of the NSIEg lens model parameters projected on the sky with coordinates (θ_1, θ_2) and origin given by the center of mass of the deflecting galaxy D. The source and image position being respectively given by S and I, as usual. The NSIEg lens model is characterized by an isothermal mass distribution having an ellipse shape with an axis ratio of f. An isodensity contour being depicted here as a gray ellipse having a semi-major axis corresponding to the scale factor θ_e, called the Einstein radius. The singularity of the SIE model is solved by introducing a cut-off in the density profile corresponding to the angular radius of the core, θ_c, depicted here as a filled gray circle. These parameters allow us to compute the scaled deflection angle(s) $\boldsymbol{\alpha}(\boldsymbol{\theta})$ associated with a given source position $\boldsymbol{\theta_s}$. The presence of an external massive object has as an effect to introduce a shear of the image(s) translating into a counter clockwise rotation of the scaled deflection angle corresponding to the shear orientation ω and into a scaling of $\boldsymbol{\alpha}(\boldsymbol{\theta})$ according to the shear strength γ. The resulting deflection angle $\boldsymbol{\alpha_\gamma}(\boldsymbol{\theta})$ then provides the final image position through $\boldsymbol{\theta} = \boldsymbol{\theta_s} + \boldsymbol{\alpha_\gamma}(\boldsymbol{\theta})$.

Table 1.3: Parameters of the non-singular isothermal ellipsoid lens model in presence of an external shear (NSIEg). Image(s) positions, $\boldsymbol{\theta}$, can be retrieved from the lens equation using the scaled deflection angle provided in equations 1.44 and 1.45. The position of the deflector is here supposed to stand at the origin of the coordinate system.

Parameter	Description
$\boldsymbol{\theta_s}$	Source position
θ_e	Scaling factor (angular Einstein radius)
θ_c	Angular size of the radius of the core of the deflector
f	Axis ratio of the isodensity contour of the mass distribution of the deflector ($0 < f \leq 1$) where $f' = \sqrt{1-f^2}$ is the associated eccentricity
γ	Shear intensity
ω	Shear orientation

where γ is a unitless constant representing the *shear strength* and where $\mathbf{R}(\omega)$ is a rotation matrix whose parameter ω is the *shear orientation*. We summarized the parameters of the NSIE lens model in presence of an external shear (NSIEg) in Table 1.3 as well as in figure 1.20.

1.3.3 Applications

GL offers a wide range of applications out of which we will cite those that are, according to the author, the most relevant in the current astrophysical context. GLs are caused by the bending of light owing to the presence of mass. Whether this mass is luminous or not does not matter: gravitational light deflection is independent of the nature of the matter and of its state. A precise modelling of a given lens can be consequently used so as to study the mass distribution of the deflecting galaxy (see equation 1.35) as well as to subsequently infer its total mass (e.g. Sonnenfeld et al. 2012). The comparison of this mass distribution against those that are derived from luminous matter based on alternative techniques (e.g. rotation curves) allows us to precisely determine the contribution and repartition of the dark matter in this galaxy.

Another point we did not yet mention is the time delay occurring in GL having multiple lensed images. Briefly, as light rays from multiples images propagate along different paths from the source to the observer, there will result in general a time delay between these lensed images. Assuming that the lensed source is variable, like quasars are, one can cross-correlate their photometric light curves so as to compute the time delays occurring between each pair of lensed images. Refsdal (1964) further proved that these time delays are inversely proportional to the Hubble constant H_0. While the classical distance ladder that is used so as to determine the distances of astronomical objects is hampered by the errors propagating from the very nearby objects, through parallaxes, up to very distant objects, through type Ia supernovae, and ultimately provides strong uncertainties on H_0, GLs on their side provide a straight way of deriving H_0 that do not rely on this distance ladder which leads to potentially fairer predictions. While not developed here, let us still note that the other cosmological parameters can be obtained based on GL statistics (e.g. Surdej et al. 1993; Kochanek 1996). This subject being extensively covered in (Schneider 1992, part2, chapter 6) for the interested reader.

Finally, thanks to the light amplification occurring in GLs, these act as giant natural telescopes that allow faint objects to be seen along with a very high signal-to-noise ratio that would otherwise have been unattainable. This magnification reveals its critical importance in identifying high redshift quasars/galaxies and provided spectra of the Lyα forest with an unprecedented quality for some of these. Whereas the host galaxy of quasars are currently barely resolvable, GL offers a unique view to these galaxies as they will form arcs upon favourable configurations. Depending on the proper modelling of the GL system, these arcs can then be back-propagated into the source plane ultimately revealing the host galaxy of the quasar.

2

Methods for the characterization of quasars

The main concern of the present work is the development of a software solution targeting the characterization of quasars and the detection of gravitational lenses (GLs) within the Gaia mission. However, we will mainly concentrate here on the primary objective of this thesis, that is the determination of the astrophysical parameters (APs) of quasars within the Gaia mission. GLs identification being further covered in Chapter 4 though the machine learning (ML) method we used for this identification is still described in section 2.4 as it concurrently yields fair predictions on the APs of quasars, and will consequently be used for comparison purpose (see chapter 3). In the following, sections 2.2 and 2.3 describe the methods that were specifically developed in the framework of the quasar classifier (QSOC) work package (WP), responsible for the characterization of quasars within the CU8 data processing, and being first outlined in the next section.

2.1 The quasar classifier work package

The QSOC WP, abbreviated here as QSOC, is one of the software module constituting the astrophysical parameters inference system (APSIS) dealing with the classification and characterization of more than one billion sources that Gaia will observe according to the CU8 commitments towards the Gaia data processing and analysis consortium (DPAC). Further informations on the structures underlying the Gaia data processing can be found in section 1.1.4. We find interesting to provide here an exhaustive list of the WPs constituting the APSIS chain as it allows the reader to have a global view of the data processing that is taking place within CU8. Table 2.1 accordingly enumerates the APSIS WPs, each having well defined tasks that are explicitly stated in their name. See however Bailer-Jones et al. (2013) for a detailed description of the latter at the time Gaia was launched.

The present section describes the general objectives of QSOC as well as its technical requirements in section 2.1.1 whereas section 2.1.2 presents the solution we envisaged so as to fulfil these objectives and the issues we encountered upon a straight implementation of this solution.

Table 2.1: Modules constituting the astrophysical parameter inference system (APSIS).

Acronym	Name
DSC	Discrete Source classifier
GSP	Generalized Stellar Parametrizer
GSP-Phot	GSP – Photometry
GSP-Spec	GSP – Spectroscopy
ESP	Extended Stellar Parametrizer
ESP-CS	ESP – Cool Stars
ESP-ELS	ESP – Emission Line Stars
ESP-HS	ESP – Hot Stars
ESP-UCD	ESP – Ultra Cool Dwarfs
FLAME	Final Luminosity Age and Mass
MSC	Multiple Star Classifier
TGE	Total Galactic Extinction
UGC	Unresolved Galaxy Classifier
QSOC	Quasar Classifier
OA	Outlier Analysis
OCA	Object Clustering Algorithm

2.1.1 Objectives and software requirements

The objectives of QSOC are compiled in two technical documents, that are the DPAC software and system specification (GAIA-C1-SP-DPAC-WOM-018, hereafter SSS) and the CU8 software requirement specifications (GAIA-C8-SP-MPIA-CBJ-032, hereafter SRS) [1]. While the SSS defined the top-level requirements against which the overall Gaia data processing must be developed, SRS instead concentrates on the requirements that are specific to CU8 and more particularly to the APSIS chain. The SRS being hence fully compliant with the SSS. The SRS is prepared by the CU8 WP managers and later approved by the CU8 configuration control board whereas the SSS is written by CU1 and all CU leaders and further endorsed by the DPAC executive (DPACE).

According to the SSS and the SRS, the objectives of the QSOC module can be respectively summarized as:

> "Where possible, CU8 should provide estimates of the physical parameters of non-stellar sources, in particular quasars and galaxies."

and

> "The QSO Classifier (QSOC) will make use of BP/RP photometry and errors calibrated by CU5 to estimate the APs and their associated errors of objects classified as QSOs by the DSC algorithm. The class-probabilities delivered by the DSC thus act as a control parameter on the algorithm through an internal test. The algorithm also makes use of an internal table of BP/RP photometry of reference QSOs, called a reference library and defined as auxiliary data."

[1] The access to both these documents is restricted to the DPAC members only, although these are provided in the bibliography

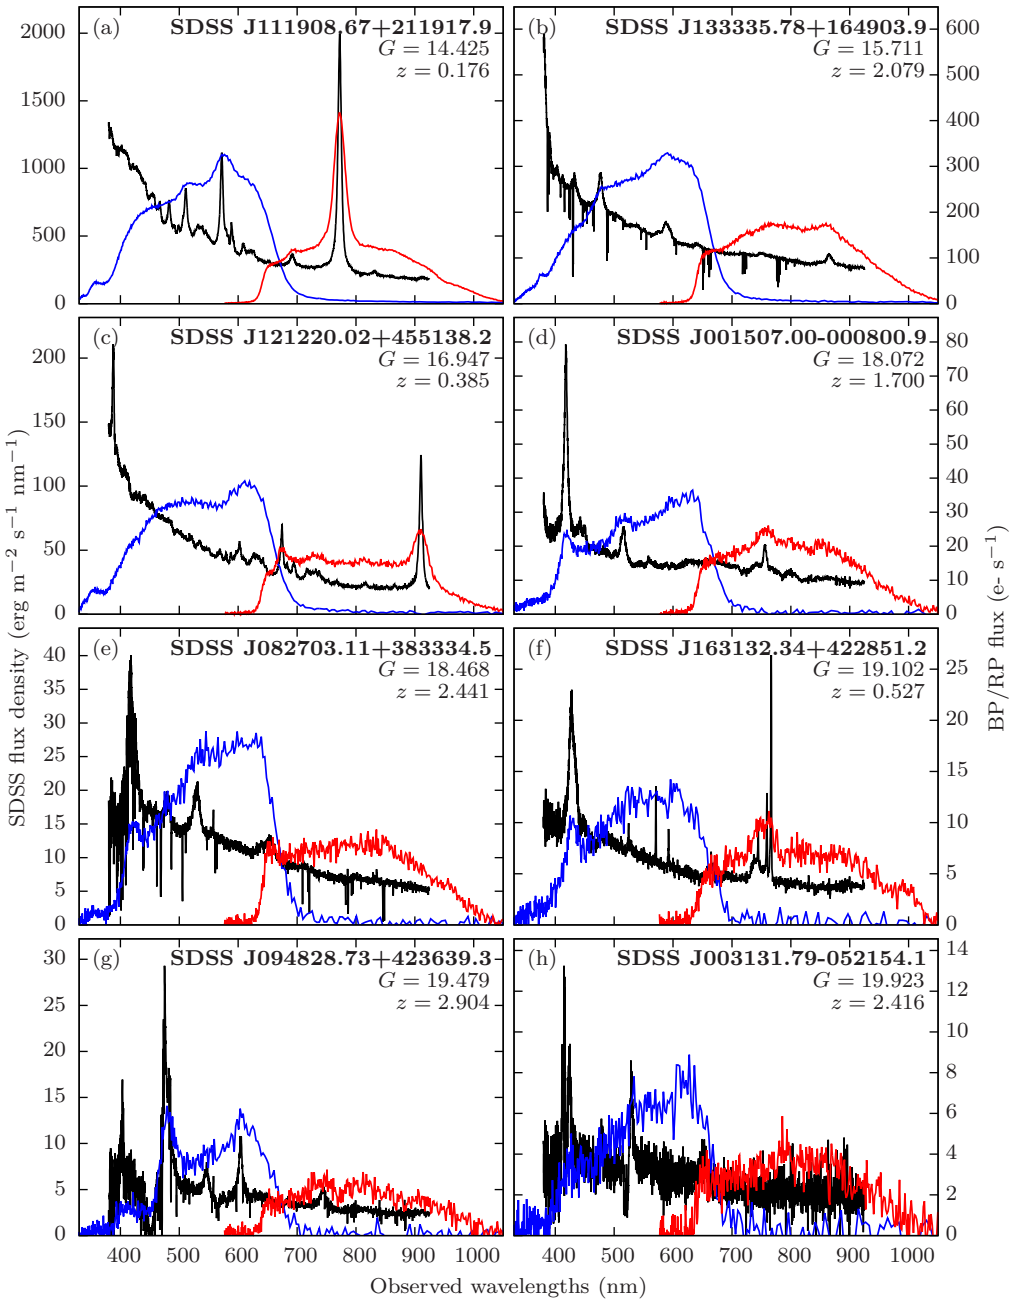

Figure 2.1: Examples of simulated BP/RP spectra of quasars of various G magnitudes. Black lines: SDSS spectra, blue/red lines: BP/RP spectra. Note that since both the SDSS flux densities (in erg m^{-2} s^{-1} nm^{-1}) and the BP/RP fluxes (in e- s^{-1}) are of the same order of magnitude, these were accordingly set on a common scale in each plot. Although provided here as a comparison point, spectra brighter than $G = 18$ will be very rare in practice (e.g. from the twelfth data release of the SDSS quasar catalogue of Pâris et al. (2017), 0.97% of spectra have $G < 18$ mag whereas 0.096% have $G < 17$ mag). The overall bell-shape of the BP/RP spectra, their low resolution, their limited signal-to-noise ratio (at $G > 18$ mag) as well as the broadening of their emission lines owing to the BP/RP spectrometer optics constitute hurdles towards the correct determination of their APs.

Consequently, the QSOC module, once operational, will be in charge of the determination of the APs of objects having a high probability to be quasars according to the discrete source classifier module (DSC). For this purpose, it will rely on the BP and RP spectra and errors coming from CU5 (as illustrated in figure 2.1), possibly using auxiliary data. The quasar identification is accordingly left to the DSC module. The aforementioned objectives are further particularized into software requirements, that are translating these informal objectives into computer assimilable objectives. We cite here a subset of these requirements we found the most relevant regarding the present work. Unmentioned requirements encompassing more technical subjects like, for example: the coding style, the documentation rules, the software interfaces, the data access policy, ...

Based on the SSS and SRS, these (non-exhaustive) requirements can be listed according to their identifiers as

CU8-QSOC-S-FUN-020 This algorithm only operates on those objects assigned a QSO probability by DSC above some threshold.

CU8-QSOC-S-FUN-040 Epoch data processing: the algorithm shall provide classifications and/or APs for QSOs based on epoch BP/RP. This includes at least the photometric redshift and some characterization of the rest-frame spectrum. Multiple solutions should be reported as appropriate.

CU8-QSOC-S-FUN-060 Combined data processing: the algorithm shall provide classifications and/or APs for QSOs based on combined BP/RP, as well as astrometry and variability. This includes at least the photometric redshift and some characterization of the rest-frame spectrum. Multiple solutions should be reported as appropriate.

CU8-QSOC-S-FUN-080 Provide uncertainty estimates in the derived APs. Where appropriate, covariances should also be reported. Very poor fits (i.e. potential misclassifications) should be flagged.

CU8-GEN-S-FUN-160 Algorithms shall be robust to missing data. That is, they must still function and return sensible results if provided with one or more undefined inputs.

CU8-GEN-T-PRF-020 The CPU resources required by the CU8 six-months data processing shall not exceed the resources allocated to the CU8 subsystem.

A particularly stringent point last cited is the CPU consumption. The CPU resources allocated to the whole APSIS chain by the data processing center at CNES (DPCC) is of 1.5×10^{12} floating-point operations (flops)[2] s^{-1}. For a 6 months operation cycle, assuming that all data shall be processed within 100 days owing to operational overheads (e.g. data storage, synchronization, reprocessing upon errors, ...) as well as to subsequent validation purposes, this gives 1.3×10^{19} flops to process 6 months of data (see GAIA-C8-SP-MPIA-CBJ-032, appendix B). Ideally, these resources are shared amongst the 14 WPs composing APSIS (see Table 2.1) according to the number of sources they have to process. A distinction should be made here between the DSC and GSP-Phot WPs that are run over the whole billion of sources of Gaia, the GSP-spec WP that is run on 10^8 sources and other WPs that are run on 10^4–10^6 sources on average (Robin et al. 2012). The CPU resources that are allocated to

[2]Specific computer instructions dealing with the floating point representation of number in computers (i.e. the IEEE 754 nom).

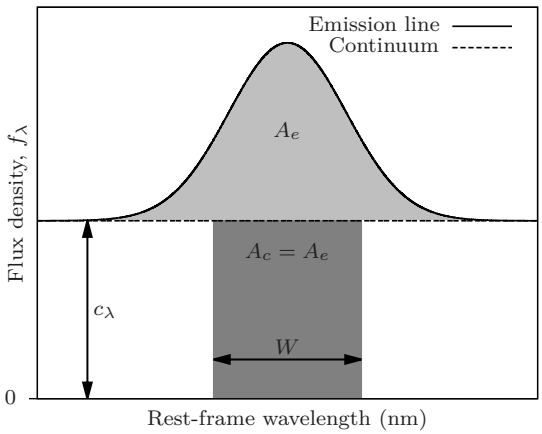

Figure 2.2: Illustration of the equivalent width of an emission line. Considering an emission line having a flux density given by f_λ and a corresponding continuum flux density of c_λ, we have that the total equivalent width of the emission line is given by $W = \int (f_\lambda - c_\lambda)/c_\lambda \mathrm{d}\lambda$. In case where c_λ is constant (i.e. we have a flat continuum), W can be seen as the width that would be required for the area under the continuum, $A_c = W c_\lambda$, to be equal to the area between the emission line and the continuum, $A_e = \int f_\lambda - c_\lambda \mathrm{d}\lambda$.

QSOC are hence estimated to be in the order of 10^9 flops per source or equivalently to a processing time of 0.6 ms per quasar in DPCC, in agreement with Bailer-Jones et al. (2013). Such a consumption being often reached by complex computer algorithms, as we will see in the next section. This somewhat enforces the APSIS WPs to use algorithms nearly exclusively coming from the field of ML methods, that frequently do not have such a high complexity while still providing rather fair predictions.

Another critical point is the ability the software shall have so as to deal with both the epoch spectra (i.e. single transit spectra) as well as with combined spectra (requirements CU8-QSOC-S-FUN-040 and CU8-QSOC-S-FUN-060) given that both have different noise levels as well as a potentially very different number of samples (e.g. epoch spectra have 60 samples for each of the BP/RP spectra while combined spectra may have up to 8 times more samples depending on the oversampling that is used, see Paper III for example). Also from CU8-GEN-S-FUN-160, spectra having some corrupted fluxes must still be treated with the highest possible efficiency.

The previously mentioned objectives and requirements along with the intrinsic BP/RP spectra shortcomings that are highlighted in figure 2.1 together reveal the intricacy of the posed problem and its need for a detailed analysis. Before going further, we have to note that the APs we shall retrieve are not explicitly specified neither by the SSS, nor by the SRS at the exception of the redshift. It is hence a matter of choice to include some additional APs that will be extracted out of the BP/RP spectra of quasars for further publication in the Gaia catalogues. Unsurprisingly, we isolated some rather classical APs applying to the widest possible range of quasars, namely:

– **The redshift of the quasar**, that we have already encountered in section 1.2.4 and equation 1.20. The latter being due to the expansion of the Universe and results in a shift of the wavelengths emitted in the quasar rest-frame, λ_{rest}, to longer wavelengths that are finally observed at wavelengths λ_{obs} such that the

redshift is defined as

$$z = \frac{\lambda_{\text{obs}} - \lambda_{\text{rest}}}{\lambda_{\text{rest}}}. \tag{2.1}$$

This parameter is the most important one as virtually all astronomical applications using extragalactic objects rely on it. It is hence critical to have the fairest prediction of the latter along with strong physical uncertainties as well as a secure identification of the unfair predictions.

– **The BAL nature of the quasar**. The latter depends on the presence of absorption features bluewards of some characteristic emission lines (e.g. C IV being the most common). The presence of broad absorption line (BAL) in $\sim 10\%$ of the quasar population currently lacks some convincing physical explanations as outlined in section 1.2.3. A statistically significant sample of these objects with a very low selection bias (e.g. Gaia will be complete down to an observed G magnitude of 20 for BP/RP spectra) is hence required so as to perform higher resolution follow-ups that may shed the light on the physical processes causing these BAL features. Some finer by-products of the BAL classification, like the Balnicity index (Weymann et al. 1991) or the absorption index (Hall et al. 2002), are not considered in this study owing to the limited capability of the BP/RP spectra regarding the characterization of the BAL troughs (see Paper III, figure 3.11). An example of this limitation stands in the spectrum (h) from figure 2.1. This object effectively encompasses a moderate BAL feature with FWHM ~ 1200 km s^{-1} bluewards of the C IV emission line ($\lambda \lesssim 527$ nm) that is largely suppressed from the BP/RP spectra.

– **The slope of the quasar continuum** can be used as a discriminant for the detection of (strong) relativistic jets amongst quasars. Spectra having a steep continuum slope (e.g. $\alpha_\nu \gtrsim -0.5$ where α_ν is later defined in equation 2.2) being effectively more prone to harbour relativistic jets producing the observed synchrotron radiation[3]. We will then have that the statistics based on this *spectral index*, α_ν, as well as dedicated surveys can bring extremely valuable informations on the physical mechanisms raising these jets as these are still obscure at the present day (see section 1.2.3). From equation 1.1, we have that the spectral index is defined through

$$\nu^{\alpha_\nu} \propto f_\nu \tag{2.2}$$

where f_ν is the flux density that is observed at rest-frame frequency ν. Furthermore owing to the fact that $\mathrm{d}\lambda/\mathrm{d}\nu = -c/\nu^2$ with $f_\nu = f_\lambda |\frac{\mathrm{d}\lambda}{\mathrm{d}\nu}|$ we have that the spectral index can be defined in terms of the rest-frame wavelength, λ, and associated flux density, f_λ, according to

$$\lambda^{\alpha_\lambda} \propto f_\lambda, \qquad \alpha_\lambda = -\alpha_\nu - 2. \tag{2.3}$$

– **The total equivalent width of the emission lines**. Likewise the slope of the QSO continuum acts as a proxy for the presence of relativistic jets, the (total) equivalent width of the quasar emission line(s) correlates with the absolute luminosity of the quasar continuum. The equivalent width (EW) of an emission line, presented in details in figure 2.2, is a measure of the strength of this emission line as compared to the underlying continuum. From the work of Baldwin (1977), we know that the highest the EW of an emission line is, the lowest the

[3] Note that this does not necessarily imply that these jets are extended

absolute continuum luminosity of the quasar is. The total EW of the emission lines is a straight generalization of the EW presented in figure 2.2 applied over the whole observed spectrum. Mathematically, it is defined as

$$W = \int \frac{e_\lambda}{c_\lambda} \mathrm{d}\lambda, \qquad \text{with } e_\lambda = \begin{cases} f_\lambda - c_\lambda & \text{if } f_\lambda > c_\lambda \\ 0 & \text{otherwise} \end{cases} \tag{2.4}$$

where f_λ and c_λ are the flux densities associated with the emission line and the continuum, respectively (see figure 2.2).

2.1.2 Envisaged solution and associated issues

Various solutions can be envisaged so as to guess the APs of quasars based on their BP/RP spectra. Still, the determination of all these APs primarily relies on a correct prediction of the redshift, whence the importance to first have a very reliable estimation of the latter. Solutions that are the most commonly implemented for this purpose are:

(i) **The visual inspection** of the whole set of quasars (e.g. Pâris et al. 2017), though this solution is infeasible regarding the 5×10^5 quasars that should be treated within the six month cycle of the Gaia processing.

(ii) **The matching of spectral lines** (e.g. Machado et al. 2013), consisting in the recognition of some patterns within spectra, mostly emission lines, that are subsequently matched to theoretical predictions (e.g. rest-frame emission lines from Table 1.2). This method is however restricted to medium resolution spectra having a relatively good signal-to-noise ratio that together enable the spectral lines to be individually isolated. The BP/RP spectra are not fulfilling these requirements in general (see figure 2.1 for example). Supplemental shortcomings of this method are that nor the line heights, nor their widths, nor their covariances[4] are taken into account here. This point turns out to be extremely harmful in case of quasars as most emissions lines have rather constant covariances (e.g. the Balmer series Hα, Hβ, Hγ, ...) and more particularly in case of BP/RP spectra as, for example, the Lyα emission line will be systematically blended with the N V emission line as does the Hβ emission line with the [O III] doublet. The presence of such very extended emission lines being a first hint towards the occurrence of such a line mixing that would be regrettable not to use.

(iii) **Methods based on machine learning** (e.g. Claeskens et al. 2006) are plausible alternatives that were effectively used in past versions of QSOC (see Bailer-Jones et al. 2013). However, these act as black-box algorithms that can lead to unphysical and/or suboptimal predictions either coming from an incomplete coverage of the input space of parameters or from an inappropriate choice of the model parameter(s) (see bias/variance trade-off, section 2.4). Moreover, these provide a limited capability in the estimation of the uncertainties as those they return are rather measures of the inadequacy of the ML model to match the provided sources. Accordingly, these methods should be privileged when no other (fast) solution is available or when the set of input parameters is fairly well covered by the learning set of observations. In case of quasars, the latter point translates into the need for a spectral library covering most of the characteristics encountered in quasars over the whole range of redshifts that is explored, hence requiring an extremely huge learning dataset.

[4] The tendency that a line has to vary along with another line

(iv) **The phase correlation algorithm** (e.g. Bolton et al. 2012) that turns out to be the solution that is the most adapted to our problem given the shortcomings of the alternative methods.

The phase correlation algorithm stands to be the most widespread method for determining the redshift of quasars and galaxies. It was first introduced in astronomy through the work of Brault & White (1971) though the weighted version we outline here is based on the specific needs that emerged out of the Baryon Oscillation Spectroscopic Survey (BOSS) of the SDSS (Bolton et al. 2012). The interested reader may find in the introduction of Paper II a more complete historical background of this technique applied to astrophysics. The idea behind the weighted phase correlation (WPC) algorithm is to recognize that the redshift from equation 2.1 turns into a simple offset if we consider the wavelength to stand on a logarithmic scale, that is

$$\log \lambda_{\mathrm{obs}} = \log \lambda_{\mathrm{rest}} + \log(z+1) = \log \lambda_{\mathrm{rest}} + Z \tag{2.5}$$

where, in the last step, we defined $Z \equiv \log(z+1)$. Let us assume that we have rest-frame quasar templates that are sampled on such a logarithmic scale, $\mathbf{T}$, out of which $\mathbf{T}_{ij}$ is the ith sample from the jth template. The ith sample being then associated with the (logarithmic) rest-frame wavelength

$$\lambda_{\mathrm{rest}} = (1+\Delta\lambda)^i, \qquad \log \lambda_{\mathrm{rest}} = i \log(1+\Delta\lambda), \tag{2.6}$$

where $\Delta\lambda$ is the wavelength sampling we choose to use. If the observed spectrum, $\boldsymbol{s}$, and the associated uncertainties, $\boldsymbol{\sigma}$, are both sampled over the same logarithmic scale (i.e. same $\Delta\lambda$), then the determination of the redshift of $\boldsymbol{s}$ can be formulated as a chi-square minimization problem through

$$\chi^2(k) = \sum_i \frac{1}{\sigma_i^2} \left(s_i - \sum_j a_j(k)\, T_{(i+k)\,j} \right)^2 \tag{2.7}$$

where $a_j(k)$ is the optimal linear coefficient allowing to fit the jth template to the observation if this template is shifted by k samples. The optimal redshift being then associated with the minimum of the χ^2 *curve* as defined through equation 2.7 and illustrated in figure 2.3.

The first issue raised by this solution if applied to Gaia is the unavailability of a spectral library of quasars as seen with the BP/RP instruments. This prevents us from building the templates upon which the χ^2 curve technique relies. These spectra can however be produced through the conversion of higher resolution spectra along with the instrument model provided by Airbus Defence and Space (ADS) during the commissioning phase of the satellite (see Paper III, section 4). A spectral library of quasars having the required resolution and wavelength coverage (i.e. $\sim$ 300–1100 nm) being unfortunately still non-existent at the present day, we had to extrapolate an existing library using the method we will later describe in Paper I. Going further into this issue, we have to note that the χ^2 curve technique requires the templates to span the whole range of rest-frame wavelengths we expect to observe. The production of such rest-frame templates being mainly hampered by the very different wavelength coverage of the individual quasar spectra, as illustrated in figure 2.4, as well as by the presence of noise in these spectra.

An efficient method often employed for the production of these quasar rest-frame templates stands in the principal component analysis (PCA, e.g. Yip et al. 2004; Bailey 2012; Tsalmantza et al. 2012, and references therein). The simplest version of

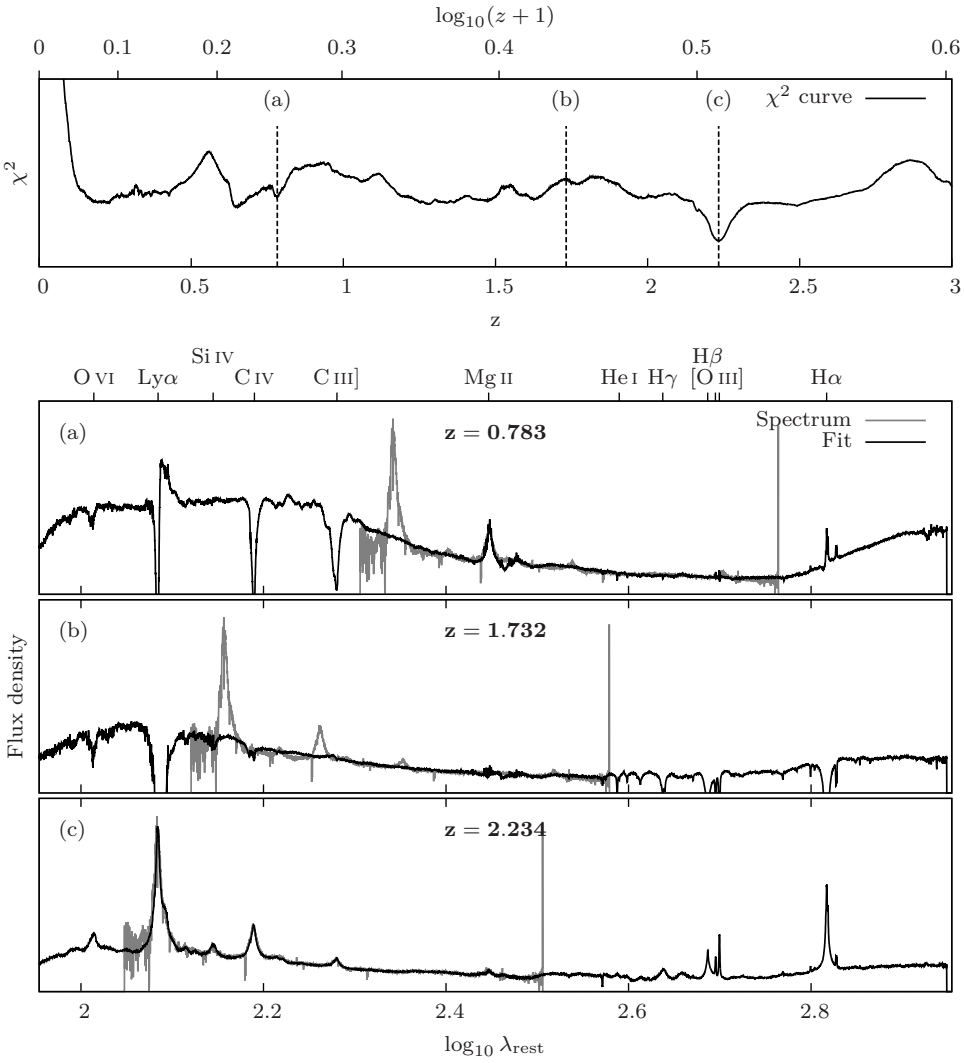

Figure 2.3: χ^2 curve associated with the SDSS J000313.08+274044.9 quasar standing at $z = 2.233$. The curve is built by repeatedly fitting shifted versions of the templates to the observed spectrum while reporting the χ^2 we obtain this way into the graph (see equation 2.7). 15 emission line templates as well as 5 continuum templates out of the online material from Paper III were used for the purpose of the present illustration. Three specific fits are further provided, each corresponding to: (a) a local minimum, (b) a local maximum and (c) the global minimum of the χ^2 curve. The discrepancy existing between the actual redshift (z=2.2329) and the redhift we derived using the χ^2 curve approach (z=2.2338) being due to the discrete sampling we used. A sub-sampling precision on the redshift being gained by fitting a quadratic curve in the vicinity of the global minimum of the χ^2 curve.

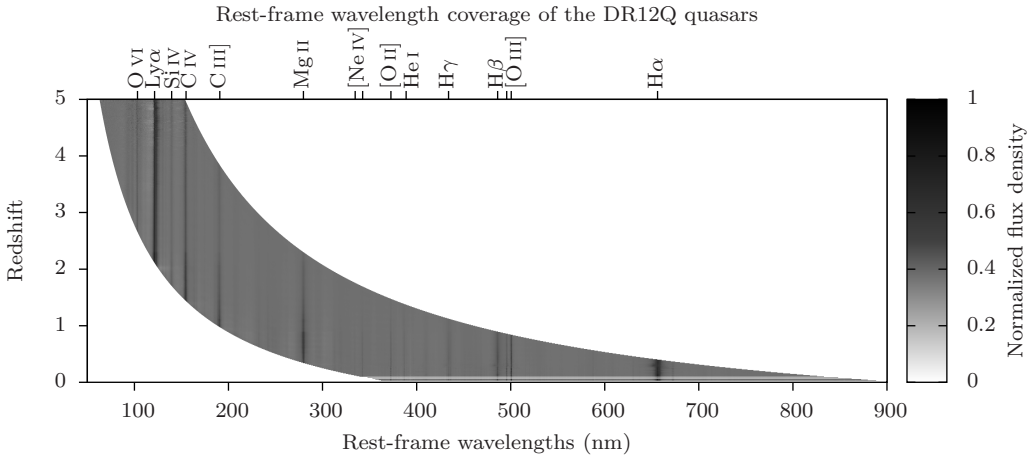

Figure 2.4: Rest-frame wavelength coverage of quasar spectra from the twelfth data release of the Sloan digital sky survey quasar catalogue (Pâris et al. 2017) according to redshift. Spectra are averaged over redshift bins of $\Delta z = 0.01$ and over wavelength bins of $\Delta\lambda = 0.1$ nm.

this PCA technique (see Paper I, section 3) is however unable to deal with missing and noisy data like those encountered in quasar spectra. Proposed solutions to this problem encompass the use of an iterative procedure so as to fill the gaps of the rest-frame spectra (Yip et al. 2004) or the use of a weighted least-squares formulation to the PCA problem (Bailey 2012; Tsalmantza et al. 2012). In the following, we will solely consider weighted versions of the PCA method as these do not rely on iterative procedures where any error made accumulates up to the final templates as well as because these concurrently allow us to take the noise straightly into account. The PCA implementations of Bailey (2012) and of Tsalmantza et al. (2012) are based on iterative approaches to a weighted least-squares formulation of the PCA problem. These however suffer from strong numerical instabilities in cases where a large number of samples ($\sim 30\%$) is missing and/or in cases of very noisy input data (see Paper I, figure 2.9). Furthermore, the chi-square approach they undertake aims to optimize the reconstruction of the input dataset along with a minimal number of templates which do not necessarily extract the most significant patterns out of the input spectra, as we will see in section 2.2.

Last but not least, the χ^2 curve method requires each of the $a_j(k)$ coefficients to be computed in order for the templates to match at best the observation at given shift k. The fastest method designed to solve such a weighted least squares problem is outline in the appendix (section A.1) and results in a numerical complexity of $\mathcal{O}\left(mn^2\right)$ flops where m is the number of samples in our spectra and where n is the number of templates we choose to use. Given that the linear coefficients $a_j(k)$ must be computed for all shifts $k \in \{0, m-1\}$, we have that the overall algorithmic complexity of the χ^2 curve approach is given by $\mathcal{O}\left(m^2n^2\right)$ flops. This quadratic time dependence being a major hurdle towards the use of this method. Indeed, assuming we would like to have a precision on the predicted redshift that is equivalent to or better than the human expertise ($\sigma_z = 0.003$ according to Pâris et al. (2017)) for all redshift up to $z_{\max} = 6$, we hence have that the logarithmic wavelength sampling we should use must satisfy

$$\Delta\lambda \leq \frac{\sigma_z}{z_{\max} + 1} \lesssim 4 \times 10^{-4}. \tag{2.8}$$

From equation 2.6, we also have that this sampling can be expressed as a linear

sampling on a logarithmic scale through $\Delta \log_{10} \lambda = \log_{10}(\Delta\lambda + 1) \approx 1.75 \times 10^{-4}$. The rest-frame wavelength coverage of the BP/RP spectra extending from $\lambda_0 = 300/(z_{\max} + 1) \sim 47$ nm up to $\lambda_1 \sim 1050$ nm at $z = 0$, the number of samples in our templates spectra is hence given by

$$N = \frac{\log_{10} \lambda_1 - \log_{10} \lambda_0}{\Delta \log_{10} \lambda} = 7709. \tag{2.9}$$

If we suppose that a reasonable amount of 3 emission line templates and 3 continuum templates are used so as to fit the observations, we have that the expected number of computer instructions that would be required upon a straight implementation of the χ^2 curve method is approximately given by $\mathcal{O}\left(m^2 n^2\right) \approx 2.14 \times 10^9$ flops, which exceeds the 10^9 flops allocated to QSOC for the processing of a single source whereas only the redshift would have been derived.

2.2 Weighted principal component analysis

As previously mentioned, the PCA implementations of Bailey (2012) and Tsalmantza et al. (2012) suffer from strong numerical instabilities if used along with rest-frame spectra of quasars having large redshift discrepancies or, equivalently, having a large fraction of rest-frame wavelengths that are unobserved. Also, they do not efficiently take into account the covariances that are noticed amongst the samples of the input datasets such that they do not necessarily extract the most significant patterns out of these. For completeness, we now take some time to make the latter point a little more explicit. Consider an input matrix $\mathbf{X}$ out of which the mean observation $\bar{\boldsymbol{x}}$ has been subtracted. The removal of this mean observation being carried out here for conciseness and without any loss of generality. As later explained in Paper I, the PCA decomposition of an input matrix $\mathbf{X}$ can be computed through its singular value decomposition (SVD):

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathrm{T}} \tag{2.10}$$

where $\mathbf{U}$, $\mathbf{V}$ are orthogonal matrices[5] and where $\mathbf{\Sigma}$ is a diagonal matrix containing the eigenvalues of $\mathbf{X}$ ordered according to their magnitudes. Such a decomposition being always existent. The first columns of $\mathbf{U}$ are then the searched principal components $\mathbf{P}$ while the principal coefficients $\mathbf{C}$ are associated with the first rows of the matrix $\mathbf{\Sigma}\mathbf{V}^{\mathrm{T}}$ (i.e. $\mathbf{\Sigma'}\mathbf{V'}^{\mathrm{T}}$). Fitting $\mathbf{X}$ with $\mathbf{P}$ such that $|\mathbf{X} - \mathbf{P}\mathbf{A}|^2$ is minimal in a least-squares sense yields to the normal equations

$$\mathbf{A} = \left(\mathbf{P}^{\mathrm{T}}\mathbf{P}\right)^{-1} \mathbf{P}^{\mathrm{T}}\mathbf{X} = \mathbf{P}^{\mathrm{T}}\mathbf{X} = \mathbf{C}. \tag{2.11}$$

Meaning that $\mathbf{C}$ contains the optimal linear coefficients so as to fit $\mathbf{X}$ with $\mathbf{P}$ in a least squares sense[6]:

$$\mathbf{X} \overset{\chi^2}{=} \mathbf{P}\mathbf{C}. \tag{2.12}$$

Furthermore, we have that the covariance matrix associated with $\mathbf{C} = \mathbf{P}^{\mathrm{T}}\mathbf{X}$,

$$\boldsymbol{\sigma}^2 = \mathbf{C}\mathbf{C}^{\mathrm{T}} = \mathbf{\Sigma'}\mathbf{V'}^{\mathrm{T}}\mathbf{V'}\mathbf{\Sigma'} = \mathbf{\Sigma'}^2, \tag{2.13}$$

is diagonal, meaning that the linear transformation $\mathbf{P}^{\mathrm{T}}\mathbf{X}$ allows us to have null covariances between the samples (i.e. all covariances of $\mathbf{C}$ are equal to zero). Whereas both

[5] An orthogonal matrix $\mathbf{U}$ has the property $\mathbf{U}^{\mathrm{T}}\mathbf{U} = \mathbf{I}$.

[6] A common misconception here is that $\mathbf{P}$ is the minimal set of templates that allow any observation out of $\mathbf{X}$ to be fitted at best in a linear least squares sense whereas this is only valid regarding the $\mathbf{X}$ matrix as a whole.

the diagonalization of the covariance matrix $\boldsymbol{\sigma^2}$ as well as the least-squares decomposition of the input matrix $\mathbf{X}$ into $\mathbf{PC}$ are concurrently achieved in case of classical PCA methods, the introduction of weights associated with each samples of $\mathbf{X}$ will break this simultaneity such that any related weighted PCA method must hence choose which objective it will accomplish. While Bailey (2012) and Tsalmantza et al. (2012) privileged the weighted least squares decomposition of the input matrix $\mathbf{X} \overset{\chi^2}{=} \mathbf{PC}$, we instead favoured here the diagonalization of the weighted covariance matrix defined by

$$\boldsymbol{\sigma^2} = \frac{(\mathbf{X} \circ \mathbf{W}) (\mathbf{X} \circ \mathbf{W})^{\mathrm{T}}}{\mathbf{W}\mathbf{W}^{\mathrm{T}}}. \tag{2.14}$$

The method we present was initially developed so as to extrapolate the SDSS spectra of quasars, covering the wavelength range 400–1000 nm, to a wavelength range that enables their conversion into BP/RP spectra using the instrument model from ADS (i.e. 300–1100 nm). It has been subsequently shown that the χ^2 curve technique yields the fairest predictions if using PCA templates coming of this method. As compared to the implementations of Bailey (2012) and Tsalmantza et al. (2012), our method was proven to be the method of choice regarding the extraction of the most significant patterns out of input datasets. These allow, for example, the extrapolation of missing samples and the modelling of unseen but similar observations to be performed in an efficient way. It however fails to reproduce at best (i.e. in a least squares sense) an input dataset given a minimal number of components.

Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **446,** 3545–3555 (2015) doi:10.1093/mnras/stu2219

Weighted principal component analysis: a weighted covariance eigendecomposition approach

L. Delchambre★

Institut d'Astrophysique et de Géophysique, Université de Liège, Allée du 6 Août 17, B-4000 Sart Tilman (Liège), Belgium

Accepted 2014 October 12. Received 2014 October 7; in original form 2014 July 4

ABSTRACT

We present a new straightforward principal component analysis (PCA) method based on the diagonalization of the weighted variance–covariance matrix through two spectral decomposition methods: power iteration and Rayleigh quotient iteration. This method allows one to retrieve a given number of orthogonal principal components amongst the most meaningful ones for the case of problems with weighted and/or missing data. Principal coefficients are then retrieved by fitting principal components to the data while providing the final decomposition. Tests performed on real and simulated cases show that our method is optimal in the identification of the most significant patterns within data sets. We illustrate the usefulness of this method by assessing its quality on the extrapolation of Sloan Digital Sky Survey quasar spectra from measured wavelengths to shorter and longer wavelengths. Our new algorithm also benefits from a fast and flexible implementation.

Key words: methods: data analysis – quasars: general.

1 INTRODUCTION

Principal component analysis (PCA) is a well-known technique initially designed to reduce the dimensionality of a typically huge data set while keeping most of its variance (Pearson 1901; Hotelling 1933). PCA is intimately related to the singular value decomposition (SVD) since the principal components of a data set, whose arithmetic mean is zero, will be equal to the eigenvectors of the covariance matrix sorted by their corresponding eigenvalue; or equivalently by the variance they account for. The principal coefficients are the linear coefficients allowing us to reconstruct the initial data set based on the principal components. Further details about PCA will be given in Section 3 of this paper. Interested readers are also invited to read Schlens (2009) for an accessible tutorial on this technique or Jolliffe (2002) for a deeper analysis.

PCA has many applications in a wide variety of astronomical domains from the classification of the Sloan Digital Sky Survey (SDSS) quasar spectra and their redshift determination (Yip et al. 2004; Pâris et al. 2014) to the study of the point spread function variation in lensing surveys (Jarvis & Jain 2004). The method described hereafter was originally developed in the framework of the *Gaia* astrophysical parameters inference system (Bailer-Jones et al. 2013) where it is used to provide learning data sets of spectrophotometric data based on SDSS quasar catalog spectra (Pâris et al. 2014). The latter cover the observed wavelength range 4000–10 000 Å and are extrapolated by our algorithm to the wavelength range 3000–11 000 Å covered by *Gaia*. Even if developed for an astronomical purpose, it can be used in any problems requiring PCA decomposition of weighted data. The case of missing data being simply the limiting case of weights equal to zero.

Classical PCA is a mature tool whose performance in dimensionality reduction and pattern recognition has been assessed for a long time. Nevertheless, its main limitation comes from the fact that it is not adapted to the case of samples having weighted and/or missing data. The inherent consequence is that the classical PCA implementations made no difference between variance coming from a genuine underlying signal and variance coming from measurement noise.

Most of the previous works cope with these limitations mainly by focusing on bypasses to the problem of noisy and/or missing data; or deal explicitly with particular cases. These encompass, for example, the interpolation of missing data (Beale & Little 1975) or cases where the weight matrix can be factorized into per-observation and per-variable weight matrices (Greenacre 1984). Jolliffe (2002) in sections 13.6 and 14.2 makes the point about these proposed solutions.

At the present time, some methods are still able to deal with weight matrices having the same size as the corresponding data set (Gabriel & Zamir 1979; Wentzell et al. 1997; Tipping & Bishop 1999; Srebro & Jaakkola 2003). Nevertheless, none of these are able to provide the orthogonal principal components ordered by the data set variance it accounts for. Rather, they provide an unsorted set of not-necessary orthogonal vectors whose linear combination is optimized to describe the underlying variance but whose goal is

★E-mail: ldelchambre@ulg.ac.be

not compatible with the explanation of the variance given a minimal number of vectors.

Based on the idea of Roweis (1997), Bailey (2012) and Tsalmantza & Hogg (2012) have recently proposed methods based upon the expectation-maximization (EM) algorithm (Dempster, Laird & Rubin 1977) in order for the PCA to include weights associated with each variable within each observation. The objective of these methods is globally similar to the one of this paper. Differences mainly come from the fact that Tsalmantza & Hogg (2012) aim at finding an orthogonal decomposition of the original matrix such that the whole data set variance is the best accounted for. Instead, our implementation focuses on finding the orthogonal vectors that are the best at individually describing the data set variance at the expense of a lower explained global variance. This trade-off comes from the fact that in a weighted case, the solution to the problem of finding the set of N components explaining most of the variance of a data set is not guaranteed to contain the eigenvectors that are the best at individually describing this variance. The implementation of Bailey (2012) takes benefits of the flexibility of the EM algorithm in order to interpolate between these two solutions.

In Section 2, we explain the notation used in this paper. We summarize the properties of the classical PCA in Section 3. We see in details two current alternative implementations of weighted PCAs in Section 4. In Section 5, we describe our new algorithm while in Section 6, we see its application on simulated data and real cases and compare it against other algorithms. Finally, some properties and extensions are discussed in Section 7 and we conclude in Section 8.

2 NOTATION

This paper uses the following notations: vectors are in bold italic, $\boldsymbol{x}$; x_i being the element i of the vector $\boldsymbol{x}$. Matrices are in uppercase boldface or are explicitly stated; i.e. $\mathbf{X}$ from which the ith row will be denoted $\mathbf{X}_i^{\mathrm{row}}$ and the jth column by $\mathbf{X}_j^{\mathrm{col}}$, element at row i, column j will then be $\mathbf{X}_{ij}$. Amongst matrix operators, $\mathbf{a} \circ \mathbf{b}$ denotes the element-wise product (Hadamard product) of $\mathbf{a}$ and $\mathbf{b}$ and $\|\mathbf{a}\|$ denotes the Euclidian matrix norm of $\mathbf{a}$.

Consider a problem where we have N_{obs} observations each containing N_{var} variables, from which we want to retrieve N_{comp} principal components. For reference, here are the often used matrices along with their corresponding sizes: $\mathbf{X}$ the data set matrix ($N_{\mathrm{var}} \times N_{\mathrm{obs}}$) from which we have subtracted the mean observation $\bar{\mathbf{y}}$ ($\mathbf{X}_i^{\mathrm{col}} = \mathbf{Y}_i^{\mathrm{col}} - \bar{\mathbf{y}}$); $\mathbf{W}$ the weight of each variable within each observation ($N_{\mathrm{var}} \times N_{\mathrm{obs}}$); $\mathbf{P}$ the orthogonal matrix of principal components ($N_{\mathrm{var}} \times N_{\mathrm{comp}}$); $\mathbf{P}_i^{\mathrm{col}}$ being the ith principal component; $\mathbf{C}$ the principal coefficient matrix ($N_{\mathrm{comp}} \times N_{\mathrm{obs}}$); $\boldsymbol{\sigma}^2$ the symmetric matrix of variance–covariance ($N_{\mathrm{var}} \times N_{\mathrm{var}}$) associated with $\mathbf{X}$.

Finally, $\mathbf{A} \overset{\chi^2}{=} \mathbf{B}$ means that $\mathbf{A}$ is the nearest matrix from $\mathbf{B}$ in a – potentially weighted – least-squares sense. Mathematically, this is equivalent to have $\mathbf{A}$ and $\mathbf{B}$ such that

$$\chi^2 = \sum_{ij} \mathbf{W}_{ij}^2 \left(\mathbf{A}_{ij} - \mathbf{B}_{ij}\right)^2 \tag{1}$$

is minimized.

3 PRINCIPAL COMPONENT ANALYSIS

Regarding classical PCA and assuming – without any loss of generality – that we would like to retrieve as many principal components as the number of variables (ie. $N_{\mathrm{comp}} = N_{\mathrm{var}}$), then the goal of the PCA will be to find a decomposition

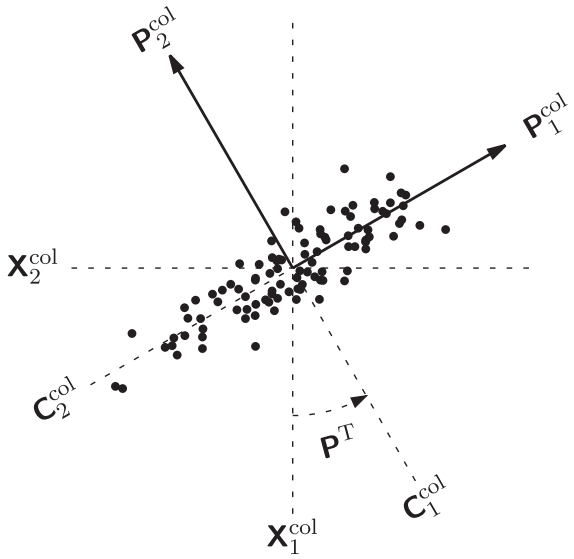

Figure 1. A two-dimensional PCA example: $\mathbf{P}^{\mathrm{T}}$ can be seen as an orthogonal coordinate transformation from $(\mathbf{X}_1^{\mathrm{col}}; \mathbf{X}_2^{\mathrm{col}})$ to $(\mathbf{C}_1^{\mathrm{col}}; \mathbf{C}_2^{\mathrm{col}})$ such that the data set variance is maximized along the vector $\mathbf{P}_1^{\mathrm{col}}$ and $\mathbf{P}_2^{\mathrm{col}}$. Note that, for a didactical purpose, we chose $\mathbf{P}$ to be a rotation matrix but practically it can be any orthogonal matrix.

$$\mathbf{X} = \mathbf{PC}, \tag{2}$$

such that

$$\mathbf{D} = \mathbf{P}^{\mathrm{T}}\boldsymbol{\sigma}^2\mathbf{P} = \mathbf{P}^{\mathrm{T}}\mathbf{X}\mathbf{X}^{\mathrm{T}}\mathbf{P} \tag{3}$$

is diagonal and for which

$$\mathbf{D}_{ii} \geq \mathbf{D}_{jj}; \quad \forall i < j. \tag{4}$$

Note that based on equation (3) and according to the spectral theorem,[1] $\mathbf{P}$ will be orthogonal.

Intuitively, the matrix $\mathbf{P}$ can be seen as a change of basis allowing us to maximize the variance within $\mathbf{D}$ and thus minimizing the off-diagonal elements corresponding to the covariance. Differently stated, each $\mathbf{P}_i^{\mathrm{col}}$ defines a privileged direction along which the data set variance is the best explained. The fact that $\mathbf{D}$ is ordered implies that for $i < j$, the principal component $\mathbf{P}_i^{\mathrm{col}}$ accounts for more – or equal – variance than $\mathbf{P}_j^{\mathrm{col}}$. For the sake of clarity, a comprehensive PCA example is given in Fig. 1.

A common solution to such a classical PCA is based on the SVD of $\mathbf{X}$:

$$\mathbf{X} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^{\mathrm{T}}, \tag{5}$$

where $\mathbf{U}$, $\mathbf{V}$ are orthogonals, $\boldsymbol{\Sigma}$ is diagonal and for which $|\boldsymbol{\Sigma}_{ii}| \geq |\boldsymbol{\Sigma}_{jj}|$; $\forall i < j$. By setting $\mathbf{P} = \mathbf{U}$ and $\mathbf{C} = \boldsymbol{\Sigma}\mathbf{V}^{\mathrm{T}}$, we find that equation (3) becomes

$$\mathbf{P}^{\mathrm{T}}\mathbf{X}\mathbf{X}^{\mathrm{T}}\mathbf{P} = \mathbf{C}\mathbf{C}^{\mathrm{T}} = \boldsymbol{\Sigma}^2, \tag{6}$$

that fulfils the conditions of equations (3) and (4). Note that in equation (3) the exact variance–covariance matrix should be normalized by N_{obs} but since we are solely interested in the diagonalization of $\boldsymbol{\sigma}^2$, we drop it.

[1] Any real symmetric matrix is diagonalized by a matrix of its eigenvectors.

4 WEIGHTED EXPECTED MAXIMIZATION PCA

As already mentioned in Section 1, the current methods efficiently dealing with the problem of weighting PCA (Bailey 2012; Tsalmantza & Hogg 2012) aim at best explaining the whole data set variance according to a given number of principal components, which is equivalent to minimize

$$\chi^2 = \sum_{ij} \mathbf{W}_{ij}^2 \left(\mathbf{X}_{ij} - [\mathbf{PC}]_{ij}\right)^2, \tag{7}$$

where the weighted mean observation we subtracted (cf. Section 2) is given by

$$\bar{\mathbf{y}} = \sum_i \mathbf{W}_i^{\rm col} \mathbf{Y}_i^{\rm col} \Big/ \sum_i \mathbf{W}_i^{\rm col}. \tag{8}$$

We notice that equation (7) has latent variables such that it has to rely on an iterative procedure to be solved.

The EM algorithm is a statistical tool specifically designed to optimize a likelihood function for models having latent – or unknown/hidden – variables (Dempster et al. 1977). This iterative procedure is composed of two steps.

(i) E-step: find the expectation value of the latent variables given the parameters of the current model.
(ii) M-step: find the parameters of the model such that the likelihood function is optimized.

Based on the latter, Roweis (1997) has developed a fast and straightforward classical PCA algorithm for which the conditions in equations (2), (3) and (4) are all fulfilled. Regarding weighted PCA and more specifically the χ^2 described by equation (7), we will have the following weighted expected maximization PCA (WEMPCA) algorithm:

$\mathbf{P} \longleftarrow$ Random orthogonal matrix
While $\mathbf{P}$ and $\mathbf{C}$ have not converged
  (E-step) Find $\mathbf{C}$ that minimizes χ^2 given $\mathbf{P}$.
  (M-step) Find $\mathbf{P}$ that minimizes χ^2 given $\mathbf{C}$.

Note that the convergence criterion is still relative. It can be based on the χ^2 – or the change in the principal components $\Delta\mathbf{P}$ – falling under a given threshold, the fact that the algorithm has reached a given number of iterations or whatever criterion we consider as relevant.

4.1 Tsalmantza's implementation

Tsalmantza & Hogg (2012) designed a general approach to the modelling and dimensionality reduction of the SDSS spectra called 'Heteroskedastic Matrix Factorization'. More specifically, it attempts to minimize

$$\chi_\epsilon^2 = \chi^2 + \epsilon \sum_{i>1} \sum_j \left[\mathbf{P}_{ij} - \mathbf{P}_{(i-1)j}\right]^2, \tag{9}$$

subject to

$$\left.\begin{array}{l} \mathbf{P}_{ij} \geq 0 \\ \mathbf{C}_{jk} \geq 0 \end{array}\right\} \forall i, j, k. \tag{10}$$

We recognize the first part of equation (9) as being equation (7) while the second part is a smoothing regularization term whose scalar ϵ defines the strength. Non-negativity constraints reflect a particular need to have a meaningful physical interpretation of the resulting spectra.

Regarding the fact that we would like to model the widest variety of data sets, we will drop the non-negativity constraints that otherwise would have restricted our search space. Concerning the smoothing regularization factor, we have to note that it will be highly problem-dependent and that it can be tricky to optimize, this will result in a potential unfair comparison with other methods. We will then consider the case $\epsilon = 0$. Moreover, as we will see in Section 7.3, our method can deal with principal components smoothing as well, consequently ignoring it will not constitute a major drawback to our implementation.

The resulting function to optimize will then be reduced to the sole equation (7). Nevertheless, and before going further, we have to note that minimizing equation (7) will provide us a lower-rank matrix approximation of $\mathbf{X}$ but it is not a sufficient condition for the resulting matrices $\mathbf{P}$ and $\mathbf{C}$ to be considered as a PCA decomposition. According to Tsalmantza & Hogg (2012), the solution to this problem can be solved in two steps.

First consider a lower-rank matrix decomposition of $\mathbf{X}$, similar to the one produced by the solution of equation (7),

$$\mathbf{X} \overset{\chi^2}{=} \mathbf{AB}, \tag{11}$$

where for clarity, the sizes of these matrices are $\mathbf{A}(N_{\rm var} \times N_{\rm comp})$ and $\mathbf{B}(N_{\rm comp} \times N_{\rm obs})$. Now suppose an orthogonal basis $\mathbf{P}_0$ of $\mathbf{A}$; such an orthogonal basis always exists for full-rank matrices and can be retrieved through a straightforward Gram–Schmidt process for example. The associated coefficients matrix $\mathbf{C}_0$ is then directly retrieved by

$$\mathbf{C}_0 = \mathbf{P}_0^{\rm T}\mathbf{AB}. \tag{12}$$

Secondly, in a way similar to equation (6), we will take the classical PCA decomposition of $\mathbf{C}_0$,

$$\mathbf{C}_0 = \mathbf{P}_c\mathbf{C}, \tag{13}$$

such that $\mathbf{CC}^{\rm T}$ is diagonal. The resulting principal coefficient matrix will then be given by

$$\mathbf{P} = \mathbf{P}_0\mathbf{P}_c, \tag{14}$$

that will be orthogonal and that will provide us with the final decomposition

$$\mathbf{X} \overset{\chi^2}{=} \mathbf{PC}. \tag{15}$$

The above-mentioned steps have thus to be performed after the EM algorithm minimizing equation (7) in order for $\mathbf{P}$ to be orthogonal and for the covariance matrix $\mathbf{CC}^{\rm T}$ to be diagonal. Sections 4.1.1 and 4.1.2 will now focus on details of the EM algorithm.

4.1.1 E-step

As stated at the beginning of this section, the expectation step regarding WEMPCA will be given by the retrieval of the coefficient matrix $\mathbf{C}$ that minimizes equation (7) – or equivalently that optimizes equation (15) – while considering the principal component matrix $\mathbf{P}$ being held fixed. Since each observation – column of $\mathbf{X}$ – is a linear combination of the principal components, finding the solution of equation (15) is equivalent to solving

$$\mathbf{X}_i^{\rm col} \overset{\chi^2}{=} \mathbf{PC}_i^{\rm col};\ \forall i, \tag{16}$$

whose solutions are given by the 'Normal Equations':[2]

$$\mathbf{C}_i^{\rm col} = \left(\mathbf{P}^{\rm T}\mathbf{w}^2\mathbf{P}\right)^{-1}\mathbf{P}^{\rm T}\mathbf{w}^2\mathbf{X}_i^{\rm col}, \tag{17}$$

with $\mathbf{w} = \mathrm{diag}(\mathbf{W}_i^{\rm col})$.

4.1.2 *M-step*

Similarly to the E-step, solution to the M-step – that is the retrieval of $\mathbf{P}$ that optimizes equation (15), given $\mathbf{C}$ – can be decomposed by noting that within each observation, a given variable is the linear combination of the corresponding principal components variables. That is,

$$\mathbf{X}_i^{\rm row} \overset{\chi^2}{=} \mathbf{P}_i^{\rm row}\mathbf{C};\ \forall i, \tag{18}$$

whose solutions are

$$\mathbf{P}_i^{\rm row} = \mathbf{X}_i^{\rm row}\mathbf{w}^2\mathbf{C}^{\rm T}\left(\mathbf{C}\mathbf{w}^2\mathbf{C}^{\rm T}\right)^{-1}, \tag{19}$$

with $\mathbf{w} = \mathrm{diag}\left(\mathbf{W}_i^{\rm row}\right)$.

4.2 Bailey's implementation

As we have seen in Section 4.1, the implementation of Tsalmantza & Hogg (2012) focuses on the solution of equation (7) while the PCA decomposition can be seen as a supplemental step that is external to the EM algorithm. Moreover, a single iteration of the algorithm requires the solution of $N_{\rm obs} + N_{\rm var}$ systems of linear equations, each of size $(N_{\rm comp} \times N_{\rm comp})$, at each iteration of the EM algorithm. This becomes quickly unmanageable regarding huge data sets.

Bailey (2012) takes the pragmatic approach that the sole solution of equation (7) can lead to good insights about the principal components if the latter were fitted individually. This hypothesis is reasonable since this will allow each individual principal component to maximize the variance it accounts for.

The resulting implementation will be similar to the one of Tsalmantza & Hogg (2012) apart from the optimization function of the M-step. Indeed, in order for the principal components to be fitted separately, we have to consider the cross-product decomposition of equation (15), that is

$$\mathbf{X} \overset{\chi^2}{=} \sum_j \mathbf{P}_j^{\rm col}\mathbf{C}_j^{\rm row}, \tag{20}$$

from which each $\mathbf{P}_j^{\rm col}$ has to be individually fitted.

Suppose that we already retrieved the $(j-1)$ first principal components. Let us also assume that the data projection along these $(j-1)$ principal components was already subtracted from the data set, that is

$$\mathbf{X}' = \mathbf{X} - \sum_i^{j-1} \mathbf{P}_i^{\rm col}\mathbf{C}_i^{\rm row}. \tag{21}$$

Then the retrieval of $\mathbf{P}_j^{\rm col}$ based on equation (20) can be decomposed in a way similar to equation (18) as

$$\mathbf{X}_i'^{\rm row} \overset{\chi^2}{=} \mathbf{P}_{ij}\mathbf{C}_j^{\rm row}, \tag{22}$$

[2] This method is known to suffer from numerical instabilities (Press et al. 2002) and is provided for a didactical purpose only. Methods such as SVD for linear least squares must be preferred in order to solve equation (16).

whose solution is straightly given by

$$\mathbf{P}_{ij} = \frac{\mathbf{X}_i'^{\rm row}\mathbf{w}^2\mathbf{C}_j^{\rm row^T}}{\mathbf{C}_j^{\rm row}\mathbf{w}^2\mathbf{C}_j^{\rm row^T}}, \tag{23}$$

with $\mathbf{w} = \mathrm{diag}(\mathbf{W}_i^{\rm row})$.

Equation (21) theoretically ensures that the last-retrieved component, $\mathbf{P}_j^{\rm col}$ will be orthogonal to any previous one. Nevertheless, due to machine round-off errors, this has to be manually checked.

Finally, we have to note that solving equation (20) will not minimize the global χ^2 – as defined by equation (7) – such that the algorithm has to rely on a last E-step at the end of the main EM algorithm.

5 NEW IMPLEMENTATION

Though both mentioned algorithms (Bailey 2012; Tsalmantza & Hogg 2012) correctly find lower-rank orthogonal decompositions that are suitable to explain the whole data set variance at best, none of them assures us that the retrieved principal components will be those that maximize the individual variance described by each of them. These principal components are then efficient at reconstructing the initial data set but are not the best at individually describing the underlying data structure.

The basic idea of this new algorithm is to focus on the maximization of the weighted variance explained by each principal component through the diagonalization of the associated weighted covariance matrix. The resulting principal components will then be those that are the most significant – under the assumption that the definition of the used weighted variance is relevant – in identifying pattern within the data set even if their linear combination is not necessarily the best at explaining the total data set variance as described by equation (7).

In the following, we will consider that the weighted variance of a given discrete variable $\boldsymbol{x}$ having weights $\boldsymbol{w}$ is given by

$$\sigma_x^2 = \frac{\sum_i w_i^2\,(x_i - \bar{x})^2}{\sum_i w_i^2}, \tag{24}$$

where $\bar{x} = \sum_i w_i x_i / \sum_i w_i$ and with the convention that $0/0 = 0$. The latter can be straightforwardly extended to the definition of the weighted covariance between two discrete variables, $\boldsymbol{x}$ and $\boldsymbol{y}$, for which the weights are given by $\boldsymbol{w}^x$ and $\boldsymbol{w}^y$ respectively, that is

$$\sigma_{x,y}^2 = \frac{\sum_i (x_i - \bar{x})\, w_i^x w_i^y\, (y_i - \bar{y})}{\sum_i w_i^x w_i^y}. \tag{25}$$

Based on these definitions, we can write the weighted covariance matrix of a data set $\mathbf{X}$ with associated weights $\mathbf{W}$ as

$$\boldsymbol{\sigma}^2 = \frac{(\mathbf{X}\circ\mathbf{W})\,(\mathbf{X}\circ\mathbf{W})^{\rm T}}{\mathbf{W}\mathbf{W}^{\rm T}}. \tag{26}$$

We know, from the spectral theorem, that there exists an orthogonal matrix $\mathbf{P}$ such that $\boldsymbol{\sigma}^2$ is diagonalized, and consequently that equation (3) is fulfilled (as well as equation 4 if $\mathbf{P}$ is ordered accordingly). This matrix $\mathbf{P}$ will then constitute the principal components of our implementation.

Fig. 2 shows a two-dimensional example of classical principal components $\mathbf{P}'$ that are unable to individually describe the underlying data variance. In this example, we have supposed – for didactical purpose – that the observations corresponding to the bright points have far lower weights associated with the x variable. Weighted principal components, $\mathbf{P}$, diagonalizing $\boldsymbol{\sigma}^2$ as described by equation (26) are given along with the variance it explains (that

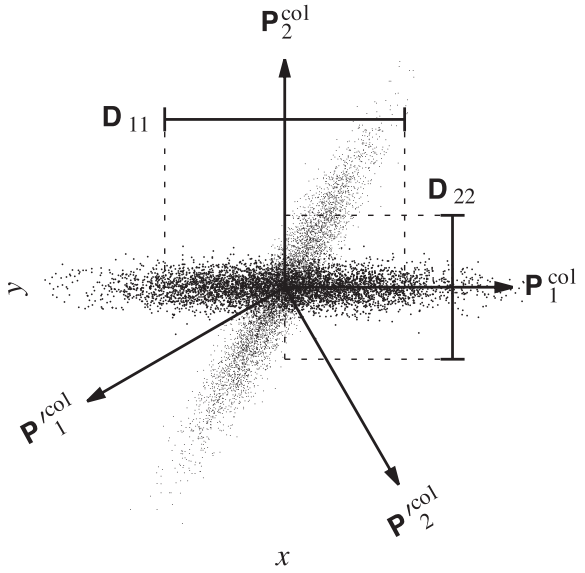

Figure 2. Comparison between weighted principal components $\mathbf{P}$ and classical principal components $\mathbf{P}'$ in presence of underweighted observations corresponding to the brighter points. $|\mathbf{D}_{11}|$ and $|\mathbf{D}_{22}|$ correspond to the variances respectively explained by $\mathbf{P}_1^{\mathrm{col}}$ and $\mathbf{P}_2^{\mathrm{col}}$.

is the diagonal elements of $\mathbf{D} = \mathbf{P}^{\mathrm{T}}\boldsymbol{\sigma}\mathbf{P}$). Note that such principal components maximize the variance explained by each variable as described by equation (24) and consequently set to zero the associated covariance as described by equation (25).

The goal of the algorithm is then to retrieve the dominant eigenvector $\boldsymbol{p}$ from the covariance matrix $\boldsymbol{\sigma}^2$ along with its associated eigenvalue λ, that is the dominant eigenpair $\langle \boldsymbol{p}, \lambda \rangle$. $\boldsymbol{p}$ will then be the principal component explaining most of the data set variance, λ. That is equivalent to find $\boldsymbol{p}$ in

$$\boldsymbol{\sigma}^2 \boldsymbol{p} = \lambda \boldsymbol{p}, \tag{27}$$

such that λ is maximized.[3]

Equation (27), corresponding to the eigenvector definition, is a well-studied problem for which many methods already exist. The reference in the domain is Golub & Van Loan (1996) where the interested reader may find a rather exhaustive list of such methods as well as proofs of the algorithms described hereafter.

Unsurprisingly, in the context or our implementation, we choose the fastest and simplest algorithm called the power iteration method. The idea behind this algorithm is to recognize that given a diagonalizable square matrix $\mathbf{A}$ and a vector $\boldsymbol{u}^{(0)}$ having a non-zero component in the direction of $\boldsymbol{p}$, the iterative relation

$$\boldsymbol{u}^{(k)} = \mathbf{A}\boldsymbol{u}^{(k-1)} = \mathbf{A}^k \boldsymbol{u}^{(0)} \tag{28}$$

will converge to a vector that is proportional to the dominant eigenvector $\boldsymbol{p}$ as $k \to \infty$. Note that in practice, each vector $\boldsymbol{u}^{(k)}$ is normalized to unity in order to avoid numerical round-off errors inherent to the computer representation of large numbers. The final eigenvector will then be given by $\boldsymbol{p} = \boldsymbol{u}^{(k)}/\|\boldsymbol{u}^{(k)}\|$ and the associated eigenvalue by the Rayleigh quotient:

$$\mathcal{R}(\mathbf{A}, \boldsymbol{p}) = \boldsymbol{p} \cdot \mathbf{A}\boldsymbol{p} = \boldsymbol{p} \cdot \lambda \boldsymbol{p} = \lambda. \tag{29}$$

[3] In fact, $\langle -\boldsymbol{p}, -\lambda \rangle$ is also solution of equation (27) but for the sake of clarity, we will only consider the case of positive eigenvalues.

Convergence and assumptions made about this algorithm will be discussed in Section 7.1. Further principal components can be retrieved by considering application of the above-mentioned algorithm to

$$\mathbf{A}' = \mathbf{A} - \lambda \boldsymbol{p} \otimes \boldsymbol{p}, \tag{30}$$

that is the matrix obtained by subtracting the data variance along the found principal components.

5.1 Refinement

As we will see in Section 7.1, the power iteration method may have a slow convergence rate under some particular conditions. Consequently it may be that some vectors did not effectively converge to an eigenvector that would have diagonalized the covariance matrix. Nevertheless, Parlett (1974) proposed an algorithm – called Rayleigh quotient iteration – designed to tackle this kind of problem.

Even if the proof of this algorithm is beyond the scope of this paper, we still mention two basic facts to enable the reader to have a minimal understanding of how it works. First, as we have seen in equation (29), the Rayleigh quotient of a matrix with one of its eigenvector is equal to its associated eigenvalue. Secondly, given a matrix $\mathbf{A}$ with eigenvalues $\lambda_1, \ldots, \lambda_n$, we find that the eigenvalues of the matrix $(\mathbf{A} - d\mathbf{I})^{-1}$ will be $(\lambda_1 - d)^{-1}, \ldots, (\lambda_n - d)^{-1}$. Based upon these facts, we will have that the sequence

$$\boldsymbol{u}^{(k)} = (\mathbf{A} - d^{(k-1)}\mathbf{I})^{-1}\boldsymbol{u}^{(k-1)} \tag{31}$$

$$d^{(k)} = \mathcal{R}(\mathbf{A}, \boldsymbol{u}^{(k)}), \tag{32}$$

where each $\boldsymbol{u}^{(k)}$ is normalized to the unit length, will converge cubically to the eigenpair $\langle \boldsymbol{p}, \lambda \rangle$ that is the nearest – regarding the absolute value of their Rayleigh quotient – from a starting point $\langle \boldsymbol{u}^{(0)}, d^{(0)} \rangle$ as $k \to \infty$. Note that equation (31) is a power iteration towards the eigenvector for which $|\lambda - d^{(k-1)}|$ is minimized.

Finally, the principal coefficients are retrieved by solving $\chi^2 \overset{\chi^2}{=} \mathbf{PC}$ whose 'Normal Equation' solution is given by equation (17).

5.2 Variance regularization

Real-world data often have sparse and unevenly distributed weights such that it may happen for some variables to have their corresponding variances to be based only on a small number of observations. Such situations may become problematic since these few variables will have a strong impact on the resulting first principal components in a way that is irrespective to their total weight.

Such 'overweighted' variables can be damped by using a regularization factor within the expression of the weighted covariance, $\sigma^2_{x,y}$, as defined by equation (25). The resulting regularized weighted covariance will take the form

$$\sigma^2_{x,y}(\xi) = \left[\sum_i w_i^x \sum_i w_i^y \right]^{\xi} \sigma^2_{x,y}, \tag{33}$$

where the regularization parameter ξ allows us to control the damping strength.

The typical value of the regularization parameter, ξ, goes from zero, where we get back to the classical behaviour of the algorithm, to two for a strong damping of these rare variables. Conversely one might want to highlight such underrepresented variables by setting the regularization parameter to a negative value.

6 COMPARISON

The performance of the method described in this paper was assessed regarding the two previously described algorithms, namely the one of Bailey (2012) and the one of Tsalmantza & Hogg (2012). The choice of these algorithms comes from the fact that they are fairly competitive and have goals that are comparable to those of the new algorithm. All methods were tested on both simulated data as well as on real observational ones.

6.1 Simulated data

Simulated data consist in random linear combinations of 10 orthogonal basis functions. These bases are produced by taking 10 shifted sine functions having periods between 0.2π and 2π and by applying a Gram–Schmidt orthogonalization process to the latter. Resulting observations are then sampled over 100 evenly spaced points in the interval $[0, 2\pi]$. To each variable, x, within each observation, $\boldsymbol{x}$, we also add a Gaussian noise having a standard deviation given by

$$\sigma_x = \sigma_{\text{in}} (1 + \sigma_{\text{obs}}) (1 + \mathcal{U}_\sigma) \max |\boldsymbol{x}| , \tag{34}$$

where σ_{in} is a user-provided parameter corresponding to the desired noise amplitude, σ_{obs} is an observation-specific noise ponderation uniformly drawn from $[-0.1, 0.1]$ and $\mathcal{U}_\sigma$ is a uniform random variable corresponding to the noise dispersion within observations and taking values in the range $[-0.1, 0.1]$. The weight associated with the variable x will then be set to $1/\sigma_x$.

Finally, we discard N_{bad} contiguous and randomly positioned variables from each observation. The latter will be used to assess performances of the various algorithms on data extrapolation while having their weights equal to zero during the PCA retrieval phase. Examples of such simulated data are illustrated in Fig. 3.

In order to perform the comparison with other algorithms, we built a given number of data sets, N_{set}, each containing 1000 observations. Each of these sets was duplicated and altered through the realization of various values of the simulation parameters σ_{in} and N_{bad}. We then retrieved, for each of the mentioned algorithms, the five first principal components out of the resulting altered data sets and computed estimators based on the following χ^2 definition:

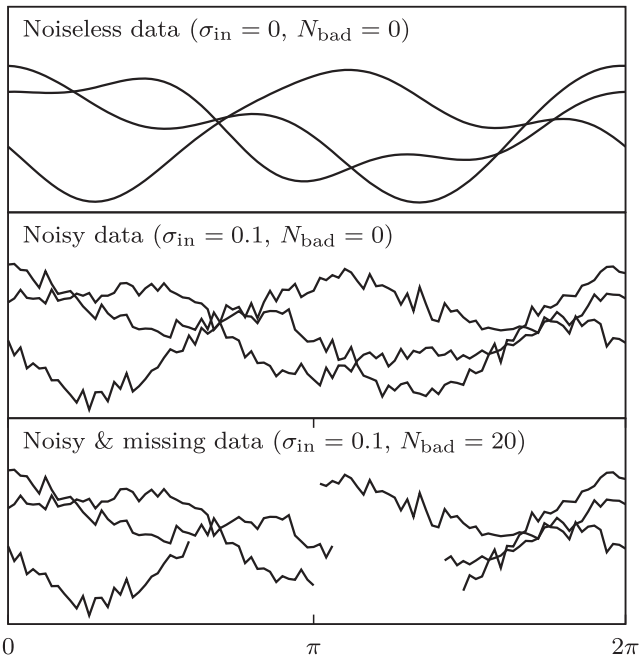

Figure 3. Examples of simulated data illustrating the effect of the simulation parameters σ_{in} and N_{bad}.

$$\chi^2 = \frac{\sum_{ij} [\mathbf{W} \circ (\mathbf{X} - \mathbf{PC})]^2_{ij}}{\sum_{ij} \mathbf{W}^2_{ij}} . \tag{35}$$

The following estimators were computed: χ^2_{fit}, the chi-square of the data set for which weights associated with the discarded variables are set to zero and χ^2_{test} where only rejected variables are considered and for which weights associated with the unrejected variables are set to zero. Let us note that χ^2_{fit} will typically account for the quality of the fit while χ^2_{test} will account for the quality of the extrapolation. If these estimators are to be computed based on the number of data set, N_{set}, we use a 3σ-clipped mean over all the χ^2_{fit} and χ^2_{test} associated with each data set, namely $\langle \chi^2_{\text{fit}} \rangle$ and $\langle \chi^2_{\text{test}} \rangle$.

For completeness, only data sets having $N_{\text{bad}} \leq 50$ will be discussed here. This decision comes from the fact that efficiently estimating the principal components of such sparse data sets while having 'only' 1000 observations is a really tricky task strongly depending on the design of these data sets. Furthermore, as we will see in Section 7.4, the Bailey and Tsalmantza algorithms are big time consumers such that dealing with bigger data sets quickly becomes unmanageable.

Since all the studied algorithms are based on iterative procedures, we have to take into account the convergence criterion for each of them. To this aim, we performed a preliminary study whose goal is to determine the minimal number of iterations needed by each algorithm in order to reach convergence. This was assessed by running 100 times each algorithm on data sets similar to those previously described and by setting the initial eigenvectors estimates to random values within each run. In order to make sure we can model unseen and potentially more complex data sets, the number of iterations was set to twice the found number of iterations needed to converge, giving respectively 500, 500 and 10^4 iterations for Bailey, Tsalmantza and our algorithm (without refinement).

Regarding the quality of the fit, χ^2_{fit}, differences between the various algorithms are fairly low even if – as expected – our algorithm is proven to have somewhat larger χ^2_{fit} with a higher dispersion over all the data sets. Fig. 4 shows the behaviour of the mean χ^2_{fit} regarding two common cases, namely the case of increasing noise and no missing data and the case of moderate noise ($\sigma_{\text{in}} = 0.1$) with increasing number of missing data. Each point on these graphs is averaged over $N_{\text{set}} = 1000$ data sets. Practically, one cannot distinguish the various algorithms if all data are present since differences are in $\mathcal{O}(10^{-6})$. In presence of missing data, differences start to be noticeable but still reasonable with differences in $\mathcal{O}(10^{-3})$. For a didactical purpose, let us note that if only the first component – out of the five retrieved – was considered then our algorithm would have had a better or equal χ^2_{fit} than Bailey and Tsalmantza implementations over all the values of the simulation parameters σ_{in} and N_{bad}.

Regarding now the quality of the extrapolation, differences are more noticeable. Fig. 5 illustrates these discrepancies according to the averaged χ^2_{test} over $N_{\text{set}} = 100$ data sets having reasonable χ^2_{fit} in each of the σ_{in} and N_{bad} simulation parameters. At first glance, our new algorithm shows a globally better $\langle \chi^2_{\text{test}} \rangle$ while suffering less from missing data: the latter being dominated by noise within the data. The Tsalmantza's algorithm shows good performances but has a stronger dependence on missing data, reaching a maximum $\langle \chi^2_{\text{test}} \rangle = 0.372$ while Bailey's algorithm shows strong numerical instabilities for $N_{\text{bad}} > 40$ that makes it unable to converge and

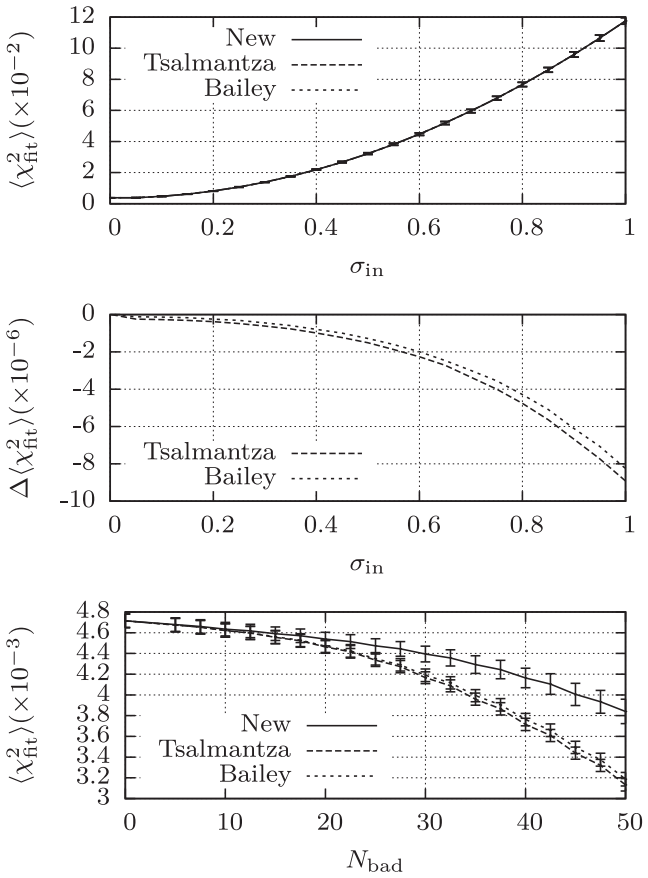

Figure 4. Averaged $\chi^2_{\rm fit}$ for increasing noise without missing data (top). Subtraction of Bailey and Tsalmantza $\chi^2_{\rm fit}$ from our algorithm $\chi^2_{\rm fit}$ (middle). Averaged $\chi^2_{\rm fit}$ with moderate noise and an increasing number of missing data (bottom).

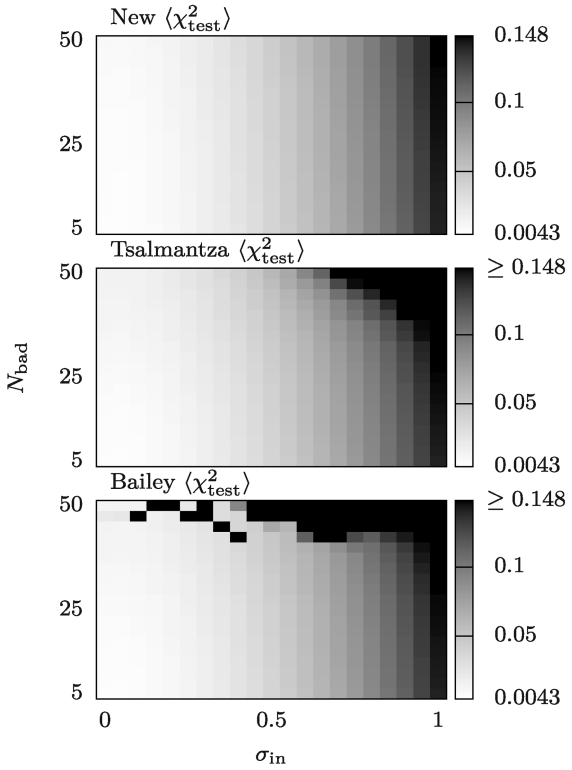

Figure 5. χ^2 maps of averaged $\chi^2_{\rm test}$ regarding $\sigma_{\rm in}$ and $N_{\rm bad}$ simulation parameters.

makes it reach a maximum $\langle\chi^2_{\rm test}\rangle \approx 3\times10^9$ for $\sigma_{\rm in}=0.9$ and $N_{\rm bad}=50$. Fig. 6 illustrates in more detail the behaviour of $\langle\chi^2_{\rm test}\rangle$ averaged over $N_{\rm set}=1000$ data sets in the common case of moderate noise ($\sigma_{\rm in}=0.1$) with an increasing number of missing data. For clarity, the plot of Bailey having $N_{\rm bad}>40$ has been removed from the graph.

Finally, let us note that data sets having $\sigma_{\rm in}=0$ and $N_{\rm bad}=0$ can be solved using a classical PCA algorithm. Consequently, both the variance explained by each individual component and the total variance can be simultaneously optimized. We will thus find – in this particular case – that all algorithms will provide us with identical results as it was already suggested in Fig. 4. Similarly, if we choose to retrieve a single component, even with noisy and/or missing data, the algorithms of Bailey and Tsalmantza will maximize the variance explained by this component, the latter will thus match the first component that would have been retrieved by our algorithm.

6.2 Observational data

Comparisons against a concrete case were performed using the SDSS DR10Q quasar catalog from Pâris et al. (2014). Out of the 166 583 QSO spectra present in the initial data release, 18 533 were rejected either due to spectra bad quality, strong uncertainties in redshift determination, presence of BAL or insufficient number of high-S/N points. The remaining spectra were set to the rest frame; the continuum was then subtracted and the spectra were normalized such as to have a zero mean and a variance of 1. Finally, visual inspection showed that in some cases the continuum was badly fitted such that the variance within these spectra can mainly be attributed to this error; these regions were removed using a k-sigma clipping algorithm for each variable among all observations. There remain 148 050 spectra having observed wavelengths between 4000 and 10 000 Å and for which the variance within each spectrum is thought to be mainly caused by genuine signals. In the following tests, the 10 first components will be retrieved from each algorithm. We also suppose, as previously, that $\xi=0$.

The number of iterations associated with each algorithm was assessed by using a subset of the above-described data set ($1000 \leq \lambda_{\rm rest} \leq 2000$ Å, $3 \leq z \leq 4$) and by running 10 times the various algorithms on it with random initial principal components. Convergence was assessed by minimizing the variance amongst the final principal components within these 10 runs. The fact that only a subset of the above-described data set was used can be explained by the large amount of time needed by each algorithm to run as well as by the fact that – as we will see – Tsalmantza and Bailey algorithms often fail to converge in presence of a large amount of missing data. The results show that the number of iterations chosen in Section 6.1 also match the complexity of this problem.

The initial test to be performed is similar to tests performed in Section 6.1 in the sense that PCA were retrieved for all algorithms using the region 5000–9000 Å in observed wavelength while regions 4000–5000 and 9000–10 000 Å are rejected and kept to assess the quality of the extrapolation. Fig. 7 shows an example of such a spectrum along with the fits by the various algorithms. The resulting SDSS DR10Q data set $\chi^2_{\rm fit}$ and $\chi^2_{\rm test}$, as defined by equation (35), are given in Table 1.

We see that $\chi^2_{\rm fit}$ is quite stable for all the algorithms while being a bit higher for our algorithm as expected. Now watching at the

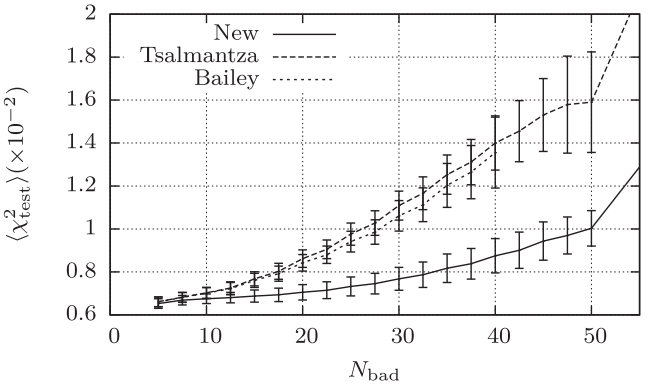

Figure 6. Averaged $\chi^2_{\rm test}$ for moderate noise with an increasing number of missing data.

extreme differences of $\chi^2_{\rm test}$ for the various algorithms and given the fact that we know it to be an estimator that is quite sensitive to outliers, we found relevant to see how these discrepancies are distributed among the observations. To this aim, we computed a $\chi^2_{\rm test}$ per-observation whose distribution is summarized in Fig. 8.

Given the significant number of observations used in this test, we can already draw some general trends. Beside the fact that the other algorithms have a better $\chi^2_{\rm fit}$, they often fail to satisfactorily extrapolate the spectra. Indeed, for the Bailey algorithm, 81 per cent of observations have $\chi^2_{\rm test} \geq 5$ with a peak up to 9.59×10^{15} (median = 8.24×10^4); the Tsalmantza algorithm has respectively 33 per cent of observations having $\chi^2_{\rm test} \geq 5$ and a peak up to 5.78×10^8 (median = 1.789) while our new algorithm has only 1.4 per cent of 'outliers' with a maximal peak of 172 (median = 1.021). If the observations are individually compared, our algorithm has a lower $\chi^2_{\rm test}$ in 90 per cent of the time regarding Bailey algorithm and 68 per cent of the time regarding Tsalmantza algorithm while in other cases differences remain quite moderate with a median $\Delta\chi^2_{\rm test} \approx 0.3$. For completeness, we have to note that these huge discrepancies can be mainly attributed to the large amount of missing data (76 per cent) within the resulting rest-frame DR10Q data set.

Another quality one would often desire is the ability to have the most general and representative set of principal components able to model unseen observations. In this optic, two tests were performed.

First, we split the initial DR10Q data set into two subsamples each of 74 025 spectra spanning from 4000 to 10 000 Å (hereafter

Table 1. SDSS DR10Q data set fit and extrapolation chi-squares followed by per-observation median chi-squares of extrapolation and associated ratio of outliers.

	New	Tsalmantza	Bailey
Data set $\chi^2_{\rm fit}$	0.107	0.088	0.094
Data set $\chi^2_{\rm test}$	1.064	2×10^5	8×10^{12}
Median $\chi^2_{\rm test}$	1.021	1.789	8×10^4
Ratio of observations having $\chi^2_{\rm test} \geq 5$	0.014	0.33	0.81

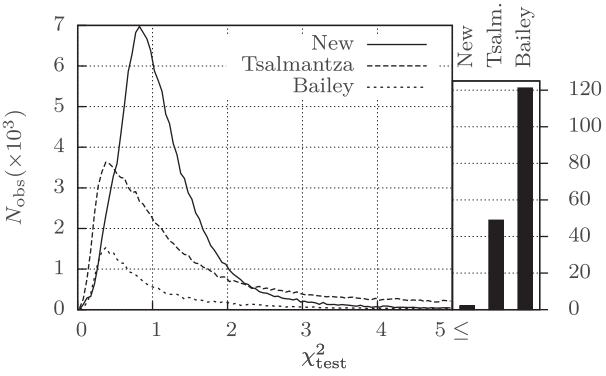

Figure 8. Distribution of the observation's $\chi^2_{\rm test}$. Right-hand part represents the number of QSOs ($\times 10^3$) having $\chi^2_{\rm test} \geq 5$.

LS1 and LS2). From each of these subsamples, we extract ∼8600 spectra for which the rest-frame region from 1000 to 2000 Å is entirely covered – and consequently 'without' missing data – and ∼62 000 spectra for which the mentioned region is only partially covered–and thus 'with' missing data. Note that according to the selected observed wavelengths and the previous definitions, spectra without missing data will correspond to those having $3 \leq z \leq 4$ and spectra with missing data to those having $1 \leq z \leq 9$.

The consistency and convergence of each algorithm were then tested by retrieving 10 principal components out of LS1 and LS2 data sets taken with and without missing data. Prior to discussing the results, we have to mention that the components may be swapped from data sets with missing data to the corresponding ones without missing data. This is easily explained by the fact that some patterns may be highlighted due to the uneven coverage of the wavelength range while being damped in the total variance if a full coverage

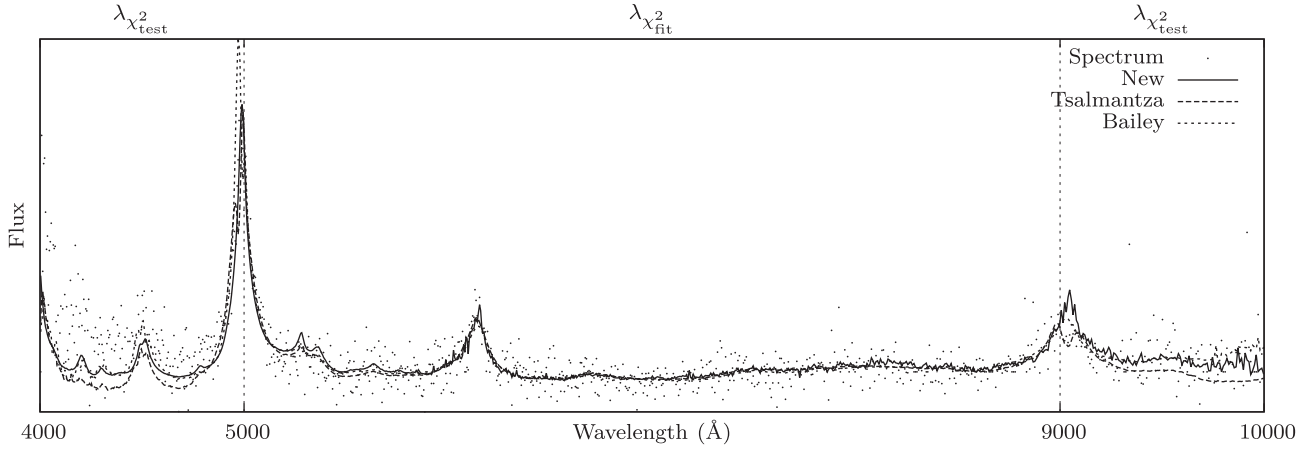

Figure 7. Example of SDSS DR10Q spectrum fits. $\lambda_{\chi^2_{\rm test}}$ and $\lambda_{\chi^2_{\rm fit}}$ denote regions used to compute, respectively, $\chi^2_{\rm test}$ and $\chi^2_{\rm fit}$ for this observation.

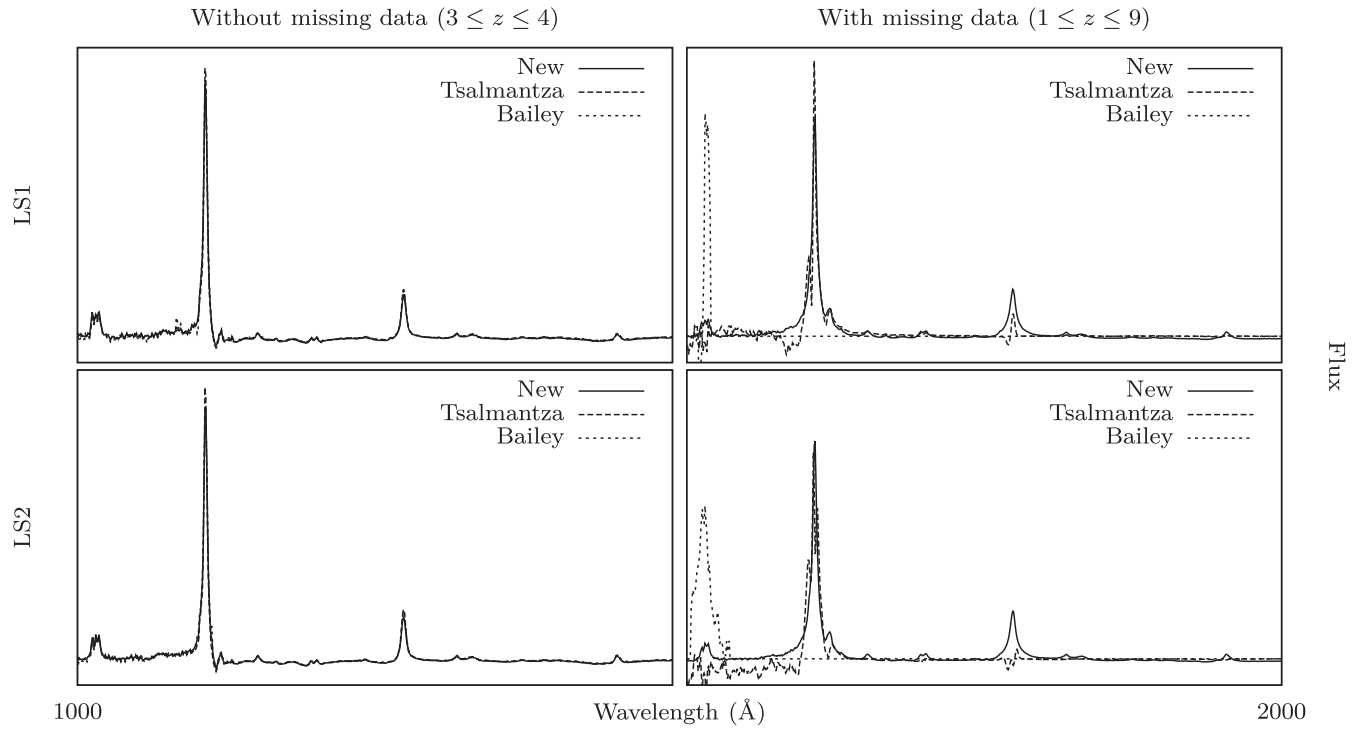

Figure 9. Principal components retrieved from two independent data sets (LS1 and LS2) coming from the DR10Q catalogue. The left-hand part corresponds to the dominant eigenvectors extracted from spectra having no missing data while the right-hand part are the associated components coming from spectra having missing fluxes.

is considered. Fig. 9 illustrates the results of this test regarding the first component of the data set without missing data and associated component with missing data. We see that in absence of missing data, all algorithms are consistent and succeed in converging towards the dominant eigenvector while in the case of missing data only our algorithm shows both a good consistency and a good convergence. More precisely, regarding our algorithm, the mean differences per-point between LS1 and LS2 are in $\mathcal{O}(10^{-6})$ for both subsets with and without missing data. Differences noticed between subsets, that is mainly the general larger equivalent width and the stronger N V λ1240 Å emission line of the subset with missing data, are consistent with the underlying data set and make it able to model the larger variety of spectra coming from the larger redshift range. Let us mention, that no significant inconsistencies have been noticed up to the sixth component for LS1 and LS2 on both subsets for our algorithm. Concerning the Bailey's algorithm, we see, as already suggested in Section 6.1, that it fails to converge if a large amount of data is missing. Tsalmantza's implementation shows a better convergence but fails to correctly reproduce some emission lines (O VI λ1033, N V λ1240, C IV λ1549 and C III λ1908 Å) as well as the region ≤1200 Å in subsets with missing data. Similarly to Bailey's algorithm, analysis show that these are due to a convergence problem occurring on sparse data sets and arising from numerical instabilities.

In the second test performed, 10 principal components were retrieved from a set of 98 700 spectra spanning from 4000 to 10 000 Å (hereafter LS3). We tested the fact that the principal components are the best at individually describing the underlying data variance by fitting the remaining 49 350 spectra (hereafter VS) using a subset of five components retrieved from LS3. That is, if the variance explained by each individual component is maximized then the fit of a smaller number of components to a similar data set should be minimal. In this case, if we consider computing χ^2 of VS data set as defined by equation (35) for each of the algorithms, we have respectively $\chi^2 = 0.372$ for the Bailey algorithm, $\chi^2 = 0.316$ for the Tsalmantza algorithm and $\chi^2 = 0.309$ for our algorithm, thus supporting our last quality criterion. Note again that if all components were considered, Bailey and Tsalmantza algorithms would have had a better resulting χ^2.

7 DISCUSSION

7.1 Convergence and uniqueness

At first glance it is surprising that the power iteration algorithm works at all, but in practice, it is easily demonstrated. Consider a diagonalizable square matrix $\mathbf{A}$ of size $(n \times n)$ having eigenvectors $\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n$ and associated eigenvalues $\lambda_1, \ldots, \lambda_n$ where $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_n|$. Since the eigenvectors are orthogonal between each other, we can write the starting vector of the power iteration algorithm as

$$\boldsymbol{u}^{(0)} = c_1 \boldsymbol{p}_1 + \cdots + c_n \boldsymbol{p}_n. \tag{36}$$

Then we will have that the vector at iteration k given by

$$\begin{aligned} \boldsymbol{u}^{(k)} &= \mathbf{A}^k \sum_{i=1}^{n} c_i \boldsymbol{p}_i = \sum_{i=1}^{n} c_i \lambda_i^k \boldsymbol{p}_i \\ &= \lambda_1^k \left[c_1 \boldsymbol{p}_1 + \sum_{i=2}^{n} c_i \left(\frac{\lambda_i}{\lambda_1} \right)^k \boldsymbol{p}_i \right]. \end{aligned} \tag{37}$$

We see that $\boldsymbol{u}^{(k)}$ will converge to $\lambda_1^k c_1 \boldsymbol{p}_1$ as $k \to \infty$ under the following conditions: $c_1 \neq 0$, i.e. the starting vector has a non-zero component in the direction of the dominant eigenvector and $|\lambda_1| > |\lambda_2|$, i.e. the data set has only one dominant eigenvalue. Also note that, in this case, the rate of convergence to the dominant eigenvector will be principally given by λ_1/λ_2 and that in the case of $\lambda_1 = \lambda_2$, the uniqueness of the solution is not guaranteed as it depends on the starting vector $\boldsymbol{u}^{(0)}$. Refinement seen in Section 5.1 obeys to the same conditions. Additionally, the dominant eigenvalue of $(\mathbf{A} - d\mathbf{I})^{-1}$ will tend to ∞ as $d \to \lambda$, where $\mathbf{A}$ is the covariance matrix, d the current 'eigenvalue' (i.e. the Rayleigh quotient of $\mathbf{A}$ and $\boldsymbol{u}$, the current 'eigenvector') and λ a real eigenvalue of $\mathbf{A}$. This may thus lead to numerical instabilities that can strongly deteriorate the vector used in the next step of the Rayleigh quotient iteration.

That being said, failures against convergence can be easily checked. For example, a satisfactory solution should have $\mathbf{A}\boldsymbol{p} \approx \lambda \boldsymbol{p}$; and if not, the power iteration algorithm can be resumed with another starting vector $\boldsymbol{u}^{(0)}$. Under the condition that we have $\lambda_1 = \lambda_2$, we will have an infinite number of eigenvectors that are equally good at describing the data set variance, thus choosing one over another is irrelevant. Concretely, during the ∼100 000 tests performed in the context of this paper, no such problems arose. Checks performed on eigenvectors from Section 6.1 and 6.2 show that the latter are orthogonal to a machine precision of $\mathcal{O}(10^{-16})$ and that the weighted covariance matrix (as described by equation 26) is diagonalized with the same precision. Nevertheless, people wanting extreme reproducibility and/or secure convergence may still use the SVD in order to extract all eigenvectors from equation (26) at the expense of a lower flexibility (see Sections 7.2 and 7.3).

7.2 A priori eigenvectors

In some situations, one may already have an approximation of the wanted principal components corresponding to a given data set. Typical examples include principal components update according to new observations added to the data set or in a real case we encountered, the use of SDSS DR9Q eigenvectors to the DR10Q data set. These will constitute a priori eigenvectors that it would be regrettable not to use. The design of our algorithm allows us to easily take benefits of these a priori vectors. Instead of using a random vector, $\boldsymbol{u}^{(0)}$, in equation (28) one can straightforwardly substitute the known vectors to the random starting vector usually used. Doing so will typically decrease the number of iterations needed to converge towards the new vectors. Beware that vectors may be swapped between a priori eigenvectors and effective data set eigenvectors – as encountered in Section 6.2 – and consequently iterations needed to converge will be the same as if we had used a random starting vector. In this case, one can perform no power iteration and use only the refinement seen in Section 5.1 in order to converge to the nearest eigenvector.

7.3 Smoothing eigenvectors

It sometimes happens for some variables within eigenvectors to exhibit some sharp features that we know to be artefacts. These occur mainly in cases where we have noise that is comparable in amplitude to the data variance, data regions covered only by a few numbers of observations or data sets containing corrupted observations. Again the flexibility of our algorithm allows us to efficiently deal with these drawbacks. Suppose that we retrieved such a 'noisy' eigenpair $\langle \boldsymbol{p}, \lambda \rangle$ from a given data set whose covariance matrix is given by $\mathbf{A}$, then, before removing the variance in the direction of the found component through equation (30), one can filter $\boldsymbol{p}$ thanks to any existing smoothing function. Assuming that the resulting vector, $\boldsymbol{p}'$, is near $\boldsymbol{p}$ regarding the norm of their difference, we can suppose that the variance accounted for by $\boldsymbol{p}'$ (i.e. its 'eigenvalue') is similar to the one of $\boldsymbol{p}$ and then subtract it along the direction of $\boldsymbol{p}'$ by

$$\mathbf{A}' = \mathbf{A} - \lambda \boldsymbol{p}' \otimes \boldsymbol{p}'. \tag{38}$$

Note that the kind of smoothing function to use is highly data-dependent and no rule of thumb exists. Nevertheless, a quite general and commonly used filter producing efficient results in the field of QSOs can be found in Savitzky & Golay (1964). Finally, let us note that, when applying a filter, the orthogonality of the components has to be manually checked and that they will obviously not diagonalize the covariance matrix anymore.

7.4 Scaling performance

PCA is often used in cases where we have a lot of observations and a reasonable number of variables, typically we have $N_{\rm obs} \gg N_{\rm var}$. We know, for example, the classical algorithm to require $\mathcal{O}(N_{\rm var}^3)$ basic operations in order to solve for the eigenvectors of the covariance matrix and $\mathcal{O}(N_{\rm var}^2 N_{\rm obs})$ operations to build this matrix. In the following, we will compute the algorithmic complexity of the various explored algorithms in a similar way.

The Tsalmantza algorithm requires for each iteration (hereafter $N_{\rm iter}$) the solution of linear systems of equations for each of the E and M-step, respectively, in each observation and in each variable. We will then have that its algorithmic complexity is given by $\mathcal{O}(N_{\rm iter} N_{\rm comp}^3 (N_{\rm obs} + N_{\rm var}))$. The Bailey algorithm will be identical except for the M-step that will be in $\mathcal{O}(N_{\rm var} N_{\rm obs} N_{\rm comp})$ and can thus be discarded, giving $\mathcal{O}(N_{\rm iter} N_{\rm comp}^3 N_{\rm obs})$. Finally, our algorithm mainly requires the building of the covariance matrix, $N_{\rm iter}$ matrix multiplications, potentially $N_{\rm refine}$ matrix inversions in order to refine the eigenvectors and a final single E-step similar to the one for the previous algorithm (discarded here), thus giving $\mathcal{O}(N_{\rm var}^2 N_{\rm obs} + N_{\rm iter} N_{\rm var}^2 N_{\rm comp} + N_{\rm refine} N_{\rm var}^3)$.

We see that in case $N_{\rm obs} \gg N_{\rm var}$, our algorithm is much faster than the other ones as the computing time is mainly spent in the covariance matrix building. As an illustration, if we take data similar to those described in Section 6.1 with $N_{\rm obs} = 10\,000$, PCA retrieval takes ∼140 s for the Bailey and Tsalmantza algorithms (with $N_{\rm iter} = 100$) and ∼3 s for the new one (with $N_{\rm iter} = 10\,000$) on a 2.4 Ghz CPU.

8 CONCLUSIONS

We presented a new method for computing principal components based on data sets having noisy and/or missing data. The underlying ideas are intuitive and lead to a fast and flexible algorithm that is a generalization of the classical PCA algorithm. Unlike existing methods, based on lower-rank matrix approximations, the resulting principal components are not those that aim to explain at best the whole variance of a given data set but rather those that are the most suitable in identifying the most significant patterns out of the data set while explaining most of its variance. The main benefits of the current implementation are a better behaviour in presence of missing data as well as faster run times on data set having a large amount of observations. Privileged problems encompass data set extrapolation, patterns analysis and principal component usage over similar data sets. We assessed the algorithm performance on simulated data as

well as on QSO spectra to which many applications are already foreseen in the field of the *Gaia* mission.

ACKNOWLEDGEMENTS

The author acknowledges support from the ESA PRODEX Programme 'Gaia-DPAC QSOs', and from the Belgian Federal Science Policy Office.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation and the U.S. Department of Energy Office of Science. The SDSS-III web site is http://www.sdss3.org/.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofisica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington and Yale University.

REFERENCES

Bailer-Jones C. A. L. et al., 2013, A&A, 559, A74
Bailey S., 2012, PASP, 124, 1015
Beale E. M. L., Little R. J. A., 1975, J. R. Stat. Soc. B, 37, 129
Dempster A. P., Laird N. M., Rubin D. B., 1977, J. R. Statist. Soc. B, 39, 1
Gabriel K. R., Zamir S., 1979, Technometrics, 21, 489
Golub G. H., Van Loan C. F., 1996, Matrix Computations, 3rd edn. The Johns Hopkins Univ. Press, London
Greenacre M. J., 1984, Theory and Applications of Correspondence Analysis. Academic Press, London
Hotelling H., 1933, J. Educ. Psychol., 24, 417
Jarvis M., Jain B., 2004, ApJ, submitted
Jolliffe I. T., 2002, Principal Component Analysis, 2nd edn. Springer, New York
Pâris I. et al., 2014, A&A, 563, A54
Parlett B. N., 1974, Math. Comput., 28, 679
Pearson K., 1901, Phil. Mag., 2, 559
Press W. H., Tuekolsky S. A., Vetterling W. T., Flannery B. P., 2002, Numerical Recipes in C++: The Art of Scientific Computing, 2nd edn. Cambridge Univ. Press, New York
Roweis S., 1997, CNS Technical Report CNS-TR-97-02, EM Algorithm for PCA and SPCA. Available at: http://cs.nyu.edu/~roweis/papers/empca.pdf
Savitzky A., Golay M. J. E., 1964, Anal. Chem., 36, 1627
Schlens J., 2009, A Tutorial on Principal Component Analysis, preprint (arXiv:1404.1100)
Srebro N., Jaakkola T., 2003, in AAAI Press, ed., Proc. 20th ICML, Weighted Low Rank Approximation. Menlo Park, p. 720
Tipping M. E., Bishop C. M., 1999, J. R. Stat. Soc. B, 21, 611
Tsalmantza P., Hogg D. W., 2012, ApJ, 753, 122
Wentzell P. D., Andrews D. T., Hamilton D. C., Faber K., Kowalski B. R., 1997, J. Chemometr., 11, 339
Yip C. W. et al., 2004, AJ, 128, 2603

APPENDIX A: MATLAB/OCTAVE SOURCE CODE

```
%   Function designed to find the weighted PCA of a
% given dataset such that the resulting principal
% components diagonalize the associated weighted
% covariance matrix and that the associated principal
% coefficients are the linear coefficients of these
% principal components that best match the dataset
% in a least-squares sense.
%
% INPUT:
%   X       : Dataset matrix from which the weighted
%             mean was already subtracted. Rows
%             correspond to variables, columns to
%             observations.
%   W       : Weight matrix (same size as X).
%   ncomp   : Number of principal components to
%             retrieve.
%   niter   : Number of iterations (allows
%             convergence to the dominant principal
%             component)
%   nrefine : Number of refinments (allows refinement
%             of the dominant principal component)
%   xi      : Regularization factor controlling the
%             influence of rare variables
% OUTPUT:
%   P       : Principal components arranged in columns
%   C       : Principal coefficients such that X = P*C
%             in a least-squares sense
% AUTHOR:
%   Ludovic Delchambre
%   Extragalactic Astrophysics and Space Observations,
%   Institute of Astrophysics and Geophysics,
%   University of Liege, Belgium
function [P, C] = WPCA(X, W, ncomp, niter, nrefine, xi)
  nvar = size(X,1); nobs = size(X,2);
  P=zeros(nvar,ncomp);C=zeros(ncomp,nobs);
  ws = sum(W,2);
  covar = ((ws * ws').^xi).*((X.*W)*(X.*W)')./(W*W');
  covar(isnan(covar)) = 0;
  for i=1:ncomp
    u = ones(nvar,1) / nvar;
    for j=1:niter
      u = covar*u;
      d = sqrt(u'*u);
      u = u / d;
    end
    d = u'*covar*u;
    for j=1:nrefine
      u = inv(covar - d * eye(nvar)) * u;
      u = u / sqrt(u'*u);
      d = u'*covar*u;
    end
    covar = covar - u*d*u';
    P(:,i) = u;
  end
  for i=1:nobs
    w = diag(W(:,i)).^2;
    C(:,i) = inv(P'*w*P) * (P'*w*X(:,i));
  end
end
```

This paper has been typeset from a TEX/LATEX file prepared by the author.

2.3 Weighted phase correlation

The major drawbacks of the χ^2 curve approach we described in section 2.1.2 stand both in its quadratic time dependency as well as in the numerical instabilities that may arise once the normal equations are used so as to fit the templates to the observations. Beside these weaknesses, we still have to note that this technique provides extremely fair predictions on the redshift and that it allows a straight interpretation of the errors that may arise during predictions (e.g. multiple solutions having a similar χ^2 are readily identified). This ultimately yields strong diagnostic tools on these errors. For these reasons, it has a long history in astronomy and more particularly in the processing of spectroscopic data (e.g. Simkin 1974). It has been consequently employed in numerous spectral reduction pipelines out of which the baryon oscillation spectroscopic survey (BOSS) pipeline of the SDSS (Bolton et al. 2012) stands to be the most famous one. In the latter implementation an SVD decomposition (equation 2.10) is advantageously replacing the Cholesky decomposition of the normal equation so as to gain numerical stability. It is effectively known that the use of such orthogonal decomposition methods leads to a high numerical stability at the expense of higher execution times (Golub & Van Loan 1996; Press et al. 2002).

With this aim in mind, we first re-write the χ^2 curve equation (equation 2.7) in terms of matrix operations as

$$\chi^2(k) = \left\| \mathbf{W}\boldsymbol{s} - \mathbf{W}\tilde{\mathbf{T}}\boldsymbol{a}(k) \right\|^2 = \left\| \boldsymbol{y} - \tilde{\mathbf{X}}\boldsymbol{a}(k) \right\|^2 , \tag{2.15}$$

where $\mathbf{W}$ is the diagonal matrix of weights associated with the observed spectrum $\boldsymbol{s}$ (i.e. $\mathbf{W}_{ii} = 1/\sigma_i$), $\tilde{\mathbf{T}}_{ij} = \mathbf{T}_{(i+k)j}$ is the matrix of shifted templates and $\tilde{\mathbf{X}} = \mathbf{W}\tilde{\mathbf{T}}$. For reasons that will become clear in Paper II, we consider here an orthogonal decomposition of the form

$$\tilde{\mathbf{X}} = \mathbf{Q}\mathbf{R} \tag{2.16}$$

where $\mathbf{Q}$ is orthogonal and $\mathbf{R}$ is upper triangular. Substituting the latter decomposition into 2.15 yields

$$\chi^2(k) = \|\boldsymbol{y} - \mathbf{Q}\mathbf{R}\boldsymbol{a}(k)\|^2 = \|\boldsymbol{y} - \mathbf{Q}\boldsymbol{b}(k)\|^2 = \|\boldsymbol{y}\|^2 - \|\boldsymbol{b}(k)\|^2 . \tag{2.17}$$

The χ^2 curve can hence be computed (up to a constant term) through the sole knowledge of $\|\boldsymbol{b}(k)\|^2 = \|\mathbf{R}\boldsymbol{a}(k)\|^2$ as $\|\boldsymbol{y}\|^2$ is shift-independent (i.e. it is constant over k). The global maximum of $\|\boldsymbol{b}(k)\|^2$ being then associated with the global minimum of the χ^2 curve, we will accordingly term $\mathrm{ccf}(k) \equiv \|\boldsymbol{b}(k)\|^2$ the *cross correlation function* (CCF) of $\boldsymbol{s}$ against $\mathbf{T}$ in the following.

In Paper II, we show how the CCF can be computed over all shifts $k \in \{0, m-1\}$ in $\mathcal{O}(m \log m)$ flops. A rigorously derivation of this numerical complexity being also provided in section A.3 of the appendices. Along with this low numerical complexity, one can finally afford to use this weighted phase correlation (WPC) algorithm in the field of the Gaia mission. Indeed, using the numerical complexity we derived in section A.3 as well as the number of samples and templates we probe in section 2.1.2 and equation 2.9 (i.e. $N = 7709$ samples with 6 templates), we have that the numerical complexity of this algorithm will be in the order of 3×10^6 flops whereas 10^9 flops are allocated for the processing of a single source. This equivalently translates in an execution time of ~ 2 μs per source in DPCC out of 0.6 ms available. Paper II further derives the uncertainties associated with the redshift predictions as well as it explores the conditions for rising warning flags.

Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **460,** 2811–2821 (2016) doi:10.1093/mnras/stw1025
Advance Access publication 2016 May 1

Redshift determination through weighted phase correlation: a linearithmic implementation

L. Delchambre★

Institut d'Astrophysique et de Géophysique, Université de Liège, Allée du 6 Août 17, B-4000 Sart Tilman (Liège), Belgium

Accepted 2016 April 28. Received 2016 April 25; in original form 2016 January 7

ABSTRACT

We present a new algorithm having a time complexity of $\mathcal{O}(N \log N)$ and designed to retrieve the phase at which an input signal and a set of not necessarily orthogonal templates match best in a weighted chi-squared sense. The proposed implementation is based on an orthogonalization algorithm and thus also benefits from high numerical stability. We apply this method successfully to the redshift determination of quasars from the twelfth Sloan Digital Sky Survey (SDSS) quasar catalogue and derive the proper spectral reduction and redshift selection methods. Derivations of the redshift uncertainty and the associated confidence are also provided. The results of this application are comparable to the performance of the SDSS pipeline, while not having a quadratic time dependence.

Key words: methods: data analysis – galaxies: distances and redshifts – quasars: general.

1 INTRODUCTION

The advent of extremely large spectroscopic surveys like the Sloan Digital Sky Survey (SDSS), which includes more than 2×10^6 high-resolution spectra over 5200 deg^2 of the sky (Alam et al. 2015), or the *Gaia* space mission, which will provide, by the end of 2018, 150×10^6 low-resolution spectra (de Bruijne 2012), provides us with unique opportunities to obtain a statistical view of the kinds of object present in the Universe, along with some of their fundamental characteristics. These play a key role in the answers to some of the currently most important astrophysical questions, like the evolution scenarios of the Galaxy and the Universe or its accelerated expansion (Perryman et al. 2001; Aubourg et al. 2015).

Along with these large surveys comes an impressive continuous flow of data that has to be treated right on time through huge dedicated processing centres. One of the most important tasks amongst the spectral reduction processes comprises the classification of objects and determination of their astrophysical parameters (APs). More specifically, in the case of extragalactic objects, this information depends critically on the availability of reliable redshift estimates.

Redshift determination, even if apparently straightforward, is in practice a challenging problem, for which numerous solutions have been proposed.

(i) Visual inspection procedures: a skilled observer can guess the APs of any object efficiently and can deal with any unexpected cases like corrupted/missing emission lines, spectra superposition or non-physical solutions. Obviously, this choice is infeasible for large surveys, though the analysis of any sufficiently large subset is invaluable, as it can serve as input to sophisticated computer algorithms that will try to mimic this human expertise. This is the solution undertaken by Pâris et al. (in preparation) regarding the redshift of quasars and, accordingly, it will be used within this article as the default quasar spectral library.

(ii) Matching of spectral lines: this method consists of extracting some significant patterns from the input spectra and then trying to match them to known emission/absorption lines. This procedure has been used for a long time, but has been shown to be restricted to relatively high signal-to-noise ratio (SNR) spectra (Machado 2013).

(iii) Computer learning methods: the goal here is to make the algorithm guess the relations that exist between some characteristics of already-reduced objects (e.g. observed wavelengths and fluxes) and the parameters of interest (e.g. redshift coming from a visual inspection procedure). The aim is then to apply these relations to the case of objects with parameters that are still unknown. Interested readers may find the descriptions of many such algorithms in Bishop (2006). Note that, depending on its complexity, the relation guessed may be non-physical and hard to interpret, leading to suboptimal or potentially unrealistic predictions. This is the reason why these should preferably be used for the case of highly non-linear problems for which no other (fast) solution exists.

(iv) Phase correlation: the idea here is to find the optimal correlation of a given observation against one or more templates in order to determine its redshift. Based upon the ability of these templates to match the observations, from the physical nature of this solution and the shortcomings of the previously mentioned alternatives, we will consider this to be the most trustworthy automated procedure for redshift determination.

Based on the work of Brault & White (1971), Simkin (1974) first suggested the use of the fast Fourier transform (FFT) as an

★ E-mail: ldelchambre@ulg.ac.be

efficient way of finding the redshift of galaxies based on their cross-correlation with a single template. Tonry & Davis (1979) later derived the formulation associated with the resulting redshift uncertainties, which were refined further by Heavens (1993). Finally, Glazebrook (1997) generalized the cross-correlation technique to the case of templates coming from the principal components analysis (PCA) decomposition of spectral libraries. Although it is currently the most widespread technique for redshift determination, the latter actually suffers from some well-known drawbacks (see Section 3.2.1). The solution to these problems comes from the use of a weighting scheme associated with the observed spectrum, as implemented in Bolton et al. (2012). Unfortunately, this solution has a quadratic time dependence that makes it fairly time-consuming.

The method proposed in the present work overcomes this high numerical complexity and was developed in the framework of the *Gaia* astrophysical parameters inference system (Bailer-Jones et al. 2013) and, more specifically, within the field of the quasar classification module (QSOC), the goal of which is to find the APs associated with the quasars that *Gaia* will detect. In this domain, the time constraints imposed by the *Gaia* mission restricted us to the use of computer learning methods, but in the end the advent of this new method will allow us to predict fair and fast redshift estimates for the upcoming *Gaia* data releases.

Section 2 explains the conventions used within this article. Section 3 gives a brief review of the phase-correlation and PCA techniques, aimed at better understanding their main limitations. We have developed a fast solution to the problem of the weighted phase correlation in Section 4. Tests against real cases are then performed within Section 5, while extensions of the algorithm presented are discussed in Section 6. Finally, we conclude in Section 7.

2 NOTATION

This article uses the following notations: vectors are in bold italic, $\boldsymbol{x}$, with x_i being the element i of the vector $\boldsymbol{x}$. Matrices are in uppercase boldface or are explicitly stated; i.e. $\mathbf{X}$, from which the ith column will be denoted by $\mathbf{X}_i^{\mathrm{col}}$ and the element at row i, column j by $\mathbf{X}_{ij}$. In the following, we will consider the problem of finding the optimal offset between an observed spectrum composed of N_s samples and N_T templates of size N_p by probing various shift estimates, Z. By considering the zero-padding necessary in order for these to be properly used within the Fourier domain, we see that the template matrices, $\mathbf{P}$ and $\mathbf{T}$, will be of size $(N \times N_\mathrm{T})$ with $N = N_\mathrm{s} + N_p$. Similarly, the observation vector, $\boldsymbol{s}$, will be of size N as well. Note that, in order for the redshift to turn into a simple offset, we will have to consider a logarithmic wavelength scale. If not stated otherwise, matrices and vectors having a tilde on top of them (e.g. $\tilde{\mathbf{T}}$) will be specific to a given shift try, $Z \in 0 \cdots N-1$. Amongst commonly used operators, $\boldsymbol{a} \cdot \boldsymbol{b}$ denotes the inner product of $\boldsymbol{a}$ and $\boldsymbol{b}$ and $\boldsymbol{a} \otimes \boldsymbol{b}$ their outer product; $\|\boldsymbol{a}\|$ the Euclidian norm of $\boldsymbol{a}$ and $\boldsymbol{a}^*$ its complex conjugate. Finally, $\mathcal{F}[\boldsymbol{x}]$ and $\mathcal{F}^-[\boldsymbol{x}]$ correspond respectively to the discrete Fourier transform (hereafter DFT) and inverse DFT of $\boldsymbol{x}$.

3 PHASE CORRELATION USING PCA

As already stated, the most commonly used technique for quasi-stellar object (QSO) redshift determination consists of finding the best correlation of an observed spectrum against templates coming from the PCA decomposition of a rest-frame spectral library. More specifically, these are based on spectra sampled on a uniform logarithmic wavelength scale, such that the observed wavelength, λ_obs, can be related to the rest-frame wavelength, λ_rest, through the QSO redshift, z, as a simple offset:

$$\log \lambda_\mathrm{obs} = \log \lambda_\mathrm{rest} + \log(z+1). \tag{1}$$

In the following, we make a brief review of the two above-mentioned techniques, which should provide the reader with insights about their way of working and is aimed at better understanding their main limitations regarding the redshift estimation of QSOs.

3.1 Principal components analysis

PCA is a well-known technique designed to extract a set of templates – the principal components – from a typically huge set of data while keeping most of its variance (Pearson 1901). These principal components will hence be those that are the best suited, in order to highlight the most important patterns out of the input data set.

Mathematically, the goal of the PCA is to find a decomposition of an input matrix $\mathbf{X}$, from which we have subtracted the mean observation, into

$$\mathbf{X} = \mathbf{PC}, \tag{2}$$

such that

$$\mathbf{D} = \mathbf{P}^\mathrm{T}\mathbf{X}\mathbf{X}^\mathrm{T}\mathbf{P} = \mathbf{P}^\mathrm{T}\boldsymbol{\sigma}^2\mathbf{P} \tag{3}$$

is diagonal and for which

$$\mathbf{D}_i \geq \mathbf{D}_j;\ \forall i < j. \tag{4}$$

$\mathbf{P}$, the matrix of the eigenvectors of $\boldsymbol{\sigma}^2$, is then called the matrix of principal components; $\mathbf{C}$ is the associated matrix of principal coefficients and $\mathbf{D}_i$ are the eigenvalues of the covariance matrix, $\boldsymbol{\sigma}^2$. Note that, according to the spectral theorem,[1] $\mathbf{P}$ will be orthonormal, such that we have

$$\mathbf{C} = \mathbf{P}^\mathrm{T}\mathbf{X}. \tag{5}$$

From this orthonormality and from equation (4), we have that the linear combination of the first principal components of $\mathbf{P}$ with the associated principal coefficients of $\mathbf{C}$ will constitute the best linear combination in order to fit $\mathbf{X}$ in a least-squares sense.

An illustrative example of PCA decomposition is given in Fig. 1. This figure is based on spectra covering the rest-frame wavelength range 1100–2000 Å coming from the SDSS DR12 quasar catalogue (Pâris et al. in preparation). Notice how the main QSO emission lines are modelled by the various components as a way to grab the variance coming from the great diversity of shapes encountered within the spectral library. Readers wishing more information on PCA decomposition are invited to read Jolliffe (2002) for a deep analysis of the technique or Schlens (2014) for an accessible tutorial.

The application of this technique to the analysis of QSO spectra was first covered by Francis et al. (1992); Yip et al. (2004) later adapted it to the case of SDSS DR1 quasar classification and redshift determination, while Cabanac, de Lapparent & Hickson (2002) performed a similar work based upon spectra coming from the Large Zenith Telescope survey, the spectral resolution of which ($\lambda/\Delta\lambda \sim 40$) is of the same order of magnitude as that of the red and blue photometers of *Gaia* (Bailer-Jones et al. 2013).

[1] Any real symmetric matrix is diagonalized by a matrix of its eigenvectors.

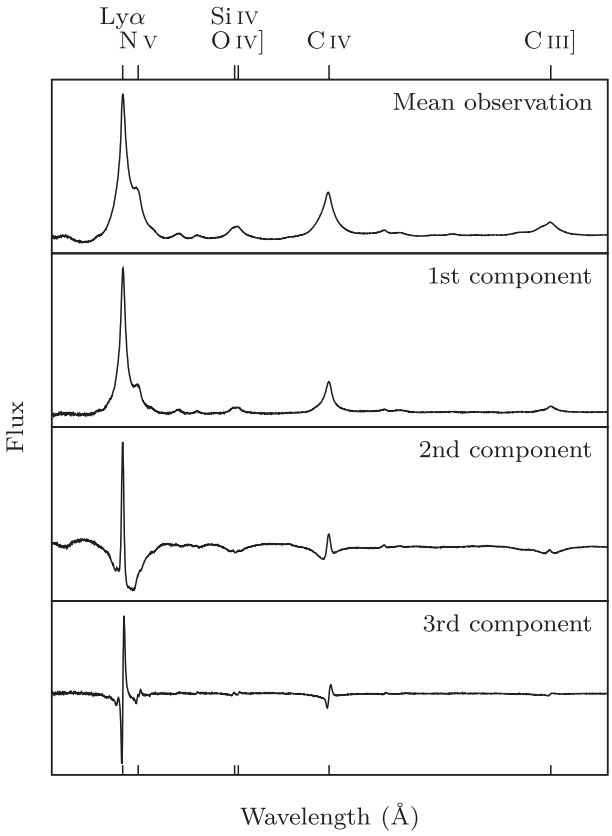

Figure 1. Rest-frame mean observation and first principal components coming from the PCA decomposition of TypeI/II SDSS DR12 quasar spectra having $2.3 \leq z \leq 2.4$ (24 939 spectra).

3.1.1 Weighted PCA

One of the main limitations of the classical PCA method is that it does not make any distinction between variance coming from noise and variance coming from a genuine signal. Furthermore, in its naive form, it does not know how to deal with missing data. This last point is particularly crucial in the field of high-redshift surveys, where the observed wavelength ranges may not overlap from object to object.

A straightforward approach so as to avoid these shortcomings stands in the use of a weighting scheme that allows each flux within each spectrum to retain its own uncertainty while performing the PCA decomposition. Such a fully-weighted PCA (WPCA) method was first described in the astronomical literature by Tsalmantza & Hogg (2012) and was later refined by Bailey (2012). In the field of the present study, we will use the implementation described in Delchambre (2015); this choice comes mainly from its high numerical stability. This method is based on the diagonalization of the weighted variance–covariance matrix as defined by

$$\boldsymbol{\sigma^2} = \frac{(\mathbf{X} \circ \mathbf{W})\,(\mathbf{X} \circ \mathbf{W})^{\mathrm{T}}}{\mathbf{W}\mathbf{W}^{\mathrm{T}}}, \tag{6}$$

where $\circ$ represents the elementwise product of two matrices and where $\mathbf{X}$ is supposed to have a weighted mean observation of zero. The decomposition of $\boldsymbol{\sigma^2}$ into a diagonal matrix of eigenvalues, $\mathbf{D}$, and a matrix of orthonormal principal components, $\mathbf{P}$, is then performed using either a combination of two spectral decomposition methods, namely the power iteration method followed by the Rayleigh quotient iteration one, or the singular value decomposition (SVD). This technique allows us to retrieve the fairest components (i.e. those for which uncertainties are taken into account) without having to worry about missing data: this case is the limiting case of weights equal to zero. Consequently, this method will be used through the rest of this article as the default process, in order to retrieve the principal components.

3.2 Phase correlation

The goal of the phase-correlation algorithm is to find the optimal shift between a set of orthonormal templates – or a sole unit-length template – $\mathbf{P}$ and a given observation, $\boldsymbol{s}$, that has been shifted relative to $\mathbf{P}$. The way to proceed is to compute for each potential shift, Z, the linear least-squares solution of the shifted templates, $\tilde{\mathbf{P}}_{ij} \equiv \mathbf{P}_{(i+Z)j}$, against the observation, so as to find the offset having the minimal resulting chi-square. More concisely, this is equivalent to finding the minimal shift-dependent chi-square, as defined by

$$\chi^2(Z) = \left\| \boldsymbol{s} - \tilde{\mathbf{P}}\boldsymbol{a}(Z) \right\|^2, \tag{7}$$

where $\boldsymbol{a}(Z)$ contains the optimal linear coefficients in order to fit $\boldsymbol{s}$ based on $\tilde{\mathbf{P}}$.

Extending the work of Simkin (1974), Glazebrook (1997) noticed that in the case of orthonormal templates, like the PCA principal components, equation (7) becomes

$$\chi^2(Z) = \|\boldsymbol{s}\|^2 - \|\boldsymbol{a}(Z)\|^2. \tag{8}$$

Consequently, equation (7) will be minimal for an associated maximal $\|\boldsymbol{a}(Z)\|^2$. Moreover, due to the orthonormality of $\mathbf{P}$, we have

$$\boldsymbol{a}(Z) = \tilde{\mathbf{P}}^{\mathrm{T}}\boldsymbol{s}. \tag{9}$$

More specifically, regarding the ith linear coefficient, $a_i(Z)$, we have

$$a_i(Z) = \sum_j \mathbf{P}_{(j+Z)i}\, s_j = \left(\mathbf{P}_i^{\mathrm{col}} \star \boldsymbol{s} \right)_Z. \tag{10}$$

We recognize equation (10) as being the correlation of the vector $\mathbf{P}_i^{\mathrm{col}}$ with $\boldsymbol{s}$, which can hence be computed efficiently in the Fourier domain. In Brault & White (1971), interested readers will find exhaustive hints about the practicalities surrounding the Fourier implementation of equation (10). Let us simply point out that both vectors, $\mathbf{P}_i^{\mathrm{col}}$ and $\boldsymbol{s}$, have to be extended and zero-padded so as to deal with the periodic nature of the DFT. Note that, in the remainder of this article, the curve obtained after evaluating $\|\boldsymbol{a}(Z)\|^2$ at each Z will be termed the *cross-correlation function* (CCF).

A sub-sampling precision on the offset can be gained by considering the fit of a continuous function in the vicinity of the maximal peak of the discrete CCF. Simkin (1974) supposed this peak to be Gaussian-profiled, but, with the aim of having a model-independent estimate of Z, we will follow Tonry & Davis (1979) and use a quadratic curve fitting that will allow us to take potential asymmetries in the fitted peak into account.

3.2.1 Practicalities

Some issues highlighted in Glazebrook (1997) are the subtraction of the QSO continuum and rest-frame mean spectrum from the observed spectrum. Here, the first issue was solved by the use of a dedicated method that allows us to fit the QSO continuum in a fast and redshift-independent way. This method will be described further in Section 5.1. The second issue is often overcome by omitting the subtraction of the mean spectrum from the input data set. We have to note that this omission typically degrades the ability of the PCA

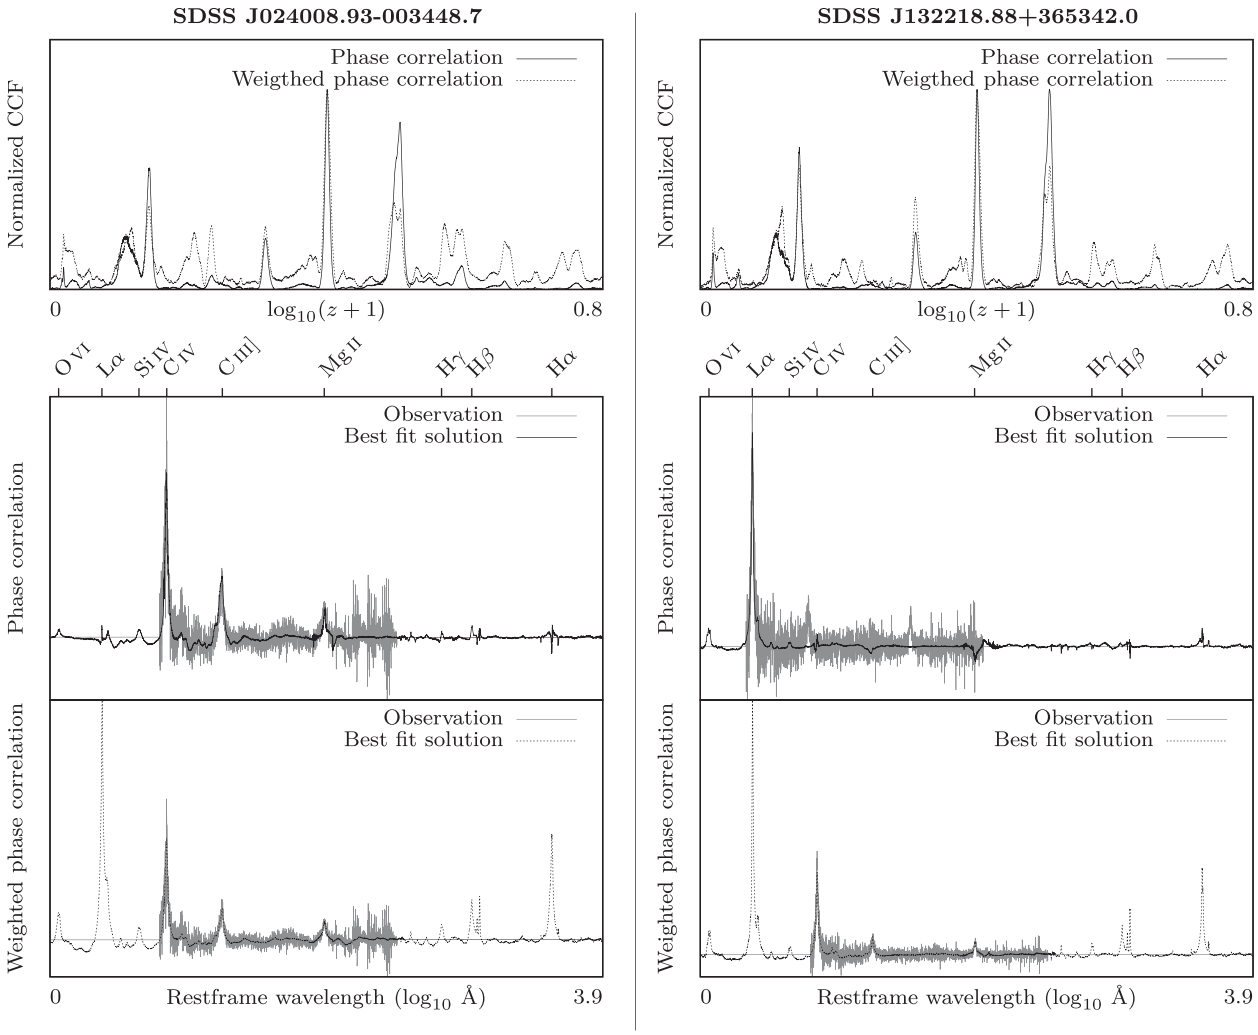

Figure 2. Examples of cross-correlation functions coming from the phase correlation and weighted phase correlation of two quasar spectra against the first ten SDSS DR12 principal components plus mean observation (top). The best-fitting solutions associated with the maximal peak of the phase correlation (middle) and weighted phase correlation (bottom) are also provided.

decomposition to extract the most significant patterns out of this input data set. Another solution would have been to alter the mean spectrum so as to make it orthonormal to the template components, $\mathbf{P}$, thanks to the use of a Gram–Schmidt orthogonalization process (Press et al. 2002), for example, and furthermore to consider it as being an additional template. This solution will be adopted here for the use of the phase-correlation algorithm.

Finally, the major drawback of the implementation of Glazebrook (1997) stands in the fact that the observed spectra typically span only a small part of the template spectra, such that the CCF will be computed over a substantial number of unknown points. As a consequence, the fit of the input spectra will be disrupted by the 'flattening' of the principal components over the unobserved wavelengths. Fig. 2 illustrates the result of the phase-correlation algorithm, along with the best-fitting solution associated with the maximal peak of the CCF. Notice how the solutions are flattened over unobserved wavelengths. More precisely, considering the observation of SDSS J024008.93−003448.7, the Lyα, Hα and Hβ emission lines are strongly damped, despite the fact that the optimal shift was found, while for the observation of SDSS J132218.88+365342.0 this 'flattening' has led to an ambiguity in the CCF that leads to an erroneous shift estimate. Additionally, uncertainties for the observed fluxes are often available and will not be used within this implementation.

4 WEIGHTED PHASE CORRELATION

With the aim of dealing efficiently with the previously mentioned problem of unobserved wavelengths and neglected uncertainties, we will use a χ^2 formulation similar to equation (7), but with fluxes weighted according to the observed spectrum wavelengths. Also, we will drop the orthonormality constraint on the fitted templates, since in any case the previously mentioned weighting will break it down. We will then have the following objective formula:

$$\chi^2(Z) = \left\| \mathbf{W}\boldsymbol{s} - \mathbf{W}\tilde{\mathbf{T}}\boldsymbol{a}(Z) \right\|^2 = \left\| \boldsymbol{y} - \tilde{\mathbf{X}}\boldsymbol{a}(Z) \right\|^2, \tag{11}$$

where $\mathbf{W}$ is the diagonal matrix of weights associated with $\boldsymbol{s}$ and $\tilde{\mathbf{T}}_{ij} \equiv \mathbf{T}_{(i+Z)j}$ is the shifted matrix of (not necessarily orthonormal) template observations. The fastest solution in order to minimize equation (11) for a given Z stands in the use of a Cholesky decomposition of the design matrix, $\tilde{\mathbf{X}}^{\mathrm{T}}\tilde{\mathbf{X}}$, followed by a

forward–backward substitution associated with the image vector $\tilde{\mathbf{X}}^{\mathrm{T}}\boldsymbol{y}$ (Press et al. 2002). We have to note that this approach is known to suffer from numerical instabilities (Golub & Van Loan 1996; Press et al. 2002) and is solely provided here as a comparison point regarding its computational performances. From a practical viewpoint, slower but more stable methods based on the orthogonalization of $\tilde{\mathbf{X}}$ should be preferred.

From a computational point of view, the evaluation of equation (11) for each Z will require $\mathcal{O}\left(N^2\right)$ floating operations (flops),[2] these being dedicated mainly to the building of the design matrices. This relatively high complexity constitutes the main limitation of this implementation and makes it unaffordable for the tight processing of a large survey like *Gaia*. Nonetheless, it has been proven to provide fair redshift estimates and is currently being used effectively in the SDSS-III spectral classification redshift measurement pipeline with a singular value decomposition (SVD) of $\tilde{\mathbf{X}}$ advantageously replacing the Cholesky decomposition of the design matrix (Bolton et al. 2012).

4.1 Orthogonal decomposition approach

The previous section points out the risks encountered while using a naive approach for solving the normal equations associated with equation (11). In this optics, let us explore the effect of the orthogonalization of $\tilde{\mathbf{X}}$ on the latter equation. For this purpose, let us detail the QR decomposition of $\tilde{\mathbf{X}} = \mathbf{QR}$,[3], which is such that

$$\mathbf{Q}^{\mathrm{T}}\tilde{\mathbf{X}} = \mathbf{Q}_{\mathbf{N_T-1}}\cdots\mathbf{Q_1}\tilde{\mathbf{X}} = \mathbf{Q}_{\mathbf{N_T-1}}\cdots\mathbf{Q_i}\tilde{\mathbf{X}}_{\mathbf{i}} = \mathbf{R}, \tag{12}$$

where $\mathbf{R}$ is an upper triangular matrix of size $(N \times N_{\mathrm{T}})$ and where each $\mathbf{Q_i}$ is a Householder reflection matrix designed to annihilate the elements below the ith row of the ith column of $\tilde{\mathbf{X}}_{\mathbf{i}}$ (Press et al. 2002). More precisely, given $\tilde{\mathbf{X}}'_{\mathbf{i}}$, the not-already upper-triangular part of $\tilde{\mathbf{X}}_{\mathbf{i}}$, we have

$$\mathbf{Q}_i = \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} - 2\boldsymbol{v}_i \otimes \boldsymbol{v}_i / \|\boldsymbol{v}_i\|^2 \end{pmatrix} = \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}'_i \end{pmatrix}, \tag{13}$$

with

$$\boldsymbol{v}_i = \boldsymbol{x}_i \pm \|\boldsymbol{x}_i\|\,\boldsymbol{e}_1; \tag{14}$$

here, $\boldsymbol{x}_i$ is the first column of $\tilde{\mathbf{X}}'_{\mathbf{i}}$ and $\boldsymbol{e}_1$ the first row of the identity matrix. For numerical stability reasons, the choice between subtraction and addition in equation (14) should be matched to the sign of the first element of $\boldsymbol{x}_i$ (Press et al. 2002).

By using such a decomposition, we find that equation (11) becomes

$$\chi^2(Z) = \left\|\boldsymbol{y} - \mathbf{QR}\boldsymbol{a}(Z)\right\|^2 = \left\|\boldsymbol{y} - \mathbf{Q}\boldsymbol{b}(Z)\right\|^2, \tag{15}$$

with the last $N - N_{\mathrm{T}}$ elements of $\boldsymbol{b}(Z)$ being zeros. The point is now to recognize equation (15) as being the weighted counterpart of equation (7), such that the first N_{T} elements of $\boldsymbol{b}(Z)$ will be equal to the first N_{T} elements of $\mathbf{Q}^{\mathrm{T}}\boldsymbol{y}$, the computation of which can be performed efficiently by successive multiplication of each of the $\mathbf{Q_i}$ with the associated $\boldsymbol{y}_i \equiv \mathbf{Q_{i-1}}\cdots\mathbf{Q_1}\boldsymbol{y} = \left(b_1(Z)\cdots b_{i-1}(Z)\boldsymbol{y}'_i\right)$, rather than by computing the general $\mathbf{Q}$ matrix explicitly. This efficiency comes mainly from the following facts.

(i) We do not need to compute any $\mathbf{Q}'_{\mathbf{i}}$ explicitly, since we have that the jth column of the product $\mathbf{Q}'_{\mathbf{i}}\tilde{\mathbf{X}}'_{\mathbf{i}}$ will be given by

$$\left(\mathbf{Q}'_{\mathbf{i}}\tilde{\mathbf{X}}'_{\mathbf{i}}\right)^{\mathrm{col}}_j = \left(\tilde{\mathbf{X}}'_{\mathbf{i}}\right)^{\mathrm{col}}_j - 2\frac{\boldsymbol{v}_i \cdot \left(\tilde{\mathbf{X}}'_{\mathbf{i}}\right)^{\mathrm{col}}_j}{\boldsymbol{v}_i \cdot \boldsymbol{v}_i}\boldsymbol{v}_i \tag{16}$$

and, similarly,

$$\mathbf{Q}'_{\mathbf{i}}\boldsymbol{y}'_i = \boldsymbol{y}'_i - 2\frac{\boldsymbol{v}_i \cdot \boldsymbol{y}'_i}{\boldsymbol{v}_i \cdot \boldsymbol{v}_i}\boldsymbol{v}_i. \tag{17}$$

In other words, the computation of $\mathbf{Q}'_{\mathbf{i}}\boldsymbol{y}'_i$ and of any column of the products $\mathbf{Q}'_{\mathbf{i}}\tilde{\mathbf{X}}'_{\mathbf{i}}$ is now reduced to a single inner product (the product $\boldsymbol{v}_i \cdot \boldsymbol{v}_i$ being common to all multiplications, it can be pre-computed) and to a single vector subtraction.

(ii) We do not need to compute any $\mathbf{R}_{ij}$. Differently stated, we do not need to compute the first row, nor the first column, of any $\mathbf{Q}'_{\mathbf{i}}\tilde{\mathbf{X}}'_{\mathbf{i}}$.

This implementation, termed the 'factorized QR algorithm', has a total complexity that can compete with the Cholesky solution of the normal equations while gaining in numerical stability. In practical terms, it is of low interest for us, since it remains a quadratic problem that is consequently outside the limits of the time processing required by the *Gaia* tight data reduction.

Let us note that equation (15) still provides us with a weighted formulation of the CCF, i.e. $\|\boldsymbol{b}(Z)\|^2$, such that we can already investigate the effects of the weighting on the best-fitting solutions at its maximal peak and on the CCF itself. As illustrated in Fig. 2, the fitted spectra no longer exhibit border flattening and, thanks to this, the maximal peaks are now clearly identified. More particularly, regarding the observation of SDSS J132218.88+365342.0, the optimal peak of the CCF turns out to be unambiguously identified thanks to the use of this weighted formulation of the phase correlation.

4.1.1 *Factorized QR algorithm with look-up tables*

The quadratic nature of the factorized QR algorithm comes from the large number of inner products involved in the computation of the first N_{T} elements of each $\boldsymbol{b}(Z)$. More specifically, by developing each inner product coming from equations (16) and (17) in the case of the initial reduction $\tilde{\mathbf{X}}_{\mathbf{2}} \equiv \mathbf{Q_1}\tilde{\mathbf{X}}$ and associated image production $\boldsymbol{y}_2 \equiv \mathbf{Q_1}\boldsymbol{y}$, we obtain

$$\boldsymbol{v_1} \cdot \boldsymbol{v_1} = 2\alpha\left(\alpha + \tilde{\mathbf{X}}_{11}\right), \tag{18}$$

$$\boldsymbol{v_1} \cdot \boldsymbol{y} = \alpha y_1 + \tilde{\mathbf{X}}^{\mathrm{col}}_1 \cdot \boldsymbol{y}, \tag{19}$$

$$\boldsymbol{v_1} \cdot \tilde{\mathbf{X}}^{\mathrm{col}}_j = \alpha\tilde{\mathbf{X}}_{1j} + \tilde{\mathbf{X}}^{\mathrm{col}}_1 \cdot \tilde{\mathbf{X}}^{\mathrm{col}}_j, \tag{20}$$

with $\alpha = \mathrm{sgn}\left(\tilde{\mathbf{X}}_{11}\right)\left(\tilde{\mathbf{X}}^{\mathrm{col}}_1 \cdot \tilde{\mathbf{X}}^{\mathrm{col}}_1\right)^{1/2}$. At this point, it should be noted that

$$\tilde{\mathbf{X}}^{\mathrm{col}}_i \cdot \tilde{\mathbf{X}}^{\mathrm{col}}_j = \boldsymbol{w}^2 \cdot \left(\tilde{\mathbf{T}}^{\mathrm{col}}_i \circ \tilde{\mathbf{T}}^{\mathrm{col}}_j\right) = \sum_{k=1}^{N} w_k^2\left(\mathbf{T}^{\mathrm{col}}_i \circ \mathbf{T}^{\mathrm{col}}_j\right)_{k+Z} \tag{21}$$

and that

$$\tilde{\mathbf{X}}^{\mathrm{col}}_i \cdot \boldsymbol{y} = \left(\boldsymbol{w}^2 \circ \boldsymbol{s}\right) \cdot \tilde{\mathbf{T}}^{\mathrm{col}}_i = \sum_{k=1}^{N} w_k^2 s_k \mathbf{T}_{(k+Z)i}, \tag{22}$$

with $\boldsymbol{w} \equiv \mathrm{diag}\left(\mathbf{W}\right)$. We can readily see that equations (21) and (22) can be computed efficiently in the Fourier domain. In order to take

[2] The interested reader will find information and references regarding the various algorithmic complexities used within this article in Golub & Van Loan (1996).

[3] Note that we dropped the upper tilde for the purpose of clarity.

benefits from it, let us define the *look-up table* of the inner products of $\tilde{\mathbf{X}}$ with itself as

$$\tilde{\mathbf{L}}_{ij} = \tilde{\mathbf{X}}_i^{\mathrm{col}} \cdot \tilde{\mathbf{X}}_j^{\mathrm{col}} = \mathcal{F}^- \left[\mathcal{F}\left[\mathbf{T}_i^{\mathrm{col}} \circ \mathbf{T}_j^{\mathrm{col}}\right]^* \circ \mathcal{F}[\mathbf{W}^2] \right]_Z \tag{23}$$

and the one containing the inner products of $\tilde{\mathbf{X}}$ with $\boldsymbol{y}$ as

$$\tilde{\boldsymbol{l}}_i = \tilde{\mathbf{X}}_i^{\mathrm{col}} \cdot \boldsymbol{y} = \mathcal{F}^- \left[\mathcal{F}\left[\mathbf{T}_i^{\mathrm{col}}\right]^* \circ \mathcal{F}[\mathbf{W}^2 \boldsymbol{s}] \right]_Z. \tag{24}$$

Note that in the latter equations, $\mathcal{F}\left[\mathbf{T}_i^{\mathrm{col}} \circ \mathbf{T}_j^{\mathrm{col}}\right]^*$ and $\mathcal{F}\left[\mathbf{T}_i^{\mathrm{col}}\right]^*$ are template-specific and hence can be computed in advance.

Explicitly stated, these look-up tables allow us to have, for any shift estimates, Z, an instantaneous evaluation of all the inner products associated with the initial reduction process. Furthermore, thanks to the $\mathbf{Q_1}$ orthonormality, we see that the look-up tables associated with $\tilde{\mathbf{X}}_2$ are also given by $\tilde{\mathbf{L}}$ and $\tilde{\boldsymbol{l}}$. Consequently, we can easily compute the inner product of $\tilde{\mathbf{X}}_2'$ with itself based on $\tilde{\mathbf{L}}$ as

$$\left(\tilde{\mathbf{X}}_2'\right)_i^{\mathrm{col}} \cdot \left(\tilde{\mathbf{X}}_2'\right)_j^{\mathrm{col}} = \tilde{\mathbf{L}}_{ij} - \mathbf{R}_{1i}\mathbf{R}_{1j} \qquad \forall i, j \tag{25}$$

and, in the same way, we can compute the inner products of $\tilde{\mathbf{X}}_2'$ with $\boldsymbol{y}_2'$ based on $\tilde{\boldsymbol{l}}$ as

$$\left(\tilde{\mathbf{X}}_2'\right)_i^{\mathrm{col}} \cdot \boldsymbol{y}_2' = \tilde{\boldsymbol{l}}_i - \mathbf{R}_{1i} b_1(Z) \qquad \forall i. \tag{26}$$

Equations (25) and (26) allow us to process each subsequent $\tilde{\mathbf{X}}_i'$ and $\boldsymbol{y}_i'$ recursively in a way similar to the one used to produce $\tilde{\mathbf{X}}_2$ and $\boldsymbol{y}_2$; they will therefore be referred to as the *look-up tables update equations*. Finally, let us note that, once these look-up tables have been computed, only the first N_{T} rows of $\tilde{\mathbf{X}}$ and $\boldsymbol{y}$ are now needed for the algorithm to run.

If we suppose now that $N_{\mathrm{T}} \ll N$, then most of the computation time will be spent in the building of the initial values of the look-up tables (equations 23 and 24). More precisely, these will correspond crudely to the DFT of $\boldsymbol{w}^2$ and of $\boldsymbol{w}^2 \circ \boldsymbol{s}$; their vector multiplication with each combination of the templates plus the inverse transforms leads to these initial values. Despite the fact that the previous derivation is a bit coarse, it still assesses the linearithmic (i.e. $\mathcal{O}(N \log N)$) behaviour of the algorithm presented. Regarding the specific problem of QSO redshift determination within the *Gaia* mission (expected to be $N = 10^4$, $N_{\mathrm{T}} = 10$), tests performed on a 2.4-GHz CPU provide execution times of 180.35 ± 6.76 s for the normal equation solution, compared with 0.173 ± 0.002 s for our implementation; these become respectively 4.95 ± 0.19 h compared with 1.81 ± 0.02 s for the case of $N = 10^5$ and $N_{\mathrm{T}} = 10$. Finally, let us note that the proposed algorithm can easily be implemented in parallel, given the fact that the estimation of each $\chi^2(Z)$ can be performed separately. As a consequence, the execution time can be scaled by an arbitrary factor that is inversely proportional to the number of running processes.

5 APPLICATION

Unsurprisingly, the performance of the method presented was assessed on type I/II QSOs coming from the SDSS DR12 quasar catalogue (Pâris et al. in preparation). The choice of this catalogue comes from the fact that all spectra contained therein were inspected visually and hence can be considered as being extremely reliable regarding their redshift. Additionally, it is also interesting to note that the latter contains a non-negligible number of 297 301 QSOs, which is adequate to derive strong statistics.

Due to time constraints and to the need for the WPCA algorithm to have a well-covered input space of parameters (i.e. numerous observations), we used a twofold cross-validation in order to test our method. We split our input catalogue into two randomly drawn parts, out of which we extract the principal components; we then compute the redshift of spectra belonging to each part based on both the weighted and classical phase-correlation algorithms, the inputs of which are the principal components built on the alternative part. A detailed description of the processes leading to this cross-validation follows.

5.1 Procedure description

Raw spectra are generally not readily exploitable. Rather, we have to reduce them so as to get rid of most of the contaminating signals, encompassing, for the specific case of this study, deviant points (amongst which are night-sky emission lines and spectrograph edge effects) and QSO continuum. Note that since the SDSS DR12 spectra are already sampled on a uniform logarithmic scale, nothing has to be done in order for equation (1) to be fulfilled, but usually spectra have to be resampled.

The estimation of the QSO continuum turns out to be a challenging problem upon which the quality of the principal components and redshift prediction depend strongly (Machado 2013). Four broad kinds of approaches have been investigated so far in order to estimate this continuum: (1) the fit of a 'damped' power-law function to the observed spectra (Ferland 1996); (2) the use of PCA so as to predict the shape of the Lyα forest continuum based on the red part of the spectrum (Suzuki et al. 2005; Pâris et al. 2011; Lee, Suzuki & Spergel 2012); (3) the modelling of the dependence between the intrinsic QSO continuum and the absorption that it encounters as a mean to extrapolate the continuum (Bernardi et al. 2003); (4) through the use of techniques related to the multiresolution analysis (Dall'Aglio, Wisotzki & Worseck 2008; Machado 2013). We choose to use this last alternative, based on the fact that we do not require the resulting continuum to have a physical basis (i.e. the continuum subtraction is instead used as a normalization) and on the fact that we would like to have the most empirical estimation of this continuum. Following Machado (2013), we found that the signature of the continuum clearly stands within the low-frequency components of the pyramidal median transform (hereafter PMT: Starck 1996) of the input spectrum. In practice, the PMT is computed on a flipped version of the spectrum concatenated with the original version and another flipped version, so as to ensure continuity at the border. After taking the inverse transform through a third-degree fitting polynomial, we enforce the smoothness of the solution by convolving it with a thousand-points-wide Savitzky–Golay filter, so as to provide the final continuum. Besides its accuracy, we have to note that the PMT, from its pyramidal nature, has an algorithmic complexity of $\mathcal{O}(N \log N)$ and consequently will not degrade the performance of the global process.

After having subtracted the derived continuum, we discard border regions for which either $\lambda < 3800$ Å or $\lambda > 9250$ Å; we reject 4-Å regions around each significant night-sky emission line and finally we perform a k–σ clipping ($k = 3$, $\sigma = 4$) on the two first scales of the PMT, so as to remove extremely deviant points. Finally, we obtain an estimate of the SNR of each continuum-subtracted spectrum through the computation of a 'noiseless' spectrum, using the hypothesis that the noise within these spectra is contained entirely within the five first scales of the biorthogonal spline stationary wavelet transform of each spectrum (Cohen & Daubechies 1992; Burrus 1997). In practice, a spline of third degree was used for both analysis and synthesis. Fig. 3 illustrates the result of the initial reduction process.

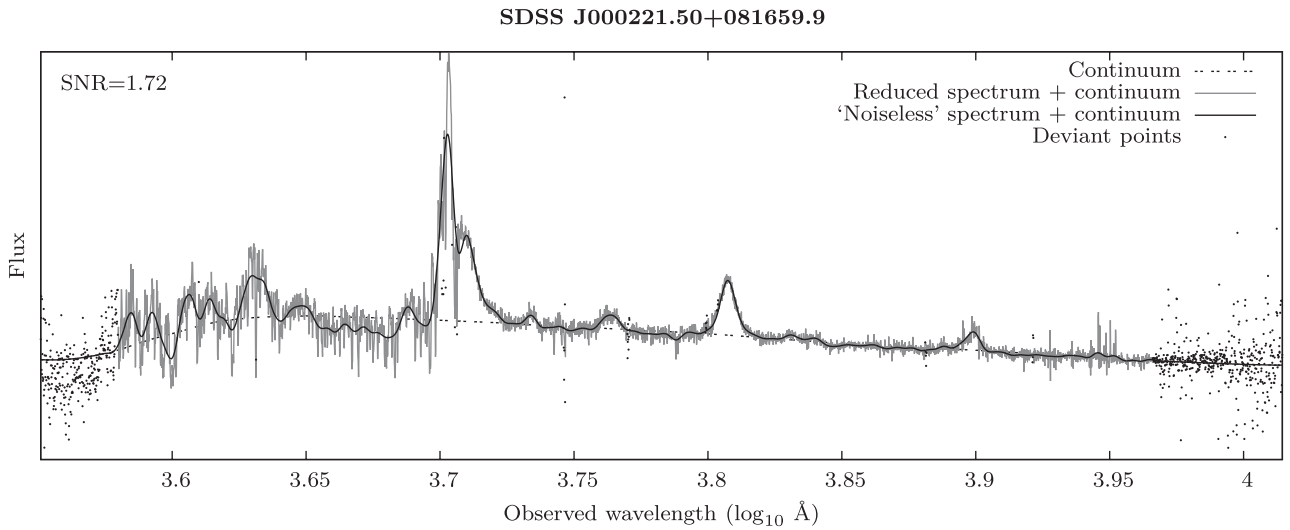

Figure 3. Result of the initial reduction process, designed to decompose any input spectrum into the following: a continuum spectrum through the smoothing of the low-frequency components of the PMT, a 'noiseless' spectrum through the removal of high-frequency components of the biorthogonal spline stationary wavelet transform and a set of deviant points through k–σ clipping of the first two scales of the PMT as well as through the removal of night-sky emission lines and border regions. Also computed is the SNR of the continuum-subtracted spectrum. The continuum drop occurring at wavelengths shorter than the Lyα limit can be explained by the absorption induced by the intergalactic medium – mainly composed of hydrogen atoms – located along the line of sight towards the QSO under study (Petitjean et al. 1993).

Spectra having an estimated SNR greater than 1 are then set on a common logarithmic wavelength scale with a uniform sampling of $\Delta \log_{10}\lambda = 10^{-4}$, equal to the original sampling of the spectra. The 116 374 resulting spectra are then divided into two equal parts – called learning sets – each of which is being used to produce the principal components and mean observations associated with each part of the cross-validation process. Resulting from this subdivision, the input catalogue will be split into two parts – the test sets – each consisting of 133 860 observations. Note that, given the fact that the broad absorption-line QSOs are discarded, the combined sets do not sum up to 293 301 QSOs.

We then compute the classical and weighted CCF of each spectrum contained within the two test sets, based on the mean observations and first ten principal components coming from the alternative learning set. Out of these CCFs, we extract the five most significant peaks – having a separation of at least 15 000 km s^{-1} – and fit them with a second-order polynomial, so as to gain a subsampling precision for the predicted peak position. Note that we choose to consider multiple solutions, based on the fact that the most significant peaks may not always lead to a physical basis. For example, we might have deep absorption lines coming from either the host galaxy of the quasars or extragalactic objects located along the line of sight during acquisition and leading to 'negative' fitted emission lines. These can definitely prevent the highest peak – the one with the associated minimal χ^2 – from being the effective one. In order to discriminate between these five selected solutions, we define two score measures: $\chi^2_r(z)$, defined as the ratio of the value of the peak associated with the redshift z to the value of the maximal peak and $Z_{\rm score}(z)$, defined as

$$Z_{\rm score}(z) = \prod \frac{1}{2}\left[1 + {\rm erf}\left(\frac{e_\lambda}{\sigma(e_\lambda)\sqrt{2}}\right)\right], \quad (27)$$

where e_λ are the mean values of the emission lines covered by the observed spectrum if we consider it to be at redshift z and $\sigma(e_\lambda)$ are the associated uncertainties. Note that both e_λ and $\sigma(e_\lambda)$ are computed over a range of 11 points surrounding each emission line. We can recognize each term of equation (27) as being the cumulative distribution function of a normal distribution of mean zero and variance $\sigma^2(e_\lambda)$ evaluated at e_λ. The use of equation (27) allows us to have a numerical estimate of the ability for a given redshift, z, to grab the following chosen QSO emission lines: O VI λ1033; Lyα λ1215; N V λ1240; Si IV λ1396; C IV λ1549; C III] λ1908; Mg II λ2797; Hγ λ4340; Hβ λ4861 and Hα λ6562 Å. Typical values of $Z_{\rm score}(z)$ range from $\sim$1 for a solution with a clear match of all positive emission lines; solutions with a match of at least one 'negative' emission line are voluntarily penalized by giving them a $Z_{\rm score}(z) \sim 0$, while values in between often occur in low SNR spectra or spectra with strongly damped emission lines. Finally, an error in each estimated peak position is derived and this will be described further in Section 6.1.

For each spectrum, the selection of the optimal redshift out of the five potential ones, $z_1, \cdots, z_5$ for which $1 = \chi^2_r(z_1) \geq \chi^2_r(z_2) \geq \cdots \geq \chi^2_r(z_5)$, coming from either the classical CCF or the weighted CCF, is done in the following way: if $Z_{\rm score}(z_1) > 0.8$, then select z_1; otherwise choose the shift having the highest χ^2_r and for which both $Z_{\rm score}(z_i) > 1 - 10^{-6}$ and $\chi^2_r(z_i) > 0.8$; otherwise choose the shift having the highest $Z_{\rm score}$ and for which $\chi^2_r(z_i) > 0.9$. Note that the previous selection and constants therein are purely empirical and based on an iterative visual inspection of the misclassified spectrum. This final step provides us with what we thought to be the most probable redshift estimate for a given input spectrum, along with the associated uncertainty and a warning flag notifying us of a failure and/or imprecision in the CCF computation, peak identification or redshift selection (e.g. all fluxes to zero, low $Z_{\rm score}$ or less precise uncertainties).

5.2 Results

Fig. 4 illustrates the result of the cross-validation process for both the classical phase-correlation and weighted phase-correlation algorithms and illustrates a comparison with the redshift predicted by the SDSS-III pipeline. We can readily see that the performance

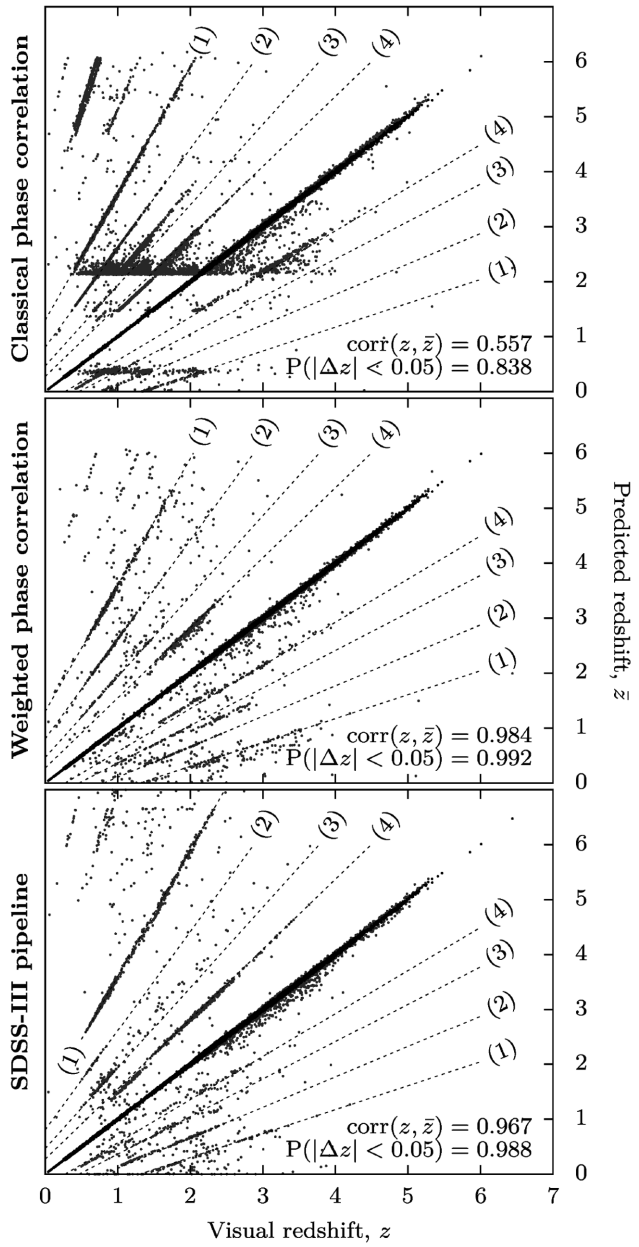

Figure 4. Results of the classical phase-correlation and weighted phase-correlation algorithms based on the cross-validation of observations coming from the SDSS DR12Q quasar catalogue plus predictions coming from the SDSS-III pipeline. The information provided comprises the correlation factor between z and $\bar{z}$ and $P(|\Delta z| < 0.05)$, the ratio of observations having an absolute error lower than 0.05 (depicted as black dots). Dotted numbered lines correspond to known mismatches between common emission lines: (1) Mg II with Lyα; (2) Mg II with C IV; (3) Mg II with C III] and (4) C IV with Lyα (see Section 5.2).

of the classical phase-correlation algorithm is strongly degraded compared with the weighted version, with a correlation factor of 0.557 compared with 0.984 and a ratio of observations having $|\Delta z| < 0.05$ of 0.838 compared with 0.992, respectively. These differences come mainly from the previously mentioned problem of border flattening, which translates into frequent emission-line mismatches and errors coming from the difficulty of the algorithm in extrapolating the regions surrounding the Lyα and Hα emission lines. This difficulty arises because of the prominence of these lines, as well as the high correlation they have with the other emission lines (Yip et al. 2004). As a consequence, the algorithm is often constrained to consider the Lyα or Hα lines to be embedded within the observed spectra, which results graphically in a gap around $0.4 < \bar{z} < 2.12$. Note that the systematic errors occurring at $\bar{z} \sim 0.4$ and at $\bar{z} \sim 2.12$ can be attributed to the fitting of these specific emission lines to the residual spectrograph edge effects – particularly significant within the low SNR spectra – and that these errors account for ∼2 per cent of the observations with $|\Delta z| \geq 0.05$.

Investigation of the most significant errors coming from the emission-line mismatch, illustrated in Fig. 4, shows that the latter can be modelled as a linear relation between the predicted redshift and the effective redshift. Indeed, if we consider an emission line observed at wavelength λ and falsely considered to stand at a rest-frame wavelength λ_f instead of λ_t, we find that the predicted redshift, z_f, can be related to the effective redshift, z_t, through

$$\frac{z_f + 1}{z_t + 1} = \frac{\lambda_t}{\lambda_f}. \tag{28}$$

These mismatches do not constitute, in themselves, real cases of degeneracy regarding our χ^2_r and $Z_{\rm score}$ selection criteria. Indeed, each of the configurations mentioned within Fig. 4 has unconfused emission lines that make the resulting redshift unambiguous. Rather, the observed degeneracies come also from the low SNR of the observed spectra. Fig. 5 illustrates the distribution of the SNR of both observations having $|\Delta z| \geq 0.05$ and those having $|\Delta z| < 0.05$ for our three cases of study. We notice that, for all three cases, the SNR of the maximal peak of the fair redshift estimate is approximately twice that of the erroneous ones; this is especially significant in the cases of the weighted phase correlation and the SDSS-III pipeline, where the errors come nearly exclusively from this line mismatch problem. Furthermore, a visual inspection of these degenerated spectra shows both potential redshifts to be indistinguishable from one another in most cases, thus constituting *in fine* effective cases of degeneracy. Consequently, some of the low SNR spectra will unavoidably have ambiguous redshift estimates that will stand in very specific regions defined by equation (28). Nevertheless, these will be easily identified as having a low $Z_{\rm score}$ and/or a low redshift confidence (see Section 6.1).

Finally, we notice that our implementation seems to have a better tolerance to noise compared with the SDSS-III implementation (within Fig. 5, see the lower peak of the erroneous curve as well as its globally smaller width). This higher tolerance does not come from differences in the algorithms, since both implementations are based on the sole solution to equation (11), but rather from either (1) the higher number of PCA components we used (11 compared with 4), (2) the fact that the components we used were more suitable to represent the observed spectra or (3) the fact that the redshifts coming from the visual inspection procedure are also subject to errors, especially since we are concerned with low SNR spectra, where degeneracy may occur.

In order to reject the fact that this higher tolerance comes from the larger number of components we used, we repeat the cross-validation procedure described by using only three components (plus mean observation) instead of ten. The results of this configuration lead us to the same conclusions, with a correlation factor of 0.976 (compared with 0.967) and a ratio of observations having $|\Delta z| < 0.05$ of 0.989 (compared with 0.988). Although the differences in the erroneous SNR curves are less perceptible, the curve coming from our implementation still remains globally sharper. Furthermore, we should mention that, within the SDSS-III pipeline, no more than four principal components were used, because any larger

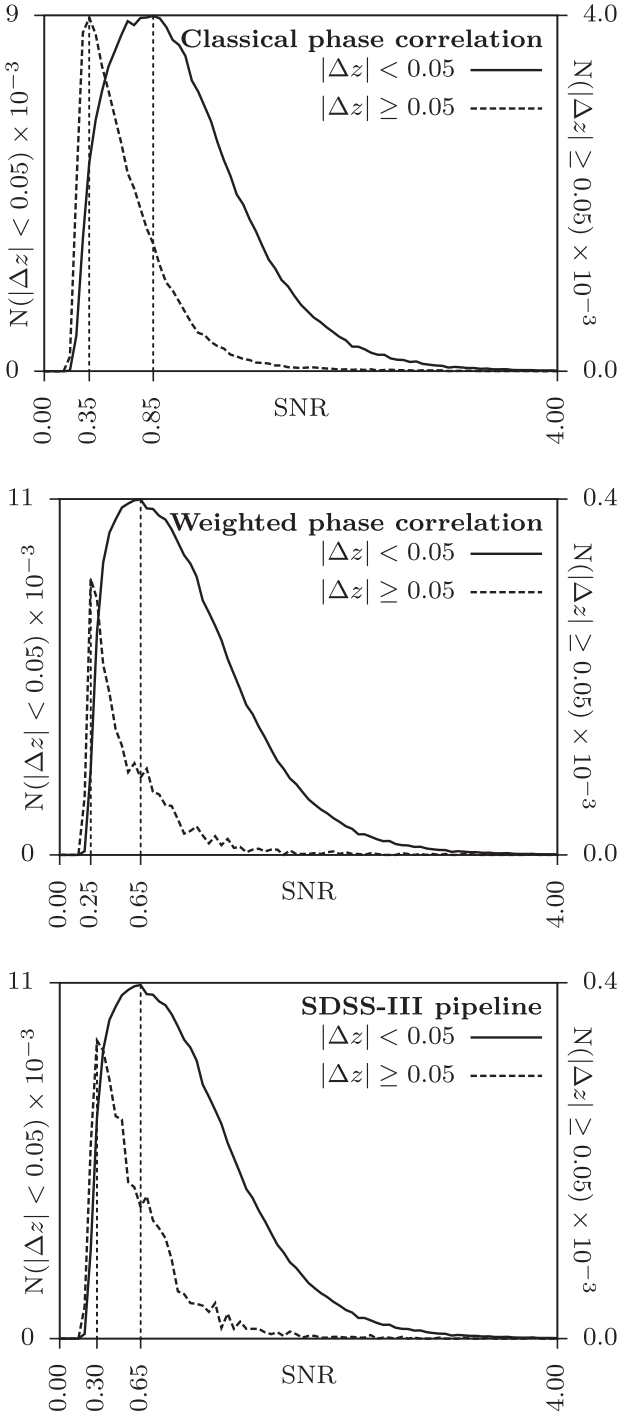

Figure 5. Histograms of the SNR distribution for observations having $|\Delta z| < 0.05$ and those having $|\Delta z| \geq 0.05$, for cases of classical phase correlation, and also of the weighted phase correlation and SDSS pipeline output. A bin width of $\Delta\mathrm{SNR} = 0.1$ was used in each of these graphs. The position of the maximal peak of each histogram is also highlighted.

number of components would make the error higher. In regard to this point and the fact that we succeeded in getting good predictions using 11 components, we might suppose, in any case, that the components we used were of higher quality in order to model this specific data set. Nevertheless, let us mention that we cannot reject totally the hypothesis according to which this better tolerance comes from a fortuitous statistical fluctuation itself produced by the degeneracy occurring during the visual inspection of some low SNR spectra.

6 DISCUSSION

6.1 Redshift confidence and uncertainty estimation

In order for the derived redshift to be used effectively within subsequent scientific applications, it is mandatory for it to have an estimation of its uncertainty and a confidence level such that the chosen redshift is indeed in the vicinity of the real redshift. To make this clear, we may have a redshift estimation with a reasonable uncertainty (e.g. $z = 2.31 \pm 10^{-3}$) but degenerated in such a way that we are not sure whether it stands in the neighbourhood of the effective redshift. Fortunately, the computed CCF offers us simple and efficient ways to evaluate both the redshift uncertainty and the confidence we can set on it.

Generally speaking, we know that, for a sufficiently large sample of observed points, the χ^2 map defined in the parameters $\{a_1, \cdots, a_n\}$ can be approximated in the neighbourhood of the global minimum, $\{a_1^\star, \cdots, a_n^\star\}$, as

$$\chi^2(a_i) \approx \frac{\left(a_i - a_i^\star\right)^2}{\sigma^2(a_i^\star)} + C, \tag{29}$$

where C is a function depending on $a_j, j \neq i$ and thus is considered here as a constant. In other words, the approximation of the χ^2 map near a global minimum can be evaluated for each of the parameters independently from the others as a simple quadratic curve, the curvature of which depends on the uncertainty of the varying parameter. As a consequence, if $\chi^2(a_i)$ increases by one compared with the optimal χ^2, then we have $\sigma^2(a_i) = \left(a_i - a_i^\star\right)^2$. The reader will find more information about the variation of χ^2 near the optimum and more particularly about the derivation of equation (29) in Bevington & Robinson (2003, section 8.1). Regarding the uncertainty in the predicted redshifts, we used a second-order polynomial such as to fit the optimal peak of the CCF, Z, and derived its associated uncertainty,[4] $\sigma(Z)$. We then use the propagation of the uncertainty, so as to obtain the error in the estimated redshift:

$$\sigma(z) = (z+1)\sigma(Z)s \log b, \tag{30}$$

where b is the base of the logarithmic scale we used (in our case $b = 10$) and s is the sampling of the spectra on this logarithmic scale (in our case $s = \Delta \log_{10}\lambda = 10^{-4}$).

Secondly, we have to evaluate the confidence we have in the predicted redshift. First estimators of this confidence are the already mentioned $Z_{\mathrm{score}}(z)$ and $\chi_r^2(z)$ (see Section 5.1). Indeed, a secure estimate will typically have $Z_{\mathrm{score}}(z) \approx \chi_r^2(z) \approx 1$. Unfortunately, these estimators do not take into account the potential ambiguity that might be present during the selection of the CCF peak associated with the predicted redshift. In order to tackle this lack, we defined the chi-squared difference associated with a redshift estimate, z_i, as

$$\Delta\chi_r^2(z_i) = \min\left|\chi_r^2(z_i) - \chi_r^2(z_j)\right| \qquad \forall j \neq i, \tag{31}$$

where each z_j corresponds to a redshift associated with a peak selected within the CCF. We notice that any redshift that is unsure due to the ambiguity in the CCF peak selection will now be marked

[4] Beware that the shift value corresponding to the uncertainty will have an associated decrease by one, compared with the maximal peak of the CCF.

as having $\Delta\chi_r^2(z) \approx 0$. Also note that, compared with Bolton et al. (2012), we use the distance between all χ_r^2 and not only those for which $\chi_r^2(z_i) \geq \chi_r^2(z_j)$, because we adopt the hypothesis that any solution having $\chi_r^2(z_i) < \chi_r^2(z_j)$ might have been falsely rejected while being a valid solution.

6.2 Dealing with zero weights

It commonly happens that the weight matrix, $\mathbf{W}$ from equation (11), has a lot of successive weights set to zero; this is especially true if we consider that the observation can be padded so as either to match the size of the templates or to deal with the periodic nature of the phase correlation. Additionally, nothing prevents us from shifting the observation in a circular way, so as to have this set of successive zeros in the first rows of the weights matrix, and to get rid of this artificial shift later by sliding back the CCF. This is particularly interesting if we have a number of zeroed weights equal to – or greater than – the number of components we used, N_T. In this case, the N_T first rows of the matrix $\tilde{\mathbf{X}}$ and of the vector $\boldsymbol{y}$ used within the factorized QR algorithm with look-up tables will all be equal to zero and, as a consequence, none of the $\tilde{\mathbf{X}}_\mathbf{i}'$, as well as none of the $\boldsymbol{y}_i'$, has to be computed. Differently stated, in addition to the building of the look-up tables and their updates, we solely have to compute $\mathbf{R}_{ij}$ and $b_i(Z)$ through equations (16) and (17):

$$\mathbf{R}_{ij} = -\frac{\tilde{\mathbf{L}}_{ij}}{\sqrt{\tilde{\mathbf{L}}_{ii}}} \tag{32}$$

and

$$b_i(Z) = -\frac{\tilde{\boldsymbol{l}}_i}{\sqrt{\tilde{\mathbf{L}}_{ii}}}. \tag{33}$$

This allows us to simplify our algorithm greatly and leads to execution times of 0.082 ± 0.001s for the case $N = 10^4$, $N_\mathrm{T} = 10$ and 0.912 ± 0.026s for the case $N = 10^5$, $N_\mathrm{T} = 10$. A rough comparison shows these execution times to be twice as fast as those presented at the end of Section 4.1.1.

6.3 Template weighting

Although the weighting of the observed spectra is most important regarding the redshift determination of QSOs, one might also want to have a template weighting such as, for example, to highlight some patterns or to reflect the fact that these templates often come with their own uncertainties. To this aim, we plug into equation (11) the diagonal matrix of weights associated with the template observations, $\mathbf{W_T}$, i.e.

$$\chi^2(Z) = \left\| \tilde{\mathbf{W}}_\mathbf{T}\mathbf{W}\boldsymbol{s} - \mathbf{W}\tilde{\mathbf{W}}_\mathbf{T}\tilde{\mathbf{T}}\boldsymbol{a}(Z) \right\|^2 = \left\| \tilde{\boldsymbol{y}} - \tilde{\mathbf{X}}\boldsymbol{a}(Z) \right\|^2. \tag{34}$$

After orthogonalization of the matrix $\tilde{\mathbf{X}} = \tilde{\mathbf{Q}}\tilde{\mathbf{R}}$, we obtain

$$\chi^2(Z) = \|\tilde{\boldsymbol{y}}\|^2 - \|\boldsymbol{b}(Z)\|^2, \tag{35}$$

with the first N_T elements of $\boldsymbol{b}(Z)$ being equal to the first N_T elements of $\tilde{\mathbf{Q}}^\mathrm{T}\tilde{\boldsymbol{y}}$. We can already note that, since $\tilde{\boldsymbol{y}}$ is now shift-dependent, knowledge of $\boldsymbol{b}(Z)$ alone is no longer sufficient to find the optimal shift, so that $\chi^2(Z)$ must be evaluated explicitly through equation (35).

Computation of the first N_T elements of $\boldsymbol{b}(Z)$ is done straightforwardly using the procedure described in Section 4.1.1, with $\tilde{\boldsymbol{y}}$ replacing $\boldsymbol{y}$ and both look-up tables given by

$$\tilde{\mathbf{L}}_{ij} = \mathcal{F}^- \left[\mathcal{F}\left[{\mathbf{W_T}}^2 \left(\mathbf{T}_i^\mathrm{col} \circ \mathbf{T}_j^\mathrm{col} \right) \right]^* \circ \mathcal{F}[\mathbf{W}^2] \right]_Z \tag{36}$$

and

$$\tilde{\boldsymbol{l}}_i = \mathcal{F}^- \left[\mathcal{F}\left[{\mathbf{W_T}}^2 \mathbf{T}_i^\mathrm{col} \right]^* \circ \mathcal{F}[\mathbf{W}^2\boldsymbol{s}] \right]_Z. \tag{37}$$

Finally, we see each $\|\tilde{\boldsymbol{y}}\|^2$ will be given by

$$\|\tilde{\boldsymbol{y}}\|^2 = \mathcal{F}^- \left[\mathcal{F}\left[{\mathbf{W_T}}^2 \right] \circ \mathcal{F}[\mathbf{W}^2\boldsymbol{s}^2] \right]_Z. \tag{38}$$

7 CONCLUSIONS

We have presented a new method for computing the weighted phase correlation of an observed input signal against templates that are not necessarily orthogonal. This method is found to be the preferred alternative to classical phase correlation in the case of input observations having a limited coverage and/or very distinct weights. The implementation of this method is based on a weighted chi-squared problem solved through a highly modified version of the QR orthogonalization algorithm, designed to benefit from the performance of the fast Fourier transform, so as to compute the numerous inner products present within the original QR algorithm. This implementation provides us with a numerically stable algorithm, which has a linearithmic time complexity that makes it affordable for the tight spectral processing of QSOs within the *Gaia* mission.

We have presented a complete application of this method to the case of the redshift determination of type I/II QSOs coming from the SDSS DR12 quasar catalogue through a twofold cross-validation procedure. This application is based on templates coming from the weighted principal components analysis decomposition of independent spectra from the same catalogue. We described in detail the reduction of those input spectra, as well as the method we used in order to select the most probable redshift from amongst the set of possible ones. Results of this cross-validation show our method to be the one of predilection for QSO redshift determination and show that it is comparable with the SDSS-III pipeline output, while not being an $\mathcal{O}\left(N^2\right)$ process.

Finally, we showed how we can obtain both the uncertainty in the predicted redshift and the confidence we can set on it. We further discuss two extensions of our method, namely the time saving we can achieve if we have a sufficient number of successive zeroed weights and the weighting of the template observations.

A free implementation of the algorithm described has been released under the GNU Public License[5] and can be downloaded freely at https://github.com/ldelchambre/wcorrQRL.

ACKNOWLEDGEMENTS

The author acknowledges support from the ESA PRODEX Programme ‘Gaia-DPAC QSOs’ and from the Belgian Federal Science Policy Office.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation and the US Department of Energy Office of Science. The SDSS-III web site is http://www.sdss3.org/.

[5] http://www.gnu.org/licenses/gpl-3.0.txt

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofisica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington and Yale University.

REFERENCES

Alam S. et al., 2015, ApJS, 219, 12
Aubourg É. et al., 2015, Phys. Rev. D, 92, 123516
Bailer-Jones C. A. L. et al., 2013, A&A, 559, A74
Bailey S., 2012, PASP, 124, 1015
Bernardi M. et al., 2003, AJ, 125, 32
Bevington P. R., Robinson D. K., 2003, Data reduction and error analysis for the physical sciences, 3rd edn. McGraw-Hill, New York
Bishop C., 2006, Pattern Recognition and Machine Learning, 1st edn. Springer-Verlag, New York
Bolton A. S. et al., 2012, AJ, 144, 144
Brault J. W., White O. R., 1971, A&A, 13, 169
Burrus C. S., Gopinath R. A., Guo H., 1997, Introduction to Wavelets and Wavelet Transforms: A Primer, 1st edn. Prentice Hall, London
Cabanac R. A., de Lapparent V., Hickson P., 2002, A&A, 389, 1090
Cohen A., Daubechies I., Feauveau J. C., 1992, Communications on Pure and Applied Mathematics, 45, 485
Dall'Aglio A., Wisotzki L., Worseck G., 2008, A&A, 491, 465
de Bruijne J. H. J., 2012, Ap&SS, 341, 31
Delchambre L., 2015, MNRAS, 446, 3545
Ferland G. J., 1996, University of Kentucky Internal Report, 565
Francis P. J., Hewett P. C., Foltz C. B., Chaffee F. H., 1992, ApJ, 398, 476
Glazebrook K., Offer A. R., Deeley K., 1997, ApJ, 492, 98
Golub G. H., Van Loan C. F., 1996, Matrix Computations, 3rd edn. The Johns Hopkins Univ. Press, London
Heavens A. F., 1993, MNRAS, 263, 735
Jolliffe I. T., 2002, Principal Component Analysis, 2nd edn. Springer-Verlag, New York
Lee K.-G., Suzuki N., Spergel D. N., 2012, AJ, 143, 51
Machado D. P., Leonard A., Starck J.-L., Abdalla F. B., Jouvel S., 2013, A&A, 560, A83
Pâris I. et al., 2011, A&A, 530, A50
Pearson K., 1901, Phil. Mag., 2, 559
Perryman M. A. C. et al., 2001, A&A, 369, 339
Petitjean P., Webb J. K., Rauch M., Carswell R. F., Lanzetta K., 1993, MNRAS, 262, 499
Press W. H., Tuekolsky S. A., Vetterling W. T., Flannery B. P., 2002, Numerical recipes in C++: The Art of Scientific Computing, 2nd edn. Cambridge Univ. Press, New York
Schlens J., 2014, preprint (arXiv:1404.1100)
Simkin S. M., 1974, A&A, 31, 129
Starck J. L., Murtagh F., Pirenne B., Albrecht M., 1996, PASP, 108, 446
Suzuki N., Tytler D., Kirkman D., O'Meara J. M., Lubin D., 2005, ApJ, 618, 592
Tonry J., Davis M., 1979, ApJ, 84, 1511
Tsalmantza P., Hogg D. W., 2012, ApJ, 753, 122
Yip C. W. et al., 2004, AJ, 128, 2603

This paper has been typeset from a TEX/LATEX file prepared by the author.

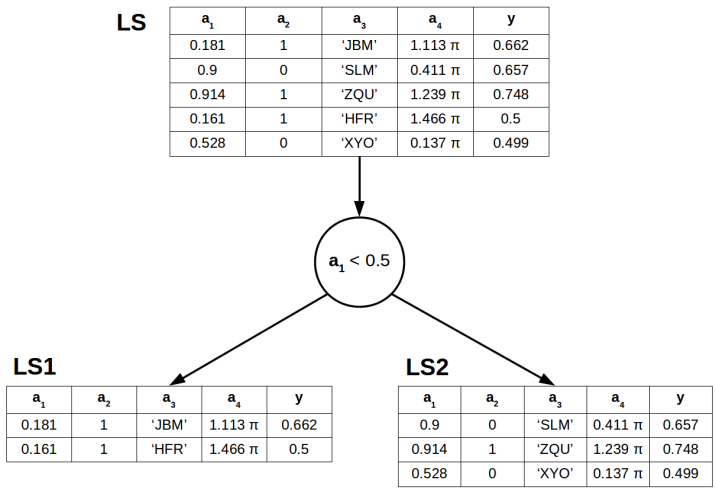

Figure 2.19: Example of a binary tree node working in regression mode. Given an input learning set of observations (LS) whose instances are composed of input attributes $\boldsymbol{a} = \{a_1, \ldots, a_n\}$ and of a numerical output variable y, the goal of the learning phase is to find an input attribute and an associated threshold among the values of this attribute such that LS is split at best according to a given score measure (e.g. the reduction of the variance amongst the output variable being a commonly used score measure in regression mode). The henceforth produced learning subsets (LS1 and LS2) being then recursively processed until an end condition is met or until no further split is possible.

2.4 Extremely randomized trees

We briefly outline here the supervised learning method we used so as to compare the results of the WPC method based on the predictions of the redshift and of the BAL nature of quasars out of BP/RP spectra of Gaia (see Paper III). This method reveals a particular interest as it also provides a fair identification of the gravitational lenses (GLs) that will be observed by Gaia (see section 4) and allows a good estimation of the parameters of the GL model if we assume these can be represented through a non-singular isothermal ellipsoid lens model in presence of an external shear (NSIEg lens model, see section 1.3.2). Supervised learning is an extremely extended field of research such that giving it a detailed overview would require more pages than we can afford in the present thesis. A reference book covering most of the machine learning (ML) aspects is to be found in Bishop (2006). We rather provide here a concise description of the supervised learning methods based on classification and regression trees (CART).

The goal of supervised learning methods can be summarized as follow:

> Supervised learning methods aim to guess the relations that may exist between a set of input parameters and a corresponding ensemble of output variables out of a learning set of observations so as to generalize these relations to cases where output variables are unknown.

This objective is attained in CART through the recursive splitting of the input learning set of observations (LS) according to the input attributes and associated thresholds

within these attributes that together allow to optimize a given score measure for each split, as depicted in figure 2.19.

More specifically, given a LS whose instances are composed of numerical input attributes $\boldsymbol{a}$ and of an output variable y, we have that the split performed according to the attribute a_i and associated threshold t_i is such that

$$\boldsymbol{y_1} = \{y(j) \mid a_i(j) < t_i\}, \qquad \boldsymbol{y_2} = \{y(j) \mid a_i(j) \geq t_i\}, \tag{2.18}$$

where j is used as an index within LS (e.g. $a_i(j)$ is the value of the ith attribute within the jth observation). The score that is associated with this split is commonly expressed in terms of the reduction of the impurity of the output variables, $\boldsymbol{y}$, as

$$S(\boldsymbol{y}, a_i, t_i) = I(\boldsymbol{y}) - \frac{|\boldsymbol{y_1}|}{|\boldsymbol{y}|} I(\boldsymbol{y_1}) - \frac{|\boldsymbol{y_2}|}{|\boldsymbol{y}|} I(\boldsymbol{y_2}) \tag{2.19}$$

where $I(\boldsymbol{y})$ is a measure of the impurity of $\boldsymbol{y}$ and $|\boldsymbol{y}|$ corresponds to the number of observations contained in it. The signification of $I(\boldsymbol{y_1})$, $I(\boldsymbol{y_2})$, $|\boldsymbol{y_1}|$, $|\boldsymbol{y_2}|$ being then trivially extrapolated from the latter definitions. Depending on the nature of the output variable, we have that the measure of the impurity $I(\boldsymbol{y})$ is either associated with the sample variance of $\boldsymbol{y}$,

$$I(\boldsymbol{y}) = \frac{1}{|\boldsymbol{y}|} \sum_j (y(j) - \bar{y})^2, \qquad \bar{y} = \frac{1}{|\boldsymbol{y}|} \sum_j y(j), \tag{2.20}$$

in case of numerical output variables or with the circular variance (Berens 2009)

$$I(\boldsymbol{y}) = 1 - R, \qquad R = \frac{1}{|\boldsymbol{y}|} \left[\left(\sum_j \sin y_j \right)^2 + \left(\sum_j \cos y_j \right)^2 \right]^{\frac{1}{2}}, \tag{2.21}$$

in case of cyclic output variable[7] having a period of 2π. In case of discrete output variable, the measure of the impurity is frequently associated with the concept of entropy from information theory (MacKay 2003)

$$I(\boldsymbol{y}) = -\sum_i p_i \log p_i, \qquad p_i = \frac{|\{y(j) \mid y(j) = y_i\}|}{|\boldsymbol{y}|}, \tag{2.22}$$

where p_i represents the probability that a randomly picked up element of $\boldsymbol{y}$ contains the value y_i. For completeness, we have to note that the Gini index is also a frequently used alternative for measuring the impurity of a discrete output variable and is defined as $I(\boldsymbol{y}) = \sum_i p_i(1 - p_i)$.

As long as identical input attributes lead to similar output values, one can exactly match the input attributes to the output variables through the recursive splitting of LS. Now, real-world observations always come with some noise that is present on both the input attributes as well as on the output variables such that a perfect fit of LS is often not desirable. The overfitting of LS would effectively lead to the modelling of noise that is deleterious for subsequent predictions and conversely, the use of a tree having an insufficient number of nodes may fail to reproduce the structure underlying the true function. The trade-off existing between models being very specific to LS while providing poor predictions over unseen observations and simpler models giving less accurate predictions but having a higher generalization potential is termed the *bias/variance trade-off* and stands at the heart of ML techniques as illustrated in figure 2.20.

[7] Any cyclic variable being easily converted so as to have a period of 2π, only the latter case will be covered here

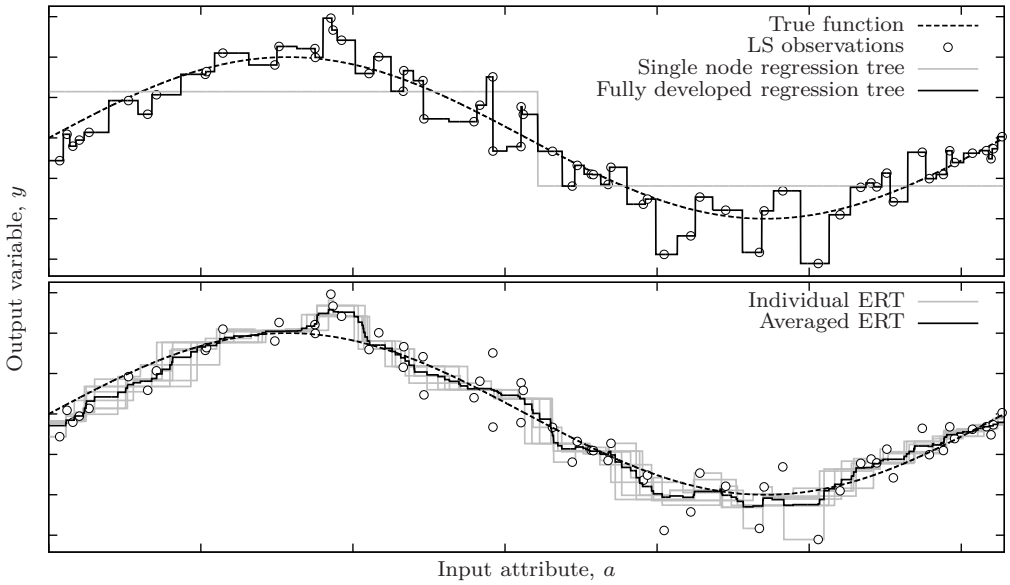

Figure 2.20: Modelling of a noisy sine function using: (up) a single node and a fully developed regression tree, (down) individual extremely randomized trees (ERT) and averaged predictions out of these individual ERT. One can note that the single node regression tree is not complex enough so as to allow the overall shape of the sine function the be correctly reproduced (i.e. it has a high prediction bias). At the opposite, a fully developed regression tree considers the noise as being part of the underlying function and will accordingly replicate it in all subsequent predictions (i.e. it has a high variance). Individual ERT, similarly to fully developed trees, typically have a low bias coming along with a high variance. This high variance is however lessened once the individual predictions are averaged.

Table 2.3: The confusion matrix is a special kind of contingency table that is particularly well-suited for the analysis of the results of classification problems. The mentioned condition corresponds to whether or not an observation belongs to a given class we choose to consider. Similarly, the predicted condition is positive only if the observation is predicted to belong to that class according to the ML model we used. The confusion matrix then compares the number of observations having positive/negative conditions and corresponding positive/negative predicted conditions. In this table, P and N are the total number of observations having a positive/negative condition, respectively, whereas TP is defined as the number of observations having both a positive condition and a positive predicted condition. Other quantities being straightforwardly derived in the same way.

Predicted condition	Condition	
	Positive (P)	**Negative** (N)
Positive	True positive (TP)	False positive (FP)
Negative	False negative (FN)	True Negative (TN)

Derived quantities: True positive rate (TPR = TP/P)
False positive rate (FPR = FP/N)
Accuracy (TP+TN)/(P+N)

Given that we would ideally like to have the best compromise between the bias and the variance or equivalently we would like to have the fairest predictions on observations having missing output variable(s), the performances of the ML methods are accordingly assessed on an independent test set of observations (TS) rather than on LS. The idea is to build the ML model based on LS, potentially using a validation set of observations (VS) so as to tune the model parameters, and to further retrieve output variables based on TS, these predictions being then compared to the effective output variables of TS in order to assess the model performances. In case of numerical variables, classical estimators are used so as to measure these performances as, for example, the correlation factor, the mean absolute error or the circular–circular correlation (Berens 2009) whereas in case of discrete output variables (i.e. in classification mode), one can use the *confusion matrix* we present in Table 2.3. Out of the latter table, we have that TPR corresponds to the probability of a positive observation to be predicted as positive (i.e. the probability it is detected), FPR corresponds to the probability that a negative observation is classified as positive (i.e. the probability of false alarm) while the accuracy corresponds to the ratio of correctly classified observations in this binary positive/negative classification scheme. To make it clear, consider a ML model having a TPR of 95% regarding GL identification coming along with a FPR of 1% and an accuracy of 85%. Assuming we used well-behaved LS and TS, this implies that a genuine GL being processed by this ML model has a probability of 95% to be detected as such whereas a fortuitous cluster of stars will be rejected at the 99% level, 85% of the observations being correctly classified in general. Finally, one has to note that most ML classifiers provide a probability of the processed observations to belong to each class rather than a straight classification of the latter. The belonging of the observations to these classes hence primarily relies on a threshold value we set on this probability. All values from the confusion matrix can hence be adjusted through the tuning of these probability thresholds. The analysis of the effect of this tuning on the TPR and on the FPR being easily performed through the receiver operating characteristic (ROC) curve that is described in Paper III.

Despite their ease of interpretation, CART are often not as accurate as alternative methods like artificial neural networks or support vector machines, for example (Bishop 2006). These low performances arising mainly because of their limited capability in generalizing the LS upon which these are built. Furthermore, owing to their top-down development strategy, any small change in LS can result in a very different resulting tree, itself leading to very distinct predictions. Solutions to these shortcomings are principally found in variance reduction techniques that aim at building ML models that are not too specific to the used LS (i.e. that have generalization capability) while keeping the associated bias as low as possible. One such technique standing, for example, in the bootstrap aggregating (bagging) method where many trees are constructed based on a random re-sampling of the initial LS (with replacement) and whose predictions are further combined so as to dim the final resulting variance. Another related solution is to disturb the learning phase of many trees built on the same LS before combining their individual predictions. This is the solution undertaken by the extremely randomized trees (Geurts et al. 2006, hereafter ERT) ML method. Instead of choosing at each node the optimal attribute/threshold allowing to maximize a given score measure, ERT rather select K random attributes as well as a random threshold amongst each of these attributes in order to isolate the one maximizing the chosen score measure. The building of N trees then allows to reduce the associated variance as in bagging. The individual tree development being either stopped once no more split is possible or in case where the number of observations in all leaf nodes falls under a given limit, $n_{\min}$. An illustration of the generalization

capability of the ERT is provided in figure 2.20 and compared to the prediction of simple regression trees.

Though the combination of the ERT predictions first aims to reduce the variance that is otherwise associated with each individual tree, it concurrently allows uncertainties on these predictions to be derived. Considering the predictions of N individual ERT, $y_1, \ldots, y_N$, for a given observation. We have that the final prediction of the ERT for this observation in case of a numerical output variable is given by the mean prediction, y, while the associated uncertainty, σ_y, is given by the standard deviation around y as

$$y = \frac{1}{N}\sum_{i=1}^{N} y_i, \qquad \sigma_y = \left[\frac{1}{N}\sum_{i=1}^{N}(y_i - y)^2\right]^{\frac{1}{2}}. \tag{2.23}$$

In case of a cyclic output variable, the final prediction turns out to be

$$y = \arctan2\left(\sum_{i=1}^{N}\sin y_i, \sum_{i=1}^{N}\cos y_i\right) \tag{2.24}$$

while the resulting uncertainty is associated with the 1σ confidence interval of the Von Mises distribution[8] as given by Upton (1986). Finally, the prediction and uncertainty of discrete output variables are associated with the frequency of appearance of each discrete class amongst the individual predictions, normalized such that these sum up to one.

[8]The Von Mises distribution being the circular counterpart of the normal distribution. Note that given the lengthy description of the associated confidence intervals, these are not reproduced here.

3

Determination of the astrophysical parameters of quasars within the Gaia mission

Although most of the issues we highlighted in section 2.1.2 are presently solved through the introduction of a weighted principal component analysis (WPCA) technique described in Paper I and of a fast weighted phase correlation (WPC) algorithm described in Paper II, the determination of the astrophysical parameters (APs) of quasars using these techniques however remains inapplicable to BP/RP spectra like those illustrated in figure 2.1. The main reasons for this inability are the facts that:

- BP and RP spectra have characteristic bell shapes coming from the instrumental response of the BP/RP spectrometers. This prevents the production of rest-frame WPCA templates out of these spectra. Similarly, the WPC algorithm is also not able to deal with this instrumental response as the latter depends on the observed wavelengths.

- BP/RP spectra are not sampled on a logarithmic wavelength scale such that the redshifts do not translate into simple offsets, as depicted in equation 2.5. This would otherwise had enable the WPC algorithm to be used.

- BP/RP spectra are distinct although these are overlapping. Considering the extraction of WPCA templates out of the sole BP or RP spectra and/or performing the determination of the redshift based on these two separate spectra would effectively lead to a tremendous loss of efficiency as, for example, characteristic emission lines often span over both the wavelength range covered by these spectra (e.g. at $z \gtrsim 2.35$, the Lyα emission line will always be found in BP whereas the C III] emission line will stand in RP). This is even more damageable given that the separation between these emission lines is one of the most important criterion for determining redshift, that is the presence of a single emission line (or a few emission lines) in a given spectrum inevitably leads to frequent ambiguities as it can be fitted by virtually all emission lines from the templates.

- Finally, a library of BP/RP spectra of quasars being still non-existent at present day, the production of WPCA templates that would be used in WPC can not hence be carried out. Also the final performances of the implemented approach can not be assessed without such a library.

This section addresses these aforementioned problems and details the exact procedures we employed so as to guess the APs of quasars. These encompass: the redshift of the quasar, their BAL nature, the slope of their continua and the total equivalent width of their emission lines (see section 2.1.1 for more information on these APs). The performances in predicting these APs is finally provided for a range of G magnitudes where we expect quasars. These ultimately correspond to the end of mission performances of the quasar classifier (QSO) work package.

Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **473,** 1785–1800 (2018) doi:10.1093/mnras/stx2417
Advance Access publication 2017 September 22

Determination of astrophysical parameters of quasars within the *Gaia* mission

L. Delchambre★
Institut d'Astrophysique et de Géophysique, Université de Liège, Allée du 6 Août 17, B-4000 Sart Tilman (Liège), Belgique

Accepted 2017 September 15. Received 2017 September 15; in original form 2017 July 18

ABSTRACT

We describe methods designed to determine the astrophysical parameters of quasars based on spectra coming from the red and blue spectrophotometers of the *Gaia* satellite. These methods principally rely on two already published algorithms that are the weighted principal component analysis and the weighted phase correlation. The presented approach benefits from a fast implementation, an intuitive interpretation as well as strong diagnostic tools on the potential errors that may arise during predictions. The production of a semi-empirical library of spectra as they will be observed by *Gaia* is also covered and subsequently used for validation purpose. We detail the pre-processing that is necessary in order for these spectra to be fully exploitable by our algorithms along with the procedures that are used to predict the redshifts of the quasars, their continuum slopes, the total equivalent width of their emission lines and whether these are broad absorption line (BAL) quasars or not. Performances of these procedures were assessed in comparison with the extremely randomized trees learning method and were proven to provide better results on the redshift predictions and on the ratio of correctly classified observations though the probability of detection of BAL quasars remains restricted by the low resolution of these spectra as well as by their limited signal-to-noise ratio. Finally, the triggering of some warning flags allows us to obtain an extremely pure subset of redshift predictions where approximately 99 per cent of the observations come along with absolute errors that are below 0.1.

Key words: methods: data analysis – quasars: general.

1 INTRODUCTION

Gaia is one of the cornerstone space missions of the Horizon 2000+ science program of the ESA that aims to bring a consensus on the history and evolution of our Galaxy through the survey of a billion celestial objects (Perryman et al. 2001). This objective is achieved by capturing a 'snapshot' of the present structure, dynamic and composition of the Milky Way by means of precise astrometric and photometric measurements of all the observed objects as well as by the determination of the distances, proper motions, radial velocities and chemical compositions of a subset of these objects (Gaia Collaboration et al. 2016).

The on-board instrumentation is principally composed of two 1.45×0.50 m telescopes pointing in directions separated by a basic angle of 106.5°, the light acquisition being then carried out by slowly rotating the satellite on its spin axis while reading each CCD column at the same rate as the objects cross the focal plane (i.e. in the so-called Time Delay Integration mode, hereafter TDI mode). The high astrometric precision of *Gaia* then comes from: (i) its lack of atmospheric perturbations, (ii) its large focal length of 35 m, (iii) the combination of the beams of light coming from both telescopes on to a single focal plane composed of a patchwork of 106 CCDs that allows to relate the positions of the objects coming from the two fields of view with an extremely precise angular resolution and (iv) its scanning law that enables to maximize the number of observed objects as well as the number of positional relations arising from the previous point (Lindegren et al. 2012).

In addition, a high-resolution ($R = \lambda/\Delta\lambda = 11\,700$) spectrometer, called Radial Velocity Spectrometer (RVS), centred around the Ca II triplet (845–872 nm) will allow us to determine the radial velocities of some of the most luminous stars ($G_{\rm RVS} < 16$ mag), while two low-resolution spectrophotometers, namely the Blue Photometer (BP) observing in the range 330–680 nm ($13 < R < 85$) and the Red Photometer (RP) observing in the range 640–1050 nm ($17 < R < 26$), will allow us to classify and characterize the objects having $G < 20$ mag. The interested reader is invited to read de Bruijne (2012) and Gaia Collaboration et al. (2016) for a more complete description of the *Gaia* spacecraft and of its payload.

★ E-mail: ldelchambre@ulg.ac.be

The previously described instrumentation, combined with the fact that *Gaia* is a full-sky survey where each object will be observed 70 times on average, is a unique opportunity in order to achieve some additional objectives. A non-exhaustive list of such applications consists of a finer calibration of the whole cosmological distance ladder (i.e. through parallaxes, Cepheids & RR Lyrae stars, type-Ia supernovae,...) (de Bruijne 2012), a better understanding of the stellar physics and evolution through the refinement of the Hertzsprung-Russell diagram (Jordi et al. 2008), the discovery of thousands of high-mass exoplanets (Perryman et al. 2014) as well as new probes regarding fundamental physics (Mignard 2009).

Amongst the most peculiar objects that *Gaia* will observe stand quasars also termed as quasi-stellar objects (QSOs) for historical reasons. Quasars are active galactic nuclei originating from the matter accretion that was occurring in the vicinity of supermassive black holes being at cosmological distances. Due to their high luminosity ($L > 10^{12} L_{\odot}$) and their large redshift ($0 < z < 7$), these objects play a key role in fixing the celestial reference frame used by *Gaia*, but they also have their own intrinsic interest in various cosmological applications like in the evolution scenarios of the galaxies (Hamann & Ferland 1999), as discriminants over the various universe model and their parametrization (López-Corredoira et al. 2016), as tracers of the large-scale distribution of Baryonic matter at high redshift (Yahata et al. 2005) or as a means to independently constrain the Hubble constant if the latter are gravitationally lensed (Schneider & Sluse 2013).

The identification and characterization of the 500 000 quasars that *Gaia* is expected to observe take place in the framework of the Data Processing and Analysis Consortium (DPAC), which is responsible for the treatment of the *Gaia* data in a broad sense, that is: from data calibration and simulation to catalogue publication through intermediate photometric/astrometric/spectroscopic processing, variability analysis and astrophysical parameters (APs) determination. The DPAC is an academic consortium composed of nine coordination units (CUs), each being in charge of a specific part of the data processing (O'Mullane et al. 2007). One of these, the CU8 'Astrophysical Parameters', is dedicated to the classification of the objects observed by *Gaia* and to the subsequent determination of their APs (Bailer-Jones et al. 2013).

This paper describes the algorithms that are to be implemented within the CU8 Quasar Classifier (QSOC) software module in order to determine the APs of the objects classified as QSOs by the CU8 Discrete Source Classifier (DSC) module while relying on their low-resolution BP/RP spectra. The collected APs aiming to be published within the upcoming *Gaia* data release 3 catalogue. The covered APs encompass the redshift, the QSO type (i.e. type I/II QSO or Broad Absorption Line QSO, hereafter BAL QSO), the slope of the QSO continuum and the total equivalent width of the emission lines.

In the following, Section 2 explains the conventions we used along this paper. Section 3 makes a brief review of the methods that were specifically developed in the field of this study. We present the production of a semi-empirical library of BP/RP spectra of QSOs used in order to train/test our models in Section 4. The AP determination procedures are covered within Section 5, while their performances are assessed in Section 6. Some discussion on the latter takes place in Section 7. Finally, we conclude in Section 8.

2 CONVENTIONS

This paper uses the following notations: vectors are in bold italic, $\boldsymbol{x}$; x_i is the element i of the vector $\boldsymbol{x}$. Matrices are in uppercase boldface or are explicitly stated; i.e. $\mathbf{X}$ from which element at row (variable) i, column (observation) j will be denoted by $\mathbf{X}_{ij}$.

Flux will here denote the spectral power received per unit area (a derived unit of W m^{-2}), while *flux density* will represent the received flux per wavelength unit (derived units of W m^{-3}). If not stated otherwise, input spectral energy distributions (SED) will be considered to be expressed in terms of flux density while BP/RP spectra will be expressed in terms of flux by convention.

3 METHODS

One of the main characteristics of the *Gaia* data processing is the large amount of information that has to be handled (e.g. about 40 GB of compressed scientific data are received from the satellite each day) and the consequent need for fast and reliable algorithms in order to reduce those data. These last requirements led the CU8 scientists to use techniques coming nearly exclusively from the field of the supervised learning methods whose underlying principle is to guess the APs of each observed object, which are unknown, based on the interpolation of the APs of some similar template objects (Bailer-Jones et al. 2013).

These methods have been proven to be fairly fast and reliable but often consist in black-box algorithms having no physical significance and having only basic diagnostic tools in order to identify the potential problems that may occur during the APs retrieval. This last point is particularly crucial in the case of medium- to low-quality observations, like BP/RP spectra, or in the case where the problem is itself prone to error, like the existing degeneracy in the redshift determination of low signal-to-noise ratio (SNR) QSOs (Delchambre 2016).

With these constraints and shortcomings in mind, we have developed two complementary algorithms that are specifically designed to gather the quasar APs within the *Gaia* mission based on the object BP/RP spectra while providing a clear diagnostic tool and ensuring an execution time that is limited to $\mathcal{O}(N \log N)$ floating point operations (i.e. conventionally considered as 'fast' algorithms).

3.1 Weighted principal component analysis

Principal component analysis (PCA) aims to extract a set of templates – the principal components – from a set of observations while retaining most of its variance (Pearson 1901). Mathematically, it is equivalent to find a decomposition of the covariance matrix associated with the input data set,

$$\boldsymbol{\sigma}^2 = \mathbf{P}\mathbf{D}\mathbf{P}^{\mathrm{T}}, \tag{1}$$

that is such that $\mathbf{P}$ is orthogonal and $\mathbf{D}$ diagonal and for which

$$\mathbf{D}_i \geq \mathbf{D}_j; \forall i < j. \tag{2}$$

The first columns of $\mathbf{P}$ are then the searched principal components. A decomposition such as the one of equation (1) is straightforwardly given by the singular value decomposition of the covariance matrix (Press et al. 2002).

Consider now the building of a set of rest-frame quasar templates based on a spectral library having a finite precision on the fluxes and a limited wavelength coverage. From the definition of the redshift, we will have that the observed wavelength, λ_{obs}, can be related to the rest-frame wavelength, λ_{rest}, through

$$\lambda_{\mathrm{obs}} = (z+1)\lambda_{\mathrm{rest}}, \tag{3}$$

and as a consequence, we will have that every quasar within the input library will cover a specific rest-frame wavelength range that

depends on its redshift. Furthermore, the measurement of the quasar fluxes often comes along with an estimation of their associated uncertainties originating, for example, from the Poisson nature of the photons counting; from the CCD readout noise; from the sky background subtraction or from spectra edge effects. These uncertainties are not taken into account within the classical PCA implementation.

In Delchambre (2015), we solved the previously mentioned issues by considering the use of a weighted covariance matrix inside equation (1). For this purpose, we defined the weighted covariance of two discrete variables, $\boldsymbol{x}$ and $\boldsymbol{y}$ having weights, respectively, given by $\boldsymbol{w}^x$ and $\boldsymbol{w}^y$ and weighted mean values given by $\bar{x}$ and $\bar{y}$ as

$$\sigma^2_{x,y} = \frac{\sum_i (x_i - \bar{x})\, w_i^x w_i^y\, (y_i - \bar{y})}{\sum_i w_i^x w_i^y}. \tag{4}$$

The suggested implementation relies on two spectral decomposition methods, namely the power iteration method and the Rayleigh quotient iteration, that allow us to gain flexibility, numerical stability as well as lower execution times[1] when compared to alternative weighted PCA methods (Bailey 2012; Tsalmantza & Hogg 2012).

3.2 Weighted phase correlation

The redshift has a particular importance over the whole quasar APs because any error committed on the latter would make the other APs diverge. It is then critical to have the most precise estimation of it along with a strong diagnostic tool in order to flag the insecure predictions. A technique fulfilling these requirements stands in the phase correlation algorithm (Glazebrook, Offer & Deeley 1998) whose goal is to find the phase at which an input signal and a set of templates match at best in a chi-square sense.

For reasons already enumerated within Section 3.1, we will consider here a weighted version of the previously mentioned algorithm. We are then searching for the shift at which an input spectrum, $\boldsymbol{s}$, associated with a weight vector, $\boldsymbol{w}$, and a set of templates, $\mathbf{T}$, matches at best in a weighted chi-square sense. Mathematically, it is equivalent to find the shift, Z, for which

$$\chi^2(Z) = \sum_i w_i^2 \left(s_i - \sum_j a_j(Z) \mathbf{T}_{(i+Z)j} \right)^2 \tag{5}$$

is minimal given that $\boldsymbol{a}(Z)$ are the linear coefficients minimizing equation (5) for a specific shift.

In Delchambre (2016), we showed that the latter equation can be re-written as

$$\chi^2(Z) = \sum_i w_i^2 s_i^2 - \mathrm{ccf}(Z), \tag{6}$$

where $\mathrm{ccf}(Z)$ is the so-called *cross-correlation function* (CCF) at shift Z, which can be evaluated for all Z in $\mathcal{O}(N \log N)$ floating point operations and N is the number of samples we used. Given that the first term of equation (6) is independent of the shift, we will simply have that the minimum of equation (5) will be associated with the maximum of the CCF.

Practically, both $\boldsymbol{s}$ and $\mathbf{T}$ must be sampled on a uniform logarithmic wavelength scale in order for the redshift to turn into a simple linear shift (i.e. $\log \lambda_{\mathrm{obs}} = \log(z+1) + \log \lambda_{\mathrm{rest}}$) and must be extended and zero-padded such as to deal with the periodic nature of the phase correlation algorithm. A sub-sampling precision on the shift can then be gained by fitting a quadratic curve in the vicinity of the optimum of the CCF while the curvature of this quadratic curve will be used as an approximation of the uncertainty associated with the found shift.

The described weighted phase correlation algorithm relies on the assumption that the most probable redshift is associated with the maximal peak of the CCF, which is not always verified in the case of QSOs. The reason for this is twofold:

(i) The highest peak of the CCF may not always lead to a physical solution like the omission of some characteristic emission lines (e.g. Ly $\alpha\,\lambda 121$, Mg II $\lambda 279$ or H $\alpha\,\lambda 656$ nm) or the fit of a 'negative' emission line coming from the presence of matter being in the line-of-sight towards the observed QSO. The origin of this issue mainly stands in the imperfections of the templates we used as well as in the assumption we made that quasar spectra can be modelled as a linear combination of these templates.

(ii) In the case of low-SNR spectra, it may also happen that the signal of some emission lines starts to be flooded within the noise such that these will not be recognized as a genuine signal but rather as a variance coming from noise. As a result, ambiguities can emerge within the CCF (i.e. multiple equivalent maxima) and hence within the redshift determination.

In order to identify these sources of errors, we defined two complementary score measures associated with each redshift candidate: (i) $\chi^2_r(z)$, defined as the ratio of the value of the CCF evaluated at z to the maximum of the CCF and (ii) $Z_{\mathrm{score}}(z)$ defined as

$$Z_{\mathrm{score}}(z) = \prod_\lambda \left[\frac{1}{2} \left(1 + \mathrm{erf} \frac{e_\lambda}{\sigma(e_\lambda)\sqrt{2}} \right) \right]^{I_\lambda}, \tag{7}$$

where e_λ is the mean value of the continuum-subtracted emission line standing at rest-frame wavelength λ if we consider the observed spectrum to be at redshift z, $\sigma(e_\lambda)$ is the associated uncertainty and I_λ is the theoretical intensity of the emission line standing at λ normalized such that all the covered emission lines intensities sum up to 1. Equation (7) can then be seen as the weighted geometric mean of a set of normal cumulative distribution functions of mean zero and standard deviations $\sigma(e_\lambda)$ evaluated in e_λ. Table 1 summarizes the various emission lines and theoretical intensities we used in the context of this study.

We can already notice that these two score measures can easily highlight the sources of errors that may occur within the CCF peak selection, namely the choice of an unphysical solution and the choice of an ambiguous candidate, respectively, through a low Z_{score} and through a low absolute difference between the chosen candidate's χ^2_r and the one from another candidate. These will constitute *in fine* strong diagnostic tools regarding our implementation.

Table 1. Rest-frame wavelengths and relative intensities of the emission lines used in the computation of $Z_{\mathrm{score}}(z)$. Listed values come from the PCA mean spectra described in Section 5.

Emission line(s)	Rest-frame wavelength nm	Relative intensity compared to Lyα
O VI	103.07	0.13161
Lyα	121.81	1.00000
C IV	154.63	0.30834
C III]	190.24	0.16413
Mg II	280.18	0.21778
H β+[O III]	488.06	0.42182
H α	656.86	1.17143

[1] Under the condition that the number of observations within the input data set is much larger than the number of variables.

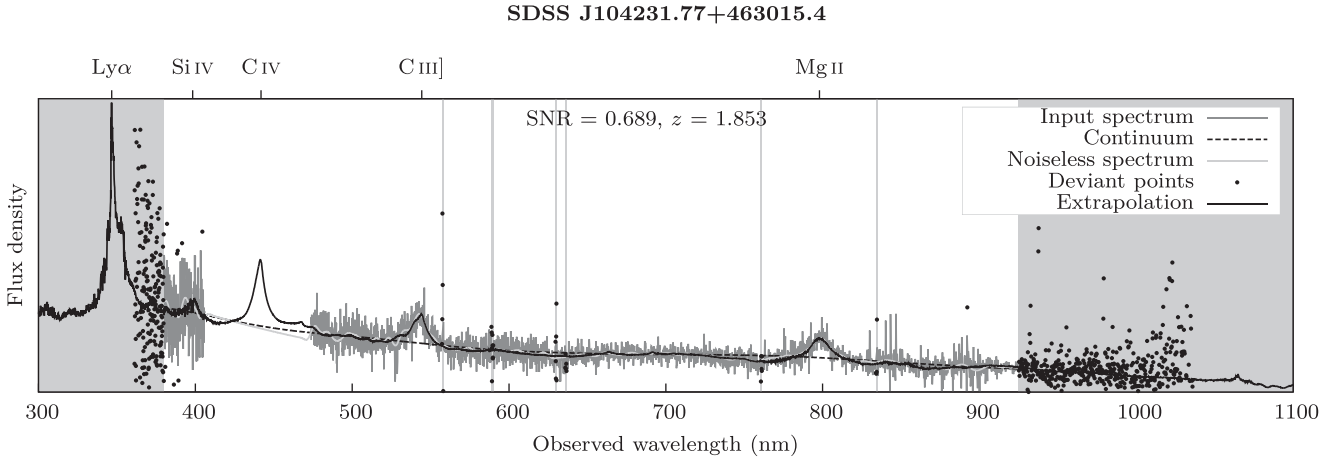

Figure 1. Results of the pre-processing and extrapolation procedures. The input spectrum is decomposed into a continuum part, a noiseless spectrum and a set of deviant points (night sky emission lines and spectrograph edge effects being here depicted as shaded regions). This decomposition allows us to compute the SNR associated with each spectrum and to set those having an SNR > 1 on a common logarithmic rest-frame wavelength scale so as to extract PCA templates out of them. The fit of these templates to each input spectrum provides the final extrapolation. Note that the illustrated input spectrum has some corrupted samples in the region encompassing the C IV emission line which are successfully recovered through this extrapolation procedure. Similarly, the unobserved Lyα emission line seems to be consistently reproduced as well.

4 SEMI-EMPIRICAL BP/RP SPECTRAL LIBRARY BUILDING

Similarly to supervised learning methods, the undertaken approach is based on the availability of a learning library of BP/RP spectra for which the various APs are known. Such a library being non-existent at the present time, we had to convert an already-released spectral library of QSOs into BP/RP spectra according to the most up-to-date instrument model. We focused, for this purpose, on the 12th data release of the Sloan Digital Sky Survey quasar catalogue (Pâris et al. 2017, hereafter DR12Q). The choice of this specific catalogue comes from (i) the fact that each spectrum in it was visually inspected yielding extremely secure APs, (ii) the large number of 297 301 QSOs it contains (amongst which 29 580 BAL QSOs), (iii) its medium resolution of $1300 < R < 2500$ and (iv) its spectral coverage which is comparable to the one of the *Gaia* BP/RP spectrophotometers ($360 < \lambda < 1000$ nm).

This spectral library will then have to be extended so as to match the wavelength range covered by the *Gaia* BP/RP spectra and be subsequently convolved with the instrumental response of the BP/RP spectrophotometers so as to provide the final library.

4.1 Spectra extrapolation

Besides the fact that the DR12Q spectra have a narrower wavelength coverage ($360 < \lambda < 1000$ nm) when compared to the BP/RP spectra ($330 < \lambda < 1050$ nm), we also have to mention that the regions where $\lambda < 380$ and where $\lambda > 925$ nm often tend to be unreliable because of spectrograph edge effects while some other inner regions might be discarded because of bad CCD columns, cosmic rays, significant scattered light or sky background subtraction problems, for example (Bolton et al. 2012; Dawson et al. 2013). In order to solve these shortcomings, we have extracted a set of rest-frame PCA templates out of the DR12Q library that were later fitted to each individual spectrum as a mean to extrapolate them. Note that since the DR12Q spectra are already sampled on a logarithmic wavelength scale (at a sampling rate of $\Delta \log_{10} \lambda = 10^{-4}$), no re-sampling will be needed before extrapolation.

Raw spectra are not readily exploitable, rather they have first to be pre-processed so as to get rid of contaminating signals and to have some insights about their usability. For this purpose, we used a procedure that is identical to the one described in Delchambre (2016, section 5.1). We will hence concentrate here on the results of this procedure rather than on the underlying implementation details. We will then get, for each spectrum: (i) the set of deviant points coming from a k–σ clipping algorithm applied to the high-frequency components of the spectrum as well as from the removal of the night sky emission lines and spectrograph edge effects, (ii) an empirical estimation of the QSO continuum coming from the low-frequency components of a multiresolution analysis of the spectrum, (iii) a smoothed version of the provided spectrum that we will consider here as being noiseless and (iv) an evaluation of its SNR coming from the ratio of the variance that is present within the noiseless and continuum-subtracted spectrum over the variance that can be attributed to noise (i.e. raw fluxes from which we subtracted the deviant points; the QSO continuum and the noiseless spectrum). Fig. 1 illustrates the results of the previously described procedure and provides a self-explanatory example of the necessity we have to pre-process our input spectra.

Spectra having an SNR larger than 1 are then normalized so as to have a weighted norm of 1 and are subsequently set on a common logarithmic rest-frame wavelength scale. These will constitute the input data set upon which we will extract our PCA templates. We choose to consider the retrieval of the BAL QSO templates separately from the type I/II QSOs as a way to ensure that the characteristic features of the BAL QSOs are correctly reproduced within our extrapolated spectra. Doing otherwise would have required a much larger number of PCA components to be fitted in order to accurately model those features (the latter being not seen in the vast majority of QSOs, they would have been omitted from the dominant PCA components). Also, we required the continuum to rely on an empirical basis so as to restrain at most any unphysical behaviour within our extrapolation. We will consequently subtract each empirical continuum from each input spectrum as a way to extract the PCA components from both these subtracted spectra as well as from the continua themselves. By way of

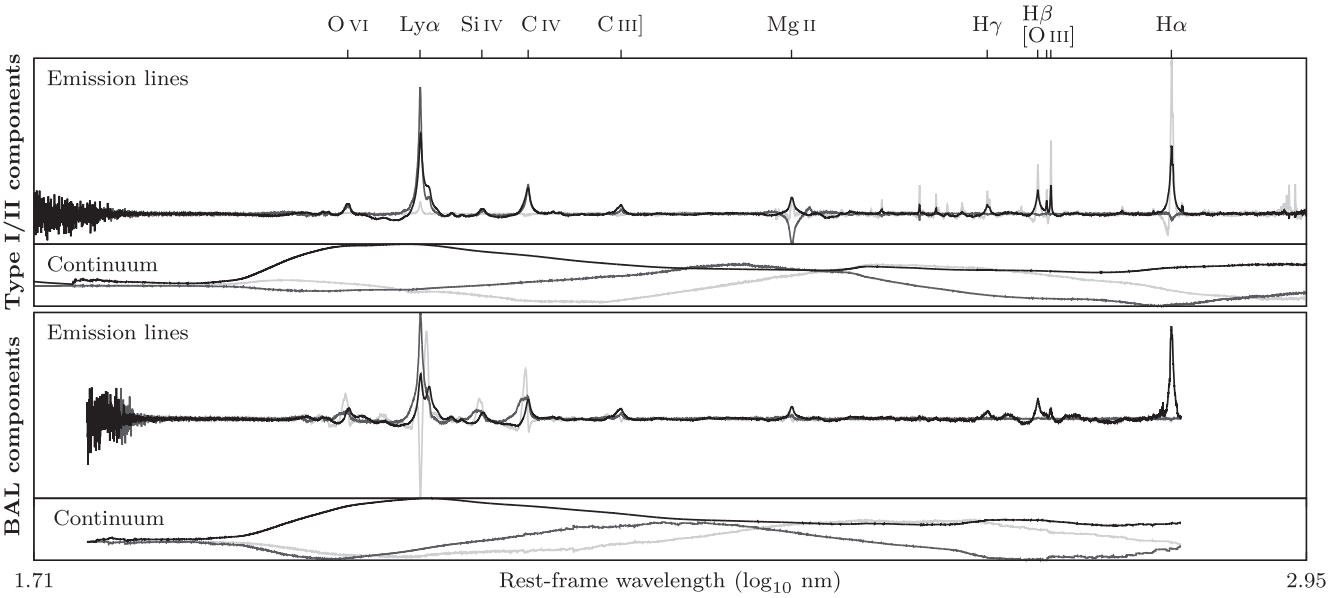

Figure 2. Mean observations (black lines) and first two principal components coming from the PCA decompositions of the DR12Q emission lines and continuum spectra regarding type I/II QSOs (up) and BAL QSOs (down). Notice how the entire emission lines are modelled by the type I/II QSOs templates while the BAL QSOs templates rather focus on the O VI; Ly α; Si IV and C IV emission lines where most of the BAL characteristics are found.

comparison, continuum templates are frequently taken as being a combination of power-law and exponential functions (Claeskens et al. 2006) that often succeed in reproducing the observed spectrum but that tend to diverge over the unobserved wavelengths. Consequently, four sets of PCA templates were built based upon the algorithm described in Section 3.1: one set of templates for the type I/II QSO emissions lines; one similar set associated with the BAL QSOs and two corresponding sets of continuum templates. Fig. 2 provides the mean observation and two first principal components for these four sets of templates. For the sake of completeness, let us also mention that, during PCA retrieval, weights were taken as the inverse standard deviation on the fluxes regarding the emission line PCAs while these were simply set to 1 if the continuum we computed was associated with some observed fluxes and zero otherwise.

Finally, we used 15 emission line templates (the mean observation and the first 14 PCA components) for both the fit of the type I/II and BAL observations along with 5 continuum templates for the type I/II QSOs and 6 continuum templates for the BAL QSOs. These fits ultimately provide the extrapolated spectra, as illustrated in Fig. 1. The number of templates we used allows us to explain 68.27 per cent of the weighted variance[2] that is present within the type I/II emission lines and 99.82 per cent of their continuum variance. These become, respectively, 66.1 per cent and 99.85 per cent in the case of BAL QSOs. Even if apparently low, these ratios practically reflect the fact that the spectra we used come along with noise that will not be grabbed by the dominant PCA components. Consequently, the produced extrapolation will be considered here as being noiseless.

[2] Weighted variance naturally results from equation (4) in the case where $\boldsymbol{x} = \boldsymbol{y}$ while having expected values, $\bar{x} = 0$ in the case of the weighted variance of the input data set or equal to the optimal linear combination of the templates in the case of the explained weighted variance.

4.2 BP/RP instrumental convolution

The optical system of the *Gaia* BP/RP spectrophotometers consists for each in six mirrors, a dispersing prism and a set of dedicated CCDs (either blue- or red-enhanced) that can be modelled as

$$S(x) = \int N_\lambda R_\lambda L_\lambda (x - x_\lambda) d\lambda, \tag{8}$$

where $S(x)$ is the dispersed flux in the one-dimensional comoving[3] focal plane position x, N_λ is the input SED at the observed wavelength λ, R_λ is the global instrumental response at λ and where $L_\lambda(x - x_\lambda)$ is the monochromatic line spread function (LSF) standing at λ and being evaluated at $x - x_\lambda$, x_λ is the comoving focal plane position associated with λ. In more details, the global instrumental response R_λ encompasses the mirrors reflectivity, the attenuation that is due to particulate and molecular contamination, the attenuation coming from the mirror roughness, the prism transmissivity and the CCD quantum efficiency at the observed wavelength λ. Note that in the following, we will consider a mean instrumental model averaged over each field-of-view and over each row of CCD within the focal plane given that our goal is to simulate end-of-mission (combined) spectra that will consist in an aggregation of individual (epoch) observations.

The BP/RP sampled fluxes, $\boldsymbol{s}$, will then be given by

$$s_i = \int_{-0.5}^{+0.5} S\left(x + \frac{i}{N_{\text{over}}}\right) \mathrm{d}x, \tag{9}$$

where N_{over} is the oversampling we choose to use. This oversampling arises from the higher SNR that is gained by the combination of the epoch spectra, which, in turn, provide the opportunity to reach a higher sampling rate when compared to the initial 60 pixels provided by the BP/RP acquisition window (e.g. through flux interpolation). A common consensus within the CU8 is to consider an oversampling

[3] Sliding at the same rate as the TDI mode drift each CCD column.

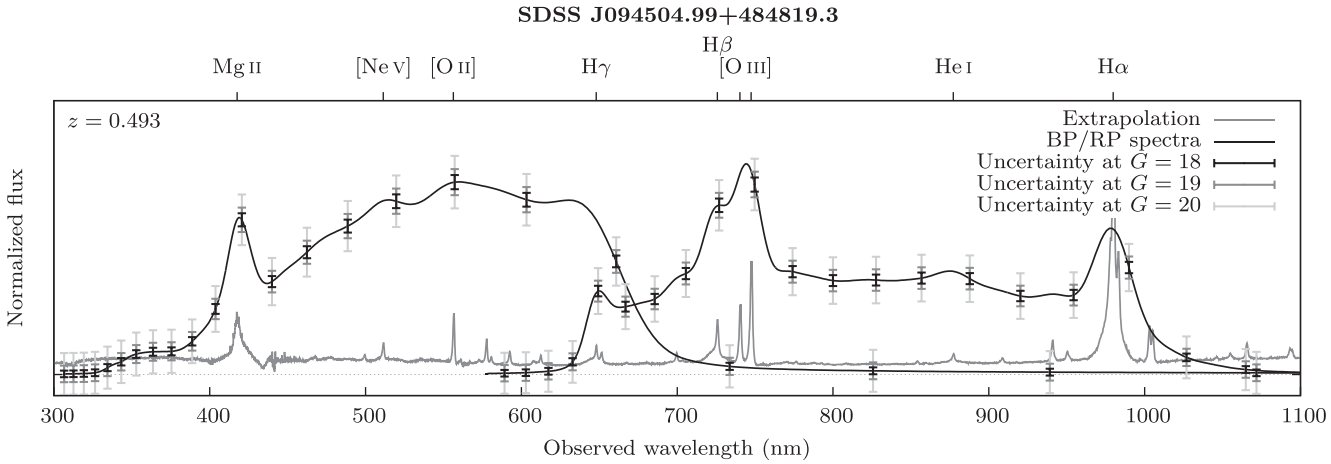

Figure 3. Example of produced BP/RP spectrum and associated flux uncertainties for G magnitudes of 18, 19 and 20. Spectra normalization allows us to fairly compare the noise levels that are present amongst the various provided magnitudes. We can straightly figure out the effect of the global instrumental response and of the LSF convolution applied to the extrapolated spectrum as respectively producing characteristic bell shapes as well as very broadened emission lines within the resulting BP/RP spectra. The noticed non-uniform spectral resolution arises from the varying wavelength dispersion function of each of the BP/RP spectrophotometers.

of $N_{\text{over}} = 8$, which results in 480 samples for each of the BP and RP spectra. This convention will be adopted here.

The instrument model described so far is not able to deal with the various sources of noise that will contaminate our actual observations. Instead, the spectra produced through equations (8) and (9) will consist in the approximated noise-free counterparts of what *Gaia* will observe. Extending the aperture photometry approach developed in Jordi et al. (2010), we can still have an estimation of the noise variance that is associated with each sampled flux, s_i, as

$$\sigma_i^2 = m^2 \frac{\sigma_{\text{epoch}}^2}{N_{\text{epoch}}/N_{\text{over}}} + \sigma_{\text{cal}}^2, \tag{10}$$

where m is an overall mission safety margin designed to take into account the potentially unknown sources of errors ($m = 1.2$, by convention); σ_{epoch}^2 is the variance of the noise associated with s_i if the latter was coming from a single epoch observation; N_{epoch} is the number of epoch observations used to compute the combined spectra and σ_{cal}^2 is the uncertainty arising from the flux internal calibration. The scaling of the single epoch variance, σ_{epoch}^2, reflects the assumption we made that each flux within the combined spectra comes from the mean value of a set of $N_{\text{epoch}}/N_{\text{over}}$ epoch fluxes. In extreme cases, for example, we will have that each combined flux is averaged over the whole epoch observations (i.e. in the case of $N_{\text{over}} = 1$) while in the case of $N_{\text{over}} = N_{\text{epoch}}$, each flux within the combined spectra can be seen as gathered directly from the epoch spectra. Next, the variance coming from the uncertainties in the flux internal calibration, σ_{cal}^2, is taken to be equal to the inner product of the fluxes that are present within the pixels surrounding each sample with a linear function that is inversely proportional to the global instrumental response, R_λ, evaluated in those pixels. Its objective is to take into account the fact that the precision on the flux calibration will principally depend on the instrumental response in the vicinity of the pixel of interest. Because of the intricacy that is inherent to the modelling of these calibration errors, the latter were voluntarily tuned so as to stand within a moderate range of values.

We can then decompose the epoch variance, σ_{epoch}^2, into variance coming from photon and CCD noise, σ_{flux}^2, and variance coming from the uncertainties in the background estimation, σ_{bg}^2, as

$$\sigma_{\text{epoch}}^2 = \frac{\sigma_{\text{flux}}^2 + \sigma_{\text{bg}}^2}{\tau^2}, \tag{11}$$

where τ is the effective CCD exposure time, and the latter is dependent on the G magnitude of the objects based on the activation of bypasses within some specific CCD columns which aim to prevent luminous objects from saturating (de Bruijne 2012). Both terms within equation (11) can then be simply extended as

$$\sigma_{\text{flux}}^2 = (s_i + b)\tau + r^2, \tag{12}$$

$$\sigma_{\text{bg}}^2 = \frac{b\tau + r^2}{N_{\text{bg}}}, \tag{13}$$

where r is the total CCD detection noise including, amongst other, the CCD readout noise and the CCD dark noise; b is the background flux we subtracted from our observation and that we will consider here as a constant based on a typical sky-background surface brightness and where N_{bg} is the number of pixels we used in order to estimate b, that is taken here as being equal to the width of the BP/RP acquisition window (i.e. $N_{\text{bg}} = 12$ pixels).

Being now able to model the entire instrumental response along with the associated uncertainties, we first choose to normalize our input spectra to G magnitudes where we expect QSOs to be observed, that is at $G = \{18, 18.5, 19, 19.5, 20\}$. This normalization allows us to study the behaviour of the implemented methods under an increasing level of noise. BP/RP spectra were then produced based on the most up-to-date instrument model coming from tests carried out by EADS Astrium (later renamed Airbus Defence and Space) during the commissioning phase of the satellite. We have to note that this instrument model still has to be updated in order to match the actual operational condition of the satellite even if the latter is not expected to vary too much from the model we used. Noisy spectra can then be obtained by adding the appropriate random Gaussian noise (i.e. having a variance of σ_i^2) to each of the (noise-free) spectral fluxes, s_i. An example of produced BP/RP spectra is illustrated in Fig. 3.

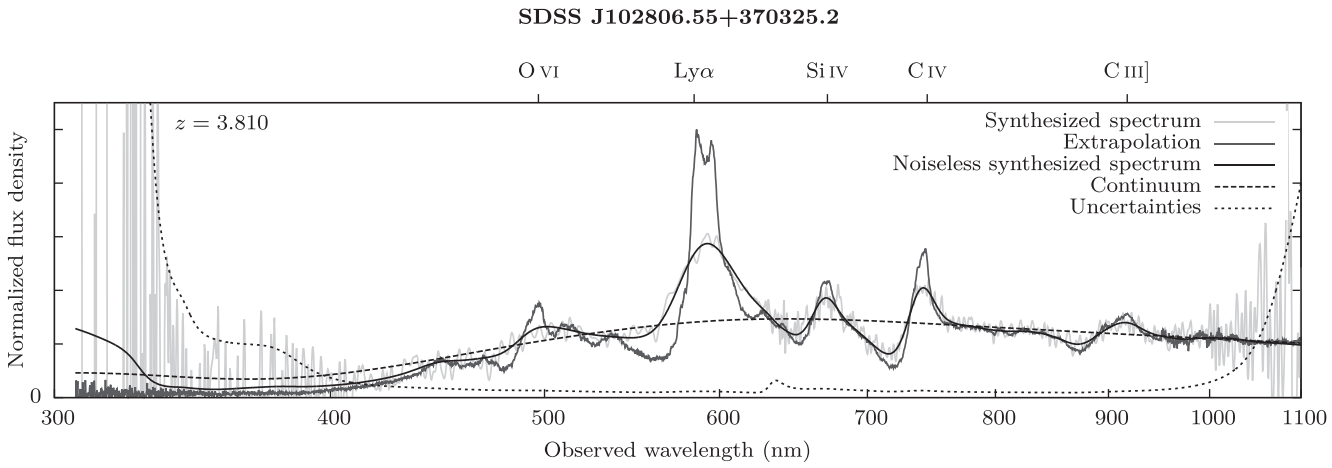

Figure 4. Illustrative synthesized spectrum coming from the linear combination of the interpolated and flat-fielded BP/RP spectra of magnitude $G = 19$. The observed BAL features allow us to highlight the capability of pre-processing we used in order to recover the general shape of the extrapolated spectrum. Flux uncertainties were also pre-processed accordingly, and the BP/RP transition region was unequivocally recognized as a bump around 635 nm. We have to note that the noiseless synthesized spectrum comes from a similar pre-processing applied to a noise-free version of the BP/RP spectra.

5 ASTROPHYSICAL PARAMETER DETERMINATION

The bell shape of the BP/RP spectra prevents us from using both algorithms described in Section 3. This is even more damageable given the fact that these are not sampled on a logarithmic wavelength scale and that the wavelength coverage of each pixel is not uniform. In order to tackle these problems, a resampling of the spectra fluxes and uncertainties was first performed using cubic spline interpolation (Press et al. 2002). The uniform logarithmic sampling we used, $\Delta\log_{10}\lambda = 1.75 \times 10^{-4}$, ensuring a sampling on the redshift that is better than 0.003, which is comparable to human expertise, while producing a reasonable amount of 7.7×10^3 sampled points in the final templates (assuming that $z \leq 6$). Note that such a logarithmic interpolation will obviously introduce covariances between the resulting samples. Nevertheless, given that the specific resampling we used stands in a wavelength range where the corresponding logarithmic function is approximately linear and that both the numbers of samples within the BP/RP spectra and within the synthesized spectra are of the same order of magnitude, we will have that these covariances will be restricted to close neighbouring samples while having moderate magnitudes. These will consequently have a limited impact on the resulting predictions and will hence be ignored in the following. The division of these interpolated spectra by a flat BP/RP spectrum (i.e. coming from a flat input SED) concurrently fixes the bell-shape issue, that is mostly due to the global instrumental response, as well as the problem of the non-uniform wavelength coverage of the pixels, that is due to the inconstant wavelength dispersion function of the BP/RP spectrophotometers. Accordingly, these flat-fielded spectra can then be considered as being approximately proportional to the convolution of the input SED by the LSF over a linear wavelength scale, plus noise. Now, we will have that the resulting spectra will be disjoint although they are overlapping, which would yield to a tremendous loss of efficiency if these were to be considered individually. A more interesting solution stands in the linear combination of the flat-fielded BP/RP spectra according to a given weighting scheme such as to produce a single synthesized spectrum. In more detail, if we consider $s_\lambda^{\rm bp}$ and $s_\lambda^{\rm rp}$ as being the interpolated fluxes of the BP and RP spectra; $\sigma_\lambda^{\rm bp}$, $\sigma_\lambda^{\rm rp}$, as their associated uncertainties; $F_\lambda^{\rm bp}$, $F_\lambda^{\rm rp}$, as their corresponding flat BP/RP fluxes and $w_\lambda^{\rm bp}$, $w_\lambda^{\rm rp}$ as the weighting coefficient used to join these spectra, then we have that the synthesized fluxes, f_λ, can be represented as

$$f_\lambda = w_\lambda^{\rm bp} \frac{s_\lambda^{\rm bp}}{F_\lambda^{\rm bp}} + w_\lambda^{\rm rp} \frac{s_\lambda^{\rm rp}}{F_\lambda^{\rm rp}}, \tag{14}$$

and their associated uncertainties, σ_λ, as

$$\sigma_\lambda = \left[\left(w_\lambda^{\rm bp} \frac{\sigma_\lambda^{\rm bp}}{F_\lambda^{\rm bp}} \right)^2 + \left(w_\lambda^{\rm rp} \frac{\sigma_\lambda^{\rm rp}}{F_\lambda^{\rm rp}} \right)^2 \right]^{\frac{1}{2}}. \tag{15}$$

In the context of this study, the weighting coefficients we selected are given by

$$w_\lambda^{\rm rp} = 1 - w_\lambda^{\rm bp} = {\rm tr}_{620}^{650}(\lambda), \tag{16}$$

where

$${\rm tr}_{\lambda_0}^{\lambda_1}(\lambda) = \frac{1}{2} \tanh \left[2\pi \left(\frac{\lambda - \lambda_0}{\lambda_1 - \lambda_0} - \frac{1}{2} \right) \right] + \frac{1}{2} \tag{17}$$

is the hyperbolic tangent transition function from λ_0 to λ_1. The specific weighting used in equation (16) ensures a smooth transition between the flat-fielded BP and RP spectra while keeping most of their significant regions. A continuum spectrum was then gathered and subsequently subtracted from each synthesized spectrum using a procedure similar to the one described in Section 4.1. An illustrative synthesized spectrum produced through equations (14) and (15) is shown in Fig. 4.

Having now exploitable input spectra, we decided to split our input spectral library into two parts. The first part is used as a 'learning set' (LS1) in order to produce the PCA templates upon which we will base the analysis of the second part, the 'test set' (TS1). Conversely, the second part will be subsequently used as a learning set (LS2) for the analysis of the first part (TS2). This twofold cross-validation procedure will finally provide the APs for the whole set of observations as if these were gathered based upon a totally independent data set. The splitting criterion we used relies on the uniform selection of half the input library sorted according to the QSOs redshifts so as to ensure an even repartition of the latter amongst these two data sets.

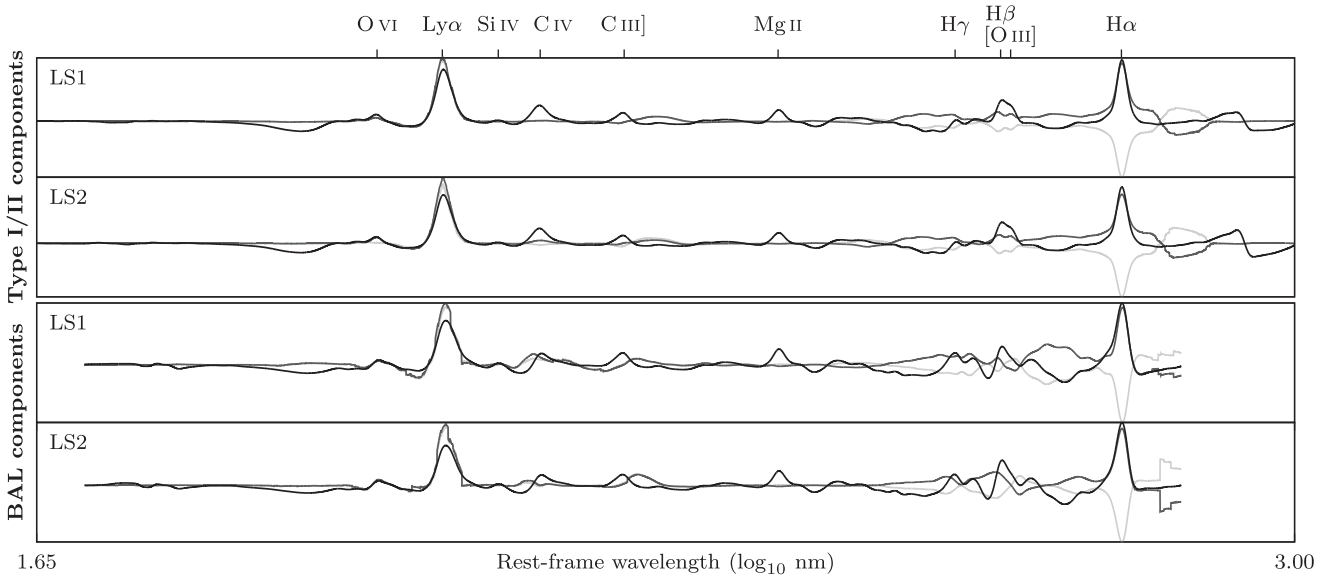

Figure 5. Mean observations (black lines) and first two principal components coming from the PCA decompositions of BP/RP synthesized spectra out of LS1 and LS2 regarding type I/II QSOs (up) and BAL QSOs (down). Both learning sets show extremely good agreement for both QSO types, the noticed discrepancies in the BAL components at wavelength longer than the Hγ emission line being explained by the small number of BAL QSOs having $z \leq 1.42$ ($\sim$1 per cent of the learning sets). We can further notice that most of the BAL features standing between the O VI and C IV emission lines were retained from the DR12Q library (see Fig. 2) while the removal of the fine QSO structures by the BP/RP instrumental convolution now allows the first BAL templates to model the entire emission lines.

Two sets of rest-frame PCA templates were then produced for each of the learning set according to the QSO type (type I/II or BAL). These were based on the noise-free, continuum-subtracted and synthesized BP/RP spectra having both un-normalized G magnitudes and SNR > 1 within the extrapolation procedure. From the noiseless nature of these input spectra, we had to select custom weights associated with the synthesized fluxes, f_λ, as

$$w_\lambda = \mathrm{tr}_{330}^{380}(\lambda)\mathrm{tr}_{1050}^{925}(\lambda)\left[0.7\left(2\mathrm{tr}_{620}^{650}(\lambda)-1\right)^2+0.3\right], \quad (18)$$

where the first two terms practically reflect the limited confidence we set on the spectra edges due, for example, to the potential inaccuracies in the spectra extrapolation or to low fluxes within the flat BP/RP spectra leading to numerical instabilities and where the last term stands for the uncertainties introduced in the BP/RP spectra combination. Fig. 5 provides the mean observation and first two principal components of the synthesized spectra of the type I/II and BAL QSOs for both learning sets.

It is worth to mention that, at first glance, it might seem misleading from the point of view of the validation process to retrieve PCA components from synthesized spectra which are themselves based on the linear combination of templates. Nevertheless, let us first remind that it is one of our assumptions that any (noiseless) DR12Q quasar spectrum can be fairly represented as the linear combination of a sufficient number of such templates. Secondly, we have to note that BP/RP spectra come from the instrumental convolution of the extrapolated spectra in the observed wavelengths. The latter is then set on rest frame, and we will have that the resulting PCA components will have to reflect the averaged convolution applied over the whole observed wavelengths. Finally, this convolution will have the effect of smoothing the high-frequency components from the extrapolated spectra. These are concurrently the main source of unexplained variance within the DR12Q templates, and we expect the produced library to be consistent regarding an hypothetical real noise-free BP/RP spectral library.

Table 2. Warning flags used in the redshift selection procedure. These can be combined through bitwise OR operator.

Warning flag	Value	Condition(s) for rising
`Z_AMBIGUOUS`	1	More than one peak have both $0 < z < 6$ and $\chi_r^2(z) > 0.85$
`Z_LOWCHI2R`	2	$\chi_r^2(z) < 0.9$
`Z_LOWZSCORE`	4	$Z_{\mathrm{score}}(z) < 0.9$
`Z_NOTOPTIMAL`	8	We did not choose the optimal peak (i.e. $\chi_r^2(z) < 1$)

The extracted PCA components were then used in order to produce their CCF against the noisy synthesized BP/RP spectra of magnitude $G = \{18, 18.5, 19, 19.5, 20\}$ through the algorithm described in Section 3.2. The redshift identification is based on the CCF peak having a corresponding redshift in the range $0 < z < 6$; a $\chi_r^2(z) > 0.85$ and a minimal scaled distance from the ideal point $(1, 1) \in (\chi_r^2(z), Z_{\mathrm{score}}(z))$ as given by

$$d(z) = \sqrt{\left(0.8\left[1-\chi_r^2(z)\right]\right)^2 + \left(0.2\left[1-Z_{\mathrm{score}}(z)\right]\right)^2}. \quad (19)$$

The selected redshift was then flagged for potential inaccuracies in the peak selection according to the values provided within Table 2. Constants used in the peak selection procedure as well as within Table 2 are purely empirical and based on a visual inspection procedure.

The optimal number of PCA components to use was chosen as a trade-off between the ratio of explained variance, the ability of the templates to model BAL QSOs and the potential overfit of the observations coming from the use of a too large number of templates. This overfitting is being characterized by frequent ambiguities in the corresponding CCFs that eventually results in a large number of erroneous redshift predictions (though these will have a non-zero warning flag). Tests performed on each learning set show that the use of 3 PCA components is a satisfactory compromise

between these constraints that ultimately lead to a ratio of explained variance of 94.6 per cent (LS1) and 93.42 per cent (LS2) regarding the type I/II QSOs and of 86.22 per cent (LS1) and 88.14 per cent (LS2) regarding the BAL QSOs.

The BAL QSO identification is based on the comparison of the value of the CCF peak we selected using type I/II templates, $y(z)$, against the value of the CCF peak selected from BAL templates, $y_b(z_b)$, through

$$p_b = \frac{y_b(z_b)}{y_b(z_b) + y(z)}, \tag{20}$$

where z is the redshift selected from type I/II templates and z_b is the redshift selected from BAL templates. Though straight classification between these two types of QSOs is sometimes practical, it is however commonly motivated by the specific needs of the end-users. As an example, studying the physics of BAL QSOs will require an extremely pure subset of observations (e.g. with $p_b > 0.7$) while a re-observation survey can easily deal with a 'hint' on the BAL nature of the observed QSOs (e.g. with $p_b > 0.5$). Still, the frequent discrepancies observed between the redshift predicted based upon these two kinds of templates enforce us to use such a classification. Accordingly we will consider, in the following, z_b as the default reshift whenever $p_b > 0.55$ and $G \leq 19$ while keeping p_b as a discriminant value for further application-specific classification. The effect of the thresholding of this discriminant value on the resulting ratio of correctly/incorrectly classified observations will be deferred to Section 6.2.

The slope of the QSO continuum corresponds to the spectral index, α_ν, as defined by

$$f_\lambda \propto \lambda^{-\alpha_\nu - 2} \tag{21}$$

or more compactly expressed in terms of frequency, ν, as $f_\nu \propto \nu^{\alpha_\nu}$. This index is obtained from the fit of a power-law function to the observations over wavelength regions that are commonly devoid from emission/absorption features, that is: 145–148; 170–180; 200–260; 325–470 and 525–625 nm. The exact procedure employs a k-sigma clipping algorithm (with $k = 3$, $\sigma = 1$) so as to underweigh iron emission blends as well as other fortuitous absorption/emission structures by a factor of 100. This procedure was applied to both the input DR12Q spectra and the synthesized BP/RP spectra as a way to fairly compare the resulting predictions while discarding any bias that can be due to the differences in the used algorithms. Because of their high numerical complexities, non-linear optimization algorithms were not used for the least-squares solution of equation (21). Rather, each power-law function was fitted through a linear regression of the wavelengths against the fluxes by taking the logarithm of both sides of the latter equation. Although this choice seems to be harmless from the point of view of the DR12Q spectra, synthesized spectra will have to cope with the large amount of discarded samples coming from the frequent negative fluxes encountered within the spectra edges. With these discarded samples leading to a large bias towards positive fluxes (see Fig. 4, for example), we consequently decided to reject samples standing outside the observed region 350–950 nm for these specific spectra.

Finally, the total equivalent width of the emission lines can be represented as

$$W = \int \frac{e_\lambda}{c_\lambda} d\lambda \tag{22}$$

where c_λ is the continuum slope we fitted based upon equation (21) and e_λ are the emission lines fluxes, with the latter being set to $f_\lambda - c_\lambda$ if λ belongs to an emission line region and to zero otherwise. The identification of these emission line regions is based on parts of the spectrum where smoothed fluxes coming from a 45 points wide Savitsky–Golay filtering (Press et al. 2002) stands higher than the continuum. In agreement with what was previously done, equation (22) was integrated over the interval 380–925 nm regarding DR12Q spectra and over the interval 350–950 nm regarding the synthesized spectra. The rest-frame total equivalent width is hence straightly given by $W_{\rm rest} = W/(z + 1)$.

Table 3. Mean predicted continuum slope, α_ν, and total emission line equivalent width, W, based on DR12Q spectra and synthesized spectra with various G magnitudes.

	Mean α_ν	Mean $\log_{10} W$
DR12Q spectra	-0.697 ± 0.626	1.74 ± 0.236
Synthesized spectra		
$G = 18.0$	-0.561 ± 0.721	1.668 ± 0.298
$G = 18.5$	-0.560 ± 0.722	1.662 ± 0.306
$G = 19.0$	-0.561 ± 0.724	1.648 ± 0.321
$G = 19.5$	-0.569 ± 0.726	1.626 ± 0.344
$G = 20.0$	-0.584 ± 0.723	1.590 ± 0.377

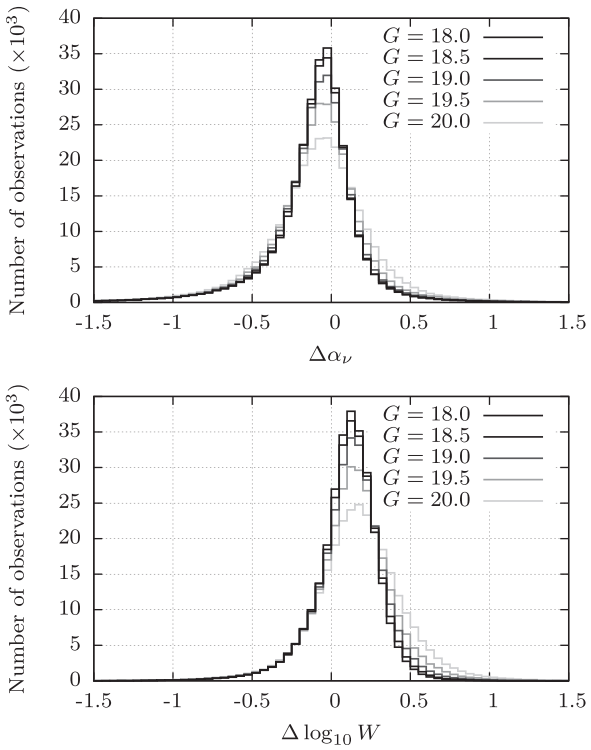

Figure 6. Histograms of the differences between the continuum slopes, α_ν, and total emission line equivalent widths, W, predicted based upon DR12Q spectra and synthesized spectra regarding various normalization magnitudes.

The results of these continuum slope and total emission line equivalent width determination procedures are summarized within Table 3 and Fig. 6, regarding DR12Q spectra and synthesized spectra of various normalizing magnitudes. We can easily see that the continuum slopes predicted from synthesized spectra tend to be bluer than those of the DR12Q spectra. This bias comes from the spread of the Si IV and C IV emission lines over the continuum region 145–148 nm and, to a lesser extent, from a similar spread of the C IV and C III] emission lines over the continuum region 170–180 nm as depicted within Fig. 7.

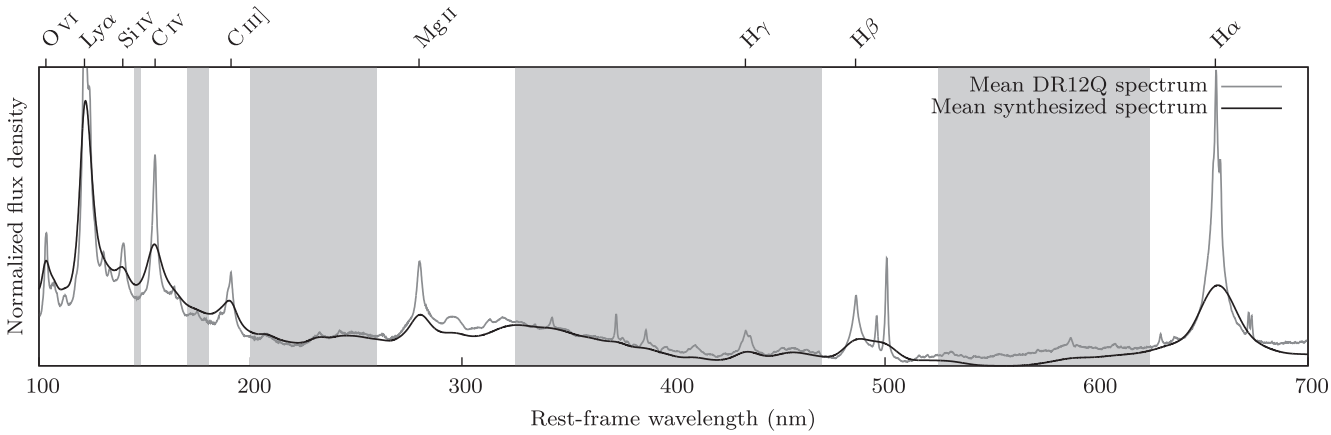

Figure 7. Mean spectra coming from the aggregation of DR12Q spectra normalized so as to have a weighted norm equal to 1 and from synthesized spectra being normalized in the same way. Flux densities were normalized between 200 and 260 nm. Shaded regions correspond to wavelength ranges that were used in order to fit the QSO continuum slopes through equation (21). One can notice the higher flux densities of the synthesized mean spectrum within the continuum regions 145–148 and 170–180 nm when compared to the DR12Q mean spectrum while the lower flux encountered within regions 325–470 and 525–625 nm often yield a red bias during the continuum fitting procedure.

The noticed relation between the flattening of the continuum slope and the normalizing magnitude comes from the increasing number of negative samples that are rejected within the continuum regions 325–470 nm and 525–625 nm where faint fluxes are usually found and which ultimately tend to artificially redden those regions. Also, we can observe that the total equivalent widths of the emission lines predicted based upon synthesized spectra are underestimated compared to the ones predicted using DR12Q spectra. The reason for this similarly stands within the globally overestimated continuum flux as well as from the reddening of the spectrum at wavelengths longer than 300 nm according to the magnitude. Note that the miss of some narrow emission lines because of the LSF convolution also tends to lessen the predicted W. The reader should hence pay a careful attention to these systematic effects once using these measurements.

Given these shortcomings, one might rightfully wonder whether the use of non-linear optimization algorithms is worth to be envisaged in order to predict the continuum slope of the QSOs at the expense of a 10 times longer execution time. Doing so will provide us with a mean value of the continuum slopes of -0.691 ± 0.657 for the DR12Q spectra and a correlation factor of 0.966 if these are compared to the results of our approach. In the case of synthesized spectra, these numbers become, respectively, -0.563 ± 0.721, -0.562 ± 0.737, -0.564 ± 0.737, -0.575 ± 0.739 and -0.601 ± 0.738 for the mean continuum slopes of magnitudes $G = \{18, 18.5, 19, 19.5, 20\}$ with associated correlation factors of 0.988 for $G \leq 19$ and of 0.985, 0.975 for $G = 19.5$ and 20, respectively. The observed flattening of the predicted continuum slopes at $G > 19$ ironically comes from the non-rejection of the negative fluxes from the red part of the spectra which tends to give larger weights to these regions (i.e. the fraction of red fluxes is then more significant). While this effect will have a negligible (but still noticeable) impact on the $G \leq 19$ predictions because of the sufficient SNR of the red part of the spectra at these magnitudes (e.g. the fit of the red part of the spectra is providing a good approximation of the continuum slopes at these magnitudes), it will have a deleterious impact on fainter magnitudes where the red part of the spectra is often better approximated by a flat curve. This effect gets further amplified through the subsequent rejection of the blue fluxes by the k-sigma clipping algorithm. Let us still mention that this non-linear approach remains the most rigorous in a statistical point of view though the strong similarities noticed in both approaches and their common difficulties in predicting the continuum slopes of faint sources do not justify its use regarding its larger time consumption.

6 PERFORMANCE COMPARISON

The performances of our approach were assessed in comparison with the Extremely Randomized Trees learning method (Geurts, Ernst & Wehenkel 2006, hereafter ERT). While classical tree-based learning methods usually try to find, at each node, a splitting criterion (i.e. an attribute and a threshold within this attribute) that is such that the learning set of observations associated with this node is split at best with respect to a given score measure (e.g. variance reduction in regression problem or information gain in classification problem), the ERT instead picks up K random attributes as well as a random threshold associated with each of these attributes in order to select the one maximizing the provided score measure. This procedure is then recursively repeated until the number of learning set observations in all leaf nodes falls under a given limit, $n_{\min}$. The averaged prediction of a set of N trees then allows us to subsequently lessen the variance of the model (i.e. the sensitivity of each individual tree to the used learning set). The choice of this specific method mainly comes from both its fast learning phase as well as from its high performances regarding other competing methods like Artificial Neural Networks or Support Vector Machine while having only a few numbers of parameters to tune. Let us also note that this method is the one that is presently in use within the QSOC software module in order to predict most of the QSO APs (Bailer-Jones et al. 2013).

First of all, let us mention that the QSO continuum slope and that the total equivalent width of the emission lines will not be considered within this performance comparison because these can be straightly predicted based on observable quantities. Regarding the adjustment of the parameters of the ERT models, tests have shown that the prediction of the QSO redshift and type are rather insensitive to the K and $n_{\min}$ parameters if these stand within reasonable ranges of values. Consequently and according to Geurts et al. (2006), the default values of $K = N_{\text{attr}}$, $n_{\min} = 5$ and $K = \sqrt{N_{\text{attr}}}$, $n_{\min} = 2$

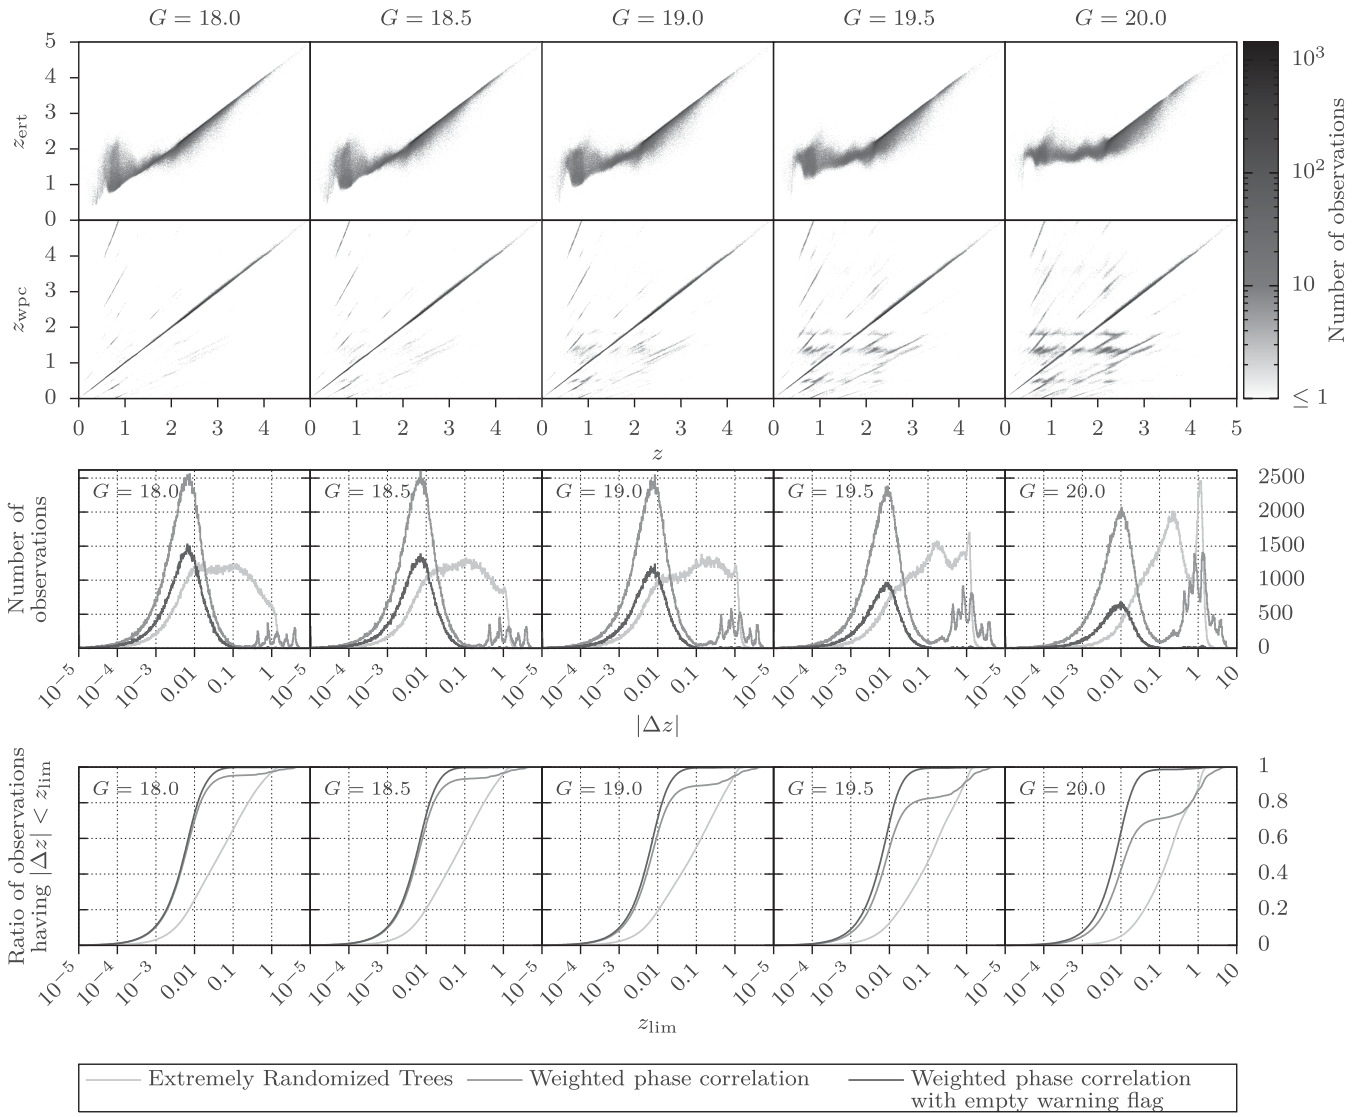

Figure 8. (top) Distribution of the predicted redshifts coming from ERT models, z_{ert}, and WPC model, z_{wpc}, for quasars with G magnitudes equal to 18, 18.5, 19, 19.5 and 20 with respect to the DR12Q redshift, z. (middle) Histogram of the absolute error, $|\Delta z|$, between DR12Q redshifts and redshift predicted based upon ERT models, WPC models and WPC models while having an empty warning flag (see Table 2). (bottom) Ratio of observations having an absolute error on the redshift that is lower than some threshold limit as given by z_{lim}.

were accepted, respectively, for the redshift regression problem and BAL classification problem, and N_{attr} is here the number of points contained within our BP/RP spectra (i.e. $N_{attr} = 960$ if $N_{over} = 8$). The number of trees to build, N, should be ideally as large as possible. Nevertheless, based on time and memory constraints we have, the latter was set to $N = 1000$.

In the following, the ERT models will be built based upon the noisy learning sets LS1 and LS2 where the observations having an SNR > 1 are selected and normalized such as to have a unit norm for the whole set of magnitudes. Their predictions are then gathered from the associated test sets of corresponding magnitude within TS1 and TS2. We have to note that because of selection effects and observational bias within the DR12Q catalogue, neither LS1 nor LS2 will follow a realistic distribution of the redshift (Pâris et al. 2017). Similarly, these will not contain a genuine fraction of BAL QSOs (Reichard et al. 2003; Knigge et al. 2008; Gibson et al. 2009). Consequently, the ERT models that will be built based on these learning sets will be particularly suited for the prediction of the QSO redshifts and type that are the most frequently encountered within LS1 and LS2. In that sense, these will constitute data-oriented models whose predictions on TS1 and TS2 will be optimistic when compared to those based on real observations. Finally, we can further note that the weighted phase correlation algorithm (hereafter WPC) is not sensitive to this data unbalancing and that the associated results will remain valid irrespectively of the actual APs distribution we will encounter.

6.1 Redshift determination

The distribution of the predicted redshifts against DR12Q redshifts is given within the upper part of Fig. 8 for the case of the ERT predictions as well as for the case of the WPC predictions regarding the various normalizing magnitudes. We can already notice a trend of the ERT predictions standing at $z \leq 2$ to be driven towards $z_{ert} \approx 2.3$ where most of our learning set observations stand. This effect is particularly noticeable at $z \approx 0.8$ where our second most numerous

source of QSOs stands and it further tends to strengthen along with an increasing magnitude (though these misclassified observations will typically have strong associated uncertainties). To a lesser extent, we may also note an opposite trend where the observations standing between $z = 2$ and $z = 3$ tend to be underestimated. These effects potentially reflect the inability of our models to fully grab the information that is present within our learning sets and/or the incompleteness of the latter. The fact that very high redshift objects ($z > 3$) get correctly predicted presumably comes from the entrance of the Lyα forest within the observed wavelength range where the extremely faint fluxes found therein allow us to unequivocally characterize these observations.

Although roughly performed here, the analysis of the results coming from machine learning methods often suffer from a lack of physical significance and interpretation that mostly arises because of their underlying complexity. Furthermore, these methods strongly depend on the completeness of the learning set we used in order to build them. In an illustrative purpose, let us consider that a given QSO spectrum is getting a correct redshift prediction from such a model, suppose now that a similar spectrum has a slightly higher redshift which results in a mean shift by a few pixels in the observed spectrum, then nothing ensures that this shifted spectrum will get a correct prediction from the previous model since this ultimately depends on whether or not a somewhat similar spectrum was encountered within the used learning set. According to this, a learning method dedicated to the redshift determination of QSOs should ideally be based on a learning set of observations covering the vast majority of QSO shapes and characteristics over the entire range of redshift we are looking for. With these arguments in mind, we may still suppose that the ERT models we used (and in a broader sense, any model based on machine learning method) are not the best suited in predicting the redshift of QSOs.

Regarding the redshift distribution from WPC (see Fig. 8, up), we can readily notice a tighter dispersion of the errors when compared to the ERT predictions with median absolute errors of 0.0057, 0.0061, 0.0069, 0.0088 and 0.0130 for the WPC predictions of magnitudes $G = \{18, 18.5, 19, 19.5, 20\}$, respectively, and corresponding ERT median absolute errors of 0.0419, 0.0586, 0.0755, 0.1164 and 0.1723 (see Fig. 8, middle). Similarly, 4.83 per cent, 6.66 per cent, 10.56 per cent, 17.6 per cent and 28.93 per cent of the observations have a catastrophic prediction on their redshift (i.e. $|\Delta z| > 0.1$) within the WPC for the same set of magnitudes while the corresponding ratios of ERT observations become, respectively, 35.04 per cent, 40.37 per cent, 45.09 per cent, 52.97 per cent and 63.19 per cent (see Fig. 8, down).

Most of the WPC errors come from mismatches between emission lines. These mainly consist in the confusions of Hβ with C III]; Mg II with Lyα; Mg II with C IV; Mg II with C III] and C IV with Lyα. We have to note that these mismatches do not constitute by themselves real cases of degeneracy but rather arise because of the effect of noise on the emission lines identification as we will soon see. Still, this effect is already unambiguously depicted in Fig. 8 (top), where the number of observations suffering from such an emission line mismatch problem tends to increase along with an increasing magnitude. In the same figure, we may also note that low SNR spectra ($G \geq 19$) tend to produce constant predictions at $z_{\rm wpc} \approx 0.5$, 1.35 and 1.8. These correspond to the fit of deviant fluxes from spectra edges by the Hα, C IV and Lyα emission lines, respectively. Though unavoidable, most of these errors will come along with a non-empty warning flag (see Table 4) which offers the possibility to discard these insecure predictions. By doing so, we rejected 46.2 per cent, 50.99 per cent, 57.72 per cent, 67.09 per cent and 77.58 per cent of the total number of observations regarding magnitudes $G = \{18, 18.5, 19, 19.5, 20\}$, respectively. This leads to corresponding median absolute errors of 0.0053, 0.0055, 0.0058, 0.0063 and 0.0072 and associated ratios of catastrophic redshift predictions of 0.23 per cent, 0.28 per cent, 0.43 per cent, 0.64 per cent and 1.42 per cent for the same set of magnitudes.

Table 4. Ratio of observations triggering warning flags and associated ratio of observations having $|\Delta z| > 0.1$ while triggering these warning flags.

Triggered warning flag	G magnitude 18.00	18.50	19.00	19.50	20.00
Ratio of observations triggering warning flags					
Z_AMBIGUOUS	0.1125	0.1393	0.1853	0.2882	0.3935
Z_LOWCHI2R	0.0084	0.0107	0.0149	0.0310	0.0440
Z_LOWZSCORE	0.4198	0.4602	0.5153	0.5833	0.6733
Z_NOTOPTIMAL	0.0330	0.0427	0.0594	0.1115	0.1636
Any	0.4620	0.5099	0.5772	0.6709	0.7758
None	0.5380	0.4901	0.4228	0.3291	0.2242
Ratio of observations having $\|\Delta z\| > 0.1$					
Z_AMBIGUOUS	0.2934	0.3221	0.3644	0.4290	0.5071
Z_LOWCHI2R	0.5607	0.5744	0.5795	0.6085	0.6474
Z_LOWZSCORE	0.0964	0.1240	0.1814	0.2684	0.3884
Z_NOTOPTIMAL	0.4434	0.4690	0.5032	0.5478	0.6077
Any	0.1018	0.1279	0.1798	0.2592	0.3688
None	0.0023	0.0028	0.0043	0.0064	0.0142

From our previous discussion, we can notice that the performances we gained were achieved at the expense of a very high rejection rate of the observations having a non-empty warning flag. The distribution of these warning flags amongst the observations is given in Table 4 along with their associated ratio of catastrophic redshift prediction once triggered. We first have to note that Z_LOWZSCORE is the most frequently triggered warning flag amongst these observations and is hence the one that contributes at most in their removal. The reason for this stands in the fact that the $Z_{\rm score}$ measure is primarily designed such as to be sensitive to the presence of all the emission lines that are theoretically covered at a given redshift estimate. The miss of one such a line can be attributed either to the wrong redshift estimate we made or to its misidentification owing to its noise or its strong damping by the instrumental convolution, though the right redshift was selected. This misleading distinction is clearly depicted in Table 4, where solely 9.64 per cent of the observations having $G = 18$ and Z_LOWZSCORE flag set comes along with $|\Delta z| > 0.1$, thus arguing for the frequent misidentification of some emission lines while at $G = 20$ this ratio becomes 38.84 per cent, hence consisting in a larger fraction of effective redshift confusions. Secondly, the Z_AMBIGUOUS flag is frequently set because of the intrinsic degeneracy existing in the prediction of the redshift of quasars albeit we can assess from Table 4 that in 70.66 per cent of the cases the right redshift is selected amongst these ambiguous solutions at $G = 18$ while this ratio drops to 49.29 per cent at $G = 20$. For completeness, we have to mention that 49.29 per cent of successful identifications is still better than the ratio that would be obtained from a random selection of the solution, given that the observations having Z_AMBIGUOUS flag set often consist in more than two ambiguous solutions (see Fig. 8, for example). We can further note that, once an ambiguity is detected (i.e. Z_AMBIGUOUS warning flag triggered), the optimal peak of the CCF is commonly selected as the most probable redshift estimate as these do not additionally trigger a Z_NOTOPTIMAL warning flag. Now, if a sub-optimal peak of the CCF is selected, then 55.66 per cent of the observations come along with $|\Delta z| \leq 0.1$ at

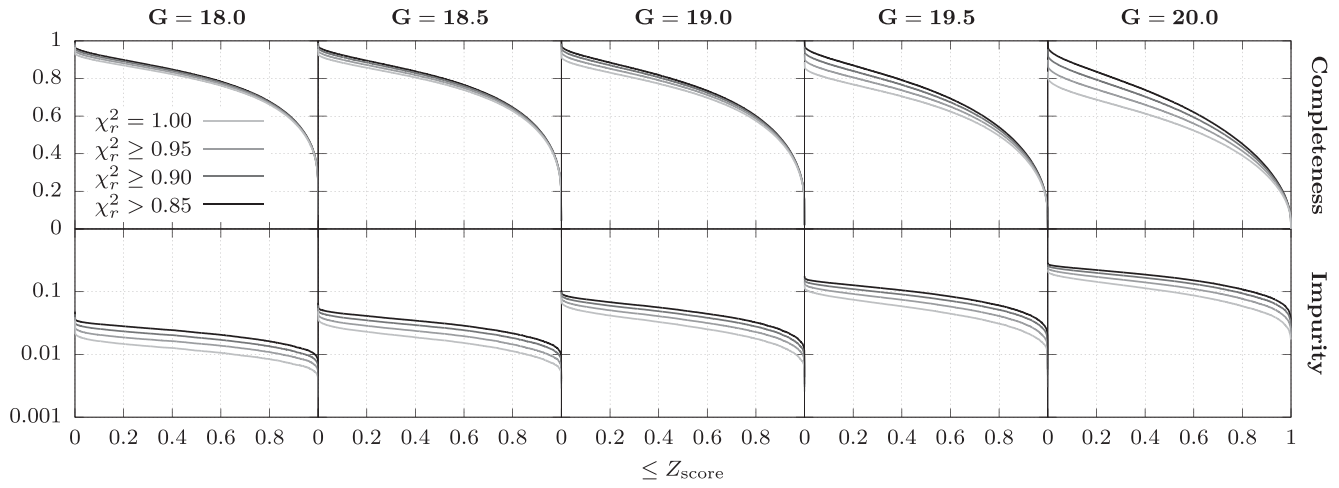

Figure 9. Ratio of observations having Z_{score} and χ_r^2 greater than or equal to the provided values (i.e. completeness) and ratio of observation having $|\Delta z| > 0.1$ amongst the latter (i.e. impurity).

$G = 18$ while this ratio becomes 39.23 per cent at $G = 20$. Given that the latter observations must have a Z_{score} that is greater than or equal to the one associated with the optimal peak of the CCF in order to be selected, these eventually reveal the effective degeneracy that exists in the redshift prediction of quasars once based on low SNR spectra. Finally, the `Z_LOWCHI2R` warning flag is rarely triggered given the strong constraints we set on it (see equation (19)). This decision is further supported by the fact that the associated ratio of catastrophic redshift prediction stands to be the highest amongst the whole set of warning flags.

As pointed out within Section 5, the thresholds that were used in order to trigger the `Z_LOWZSCORE` and `Z_LOWCHI2R` warning flags are somewhat arbitrary and other values might be better suited regarding the specific needs of the end user. This is particularly true given that these were shown to have a strong impact on the trade-off between the completeness and the impurity of our predictions as we have just seen. These ratios of completeness and impurity are given in Fig. 9 for varying thresholds on χ_r^2 and Z_{score}. Note that we do not consider the `Z_AMBIGUOUS` and `Z_NOTOPTIMAL` warning flags in this analysis, given that $\chi_r^2 < 1$ automatically implies that both these flags are set. Also remind that we required $\chi_r^2 > 0.85$ so as to limit the number of ambiguous solutions that are potentially associated with each observation.

6.2 BAL binary classification

The data unbalancing has a particularly insidious impact on the analysis of the results coming from the binary classification of BAL QSOs. Indeed, based on the fact that solely 9.95 per cent of the DR12Q observations are BAL QSOs, a model that will systemically classify the observations as type I/II QSO would then provide a satisfactory ratio of correctly classified observations (i.e. an *accuracy*) of 90.05 per cent while no BAL QSO will be identified. Consequently, this ratio will not constitute an objective analysis tool if considered alone. We will hence use two additional and complementary statistical measures that were specifically designed for the analysis of the performance of binary classifiers. First, the *true positive rate* (hereafter TPR) will here denote the fraction of BAL QSOs that are correctly identified by a given model. It constitutes an estimator of the probability of detection of the BAL QSOs by this model. Secondly, the *false positive rate* (hereafter FPR) will denote the fraction of type I/II QSOs that are wrongly classified as BAL QSOs. A perfect binary classifier should hence have TPR $= 1$ along with FPR $= 0$. Note that both these statistical measures can be adjusted by varying the user-defined threshold that was set either on p_{b}, for the case of the WPC, or on the number of trees that voted for the BAL class regarding the ERT. By doing so and reporting the corresponding TPR against FPR, we obtain the so-called *Receiver Operating Characteristics* (ROC) curve as depicted within Fig. 10 for the case of the WPC and ERT models for quasars with magnitudes $G = \{18, 18.5, 19, 19.5, 20\}$. These curves allow us to straightly compare the performances of these two competing models while depending neither on the data unbalancing, nor on the specific thresholds we used. The area under the ROC curve is then often taken as a fair indicator of their global performances.

Now, like many data-reduction pipelines, our primary objective will be to optimize the accuracy of our model with respect to the fraction of BAL QSOs that will be encountered amongst the real observations. We will then have to take into account the potential unbalancing that will be present within the *Gaia* observations. However, because of the uncertainties surrounding the selection effect from DSC as well as the observational bias, this unbalancing is not known a priori. Consequently, we decided to consider a fraction of BAL QSOs, r_{b}, equal to the one that is present within the DR12Q catalogue (i.e. $r_{\mathrm{b}} = 0.0995$). The presented accuracies should hence be updated once a realistic ratio will be available though the general conclusions drawn out of these are not supposed to change (assuming that r_{b} remains small). We can then easily figure out that the regions of the ROC curves where the accuracy is constant correspond to lines whose equations are given by

$$\mathrm{TPR} = \frac{1 - r_{\mathrm{b}}}{r_{\mathrm{b}}} \times \mathrm{FPR} + C. \tag{23}$$

Our goal will then be to find the point(s) of the ROC curve that intersect such a line while maximizing C. Note that the trivial case where $C = 0$ corresponds to the accuracy that would be obtained by a constant type I/II classifier which is thereby always achievable. Stated otherwise, the point(s) of the ROC curve having an optimal associated accuracy correspond(s) to the one (or those) whose distance to the line of constant type I/II accuracy is the greatest while being on its left side.

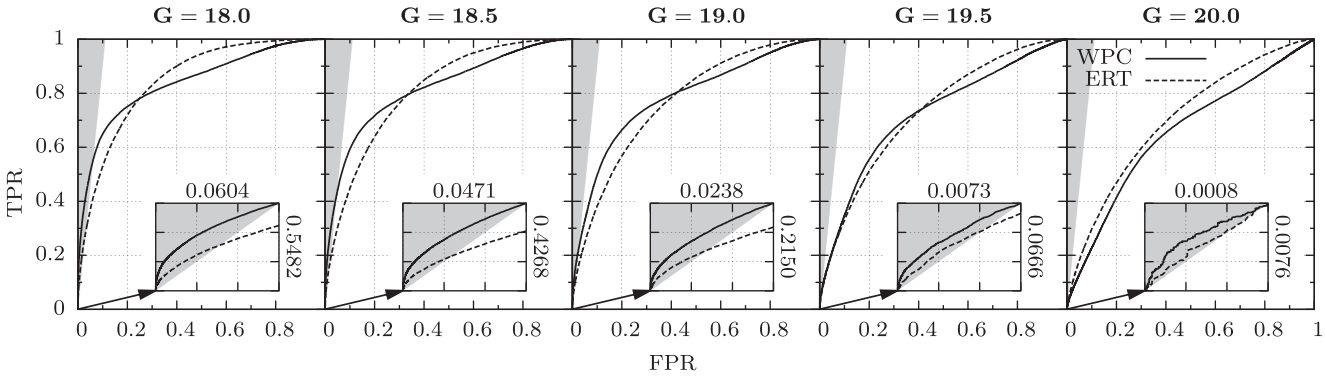

Figure 10. ROC curves of BAL binary classification regarding the WPC and ERT models for quasars with G magnitudes equal to 18, 18.5, 19, 19.5 and 20. Shaded areas correspond to regions where the accuracy stands higher than or equal to the one coming from a constant type I/II classifier. The points of the ROC curves associated with the optimal accuracies stand at the highest distances from the border of these regions while being contained in it.

Table 5. Optimal accuracies of the ERT and WPC models for quasars with G magnitudes of 18, 18.5, 19, 19.5 and 20 along with their associated thresholds, TPR and FPR.

G magnitude	Threshold	TPR	FPR	Accuracy
		ERT		
18.00	552	0.15553	0.01048	90.644%
18.50	604	0.07920	0.00448	90.428%
19.00	604	0.05204	0.00345	90.252%
19.50	598	0.02857	0.00235	90.119%
20.00	617	0.00352	0.00029	90.058%
		WPC		
18.00	0.5341	0.31992	0.01666	91.725%
18.50	0.5517	0.21295	0.01105	91.169%
19.00	0.5843	0.09140	0.00437	90.562%
19.50	0.6102	0.03348	0.00203	90.198%
20.00	0.6486	0.00413	0.00023	90.069%

Fig. 10 focuses on the regions of the ROC curves where the accuracy stands higher than the one of a constant type I/II classifier. From our previous discussion, we can readily see that the WPC models have overall better accuracies when compared to the ERT models for the whole set of G magnitudes. The best achievable accuracies are summarized in Table 5 along with their associated thresholds, FPR and TPR for G magnitudes of 18, 18.5, 19, 19.5 and 20. The extremely low TPR found therein can be explained by both the relatively low SNR of the synthesized spectra as well as by the removal of most of the narrow absorption features by the LSF convolution and/or by the undersampling of these spectra. The effect of noise can be readily recognized based on Fig. 10, where the ROC curves tend to match the ones that would be obtained from a random classifier (i.e. a diagonal line) as we increase the G magnitude. This translates as a drop of the point of optimal accuracy along the ROC curves which consists in both a lower TPR and a compensating lower FPR (see Table 5). In extreme cases, BAL QSOs become barely identifiable with a probability of detection within the WPC of 3.348 per cent for G magnitude of 19.5 and of 0.413 per cent for the case of $G = 20$. Fig. 11 compares the TPR of the WPC with the Balnicity index of the C IV trough (Weymann et al. 1991, hereafter BI) for the various normalizing magnitudes. This BI can be seen as a modified equivalent width of the BAL absorption occurring in the blue part of the C IV emission line. We can notice a strong dependence of the TPR according to BI, which reflects the difficulty

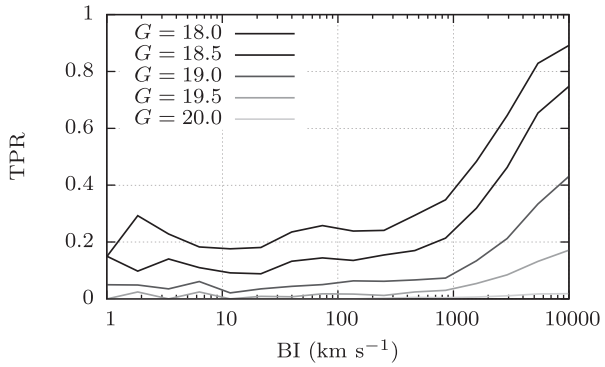

Figure 11. Comparison of the Balnicity index of the C IV trough with respect to the TPR of the WPC models regarding BAL QSOs with various magnitudes.

in identifying BAL QSOs having narrow absorption features. We can finally notice that if one can afford to have a high FPR, then the ERT provides a better TPR than the WPC. This would be the case, for example, if we would like to filter the *Gaia* catalogue by keeping most of the the BAL QSOs while rejecting a still significant number of type I/II QSOs.

7 DISCUSSION

Although already fully operational, the presented software module may still experience some minor improvements that will be summarized in the remainder. First, we did not consider any extinction by the interstellar medium. The associated correction relies on the availability of a wavelength-dependent extinction law such as the one of Fitzpatrick (1999) as well as on a map of galactic extinction like the one that will be produced by the Total Galactic Extinction software module from CU8 (Bailer-Jones et al. 2013). The total equivalent width of the emission lines as well as the continuum slope from equation (21) might benefit from this correction. Nevertheless, due to the fact that the continuum slopes we subtracted from our synthesized spectra are purely empirical (see Section 5), these will also contain most of the encountered extinction. Accordingly, this correction is not expected to bring any major improvement neither on the prediction of the redshift of QSOs nor on the subsequent calculation of the BAL discriminant value, p_b. Furthermore, based on the fact that most of the DR12Q spectra stand at relatively

high galactic latitude (i.e. $|b| > 30^{\circ}$) where the extinction is weak, the spectra we used in this study were not much affected by this extinction. A more challenging objective would be to enable the prediction of this extinction based on the BP/RP spectra of quasars. This problem is currently being investigated but seems to be hardly attainable because of the degeneracy existing between the extinction curve and the intrinsic continuum slope of the QSOs. Secondly, the computation of a χ^2 value from the optimal point of the CCF (see equation (6)) can straightly allow us to send feedback about the potential misclassification of the quasars we received from the DSC module.

8 CONCLUSION

We have described in this work the processing of the BP/RP spectra coming from the *Gaia* satellite in order to determine the APs of quasars within the QSOC module of the CU8 CU from the DPAC. These APs encompass: the redshift of the QSOs, their continuum slopes, the total equivalent width of their emission lines and whether or not these are broad absorption lines (BAL) QSOs. We have highlighted the necessity to have fast and reliable algorithms so as to deal with the huge amount of spectra that *Gaia* will provide as well as with their limited signal-to-noise ratio and resolution. We have introduced two already developed algorithms, namely the weighted principal component analysis and the weighted phase correlation, that were specifically designed in order to fulfil both these mentioned objectives and whose combination allows both to securely predict the redshift of the QSOs and to set a discriminant on their type. We have presented the construction of a semi-empirical library of BP/RP spectra based on the *Gaia* instrumental convolution of the observations coming from the Sloan Digital Sky Survey which were extrapolated in order to cover the wavelength range of the BP and RP spectra. We saw the pre-processing that is required in order for these BP/RP spectra to be fully exploitable by our algorithms as well as the methods we used for predicting the various APs. Some systematic bias was noticed within the prediction of the continuum slopes and of the total equivalent width of the emission lines. This bias can be mostly explained by both the spread of the Si IV, C IV and C III] emission lines over the continuum regions situated between 145–148 nm and 170–180nm as well as by the rejection of the negative fluxes that are usually found within the red part of the pre-processed spectra.

A comparison with the currently used machine learning method showed that our approach is the one of predilection for the determination of the redshift of the quasars while benefiting from a straight physical significance as well as from strong diagnostic tools on the potential errors that may arise during predictions. Cross-validation tests showed that 95.17 per cent, 93.34 per cent, 89.44 per cent, 82.4 per cent and 71.07 per cent of the observations come along with an absolute error on the predicted redshift that is lower than 0.1 for the case of quasars with G magnitudes equal to $G = \{18, 18.5, 19, 19.5, 20\}$. These ratios become, respectively, 99.77 per cent, 99.72 per cent, 99.57 per cent, 99.36 per cent and 98.580 per cent once the insecure predictions are discarded based on the triggering of some warning flags. We explored the repartition of these warning flags amongst the observations and studied the effect of setting customized warning thresholds on the trade-off between the completeness and the impurity of our predictions. Our methods were proved to yield the best ratio of correctly classified observations regarding the identification of BAL QSOs assuming that these will be observed much less frequently than the type I/II QSOs. Machine learning methods may still provide a better probability of detection of these BAL QSOs at the expense of much higher contamination rates. Finally, we have that 91.725 per cent, 91.1069 per cent, 90.562 per cent, 90.198 per cent and 90.069 per cent of the observations were correctly classified by our methods regarding quasars with G magnitudes of 18, 18.5, 19, 19.5 and 20, respectively.

ACKNOWLEDGEMENTS

The author acknowledges support from the ESA PRODEX Programme 'Gaia-DPAC QSOs' and from the Belgian Federal Science Policy Office.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation and the U.S. Department of Energy Office of Science. The SDSS-III web site is http://www.sdss3.org/.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofisica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington and Yale University.

REFERENCES

Bailer-Jones C. A. L. et al., 2013, A&A, 559, A74
Bailey S., 2012, PASP, 124, 1015
Bolton A. S. et al., 2012, AJ, 144, 144
Claeskens J.-F., Smette A., Vandenbulcke L., Surdej J., 2006, MNRAS, 367, 879
Dawson K. S. et al., 2013, AJ, 145, 10
de Bruijne J. H. J., 2012, Ap&SS, 341, 31
Delchambre L., 2015, MNRAS, 446, 3545
Delchambre L., 2016, MNRAS, 460, 2811
Fitzpatrick E. L., 1999, PASP, 111, 63
Gaia Collaboration et al., 2016, A&A, 595, A1
Geurts P., Ernst D., Wehenkel L., 2006, Machine Learn., 63, 3
Gibson R. R. et al., 2009, ApJ, 692, 758
Glazebrook K., Offer A. R., Deeley K., 1998, ApJ, 492, 98
Hamann F., Ferland G., 1999, ARA&A, 37, 487
Jordi C., Fabricius C., Carrasco J. M., Figueras F., Masana E., Voss H., Luri X., 2008, in Jin W. J., Platais I., Perryman M. A. C., eds, IAU Symposium Vol. 248, A Giant Step: from Milli- to Micro-arcsecond Astrometry. p. 500
Jordi C. et al., 2010, A&A, 523, A48
Knigge C., Scaringi S., Goad M. R., Cottis C. E., 2008, MNRAS, 386, 1426
Lindegren L., Lammers U., Hobbs D., O'Mullane W., Bastian U., Hernández J., 2012, A&A, 538, A78
López-Corredoira M., Melia F., Lusso E., Risaliti G., 2016, Int. J. Mod. Phys. D, 25, 1650060
Mignard F., 2009, in IAU Symposium #261, American Astronomical Society. p. 891
O'Mullane W. et al., 2007, in Shaw R. A., Hill F., Bell D. J., eds, ASP Conf. Ser. Vol. 376, Astronomical Data Analysis Software and Systems XVI. Astron. Soc. Pac., San Francisco, p. 99

Pâris I. et al., 2017, A&A, 597, A79
Pearson K., 1901, Phi. Mag., 2, 559
Perryman M. A. C. et al., 2001, A&A, 369, 339
Perryman M., Hartman J., Bakos G. Á., Lindegren L., 2014, ApJ, 797, 14
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 2002, Numerical Recipes in C++ : The Art of Scientific Computing, 2nd edn. Cambridge Univ. Press, New York
Reichard T. A. et al., 2003, AJ, 126, 2594
Schneider P., Sluse D., 2013, A&A, 559, A37
Tsalmantza P., Hogg D. W., 2012, ApJ, 753, 122
Weymann R. J., Morris S. L., Foltz C. B., Hewett P. C., 1991, ApJ, 373, 23
Yahata K. et al., 2005, PASJ, 57, 529

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

DR12Q_PCA.fits
DR12Q_PCA.tar.gz

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

This paper has been typeset from a TEX/LATEX file prepared by the author.

4

Identification of gravitational lens candidates within the Gaia mission

As outlined in section 1.3, quasars reveal a supplemental interest once their light rays get deflected by a foreground galaxy being in the line-of-sight. Upon a favourable alignment between the observer, the deflecting galaxy and the background quasar, multiple images of this quasar can be produced that may lead to some of the currently most important cosmological applications, as already explained in section 1.3.3. As a reminder, an immediate benefit of such a gravitational lensing (GL) phenomenon stands in the fact that some of these images are magnified versions of what would be observed if no deflector was present. This allows high quality observations of these very distant objects to be obtained. Also, independent estimations of the cosmological parameters can be gained if the time delays occurring between the photometric variability curves of each of these images is taken into account. Gaia, on its side, is expected to have an astrometric precision of $\sigma_0 = 462$ μas along with an angular separation power of 200 mas and a photometric precision of $\sigma_G = 0.0312$ mag regarding quasars of magnitude $G = 20$ mag, assuming these have an observed color $V - I = 0.225$ mag on average (Gaia Collaboration et al. 2016). Owing to these high performances and to its full sky coverage, Gaia hence offers a unique opportunity to identify strong gravitational lens (GL) systems. This problematic being explored in the present chapter.

The goal of the GL identification is then to predict whether compact configurations of three or four point-like images constitute plausible GL candidates based on the sole positional and photometric informations of these images. A GL candidate being here defined as a configuration of images that is compatible with the one that would be produced through a non-singular isothermal ellipsoid in presence of an external shear (NSIEg) lens model. The reason for considering this particular model stands in the fact that it is rather general as it allows most of the commonly observed GLs to be modelled. GLs whose deflectors are a cluster of galaxies will accordingly not be addressed here as these require more complex mass distributions to be considered (see however Keeton 2001, for an overview of these methods).

This effort is conducted through the gravitational lenses in Gaia (GraL) working group composed of a dozen scientists, mainly coming from the data processing and analysis consortium (DPAC). The GraL activities are carried out in parallel through

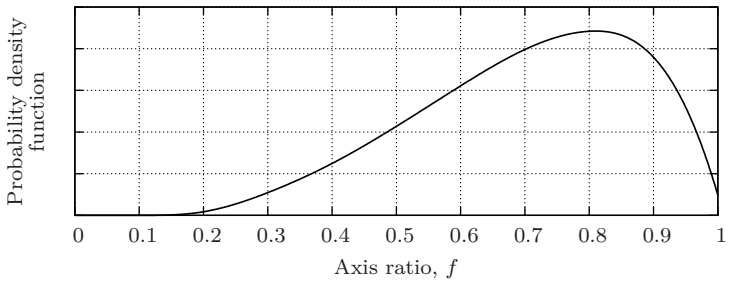

Figure 4.1: Probability density function of the axis ratio of early type galaxies (Choi et al. 2007).

the GravLens chain out of the DPAC CU4 as well as through the systematic exploration of lenses from astrometry (SELenA) pipeline. This distinction is made here because the GravLens chain, although having access to privileged informations like sky mappers or IDTs under the DPAC data access policy, is contractually restricted to provide its results in the final data release of the Gaia catalogue. Still, as earlier releases are already fully exploitable, a concurrent implementation (i.e. the SELEnA pipeline) was developed so as to accomplish similar objectives while only relying on public data releases.

Two approaches can be envisaged to identify GLs:

- Quasars represent only a small fraction of the celestial objects at $G \lesssim 22$ mag, the presence of a companion object in the near vicinity of a known quasar being hence a first hint of a GL candidate, this probability is furthermore reinforced if this companion object is itself a quasar. A cross-match performed between the Gaia data release 1 (GDR1) and the large quasar astrometric catalog (Souchay et al. 2015, hereafter LQAC) accordingly provided ~ 1300 configurations composed of two images, ~ 70 configurations having three images and 10 configurations having four images. Although most of these configurations correspond to fortuitous combinations of foreground star(s) and quasar(s) projected on the celestial sphere, some well-known GLs as well as some GL candidates are still recovered thanks to this technique.

- Secondly, as an important fraction of the quasars that Gaia will identify will be new discoveries, the search around known quasars from the previous point suffers from a critical lack of completeness. This is particularly true as the classification of the objects that will be detected by Gaia through the DSC module only aims to be published in the third data release of the Gaia catalogue. A blind identification of GLs performed over the whole sky can accordingly bring a non-negligible additional set of GL candidates before we get the whole quasar identifications from DSC.

The method we describe below was primarily designed in order to provide an estimation of the capability of the NSIEg lens model to reproduce the observed configurations of the lensed images without having to rely on slow procedures that are commonly used for this purpose as, for example, the sampling of the posterior probability function based on a Bayesian approach to the problem (Keeton 2010). This identification being particularly well suited for an analysis through machine learning (ML) methods as those described in section 2.4.

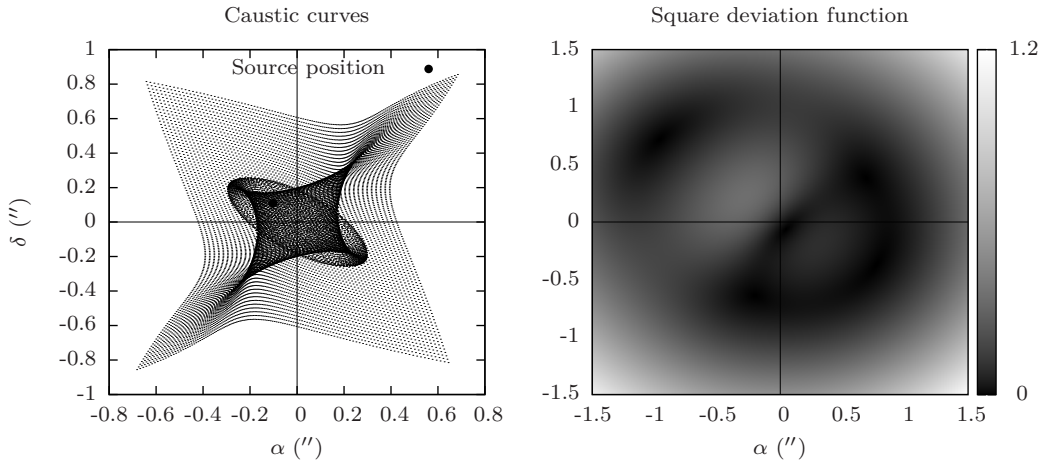

Figure 4.2: Illustration of the caustic curves obtained by projecting a uniform grid of points from the image plane back into the source plane for a given GL model (left) and associated square deviation function (right). The parameters of the NSIEg GL model we consider here are given by $\theta_e = 1.23\ ''$, $\theta_c = 0.19\ ''$, $f = 0.56$, $\gamma = 0.07$ and $\omega = 1.89$.

Supervised learning methods however require a rather exhaustive learning set of observations (LS) to reliably guess the relations that may exist between a set of input attributes and an output variable (see section 2.4). This output variable being here associated with a binary value depending on whether a given configuration of images can be modelled through a NSIEg lens model, or equivalently if this configuration is a GL candidate. A very limited sample of GLs (~ 100) is actually available such that a simulation of these GL systems has to be performed so as to yield an arbitrary large LS. The idea is then to build a set of GL lens candidates based on a realistic sampling of the parameters of the NSIEg lens model as summarized in Table 1.3. More precisely, if we assume the Einstein radius, θ_e, to be a simple scaling factor, all these parameters turn out to follow uniform distributions at the exception of the axis ratio, f, where a Monte-Carlo Markov chain (MCMC) sampling of the probability distribution function illustrated in figure 4.1 had to be used. Other parameters were then uniformly drawn from the interval $\theta_e \in [0.1, 1]$ in the case of the Einstein radius, in the range $\theta_c \in [0, 0.1\,\theta_e]$ for the angular size of the core of the deflector, in the range $\gamma \in [0, 0.1]$ for the shear intensity and with $\omega \in [0, \pi]$ regarding the shear orientation. Once these parameters are correctly settled, we randomly probed 100 source positions for each so built NSIEg deflector while only keeping those providing more than three lensed images (one of these images being systematically strongly de-amplified it will generally not be detected and can be safely ignored).

Although apparently straightforward, the derivation of the image positions turns out to be a challenging problem as it requires the lens equation, $\boldsymbol{\theta_s} = \boldsymbol{\theta} - \boldsymbol{\alpha}(\boldsymbol{\theta})$, to be inverted. That is finding the image positions, $\boldsymbol{\theta}$, that are associated with a source position $\boldsymbol{\theta_s}$ is a complicated task for which analytical solutions rarely exist. At the opposite, finding the source position $\boldsymbol{\theta_s}$ that match an image position $\boldsymbol{\theta}$ is trivial as it only requires the lens equation to be evaluated. Numerical methods are hence frequently employed to 'invert' the lens equation that relies on this straight evaluation. Based on this idea and on the previous work of Keeton (2010), we used a tiling of the image plane we propagated back into the source plane to approximate the *caustic curves* that are associated with the GL system (see figure 4.2 for an illustration). The

overlapping of these tiles in the source plane effectively provides the number of images that will be produced for a given GL system. The positions of these images are then refined through the use of a Nelder-Mead optimization algorithm (Nelder & Mead 1965) applied to the square deviation function (Schramm & Kayser 1987, hereafter SDF) as defined by

$$S^2(\boldsymbol{\theta}) = \|\boldsymbol{\theta_s} - \boldsymbol{\theta} + \boldsymbol{\alpha}(\boldsymbol{\theta})\|^2 . \tag{4.1}$$

Let us note that this equation yields the squared distance between the source position and the projection of a point $\boldsymbol{\theta}$ from the image plane back into the source plane as illustrated in figure 4.2. More particularly, $S^2(\boldsymbol{\theta}) = 0$ implies that the position $\boldsymbol{\theta}$ is associated with an image of the source according to the current GL model and source position.

The previous simulation resulted in 8×10^6 GL candidates having more than three images, 97.4% out of these having five images as expected from the odd number theorem (Burke 1981). According to this theorem, GLs having a smooth surface mass density which decreases faster than $|\boldsymbol{\theta}|^{-1}$ always produce an odd number of images. The remaining 2.6% of GLs having four images hence corresponds to cases where we voluntarily considered two of these lensed images as being indistinguishable one from another (i.e. in cases where $|\Delta\boldsymbol{\theta}| < 0.1\,\boldsymbol{\theta_e}$). The removal of one or two of the faintest image(s) out of these 8×10^6 candidates finally provides us with two sets of simulated observations having an approximately equal size and corresponding to GL candidates having four and three lensed images, respectively. Alternative sets of contaminant observations were also subsequently built based on the Besançon model of stellar population (Robin et al. 2003), out of which three and four image configurations were randomly drawn. These have a difference in magnitude of $\Delta G \leq 2.5$ mag and a minimal image separation of $|\Delta\boldsymbol{\theta}| \geq 0.1\,\Delta\boldsymbol{\theta}_{\max}$, where $\Delta\boldsymbol{\theta}_{\max}$ is the maximal separation between any two images. Half of the simulated GL candidates as well as half of the contaminant observations were then merged for each of the three and four image datasets to yield LS, the other half being kept as test sets of observations (TS). Three noise realizations were produced out of these TS corresponding to: a perfect instrument ($\sigma_{\boldsymbol{\theta}} = 0$ mas, $\sigma_G = 0$ mag), the expected performances of Gaia on typical GLs ($\sigma_{\boldsymbol{\theta}} = 10$ mas, $\sigma_G = 0.02$ mag) and to a defeatist approach ($\sigma_{\boldsymbol{\theta}} = 100$ mas, $\sigma_G = 0.05$ mag). Finally, all observations were normalized so as to have their brightest images standing at $\boldsymbol{\theta} = (0, 0)$ along with a G magnitude of 0 and their second brightest image standing at $\boldsymbol{\theta} = (1, 0)$. Figure 4.3 illustrates the result of this simulation procedure on TS.

Before building the ML models upon which the identification of the lens candidates will be performed, one might rightly wonder what is the fraction of GL candidates as compared to the number of fortuitous clusters of celestial objects. The exact answer to this question is however not as straightforward as it may seem given that it requires strong lensing probability to be taken into account as well as the precise inhomogeneities of stars (e.g. galactic plane, globular clusters, ...) to be considered. Rather, an approximate lower limit on the fraction of GL candidates will be derived here based on simple assumptions. We suppose first that most GLs whose deflectors are galaxies are found within a radius of 2.5 ″centred around these galaxies such that we are only interested in clusters of celestial objects standing in a spherical cap of similar radius. From Poisson statistics, we have that the probability that k celestial objects stand within a given spherical cap of radius r is given by

$$P(k \in \Omega_r) = \frac{\lambda^k e^{-k}}{k!} \tag{4.2}$$

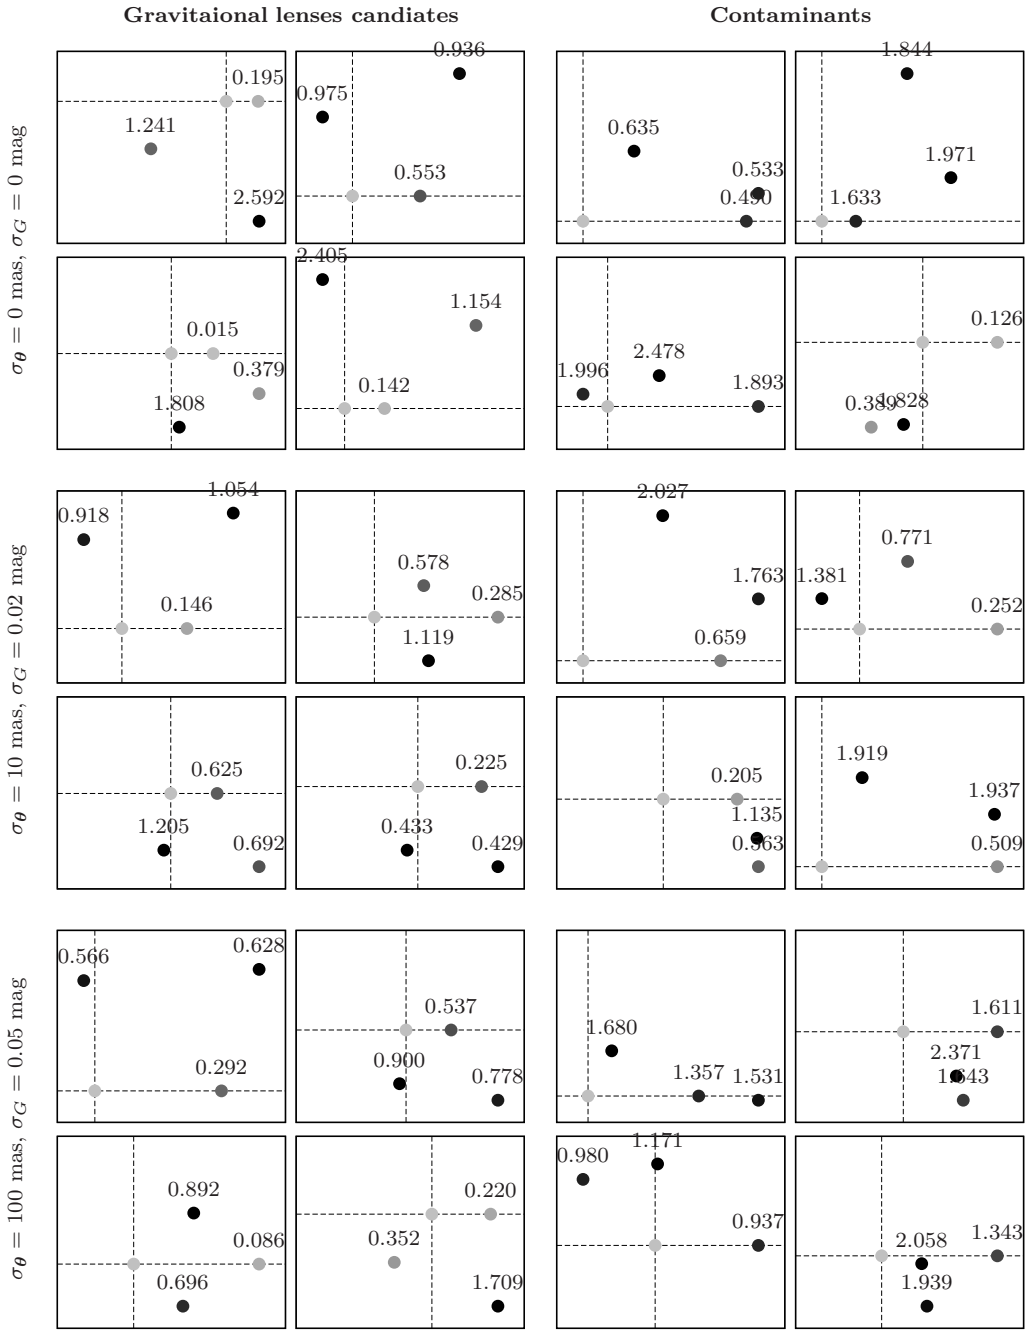

Figure 4.3: Examples of simulated GL candidates and contaminants taken from TS regarding various realizations of the noise. The origin of the system stands on the brightest image which is also normalized so as to have a magnitude of $G = 0$ mag, the second most luminous image standing then at $\boldsymbol{\theta} = (1, 0)$. Images are labelled by their relative magnitudes as compared to the brightest image. The associated sets of observations coming along with three images are then obtained by removing the faintest image out of these configurations.

Table 4.1: Extremely randomized trees (ERT) performances obtained in the identification of GL candidates once requiring the FPR to be comparable to the fraction of GL candidates that will be present in the observations (i.e. 0.1% of the configurations composed of three images will consist of GL candidates while 1% of the configurations composed of four images will be candidates, as derived in the text). This eventually yields a contamination ratio among the predicted GL candidates of $\sim 50\%$.

Number of images	$\sigma_{\boldsymbol{\theta}}$ (mas)	σ_G (mag)	**Threshold**	**TPR**	**FPR**
3	0	0	> 0.999	0.25992	0.0013024
3	10	0.02	> 0.999	0.22410	0.0012945
3	100	0.05	> 0.999	0.13693	0.0012994
4	0	0	> 0.199	0.99957	0.0098688
4	10	0.02	> 0.197	0.99940	0.0099865
4	100	0.05	> 0.237	0.98938	0.0099954

where λ is the expected number of celestial objects in a spherical cap of radius r, subtending a solid angle Ω_r on the sphere. In the case of Gaia, 1.1×10^9 sources are expected to be observed (Robin et al. 2012), supposing these are uniformly distributed over the celestial sphere[1], we have that

$$\lambda = 1.1 \times 10^9 \, \frac{\Omega_r}{4\pi}, \qquad \text{with } \Omega_r = \oint \int_0^r \sin\theta d\theta d\phi = 4\pi \sin^2 \frac{r}{2}, \tag{4.3}$$

leading to $\lambda = 0.04$ if $r = 2.5\ ''$. We hence have that $P(2 \in \Omega_r) = 10^{-4}$ and that $P(3 \in \Omega_r) = 5.5 \times 10^{-7}$ such that if we consider spherical caps centred on all of the 1.1×10^9 Gaia sources, the approximate number of fortuitous configuration of three and four images we obtain are respectively given by 1.1×10^5 and 605. These numbers have to be compared to the ~ 80 GLs having more than two lensed images as well as to the few GLs having four lensed images that Gaia is expected to observe (Finet & Surdej 2016). Accordingly, we estimate a contamination ratio among GLs candidates in the order of 99.9% in the case of three image configurations and in the order of 99% in the case of four image configurations. These ratios must however be taken with the caution they deserve given their coarse derivation as well as the strong assumptions we made.

An ensemble of $N = 500$ extremely randomized trees (ERT) was then built based on LS for each of the three and four image configurations. Tests have shown that the present problem is rather insensitive to the $n_{\min}$ parameter of the ERT that was accordingly set to its default value $n_{\min} = 2$. On the other hand, all attributes were considered regarding the K parameter (see section 2.4). The predictions of these models are assessed based on TS having varying noise realizations, as we already described. Figure 4.4 and Table 4.1, present the result of this validation procedure. Although three image GL candidates are barely identified ($13\% < \text{TPR} < 26\%$) owing to the high contamination rate among these configurations which in turn results in a required extremely low $\text{FPR} < 0.1\%$, four image GL candidates are easily identified as, for example, a TPR of 90.1% is still associated with a FPR of 0.007%.

These ERT models were then ran on GDR1 though no serious GL candidates having four images were identified. The reason for this stands in the filtering that is

[1]This assumption is obviously false but it still provide us with an upper limit on λ given that most of the GLs will stand at high galactic latitudes (i.e. $|b| \gtrsim 30°$) where the stellar density is low

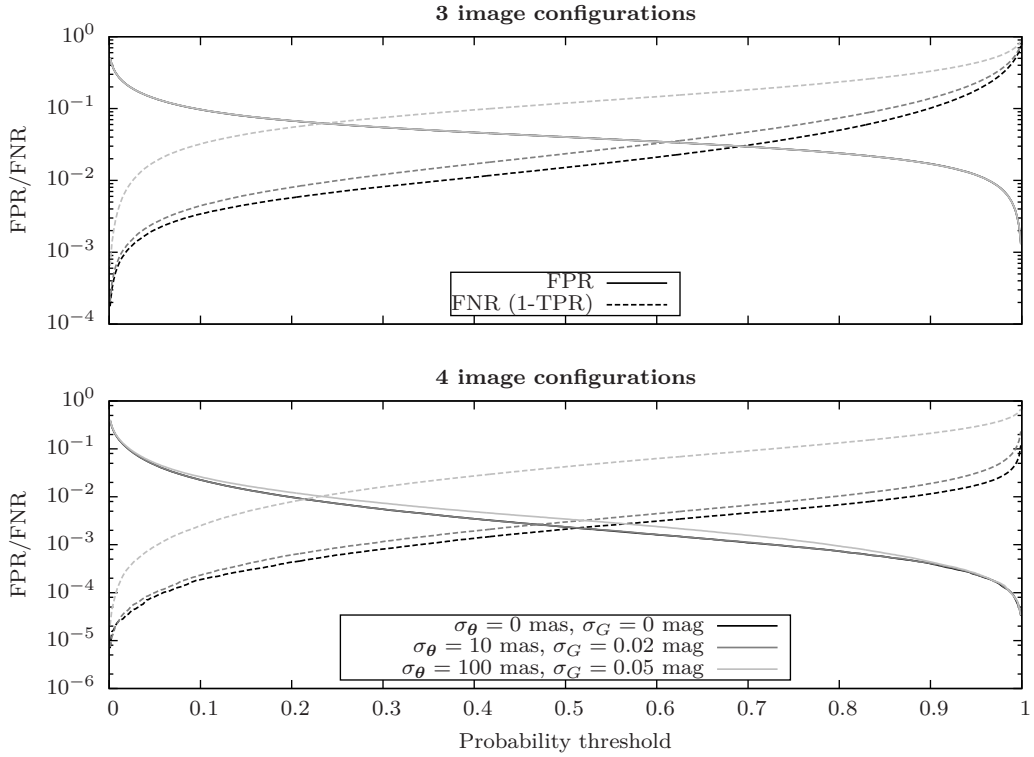

Figure 4.4: Performances of the Extremely randomized trees (ERT) on GL candidate identifications for the cases of three and four image configurations with varying noise. The false positive rate (FPR) corresponds here to the probability that a random cluster of stars will be considered as a GL candidate while the false negative rate (FNR) corresponds to the probability that a GL candidate will be missed (i.e. FNR = 1-TPR). Both these values depending on the threshold we set on the probability that is returned by the ERT so as to consider a given configuration as a GL candidate (see section 2.4).

Table 4.2: ERT performances on GL candidate characterization for $\sigma_{\boldsymbol{\theta}} = 10$ mas and $\sigma_G = 0.02$ mag. The source position is here given by $\boldsymbol{\theta_s} = (x_s, y_s)$ while the deflector position is given by $\boldsymbol{\theta_d} = (x_d, y_d)$, the overall orientation of the system being given by ϖ. Performances are assessed based on the mean absolute error (MAE) as well as on the correlation between the predicted parameters and the effective parameters of genuine GL candidates.

ERT	**3 image GL candidates**		**4 image GL candidates**	
Parameters	MAE	Correlation	MAE	Correlation
θ_e	0.235	0.9485	0.172	0.9739
f	0.057	0.8734	0.039	0.9378
ϖ	0.151	0.9746	0.092	0.9911
θ_c	0.042	0.6904	0.034	0.7827
x_s	0.040	0.9187	0.016	0.9901
y_s	0.055	0.9936	0.014	0.9996
x_d	0.042	0.7648	0.018	0.9613
y_d	0.062	0.9932	0.020	0.9992
γ	0.023	0.3100	0.017	0.7024
ω	0.533	0.7890	0.304	0.9259

applied to each source of the catalogue before it is made public. This aims to prevent variable objects to contaminate the catalogue as well as it removes sources for which the astrometric positions have large uncertainties. Accordingly, most of the known GLs come along with only one or two images and can hence not be identified by our procedure. Still 137 configurations composed of three images have a probability of being GL candidates that is higher than 90%, one of these even has a probability of 100%. A visual inspection procedure then selected 42 of these promising candidates for further observations using the Devasthal optical telescope (Sagar et al. 2011, hereafter DOT). Those observations are being planned for the beginning of 2018.

Finally, as the GL identification relies on the ability of the NSIEg model to reproduce the observed configurations, the determination of the parameters of this model should be equivalently doable. Table 4.2 accordingly lists the expected performances of the ERT in retrieving the parameters of the NSIEg lens model regarding GL candidates having a typical noise of $\sigma_{\boldsymbol{\theta}} = 10$ mas and $\sigma_G = 0.02$ mag. An important result appearing in this table is that both the source position $\boldsymbol{\theta_s}$ as well as the deflector position[2] $\boldsymbol{\theta_d}$ are determined with a relatively high accuracy. Also the axis ratios, f, have rather fair associated predictions. All these estimated parameters are intended to be used as starting conditions in a Bayesian modelling of the most promising candidates out of the Gaia and DOT observations. This point being the subject of an ongoing research.

[2]The deflecting galaxies often having smooth luminosity profiles, they will generally not be detected by Gaia

5 Conclusions

During this PhD thesis we have developed software solutions designed to solve the problems of the characterization of quasars and of the detection of gravitational lens candidates within the Gaia mission. The specific approach we have implemented regarding the former of these objectives (Delchambre 2018) relies on intuitive ideas that yield to the currently most accurate predictions of the astrophysical parameters of quasars out of the blue and red spectrophotometers of Gaia, hereafter abbreviated as BP and RP, respectively. The ease of interpretation that is associated with this implementation eventually provides strong diagnostic tools on the potential errors that may arise during predictions as, for example, a straight identification of the degenerated cases where different values of the redshift of a quasar can lead to the fits of the observed spectrum with comparable accuracies. Other possibilities offered by this technique is a fair recognition of the insecure predictions, upon which warning flags can be triggered, as well as the availability of straightly computable uncertainties for each of the estimated astrophysical parameters. A major challenge we had to face during the development of these quasar characterization methods stands in the restricted computer resources that are allocated to the processing of each quasar owing to the one billion of sources that Gaia has to manage in total. This requirement translates in a limited amount of 10^9 floating point operations (flops) that are allocated to the treatment of each of these sources, or equivalently to a processing time of 0.6 ms per source in the data processing center at CNES (DPCC).

Practically, the development of a dedicated weighted principal components analysis algorithm (Delchambre 2015), based on the diagonalization of the weighted covariance matrix, first enabled the creation of templates upon which the spectra of quasars are best modelled. This algorithm concurrently allowed quasars spectra from the Sloan digital sky survey (SDSS) to be extrapolated over wavelength ranges that are comparable to those covered by the red and blue spectrophotometers of Gaia with a view to their conversion into BP/RP spectra. The development of a weighted phase correlation algorithm (Delchambre 2016) designed to find the phase at which an input signal and a set of templates match at best in a weighted chi-squared sense subsequently allowed us to securely determine the redshift of the observed quasars. The specific implementation we elaborated is based on a highly modified orthogonal decomposition of the normal equations that are associated with the fit of the templates to the observation at each trial redshift. The resulting algorithm has a numerical

complexity of $\mathcal{O}(N \log N)$ flops that makes it affordable for the tight data processing of Gaia as the determination of the redshift of each source only requires $\sim 3 \times 10^6$ flops (assuming our template comes along with $N = 7709$ samples and that we use 6 of them). This equivalently translates into a processing time of ~ 2 μs per source in DPCC. The orthogonal decomposition we considered also ensures a high numerical stability of the whole algorithm. We detailed the building of a library of BP/RP spectra out of the SDSS observations and defined the appropriate processing that is necessary for these BP/RP spectra to be fully exploitable by the aforementioned algorithms. Finally, we properly derived the way the other astrophysical parameters were retrieved, amongst which: the slope of the quasar continuum, the equivalent width of their emission lines and whether these encompass broad absorption lines (BAL) or not.

Tests performed using a cross validation procedure show that 95.17%, 93.34%, 89.44%, 82.4% and 71.07% of the observations come along an absolute error on the predicted redshift, $|\Delta z|$, that is lower than 0.1 regarding BP/RP spectra of magnitude $G = \{18, 18.5, 19, 19.5, 20\}$, the associated mean absolute errors being then respectively given by 0.0057, 0.0061, 0.0069, 0.0088 and 0.013. Considering the removal of the predictions triggering warning flag(s), that encompass 46.2%, 50.99%, 57.72%, 67.09% and 77.58% of the total number of observations, we have that 99.77%, 99.72%, 99.57%, 99.36%, 98.58% of the remaining observations come along with $|\Delta z| < 0.1$ while the associated mean absolute errors become 0.0053, 0.0055, 0.0058, 0.0063, 0.0072 regarding the same set of magnitude. The accuracy in classifying BAL quasars being then respectively given by 91.725%, 91.169%, 90.562%, 90.198% and 90.069% along with a probability of detection of the BAL quasars of 31.992%, 21.295%, 9.14%, 3.348% and 0.413%. The slope of the quasar continua is however overestimated on average by an amount of $\Delta\alpha_\nu \approx 0.1$ owing to the spread of the Si IV and C IV emission lines over the continuum regions 145–148 nm, because of the convolution occurring in the BP/RP optics, and, to a lesser extent, from a similar spread of the C IV and C III] emission lines over the continuum region 170–180 nm. This overestimation consequently leads to an underestimation of the total equivalent width of the emission lines that is of the same order of magnitude.

The identification of gravitational lens (GL) candidates, on its side, had to rely on simulated configurations as the presently known sample of GLs is insufficient large in order to permit machine learning methods to efficiently guess the relation that may exist between the magnitudes and positions of some images and whether these are part of a GL system or not. We considered for this purpose a non-singular isothermal ellipsoid GL model in presence of an external shear (NSIEg) as it allows most of GLs whose deflector is a galaxy to be modelled. A learning set of observations was accordingly built based on a realistic distribution of the parameters of the NSIEg model as well as on fortuitous arrangements of stars coming from usual stellar population models. The determination of the lensed image positions out of a given GL model revealed a particular intricacy as it required the lens equation to be inverted. These positions were consequently obtained through a dedicated numerical method that only relies on straight evaluations of the latter equation. A supervised learning model, consisting of extremely randomized trees, was then trained on the available learning sets of observations composed of three and four image configurations and was subsequently tested using an independent test set of observations. Assuming we have typical positional and photometric uncertainties as well as given the theoretical fraction of GLs as compared to clusters of stars, we therefore estimated that 22.41% of the GLs composed of three images will be detected along with an associated probability for a cluster of stars to be identified as a GL of 0.129%. In case of four image

configurations, these numbers are respectively 99.94% and 0.009% such that these will be undoubtedly identified in the upcoming data releases of the Gaia catalogue. 42 promising GL candidates composed of three images were already identified out of the first Gaia data release and are currently awaiting confirmation through ground-based facilities.

APPENDIX

A

Numerical complexities

A.1 Weighted least-squares solution to a system of linear equations

This section aims to describe the fastest method designed to find the weighted least-squares solution to a system of linear equations and to properly derive its underlying numerical complexity. The method presented here potentially suffers from strong numerical instabilities (Press et al. 2002) and should not be used in practice, the latter being solely outlined in a comparison purpose regarding its low execution time.

Suppose we have a system of linear equations of the form

$$\mathbf{A}\boldsymbol{x} = \boldsymbol{y} \tag{A.1}$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\boldsymbol{x} \in \mathbb{R}^n$ and $\boldsymbol{y} \in \mathbb{R}^m$. The goal of the weighted least-squares solution to equation A.1 is to find a vector $\boldsymbol{x}$ such that

$$\chi^2 = \sum_i \frac{1}{\sigma_i^2} \left(y_i - \sum_j \mathbf{A}_{ij} x_j \right)^2 \tag{A.2}$$

is minimized. Note that we introduced in the latter equation $\boldsymbol{\sigma}$, the vector of the uncertainties that are associated with $\boldsymbol{y}$. Taking the partial derivatives of equation A.2 with respect to $\boldsymbol{x}$ and rearranging yields

$$\mathbf{A}^{\mathrm{T}}\mathbf{W}\mathbf{A}\boldsymbol{x} = \mathbf{A}^{\mathrm{T}}\mathbf{W}\boldsymbol{y}, \qquad \mathbf{W}_{ij} = \begin{cases} 1/\sigma_i^2, & \text{if } i = j \\ 0, & \text{otherwise} \end{cases}, \tag{A.3}$$

that are the *normal equations* associated with the least-squares minimization of equation A.1.

From basic considerations, we have that the construction of the normal equations matrix $\mathbf{B} = \mathbf{A}^{\mathrm{T}}\mathbf{W}\mathbf{A}$ requires $n^2m + mn$ flops as the matrix multiplication $\mathbf{A'} = \mathbf{W}\mathbf{A}$ consumes nm flops whereas the matrix multiplication $\mathbf{A}^{\mathrm{T}}\mathbf{A'}$ takes n^2m flops. The vector $\boldsymbol{z} = \mathbf{A}^{\mathrm{T}}\mathbf{W}\boldsymbol{y}$ is then requiring $nm + m$ flops based on the same principles. The overall numerical complexity that is needed for the construction of the normal equations being hence given by

$$t_{\mathrm{norm}} = m(n+1)^2 \tag{A.4}$$

flops.

Algorithm 1 Cholesky decomposition. Decompose a symmetric positive-definite matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$ into a lower triangular matrix $\mathbf{L}$ and an upper triangular matrix $\mathbf{L}^{\mathrm{T}}$ such that $\mathbf{A} = \mathbf{L}\mathbf{L}^{\mathrm{T}}$. The $\mathbf{L}$ matrix being stored in the lower triangular part of $\mathbf{A}$ upon return.

```
1:  for i = 1 to n do
2:      for j = 1 to i do
3:          S ← A_ij
4:          for k = 1 to j − 1 do
5:              S ← S − A_ik A_jk
6:          end for
7:          if i = j then
8:              A_ij ← √S
9:          else
10:             A_ij ← S/A_jj
11:         end if
12:     end for
13: end for
```

Assuming that $\mathbf{B}$ is positive-definite, then it can be decomposed through algorithm 1 as

$$\mathbf{B} = \mathbf{L}\mathbf{L}^{\mathrm{T}} \tag{A.5}$$

where $\mathbf{L}$ is a lower triangular matrix. The numerical complexity of algorithm 1 can be derived if we recognize the inner loop (see line 4) as being composed of $j-1$ flops[1] while the last step (either line 8 or 10) requires 1 flop. The loop at line 2 is executed i times, giving a complexity of $\sum_{j=1}^{i} j = (i^2+i)/2$ flops. Finally, the outer loop (line 1) is executed n times giving a final algorithmic complexity of

$$t_{\mathrm{chol}} = \sum_{i=1}^{n} \frac{i^2+i}{2} = \frac{n(n+1)(n+2)}{6} \tag{A.6}$$

flops.

Algorithm 2 Forward substitution. Given a lower triangular matrix $\mathbf{L} \in \mathbb{R}^{n\times n}$ and an image vector $\boldsymbol{y} \in \mathbb{R}^n$, the algorithm finds the vector $\boldsymbol{x}$ such that $\mathbf{L}\boldsymbol{x} = \boldsymbol{y}$.

```
1: for i = 1 to n do
2:     x_i ← y_i
3:     for j = 1 to i − 1 do
4:         x_i ← x_i − L_ij x_j
5:     end for
6:     x_i ← x_i/L_ii
7: end for
```

The vector $\boldsymbol{x}$ can then be retrieved through the consecutive solutions of the systems of linear equations

$$\mathbf{L}\boldsymbol{b} = \boldsymbol{y}, \qquad \mathbf{L}^{\mathrm{T}}\boldsymbol{x} = \boldsymbol{b}. \tag{A.7}$$

[1] We assume here that any operation of the form $x \leftarrow x \pm ab$ can be executed in one flop, as implemented in all modern CPUs.

that are respectively called *forward substitution* and *backward substitution*. The forward substitution algorithm (algorithm 2) consists in $i-1$ operations from the innermost loop while the last step is a single division. This gives a total algorithmic complexity of

$$t_{\text{fwd}} = \sum_{i=1}^{n} i = \frac{n^2+n}{2} \tag{A.8}$$

flops. Note that the backward substitution algorithm (with an upper triangular matrix) has trivially the same complexity, (i.e. $t_{\text{bwd}} = t_{\text{fwd}}$).

The overall numerical complexity of the weighted least-squares solution to equation A.1 being hence given by

$$t_{\text{lsqr}} = t_{\text{norm}} + t_{\text{chol}} + t_{\text{fwd}} + t_{\text{bwd}} = m(n+1)^2 + \frac{n(n+1)(n+8)}{6}. \tag{A.9}$$

Since we often have $n \ll m$, the first term of equation A.9 will be dominant (i.e. most of the CPU resources being consumed in the building of the normal equation matrix) such that the numerical complexity of the algorithm is given in general by $\mathcal{O}\left(mn^2\right)$ flops.

A.2 Fast Fourier transform

As a foreword to the analysis of this fast Fourier transform (FFT) algorithm, we have to note that many FFT algorithms actually exist. Each of these algorithms is optimized for a given input vector's size and/or for some data pattern. Consequently, a precise analysis–as previously done in this section–would not be possible as it would require the exploration of the huge variety of algorithms as well as it will necessitate a very thorough mathematical background in order to understand these. Indeed, FFT is still a very active topic in the research communities and their implementations are rarely simple (see Frigo & Johnson 2005; Johnson & Frigo 2007, for examples). Rather, this section aims at giving the reader the insights that the FFT is effectively an algorithm with a complexity of $\mathcal{O}\left(N \log N\right)$.

In algorithm 3, the loop at line 5 is used to split the input vector $\boldsymbol{x}$ into an 'even' and an 'odd' part, each of them being recursilvely processed respectively by line 9 and 10. The sorting appears explicitly here in a didactical puprose, but practically it is implemented using an appropriate indexing of $\boldsymbol{x}$ and is thus costless. Suppose now that $t_{\text{fft}}(N)$ is the complexity of the present algorithm, then lines 9 and 10 will have a complexity of $2t_{\text{fft}}(\frac{N}{2})$ flops to which we have to add the complexity of the loop at line 11. W factors (line 12) are almost always pre-computed in all modern implementation and are consequently costless, while lines 13 and 14 requires 2 flops to be executed giving the loop at line 11 a complexity of N. We will then have a total complexity of

$$\begin{aligned} t_{\text{fft}}(N) &= N + 2t_{\text{fft}}\left(\frac{N}{2}\right) \\ &= 2N + 4t_{\text{fft}}\left(\frac{N}{4}\right) \\ &\vdots \\ &= kN + 2^k t_{\text{fft}}\left(\frac{N}{2^k}\right). \end{aligned} \tag{A.10}$$

Algorithm 3 Find the discrete Fourier Transform (DFT), $\boldsymbol{y} \in \mathbb{C}^N$, of the input vector $\boldsymbol{x} \in \mathbb{C}^N$ with $N = 2^n; n \in \mathbb{Z}^+$ using a recursive Cooley-Tukey algorithm.

1: **function** fft($\boldsymbol{x}$, N)
2: **if** $N = 1$ **then**
3: $\quad \boldsymbol{y} \leftarrow \boldsymbol{x}$
4: **else**
5: $\quad$ **for** $j = 1$ to $\frac{N}{2}$ **do**
6: $\quad\quad \boldsymbol{e} \leftarrow (\boldsymbol{e}\ x_{2j-1})$
7: $\quad\quad \boldsymbol{o} \leftarrow (\boldsymbol{o}\ x_{2j})$
8: $\quad$ **end for**
9: $\quad \hat{\boldsymbol{e}} \leftarrow$fft$(\boldsymbol{e}, \frac{N}{2})$
10: $\quad \hat{\boldsymbol{o}} \leftarrow$fft$(\boldsymbol{o}, \frac{N}{2})$
11: $\quad$ **for** $j = 1$ to $\frac{N}{2}$ **do**
12: $\quad\quad W \leftarrow \exp(-2i\pi(j-1)/N)$
13: $\quad\quad y_j \leftarrow \hat{e}_j + W\hat{o}_j$
14: $\quad\quad y_{j+\frac{N}{2}} \leftarrow \hat{e}_j - W\hat{o}_j$
15: $\quad$ **end for**
16: **end if**
17: **end function**

The basic case will occur when $\frac{N}{2^k} = 1$, that is when $k = \log_2 N$, replacing within Eq. A.10 gives the non recursive complexity as

$$t_{\text{fft}}(N) = N\log_2 N + N = N\log_2(2N) \tag{A.11}$$

flops.

That being said, one often needs to compute the FFT of real-valued data. In this case, by taking advantage of the symmetry $F_i = F^*_{N-i}$, we can reach a complexity twice better that the one of the presented algorithm. Similarly, passing from a Fourier Transform having this kind of symmetry to its (real) inverse can be performed in half the time needed by the classical algorithm.

A.3 Weighted phase correlation through factorized QR decomposition and lookup tables

The weighted phase correlation algorithm we detail here is first introduced in section 2.3 as well as it is extensively covered in Paper II. Accordingly we refer the reader to these specific sections regarding the implementation details we straightly provide here. The solution we present in algorithm 4 is used so as to rigorously derive its numerical complexity, it is however fully operational and can be straightly translated in any computer language.

Out of algorithm 4, we have that lines 2, 3, 6 and 8 all correspond to (inverse) FFT of real signals (see section A.2). These hence require m flops in order to perform the vector element-wise multiplications (ie. computing $\boldsymbol{w}^2$, $\boldsymbol{w}^2 \circ \boldsymbol{s}$, $\hat{\boldsymbol{w}}^* \circ \hat{\mathbf{T}}_{\mathbf{ij}}$ and $\hat{\boldsymbol{s}}^* \circ \hat{\boldsymbol{t}}_{\boldsymbol{i}}$) plus $\frac{1}{2}t_{\text{fft}}(m)$ in order to switch between the time and the frequency domain, leading to a complexity of $\frac{m}{2}\log_2(2m) + m = \frac{m}{2}\log_2(8m)$ flops for each of these lines. Furthermore, we have that line 6 is executed $(n + n^2)/2$ times and line 8, n times. The building of the lookup tables (from line 2 to line 9) hence requires

$$t_{\text{lup}}(m, n) = \frac{m}{4}\log_2(8m)(n^2 + 3n + 4) \tag{A.12}$$

flops.

For a given shift k, the line 23 is executed $\sum_{i=1}^{n} i(i-1) = (n^3-n)/3$ times in order to update the matrix $\mathbf{X}$ based on the $(n^2-n)/2$ values of γ computed at line 21. Similarly, the line 29 is executed $(n^2+n)/2$ times so as to update $\boldsymbol{y}$. Lines 16, 19 and 27 being finally executed n times while lines 13 and 14 respectively require n^2 and n flops so as to initialize $\mathbf{X}$ and $\boldsymbol{y}$. The construction and update of the $\mathbf{X}$ matrix and of the $\boldsymbol{y}$ vector (lines from 13 to 30) hence require

$$t_{\mathrm{xy}}(n) = \frac{n+2}{6}(2n^2+8n+3) \tag{A.13}$$

flops. Now, the update of the lookup table takes $(n^2-n)/2$ flops regarding the update of $\tilde{\boldsymbol{l}}$ as well as $\sum_{i=1}^{n-1}(i^2-i)/2 = (n-2)(n-1)n/6$ flops regarding the update of $\tilde{\boldsymbol{L}}$, yielding an overall complexity for these updates of

$$t_{\mathrm{up}}(n) = \frac{n(n^2-1)}{6} \tag{A.14}$$

flops. The computation of each $\mathrm{ccf}(k)$ (line 40) is finally requiring n flops.

Given that the algorithm is executed for all $k \in \{0, m-1\}$ between lines 13 and 40, we have a final complexity for the WPC algorithm of

$$\begin{aligned} t_{\mathrm{wpc}}(n,m) &= t_{\mathrm{lup}}(m,n) + m\left[t_{\mathrm{xy}}(n) + t_{\mathrm{up}}(n) + n\right] \\ &= t_{\mathrm{lup}}(m,n) + \frac{m}{2}\left(n^3+4n^2+8n+2\right) \end{aligned} \tag{A.15}$$

flops. If we assume that $n \ll m$, then most of the time is spent in building the lookup tables such that the numerical complexity of the whole algorithm is approximately given by $m\log_2(8m)n^2$ flops.

Algorithm 4 Compute the weighted cross correlation function (CCF) of a given input vector $\boldsymbol{s} \in \mathbb{R}^m$ having weights $\boldsymbol{w} \in \mathbb{R}^m$ against a set of templates $\mathbf{T} \in \mathbb{R}^{m \times n}$ where $\mathbf{T}_{(i+m)j} = \mathbf{T}_{ij}$. The returned CCF, $\mathrm{ccf}(k)$, is such that the optimal $\chi^2(k)$ resulting from the fit of $\boldsymbol{s}$ with $\tilde{\mathbf{T}}$ ($\tilde{\mathbf{T}}_{ij} \equiv \mathbf{T}_{(i+k)j}$) using weight $\boldsymbol{w}$ satisfies $\chi^2(k) = \|\boldsymbol{w} \circ \boldsymbol{s}\|^2 - \mathrm{ccf}(k)$. We assume here that the matrices $\hat{\mathbf{T}}_{\mathbf{ij}} = \mathrm{fft}(\mathbf{T}_i^{\mathrm{col}} \circ \mathbf{T}_j^{\mathrm{col}}, m)$ and the vectors $\hat{\boldsymbol{t}}_{\boldsymbol{i}} = \mathrm{fft}(\mathbf{T}_i^{\mathrm{col}}, m)$ are computed in advance.

— Build the lookup tables $\tilde{\mathbf{L}}_{\mathbf{ij}}$ and $\tilde{\boldsymbol{l}}_{\boldsymbol{i}}$ —
$\hat{\boldsymbol{w}} \leftarrow \mathrm{fft}(\boldsymbol{w}^2, m)$
$\hat{\boldsymbol{s}} \leftarrow \mathrm{fft}(\boldsymbol{w}^2 \circ \boldsymbol{s}, m)$
for $i = 1$ to n **do**
  for $j = i$ to n **do**
    $\tilde{\mathbf{L}}_{\mathbf{ij}} \leftarrow \mathrm{ifft}(\hat{\boldsymbol{w}}^* \circ \hat{\mathbf{T}}_{\mathbf{ij}}, m)$ {i.e. $\tilde{\mathbf{L}}_{\mathbf{ij}}(k) = \sum_l w_l^2 \mathbf{T}_{(l+k)i} \mathbf{T}_{(l+k)j}$}
  end for
  $\tilde{\boldsymbol{l}}_{\boldsymbol{i}} \leftarrow \mathrm{ifft}(\hat{\boldsymbol{s}}^* \circ \hat{\boldsymbol{t}}_{\boldsymbol{i}}, m)$ {i.e. $\tilde{\boldsymbol{l}}_{\boldsymbol{i}}(k) = \sum_j w_j^2 s_j \mathbf{T}_{(j+k)i}$}
end for

— Evaluate the cross correlation function in each shift k —
for $k = 0$ to $m - 1$ **do**
  $\mathbf{X}_{ij} \leftarrow w_i \mathbf{T}_{(i+k)j} \quad \forall i, j \in \{1, n\}$
  $y_i \leftarrow w_i s_i \quad \forall i \in \{1, n\}$
  for $i = 1$ to n **do**
    $\alpha \leftarrow \mathrm{sgn}(\mathbf{X}_{ii}) \sqrt{\tilde{\mathbf{L}}_{\mathbf{ii}}(k)}$
    — Perform a Householder reflection so as to annihilate the elements —
    — below the ith row of the ith column of $\mathbf{X}$ —
    $\mathbf{X}_{ii} \leftarrow \mathbf{X}_{ii} + \alpha$
    for $j = i + 1$ to n **do**
      $\gamma \leftarrow \left(\mathbf{X}_{ij} + \tilde{\mathbf{L}}_{\mathbf{ij}}(k)/\alpha\right) / \mathbf{X}_{ii}$
      for $l = i$ to n **do**
        $\mathbf{X}_{lj} \leftarrow \mathbf{X}_{lj} - \gamma \mathbf{X}_{li}$
      end for
    end for
    — Perform a similar reflection on $\boldsymbol{y}$ —
    $\gamma \leftarrow \left(y_i + \tilde{\boldsymbol{l}}_{\boldsymbol{i}}(k)/\alpha\right) / \mathbf{X}_{ii}$
    for $j = i$ to n **do**
      $y_j \leftarrow y_j - \gamma \mathbf{X}_{ji}$
    end for
    — Update the lookup tables —
    for $j = i + 1$ to n **do**
      for $l = j$ to n **do**
        $\tilde{\mathbf{L}}_{\mathbf{jl}}(k) \leftarrow \tilde{\mathbf{L}}_{\mathbf{jl}}(k) - \mathbf{X}_{ij} \mathbf{X}_{il}$
      end for
      $\tilde{\boldsymbol{l}}_{\boldsymbol{j}}(k) \leftarrow \tilde{\boldsymbol{l}}_{\boldsymbol{j}}(k) - \mathbf{X}_{ij} y_i$
    end for
  end for
  — Compute the cross correlation function at shift k —
  $\mathrm{ccf}(k) \leftarrow \sum_{i=1}^n y_i^2$
end for

List of Figures

1.1 Gaia optical bench and mirrors . . . 6
1.2 The focal plane of Gaia . . . 6
1.3 Prisms of the blue and red spectro-photometers of Gaia . . . 9
1.4 Optical module of the radial velocity spectrometer of Gaia . . . 10
1.5 Multi-wavelength observations of the 3C273 quasar . . . 15
1.6 Radio jets observations in AGN . . . 15
1.7 Mid-infrared observations of Seyfert I/II galaxies . . . 17
1.8 Different types of AGN as seen in the optical band . . . 19
1.9 Host galaxies of AGN viewed by the HST. . . . 20
1.10 Mean X-ray spectrum of a type I AGN . . . 22
1.11 Illustration of the accretion disk and hot corona of AGN. . . . 22
1.12 Multi-wavelength variability of the 3C273 quasar . . . 24
1.13 Power density spectra of the 3C273 quasar . . . 25
1.14 A unified model for AGN . . . 27
1.15 Apparent superluminal motion in the jet of M87 . . . 28
1.16 Proper distances as a function of redshift . . . 37
1.17 Lyman absorbers in a high redshift quasar . . . 39
1.18 Geometry of the gravitational lensing phenomenon . . . 42
1.19 Gravitationally lensed image projected on the sky . . . 44
1.20 Illustration of the NSIEg lens model parameters . . . 47

2.1 Examples of simulated BP/RP spectra of quasars having various G magnitudes . . . 53
2.2 Illustration of the equivalent width of an emission line . . . 55
2.3 χ^2 curve of the SDSS J000313.08+274044.9 quasar . . . 59
2.4 Rest-frame wavelength coverage of quasar spectra according to redshift 60
2.5 Paper I: a two-dimensional PCA example . . . 64
2.6 Paper I: weighted and classical PCA in two dimensions . . . 67
2.7 Paper I: simulated data for the comparison of WPCA methods . . . 68
2.8 Paper I: averaged χ^2_{fit} of WPCA methods based on simulated data . . . 69
2.9 Paper I: averaged χ^2_{test} of WPCA methods based on simulated data . . . 69
2.10 Paper I: averaged χ^2_{test} of WPCA methods in presence of missing simulated data . . . 70
2.11 Paper I: SDSS DR10Q spectrum fit with WPCA methods . . . 70
2.12 Paper I: observation's χ^2_{test} out of the SDSS DR10Q dataset . . . 70
2.13 Paper I: components retrieved with and without missing data out of the SDSS DR10Q dataset . . . 71
2.14 Paper II: mean observation and first principal components of SDSS quasar spectra at $2.3 \leq z \leq 2.4$ . . . 77

2.15 Paper II: examples of cross-correlation functions 78
2.16 Paper II: result of the initial processing of SDSS spectra 81
2.17 Paper II: results of the phase correlation and WPC algorithms on SDSS spectra 82
2.18 Paper II: histograms of the SNR distribution of redshift according to various phase correlation methods 83
2.19 Example of a binary tree node 86
2.20 Modelling a noisy sine function using regression trees and ERT 88

3.1 Paper III: results of the pre-proceessing and extrapolation procedure on a DR12Q spectrum 96
3.2 Paper III: principal components out of DR12Q emission lines and continuum 97
3.3 Paper III: example of simulated BP/RP spectrum 98
3.4 Paper III: illustration of a synthesized spectrum 99
3.5 Paper III: principal components out of synthesized spectra 100
3.6 Paper III: histogram of the errors of predicted continuum slopes and total emission lines equivalent widths 101
3.7 Paper III: mean quasar spectrum from synthesized spectra and DR12 spectra 102
3.8 Paper III: distribution of the predicted redshift from WPC and ERT . 103
3.9 Paper III: impurity and completeness of observations according to Z_{score} and χ^2_r 105
3.10 Paper III: ROC curves of BAL binary classification 106
3.11 Paper III: comparison of the TPR of BAL classification with the balnicity index 106

4.1 Probability density function of the axis ratio of galaxies 110
4.2 Caustic curves and square deviation function of a GL model 111
4.3 Examples of GL candidates and contaminants out of TS 113
4.4 ERT performances on GL candidates identification 115

List of Tables

1.1 Coordination Units from DPAC. 10
1.2 Dominant emission lines observed in the optical spectra of AGN. . . . 18
1.3 Parameters of the NSIEg lens model 48

2.1 Modules constituting the astrophysical parameter inference system . . 52
2.2 Paper I: $\chi^2_{\rm fit}$ and $\chi^2_{\rm test}$ statistics out of the SDSS DR10Q dataset 70
2.3 Confusion matrix for classification problems 88

3.1 Paper III: emission lines used in the computation of the $Z_{\rm score}$ 95
3.2 Paper III: warning flags used in the redshift prediction procedure . . . 100
3.3 Paper III: mean predicted continuum slopes and total emission lines equivalent widths 101
3.4 Paper III: ratio of observations triggering warning flags 104
3.5 Paper III: optimal accuracies of the BAL classification 106

4.1 ERT performances on GL candidates identification 114
4.2 ERT performances on GL candidate characterization 116

List of Abbreviations

3C	Third Cambridge Catalogue of Radio Sources
3CR	Revision of the 3C catalogue
AC	ACross scan
ADS	Airbus Defence and Space
AF	Astrometric Field
AGIS	Astrometric Global Iterative Solution
AGN	Active Galactic Nucleus/Nuclei
AL	ALong scan
AP	Astrophysical Parameter
APSIS	Astrophysical Parameter Inference System
BAL	Broad Absorption Line
BAM	Basic Angle Monitor
BH	Black Hole
BLR	Broad Line Region
BOSS	Baryon Oscillation Spectroscopic Survey
BP	Blue Photometer
CART	Classification And Regression Trees
CASSIS	Combined Atlas of Sources with Spitzer IRS Spectra
CASTLES	CfA-Arizona Space Telescope LEns Survey of GLs
CCD	Charged Coupled Device
CCF	Cross Correlation Function
CNES	Centre National d'Études Spatiales
CPU	Central Processing Unit
CU	Coordination Unit
DFT	Discrete Fourier Transform
DLA	Damped Lyα Absorber
DOT	Devasthal Optical Telescope
DPAC	Data Processing and Analysis Consortium
DPACE	DPAC Executive
DPC	Data Processing Centre
DPCC	DPC at CNES
DSC	Discrete Source Classifier
EBL	Extragalactic Background Light
EO	Extended Object
ERT	Extremely Randomized Trees
ESA	European Space Agency
EW	Equivalent Width
FFT	Fast Fourier Transform
FIR	Far-infrared

FNR	False Negative Rate
FOV	Field Of View
FPR	False Positive Rate
FWHM	Full Width at Half Maximum
GASS	GAia System-level Simulator
GCRF	Gaia Celestial Reference Frame
GDR	Gaia Data Release
GIBIS	Gaia Instrument and Basic Image Simulator
GL	Gravitational Lens(ing)/ Gravitational Lens system
GOG	Gaia Object Generator
GraL	GRAvitational Lenses in Gaia
HR	Hertzsprung-Russell
HST	Hubble Space Telescope
ICRF	International Celestial Reference Frame
IDT	Initial Data Treatment
IDU	Intermediate Data Update
IGM	Inter-Galactic Medium
IR	Infra-Red
LF	Luminosity Function
LINER	Low Ionization Narrow-Line Emission Radio galaxies
LQAC	Large Quasar Astrometric Catalog
LS	Learning Set of observations
LSF	Line Spread Function
MAE	Mean Absolute Error
MCMC	Monte-Carlo Markov Chain
MDB	Main DataBase
MIR	Mid-infrared
ML	Machine Learning
NAL	Narrow Absorption Line
NIR	Near-infrared
NLR	Narrow Line Region
NSIE	Non-Singular Isothermal Ellipsoid
NSIEg	NSIE lens model in presence of an external shear
NSIS	Non-Singular Isothermal Sphere
NSS	Non-Single-Star
OVV	Optically Violent Variable
PAH	Polycyclic Aromatic Hydrocarbon
PCA	Principal Components Analysis
PSD	Power Spectral Density
PSF	Point Spread Function
QSO	Quasi-Stellar Object
QSOC	QSO Classifier
ROC	Receiver Operating Characteristic
RP	Red Photometer
RVS	Radial Velocity Spectrometer
SDF	Square Deviation Function
SDSS	Sloan Digital Sky Survey
SED	Spectral Energy Distribution
SIE	Singular Isothermal Ellipsoid
SIS	Singular Isothermal Sphere
SM	Sky Mapper

SMBH	Super Massive Black Hole
SNR	Signal-to-Noise Ratio
SQLS	SDSS Quasar Lens Search
SRS	Software Requirement Specifications
SSO	Solar System Objects
SSS	Software and System Specification
SVD	Singular Value Decomposition
TDI	Time Delay Integration
TPR	True Positive Rate
TS	Test Set of observations
UV	Ultra-Violet
VLA	Very Large Array
WFS	Wave-Front Sensor
WHIM	Warm Highly Ionized Medium
WP	Work Package
WPC	Weighted Phase Correlation
WPCA	Weighted Principal Components Analysis

List of symbols and notations

Symbol	Description	Unit	Comments
z	Redshift	–	
λ	Wavelength	m	
ν	Frequency	s^{-1}	
l	Galactic longitude	degree	
b	Galactic latitude	degree	
L	Absolute luminosity	erg m^{-2} s^{-1}	
E	Energy	eV	
T	Temperature	K	
$\dot{x}$	Time derivative of x	unit of x s^{-1}	$\dot{x} = \mathrm{d}x/\mathrm{d}t$
$\ddot{x}$	Time curvature of x	unit of x s^{-2}	$\ddot{x} = \mathrm{d}^2x/\mathrm{d}t^2$
$\boldsymbol{x}$	Vector $\boldsymbol{x}$	–	Element i is given by x_i
$\mathbf{X}$	Matrix $\mathbf{X}$	–	Element at row i, column j is given by $\mathbf{X}_{ij}$
$\nabla\boldsymbol{x}$	Gradient of $\boldsymbol{x}$	–	
$\nabla^2\boldsymbol{x}$	Laplacian of $\boldsymbol{x}$	–	
$\|\boldsymbol{x}\|$	Norm[2] of $\boldsymbol{x}$	–	.
$\circ$	Element-wise product	–	e.g. $\boldsymbol{x} \circ \boldsymbol{y}$
$\mathbb{R}$	The set of real numbers	–	$\mathbb{R}_0$: real numbers except 0
$\mathbb{C}$	The set of complex numbers	–	
$\mathbb{Z}$	The set of integer numbers	–	
$\overset{\chi^2}{=}$	Nearest in a least-squares sense	–	e.g. $\boldsymbol{x} \overset{\chi^2}{=} \boldsymbol{y}$

[2] If not stated otherwise, Euclidean norm is considered (i.e. $|\boldsymbol{x}|^2 = \sum_i x_i^2$)

List of constants and units

Symbol	Description	Value
c	Speed of light in vacuum	2.99792458×10^{8} m s^{-1}
G	Gravitational constant	6.67384×10^{-11} m^{3} kg^{-1} s^{-2}
h	Planck constant	$6.6260696 \times 10^{-34}$ J s
k_{B}	Boltzmann constant	$1.3806488 \times 10^{-23}$ J K^{-1}
eV	Electron-volt	$1.6021766 \times 10^{-19}$ J
$M_{\odot}$	Mass of the sun	1.9891×10^{30} kg
$L_{\odot}$	Luminosity of the sun	3.845×10^{26} J s^{-1}
$R_{\odot}$	Solar radius	6.957×10^{8} m
H_0	Current value of the Hubble parameter (i.e. Hubble constant)	~ 67.7 km s^{-1} Mpc^{-1}
pc	Parallax second (parsec)	3.0856776×10^{16} m
ly	Light year	$9.46073047 \times 10^{15}$ m
$^{\circ}$	Degree	$\pi/180$
$'$	Arcminute	$1/60\,^{\circ}$
$''$	Arcsecond	$1/60\,'$
mas	Milli-arcsecond	$1/1000\,''$
μas	Micro-Arcsecond	1/1000 mas

Bibliography

Alam, S., **Albareti**, F. D., **Allende Prieto**, C. **et al.**, *The Eleventh and Twelfth Data Releases of the Sloan Digital Sky Survey: Final Data from SDSS-III.* ApJS, 219, (2015), 12. doi:10.1088/0067-0049/219/1/12.

Antonucci, R., *Unified models for active galactic nuclei and quasars.* ARA&A, 31, (1993), 473. doi:10.1146/annurev.aa.31.090193.002353.

Antonucci, R. R. J. & **Miller**, J. S., *Spectropolarimetry and the nature of NGC 1068.* ApJ, 297, (1985), 621. doi:10.1086/163559.

Bailer-Jones, C. A. L., **Andrae**, R., **Arcay**, B. **et al.**, *The Gaia astrophysical parameters inference system (Apsis). Pre-launch description.* A&A, 559, (2013), A74. doi:10.1051/0004-6361/201322344.

Bailer-Jones, C. A. L., **Andrae**, R., **Astraatmadja**, T. L. **et al.** (2016). *CU8 Software Requirements Specification for APSIS.* Tech. Rep. GAIA-C8-SP-MPIA-CBJ-032, Max-Planck-Institute for Astronomy, Heidelberg. URL `http://www.rssd.esa.int/cs/livelink/open/3072504`. Access restricted to DPAC members only.

Bailey, S., *Principal Component Analysis with Noisy and/or Missing Data.* PASP, 124, (2012), 1015. doi:10.1086/668105.

Baldwin, J. A., *Luminosity Indicators in the Spectra of Quasi-Stellar Objects.* ApJ, 214, (1977), 679. doi:10.1086/155294.

Becker, R. H., **Fan**, X., **White**, R. L. **et al.**, *Evidence for Reionization at z˜6: Detection of a Gunn-Peterson Trough in a z=6.28 Quasar.* AJ, 122, (2001), 2850. doi:10.1086/324231.

Becker, R. H., **White**, R. L., **Gregg**, M. D. **et al.**, *Properties of Radio-selected Broad Absorption Line Quasars from the First Bright Quasar Survey.* ApJ, 538, (2000), 72. doi:10.1086/309099.

Beckmann, V. & **Shrader**, C. R. (2012). *Active Galactic Nuclei.* Wiley-VCH Verlag GmbH & Co. ISBN 9783527410781.

Beckmann, V., **Soldi**, S., **Ricci**, C. **et al.**, *The second INTEGRAL AGN catalogue.* A&A, 505, (2009), 417. doi:10.1051/0004-6361/200912111.

Bennett, A. S., *The revised 3C catalogue of radio sources.* MmRAS, 68, (1962), 163.

Berens, P., *Circstat: A matlab toolbox for circular statistics.* Journal of Statistical Software, Articles, 31(10), (2009), 1. ISSN 1548-7660. doi:10.18637/jss.v031.i10.

Biretta, J. A., **Sparks**, W. B. & **Macchetto**, F., *Hubble space telescope observations of superluminal motion in the m87 jet.* The Astrophysical Journal, 520(2), (1999), 621.

Bishop, C. M. (2006). *Pattern recognition and machine learning.* Springer. ISBN 9780387310732.

Blandford, R. D. & **Königl**, A., *Relativistic jets as compact radio sources.* ApJ, 232, (1979), 34. doi:10.1086/157262.

Blandford, R. D. & **Znajek**, R. L., *Electromagnetic extraction of energy from Kerr black holes.* MNRAS, 179, (1977), 433. doi:10.1093/mnras/179.3.433.

Bolton, A. S., **Schlegel**, D. J., **Aubourg**, É. **et al.**, *Spectral Classification and Redshift Measurement for the SDSS-III Baryon Oscillation Spectroscopic Survey.* AJ, 144, (2012), 144. doi:10.1088/0004-6256/144/5/144.

Brault, J. W. & **White**, O. R., *The Analysis and Restoration of Astronomical Data via the Fast Fourier Transform.* A&A, 13, (1971), 169.

Burke, W. L., *Multiple Gravitational Imaging by Distributed Masses.* ApJ, 244, (1981), L1. doi:10.1086/183466.

Cao, X., *An accretion disc-corona model for X-ray spectra of active galactic nuclei.* MNRAS, 394, (2009), 207. doi:10.1111/j.1365-2966.2008.14347.x.

Choi, Y.-Y., **Park**, C. & **Vogeley**, M. S., *Internal and Collective Properties of Galaxies in the Sloan Digital Sky Survey.* ApJ, 658, (2007), 884. doi:10.1086/511060.

Claeskens, J.-F., **Smette**, A., **Vandenbulcke**, L. **et al.**, *Identification and redshift determination of quasi-stellar objects with medium-band photometry: application to Gaia.* MNRAS, 367, (2006), 879. doi:10.1111/j.1365-2966.2006.10024.x.

Croom, S. M., **Richards**, G. T., **Shanks**, T. **et al.**, *The 2dF-SDSS LRG and QSO survey: the QSO luminosity function at 0.4 < z < 2.6.* MNRAS, 399, (2009), 1755. doi:10.1111/j.1365-2966.2009.15398.x.

Cui, X.-Q., **Zhao**, Y.-H., **Chu**, Y.-Q. **et al.**, *The Large Sky Area Multi-Object Fiber Spectroscopic Telescope (LAMOST).* Research in Astronomy and Astrophysics, 12, (2012), 1197. doi:10.1088/1674-4527/12/9/003.

Delchambre, L., *Weighted principal component analysis: a weighted covariance eigendecomposition approach.* MNRAS, 446, (2015), 3545. doi:10.1093/mnras/stu2219.

Delchambre, L., *Redshift determination through weighted phase correlation: a linearithmic implementation.* MNRAS, 460, (2016), 2811. doi:10.1093/mnras/stw1025.

Delchambre, L., *Determination of astrophysical parameters of quasars within the gaia mission.* MNRAS, 473(2), (2018), 1785. doi:10.1093/mnras/stx2417.

Delubac, T., **Bautista**, J. E., **Busca**, N. G. **et al.**, *Baryon acoustic oscillations in the Lyα forest of BOSS DR11 quasars.* A&A, 574, A59. doi:10.1051/0004-6361/201423969.

Dicken, D., **Tadhunter**, C., **Axon**, D. **et al.**, *The Origin of the Infrared Emission in Radio Galaxies. II. Analysis of Mid- to Far-Infrared Spitzer Observations of the 2Jy Sample.* ApJ, 694, (2009), 268. doi:10.1088/0004-637X/694/1/268.

DPAC CU leaders (2012). *DPAC Software and System Specification.* Tech. Rep. GAIA-C1-SP-DPAC-WOM-018, Data Processing and Analysis Consortium. URL http://www.rssd.esa.int/cs/livelink/open/2786798. Access restricted to DPAC members only.

Edge, D. O., **Shakeshaft**, J. R., **McAdam**, W. B. **et al.**, *A survey of radio sources at a frequency of 159 Mc/s.* MmRAS, 68, (1959), 37.

Elvis, M., *A Structure for Quasars.* ApJ, 545, (2000), 63. doi:10.1086/317778.

Etherington, I. M. H., *On the Definition of Distance in General Relativity.* Philosophical Magazine, 15.

Fath, E. A., *The spectra of some spiral nebulae and globular star clusters.* Lick Observatory Bulletin, 5, (1909), 71. doi:10.5479/ADS/bib/1909LicOB.5.71F.

Finet, F. & **Surdej**, J., *Multiply imaged quasi-stellar objects in the Gaia survey.* A&A, 590, (2016), A42. doi:10.1051/0004-6361/201425411.

Frigo, M. & **Johnson**, S. G., *The design and implementation of FFTW3.* Proceedings of the IEEE, 93(2), (2005), 216.

Gaia Collaboration, **Prusti**, T., **de Bruijne**, J. H. J. **et al.**, *The Gaia mission.* A&A, 595, (2016), A1. doi:10.1051/0004-6361/201629272.

Geurts, P., **Ernst**, D. & **Wehenkel**, L., *Extremely randomized trees.* Machine Learning, 63(1), (2006), 3. ISSN 1573-0565. doi:10.1007/s10994-006-6226-1.

Golub, G. H. & **Van Loan**, C. F. (1996). *Matrix computations (3rd ed.).* Johns Hopkins University Press. ISBN 0801854148.

González-Martín, O. & **Vaughan**, S., *X-ray variability of 104 active galactic nuclei. XMM-Newton power-spectrum density profiles.* A&A, 544, (2012), A80. doi:10.1051/0004-6361/201219008.

Greenstein, J. L. & **Schmidt**, M., *The Quasi-Stellar Radio Sources 3C 48 and 3C 273.* ApJ, 140, (1964), 1. doi:10.1086/147889.

Gunn, J. E. & **Peterson**, B. A., *On the Density of Neutral Hydrogen in Intergalactic Space.* ApJ, 142, (1965), 1633. doi:10.1086/148444.

Hall, P. B., **Anderson**, S. F., **Strauss**, M. A. **et al.**, *Unusual Broad Absorption Line Quasars from the Sloan Digital Sky Survey.* ApJS, 141, (2002), 267. doi:10.1086/340546.

Hubble, E., *A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae.* Proceedings of the National Academy of Science, 15, (1929a), 168. doi:10.1073/pnas.15.3.168.

Hubble, E. P., *A spiral nebula as a stellar system, Messier 31.* ApJ, 69. doi:10.1086/143167.

Inada, N., **Oguri**, M., **Shin**, M.-S. **et al.**, *The Sloan Digital Sky Survey Quasar Lens Search. V. Final Catalog from the Seventh Data Release.* AJ, 143, (2012), 119. doi:10.1088/0004-6256/143/5/119.

Institute of Electrical and Electronics Engineers, *Ieee standard for floating-point arithmetic.* IEEE Std 754-2008, 1–70. doi:10.1109/IEEESTD.2008.4610935.

Jaffe, W., **Meisenheimer**, K., **Röttgering**, H. J. A. **et al.**, *The central dusty torus in the active nucleus of NGC 1068.* Nature, 429, (2004), 47. doi:10.1038/nature02531.

Jansky, K. G., *Radio Waves from Outside the Solar System.* Nature, 132, (1933), 66. doi:10.1038/132066a0.

Jewett, J. & **Serway**, R. (2008). *Physics for Scientists and Engineers with Modern Physics.* Thomson Brooks/Cole. ISBN 9780495112402.

Johnson, S. G. & **Frigo**, M., *A modified split-radix FFT with fewer arithmetic operations.* IEEE Trans. Signal Processing, 55(1), (2007), 111.

Kaspi, S., **Smith**, P. S., **Netzer**, H. **et al.**, *Reverberation Measurements for 17 Quasars and the Size-Mass-Luminosity Relations in Active Galactic Nuclei.* ApJ, 533, (2000), 631. doi:10.1086/308704.

Keeton, C. R., *Computational Methods for Gravitational Lensing.* ArXiv Astrophysics e-prints.

Keeton, C. R., *On modeling galaxy-scale strong lens systems.* General Relativity and Gravitation, 42, (2010), 2151. doi:10.1007/s10714-010-1041-1.

Keeton, C. R. & **Kochanek**, C. S., *Gravitational Lensing by Spiral Galaxies.* ApJ, 495, (1998), 157. doi:10.1086/305272.

Kochanek, C. S., *Is There a Cosmological Constant?* ApJ, 466, (1996), 638. doi:10.1086/177538.

Kochanek, C. S., **Eisenstein**, D. J., **Cool**, R. J. **et al.**, *AGES: The AGN and Galaxy Evolution Survey.* ApJS, 200, (2012), 8. doi:10.1088/0067-0049/200/1/8.

Kormann, R., **Schneider**, P. & **Bartelmann**, M., *Isothermal elliptical gravitational lens models.* A&A, 284, (1994), 285.

Kovner, I., *The quadrupole gravitational lens.* ApJ, 312, (1987), 22. doi:10.1086/164846.

Lawrence, A., *The relative frequency of broad-lined and narrow-lined active galactic nuclei - Implications for unified schemes.* MNRAS, 252, (1991), 586. doi:10.1093/mnras/252.4.586.

Lemaître, G., *Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques.* Annales de la Société Scientifique de Bruxelles, 47, (1927), 49.

Lindegren, L. & **Perryman**, M. A. C., *GAIA: Global astrometric interferometer for astrophysics.* A&AS, 116, (1996), 579.

Machado, D. P., **Leonard**, A., **Starck**, J.-L. **et al.**, *Darth Fader: Using wavelets to obtain accurate redshifts of spectra at very low signal-to-noise.* A&A, 560, (2013), A83. doi:10.1051/0004-6361/201219857.

MacKay, D. (2003). *Information Theory, Inference and Learning Algorithms.* Cambridge University Press. ISBN 9780521642989.

Makarov, V. V., *Absolute measurements of trigonometric parallaxes with astrometric satellites.* A&A, 340, (1998), 309.

Mortlock, D. J., **Warren**, S. J., **Venemans**, B. P. **et al.**, *A luminous quasar at a redshift of z = 7.085.* Nature, 474, (2011), 616. doi:10.1038/nature10159.

Mushotzky, R. F., **Done**, C. & **Pounds**, K. A., *X-ray spectra and time variability of active galactic nuclei.* ARA&A, 31, (1993), 717. doi:10.1146/annurev.astro.31.1.717.

Nelder, J. A. & **Mead**, R., *A simplex method for function minimization.* The computer journal, 7(4), (1965), 308.

Nenkova, M., **Sirocky**, M. M., **Ivezić**, Ž. **et al.**, *AGN Dusty Tori. I. Handling of Clumpy Media.* ApJ, 685, (2008), 147. doi:10.1086/590482.

Netzer, H., *Revisiting the Unified Model of Active Galactic Nuclei.* ARA&A, 53, (2015), 365. doi:10.1146/annurev-astro-082214-122302.

Osterbrock, D. E., *Seyfert galaxies with weak broad H alpha emission lines.* ApJ, 249, (1981), 462. doi:10.1086/159306.

Padovani, P., **Alexander**, D. M., **Assef**, R. J. **et al.**, *Active galactic nuclei: what's in a name?* A&A Rev., 25, (2017), 2. doi:10.1007/s00159-017-0102-9.

Pappa, A., **Georgantopoulos**, I., **Stewart**, G. C. **et al.**, *The X-ray spectra of optically selected Seyfert 2 galaxies: are there any Seyfert 2 galaxies with no absorption?* MNRAS, 326, (2001), 995. doi:10.1046/j.1365-8711.2001.04609.x.

Pâris, I., **Petitjean**, P., **Ross**, N. P. **et al.**, *The Sloan Digital Sky Survey Quasar Catalog: Twelfth data release.* A&A, 597, (2017), A79. doi:10.1051/0004-6361/201527999.

Perryman, M. A. C., **de Boer**, K. S., **Gilmore**, G. **et al.**, *GAIA: Composition, formation and evolution of the Galaxy.* A&A, 369, (2001), 339. doi:10.1051/0004-6361:20010085.

Perryman, M. A. C., **Hassan**, H., **Batut**, T. **et al.**, eds. (1989). *The Hipparcos mission. Pre-launch status. Volume I: The Hipparcos satellite.*, vol. 1.

Peterson, B. M. (1997). *An Introduction to Active Galactic Nuclei.* Cambridge University Press. ISBN 9781139170901.

Peterson, B. M. & **Horne**, K., *Echo mapping of active galactic nuclei.* Astronomische Nachrichten, 325, (2004), 248. doi:10.1002/asna.200310207.

Press, W. H., **Teukolsky**, S. A., **Vetterling**, W. T. **et al.** (2002). *Numerical Recipes in C++: The Art of Scientific Computing.* Cambridge University Press. ISBN 0521750334.

Pringle, J. E., **Rees**, M. J. & **Pacholczyk**, A. G., *Accretion onto Massive Black Holes.* A&A, 29, (1973), 179.

Reber, G., *Cosmic Static.* ApJ, 100, (1944), 279. doi:10.1086/144668.

Refsdal, S., *On the possibility of determining Hubble's parameter and the masses of galaxies from the gravitational lens effect.* MNRAS, 128, (1964), 307. doi:10.1093/mnras/128.4.307.

Refsdal, S. & **Surdej**, J., *Gravitational lenses.* Reports on Progress in Physics, 57, (1994), 117. doi:10.1088/0034-4885/57/2/001.

Riess, A. G., **Filippenko**, A. V., **Challis**, P. **et al.**, *Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant.* AJ, 116, (1998), 1009. doi:10.1086/300499.

Robin, A. C., **Luri**, X., **Reylé**, C. **et al.**, *Gaia Universe model snapshot. A statistical analysis of the expected contents of the Gaia catalogue.* A&A, 543, (2012), A100. doi:10.1051/0004-6361/201118646.

Robin, A. C., **Reylé**, C., **Derrière**, S. **et al.**, *A synthetic view on structure and evolution of the Milky Way.* A&A, 409, (2003), 523. doi:10.1051/0004-6361:20031117.

Ryle, M. & **Smith**, F. G., *A New Intense Source of Radio-Frequency Radiation in the Constellation of Cassiopeia.* Nature, 162, (1948), 462. doi:10.1038/162462a0.

Saez, C., **Brandt**, W. N., **Shemmer**, O. **et al.**, *The X-Ray Properties of Typical High-redshift Radio-loud Quasars.* ApJ, 738, (2011), 53. doi:10.1088/0004-637X/738/1/53.

Sagar, R., **Omar**, A., **Kumar**, B. **et al.**, *The new 130-cm Optical telescope at Devasthal, Nainital.* Current Science, 101, (2011), 1020.

Sandage, A., *The Existence of a Major New Constituent of the Universe: the Quasistellar Galaxies.* ApJ, 141, (1965), 1560. doi:10.1086/148245.

Schlafly, E. F. & **Finkbeiner**, D. P., *Measuring Reddening with Sloan Digital Sky Survey Stellar Spectra and Recalibrating SFD.* ApJ, 737, (2011), 103. doi:10.1088/0004-637X/737/2/103.

Schmidt, M., *3C 273 : A Star-Like Object with Large Red-Shift.* Nature, 197, (1963), 1040. doi:10.1038/1971040a0.

Schmidt, M., *Space Distribution and Luminosity Functions of Quasi-Stellar Radio Sources.* ApJ, 151, (1968), 393. doi:10.1086/149446.

Schmitt, H. R., **Antonucci**, R. R. J., **Ulvestad**, J. S. **et al.**, *Testing the Unified Model with an Infrared-selected Sample of Seyfert Galaxies.* ApJ, 555, (2001), 663. doi:10.1086/321505.

Schneider, P. (1992). *Gravitational lenses.* Springer-Verlag. ISBN 3540970703.

Schneider, P. (2006). *Extragalactic Astronomy and Cosmology.* Springer-Verlag. ISBN 9783642540837.

Schneider, P., **Kochanek**, C. & **Wambsganss**, J. (2006). *Gravitational Lensing: Strong, Weak and Micro.* Springer-Verlag Berlin Heidelberg. ISBN 978-3-540-30310-7.

Schramm, T. & **Kayser**, R., *A simple imaging procedure for gravitational lenses.* A&A, 174, (1987), 361.

Simkin, S. M., *Measurements of Velocity Dispersions and Doppler Shifts from Digitized Optical Spectra.* A&A, 31, (1974), 129.

Soldi, S., **Türler**, M., **Paltani**, S. **et al.**, *The multiwavelength variability of 3C 273.* A&A, 486, (2008), 411. doi:10.1051/0004-6361:200809947.

Sonnenfeld, A., **Treu**, T., **Gavazzi**, R. **et al.**, *Evidence for Dark Matter Contraction and a Salpeter Initial Mass Function in a Massive Early-type Galaxy.* ApJ, 752, (2012), 163. doi:10.1088/0004-637X/752/2/163.

Souchay, J., **Andrei**, A. H., **Barache**, C. **et al.**, *The third release of the Large Quasar Astrometric Catalog (LQAC-3): a compilation of 321 957 objects.* A&A, 583, (2015), A75. doi:10.1051/0004-6361/201526092.

Stern, D., **Moran**, E. C., **Coil**, A. L. **et al.**, *Chandra Detection of a Type II Quasar at z = 3.288.* ApJ, 568, (2002), 71. doi:10.1086/338886.

Surdej, J., **Claeskens**, J. F., **Crampton**, D. **et al.**, *Gravitational lensing statistics based on a large sample of highly luminous quasars.* AJ, 105, (1993), 2064. doi:10.1086/116584.

Suyu, S. H., **Auger**, M. W., **Hilbert**, S. **et al.**, *Two Accurate Time-delay Distances from Strong Lensing: Implications for Cosmology.* ApJ, 766, (2013), 70. doi:10.1088/0004-637X/766/2/70.

Tadhunter, C., *Radio AGN in the local universe: unification, triggering and evolution.* A&A Rev., 24, (2016), 10. doi:10.1007/s00159-016-0094-x.

Terrell, J., *Luminosity Fluctuations and a Local Model of Quasi-Stellar Objects.* ApJ, 147, (1967), 827. doi:10.1086/149065.

Tristram, K. R. W., **Meisenheimer**, K., **Jaffe**, W. **et al.**, *Resolving the complex structure of the dust torus in the active nucleus of the Circinus galaxy.* A&A, 474, (2007), 837. doi:10.1051/0004-6361:20078369.

Trump, J. R., **Impey**, C. D., **Kelly**, B. C. **et al.**, *Accretion Rate and the Physical Nature of Unobscured Active Galaxies.* ApJ, 733, (2011), 60. doi:10.1088/0004-637X/733/1/60.

Tsalmantza, P., **Karampelas**, A., **Kontizas**, M. **et al.**, *A semi-empirical library of galaxy spectra for Gaia classification based on SDSS data and PÉGASE models.* A&A, 537, (2012), A42. doi:10.1051/0004-6361/201117125.

Türler, M., **Paltani**, S., **Courvoisier**, T. J.-L. **et al.**, *30 years of multi-wavelength observations of 3C 273.* A&AS, 134, (1999), 89. doi:10.1051/aas:1999125.

Upton, G. J. G., *Approximate confidence intervals for the mean direction of a von mises distribution.* Biometrika, 73(2), (1986), 525. ISSN 00063444.

Urry, C. M. & **Padovani**, P., *Unified Schemes for Radio-Loud Active Galactic Nuclei.* PASP, 107, (1995), 803. doi:10.1086/133630.

Vanden Berk, D. E., **Richards**, G. T., **Bauer**, A. **et al.**, *Composite Quasar Spectra from the Sloan Digital Sky Survey.* AJ, 122, (2001), 549. doi:10.1086/321167.

Véron-Cetty, M.-P. & **Véron**, P., *A catalogue of quasars and active nuclei: 13th edition.* A&A, 518, (2010), A10. doi:10.1051/0004-6361/201014188.

Watson, D., **Denney**, K. D., **Vestergaard**, M. **et al.**, *A New Cosmological Distance Measure Using Active Galactic Nuclei.* ApJ, 740, (2011), L49. doi:10.1088/2041-8205/740/2/L49.

Weymann, R. J., **Morris**, S. L., **Foltz**, C. B. **et al.**, *Comparisons of the emission-line and continuum properties of broad absorption line and normal quasi-stellar objects.* ApJ, 373, (1991), 23. doi:10.1086/170020.

Wu, Y., **Charmandaris**, V., **Huang**, J. **et al.**, *Spitzer/IRS 5-35 μm Low-resolution Spectroscopy of the 12 μm Seyfert Sample.* ApJ, 701, (2009), 658. doi:10.1088/0004-637X/701/1/658.

Wu, Y.-Z., **Zhang**, E.-P., **Liang**, Y.-C. **et al.**, *The Different Nature of Seyfert 2 Galaxies With and Without Hidden Broad-line Regions.* ApJ, 730, (2011), 121. doi:10.1088/0004-637X/730/2/121.

Yip, C. W., **Connolly**, A. J., **Vanden Berk**, D. E. **et al.**, *Spectral Classification of Quasars in the Sloan Digital Sky Survey: Eigenspectra, Redshift, and Luminosity Effects.* AJ, 128, (2004), 2603. doi:10.1086/425626.

Acknowledgements

Writing a thesis is an extremely demanding task that requires some sacrifices. Mine were the days and nights I spent, locked in my office whereas I was surrounded by my beloved family. I have a very particular thought for my son, Arthur, and for my daughter, Clémentine, to which this thesis is dedicated as a counterpart of the numerous moments I missed them. I would also like to thank my wife Hélène for her unconditional love, for her help and for her support. My parents should neither be forgotten for the natural curiosity and critical mind they brought me

The five years I used to work on the very interesting problems herein exposed were a palpitating adventure that would not have been possible without the intervention of some people. Obviously, I would like first to thank my supervisor, Jean Surdej, for his trust, his humanity as well as for the hundreds of hours he spent reading my drafts, articles, presentations and the final version of the present thesis. During these years I had the opportunity to carry out my research the way I wanted, with just the right advices and guidance from that reference person. This was incredibly more formative for me than any other framework in which this thesis would have been done. For all these, thanks Jean!

Although astronomy is a lively passion since I was I child, I did not follow an academic formation in this domain because I feared the lack of outlets. Éric Gosset and Yassine Damerdji, members of the jury, offered me the possibility of hanging on with this child's dream right after my master thesis. Éric, Yassine, thanks for this confidence and for introducing me in the wonderful world of astronomy.

Finally, my day by day work, would not have been the same without some of my colleagues, for the very interesting discussions we got but also for the less serious ones. I would particularly like to thank: Olivier (alias Numerobis) for the three years we shared the same office and for the great moments we got in it, Olivier (alias Numero-uno) for the ongoing writing of his reference book on whist rules, Denis for the regular organization of the pizza challenges, Maxime for his depth knowledge of Walloon, Bikram for the culinary exchanges we made, Sylvia for her assistance and the many others who will certainly recognize themselves here.