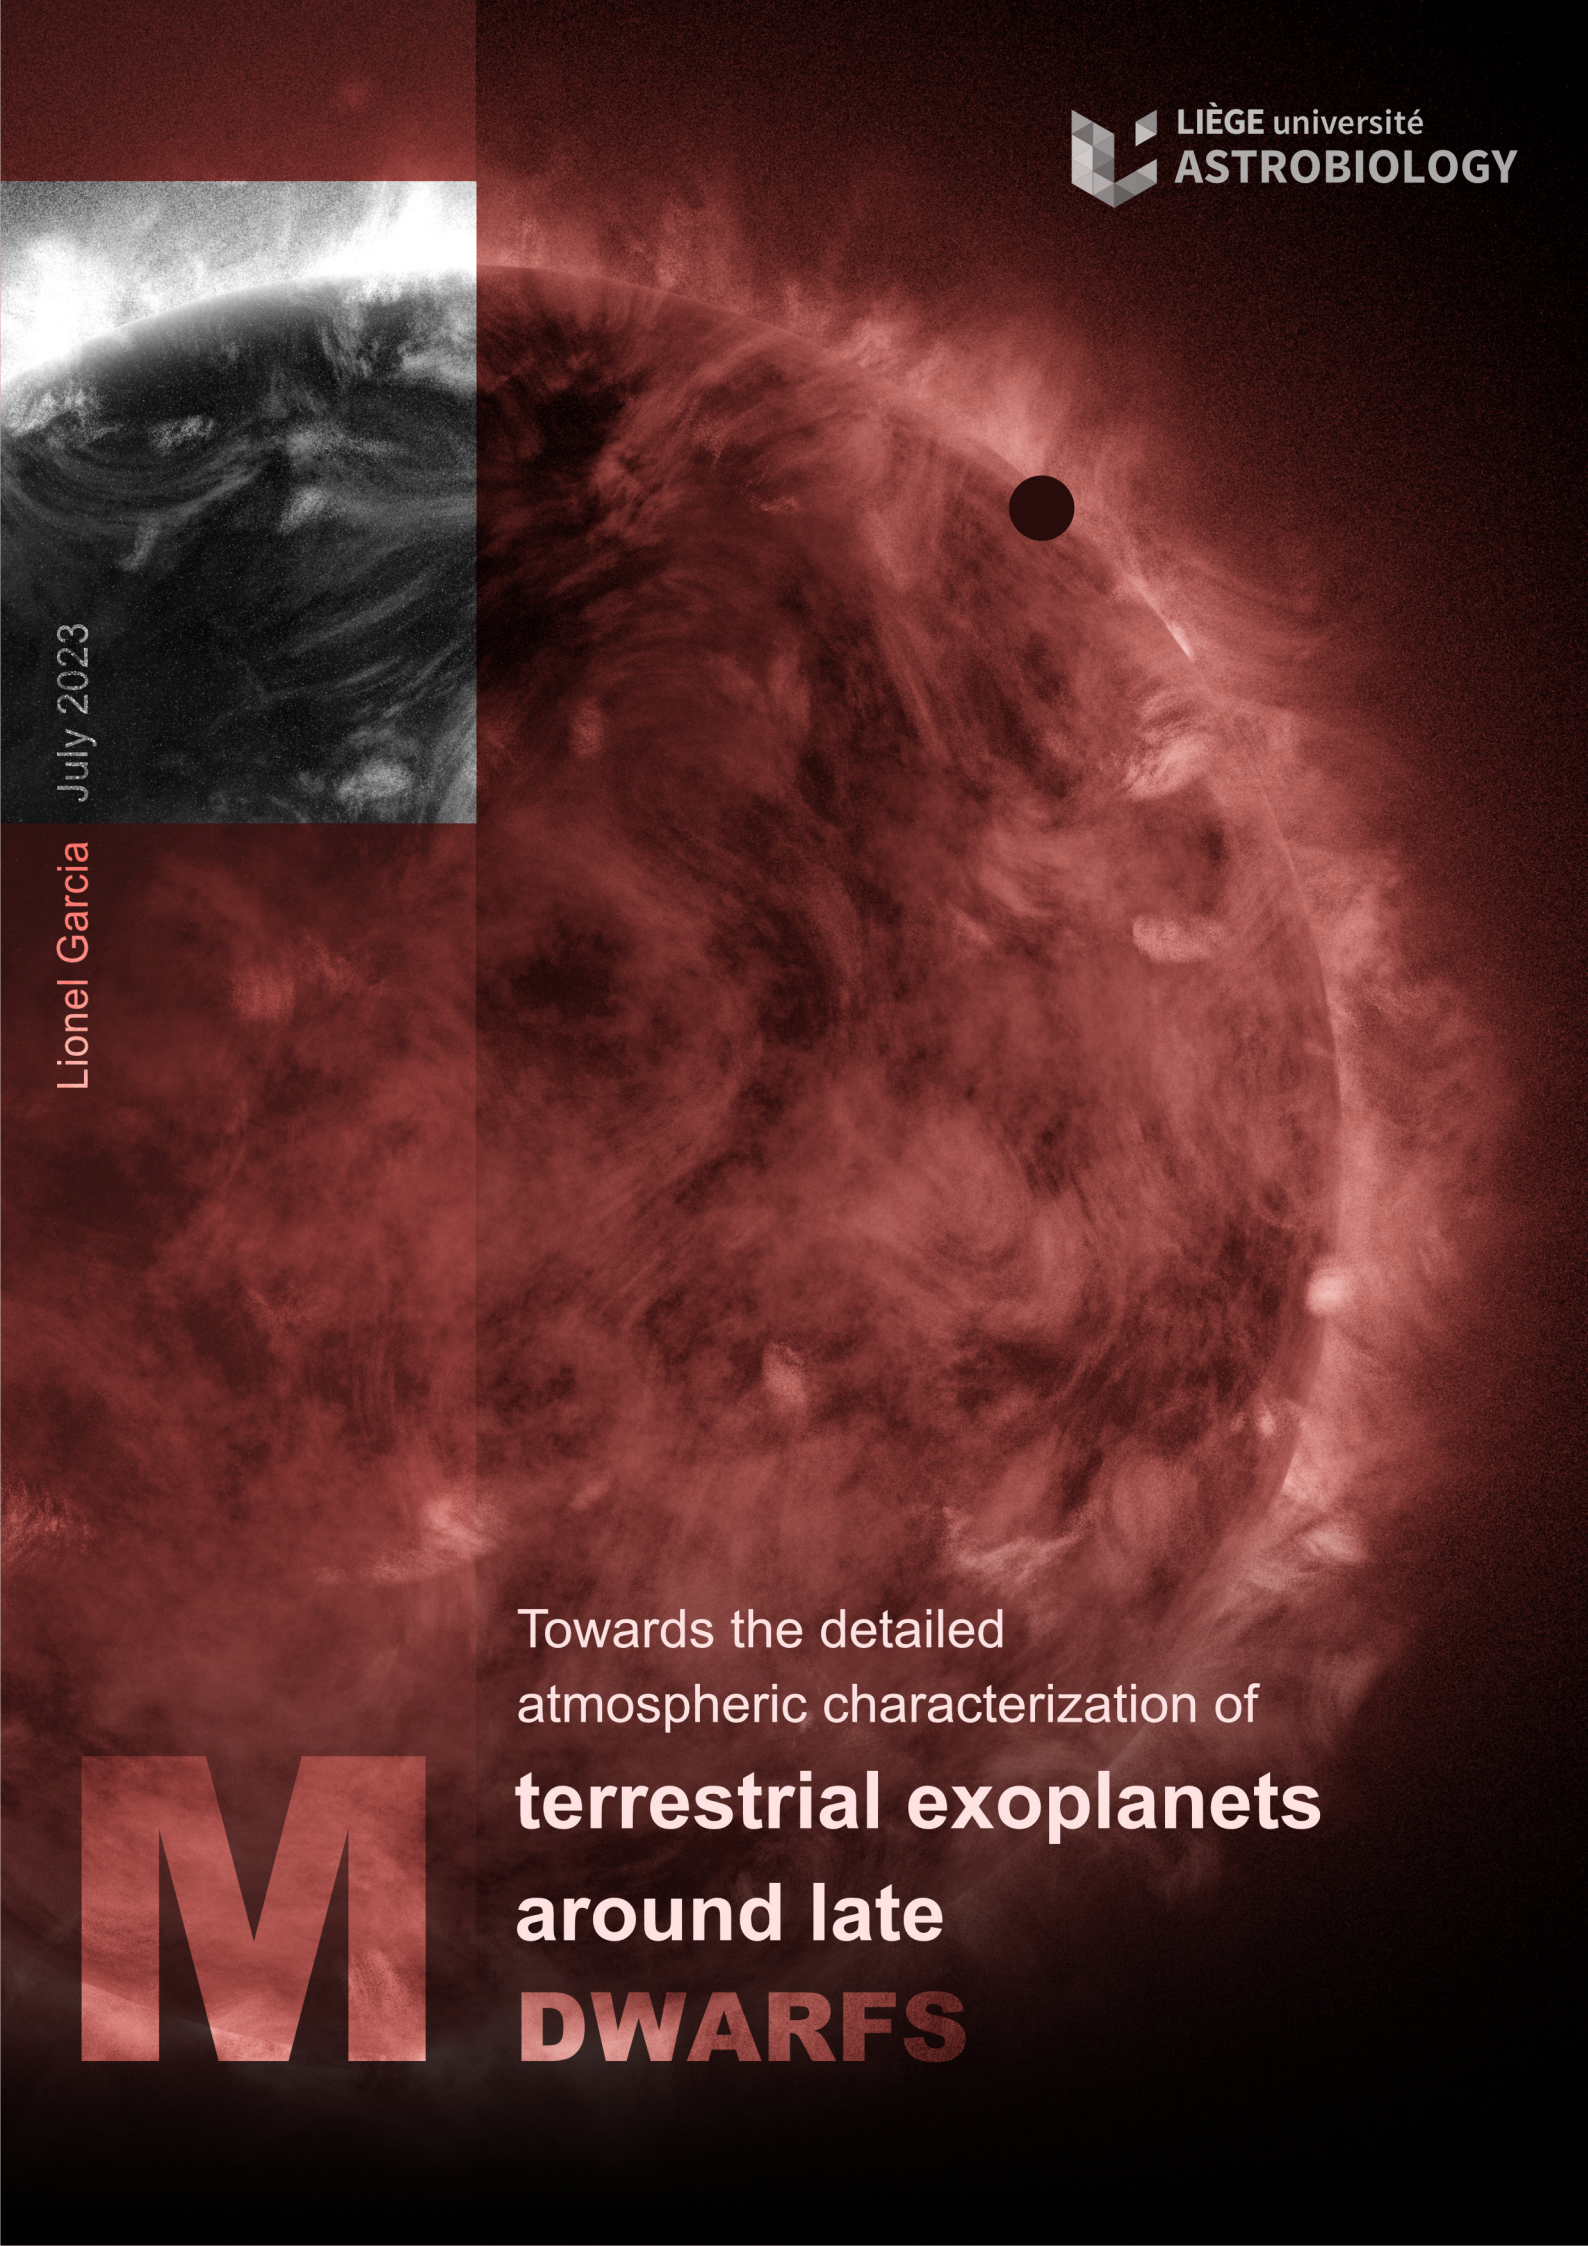


July 2023

Lionel Garcia

A large, detailed image of a red planet, likely Mars, with a prominent black dot in the upper right quadrant. The planet's surface shows various features like craters and atmospheric haze.

Towards the detailed
atmospheric characterization of
terrestrial exoplanets
around late
DWARFS

A large, stylized letter 'M' in a reddish-brown color, matching the background of the planet.

Towards the detailed atmospheric characterization of
terrestrial exoplanets around late M-dwarfs

Lionel Garcia

University of Liège
2023

Thesis supervisor

Dr. Michaël Gillon (University of Liège)

Members of the jury

Prof. Denis Grodent, President (University of Liège)

Dr. Laetitia Delrez (University of Liège)

Prof. Brice-Olivier Demory (University of Bern)

Prof. Julien De Wit (Massachusetts Institute of Technology)

Prof. Didier Queloz (University of Cambridge & ETH Zürich)

*A dissertation submitted in partial
fulfillment of the requirements for the degree of*

Doctor of Philosophy

Acknowledgement

I would like to start by expressing my deep gratitude to my supervisor, Michaël Gillon. With my background, starting a PhD in Astrophysics was a tough journey, and I would have given up on this project without his initial kick. Thanks for believing in me.

From the beginning of my PhD, I was surrounded by amazing colleagues. They made the journey very fun, and generously supported me through so many opportunities. In this regard I would like to thank everyone in the SPECULOOS consortium. A special thanks to Julien, for his very strong support and for calling me a poet. Peter, thanks for sharing my enthusiasm for technology. Besides the SPECULOOS consortium, many colleagues were absolutely instrumental to my research and the development of my ideas. In particular, I would like to thank Hannah Wakeford, for offering me a first paper and an amazing learning opportunity. No matter the project, Benjamin Rackham was someone I could freely share my ideas with, thank you so much. Among my external collaborators, I would finally like to thank Daniel Foreman-Mackey, for the immense support on my new projects and for being a great mentor.

In the University of Liège, the AGO department was the most welcoming place I have ever worked in. Starting with people in my group. Thank you Daniel, for sharing your office and being willing to discuss the details of all things. Seba, from the beer pong table to some telescope surgery, thanks for being such a great team member. Artem, it was a true pleasure to operate Artemis on your side, thank you for being so kind and available. And Khalid, thank you for sharing your passion and being so positive. During the highest heights and the deepest lows, some colleagues became true friends. Laetitia, you definitely listened to every bit of my struggles and crazy thoughts. Thank you for your kind understanding. Mathilde, I wish you could be in the next office for the rest of my career. Thank you so much for being such a friend and sharing the PhD journey with me. And finally Elsa. You made me come to Liège and supported me all the way through. I am ending up both with a PhD and an amazing friend. Thank you for everything fraté.

Beyond my team, I would like to thank everyone in the B5c. Martin, for the unforgettable laughter and drinks. Clémentine, the greatest happiness manager the department will ever have. Lynne, for the songs and the roundest of all rocks. And others: Marc-Antoine, Michaël De Beckers, Stéphane, Loïc, Sampath, Mathieu, Guillaume, Antoine and Maxime! You made the institute feel like home.

Other friends and family were essential supporters during my PhD. The next words are for them.

Merci Charly et Maryline, d'avoir été ma seconde famille pendant vos années à Liège. Merci Pauline et Juliette, pour avoir vécu cette intense dernière année avec moi. Merci Lorenzo, pour les pizzas et les Lupulus. Merci à tous mes amis de Paris: Théo, Vincent, Florian, David, Pennec, Marku, Mickael, Rémi, Magro, Juliette, Elorri, Valentine, Briec et Bouyss. Un merci tout particulier à toi Pyv, pour m'avoir toujours ouvert ta porte et pour les aventures partagées. Merci à Karine, pour me suivre malgré les années. Et merci Marine, pour le monde qu'on a refait. Les moments passés avec vous tous étaient de vraies bouffées d'oxygène

pendant ma thèse. Merci pour votre soutien durant ces 4 années.

Je voudrais finalement remercier ma famille, pour son soutien tout au long de mon parcours. Christian, merci de m'avoir transmis ta passion pour l'Astronomie. Sophie et Pauline, merci pour tout ce que vous m'avez inculqué et d'être toujours à mes côtés. Enfin, merci à mes parents. Maman, merci pour la liberté que tu m'as donnée, et pour avoir été aux premières loges de tous mes progrès. Papa, merci pour l'Univers auquel tu m'as initié, la Science, l'Art, et pour ton soutien permanent quand j'ai commencé à étudier. À vous deux, merci de m'avoir appris à échouer.

Arielle, dans la thèse comme dans la vie, merci d'être mon alliée.

Abstract

During the next decade, the atmospheric characterization of exoplanets as small as Earth will only be possible for planets transiting stars much smaller than the Sun: late M-dwarfs. While progress in instrumentation slowly enables the study of these remote worlds, many challenges remain. This PhD thesis intends to tackle some of these challenges, and provide a path towards the detailed atmospheric characterization of terrestrial exoplanets around late M-dwarfs.

Evidence from multiple fields suggest that life may have appeared and evolved in other places than Earth. From these fields, Information Theory and Thermodynamics have set themselves apart by jointly providing incentives to emergent phenomena beyond our subjective imagination, beyond life *as we know it*. Whatever form they take, such organisms should intrinsically put the environment they live in out of equilibrium, making them indirectly detectable. In this respect, the search for life in our Solar system is still ongoing, but focus is increasingly given to the search for life on *exoplanets*, planets orbiting stars other than the Sun. By measuring the composition of exoplanets atmospheres, one can assess the existence of such disequilibria and try to link their origin to the presence of a biosphere, but only to a certain degree. Indeed, distinguishing purely physical processes from potential biosignatures requires a deep understanding of the planet in its entirety, not only its atmosphere but also its oceans and crust (if terrestrial), in addition to its host star's environment. Such detailed knowledge is only available for a single planet in the whole Universe: Earth, making Earth-like planets compulsory analogues to search for life outside the solar system, on a terrestrial planet *as we know it*.

In the Introduction, I review the evidence suggesting life may have appeared and be present elsewhere in the Universe, on exoplanets. In [Chapter 1](#), I describe the architecture and observation of these remote worlds, and draw a road-map of how life could be detected on these planets in the next two decades. Doing so, I focus on the unique opportunity offered by exoplanets transiting late M-dwarfs, and on two non-instrumental challenges that remain to fully exploit their potential: the detection of transit signals in the light curves of these cool objects, and the concerning effect of stellar contamination.

In [Chapter 2](#), I expose my contribution to the observation of late M-dwarfs. I start by describing SPECULOOS, a transit survey searching for temperate terrestrial planets around 1700 nearby late M-dwarfs. Then I describe `prose`, a Python package to build modular image processing pipelines, initially developed to process the raw images of SPECULOOS.

In [Chapter 3](#), I describe the development and application of `nuance`, an algorithm to detect exoplanets around variable stars, well suited for late M-dwarfs.

In [Chapter 4](#), I focus on the second challenge posed by stellar contamination, employing the atmospheric reconnaissance of TRAPPIST-1h as a case study. Using Hubble Space Telescope WFC3 observations, I present an infrared transmission spectrum of the outermost planet of the ultra-cool M8-type star TRAPPIST-1 and show how stellar contamination hindered the interpretation of such measurements. Overall, I highlight the limitations of using

1D stellar atmosphere models and state-of-the-art techniques to solve this issue.

[Chapter 5](#) is a response to the effect of stellar contamination, and a description of new strategies to better understand late M-dwarfs. In this chapter, I describe the constraints that could be drawn from an ensemble analysis of the SPECULOOS rotational light curves, and their use in atmospheric retrievals. Finally, I expose how spot occultations by large planetary companions could help constrain the photospheres of late M-dwarfs, using the observation of TOI-3884 large polar spot as an example.

I conclude by making a summary of my contributions to the field of exoplanetary science, ultimately helping to lay a path towards the detailed atmospheric characterization of terrestrial exoplanets around late M-dwarfs, in the search for life.

Résumé

Au cours de la prochaine décennie, la caractérisation atmosphérique d'exoplanètes aussi petites que la Terre ne sera possible que pour des planètes en transit autour d'étoiles beaucoup plus petites que le Soleil : les étoiles naines de type spectral M. Alors que nos instruments sont de plus en plus aptes à étudier ces mondes lointains, de nombreux défis demeurent. Cette thèse de doctorat a pour but d'étudier certains de ces défis et fournit une voie vers la caractérisation atmosphérique détaillée d'exoplanètes terrestres en transit autour des naines M.

Plusieurs domaines suggèrent que la vie a pu apparaître et évoluer ailleurs que sur Terre. Parmi ces domaines, la théorie de l'information et la thermodynamique permettent d'expliquer l'apparition de phénomènes émergents qui dépassent notre imagination, allant au-delà de la vie telle que nous la connaissons. Quelle que soit leur forme, ces organismes placent leurs environnements hors de l'équilibre thermodynamique, ce qui les rend indirectement détectables. À cet égard, la recherche de la vie dans notre système solaire continue, en parallèle de la recherche de la vie sur des planètes en orbite autour d'étoiles autres que le Soleil. En mesurant la composition atmosphérique de ces exoplanètes, l'existence de tels déséquilibres peut être détectée, et reliée à la présence d'une biosphère. Cependant, distinguer les processus purement géophysiques des biosignatures nécessite une compréhension profonde de la planète dans son ensemble, non seulement de son atmosphère mais aussi de ses océans et de sa croûte (si elle est rocheuse), combinée à l'environnement stellaire. Une connaissance aussi détaillée n'est disponible que pour une seule planète dans tout l'Univers : la Terre, faisant des planètes de type terrestre des analogues primordiaux pour la recherche de vie ailleurs dans l'Univers.

L'introduction de ce manuscrit décrit les indices qui suggèrent que la vie a pu apparaître et être présente ailleurs dans l'Univers, sur des exoplanètes. Dans le [Chapitre 1](#), je décris l'architecture et l'observation de ces mondes lointains, et décrit la manière dont la vie pourrait être détectée sur ces planètes durant les deux prochaines décennies. Ce faisant, je me concentre sur l'opportunité unique offerte par les exoplanètes en transit autour d'étoiles de type M, et sur deux défis non-instrumentaux qui restent à relever pour exploiter pleinement leur potentiel : la détection d'exoplanètes autour de ces objets purement convectifs, et l'effet préoccupant de la contamination stellaire.

Le [Chapitre 2](#) expose ma contribution à l'observation d'étoiles de type M tardives. Je commence par décrire SPECULOOS : un survey à la recherche de planètes rocheuses en transit autour de 1700 naines tardives proches. Ensuite, afin de traiter ses images brutes, je décris `prose` : un package Python facilitant le design et l'exécution de pipelines dédiés au traitement d'images.

Dans le [Chapitre 3](#), je décris le développement et l'application de `nuance` : un algorithme pour détecter les transits d'exoplanètes autour d'étoiles variables, particulièrement adapté à la recherche de transits dans les courbes de lumière de naines tardives.

Le [Chapitre 4](#) se concentre sur la contamination stellaire, au travers de la reconnaissance atmosphérique de TRAPPIST-1h. En utilisant les observations du télescope spatial Hubble, je

présente un spectre de transmission infrarouge de la planète externe du système TRAPPIST-1, et montre comment la contamination stellaire impacte son interprétation. Dans l'ensemble, je souligne les limites de l'utilisation des modèles d'atmosphère stellaire 1D et le manque de techniques de pointe en ce qui concerne le traitement de la contamination stellaire.

Enfin, le [Chapitre 5](#) est une réponse à l'étude de TRAPPIST-1h, et présente de nouvelles stratégies pour lutter contre la contamination stellaire. Dans ce chapitre, je décris les contraintes qui pourraient être tirées d'une analyse d'ensemble des courbes de lumière rotationnelles de SPECULOOS, et comment cette étude pourrait aider les caractérisations atmosphériques à venir. Enfin, j'expose comment l'observation de planètes transitant des régions actives à la surface d'étoiles de types M pourrait aider à contraindre leurs photosphères, prenant comme exemple le cas unique de TOI-3884.

Je conclus cette dissertation par un résumé de mes contributions au domaine de l'exoplanétologie, faisant un pas de plus vers la caractérisation atmosphérique détaillée d'exoplanètes de type terrestres, à la recherche de vie ailleurs dans l'Univers.

Contents

Introduction	3
1 Exoplanets observation and characterization	6
1.1 A planet and a star	7
1.1.1 Kepler’s problem	7
1.1.2 Three dimensional orbit	8
1.1.3 Radial velocity	9
1.2 Transiting exoplanets	10
1.2.1 Transmission spectroscopy	11
1.2.2 Occultation spectroscopy	12
1.3 Transit light curves	13
1.3.1 Uniform stellar surface	13
1.3.2 Non-uniform stellar surface	14
1.4 Bayesian data analysis	20
1.4.1 Parameter estimation	20
1.4.2 Sampling algorithms	21
1.5 The M-dwarf opportunity	23
1.5.1 Signatures of molecular species	23
1.5.2 Late M-dwarfs	24
1.6 A road-map	27
2 Time-series photometry of late M-dwarfs	29
2.1 The SPECULOOS project	30
2.1.1 Telescopes	30
2.1.2 Survey	31
2.1.3 Data	31
2.2 Image processing with prose	33
2.2.1 Principle	33
2.2.2 Base photometric pipeline	37
2.2.3 TESS follow-up comparison with AstroImageJ	41
2.2.4 Other applications	46
Appendix	47
3 Detection of exoplanets around M-dwarfs	50
3.1 The issue with correlated noise and its detrending	53
3.1.1 The effect of correlated noise on transits detectability	53
3.1.2 The effect of detrending on transits detectability	54

3.2	nuance	57
3.2.1	The linear search	58
3.2.2	The periodic search	59
3.2.3	The transit search periodogram	62
3.2.4	An open-source python package	64
3.3	Control test against BLS and TLS	64
3.4	nuance <i>vs.</i> <code>biweight</code> +BLS on simulated light curves	66
3.4.1	Dataset	66
3.4.2	Results	66
3.5	Validation on TESS light curves	67
3.5.1	Injection-recovery studies	70
3.5.2	Results	72
3.6	Discussion	76
4	Atmospheric reconnaissance of TRAPPIST-1h	79
4.1	Observations	80
4.2	Light-curve modeling	82
4.2.1	Systematic models comparison	82
4.2.2	Transmission spectra inference	84
4.3	Stellar Contamination	87
4.3.1	Photosphere model	87
4.3.2	Models and data preparation	87
4.3.3	Fitting the out-of-transit spectra	89
4.3.4	Corrected transit depths	95
4.4	Planetary spectrum	96
4.4.1	Model Results with HST	97
4.4.2	Model Results with Full Wavelength Coverage	98
4.5	Conclusion	100
5	Prospects in studying the photosphere of late M-dwarfs	101
5.1	Making every transit count	102
5.2	Ensemble light curves analysis	103
5.3	Spots occultations	104
5.4	Open-source tools	106
	Conclusion	108

Introduction




Drop a stone from your hand: you can make assumptions simple enough to predict its trajectory. Drop a feather from your hand: some knowledge about aerodynamics will still allow you to make such predictions. But now, drop a living butterfly from your hand: chances are you will never be able to predict where it will be landing. While the butterfly still obeys the laws of aerodynamics, knowing what the butterfly will do requires a knowledge far more extended than the laws of physics. For a start, it requires to know what the butterfly knows about its environment: information. Add a flower to the scene and your chances of predictions are changed.

Understanding the complexity of life required the separate development of Biology, in which Physics could only explain separate and local phenomenons. But as Information Theory and Thermodynamics recently stepped into the game, our view of life as a rare and unique event is changing, and so is our view of its place in the Universe, among other physical phenomenons.

The story of our planet as we know it starts 4.5 billions years ago, in the early solar system, when Earth was a sterile rock with no moon. At this time, theories (reviewed in [Hartmann 2014](#)) suggest that Earth was involved in a dangerous dance with another body of the solar system, *Theia*, until the day they collapsed in a giant impact¹. After this dramatic event, remnants from the impact formed the Moon, while the surface of Earth was left covered with magma oceans that had to cool down. This is how astrobiologist Robert Hazen² describes the aftermaths of this dramatic impact:

"You've got a vast surface of area of Earth, you've got millions of years to play with, you've got all different types of chemical systems and fluxes. Earth was a great experimental laboratory for chemistry, with a lot of combinations of chemicals you can try. Life pops out of it."

600 millions years after the impact, Earth was striving with life ([Dodd et al., 2017](#)), and with the story of its origins comes the question of its definition. There is currently no scientific consensus on how to properly define life, but everyone share a common intuition about it: a butterfly is a living entity, and a stone is definitely not.



Outside obvious comparisons, the question becomes trickier for everything in-between, a continuum of chemical complexity where non-living systems share common characteristics with what we would call *living things*. These characteristics extend from *self-replication* ([England, 2013](#)) to *evolution* ([Solé et al., 2007](#)) up to redefining our sense of *intelligence*. For example, in a computer simulation where disk-shaped particles were placed in specific configurations ([Wissner-Gross & Freer, 2013](#)), a large disk was able to use a small disk to extract another disk from a narrow tube, mimicking what we would call a *tool*. In the same kind of simulation, two disks in different containers spontaneously synchronized to pull down a larger disk, a sense of *cooperation*. In opposition, what we accept to be characteristic of life (like *reproduction*) fails for obvious living organisms (e.g. a mule). The definition of life falls down to taxonomy.

Among complex systems, scientists recognize the peculiarity of biological structures, such as the prevalence of chirality in the molecular building-blocks of a butterfly ([Meierhenrich, 2013](#)), and the impact these systems have on their environments from a physical point of view. Using these purely physical principles, independent teams came to the conclusion that life is an inevitable outcome of processes arising from non-equilibrium thermodynamics (e.g. [Jarzynski 2006](#), [England 2013](#)³), living structures being the best at dissipating energy and

¹The *Giant Impact Hypothesis*.

²Episode *What is life?* of the podcast *The Joy of y*.

³Not without critics, see [Farmer 2005](#).

maximizing entropy. To achieve that, these systems must constantly adapt to their environment and naturally become inference machines, storing and processing information in the most energy-efficient way to make predictions and survive (Still et al., 2012). In that respect, claims about the ubiquity of life are backed and complemented⁴ by numerous studies in the field of Information Theory (e.g. Walker & Davies 2013).

When so many disciplines respond to each other presenting the emergence of life as an obvious thing, a natural question would be: *Where is everybody?*⁵

From these fields, we learn that life, complex systems being good at harvesting energy, only form in environments with the right composition (Adami & LaBar, 2015) exposed to an external source of energy putting it far from equilibrium (Smith & Morowitz, 2018). Planets in the so-called *habitable-zone* of their stars (where liquid water can exist) match these conditions, being rich reservoirs of chemicals exposed to very diverse host stars' radiations. From thermodynamics, we learn that life will arise from these non-equilibrium environments, and imprint its mark at the scale of whole planets.

As early as 1965, James Lovelock identified that non-equilibrium thermodynamics would be a direct marker for the presence of life on other worlds, and that studying the atmospheric composition of a planet would provide some mean of detection (Lovelock, 1965). In his study, he recognizes that the presence of Oxygen (in the form of O_2), being naturally very reactive hence rare, would be a significant clue for the presence of complex bio-chemistry, a term later coined in the literature as *biosignature*. As in Mathematics, proving that something exists is much easier than proving it does not, we only need one instance of it. For life elsewhere in the universe, we only need one detection. But what if O_2 can be produced by geophysical processes alone? How much false detections should we expect? Meadows et al. (2018) explore this question by studying sources of abiotic O_2 on exoplanets, highlighting that the detection of life should instead rely on combinations of biosignatures and a good understanding of the planetary geophysical and atmospheric interplay. In that respect, exporting our understanding of Earth to the study of other planets becomes compulsory. Hence, focusing the search for life to rocky exoplanets like Earth will maximize our chances of interpreting the biosignatures we are looking for.

⁴A natural reason why the two are linked is nicely illustrated in 7.5.1 *Information theory as a mirror on thermodynamics* of Eric Smith and Harold Morowitz' book *The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere* (see small review in Walker 2017).

⁵A famous but uncertain quote from Enrico Fermi during a conversation with fellow physicists Edward Teller, Herbert York, and Emil Konopinski in the 1950. This question and its implications results in what is called *the Fermi paradox*.

Chapter 1

Exoplanets observation and characterization

Some exoplanets may harbor life, complex systems that will necessarily imprint their mark at a planetary scale. In this chapter, I describe how planetary systems can be studied, in particular the atmosphere of exoplanets, where signs of life-like non-equilibrium thermodynamics could be detected.

Exoplanetary systems consist in at least one non-stellar object being gravitationally bound to a stellar system (one or several stars). This dissertation specifically deals with planets (without entering in their somehow controversial definition, taking Earth-like bodies as our main targets) in orbit around single host stars. The first mention of such systems, outside our own, can be tracked back to the 16th century, when the Italian philosopher Giordano Bruno raises the possibility that other stars may be surrounded by planets too. This possibility was later confirmed by the detection of two planets orbiting the pulsar PSR B1257+12 in 1992 (Wolszczan, 1994), and in 1995, with the detection of a planet orbiting the Sun-like star 51 Pegasi (Mayor & Queloz, 1995). By the end of the 20th century, the first multi-planetary system was discovered (Butler et al., 1999), and a total of 71 exoplanets were confirmed¹. This new era marked the beginning of exoplanetary science as we know it, with currently more than 5000 planets detected, and over 9000 candidates waiting confirmation².

¹<http://exoplanet.eu>

²<https://exoplanets.nasa.gov/discovery/discoveries-dashboard/> (2023.02.20)

1.1 A planet and a star

While many configurations are possible, let's consider the simple situation where a star of mass M_\star and a planet of mass M_p are isolated in empty space where only gravitational forces are at play (in a Newtonian setup where relativistic effects are neglected). We call \mathbf{r} the position vector of the planet with respect to the star, describing the system's motion under Newton's second law

$$\ddot{\mathbf{r}} + GM_\star M_p \frac{\mathbf{r}}{r^3} = 0, \quad (1.1)$$

where $\ddot{\mathbf{r}}$ is the acceleration vector of the planet and G is the universal gravitational constant. Treating the problem in polar coordinates, with the substitution $u = 1/r$ leads to

$$\text{Equation 1.1} \quad \Leftrightarrow \quad \frac{d^2 u}{d\theta^2} + u = G \frac{M_\star + M_p}{h^2},$$

also known as Binet's equation, where it can be shown that $h = r^2 \dot{\theta}$ is a constant. A general solution for this differential equation is

$$r = \frac{p}{1 + e \cos(\theta - \omega_0)}, \quad (1.2)$$

where e and ω_0 are constants of integration and $p = h^2/G(M_\star + M_p)$. Equation 1.2, that describes the motion of the planet, is the expression of conic sections in polar coordinates, whose specific form depends on the initial conditions of the system (see the margin Figure 1.1). Since planetary systems consist in stable elliptic orbits, I focus the rest of this dissertation on the circular and elliptic cases for which $0 \leq e < 1$ (red and grey in Figure 1.1). Hence, the trajectory of the planet in the rest frame of the star is an ellipse of ellipticity e and semi-major axis a given by $a = p/(1 - e^2)$ (see Figure 1.2), with the star located at one foci. If we define the true anomaly $f = \theta - \omega_0$ we can rewrite Equation 1.2 as

$$r = \frac{a(1 - e^2)}{1 + e \cos f}. \quad (1.3)$$

In this equation, ω_0 is called the *longitude of the periaapse* and is used to define the orientation of the ellipse in the orbital plane (based on reference points defined in Figure 1.3).

1.1.1 Kepler's problem

Although being very elegant, Equation 1.3 is rather impractical since it depends on the *true anomaly* f . As we will see in the next subsection, we seldom have access to such a direct observable and rather measure quantities along time t , making a time-dependant formulation of Equation 1.3 highly desirable. To solve this problem, we define the *eccentric anomaly* E (Figure 1.2) such that

$$r = a(1 - e \cos E), \quad (1.4)$$

and use Kepler's second and third laws to show that (Murray & Correia, 2010a, Section 3.)

$$M = E - e \sin E, \quad (1.5)$$

$$\text{with} \quad M = \frac{2\pi}{P}(t - t_0), \quad (1.6)$$

where P is the orbital period of the system and t_0 the boundary condition for $E = 0$. This corresponds to Kepler's equation, where M is defined as the *mean anomaly*. Given this expression, the procedure to find the time-dependant polar coordinates (r, θ) of the planet in the rest frame of the star at time t is:




Figure 1.1: Circular, elliptic and parabola conic sections.

1. Compute M using its definition from Equation 1.6.
2. Compute E by solving Kepler's Equation 1.5 (numerically).
3. Compute r using Equation 1.4 and f (then θ) using Equation 1.3.




Figure 1.2: Elliptic motion of a planet in orbit around its host star. O denotes the center of the ellipse and F the ellipse's focus where the star is located.

1.1.2 Three dimensional orbit

From the planet's coordinate in the plane of its orbit, we can compute its coordinates in three dimensions from any frame of reference. We define such a frame of reference, associated to an orthonormal right-handed system $(\vec{x}, \vec{y}, \vec{z})$, centered on the star, with \vec{z} pointing away from the star towards an observer (us, see Figure 1.3). Hence, we are interested in the planet's coordinates (and motion) in this frame.

The elliptical motion of the planet in the plane (\vec{x}_0, \vec{y}_0) of its orbit is known, a plane from which we define the orthonormal right-handed system $(\vec{x}_0, \vec{y}_0, \vec{z}_0)$. The transformation of the planet's coordinates from $(\vec{x}_0, \vec{y}_0, \vec{z}_0)$ to $(\vec{x}, \vec{y}, \vec{z})$ can be obtained following three rotations (Figure 1.3):

1. A rotation of angle ω about the \vec{z}_0 -axis, the *longitude of the periaapse*.
2. A rotation of angle $-i$ about the \vec{x}_0 -axis, the *orbital inclination*.
3. A rotation of angle Ω about the \vec{z} -axis, the *longitude of the ascending node*.

These angles are defined and shown in Figure 1.3, measured with respect to reference points that are specific to observations done from Earth. To compute these rotations, we use the rotation matrices $\mathbf{P}_x(\phi)$ and $\mathbf{P}_z(\phi)$ of angle ϕ respectively about the \vec{x} and \vec{z} axes:

$$\mathbf{P}_x(\phi) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \phi & -\sin \phi \\ 0 & \sin \phi & \cos \phi \end{bmatrix} \quad \text{and} \quad \mathbf{P}_z(\phi) = \begin{bmatrix} \cos \phi & -\sin \phi & 0 \\ \sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (1.7)$$

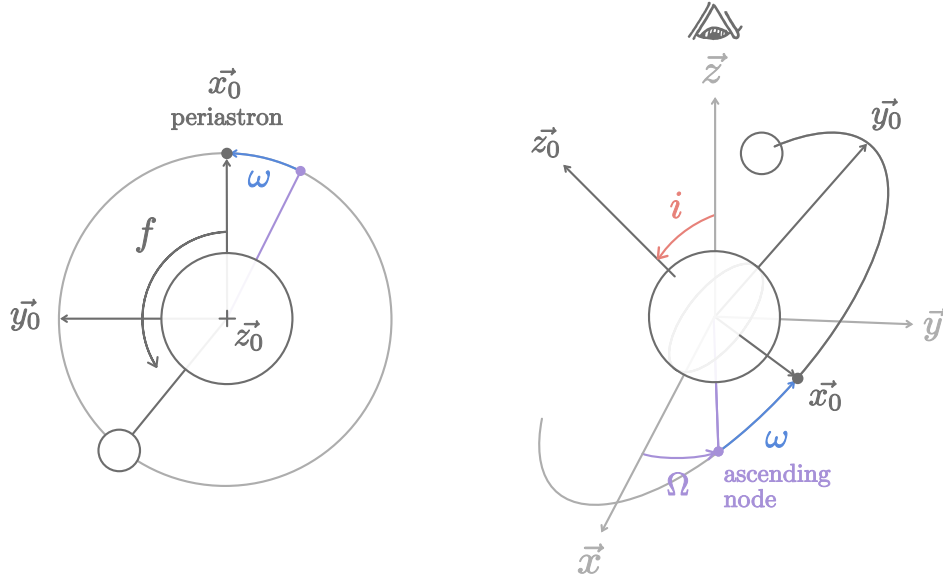


Figure 1.3: Orbital parameters of the two-body problem. The ascending node Ω is the point where the orbit of the planet crosses the \vec{x} - \vec{y} plane from below. The argument of the periastron ω is the angle between the ascending node and the periastron, the point of closest approach of the planet to the star. Finally, the orbital inclination i is the angle between the \vec{x} - \vec{y} plane and the plane of the orbit.

The position \mathbf{R} of the planet in the frame associated to $(\vec{x}, \vec{y}, \vec{z})$ is then


$$\mathbf{R} = \mathbf{P}_z(\Omega)\mathbf{P}_x(-i)\mathbf{P}_z(\omega)\mathbf{r}. \quad (1.8)$$

With these expressions, the cartesian coordinates (x, y, z) of the planet in the rest frame $(\vec{x}, \vec{y}, \vec{z})$ of the star are

$$\begin{aligned} x &= r(\cos(\Omega)\cos(\omega + f) - \sin(\Omega)\sin(\omega + f)\cos(i)), \\ y &= r(\sin(\Omega)\cos(\omega + f) + \cos(\Omega)\sin(\omega + f)\cos(i)), \\ z &= -r\sin(\omega + f)\sin(i). \end{aligned} \quad (1.9)$$

This set of equations leads to a variety of observable quantities that can be measured to infer the properties of a given planetary system.

1.1.3 Radial velocity



As an example, we can use [Equation 1.9](#) to constrain the mass M_p of a planetary companion around a known host star, using measurements of the star's *radial velocity* along \vec{z} . Indeed,

in the two-body problem, not only the two bodies orbit around each other but they also orbit around their common center of mass. Hence, the star also describes an eccentric orbit, with a semi-major axis a_* scaled as

$$a_* = a \frac{M_p}{M_p + M_*},$$

i.e. a much narrower orbit than the one of the planet if $M_p \ll M_*$. Hence, the *radial velocity* v_z of the star projected on \vec{z} , towards the observer (see [Figure 1.3](#)), can be computed ([Murray & Correia, 2010b](#)) using [Equation 1.9](#) as


$$v_z = \bar{v}_z + K(\cos(\omega + f) + e \cos \omega),$$

where \bar{v}_z is the *proper motion* (i.e. intrinsic) of the system in this direction, and

$$K = \frac{M_p}{M_p + M_*} \frac{na \sin i}{\sqrt{1 - e^2}}.$$

where $n = 2\pi/T$ is the mean motion of the planet, with T its orbital period. This variation in radial velocity, modulated by the planet's orbital motion with an amplitude K , leads to a change of its observed color due to the Doppler effect, that can be measured through spectroscopic observations. If M_* is known (thanks to spectroscopic measurements for example), measurements of K lead to an estimate of $M_p \sin i$, a lower constraint on the mass of the planet. This method was used to discover the first known exoplanet orbiting a Sun-like star ([Mayor & Queloz, 1995](#)).

1.2 Transiting exoplanets



An exoplanet with an orbital plane specifically aligned with our line of sight will eventually pass in front of its host star, blocking the light we receive from its atmosphere (see the transit of Venus in front of the Sun in [Figure 1.4](#)).


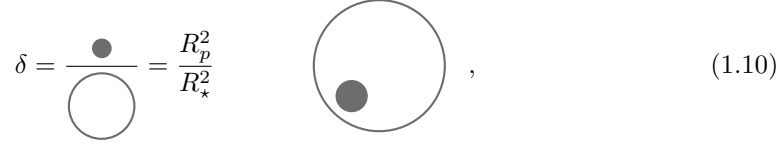


Figure 1.4: Transit of planet Venus in front of the Sun, observed by the Solar Dynamics Observatory on the 2020-06-06 in the 193 Å filter, sensitive to Fe XII/XXIV emission from the corona and hot flare plasma.

At first order, this event produces a decrease δ of the apparent flux we receive from the star corresponding to


$$\delta = \frac{\text{Area of planet}}{\text{Area of star}} = \frac{R_p^2}{R_\star^2} \quad , \quad \text{(1.10)}$$


a quantity denoted *transit depth* where R_p and R_\star are the planetary and stellar radii. Assuming that all stars host at least one exoplanet with a random orbital inclination i (i.e. that $\cos i$ is uniformly drawn from 0 to 1) the probability to observe an exoplanet transiting a star of radius R_\star is (Barnes 2007, Equation 8.)

$$p = \frac{R_\star}{a(1 - e^2)}, \quad \text{(1.11)}$$

where a and e are the semi-major axis and ellipticity of the orbit, and the planet's radius R_p is assumed to be negligible compared to the star's radius R_\star . Around 75% of known exoplanets have been discovered through the detection of transit events ³.

1.2.1 Transmission spectroscopy



During the transit of an exoplanet, part of the stellar light is filtered through the planet's atmosphere (if any). Due to the variety of chemical species it contains, the opacity of the atmosphere will be wavelength-dependant, and so will be the apparent radius of the exoplanet. This leads to a chromatic dependence of the transit depth. Modeling these changes requires the computation of radiative transfer equations, which consist in modeling how light is absorbed, emitted and scattered throughout the atmosphere, an expensive task.

A simpler estimate for the wavelength-dependant transit depth can be found by considering the simple case of an exoplanet with an atmosphere at hydrostatic equilibrium (Chamberlain & Hunten, 1987), for which the typical atmospheric scale height H is

$$H = \frac{k_b T_{eq}}{\mu g},$$

where k_b is the Boltzmann constant, T_{eq} the equilibrium temperature of the atmosphere, μ its mean molecular weight, and g the planet's surface gravity. For such a toy model, the transit depth observed at a wavelength λ is then simply


$$\delta_\lambda = \frac{(R_p + nH)^2}{R_\star^2} \quad \text{(1.12)}$$

³<https://exoplanets.nasa.gov/discovery/discoveries-dashboard/> (2023.02.20)


where n is the number of scale heights crossed at wavelengths with high opacity.

Then, transits observed at different wavelengths (*transmission spectroscopy*) can be used to infer the wavelength-dependant radius of an exoplanet and allows the detailed characterization of its atmosphere. This technique led to the first detection of an exoplanet on the gas giant HD 209458 b (Charbonneau et al., 2002) and is increasingly used in attempts to detect and characterize atmospheres on small terrestrial exoplanets (e.g. Chapter 4 and references therein).

1.2.2 Occultation spectroscopy



At opposite conjunction, the planet passes behind its star (*occultation*), leading to a decrease of the flux we receive from the system. This decrease corresponds to the thermal emission and reflection of the planet being occulted and contains precious information about the composition, temperature profile (emission) and albedo (reflection) of its atmosphere. Depending on their precision, such measurements performed at high cadence can even allow the 2D mapping of the planet's emission, layer by layer (e.g. Majeau et al. 2012; de Wit et al. 2012).



Throughout its revolution, the planet's thermal emission and reflection remains visible. Measuring the flux received from the system during a full orbit (*phase curve*) is a unique tool to measure the aforementioned quantities at different longitudes of the exoplanet, among which are the composition and physical properties of its atmosphere (e.g. Kreidberg et al. 2018).

In this dissertation, I study Earth-like planets around very small and cool stars, for which a precise-enough measurement of occultation and thermal phase curves are prohibitively expensive in terms of observing time. For this reason, I focus the rest of my thesis on the acquisition of transmission spectra, particularly informative for planet's orbiting late M-dwarfs, from the study of their atmosphere to advancing the current state of knowledge about their host star photospheres (see Chapter 5)

1.3 Transit light curves

As an exoplanet transits its star, it produces a time-varying decrease of its apparent flux: a time-series called *transit light curve*. While I explained the wavelength-dependant nature of these signals (*spectroscopic light curves*), I focus this section on describing *transit signals* observed at a given wavelength λ .

1.3.1 Uniform stellar surface

Let consider a star of radius $R = 1$ and uniform brightness $I = 1$, located at the center of the orthonormal system $(\vec{x}, \vec{y}, \vec{z})$ defined earlier, where \vec{z} points away from the star towards an observer. This star is transited by a planet of radius $r < 1$. We are interested in the total flux received from the star while its surface is occulted by the planet, i.e. $F = \pi - A$, where A is the overlap area between the planetary and stellar disks. To track the position of the planet relative to the center of the star, we introduce the distance $b = \sqrt{x^2 + y^2}$ (see Figure 1.5). In this configuration, A is the overlapping area between two disks of radii 1 and r , an area whose expression depends on the value of b .

- if $b \geq 1 + r$, the planetary disk is projected outside the stellar disk and $A = 0$
- if $b \leq 1 - r$, the planetary disk is entirely contained within the stellar disk, and $A = \pi r^2$
- if $|1 - r| < b < 1 + r$, the planetary disk projects only partially onto the stellar disk, and A corresponds to a lensed-shaped area, that can be found through its decomposition in a sum of simpler geometrical shapes (Figure 1.6)




Figure 1.5: Variable definition of the overlapping disks problem.




Figure 1.6: Simple geometric decomposition enabling the computation of the overlapping area A between two disks.

Doing so, the full set of expressions for A is

$$A = \begin{cases} 0 & 1 + r \leq b, \\ r^2 \kappa_0 + \kappa_1 - \sqrt{\frac{4b^2 - (1 + b^2 - r^2)^2}{4}} & |1 - r| < b < 1 + r, \\ \pi r^2 & b \leq 1 - r \end{cases} \quad (1.13)$$

(Agol et al., 2020), where

$$\begin{aligned}\cos \kappa_0 &= \left(\frac{(r-1)(r+1) + b^2}{2br} \right), \\ \cos \kappa_1 &= \left(\frac{(1-r)(1+r) + b^2}{2b} \right),\end{aligned}$$

and the differential flux $1 - A/\pi$ can be computed at any time, using the position (x, y) of the planet from Equation 1.9. Such a signal is shown in Figure 1.7.




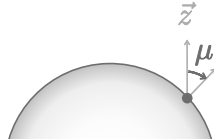
Figure 1.7: Transit light curve of star with a disk of uniform intensity, obtained by computing the area A over time using Equation 1.13 with $r = 0.1$ and b computed using Equation 1.9.

1.3.2 Non-uniform stellar surface

While the last subsection left us with a fairly simple expression of transit light curves, it is of very limited interest. Indeed, stellar disks are non-uniform by essence. The goal of this subsection is to present the causes of this non-uniformity, and to describe a mathematical framework to model the transit light curve of stars with non-uniform surfaces.

Limb darkening

Stellar atmospheres, in general, are not fully opaque, leading their disk-projected surfaces to look dimmer towards the limb. This effect, called *limb-darkening*, is due to the changing optical depth towards the edge of the disk, corresponding to a surface seen at grazing angles, exposing photons from deeper and cooler layers of the stellar atmosphere. A common approach to model these variations on the disk is to employ a quadratic formula (Claret, 2000), and express the intensity $I(\mu)$ of a point on the disk as



$$\frac{I(\mu)}{I(0)} = 1 - u_1(1 - \mu) - u_2(1 - \mu)^2. \quad (1.14)$$

where $\mu = \cos \gamma$, γ being the angle between the line of sight and the emergent intensity.

Active regions

Beyond the limb-darkening effect, the disk is a projection of the stellar surface, itself fundamentally non-uniform due to the presence of active regions. Figure 1.8 shows the surface of

the Sun for example, where small-scale *granulation* associated to convective cells of plasma can be observed. This convection generates powerful magnetic fields capable, in turn, to locally inhibit convection, leading to the large and cool sunspots (*starspots* for others stars) seen in [Figure 1.8](#). Other larger-scale structures can be present in the corona (see [Figure 1.4](#)). Overall, depending on the wavelength and the precision of the observation, these sources of non-uniformity strongly affect the shape and chromaticity of transit light curves, challenging the interpretation of transmission spectra (see [Chapter 4](#)).




Figure 1.8: Spots at the surface of the Sun (image courtesy of Alan Friedman, NASA)

Spherical Harmonics

To model high-precision transit signals, one has to model these sources of disk non-uniformity, accounting for both the effect of limb-darkening and the presence of stellar surface active regions. Since these structures lie on stellar spheres⁴, a mathematical object well suited for this application is the set of real *spherical harmonics*: special functions defined on the surface of a sphere that can serve as orthonormal basis to represent any surface map.

From [Luger et al. \(2019\)](#), the orthonormal real spherical harmonics $Y_{lm}(\theta, \phi)$ of degree $l \geq 0$ and order $m \in [-l, l]$ are defined in spherical coordinates as

$$Y_{lm}(\theta, \phi) = \begin{cases} \bar{P}_{lm}(\cos \theta) \cos(m\phi) & m \geq 0, \\ \bar{P}_{l|m|}(\cos \theta) \sin(|m|\phi) & m \leq 0, \end{cases}$$

with \bar{P}_{lm} the normalized associated Legendre functions ([Luger et al., 2019](#), Equation A2). On the surface of the unit sphere, the coordinates (x, y, z) can easily be found as

$$\begin{aligned} x &= \sin \theta \cos \phi, \\ y &= \sin \theta \sin \phi, \\ z &= \cos \theta, \end{aligned}$$

so that Y_{lm} are just polynomials of x , y and z . Defined on the surface only, we rewrite $z = \sqrt{1 - x^2 - y^2}$. So that the projected surface disk can be described by (x, y) only.

Using the spherical harmonics basis vector

$$\tilde{\mathbf{y}} = (Y_{0,0} \ Y_{1,-1} \ Y_{1,0} \ Y_{1,1} \ Y_{2,-2} \ Y_{2,-1} \ Y_{2,0} \ \dots)^\top,$$

⁴An approximation we make.

and the coefficients \mathbf{y} , the specific intensity of the disk-projected surface at point (x, y) is

$$I(x, y) = \tilde{\mathbf{y}}^\top \mathbf{y},$$

where Y_{lm} can be expressed solely using polynomials of x and y . [Figure 1.9](#) shows the disk-projected real spherical harmonics up to degree $l = 3$.




Figure 1.9: Real spherical harmonics up to the degree $l = 3$.

[Figure 1.12](#) shows a disk whose surface map is built from the real spherical harmonics basis up to order $l = 5$ associated with random coefficients.




Figure 1.10: Decomposition of a surface map into a sum of spherical harmonics.

Rotational phase curves

Stars rotate. In order to compute the disk integrated flux of a rotating stellar surface, we employ a special Wigner rotation matrix \mathbf{R} (different from Equation 1.7; Luger et al. 2019, Section 3) that can be directly applied to the spherical harmonic maps, so that the specific intensity at point (x, y) of the rotated surface map is

$$I(x, y) = \tilde{\mathbf{y}}^\top \mathbf{R} \mathbf{y}.$$

From there, the integrated flux over the projected disk of the (rotated) body is

$$F = \oint\!\!\!\oint I(x, y) dS. \quad (1.15)$$

Hence, the basis $\tilde{\mathbf{y}}$ needs to be computed only once, in opposition to the rotation matrix \mathbf{R} that needs to be computed over time, leading to the rotational phase curve of the star. To simplify the computation of the surface integral in Equation 1.15, we define the change of basis matrix \mathbf{A}_1 , transforming from the spherical harmonics basis $\tilde{\mathbf{y}}$ to the polynomial basis

$$\tilde{\mathbf{p}} = (1 \ x \ z \ y \ x^2 \ xz \ xy \ yz \ y^2 \ \dots),$$

such that a vector \mathbf{y} expressed in this basis is

$$\mathbf{p} = \mathbf{A}_1 \mathbf{y}.$$

After the surface maps have been rotated and expressed in this new basis, we can rewrite

$$I(x, y) = \tilde{\mathbf{p}}^\top \mathbf{A}_1 \mathbf{R} \mathbf{y}.$$

This way

$$\begin{aligned} F &= \oint\!\!\!\oint I(x, y) dS \\ &= \oint\!\!\!\oint \tilde{\mathbf{p}}^\top \mathbf{A}_1 \mathbf{R} \mathbf{y} dS \\ &= \left(\oint\!\!\!\oint \tilde{\mathbf{p}}^\top dS \right) \mathbf{A}_1 \mathbf{R} \mathbf{y}. \end{aligned} \quad (1.16)$$

We define

$$\mathbf{r}^\top \equiv \oint\!\!\!\oint \tilde{\mathbf{p}}^\top dS \quad \text{where the } n\text{-th component} \quad r_n = \oint\!\!\!\oint \tilde{p}_n(x, y) dS$$

is shown to have a simple analytical expression that only depends on l and m (again, described in Section 3 of Luger et al. 2019). The integrated flux over the projected disk of the rotated body is then

$$F = \mathbf{r}^\top \mathbf{A}_1 \mathbf{R} \mathbf{y}, \quad (1.17)$$

where \mathbf{r} and \mathbf{A}_1 are independent of the surface map coefficients (only their orders and degrees) and can be pre-computed for efficiency. If the map coefficients \mathbf{y} are known, only \mathbf{R} is variable, and corresponds to the rotation matrix of the surface map of the body. Like before, this matrix can be computed over time, yielding the rotation phase curve of the star.

Interestingly, the quadratic limb-darkening law from Equation 1.14 can easily be expressed in the polynomial basis $\tilde{\mathbf{p}}$, such that the rotational phase curve of a non-uniform limb-darkened

body can now be easily computed. Using Equation 1.17, Figure 1.11 shows the rotational phase curve of the surface map built in the previous subsection (Figure 1.10).




Figure 1.11: Rotation phase curve of the simulated surface map shown in Figure 1.10. The signal is plotted in arbitrary time units and each marker is the value of the disk-integrated flux of the maps shown at the top of the figure.

Transit light curves

Computing the light curve of a surface occulted by another spherical body is done thanks to the same Equation 1.16, but this time integrating the specific intensity only over the unocculted surface of the star. This integral then depends on the radius r and the position (x_0, y_0) of the occulting body (like in subsection 1.3.1) and may be difficult to evaluate.

In calculus, *Green's theorem* allows to express a double integral over a surface S of the form

$$\iint_S \left(\frac{\delta M}{\delta x} - \frac{\delta L}{\delta y} \right) dx dy$$

as the line integral over the closed contour C of S as

$$\oint_C (L dx + M dy),$$

where M and L are functions of (x, y) defined and differentiable on the surface S .

The idea to simplify the evaluation of the double integral in Equation 1.16 over the unocculted stellar disk is to apply Green's theorem (Pál, 2012), by expressing the contour of the occulting disk on this surface in a simple form, and finding a change of basis matrix \mathbf{A}_2 such that

$$\tilde{\mathbf{g}}^\top = \tilde{\mathbf{p}}^\top \mathbf{A}_2,$$

where it exists a vector function \mathbf{G} such that the n -th component of the new basis $\tilde{\mathbf{g}}$ can be written as

$$\tilde{g}_n(x, y) = \frac{\delta G_{ny}}{\delta x} - \frac{\delta G_{nx}}{\delta y},$$

decomposing $\mathbf{G}_n(x, y)$ into $G_{nx}(x, y) \vec{x} + G_{ny}(x, y) \vec{y}$. The total flux of the unocculted disk of the star is then

$$\begin{aligned} F &= \iint_S I(x, y) dS \\ &= \iint_S \tilde{\mathbf{g}}^\top \mathbf{A}_2 \mathbf{A}_1 \mathbf{R} \mathbf{y} dS \\ &= \mathbf{s}^\top \mathbf{A} \mathbf{R} \mathbf{y}, \end{aligned} \tag{1.18}$$

where $\mathbf{A} \equiv \mathbf{A}_2 \mathbf{A}_1$, and thanks to Green's theorem, the n-th component of \mathbf{s} is

$$s_n \equiv \iint_S \tilde{g}_n(x, y) dS = \oint_C \mathbf{G}_n(x, y) \cdot d\mathbf{r}. \tag{1.19}$$

In [Luger et al. \(2019\)](#), a basis $\tilde{\mathbf{g}}$ (called the *Green's basis*) obeying these properties is found, where \mathbf{G}_n can be simply expressed in terms of l and m . Given a last rotation of the disk through a matrix \mathbf{R}' , a simple expression for the geometrical contour C of the occulting body on the stellar disk is also established, so that the total flux F follows the idiomatic linear equation⁵

$$F = \mathbf{s}^\top \mathbf{A} \mathbf{R}' \mathbf{R} \mathbf{y}. \tag{1.20}$$

For a known surface described by \mathbf{y} , only \mathbf{R}' and \mathbf{R} are time-dependant, so that \mathbf{s}^\top and \mathbf{A} can be pre-computed once. This makes [Equation 1.20](#) very efficient to compute the transit light curve of a planet over a non-uniform stellar surface. The transit light curve of a planet of radius $r = 0.1$ in front of the stellar surface built in the previous subsection is shown in [Figure 1.12](#).




Figure 1.12: Transit light curve of the simulated surface map shown in [Figure 1.10](#) by a secondary companion (black disk). The signal is plotted in arbitrary time units and each marker is the value of the flux from the surface shown at the top of the figure.

We note that finding an expression for the line integral in [Equation 1.19](#) is tedious. Nonetheless, this expression exists and only involves sines, cosines and complete elliptic integrals, all being analytical. Finally, we note that the spherical harmonic basis (computed once) can serve to express [Equation 1.20](#) at different wavelengths, providing spectroscopic maps of the stellar surface.

⁵Read "starry".

1.4 Bayesian data analysis

Applying the technique of transmission spectroscopy to study the composition of an exoplanet atmosphere starts with the acquisition of spectroscopic transit light curves, and their comparison with a model from which the wavelength-dependant planet's radius can be inferred. In this section, I describe how this parameter inference can be done in practice.

1.4.1 Parameter estimation

At a given wavelength λ , let f be the observed flux of a star, sampled and arranged in the $(1 \times N)$ column-vector \mathbf{f} associated to the column-vector of times t . Given these *data*, we want to infer the planet's radius R_p ($R_{p,\lambda}$ where we omit λ for simplicity), a parameter involved as r in the expressions leading to Equation 1.20, that we assume to accurately describe the data. These data correspond to physical fluxes measurements, with associated measurement errors, so that R_p can be described as a random variable in this model. Other parameters, such as the stellar radius R_* , must also be treated as random variables, using values inferred from other type of measurements. Overall, the probabilistic modeling of these measurements and parameters (including R_p) must account for any prior knowledge that may exist on the system parameters, calling for a *Bayesian analysis* of the data.

Gaussian processes

In what follows, we assume that the flux measurements can be modeled as a Gaussian process: an extension of Gaussian distributions to functions. For discrete data, a Gaussian process is similar to a multivariate Gaussian distribution, described by a mean vector $\boldsymbol{\mu}$ and a matrix \mathbf{C} characterizing the covariance of the data. Applied to functions (i.e. at infinite dimensions) a Gaussian process can be described by a mean function $\mu(t)$ (of time t in our case) and a covariance function $k(t, t')$ that characterizes the covariance between the function values at any two times t and t' . Thanks to its relation to the multivariate Gaussian distribution, Gaussian processes can be easily sampled from.

Bayes theorem

With Bayes' theorem, the probability to observe the model parameters $\boldsymbol{\theta}$ (including R_p) given the data \mathbf{f} is

$$P = p(\boldsymbol{\theta}|\mathbf{f}) = \frac{p(\mathbf{f}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{f})}. \quad (1.21)$$

In this equation:

- $p(\boldsymbol{\theta}|\mathbf{f})$ is the *posterior* probability of the model parameters $\boldsymbol{\theta}$ given the data \mathbf{f} .
- $p(\boldsymbol{\theta})$ is the *prior* knowledge on the model parameters $\boldsymbol{\theta}$.
- $p(\mathbf{f}|\boldsymbol{\theta})$ is the *likelihood* of the data \mathbf{f} given the model parameters $\boldsymbol{\theta}$.

I assume that our fluxes measurements can be modeled by a Gaussian Process with a mean $\boldsymbol{\mu}$ and a covariance matrix \mathbf{C} , i.e.

$$\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}).$$

In other words, the data is drawn from a distribution with a mean $\boldsymbol{\mu}$ and covariance matrix \mathbf{C} . With these assumptions, the natural logarithm of the likelihood (aka *log-likelihood*) is simply

$$\log p(\boldsymbol{\theta}|\mathbf{f}) = -\frac{1}{2}(\mathbf{f} - \boldsymbol{\mu})^\top \mathbf{C}^{-1}(\mathbf{f} - \boldsymbol{\mu}) - \frac{1}{2} \log |\mathbf{C}| - \frac{1}{2} N \log 2\pi,$$

with N the number of measurements (i.e. the length of \mathbf{f}). If we assume that each measurement f_i is drawn from a normal distribution with variance σ_i^2 (the uncertainties of our measurements), we can write $\mathbf{C} = \sigma^2 \mathbf{I} + \mathbf{K}$, with $\sigma^2 \mathbf{I}$ a diagonal matrix and \mathbf{K} describing the covariance structure of the data (a condition called *homoscedasticity*).

- $p(\mathbf{f})$ is the *evidence*, acting as a normalization term, and is usually prohibitively expensive to compute.

Parameters expectation⁶

Hence, $p(\boldsymbol{\theta}|\mathbf{f})$ takes the form of a probability density function. Conditioned on this distribution, the expectation of the model parameters $\hat{\boldsymbol{\theta}}$ is

$$\hat{\boldsymbol{\theta}} = \langle \boldsymbol{\theta} \rangle = \int \boldsymbol{\theta} p(\boldsymbol{\theta}|\mathbf{f}) d\boldsymbol{\theta} = \frac{1}{Z} \int \boldsymbol{\theta} P(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (1.22)$$

where $P(\boldsymbol{\theta}) = p(\mathbf{f}|\boldsymbol{\theta})p(\boldsymbol{\theta})$, and $Z = p(\mathbf{f})$ is the evidence, independent of $\boldsymbol{\theta}$.

Usually, the expectation given by Equation 1.22 doesn't have a closed form, but an approximate solution can be found through a *Monte Carlo* algorithm which consists in:

1. Draw a set of N samples $\boldsymbol{\theta}_i$ from $p(\boldsymbol{\theta}|\mathbf{f})$
2. Compute

$$\hat{\boldsymbol{\theta}} \simeq \frac{1}{N} \sum_i^N \boldsymbol{\theta}_i$$

Unfortunately, we generally cannot draw samples directly from $P = p(\boldsymbol{\theta}|\mathbf{f})$ because 1) Z is unknown and expensive to compute and 2) even knowing Z , drawing from P would require to evaluate most of its values, an intractable task as the number of dimensions of $\boldsymbol{\theta}$ increases.

1.4.2 Sampling algorithms

Fortunately, there exists a wide variety of *sampling algorithms* to efficiently draw samples, not from P directly, but using its pointwise evaluation up to a certain factor, from $P^* = Z P$.

Metropolis-Hasting

Most sampling algorithms successively draw samples from P^* through a proposal (fixed and pre-defined) density Q , that must look like P and from which we can (unlike P) draw samples from.

In opposition, the *Metropolis-Hasting* method uses a distribution Q that does not need to look like P , and is adapted from each sample to the next. If we call $\boldsymbol{\theta}_i$ the i -th drawn sample and Q_i the probability density Q associated to $\boldsymbol{\theta}_i$, the Metropolis-Hasting method to draw $\boldsymbol{\theta}_{i+1}$ goes as follows:

⁶A presentation of the problem borrowed from Luengo et al. (2020).

1. Draw a sample θ_u from Q_i
2. Compute

$$a = \frac{P^*(\theta_u) Q_u(\theta_i)}{P^*(\theta_i) Q_i(\theta_u)},$$

and accept the new sample with a probability $\min(a, 1)$.

- If accepted, the new sample is $\theta_{i+1} = \theta_u$
- Else, the new sample is $\theta_{i+1} = \theta_i$

The density probability of the N samples formed this way converges to P as $N \rightarrow \infty$, whatever the choice of Q . This way, samples are drawn using evaluations of P^* only, without knowing Z .

Hamiltonian Monte Carlo

The Hamiltonian Monte Carlo method is a kind of Metropolis-Hasting method to sample distributions of the form

$$P(\theta) = \frac{e^{-E(\theta)}}{Z} \quad (1.23)$$

where E is differentiable. At sample θ_i , it uses the gradient of E to generate a sample θ_{i+1} based on its momentum \mathbf{p}_i simulated using Hamiltonian dynamics. With this information, the Hamiltonian Monte Carlo method draws samples towards the maximum probability, sampling a wider region of the target distribution and in a more efficient way (see [Figure 1.13](#)).

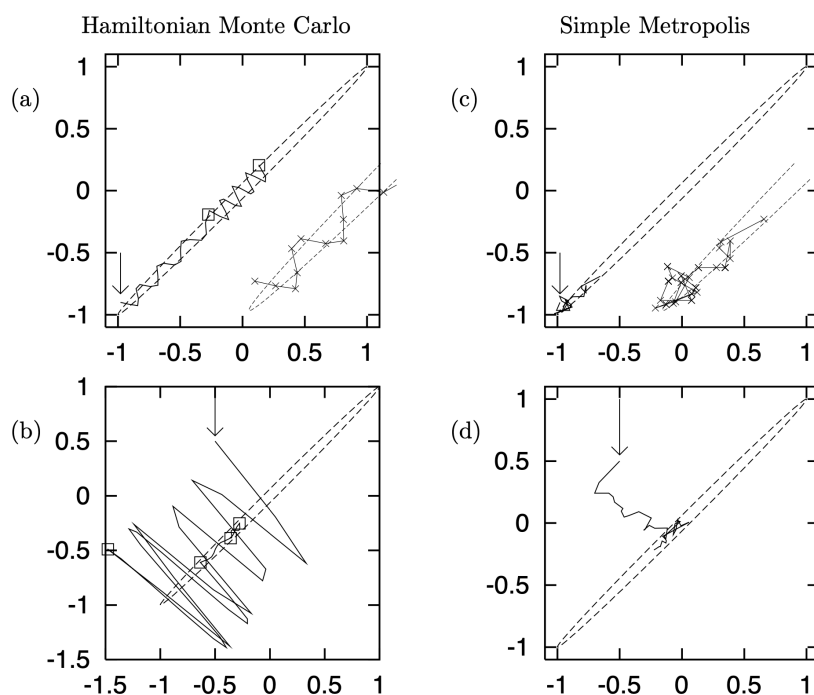


Figure 1.13: From [MacKay 2003](#): (a,b) Hamiltonian Monte Carlo used to generate samples from a bivariate Gaussian. (c,d) For comparison, a simple random-walk with the Metropolis-Hasting method, given equal computer time. Arrows mark the chains' starting points.

In this dissertation most parameters inferences are done through the Hamiltonian Monte Carlo method, exploiting the fact that the model given in [Equation 1.20](#) can be analytically differentiated.

1.5 The M-dwarf opportunity

1.5.1 Signatures of molecular species

In theory, partly thanks to the tools presented in the last section, we could use transmission spectroscopy to probe the atmospheric composition of any transiting exoplanet. In practice however, this exploration is fundamentally limited by our instrumental capabilities.

As an illustration, let consider the case of Earth being observed as an exoplanet through its transit in front of the Sun. Using [Equation 1.10](#), this signal has a depth

$$\delta_{\oplus} = \frac{R_{\oplus}^2}{R_{\odot}^2} = 8.4050179 \times 10^{-5} \simeq 84 \text{ ppm},$$

observable exactly once per year. The precision we can expect from this measurement would be dependant on the type of telescope and instrument used. Let assume that we are only limited by photon noise. If we observe N photons from this system, the signal-to-noise ratio (SNR) of the measurement is

$$S = \frac{N}{\sqrt{N}}.$$

However, as we are interested in studying the atmosphere of the planet using transit spectroscopy, we would perform this measurement at different wavelengths. Using a spectrograph covering a wave-band of width $\Delta\lambda$ with a resolution $R = \lambda/l$ (where l is the width of the resolution element), the SNR of the flux we receive from the star at wavelength λ is

$$S_{\lambda} = S \sqrt{\frac{\lambda}{R\Delta\lambda}}.$$

Hence, the larger the waveband we observe at the larger resolution R , the smaller our precision on the spectroscopic flux measurements of the system in each spectral bin. Translated into the signal-to-noise-ratio $S_{\delta_{\lambda}}$ of the transit depth δ_{λ} observed at λ ([Pont et al., 2006](#)), we have

$$S_{\delta_{\lambda}} = \delta_{\lambda} \sqrt{\frac{nN\lambda}{R\Delta\lambda}},$$

with n the number of transits. Hence, for a molecule \mathcal{M} found at scale height H ([Equation 1.12](#)), we approximate the signal-to-noise-ratio of its signature at its peak absorption

$$S_{\mathcal{M}} = \frac{H^2}{R_{\star}^2} \sqrt{\frac{nN\lambda}{R\Delta\lambda}}. \quad (1.24)$$

With these equations, we can estimate the strength of molecular species signatures found in the transmission spectrum of the planet, and assess the feasibility of their remote detections.

Using a formalism similar to the one described above, [Kaltenegger & Traub \(2009\)](#) consider a 6.5 meter telescope (such as the James Webb Space Telescope) observing an Earth-like planet transiting a Sun-like star located 10 parsec away, and explore the potential of transmission spectroscopy to detect several species of interest. Assuming a contemporary Earth

atmosphere, the signal-to-noise-ratio of O_2 is $S_{O_2} = 1.67$. One way to improve this figure is to increase the number n of observed transits (see Equation 1.24). For 200 hours of transits, S_{O_2} reaches 6.47, providing a detection of O_2 requiring the observation of 15 transits, a campaign spanning no less than 15 years for an Earth-like planet around a Sun-like star...

Looking closely at Equation 1.24, we understand that the signal-to-noise-ratio of a given species can be increased in several ways:

1. By observing at lower resolution to decrease R (i.e. to increase the resolution element l), or waiting for the development of telescopes capable of collecting an increased number of photons N (✗).
2. By focusing on extended atmospheres only, increasing the scale heights H . However, these atmospheres would not be representative of compact Earth-like atmospheres (✗).
3. By studying planets around much smaller stars. This way, not only the signal-to-noise-ratio S_M will increase through the quadratic decrease of R_* , but it will benefit from the fact that stars cooler and smaller than the Sun have a close-in habitable zone. Hence, planetary systems transiting such stars are both more likely and more frequently observed, increasing the number of observable transits n for a given time-frame (✓).

We illustrate the benefit of observing smaller stars in Figure 1.14, using some results from Kaltenecker & Traub (2009).

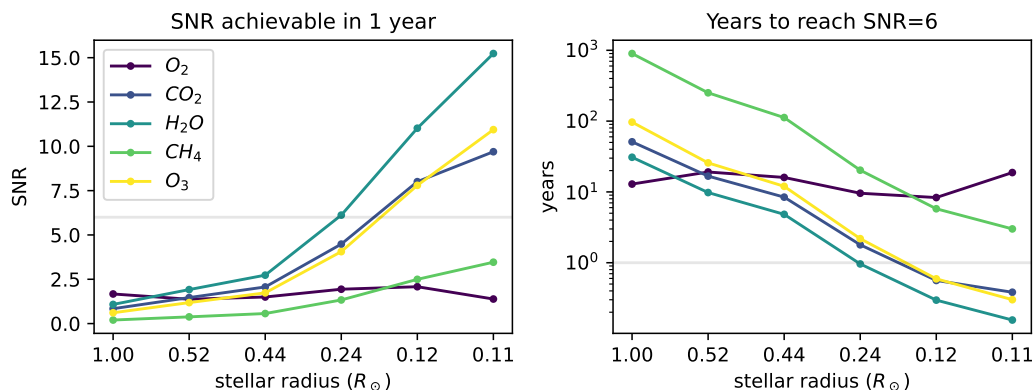


Figure 1.14: Simulation of molecular species SNR in the transmission spectrum of Earth-like planets around stars of different radii. This figure uses the SNR computed in Table 2 from Kaltenecker & Traub (2009), and the number of transits from Table 4. Radii are computed based on Mann et al. (2015) relations at different spectral types.

1.5.2 Late M-dwarfs

Stars with a radius R_* lower than $\sim 0.6 R_\odot$ are M-type stars, also called M-dwarfs, and lie at the bottom of the main sequence. These are the smallest and coolest main-sequence stars in the Universe, and the most abundant in our Galaxy ($\sim 80\%$, Henry et al. 1994). Due to their faintness and large convective layers (Reiners & Basri, 2009), these stars are both challenging to observe and model (discussed in the following chapters and, for example, in Iyer et al. (2023)). Despite this difficulty, arguments such as those previously drawn make

these small objects our best chance to characterize and understand rocky exoplanets in the next decades, by taking advantage of their short orbital periods and large transit depths when located in the habitable zone of their cool hosts. Before describing how the detailed atmospheric characterization of a rocky exoplanet around a cool M-dwarf may be achieved in the next two decades, I focus the next paragraphs on a major obstacle we will encounter when trying to achieve this goal: their potentially strong magnetic activity leading to a non-uniform photosphere.

In [subsection 1.3.2](#), we modeled the flux of a star with a non-uniform surface, partly motivated by the possible presence of stellar active regions. These regions (such as the *spots* shown in [Figure 1.8](#)) form through the complex interplay between the star’s plasma dynamics and magnetic fields. In the case of the Sun for example, the radiative pressure from the core counterbalances the inward forces of gravity, shaping the internal structure of the star as a large radiative zone surrounded by a thin convective envelope. Energy is transported differently in these layers, resulting in a slow rotation of the radiative core and a faster differential rotation of the convective layers. At the interface between the two lies the *tachocline*, where tubes of highly conductive plasma endure a strong shear and slowly transform the poloidal magnetic field of the Sun into a toroidal shape (see [Figure 1.15](#); e.g. [Charbonneau 2005](#)).




Figure 1.15: Evolution of the magnetic field lines of the Sun (in light grey, evolving from the left to right) due to shear effects induced by differential rotation at the star’s tachocline.

Plasma physics describes how the strength of these magnetic fields, carried in tubes, increase as they get elongated. This increased magnetic fields translate into a lower density plasma, that starts to rise towards the surface, directing field lines out of it. These strong fields inhibit the convective transport of heat from the core to the surface and lead to the surrounding plasma being at a cooler temperature, appearing darker. This process is at the origin of sunspots and explains how they form and evolve throughout the Solar cycle.

Late M-dwarfs are fully convective, with no radiative zone and no tachocline. However, they exhibit signs of activity such as flares (happening during the dramatic reconnection of magnetic-field lines after they emerge) and surface non-uniformity. Indeed, spots on M-dwarfs (e.g. [Bailer-Jones & Mundt 2001](#); [Libby-Roberts et al. 2023a](#)) are also believed to arise from strong magnetic fields generated in their fully-convective interiors ([Browning, 2008](#)). In addition, trends in magnetic fields topologies between partially-radiative stars (e.g. the Sun) and fully-convective stars (e.g. late M-dwarfs) are observed. Fully-convective stars tend to have strong poloidal magnetic fields while stars with masses greater than $0.5 M_{\odot}$ show weaker

magnetic fields with toroidal components (Morin et al., 2008). However, magnetic fields with different topologies are observed on stars with similar stellar parameters, a possible indication that magnetic topologies depend on stellar ages and evolve with time (in a cyclic or chaotic manner; Morin et al. 2010).


While being poorly understood, the presence of active regions at the surface of late M-dwarfs is an observational fact (cf. Chapter 5) that can have dramatic impacts on the characterization of their transiting exoplanets. Indeed, transmission spectra are built assuming that the transited chord is representative of the full stellar photosphere, enabling a differential spectrum to be constructed. However, any difference between the transited chord and the rest of the stellar photosphere may translate in stellar spectral features able to mimic or hide signatures from the planetary atmosphere. This effect, called *stellar contamination*, is expected to profoundly challenge the detailed atmospheric characterization of rocky exoplanets envisioned for the next two decades and is the focus of Chapter 4 and Chapter 5.

1.6 A road-map

Given current technologies, the detailed atmospheric characterization of exoplanets as small as Earth will only be possible for exoplanets transiting late M-dwarfs⁷. Despite being the most numerous stars in our solar neighborhood, these objects are poorly understood and require particular instrumentation and surveys to be developed. The following is a road-map of how the detailed atmospheric characterization of a rocky exoplanet around a late M-dwarf might be conducted, with instrumentation available in the next two decades.

OBSERVATION

M-dwarfs photometry




Everything starts with the search for rocky exoplanets, by making photometric observations of late M-dwarfs to search for transiting candidates. In this context, many surveys are already being conducted from the ground (e.g. [Nutzman & Charbonneau 2008](#); [Bonfils et al. 2015](#); [Gillon et al. 2017](#); [Delrez et al. 2018](#); [Dietrich et al. 2023](#)) and from space (e.g. [Ricker et al. 2015](#)). [Chapter 2](#) describes some of these observations, and the processing of their raw products.

DATA ANALYSIS

Transits detection


From large amounts of photometric light curves, transit search algorithms are used to detect potential candidates, later followed-up and accompanied by more observations, such as radial velocity measurements to have an estimate of their masses. However, due to their non-uniform surfaces, the detection of transit signals in the light curves of active M-dwarfs still poses a challenge to the community, especially for small rocky exoplanets. [Chapter 3](#) is a review of existing transit search techniques and present a new algorithm particularly suited to late M-dwarfs light curves.



⁷ Assuming that new technologies to directly image and study small rocky exoplanets remain inaccessible in the next decade.

OBSERVATION

Transmission spectroscopy




Once promising targets are identified, spectroscopic light curves can be obtained using major space-based and ground-based facilities. As current instrumentation will require many transits to be observed, this effort might take a couple of decades, even for a single target, but might result in auxiliary data that can help to understand these planetary systems as an all, including their host stars. Such measurements are presented in [Chapter 4](#).

DATA ANALYSIS & MODELING


Spectra decontamination

From the inferred transmission spectra, it will be essential to disentangle planetary signals from stellar ones, and tackle the so-called effect of stellar contamination. To deal with it, the community is currently left with no other choice but to employ 1D stellar atmosphere models, whose mutual discrepancies for M-dwarfs strongly bias the inference of planetary atmosphere properties ([Iyer et al., 2023](#)). [Chapter 4](#) demonstrates the impact of this approach on the atmospheric characterization of TRAPPIST-1h and [Chapter 5](#) presents alternative strategies and new avenues to better understand and tackle the effect of stellar contamination.



Atmospheric retrieval

Finally, atmospheric properties can be inferred from decontaminated planetary transmission spectra. Among these properties is the chemical composition of an atmosphere, and the study of the planet's surface conditions, matching or not with the presence of complex non-equilibrium thermodynamics. In this context, a detailed modeling of planetary atmospheres is required, as well as assessing the significance of biosignatures combinations: two subjects extending beyond the present dissertation.



Chapter 2

Time-series photometry of late M-dwarfs

A chapter based on the work published in [Garcia et al. 2022b](#) with Mathilde Timmermans, Francisco J. Pozuelos, Elsa Ducrot, Michaël Gillon, Laetitia Delrez, Robert D. Wells, and Emmanuël Jehin. In this chapter, my personal contributions have been distinctly emphasized and clearly differentiated from those made by co-authors.

The first step to characterize the atmosphere of a rocky exoplanet is to find one. In this quest, late M-dwarfs offer several advantages, but also come with few drawbacks. These stars are faint, and have their peak intensity in the near-infrared, requiring specific surveys with specific instruments. Such activities are ongoing ([Nutzman & Charbonneau, 2008](#); [Quirrenbach et al., 2014](#); [Artigau et al., 2014](#); [Bonfils et al., 2015](#); [Delrez et al., 2018](#); [Dietrich et al., 2023](#)) fitting into a global effort to detect and characterize the planetary systems they host. In this chapter, I describe the SPECULOOS survey ([Gillon, 2018](#)), one of the highest-precision ground-based photometric survey as of 2023, specifically designed to observe late M-dwarfs.

Once observed, whether by the SPECULOOS telescopes or other facilities, extracting the photometry of stars from full frame images is computationally demanding, and motivated the development of several codes and frameworks in the last decades. I review some of these tools and present `prose`, a Python package dedicated to astronomical images processing offering several advantages over state-of-the-art alternatives. Overall, this chapter is an update on [Garcia et al. 2022b](#), describing version 3.0.0 of the package and ending with a brief showcase of the Science it enabled since its creation in 2021.

2.1 The SPECULOOS project

2.1.1 Telescopes

Given their potential impact on the field of exoplanetary science, the SPECULOOS survey specifically focuses on late M-dwarfs, using 6 telescopes located around the globe (Figure 2.1) to search for transiting Earth-sized planets.


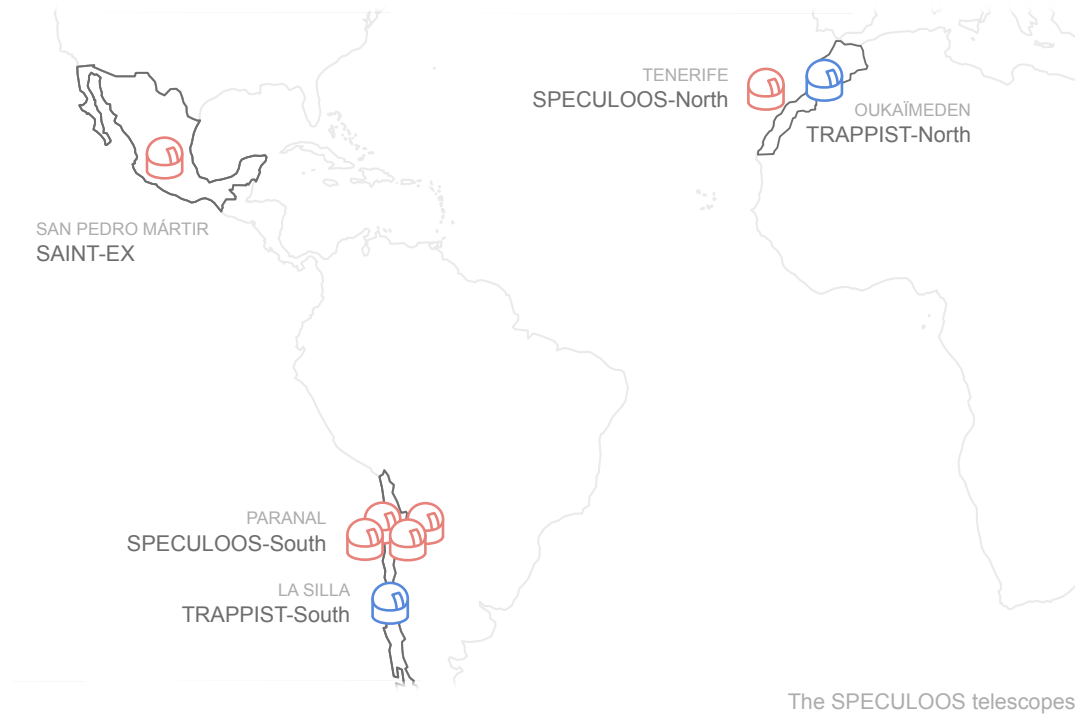


Figure 2.1: Location of the SPECULOOS telescopes. Blue symbols correspond to the TRAPPIST telescopes, used to conduct the SPECULOOS prototype survey, and red symbols to the more recent SPECULOOS telescopes.

In 2010, The TRAPPIST-South telescope (Gillon et al., 2011) was first installed and commissioned at the ESO La Silla Observatory, Chile. The TRAPPIST-South telescope is a robotic telescope with a 60 cm primary mirror with a f/8 Ritchey–Chrétien design attached to a German equatorial mount (all built and designed by the German company ASTELCO).

Searching for transits on the 50th brightest southern very-late M-dwarfs (Gillon et al., 2013), this prototype survey led to the discovery of the TRAPPIST-1 system: 7 Earth-sized planets orbiting an M8 type star, 3 of which lying in its habitable zone. In parallel to this survey, the first two SPECULOOS telescopes were installed in 2017 at Paranal, Chile. These were later followed by two more telescopes, thereby completing the SPECULOOS-South observatory (Delrez et al., 2018; Murray et al., 2020). In the same year, this network of robotic telescopes was further complemented by the TRAPPIST-North telescope located in Oukaimeden, Morocco, as well as a SPECULOOS-North telescope (Burdanov et al., 2022) installed in Tenerife, Spain. Lastly, the Saint-Ex telescope was installed in 2019 at San Pedro-Martir, Mexico, as a further extension of the SPECULOOS network. Overall, the SPECULOOS project is led by the University of Liège, in collaboration with the University of Cambridge, the University of Birmingham, The University of Bern, the Massachusetts Institute of Technology and ETH Zurich.

2.1.2 Survey

The SPECULOOS survey focuses on 1,657 ultra-cool-dwarf stars within 40 parsecs from the Sun, with spectral types ranging from M4 to L9 (Sebastian et al., 2021). The survey contains three separate programs designed to maximize its scientific return in synergy with currently deployed space and ground-based facilities:

- *Program 1* focuses on the search of temperate Earth-sized planets (with a solar irradiance of $1 S_{\oplus}$) that are amenable for atmospheric characterization using transmission spectroscopy with the James Webb Space Telescope (365 targets).
- *Program 2* focuses on targets already observed by the TESS mission (171 targets).
- *Program 3* contains all other targets with spectral types later than M6 (1,121 targets) and intend to assess the occurrence rate of exoplanets around ultra-cool-dwarf stars.

The strategy of the SPECULOOS survey is optimized to observe at least 80% of the orbital phase of the typical planet of each program, i.e. a $1 S_{\oplus}$ planet for program 1 and a $4 S_{\oplus}$ planet for program 2 and 3. This is achieved by observing each target from 100 to 200 hours depending on each program requirements.

2.1.3 Data

The aforementioned strategy leads to telescopes operating during all nights with good observing conditions, daily yielding ~ 70 Gigabytes of raw data in the form of full-frame FITS images (Wells & Greisen, 1979). From these images, photometric time-series of the survey’s targets need to be extracted (see next section) and made available for the SPECULOOS consortium to be vetted and searched for transits. In order to obtain first-hand products, an automatic pipeline was developed and deployed by Murray et al. (2020), and a web-based interface developed by Peter Pihlmann Pedersen (back-end) and myself (front-end) to ease the vetting process (see features in Figure 2.2).

The next section presents **prose**, a tool to process raw images into photometric time series, with applications extending well beyond the field of exoplanetary science.

Exploration tools
latest observations or random ones

An extensive search bar
with autocomplete features

Quick links

Centralised **documentation** for download and analysis

An interactive calendar to explore our **schedules**

The interface features a dark sidebar on the left with the SPECULOOS logo and navigation options: Observations (latest/random), Observation (random), Target (random), Search (Target, Date), Telescopes (SSO/SNO status, Trappist status, SSO webcam), Documentation (Products/API), and Schedules (calendar for NOVEMBRE 2020). The main panel displays 'Trappist-1' with 575.81 hours of observations, 145 light-curves, and 81 submissions. It shows two light-curve plots: one for 2020-10-17 and another for 2020-10-15. The 2020-10-15 plot includes a tooltip for 'Daniel S. 2020-11-28 transits 1' and interactive controls for aperture, binning, and PWV.

Target header with access to related **Slack discussions**

Interactive light-curves featuring collaborative tools (submissions)

Detailed observation view

Trappist-1b 2635476908753563008

2020-10-15 I+Z

✓ observation seems clear

Exposure	23 s
Observed	2 hours
Mean fwhm	3.30 pixels

[see all...](#)

The detailed view shows a light-curve plot for 'Trappist-1b' on 2020-10-15. The plot displays relative flux over time (JD - 2 450 000) with a zoomed-in section around JD 9,138,581. A tooltip for 'Daniel S. 2020-11-28 transits 1' is visible. The plot includes interactive controls for aperture, binning, and PWV. On the right, a 'fwhm' section shows 'sky' selected, with options for dx, dy, and airmass.

Easy vetting of common observation measurements

Live view of observation images when hovering plot(s)

Figure 2.2: SPECULOOS data portal aspect and features.

2.2 Image processing with prose

When an observation is completed, one has to process its raw products into a desired output. Observatories and surveys-specific pipelines are sometimes available (Murray et al. 2020 in the case of SPECULOOS) but general or custom software are often required, offering higher flexibility and control. These tools and their development benefit from legacy and modern algorithms, implemented into community-maintained libraries (like IRAF (Tody, 1986) or *astropy* (Astropy Collaboration et al., 2013, 2018)). However developing a custom option requires time and effort. Indeed, once the algorithms are selected, their implementations must be assembled into a complex software, ingesting raw FITS images and able to interact with many components of a given computational stack (CPUs, GPUs, file system, databases... etc.). To reach an optimized output, a fine-tuning of this code is often required, from its individual processing components up to the complete pipeline, including complex I/O operations and architectural choices. In its final version, such software are generally frozen and maintained (rather fixed) when they break. These tools, whether custom or third-party, seldom provide a complete analysis framework such that their end products still end up being exploited into scripting languages like Julia¹ or Python².

In 2020, driven by the need of the SPECULOOS survey, I proposed an alternative for these tools, fundamentally limited in their design. Pipelines are generally described as a sequence of individual data-processing units. These processing units (hereafter called *blocks*) are fundamentally common from one pipeline to another. They complete independent tasks using algorithms widely used by the whole community. Hence, in a framework where such blocks are all inter-compatible, one would only have to focus on which algorithms to use and assemble these processing units into a modular pipeline, with less need for software architecture considerations.

In this section I present *prose*³, a Python package acting as a framework to build modular image processing pipelines for astronomy, allowing rapid iterations toward specific needs and rich scientific products. Its architecture allows the creation of inter-compatible blocks, many being pre-implemented and available in the initial package. Developed in the context of transiting exoplanet observations, it contains tools specific to high-precision point-source photometry. However, the aim of *prose* is to be versatile and to adapt to a variety of astronomical datasets, something driving the development of its latest versions.

2.2.1 Principle

A *pipeline*, strictly speaking, is a series of connected tubes running a fluid. In the scientific literature, the word refers to processing pipelines in which data are flowing, going through processing units as in tubes. Usually, image processing algorithms are implemented as standalone codes which are easy to run on single images. However, assembling multiple low-level codes, successively ran on a set of images requires a complete architecture to be built and tested. The final result is often monolithic, far from where we take the word *pipeline* from. Of course, the flexibility offered by lower-level codes is essential. However, working with astronomical observations during the writing of this dissertation led to a simple idea,

¹<https://julialang.org/>

²<https://www.python.org/>

³<https://prose.readthedocs.io>

that generalizing the approach of building pipelines for astronomical image processing would come with great advantages, especially if available in a well-maintained open-source package.

`prose` contains the structure for such a modular concept to work by providing three key objects: *Image*, *Block* and *Sequence*. An *Image* encapsulates the data and metadata of an exposure frame, much like a FITS file does (Wells & Greisen, 1979). A *Block* is a single processing unit that acts on an *Image*. Finally, a *Sequence* is a series of *Block*, assembled sequentially to process images (see Figure 2.3).




Figure 2.3: Illustration of `prose` key objects

Blocks

With this architecture, one can focus on designing robust and reusable blocks, independently tested and maintained. To do that, `prose` comes with pre-implemented blocks featuring common state-of-the-art algorithms, some being listed in Table 2.1. Describing all base blocks of `prose` is beyond the scope of this paper, and the interested reader is redirected to the online documentation⁴.

⁴<https://prose.readthedocs.io>

Detection blocks	
AutoSourceDetection	Detect all sources
PointSourceDetection	Detect point sources
DAOFindStars	DAOPHOT (Stetson, 1987) stars detection
SEDetection	Source Extractor (Bertin & Arnouts, 1996) stars detection
TraceDetection	Detect trace sources
PSF blocks	
MedianEPSF	Stack cutouts into a median effective PSF
Gaussian2D	Model PSF as a 2D Gaussian profile
Moffat2D	Model PSF as a 2D Moffat profile
Trim	Image trimming.
Alignment & geometry blocks	
Cutouts	Create cutouts around all sources
Align	Align image data to a reference image
AlignReferenceSources	Align sources to reference image
AlignReferenceWCS	Compute image WCS based on a reference image
Centroiding blocks	
CentroidCOM	Center of mass centroiding
CentroidQuadratic	Quadratic model centroiding
CentroidGaussian2D	2D Gaussian model centroiding
CentroidBallet	Neural network centroiding
Photometry blocks	
AperturePhotometry	Aperture photometry
AnnulusBackground	Background annulus photometry
Other blocks	
Get	Retrieve and store properties from an image
Calibration	Flat, Bias and Dark calibration
SortSources	Sort sources
CleanBadPixels	Detect and clean bad pixels

Table 2.1: List of some base blocks currently included in `prose`

Telescope

Blocks are likely to use image characteristics to complete their tasks. In the case of FITS images, this information is located in their headers, whose definition may differ between observatories. To deal with the variety of header keywords and configurations, `prose` provides a `Telescope` object containing a semantic dictionary of keywords used to retrieve physical

quantities from image headers. Once a new *Telescope* is declared, its dictionary is saved and used whenever its name appears, filling the appropriate *Image* metadata. This mechanism, also found in general reduction tools like AIJ (Collins et al., 2017), allows *prose* to be instrument-agnostic but is only suited to single extension FITS products containing simple headers. Images from the Vatican Advanced Technology Telescope 4K camera (VATT4k), for example, are taken by two amplifiers, leading to FITS images spread over two extensions. Each of these extensions contains data as well as useful header information, while more general characteristics can be found in a separate primary header. To allow such images to be combined and processed, *prose* allows custom *Image* loaders to be developed (which in the case of the VATT4k images, for example, consists in 5 lines of Python code).

A simple example

Most blocks in *prose* use state-of-the-art methods and libraries. The novelty of the package only stands in its object-oriented structure, allowing for powerful abstraction in implementing custom pipelines. In this section, I present a simple example: a pipeline to characterize the effective point-spread-function (PSF) of images taken during a full night of observation. For every image, the pipeline starts with bias, dark and flat calibration (Howell, 2006, section 4.5). Then, bright stars are detected to build an effective PSF from cutouts, and characterized using a specific PSF model. Here, I use a 2D Moffat profile (Moffat, 1969). Finally, images are stacked and a movie of the night is generated for further vetting.

The Python code to build such a pipeline with *prose* (version 3.0.0) is shown in Figure 2.4, where blocks have been assembled into a single sequence to perform each task.

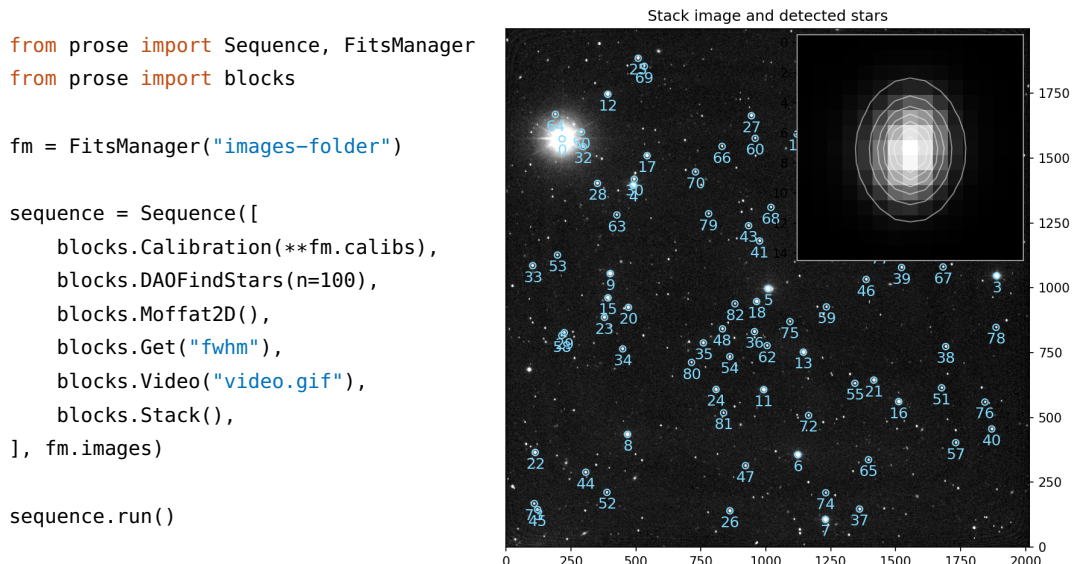


Figure 2.4: Python code of the pipeline described in 2.2.1, performing the characterization of images effective PSF over a full night of observation. The right plot shows the resulting stack image and PSF model of the observation and was produced using products directly accessible from the `sequence` object. This example is made from 742 10s exposure images taken by the TRAPPIST-South telescope on December 14, 2018.

In this pipeline, a wide range of stellar profiles or detection techniques may be experimented with through the use of different blocks. [prose](https://prose.readthedocs.io) online documentation⁵ provides many more examples as well as the description of several convenient tools, such as the `FitsManager` object, allowing to deal with highly disorganized FITS folders by using their header information.

2.2.2 Base photometric pipeline

Using the structure previously described, I developed a standard photometric reduction pipeline requiring minimal user inputs. This pipeline can be decomposed in few steps:

1. A reference image is selected on which reference stars are detected and the effective PSF of the observation characterized. This can be done by running the following sequence on the reference image:

```
from prose import Sequence, blocks

calibration = Sequence([
    blocks.Calibration(**calibs),      # calibration
    blocks.Trim(),                     # trimming image overscan
    blocks.PointSourceDetection(n=20), # stars detection
    blocks.Cutouts(21),               # making small cutouts around stars
    blocks.MedianEPSF(),              # combine cutouts into an effective PSF
    blocks.psf.Moffat2D(),             # modeling the effective PSF
])
```

The Calibration block employs a typical bias, dark and flat correction ([Howell, 2006](#), section 4.5) and PSF-related blocks are added in preparation of step 3.

2. A set of radii are defined, providing the relative sizes of the circular apertures used for the photometric extraction.
3. All images are calibrated and aligned to the reference, before the fluxes from each detected star are extracted. For each image, all apertures are scaled with the full-width at half-maximum (FWHM) of the effective PSF, allowing the photometric reduction to adapt to seeing variations throughout the observation. This processing is performed through the following sequence:

```
photometry = Sequence([
    *calibration,                      # we reuse the previous sequence
    blocks.AlignReferenceSources(ref), # alignment of sources to the reference
    blocks.CentroidBallet(),           # centroiding
    blocks.AperturePhotometry(radii=radii), # aperture photometry
    blocks.AnnulusBackground(),       # background estimate
    blocks.GetFluxes()                # storing fluxes information
])
```

Due to pointing errors, a centroiding block (`CentroidBallet`) is employed (see [section 2.2.2](#)).

The products obtained by running these steps can be accessed within a Python instance, through the values stored in the `GetFluxes` block, and saved for later use. A variety of time-series are also recorded in this block, such as the mean FWHM of the effective PSF, airmass,

⁵<https://prose.readthedocs.io>

background levels or frame misalignments, overall essential to model the extent of systematic signals in the measured fluxes.

Differential photometry

During the reduction, visible stars are automatically detected for every frame and aperture photometry is performed over a range of 40 apertures, yielding ideal measurements to leverage the power of ensemble differential photometry. For that purpose, *prose* implements a pure Python version of the algorithm described in [Broeg et al. \(2005\)](#). This algorithm consists in building an *artificial comparison star* using the weighted sum of all available stars in the field. Appropriate weights are iteratively found as to minimize variance over all differential light curves. Doing so, higher weights are given to stars displaying lower variability and higher signal-to-noise ratio (SNR), more likely to feature systematic signals shared among all sources. Finally, the best aperture is chosen as the one minimizing the white noise of the target’s light curve, estimated with the median standard deviation of points within 5 minutes bins.

The photometric pipeline described above relies on a strategy which consists in aligning the sources of each to a set of reference stars, requiring a perfect registration of the detected stars to the reference ones and a powerful centroiding algorithm to refine their positions. If processing time is a concern, general registration and centroiding algorithms are a problem. To keep *prose* fast on domestic computers, a centroiding method based on neural-networks has been developed, as well as a fast registration algorithm based on [Lang et al. \(2010\)](#), both presented in the next subsections.

Twirl: stars registration

Stars registration is a common need in astronomy. Given two sets of stars positions, it consists in identifying pairs of coordinates corresponding to unique objects, but observed from different angles, hence requiring an absolute reference on the celestial sphere. To do that, many algorithms, as well as humans, birds and seals ([Foster et al., 2018](#)), use *asterisms*: shapes produced by stars arrangements, like constellations, recognizable in the sky. [Lang et al. \(2010\)](#) provide the most frequently used method for this application, standing at the core of the widely adopted astrometric calibration service [astrometry.net](#). While it produces robust astrometric solutions, it requires a cross-match to a whole-sky catalog, that must be avoided if speed is required. Hence, a faster solution consists in avoiding the computation of a complete WCS solution but a transformation matrix allowing stars to be registered to a common frame. If a WCS solution is required and the coordinates of an image are approximately known, this lighter solution can be used in combination with local catalogs queries.

The algorithm I proposed starts by computing four-points asterisms, translated into unique hash codes consisting in four numbers as described in [Figure 2.5](#). These codes have the advantage of being invariant over any affine transformation. Then, if *set 1* must be registered to *set 2*, their hash codes can be produced and matched to find the transformation from *set 1* to *set 2*. The proposed implementation produces unique hash codes the same way as [Lang et al. \(2010\)](#) but using a simpler matching strategy. For all asterisms in *set 1*, the

closest asterism from *set 2* are found⁶, and all corresponding transformations computed and applied to *set 1*. The number of cross-matched stars between the transformed *set 1* and *set 2* are computed, serving as a metric to assess the accuracy of each transformation. The transformation maximizing the number of cross-matched stars between the transformed *set 1* and *set 2* is kept as the correct one. The overall method and its implementation in Python is referred to as *twirl*.




Figure 2.5: Demonstration of the four-points asterisms used in *twirl*. The left figure is reproduced from [Lang et al. \(2010\)](#) and shows how a unique hash code (x_C, y_C, x_D, y_D) is formed from four points coordinates: by taking the relative positions of two of the coordinates (C, D) with respect to the outermost ones (A and B). The right plot shows the matching of this same asterism with one from another set.

This method is used extensively throughout the alignment blocks provided by *prose*, and its Python implementation is made publicly available in the *twirl* Python package⁷.

ballet: Fast centroiding

Star positions being computed, as well as their alignment to a common reference, the caveat of the reduction strategy previously described is to rely on a centroiding algorithm to refine their centroids on every frame, which can be CPU intensive and lead to variable results depending on the method used. Inspired by [Herbel et al. \(2018\)](#), a deep convolutional neural network (CNN) was designed and trained to predict accurate centroid positions out of normalized stars cutouts.

The network architecture is shown in [Figure 2.6](#) and consists in 3 ReLu-activated convolutional layers of sizes 64, 128 and 256, each followed by a pooling layer downsampling their outputs by a factor 2. The 1028 resulting weights serve as inputs to 2 sigmoid-activated fully connected layers, finally leading to the centroid (x, y) coordinates.

⁶Thanks to a kd-tree structure, like in [Lang et al. \(2010\)](#) but using an implementation from [scipy.spatial.KDTree](#)

⁷<https://github.com/lgrcia/twirl>

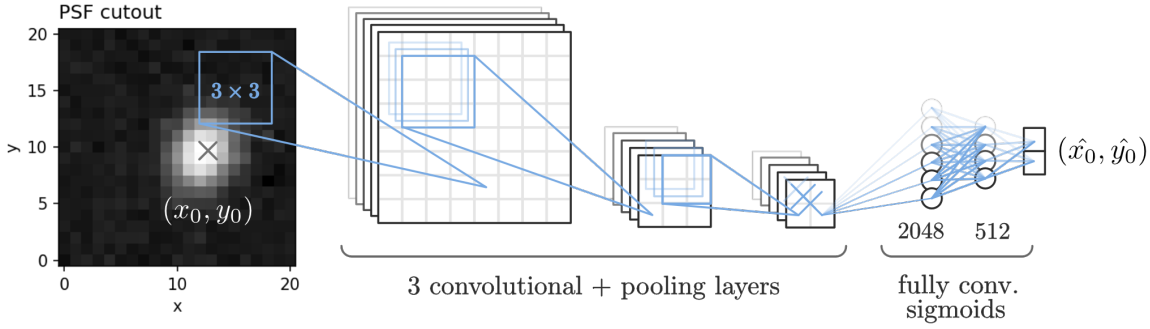


Figure 2.6: Neural network architecture

This CNN was implemented with the tensorflow Python package⁸, and trained using the Adam stochastic gradient descent algorithm (Kingma & Ba 2017) on a Huber loss function (Huber, 1964). The training samples consisted in 200 000 simulated Moffat profiles (Moffat, 1969) expressed as:

$$F(x, y) = b + \frac{1}{\left(1 + \frac{dx^2}{\sigma_x^2} + \frac{dy^2}{\sigma_y^2}\right)^\beta}$$

$$\text{where } \begin{cases} dx = (x - x_0) \cos \theta + (y - y_0) \sin \theta \\ dy = -(x - x_0) \sin \theta + (y - y_0) \cos \theta \end{cases}$$

with (x_0, y_0) the centroid position, (σ_x, σ_y) the standard deviation of the point-spread-function (PSF) in the x and y -axis, θ its relative angle with respect to the x -axis, β the concentration of the Moffat profile and b its relative background. To train the model, these parameters were randomly drawn from the distributions given in Table 2.2, representative of a wide range of PSF characteristics, including widely off-center ones.

Parameter	Distribution
x_0, y_0	$\mathcal{U}(3, 12)$
σ_x	$\mathcal{U}(2.5, 7.5)$
σ_y	$\sigma_x \cdot \mathcal{U}(0.5, 1.5)$
θ	$\mathcal{U}(0, \pi/8)$
β	$\mathcal{U}(1, 8)$
b	$\mathcal{U}(0, 0.1)$

Table 2.2: Training sample parameters distributions. $\mathcal{U}(a, b)$ is a uniform distribution bounded by a and b .

To assess the performances of the trained network, another 10 000 samples were drawn from the same distributions and their predicted centroids compared to the predictions of

⁸<https://www.tensorflow.org/>




Figure 2.7: Performance assessment of multiple centroiding algorithms against the presented model using 10 000 samples drawn from the distributions described in Table 2.2. In the right table: *accuracy* corresponds to the root-mean-square of the differences between the true centroid coordinates and their predicted values; *time* is the processing time to predict the centroids of 1000 PSF cutouts; and *robustness* is the percentage of errors under 0.1 pixels, which is linked to the occurrence of prediction outliers. This newly developed CNN is referred to as *ballet*.

three other centroiding methods. The first consisted in fitting the star cutout with a two-dimensional quadratic polynomial; the second with a fit of the marginal x and y PSF to univariate Gaussians; and the third in a complete fit of the cutout to a two-dimensional Gaussian. These are implemented in the `photutils`⁹ package (Bradley et al., 2020) under the name `centroid_quadratic`, `centroid_1dg` and `centroid_2dg`. The optimized neural network, named *ballet*, predicts centroid positions with an error of 0.015 pixels, comparable to the precision obtained from directly fitting a two-dimensional Gaussian while being two orders of magnitudes faster (see Figure 2.7). For more details about this particular application of CNNs, the interested reader is redirected to Herbel et al. (2018), in which rich references were found. Choosing and tweaking the architecture of a neural network yet remains empirical and, in this case, was the result of many manual optimizations. This pre-trained model and its `tensorflow`¹⁰ implementation is made available in the *ballet* repository¹¹, which will be complement with other CNN applications in the future. Most importantly this CNN is integrated into the `CentroidBallet` block, which makes its use straightforward in any pipeline developed with `prose`.

2.2.3 TESS follow-up comparison with AstroImageJ

The power of the pipeline described in section 2.2.2 is to rely on general-purpose blocks leading to rich sets of measurements to work on. In this section, I present a comparison of this pipeline with AstroImageJ (AIJ, Collins et al. 2017), a graphical user interface to process astronomical images into high-precision photometry. AIJ has become a reference in the field of transiting exoplanet observations and the main software for the TESS Follow-up Observing Program¹² working group (TFOP), to process ground-based transit follow-up

⁹<https://photutils.readthedocs.io/en/stable/centroids.html>

¹⁰<https://www.tensorflow.org/>

¹¹<https://github.com/lgrcia/ballet>

¹²<https://tess.mit.edu/followup/>

observations. For this reason, I will do this comparison on 26 TESS follow-up observations, all already shared within TFOP.

Observations sample and metrics

telescope	date	target	filter	type
TRAPPIST-South	2018-10-01	TOI-145	z	other
TRAPPIST-South	2018-10-18	TOI-142	I+z	transit
TRAPPIST-South	2018-11-03	TOI-169	B	transit
TRAPPIST-South	2018-11-25	TOI-207	I+z	transit
TRAPPIST-South	2018-11-28	TOI-212	z	transit
TRAPPIST-South	2018-11-29	TOI-171	B	transit
TRAPPIST-South	2018-12-14	TOI-273	I+z	other
TRAPPIST-South	2018-12-15	TOI-270	z	transit
TRAPPIST-South	2018-12-19	TOI-150	B	transit
TRAPPIST-South	2018-12-27	TOI-270	z	transit
TRAPPIST-South	2019-01-12	TOI-350	I+z	transit
TRAPPIST-South	2019-02-16	TOI-442	z	transit
TRAPPIST-South	2019-03-24	TOI-500	B	transit
TRAPPIST-South	2019-03-27	TOI-504	RC	transit
TRAPPIST-South	2019-04-17	TOI-602	z	transit
TRAPPIST-South	2019-05-13	TOI-674	I+z	transit
TRAPPIST-South	2019-05-21	TOI-674	z	transit
TRAPPIST-South	2020-11-21	TOI-370.01	z	transit
TRAPPIST-South	2020-12-02	TOI-713.02	z	transit
TRAPPIST-South	2020-12-12	TOI-2202.01	I+z	transit
TRAPPIST-South	2020-12-13	TOI-2407.01	I+z	transit
TRAPPIST-South	2020-12-19	TOI-2416.01	z	transit
TRAPPIST-South	2020-12-24	TOI-2202.01	I+z	transit
TRAPPIST-South	2021-01-07	TOI-542.01	Exo	other
TRAPPIST-South	2021-01-17	TOI-736.02	I+z	transit
TRAPPIST-South	2021-01-24	TOI-2322.01	z	other

Table 2.3: Observations used for comparison with AIJ.

The 26 observations used for this comparison are detailed in [Table 2.3](#) and came from the TRAPPIST-South telescope. They were gathered between July 2018 and March 2021. AIJ reductions and analysis were provided by several members of the TRAPPIST team and products shared within TFOP as part of the standard follow-up program. *prose* reductions were done in bulk, the only interaction from the user being the manual selection of the target star on the reference images. For each observation, I chose the reference image as the one from the middle of the night. Finally, I produced the differential light curves using [Broeg et al. \(2005\)](#).

I assessed the performances of both tools using four distinct metrics. The first was the

median standard deviation of the target light curve within 12 points bins (arbitrary), a proxy to the white noise variance. The second and third metrics were the white and red noise computed using a method inspired by Pont et al. (2006) which has the following principle: if the light curve only contains white noise, the variance of the binned light curve with bins of varying sizes would be inversely proportional to the number of points in the bins. In opposition, the variance of varying-sized bins will be independent of the bin size for a light curve with a purely correlated signal. The white and red noise can then be obtained by fitting the bin-size dependent variance of the light curve to

$$\nu(n) = \frac{\sigma_w^2}{n} + \sigma_r^2,$$

where σ_w and σ_r are the white and red noise standard deviations and n the bins size. The last metric considered was the single-transit SNR as described in Pont et al. (2006), i.e.

$$S_r = \frac{df}{\sqrt{\frac{\sigma_w^2}{n} + \sigma_r^2}}, \quad (2.1)$$

where df is the relative depth of the transit and n is the number of points in transit. Most of the light curves of the sample contained a transit signal and a meridian flip¹³ that, if being ignored, could be considered as red noise. To avoid that, the estimation of σ_r was done on light curves for which the meridian flip has been modeled and removed, simultaneously with the transit signal using orbital parameters found by peers.

Results

Figure 2.8 shows that `prose` automatically yielded light curves with lower white noise and comparable red noise, resulting in transit recovered at similar SNR on average. The measured red noise was dominated by the light curve correlation with airmass, itself driven by the choice and weights of comparison stars. This red noise also dominated the denominator of the transit SNR (Equation 3.3), that `prose` did not generally improve despite producing light curves with lower white noise. In Figure 2.9, a light curve of TOI-142 is shown, featuring the most striking difference observed between AIJ and `prose`. This difference was likely due to a non-optimal choice of aperture size and comparison stars in the manual analysis done with AIJ, illustrating the advantage of using several apertures together with an automatic differential photometry algorithm. In opposition, the same figure shows a transit of TOI-647 where similar outputs were observed between the two tools, being far more representative of this comparison.

AIJ does not require prior programming knowledge, making it a popular choice for professional astronomers as well as in the amateur astronomy community. Aperture photometry with AIJ employs a unique aperture size. In opposition, using standard blocks in `prose` photometric pipeline allows to automatically detect a larger number of stars, use a wide range of apertures and apply an automatic differential photometry algorithm, overall yielding light curves with lower white noise. Apart from alignment and centroiding aspects, AIJ and `prose` use similar algorithms (such as aperture photometry) which naturally led to comparable results. This validated the performances of `prose` default photometric pipeline against AIJ.

¹³A flip of the telescope required to track objects passing across the meridian line, required for some types of German equatorial mounts.

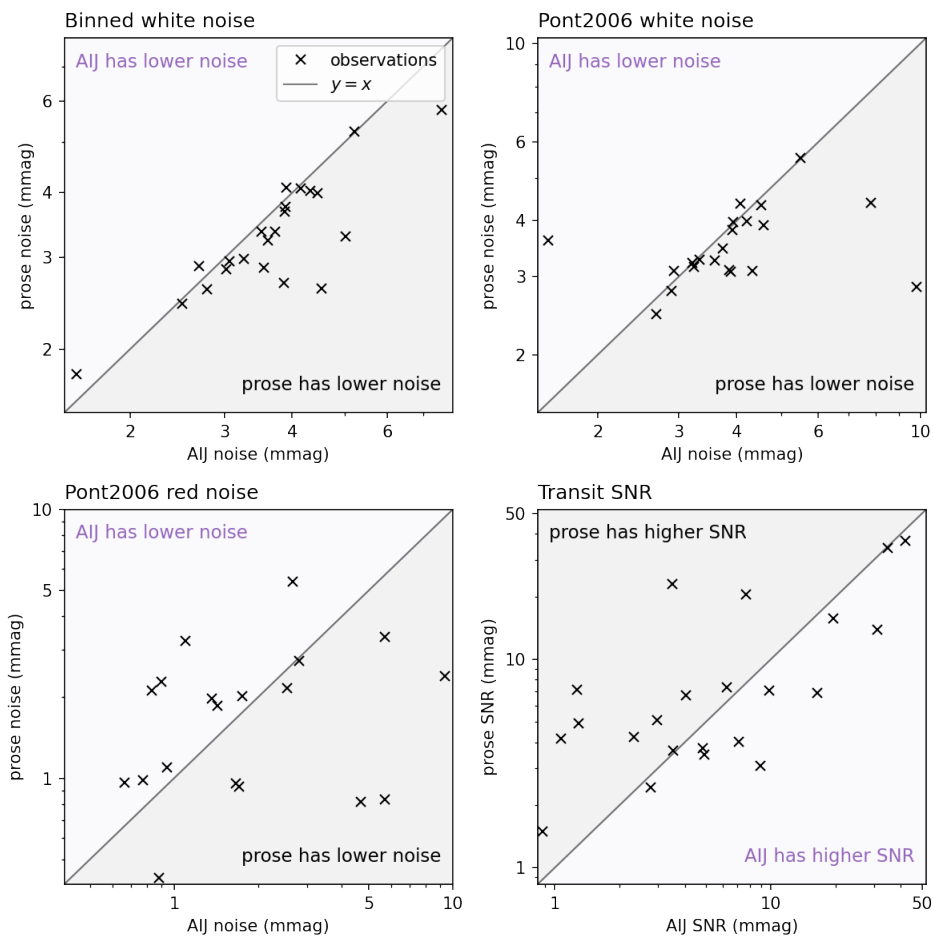


Figure 2.8: Noise and transit SNR comparison of prose against AstroImageJ (AIJ) on 26 TESS follow-up observations.




Figure 2.9: Light curves of TOI-142 (top) and TOI-647 (bottom) obtained with AstroImageJ (left) and `prose` (right).

Processing time

Comparison products from AIJ were provided by several users with no information on the time required to do their complete reduction. In addition, `prose` provides measurements on a much larger number of stars and a wide range of apertures, making the comparison with AIJ uneven. That said, `prose` pipeline was developed with domestic computers in mind, an aspect which deserved to be quantified.

Given an observation, the reduction time of `prose` photometric pipeline depends on the number of images to process, their sizes, the number of stars to extract photometry on, the number of apertures considered, their radii and of course the computer being used. Providing a figure for all these parameters is beyond the scope of this chapter, instead, a practical metric which can be easily adapted to estimate the time required to reduce a given observation can be produced. The sample for this analysis contained 2096×2080 pixels images on which I considered the 100 brightest stars using 40 apertures. I performed the reductions on a 3 Ghz Intel core i5 processor and 8Gb of RAM, taking an average of 10 *ms* per image per star, which translated into 8 minutes for a typical 500 frames observation with photometry extracted on 100 stars and 40 apertures.

Reports

TESS candidates follow-ups are meant to be reported to the TFOP community. To comply with this need, the companion Python package `verse`¹⁴ with advanced reporting capabilities was developed, in collaboration with Mathilde Timmermans. Reports produced with `verse` consist in \LaTeX templates that can be instantiated and automatically filed from Python

¹⁴<https://github.com/mathtimm/verse>

scripts. While they contain general sections, TESS-specific pages are added to provide all information related to TFOP activities. An example of such report is shown in [section 2.2.4](#).

2.2.4 Other applications

Applications of `prose` extend beyond its base photometric pipeline. Here is a list of different applications it was used for since its creation:

- The reduction of raw Hubble Space Telescope spectroscopic images of TRAPPIST-1 to produce transmission spectra of the Earth-sized outermost planet TRAPPIST-1h ([Garcia et al. 2022a, Chapter 4](#)).
- The development and deployment of an automatic pipeline for the ASTEP+ telescope (Antarctic Search for Transiting ExoPlanets, [Dransfield et al. 2022](#)).
- Development and deployment of at least two automatic pipelines for separate amateur projects targeting the observation and discovery of new variable stars (private communications).
- The publishing of more than an hundred exoplanet follow-up reports to the TESS Follow-up Observing Program¹⁵ (through the work of colleagues Khalid Barkaoui and Mathilde Timmermans)
- Overall, the reduction of images from 12 different observatories, without modification of the base photometric pipeline presented in this section, thanks to `prose` instrument-agnostic design.
- As of April 2023, the reduction of several observations used in a few dozen exoplanets discovery papers (published in peer-reviewed journals, e.g. [Schanche et al. 2022](#); [Delrez et al. 2022a](#))

In addition of these past achievements, current activities from outside the field of exoplanetary science are ongoing:

- Development of a semi-automatic pipeline for the comet observations of the TRAPPIST telescopes (Mathieu Vander Donckt).
- Development of a pipeline to monitor satellite pollution of scientific observations at the ESO Paranal Observatory (Chile, Sebastián Zúñiga-Fernández in collaboration with ESO).
- At least two similar projects led by PhD students, one initiated by a private company (private communications).
- Development of solutions to track low earth orbit satellites (private communications).

By adopting the modular approach offered by `prose`, its development team hopes to see the deployment of transparent pipelines which can live beyond individuals projects, serving the entire community and leading to reproducible results.


`prose` is open-source under the MIT license and publicly developed under a version controlled environment¹⁶.

¹⁵<https://tess.mit.edu/followup/>

¹⁶<https://github.com/lgrcia/prose>

Appendix

Figures 2.10 and 2.11 show elements of a TESS follow-up report produced with prose using \LaTeX . While the first *Summary* page is automatically generated, other pages are pre-filled and can be freely complemented with figures and text.



TIC id	251855940
Time	02:32 - 06:49 [4h16]
RA - DEC	8.65967 -32.16653
Images	119
GAIA id	2317155751208922112
Mean std - fwhm (epsf)	1.53 - 3.59 pixels
Fwhm - fwhm (target)	3.18 - 3.75 pixels
Optimum aperture	5.51 pixels
Telescope	Trappist-South
Filter	z
Exposure	120.0 s










Figure 2.10: **verse** reports start with a *Summary* page, which intends to provide a quick look to an observation and its main products. The left side of this page features a stack image with the detected stars, among which are highlighted the target star (in blue) and comparison stars (yellow) used to build the differential light curves. At the bottom left, a cutout around the target as well as a radial PSF is plotted and the corresponding aperture overlaid. The rest of the page displays the raw and differential flux of the target star, as well as external parameters time-series (e.g. airmass) and comparison stars' light curves.

TRANSIT MODEL



MODEL PARAMETERS

Parameters	Model	TESS
[u1,u2]	[0.3508,0.2223]	-
[R*]	[1.05486 R _☉]	1.05486 R _☉
[M*]	[1.04 M _☉]	1.04 M _☉
P	4.046 ± 0.0 d	4.046292 ± 0.000365 d
Rp	18.607 ± 0.342 R _⊕	18.266975 ± 0.382044 R _⊕
Tc	2459175.706 ± 0.001 BJD-TDB	2459175.7061 BJD-TDB
b	0.532 ± 0.033	-
Duration	183.17 min	180 ± 18
(Rp/R*) ²	2.61e-2	-
Apparent depth	2.89e-02	2.83e-2
a/R*	10.26	-
i	87.01°	-
SNR	40.65	30.06
RMS per bin (7.2 min)	1.24e-03	-

49

Figure 2.11: The *Transit model* report page shows a model of the transit light curve together with the inferred transit parameters (here a corner plot has been added manually). This model was inferred using the `exoplanet` Python package (Foreman-Mackey et al., 2021) and the corner plot done with `corner` (Foreman-Mackey, 2016).

Chapter 3

Detection of exoplanets around M-dwarfs

A chapter based on a paper in preparation including contributions from Daniel Foreman-Mackey and Francisco F. Pozuelos.

Although transit signals can be seen in the apparent flux received from transiting exoplanets' host stars, they are often mixed with other astrophysical and instrumental signals. If they can be disentangled from nuisance signals, transits offer a powerful way to detect exoplanets. However, disentangling these signals comes with many challenges, especially for M-dwarfs, usually observed at lower SNR and more likely to display photometric variability (e.g. [Murray et al. 2020](#)). In this dissertation, the term *correlated noise* encapsulates all nuisance signals that impact the search for transits, even if having an astrophysical origin.

Widely used transit-search algorithms (BLS types, [Kovács et al. 2002](#)) are capable of detecting transits in light curves containing only transit signals and white noise. Hence, the simplest way to find periodic transit signals is to first clean a light curve from nuisance signals before performing the search. This strategy is widely adopted by the community, both using physically-motivated systematic models like [Luger et al. \(2016, 2018\)](#), or filtering techniques ([Jenkins et al. 2010](#), [Hippke et al. 2019](#)). However, when correlated noise starts resembling transits, this cleaning step (often referred to as *detrending*) is believed to degrade their detectability. In this case, the only alternative to search for transits is to perform a full-fledged modeling of the light curve, including both transits and correlated noise, and to compute the likelihood of the data to the transit model on a wide parameter space, an approach largely avoided due to its intractable nature. Nonetheless, [Kovács et al. \(2016\)](#) ask: *Periodic transit and variability search with simultaneous systematics filtering: Is it worth it?*. The latter study discards the benefit of using a full-fledged approach in the general case, but fails at exploring the light-curves characteristics for which it becomes necessary.

While it might only represent a handful of systems, transits hidden in correlated noise are extremely valuable for the exoplanetary science community. Indeed, correlated noise associated with stellar variability may originate from surface active regions, that can be probed with the help of planetary transits. A better understanding of these structures benefit both the study of stellar atmospheres and their concerning impact on planetary atmosphere retrievals

(see [Chapter 4](#) and [Chapter 5](#)).

In this chapter I present **nuance**, an algorithm using linear models and Gaussian processes to simultaneously search for transits while modeling correlated noise in a tractable way.

Transit light curve simulations

In order to study the effect of correlated noise on transit search, this chapter relies on transit light curve simulations including a stochastically driven model of stellar variability. The following describes how such signals are modeled.




Figure 3.1: Example dataset sampled at $N = 2880$ times corresponding to an observation of 4 days with an exposure time of 2 minutes. The mean of this signal consists in a periodic transit signal of period $P = 0.7$ days, duration $D = 0.05$ days and depth of 2% (\mathcal{T} in gray) plus instrumental signals (\mathcal{S} in blue). Correlated noise in the form of stellar variability is simulated by modeling the covariance matrix of the signal with a Gaussian process (\mathcal{V} in red) including a diagonal variance of 0.001^2 corresponding to white noise. This simulated signal is not intended to be physically realistic.

Let f be the simulated flux of a star sampled and arranged in the $(1 \times N)$ column-vector \mathbf{f} associated to the column-vector of times \mathbf{t} . \mathbf{f} is such that

$$\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}),$$

i.e. that \mathbf{f} is drawn from a Gaussian process (GP) of mean $\boldsymbol{\mu}$ and covariance matrix \mathbf{C} (see [section 1.4](#) for a definition of GPs). In this equation, $\boldsymbol{\mu}$ is such that its i -th element is defined by $\mu_i = \mathcal{T}(t_i) + \mathcal{S}(t_i)$ where t_i is the i -th time of observation, \mathcal{T} is a periodic transit

function and \mathcal{S} a function describing the instrumental part of the signal (both described below). The covariance matrix \mathbf{C} is built such that $C_{i,j} = \nu(t_i, t_j)$ where ν is a covariance function accounting for correlated noise in the form of stellar variability with added white noise. An example of such signal is simulated and shown in [Figure 3.1](#).

Transit signal \mathcal{T}

The periodic transit signal \mathcal{T} is simulated using the simple model described in [Protopapas et al. \(2005\)](#), where a transit of period P , epoch T_0 , duration D and unitary depth observed at time t is given by

$$\mathcal{T}_c(t, P, T_0, D) = \frac{1}{2} \tanh\left(c\left(\theta - \frac{1}{2}\right)\right) - \frac{1}{2} \tanh\left(c\left(\theta + \frac{1}{2}\right)\right), \quad (3.1)$$

$$\text{with } \theta = \frac{P}{\pi D} \sin\left(\frac{\pi(t - T_0)}{P}\right),$$

where the dimensionless parameter c controls the roundness of the transit depth ($c \gg 1$ corresponding to a box-shaped transit as shown in [Figure 3.2](#)). This analytical model is fully empirical but easily differentiable.



Figure 3.2: Simulations of a single transit signal ([Equation 3.1](#)) shown for different values of c .

In this chapter, all transits are simulated with $c = 12$, a value arbitrarily chosen that can be fined-tuned in real applications using the limb-darkening coefficients of a given star. The periodic transit signal \mathcal{T} seen in [Figure 3.1](#) corresponds to $\mathcal{T} = 0.02 \times \mathcal{T}_{c=12}(t, P = 0.7, T_0 = 0.2, D = 0.05)$, all parameters in unit of days.

Instrumental signals \mathcal{S}

Instrumental signals are simulated as a linear model of M explanatory variables arranged in the $(M \times N)$ design matrix \mathbf{X} . Hence,

$$\mathcal{S} = \mathbf{w}\mathbf{X},$$

where the vector \mathbf{w} contains the linear coefficients of the model. The simulated flux shown in [Figure 3.1](#) contains a linear model where the $M = 4$ columns of the design matrix \mathbf{X} are given by $\mathbf{X}_i = t^i$ (i.e. \mathbf{X} is the Vandermonde matrix order 3 of time t) and $\mathbf{w} = [1.0 \ 0.0005 \ -0.0002 \ -0.0005]$.

Stellar variability ν

As this chapter focuses on stellar variability and its effect on transit detection, I employed a physically-motivated GP kernel, describing stellar variability through the covariance of

a stochastically-driven damped harmonic oscillator (SHO, [Foreman-Mackey et al. 2017](#); [Foreman-Mackey 2018](#)) taking the form

$$k(\tau) = \sigma^2 \exp\left(-\frac{\omega \tau}{2Q}\right) \begin{cases} 1 + \omega \tau & \text{for } Q = 1/2 \\ \cosh(f \omega \tau / 2Q) + \sinh(f \omega \tau / 2Q) / f & \text{for } Q < 1/2 \\ \cos(g \omega \tau / 2Q) + \sin(g \omega \tau / 2Q) / g & \text{for } Q > 1/2 \end{cases} \quad (3.2)$$

where $\tau = |t_i - t_j|$, $f = \sqrt{1 - 4Q^2}$ and $g = \sqrt{4Q^2 - 1}$

where Q is the quality factor of the oscillator, ω its pulsation and σ the amplitude of the kernel function. GP computations in this chapter use the implementation from `tinygp`¹, a Python package exposing the quasi-separable kernels from [Foreman-Mackey \(2018\)](#) and powered by `JAX`². The stellar variability signal in [Figure 3.1](#) has been sampled from a GP with an SHO kernel of parameters $\omega = \pi/6D$ (i.e. a period equal to 12 times the duration D of the simulated transit), $Q = 45$ and $\sigma = 0.02$, the depth of the simulated transit. An extra term $\sigma_f^2 = 0.001^2$ is added to the diagonal of the covariance matrix, corresponding to the variance of the simulated measurement f and leading to the white noise observed in [Figure 3.1](#).

Using these simulated signals, with an absolute control on the model parameters, I study in the next section the impact of correlated noise and its detrending on transits detectability.

3.1 The issue with correlated noise and its detrending

Two sources of correlated noise particularly justify the need for a detrending step before searching for transits: instrumental noises (such as telescope pointing errors) and stellar variability (induced by pulsations or starspots).

3.1.1 The effect of correlated noise on transits detectability

To study transits detectability, I focused on the signal-to-noise (SNR) of a unique event, reduced to the simplified expression ([Pont et al. 2006](#), Equation 12):

$$SNR = \frac{\Delta}{\sqrt{\frac{\sigma_w^2}{n} + \frac{\sigma_c^2}{N_{tr}}}} \quad (3.3)$$

where Δ is the relative transit depth, n is the number of points within transit, N_{tr} the number of transits ($N_{tr} = 1$ here since we consider a single transit), and σ_w^2 and σ_c^2 are the white and correlated noise variances. To show the impact of correlated noise on transit detectability, I simulated a unique transit signal (using [Equation 3.1](#)) and computed its SNR using [Equation 3.3](#), both in the presence and absence of correlated noise ([Figure 3.3](#)).

¹<https://github.com/dfm/tinygp>

²<https://github.com/google/jax>




Figure 3.3: Illustration of the effect of correlated noise on a single transit SNR. I simulated a 1-hour transit of depth 1% on top of white noise (variance of 0.0015^2) as part of a 24-hours observation with an exposure time of 1 minute (top). Then, in the bottom plot, correlated noise was added to the transit signal and simulated using a GP with a Matérn-3/2 kernel of scale 1 hour and sigma of 0.2%. The SNR on the right of each light curve is computed using Equation 3.3, where the white and red noise standard deviations are computed as in subsection 2.2.3.

As illustrated in Figure 3.3, the presence of correlated noise strongly decreased the transit signal SNR, limiting its detection. This issue naturally motivated the development of systematic detrending algorithms such as the Trend Filtering Algorithm (TFA, Kovács et al. 2005, in its primary use case), SYSREM (Tamuz et al. 2005), or Pixel Level Decorrelation (PLD, Deming et al. 2015; see also EVEREST from Luger et al. (2016, 2018)). Most of these methods rely on the shared nature of instrumental signals among light-curves (or neighboring pixels) such that the correction applied should not degrade the transit signal. Except for TFA, these algorithms were mostly applied to space-based continuous observations, that provided continuous stellar baselines and mostly reproducible systematic signals. This is not the case for the vast majority of sparse ground-based observations, in addition subject to periodic daytime interruptions and varying atmospheric extinction.

3.1.2 The effect of detrending on transits detectability

Instrumental signals often have the benefit to be shared among light curves of stars observed with the same instrument, strongly correlated with measurements from the experimental setup (like detector’s temperature, pointing error, sky background or airmass time series). While not being true in general, I first assumed that the detrending of the systematic signals in light curves based on their incomplete modelling, one that ignores transit signals (because unknowns), do not affect the search for transits. In opposition, stellar variability and other astrophysical correlated noise are generally unknown and harder to correlate with simultaneous measurements. This gave rise to several treatments in order to reconstruct and detrend stellar variability. One of them is physically-motivated and makes use of Gaussian processes (e.g. Aigrain et al. 2016). Another is empirical and makes use of filtering algorithms (Jenk-

ins et al. 2010, Hippke et al. 2019). In this section, I study the effect of both approaches on transit detectability, by estimating the degradation of a unique transit SNR for a wide variety of stellar variability characteristics.

To study this effect, I added to a transit signal correlated noise in the form of stellar variability, featuring different timescales and amplitudes. In case *a* (purple in Figure 3.4), I detrended the resulting light curves using the widely-adopted Tukey’s bi-weight filtering method presented in Mosteller & Tukey (1977) and its implementation from wotan³ (Hippke et al., 2019). Presented as an optimal value in the latter study, I used an optimal window size of three times the transit duration. In case *b* (red in Figure 3.4) I employed the same GP model used to simulate stellar variability to reconstruct and remove correlated noise from the light curves. In both cases, such detrending ignores the presence of potential transits (the case of unknown transits search). Once light curves were cleaned from correlated noise, I computed the remaining transit SNR using Equation 3.3.

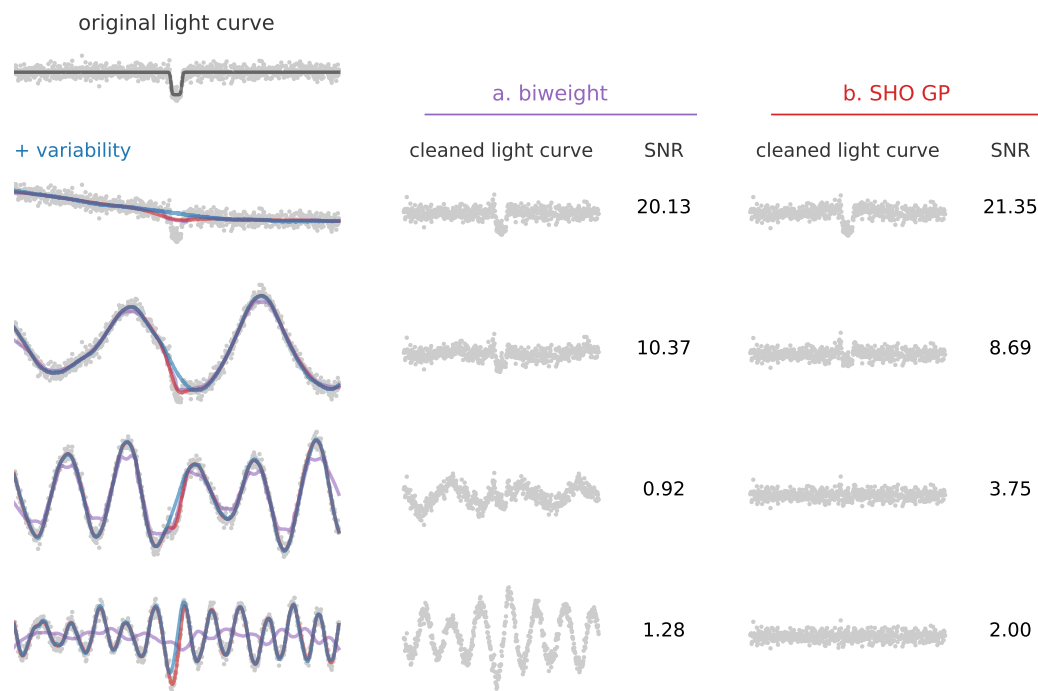


Figure 3.4: (Top left) I added a transit signal with a depth of 0.8% and a duration of 0.1 days to a 1000 points pure white noise signal of variance 0.0015^2 , representing an observation spanning 3 days. (Left) Correlated noise in the form of stellar variability with different timescales and amplitudes were added to the signal. This signal was then fitted using a bi-weight filter (Hippke et al., 2019, in purple) and a GP with an SHO kernel (in red) with the same hyperparameters used to generate the data. I then subtracted the fitted variability (central and left plots) and computed the transit SNR using Equation 3.3, with the transit depth estimated as the minimum in-transit flux.

Figure 3.4 clearly shows the effect of both detrending techniques on transits SNR, and intu-

³<https://github.com/hippke/wotan>

itively suggests that this degradation is strongly dependant on the correlated noise characteristics encountered.

In order to explore the parameter space for which detrending is the most problematic, I simulated 10 000 light curves containing correlated noise in the form of stellar variability that I detrended using a bi-weight filter. For convenience, the simulated light curves shared a common transit added on top of white noise, and a variability signal only defined by two parameters: τ , the relative timescale of the variability with respect to the transit duration; and δ , the relative amplitude of the variability against transit depth, i.e.

$$\tau \propto \frac{\text{variability timescale}}{\text{transit duration}} \quad \text{and} \quad \delta \propto \frac{\text{variability amplitude}}{\text{transit depth}}. \quad (3.4)$$

To follow this parametrization, I used a GP with a SHO kernel with the following hyperparameters:

$$\omega = \frac{\pi}{\tau D}, \quad \sigma = \delta \frac{\Delta}{2} \quad \text{and} \quad Q = 10, \quad (3.5)$$

with $\Delta = 1\%$ and $D = 0.04$ days, the depth and duration of the simulated transit. For $\tau = 1$ and $\delta = 1$, the expressions of ω and σ given in Equation 3.5 correspond to a variability signal with a period half that of the transit duration, and a standard deviation two times that of the transit depth, i.e. strongly resembling the simulated transit signal. As in Figure 3.4, I reconstructed the variability using a bi-weight filter with an optimal window length of $3 \times D$ (Hippke et al., 2019), that I subtracted from each light-curve. I then computed the resulting transit SNR, using Equation 3.3, where the transit depth was taken as the remaining in-transit minimum flux. Figure 3.5 shows the remaining SNR values of transits computed this way, after each variability signal with random (τ, δ) have been detrended.

Figure 3.5 shows that it exists an entire region of the (τ, δ) parameter space for which the bi-weight detrending degrades transit SNR to the point of no detection ($SNR < 6$). Hence, bi-weight filter detrending makes transit-search blind to many systems. While this study should be extended to other detrending techniques, it highlights the need for a more informed transit search algorithm able to deal with correlated noise.




Figure 3.5: SNR of a unique transit after detrending light curves using an optimal bi-weight filter. All light curves correspond to a 2.8 days observation with a cadence of 2 minutes, and contain a unique transit of duration 1 hour and a depth of 1%, added on top of white noise with a standard deviation of 0.5% plus stellar variability. Light curves at the top and right side of the central plot are shown with their corresponding τ and δ values. The amplitude of the variability increases with δ and the timescale of the variability increases with τ .

3.2 nuance

nuance is an algorithm capable of searching for planetary transits in light curves containing correlated noise, such as instrumental signals and stellar photometric variability.

Let assume that the flux f of a star is observed and arranged in the $(1 \times N)$ column-vector \mathbf{f} , associated to the column-vector of times \mathbf{t} . This flux contains instrumental signals, stellar variability and a periodic transit signal that we wish to uncover (see [Figure 3.1](#)). I assume that a set of M observed measurements (such as the position of the star on the detector, the FWHM of the PSF or the sky background) taken at the same time as the flux can be treated as explanatory variables for \mathbf{f} . These measurements are arranged in the $(M \times N)$ matrix \mathbf{X} , later called *design matrix*.

Ideally, we would detect the periodic transit signal in this flux by sampling the posterior likelihood of this data to a full-fledge model including stellar variability (more generally correlated noise), instrumental systematic signals (modeled with explanatory variables), and a periodic transit signal of period P , epoch T_0 , duration D and depth Δ . We would then

reduce the posterior likelihood to $p(\mathbf{f}|P)$, its marginalized version over all parameters except the period P , producing a transit search periodogram $\mathcal{Q}(P)$. However, this approach has two issues: It is highly intractable, and it may lead to multimodal distributions that are hard to interpret.

Given a period P , we instead want to compute the likelihood of a periodic transit signal at the maximum likelihood parameters \hat{T}_0 , \hat{D} and $\hat{\Delta}$, i.e the periodogram

$$\mathcal{Q}(P) = p(\mathbf{f}|P, \hat{T}_0, \hat{D}, \hat{\Delta}) \quad (3.6)$$

This is done by adopting the strategy of [Foreman-Mackey et al. \(2015\)](#), and separate the transit search into two components: the linear search and the periodic search. During the linear search, the likelihood of a single non-periodic transit is computed for a grid of epochs, durations and depths. Then, the periodic search consists in combining these likelihoods to compute the likelihood of the data given a periodic transit signal for a range of periods. These combined likelihoods yield a transit-search periodogram on which the periodic transit detection is based. `nuance` differs from [Foreman-Mackey et al. \(2015\)](#) and other existing transit search algorithms as it models the covariance of the light curve with a GP, accounting for correlated noise (especially in the form of stellar variability) while keeping the model linear and tractable. This way, `nuance` searches for transits while, at the same time, modeling correlated noise, avoiding the detrending step that degrades transit signals SNR.

The approach employed by `nuance` (a two-step approach proposed by [Foreman-Mackey et al. \(2015\)](#)), shares similarities with the approach of [Jenkins et al. \(2010\)](#), where a single event statistic is computed and combined into a multiple event statistics.

3.2.1 The linear search

During the linear search, the goal is to compute the likelihood $p(\mathbf{f}|T, D, \Delta)$ of the data given a single non-periodic transit signal of epoch T , duration D and depth Δ , for a grid of epochs, durations and depths.

To account for correlated noise, the light curve f is modeled as being drawn from a GP, like the simulated light curves described in the introduction of this chapter, such that

$$\mathbf{f} \sim \mathcal{N}(\mathbf{w}\mathbf{X}, \Sigma),$$

with mean $\mathbf{w}\mathbf{X}$ (i.e. a linear model of the M explanatory variables with coefficients \mathbf{w}) and covariance Σ . To account for the presence of a single non-periodic transit of epoch T and duration D , this signal is computed and appended as the last column of the design matrix \mathbf{X} , using the simple transit model from [Protopapas et al. \(2005\)](#) with a unitary depth ([Equation 3.1](#)). This way, the transit signal is part of the linear model and its depth Δ can be solved linearly. Under this assumption, the log-likelihood of the data given a single non-periodic transit is ([Rasmussen & Williams, 2005](#))

$$\ln p(\mathbf{f}|I) = -\frac{1}{2}(\mathbf{f} - \mathbf{w}\mathbf{X})^T \Sigma^{-1}(\mathbf{f} - \mathbf{w}\mathbf{X}) - \frac{1}{2} \ln |\Sigma| - \frac{N}{2} \ln 2\pi, \quad (3.7)$$

where the parameters vector \mathbf{w} and their errors σ are computed using the generalized least-square solution

$$\mathbf{w} = (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{X}^T \Sigma^{-1} \mathbf{f} \quad \text{and} \quad \sigma = (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1}, \quad (3.8)$$

with Σ the covariance matrix modeled using the GP. Hence, $\ln p(\mathbf{f}|I)$ can be computed on a grid of epochs and durations, the transit depth being linearly solved for any (T, D) . In this equation. This linear search leads to the set of likelihoods

$$\{\ln \mathcal{L}_{i,j}\}_{i,j} = \{\ln p(\mathbf{f}|T_i, D_j, \Delta_{i,j})\}_{i,j},$$

where $\Delta_{i,j}$ is the depth linearly solved for a given (T_i, D_j) ⁴. Figure 3.6 shows this likelihood grid computed for a simulated dataset, using the same GP and design matrix \mathbf{X} used to simulate the data.




Figure 3.6: Principle and output of the linear search. The simulated dataset (top) corresponds to the one shown and described in Figure 3.1. First, a set of durations and depths $\{T_i, D_j\}_{i,j}$ is generated. For each pair of indices (i, j) , the likelihood $\ln p(\mathbf{f}|T_i, D_j, \Delta_{i,j})$ is computed using the parameters from Equation 3.8 and the expression of Equation 3.7. This process yields the grid of log-likelihoods $\ln \mathcal{L}$ (bottom plot), as well as the $\{\Delta_{i,j}, \sigma_{i,j}\}_{i,j}$ transit depths and errors inferred linearly using Equation 3.8.

To prepare for the next step, the corresponding depths $\Delta_{i,j}$ linearly solved for any (T_i, D_j) are stored, as well as their associated uncertainties $\sigma_{i,j}$, corresponding to

$$\Delta_{i,j} = \mathbf{w}_M \quad \text{and} \quad \sigma_{i,j} = \sigma_{MM},$$

M denoting the index of the last column of the design matrix \mathbf{X} , where the transit signal is contained.

3.2.2 The periodic search

We then need to combine the likelihoods computed from the linear search to obtain

$$p(\mathbf{f}|P, T_0, D, \Delta),$$

⁴The expression of $\{\ln \mathcal{L}_{i,j}\}_{i,j}$ omits the vector $\mathbf{w}_{i,j}$ (except its last value $\Delta_{i,j}$) as it is also linearly solved for any given value of (T_i, D_j) and irrelevant in what follows.

i.e. the probability of a periodic transit of period P , epoch T_0 , duration D and depth Δ given the data \mathbf{f} .

For a given transit duration D , any combination of (P, T_0) leads to K transits, for which it is tempting to write

$$p(\mathbf{f}|P, T_0, D, \Delta) = \prod_k^K p(\mathbf{f}|T_k, D, \Delta_k), \quad (3.9)$$

where $\{T_k\}_k$ are the epochs matching (T_0, P) and $\{\Delta_k\}_k$ the corresponding depths. So that

$$\ln p(\mathbf{f}|P, T_0, D, \Delta) = \sum_k^K \ln \mathcal{L}_k.$$

This is the joint likelihood of transits belonging to the same periodic signal but with varying depths $\{\Delta_k\}_k$. However, individual transits from a periodic signal cannot be considered independent, and should instead be found periodically and share a common transit depth Δ . Let demonstrate that there is an analytical expression for the joint likelihood of K individual transits with depths and errors $\{\Delta_k, \sigma_k\}_k$ assuming a common depth Δ , corresponding to

$$\begin{aligned} \ln p(\mathbf{f}|P, T_0, D, \Delta) &= \sum_k^K \ln \mathcal{L}_k - \frac{1}{2} \sum_k^K \left(\ln(\sigma_k^2) - \ln(\sigma^2 + \sigma_k^2) + \frac{(\Delta_k - \Delta)^2}{\sigma_k^2 + \sigma^2} \right) \\ \text{with } \frac{1}{\sigma^2} &= \sum_k^K \frac{1}{\sigma_k^2} \quad \text{and} \quad \Delta = \sigma^2 \sum_k^K \frac{\Delta_k}{\sigma_k^2}. \end{aligned} \quad (3.10)$$

Proof

From the linear search, we retain and index by k the parameters of the K individual transits whose epochs $\{T_k\}_k$ are compatible with a periodic signal of period P and epoch T_0 . From the likelihoods of these transits (computed in [subsection 3.2.1](#)), we want an expression for

$$p(\mathbf{f}|P, T_0, D, \Delta) = \prod_{k \in \mathbb{T}} p(\mathbf{f}|T_k, D, \Delta),$$

i.e., given a depth D , the likelihood of the data given a periodic transit signal of period P , epoch T_0 and a common depth Δ . Since only $\{p(\mathbf{f}|T_k, D, \Delta_k)\}_k$ is known (i.e. transits with different depths), we decompose

$$p(\mathbf{f}|T_k, D, \Delta) = \int p(\mathbf{f}|T_k, D, \tilde{\Delta}) p(\tilde{\Delta}|\Delta) d\tilde{\Delta}, \quad (3.11)$$

where $p(\mathbf{f}|T_k, D, \tilde{\Delta})$ is the probability of the k -th transit to have a depth $\tilde{\Delta}$ and $p(\tilde{\Delta}|\Delta)$ the probability to observe the depth $\tilde{\Delta}$ knowing the existence of a common depth Δ . In other words, [Equation 3.11](#) involves the likelihood of the non-periodic transit k to be part of a periodic transit signal with a common depth Δ .

Since each depth Δ_k is found through generalized least square, each follow a normal distribution $\mathcal{N}(\Delta_k, \sigma_k^2)$, centered on Δ_k with variance σ_k^2 and an amplitude \mathcal{L}_k , leading to the likelihood function

$$p(\mathbf{f}|T_k, D, \tilde{\Delta}) = \mathcal{L}_k \exp\left(-\frac{(\tilde{\Delta} - \Delta_k)^2}{2\sigma_k^2}\right).$$

As for the common transit depth Δ , it can be estimated through the joint probability of all other transit depths than Δ_k , such that

$$\Delta \sim \prod_{i \neq k}^K \mathcal{N}(\Delta_i, \sigma_i^2),$$

with

$$\frac{1}{\sigma^2} = \sum_{i \neq k}^K \frac{1}{\sigma_i^2} \quad \text{and} \quad \Delta = \sigma^2 \sum_{i \neq k}^K \frac{\Delta_i}{\sigma_i^2}. \quad (3.12)$$

Hence

$$p(\tilde{\Delta}|\Delta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\tilde{\Delta} - \Delta)^2}{2\sigma^2}\right).$$

We can now rewrite [Equation 3.11](#) as

$$p(\mathbf{f}|T_k, D, \Delta) = \frac{\mathcal{L}_k}{\sqrt{2\pi\sigma^2}} \int \exp\left(-\frac{(\tilde{\Delta} - \Delta_k)^2}{2\sigma_k^2}\right) \exp\left(-\frac{(\tilde{\Delta} - \Delta)^2}{2\sigma^2}\right) d\tilde{\Delta}.$$

The integral in this equation is a product of gaussian integrals that can be obtained analytically, leading to

$$p(\mathbf{f}|T_k, D, \Delta) = \mathcal{L}_k \sqrt{\frac{\sigma_k^2}{\sigma^2 + \sigma_k^2}} \exp\left(-\frac{1}{2} \frac{(\Delta_k - \Delta)^2}{\sigma_k^2 + \sigma^2}\right).$$

Finally,

$$\ln p(\mathbf{f}|P, T_0, D, \Delta) = \sum_k^K \ln \mathcal{L}_k - \frac{1}{2} \sum_k^K \left(\ln(\sigma_k^2) - \ln(\sigma^2 + \sigma_k^2) + \frac{(\Delta_k - \Delta)^2}{\sigma_k^2 + \sigma^2} \right), \quad (3.13)$$

the log-likelihood of the data given a periodic transit signal of period P , epoch T_0 , duration D and common depth Δ . In order to reduce the number of times [Equation 3.12](#) is computed, we adopt the biased estimates

$$\frac{1}{\sigma^2} = \sum_k^K \frac{1}{\sigma_i^2} \quad \text{and} \quad \Delta = \sigma^2 \sum_k^K \frac{\Delta_i}{\sigma_i^2}, \quad (3.14)$$

so that Δ and σ are independent of k in the last sum of [Equation 3.13](#).

* * *

While [Equation 3.13](#) takes a closed form, the individual epochs matching T_0 and P are not necessarily available in the grid of epochs $\{T_k\}_k$. In [Foreman-Mackey et al. \(2015\)](#), a similar issue is solved by using the nearest neighbors in the epochs grid. Instead, to allow the efficient matrix computation of [Equation 3.13](#), the likelihood grid is linearly interpolated from $\{T_i\}_i$ to a common grid of transit phases $\{\phi_i\}_i$, leading to the periodic search log-likelihood

$$\ln \mathcal{P}(P) = \{ \ln p(\mathbf{f}|P, \phi, D) \}_{i,j}$$

shown for few periods in [Figure 3.7](#) (b). In the latter equation, $\Delta_{i,j}$ is omitted since being interpolated from the linear search using ϕ_i , D_j and $T_0 = 0$.

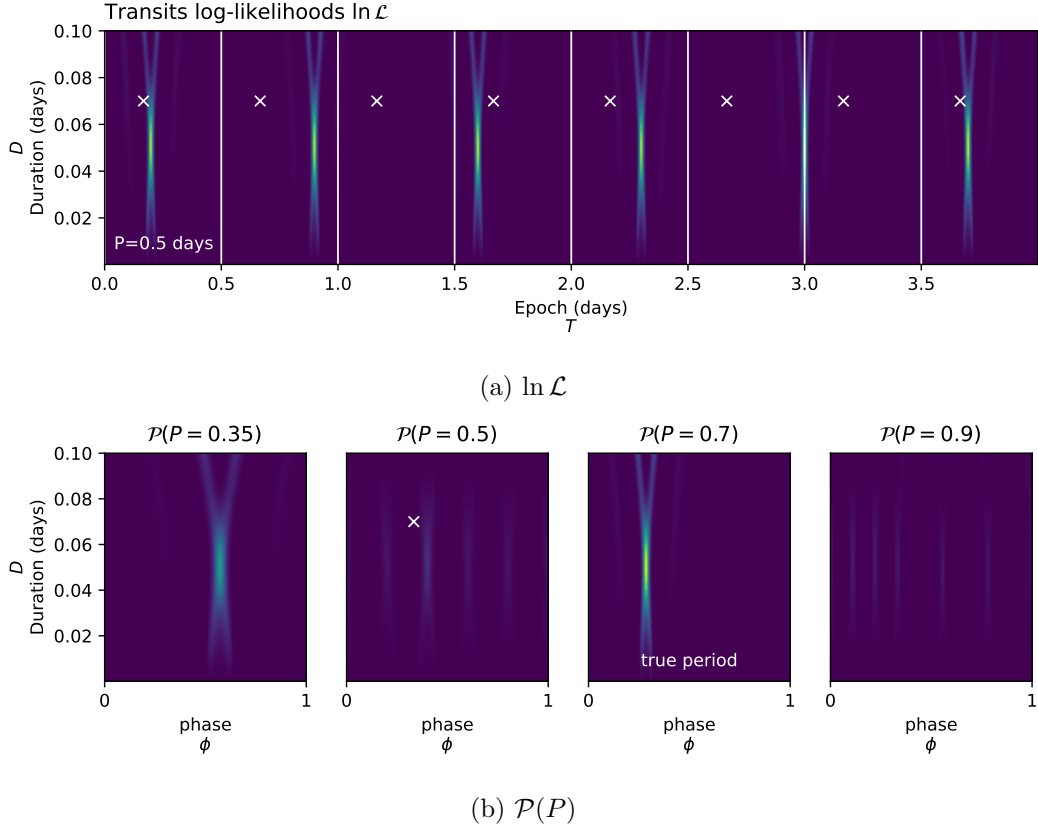


Figure 3.7: Applied to the dataset shown [Figure 3.6](#), this figure shows how the periodic search works at different periods P (including the true period $P = 0.7$ days). Given different periods P and $T_0 = 0$, the likelihood $\ln \mathcal{L}$ shown in (a) is phase-folded and interpolated onto a common grid of phases shown in (b). As an example, the white lines in (a) mark the edges of each fold for a period of $P = 0.5$ days, and the white crosses show the epochs $\{T_k\}_{k \in \mathbb{K}}$ matching a particular phase in the grid (reported in (a) and (b)) on which the corresponding $\{\ln \mathcal{L}_k\}_{k \in \mathbb{K}}$ are interpolated and combined using [Equation 3.13](#). To allow the use of efficient matrix computations, this is done for all durations $\{D_i\}_i$ so that \mathcal{P} is computed on the full grid $\{D_i, \phi_i\}_{i,j}$ at once. We understand from the folded likelihoods plots in (b) that a different choice of epoch T_0 may only shift the results in phase but do not affect the values of \mathcal{P} . For this reason, computing \mathcal{P} for $T_0 = 0$ is sufficient. We also notice how the maximum value of $\mathcal{P}(P = 0.7/2)$ (left plot of (b)) is lower than for $P = 0.7$ days, a result of combining the log-likelihoods using [Equation 3.13](#) instead of [Equation 3.9](#), in favor of individual transits matching a common depth Δ .

3.2.3 The transit search periodogram

Using [Equation 3.13](#), we can now compute $\ln \mathcal{P}$ for a range of periods and build a transit search periodogram using [Equation 3.6](#). But a final issue emerges, one that is fundamentally linked to our strategy. Each likelihood $p(\mathbf{f}|T, D, \Delta)$ estimated during the linear search is computed using N measurements. Hence, combining transits in the periodic search, through Δ_k , σ_k and the product of K likelihoods $\{\mathcal{L}_k\}_k$ (see [Equation 3.13](#)), artificially leads to a

likelihood involving $N \times K$ measurements. This leads to a normalization issue when trying to compare the joint log-likelihoods $\mathcal{P}(P)$ from one period to another, as the number of observed transits differs from one period to another. This motivates a final step to produce the transit search periodogram \mathcal{Q} . For any period P , instead of taking $\mathcal{Q}(P)$ as the maximum value of $\ln \mathcal{P}$, the maximum likelihood parameters

$$(\phi_0, D) = \arg \max_{\phi_i, D_j} \{ \ln p(\mathbf{f}|P, \phi_i, D_j) \}_{i,j} \quad (3.15)$$

are retrieved, and $\mathcal{Q}(P)$ is defined as the SNR of the transit of period P , epoch $T_0 = \phi_0 P$, duration D and depth Δ , i.e.

$$\mathcal{Q}(P) = \frac{\Delta}{\sigma},$$

where Δ and σ are obtained using Equation 3.14 with the last column of X containing a periodic transit signal of period P , epoch T_0 , duration D and depth 1. This process and the resulting periodogram \mathcal{Q} are shown in Figure 3.8.

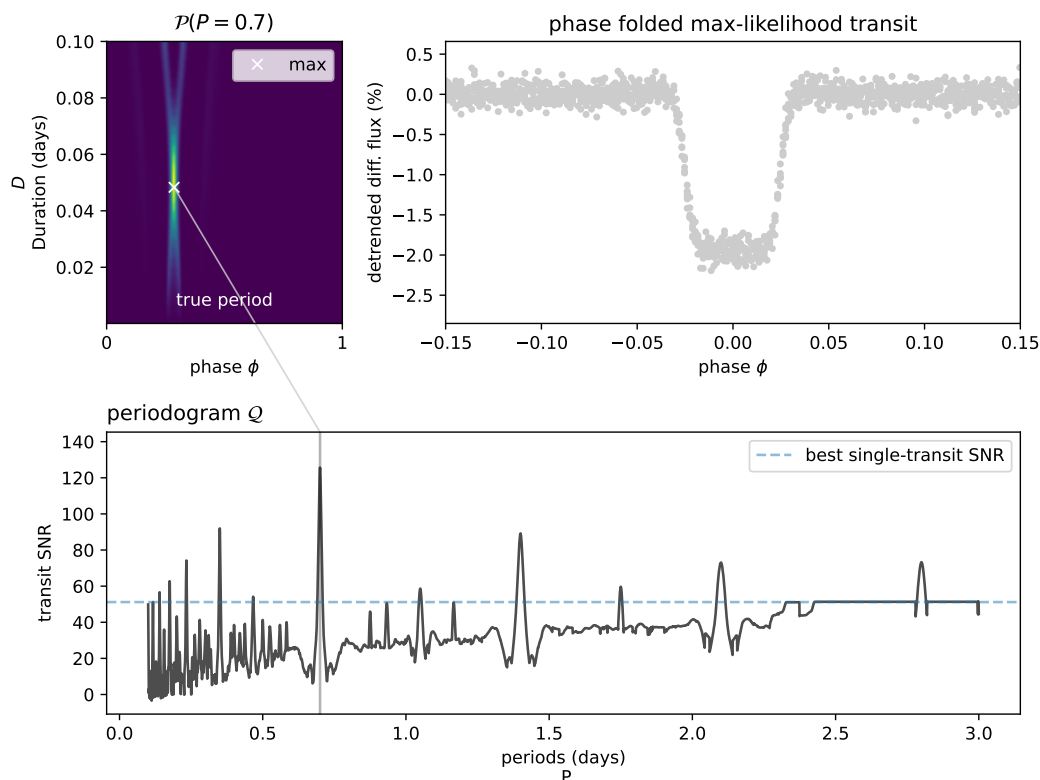


Figure 3.8: For each period P , the joint likelihood $\mathcal{P}(P)$ is computed using Equation 3.13, and the value of the maximum likelihood transit SNR retained as $\mathcal{Q}(P)$.

The periodic transit of period P with the maximum SNR, i.e. maximizing \mathcal{Q} , is adopted as the best candidate, basing the confidence in this signal through its SNR. The parameters of this transit are the period P , the epoch $T_0 = \phi_0 P$ and duration D (Equation 3.15), and the depth Δ with error σ (given by Equation 3.14).

3.2.4 An open-source python package

The methods described in this chapter are implemented in the `nuance` open-source Python package, hosted on Github⁵ and released on the Python Package Index (PyPI)⁶.

To instantiate a search, a user can start by creating a `Nuance` object with

```
from nuance import Nuance

nu = Nuance(time, flux, gp=gp, X=X)
```

where `gp` is a `tinygp` GP instance and `X` the design matrix of the linear model. `nuance` exploits the use of `tinygp`⁷, a Python package powered by `JAX`⁸, allowing for custom kernels to be built and highly tractable computations. We can then define a set of epochs `t0s` and durations `Ds` and run the linear search with

```
import numpy as np

t0s = time.copy()
# a range of 10 durations
Ds = np.linspace(0.01, 0.2, 10)
nu.linear_search(t0s, Ds)
```

Finally, the periodic search is run with

```
# a wide range of periods
periods = np.linspace(0.1, 5, 2000)
search = nu.periodic_search(periods)
```

From this `search` object, the best transiting candidate parameters can be computed (`search.best`), or the Q periodogram retrieved (`search.Q_snr`), together with valuable information about the transit search. The `Nuance` object also provides methods to perform transit search on light curves from multi-planetary systems, the advantage of `nuance` being that the linear search only needs to be performed once, and reused for the search of several transiting candidates. An extensive and maintained online documentation is provided at nuance.readthedocs.io.

3.3 Control test against BLS and TLS

To start testing `nuance` against existing methods, a simple light curve is simulated, consisting in pure white noise with a standard deviation of 5×10^{-4} (relative flux), observed for a duration of 6 days with a cadence of 2 minutes. In this signal, transits with periods ranging from 0.3 to 2.5 days are injected, with a unique transit duration of 50 minutes and depths ranging from 2.6×10^{-4} to 6.8×10^{-4} (using the simple model from [Protopapas et al. \(2005\)](#) presented in the introduction). By design, these injected signals have an SNR ranging from 5 to 30 (with $\sigma_r = 0$ in [Equation 3.3](#)). For each light curve, transits are searched using three different methods:

⁵<https://github.com/lgrcia/nuance>

⁶<https://pypi.org/project/nuance/>

⁷<https://github.com/dfm/tinygp>

⁸<https://github.com/google/jax>

- `nuance`, using its implementation from the Python package described in the previous section.
- BLS, the box-least-square algorithm (Kovács et al., 2002), using the `astropy BoxLeastSquares` implementation⁹.
- TLS, the transit-least-square algorithm (Hippke & Heller, 2019), an optimized version of the BLS algorithm implemented in the `transitleastsquares` Python package¹⁰. For this study, a modified version of the `transitleastsquares` package is employed, allowing to search for transits using a custom range of periods and durations.

For all of these methods, 3000 trial periods from 0.2 to 2.6 days are searched, with a single trial duration fixed to the unique known duration of 50 minutes. A transit signal is considered detected if the absolute difference between the injected and the recovered period is less than 0.01 day. To ease the detection criteria, transit periods recovered at half or twice the injected period (aka *aliases*) are considered as being detected. For this reason, detected transit epochs are not considered. Results from this *injection-recovery* procedure are shown in Figure 3.9.

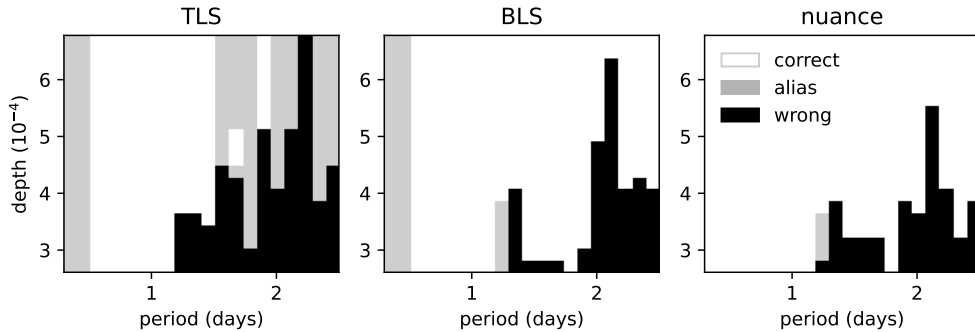


Figure 3.9: Injection-recovery of 20×20 individual transit signals in a flat light curve with only white noise (whose characteristics are described in the main text), using TLS, BLS and `nuance`. Transits recovered with the correct injected period correspond to the white area of the plot, while transits recovered at half or twice the injected period are shown in gray (aliases). Transit signals that were not detected are shown in black. Note that for each period and depth, a single transit epoch is used, explaining the sharp edges and discontinuities in the plot.

These results show that there is no perfect match between the detection capabilities of `nuance` compared to BLS and TLS. However, `nuance` and BLS seems to have comparable performances, compared to TLS that is less performant for lower SNR transits and more sensitive to period aliases (i.e. half or twice the injected periods). At periods lower than 0.5 days, there seem to be an advantage for `nuance` to detect transits with correct periods, compared to BLS and TLS detecting periods aliases.

Explaining why such differences are observed, especially between BLS and TLS is beyond the scope of this chapter, especially since BLS compares particularly well with `nuance`. Hence, in the following sections, `nuance` will be compared to methods employing BLS on light curves featuring correlated noise.

⁹<https://docs.astropy.org/en/stable/api/astropy.timeseries.BoxLeastSquares.html>

¹⁰<https://transitleastsquares.readthedocs.io>

3.4 nuance vs. biweight+BLS on simulated light curves

The [subsection 3.1.2](#) highlighted that `nuance`'s full-fledged modeling capabilities may not always be necessary and may only be beneficial for certain noise characteristics, relative to the searched transit parameters. As a result, this section focuses on evaluating the performance of `nuance` in the relative parameter space (τ, δ) described in [Equation 3.5](#) and see when it becomes necessary.

In this section, `nuance` is compared to the approach that involves removing stellar variability from the light curves before performing the search on the detrended datasets. Similar to [subsection 3.1.2](#), an optimal bi-weight filter implemented in the `wotan` Python package¹¹ is used with a window size three times that of the injected transit duration. The transit search is then performed using the BLS algorithm, like in [section 3.3](#) using the `astropy` `BoxLeastSquares` implementation¹². This transit search strategy is widely used in the community and is denoted `biweight+BLS`. In what follows, the transit detection criteria are the same as the ones used in [section 3.3](#), i.e. that a transit is considered recovered if the absolute difference between the injected and recovered period is less than 0.01 days, including periods found at half or twice the injected ones (aliases).

3.4.1 Dataset

The dataset consists in 4000 light curves simulated using the model described in the introduction of this chapter. We simulate a common periodic transit added to all light curves, of period $P = 1.1$ days, epoch $T_0 = 0.2$ days, duration $D = 0.04$ days and depth $\Delta = 1\%$. Each light curve consist in a 4 days observation with an exposure time of 2 minutes, leading to $N = 2880$ data points with a normal error of 0.1%.

For a given pair of (τ, δ) , we simulate stellar variability using a GP with an SHO kernel of hyperparameters defined by [Equation 3.5](#) (except for Q , described below), computed with respect to the injected transit parameters D and Δ . The same kernel is used for the search with `nuance`, an optimal choice on equal footing with the optimal $3 \times D$ window size of the bi-weight filter employed in the `biweight+BLS` search. 4000 pairs of (τ, δ) are generated such that

$$\tau \sim \mathcal{U}(0.1, 10), \quad \delta \sim \mathcal{U}(0.1, 25) \quad \text{and} \quad Q \sim \mathcal{U}(10, 100)$$

where $\mathcal{U}(a, b)$ denotes a uniform distribution of lower bound a and upper bound b .

3.4.2 Results

The results of this injection-recovery procedure are shown in [Figure 3.10](#) and highlight particularly well the benefit of `nuance` against the `biweight+BLS` strategy on transits with relatively low depth compared to the stellar variability amplitude, and a relatively small duration compared to the stellar variability period. This empirical statement only concerns light curves with a given amount of white noise, and may vary depending on the length of the observing window or the number of transits. For this reason, quantifying for which values of (τ, δ) `nuance` outperforms `biweight+BLS` would only apply to this specific example. While more

¹¹<https://github.com/hippke/wotan>

¹²<https://docs.astropy.org/en/stable/api/astropy.timeseries.BoxLeastSquares.html>

examples are considered in the following [section 3.5](#), proper guidelines on when to use `nuance` instead of more classic detrending techniques are currently being developed.




Figure 3.10: Injection-recovery results on 4000 simulated light curves described in [subsection 3.4.1](#). The color scale represents the fraction of recovered transits.

This injection-recovery was done in a particularly optimal setup, on light curves not all physically realistic, hence demonstrating the performance of `nuance` only on a purely experimental basis. In the next section, we perform transits injection-recovery on real space-based light curves.

3.5 Validation on TESS light curves

The parameter space (τ, δ) introduced in [subsection 3.1.2](#) is very convenient from a signal processing point of view, but lacks physical reality. Very low values of τ and δ often translate into non-physical exoplanets properties as well as unrealistic stellar variability, with a strong dependence on the host star properties.

In order to assess the performance of `nuance` on real light curves, its transit recovery capabilities are rather evaluated in the planetary (period, radius) parameter space, more often encountered in the literature (e.g. Figure 12 of [Delrez et al. 2022b](#)). In what follows, transits described by the simple model of [Protopapas et al. \(2005\)](#) are injected into light curves from the Transiting Exoplanet Survey Satellite (TESS, [Ricker et al. 2015](#)). The TESS mission, currently in its 5th year of operation, is a space-based all-sky survey searching for transiting exoplanets around more than 9 million stars¹³. As this chapter focuses on the detection of planets transiting variable M-dwarfs, targets are chosen from a list of ~ 600 M-dwarfs found to have detectable rotation signals with periods lower than one day ([Ramsay et al., 2020](#)). For each of the studied target, transits are injected and recovered in the TESS 2 min cadence SPOC Simple Aperture Photometry and Pre-search Data Conditioning light curves (PDCSAP, [Caldwell et al. 2020](#)) on a single sector (the first being observed for each target) spanning an average of 27 days.

¹³<https://exofop.ipac.caltech.edu/tess/search.php>

In the following sections, `nuance` is compared to other transit search strategies employing the BLS algorithm, each time after detrending the light curves with three different techniques:

- `bspline+BLS` employs a B-spline¹⁴ detrending, fitted using the `scipy.interpolate.splprep` function¹⁵, followed by a search with the BLS algorithm.
- `biweight+BLS` employs an optimal bi-weight filter implemented in the `wotan` Python package (Hippke et al., 2019) with an optimal window size of $3 \times D$ (i.e. three times the transit duration), followed by a search with the BLS algorithm.
- `harmonics+BLS` employs a linear harmonic detrending, where the light curve is modeled as a Fourier series including four harmonics of the stellar rotation period found by Ramsay et al. (2020), with coefficients found through ordinary least square. This detrending is followed by a search with the BLS algorithm.

Stellar variability kernel

A GP kernel with a single SHO kernel is not representative of realistic stellar variability. Instead, a mixture of two SHO kernels of period P_* and $P_*/2$ (with P_* the rotation period of the star), being representative of a wide range of stochastic variability in stellar time series from rotation to pulsations¹⁶, is adopted, with the addition of one short scale and one long term exponential term. This full rotation kernel is expressed as

$$k = k_1 + k_2 + k_3 + k_4$$

with

- k_1 a SHO kernel with hyperparameters

$$Q_1 = 1/2 + Q_0 + \delta Q, \quad \omega_1 = \frac{4\pi Q_1}{P\sqrt{4Q_1^2 - 1}} \quad \text{and} \quad S_1 = \frac{\sigma^2}{(1+f)\omega_1 Q_1}.$$

- k_2 a SHO kernel with hyperparameters

$$Q_2 = 1/2 + Q_0, \quad \omega_2 = 2\omega_1 \quad \text{and} \quad S_2 = \frac{f\sigma^2}{(1+f)\omega_2 Q_2},$$

where Q_0 is the quality factor for the secondary oscillation, δQ is the difference between the quality factors of the first and the second modes, f is the fractional amplitude of the secondary mode compared to the primary and σ is the standard deviation of the process. The kernels k_3 and k_4 are expressed as

$$k(t, t') = \sigma^2 \exp\left(-\frac{|t - t'|}{\ell}\right),$$

with ℓ and σ the scale and standard deviation of the process. In total, the rotation kernel k has 8 hyperparameters.

¹⁴<https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.BSpline.html>

¹⁵<https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.splprep.html>

¹⁶<https://celerite2.readthedocs.io/en/latest/api/python>

Light curve cleaning

As the described light curve model does not account for stellar flares, the light curve of each target is cleaned using an iterative sigma clipping approach. For each iteration, points 4 times above the standard deviation of the full light curve (previously subtracted by its median) are identified. Then, the 30 adjacent points on each side of the found outliers are masked. This way, large flare signals and their expected ingress and egress are masked, using a total of 3 iterations. As PDCSAP light curves often start with a ramp-like signal, the first 300 points (as well as the last 300 points) of each continuous observation are masked. Finally, each light curve is normalized by its median value. Such a cleaned light curve is shown in Figure 3.11.




Figure 3.11: Cleaned single-sector light curve of the target TIC 153193474. The light curve is normalized by its median value.

GP hyperparameters optimization

The hyperparameters of the employed rotation kernel are optimized on cleaned light curves, using the `scipy.optimize.minimize`¹⁷ wrapper provided by the `jaxopt` Python package, taking advantage of the JAX implementation of `tinygp`. As red noise is expected to bias the light curve uncertainty estimates performed by SPOC, the diagonal of the full covariance matrix of the data (i.e. their uncertainty, assuming homoscedasticity) is held free, increasing the number of optimized parameters to 9. The optimization is performed using the BFGS algorithm (Fletcher, 1987), minimizing the negative log-likelihood of the data as expressed in Equation 3.7 (without transit), i.e. accounting for a linear systematic model of the data in addition of stellar variability. For simplicity, and to adopt a uniform treatment for all target light curves, a design matrix \mathbf{X} with a single constant column is adopted, such that the systematic model only consists in a single parameter corresponding to the mean value of the differential flux (expected to be close to 1) solved linearly.

Injected transits

For a given orbital period and planetary radius, the corresponding transit duration and depth are computed and a random transit epoch is generated. These parameters are then used to generate a periodic transit signal using the simple model of Protopapas et al. (2005), finally added to the cleaned target light curve. For each target light curve, 10 unique orbital periods combined with 10 unique planetary radii are individually considered, for a total of 100 transit light curves. Orbital periods P are sampled on a regular grid between 0.4 and 10 days, and planetary radii R_p are sampled on a regular grid designed to lead a minimum transit SNR of 4 and a maximum of 30, using Equation 3.3 with $\sigma_r = 0$, σ_w equal to the mean uncertainty

¹⁷https://jax.readthedocs.io/en/latest/_autosummary/jax.scipy.optimize.minimize.html

estimated by the SPOC pipeline and the depth δ found using the radius of the star reported by Ramsay et al. (2020).

Detection criteria

Because we also consider planets recovered with half or twice the real period, we ease the detection criteria by ignoring the match of the recovered transit epoch with the real epoch. Note that a visual vetting ensured that the found epochs were consistent with the ones injected for cases where the planet was considered detected.

3.5.1 Injection-recovery studies

TIC 140212114: fast sinusoidal variability

TIC 140212114 has an effective temperature of $3595 \pm 157 \text{ K}$ ¹⁸ (M1.5-type star, Rajpurohit et al. 2013) with a rotation period of 0.14 days (Ramsay et al., 2020). For each technique, the left plot of Figure 3.12 shows a portion of TIC 140212114 cleaned light curve with an injected transit signal (hardly distinguishable due to correlated noise of the light curve compared to the transit depth). On the left plots, the black lines correspond to the trends computed for this injected light curve subtracted from it before being searched for transits. This detrending only applies to techniques other than nuance, where instead the black line corresponds to the mean of the GP model. For each technique, the right plot of Figure 3.12 shows the results of the transit search for each of the 10×10 injected light curves, reported in the planetary (period, radius) parameter space. A black square denotes that the transit signal has not been detected, a gray square denotes a signal detected at an alias period ($P/2$ or $2P$), and a white square denotes a transit signal detected with the correct period. On the right plots of Figure 3.12, secondary axes show the values of the injected transits in the (τ, δ) relative parameter space.

Results from Figure 3.12 are discussed in subsection 3.5.2.

¹⁸<https://exofop.ipac.caltech.edu/tess/target.php?id=140212114>

TIC 140212114 - sector 5

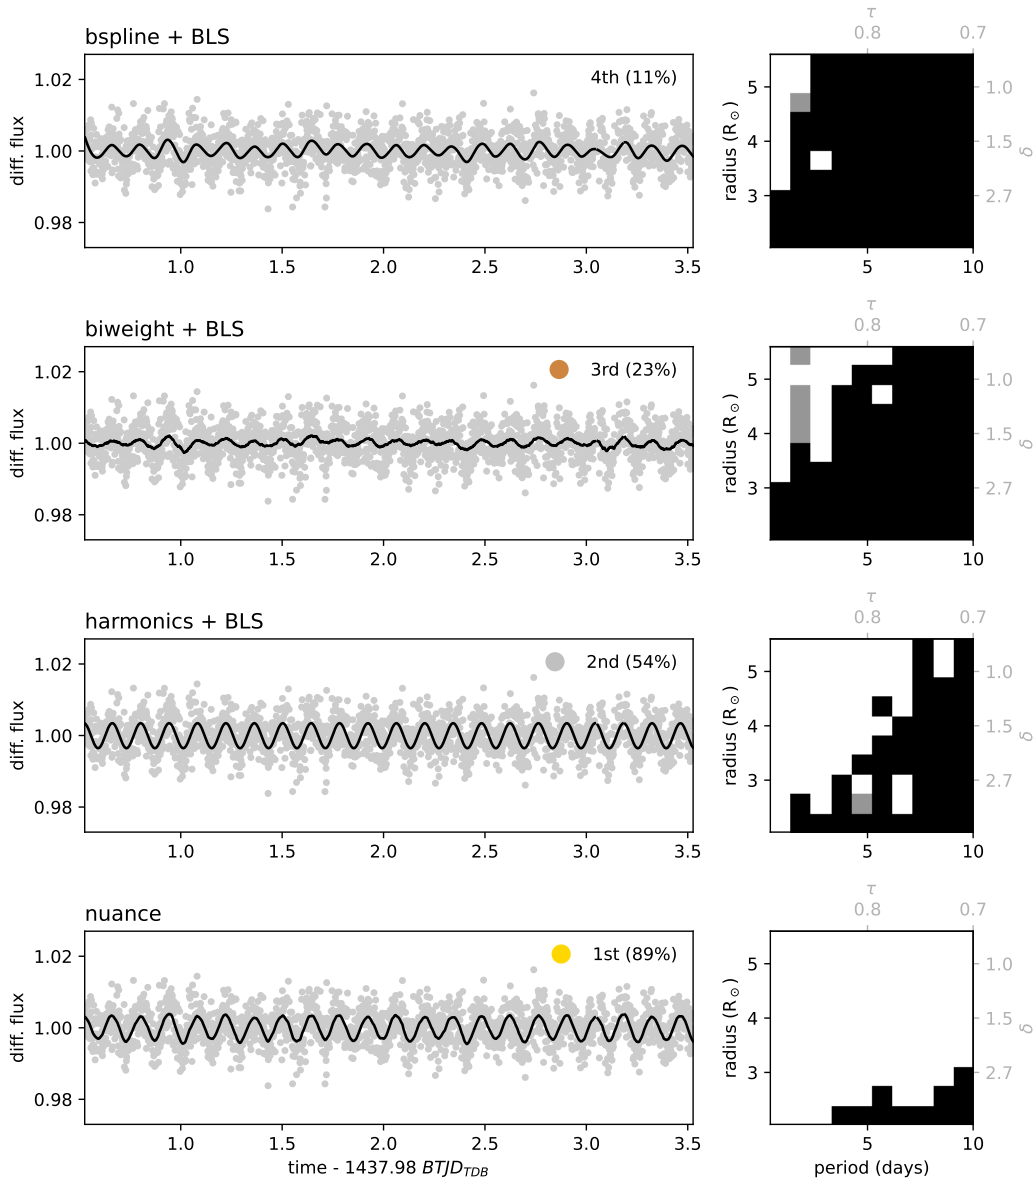


Figure 3.12: Results of the transits injection-recovery on TIC 140212114 single-sector light curve. Left: cleaned light curve with computed trend overplotted in black (except for **nuance** where it corresponds to the mean of the GP model). Right: Results of the transit search where a black square denotes a transit signal not detected, gray a signal detected at an alias period ($P/2$ or $2P$), and white a signal detected with the correct period. On the right plots, secondary axes show the (τ, δ) relative parameter space. For each method, the upper right legend on the left plot indicates its ranking based on the percent of recovered transit signals (where a transit with an aliased period counts as being detected). The **harmonics+TLS** method is the best after **nuance**, but with 35% less transits detected in total.

TIC 452793374: high-amplitude red noise

TIC 452793374 has an effective temperature of $3377 \pm 157 \text{ K}$ ¹⁹ (M2.5-type star, [Rajpurohit et al. 2013](#)) with a rotation period of 0.33 days ([Ramsay et al., 2020](#)). Injection-recovery results are shown in [Figure 3.13](#) and discussed in [subsection 3.5.2](#).

TIC 153193474: non-sinusoidal variability

TIC 153193474 has an effective temperature of $3351 \pm 157 \text{ K}$ ²⁰ (M2.5-type star, [Rajpurohit et al. 2013](#)) with a rotation period of 0.11 days ([Ramsay et al., 2020](#)). Injection-recovery results are shown in [Figure 3.14](#) and discussed in [subsection 3.5.2](#).

TIC 306331621: slow sinusoidal variability

TIC 306331621 has an effective temperature of $3067 \pm 157 \text{ K}$ ²¹ (M5-type star, [Rajpurohit et al. 2013](#)) with a rotation period of 1.00 days ([Ramsay et al., 2020](#)). It was chosen for its relatively slow rotation (compared to the three other targets) and higher white noise compared to the amplitude of its variability. Injection-recovery results are shown in [Figure 3.15](#) and discussed in [subsection 3.5.2](#).

3.5.2 Results

From the injection-recovery conducted on TIC 140212114, TIC 452793374, TIC 153193474 and TIC 306331621, it can be concluded that searching for transits directly with *nuance* offers a clear advantage over any of the presented detrending methods combined with BLS. While not fully described in this dissertation, injection-recovery on additional targets lead to same conclusion.

Targets with light curves featuring stable sinusoidal stellar variability, i.e. with a high quality factor Q (see the SHO kernel), are the ones that benefit less from *nuance*, as a set of harmonic signals can be combined linearly to robustly model and detrend the periodic signal. In the cases of TIC 452793374 and TIC 306331621, an additional $\sim 6\%$ of planets are detected by *nuance* over the probed planetary (period, radius) parameter space. From these two examples, it cannot be concluded if this difference is solely due to the difference observed between *nuance* and BLS in [section 3.3](#). While red noise was properly detrended by a bi-weight filter in the case of TIC 452793374, additional tests must be done to show that *nuance* clearly outperforms this method when any red noise is present in the light curve (as partially observed in [section 3.4](#)).

Finally, targets with non-sinusoidal or short scale variability are the ones that benefit the most from *nuance*, as no other method is able to robustly detrend their light curves without removing transit signals. This is well illustrated by the study of TIC 140212114 ([Figure 3.12](#)) and TIC 153193474 ([Figure 3.14](#)).

¹⁹<https://exofop.ipac.caltech.edu/tess/target.php?id=452793374>

²⁰<https://exofop.ipac.caltech.edu/tess/target.php?id=27734884>

²¹<https://exofop.ipac.caltech.edu/tess/target.php?id=306331621>

TIC 452793374 - sector 5

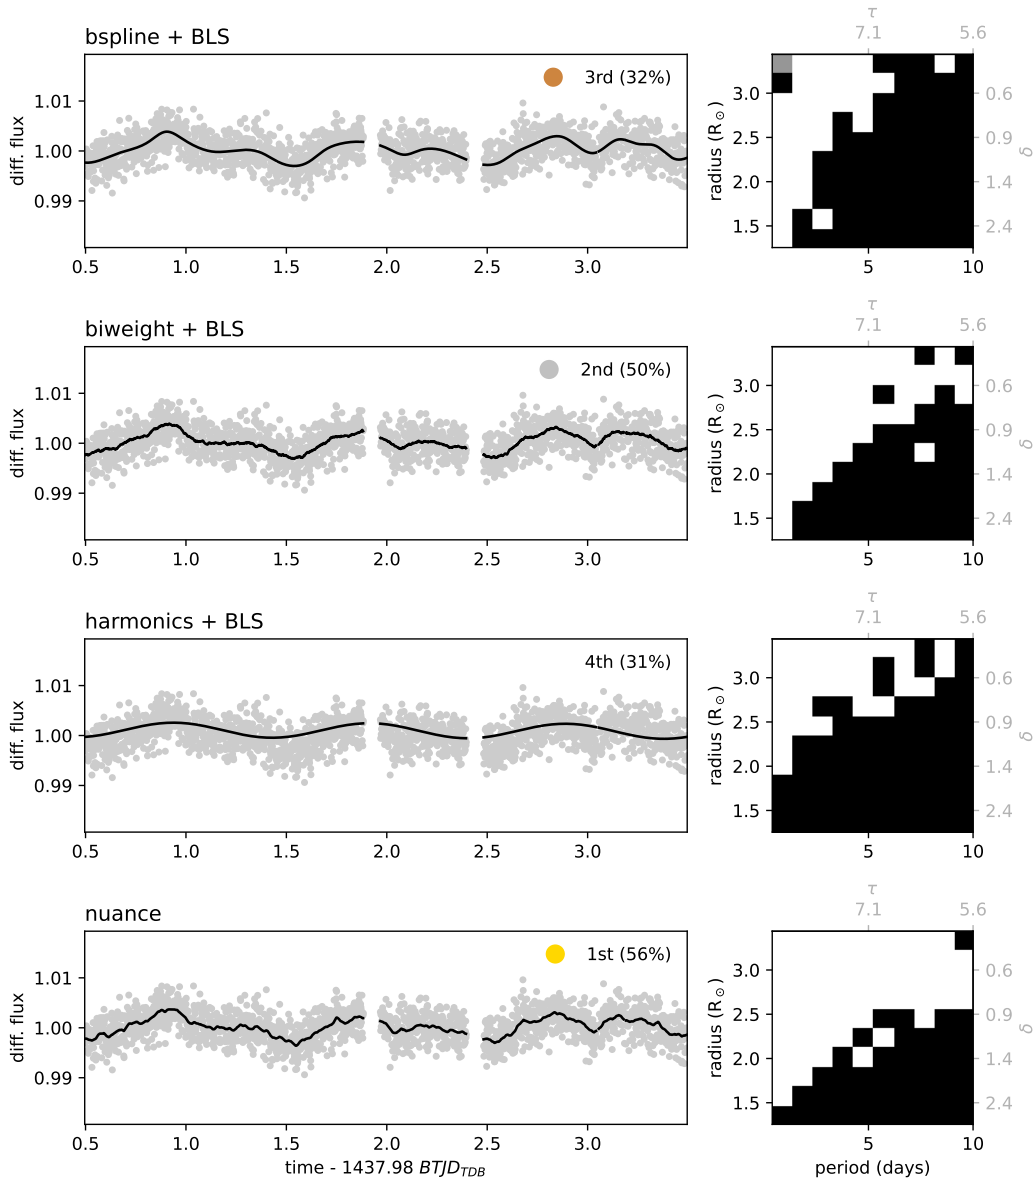


Figure 3.13: Results of the transits injection-recovery on TIC 452793374 single-sector light curve. See [section 3.5.1](#) and legend of [Figure 3.12](#) for a detailed description of this figure. The **biweight+TLS** method is the best after nuance, with 6% less transits detected in total that could be due to the sole difference between nuance and BLS discussed in [section 3.3](#). It indicates that a more flexible filter is necessary for this light curve with significantly higher red noise, making nuance not necessary.

TIC 153193474 - sector 7

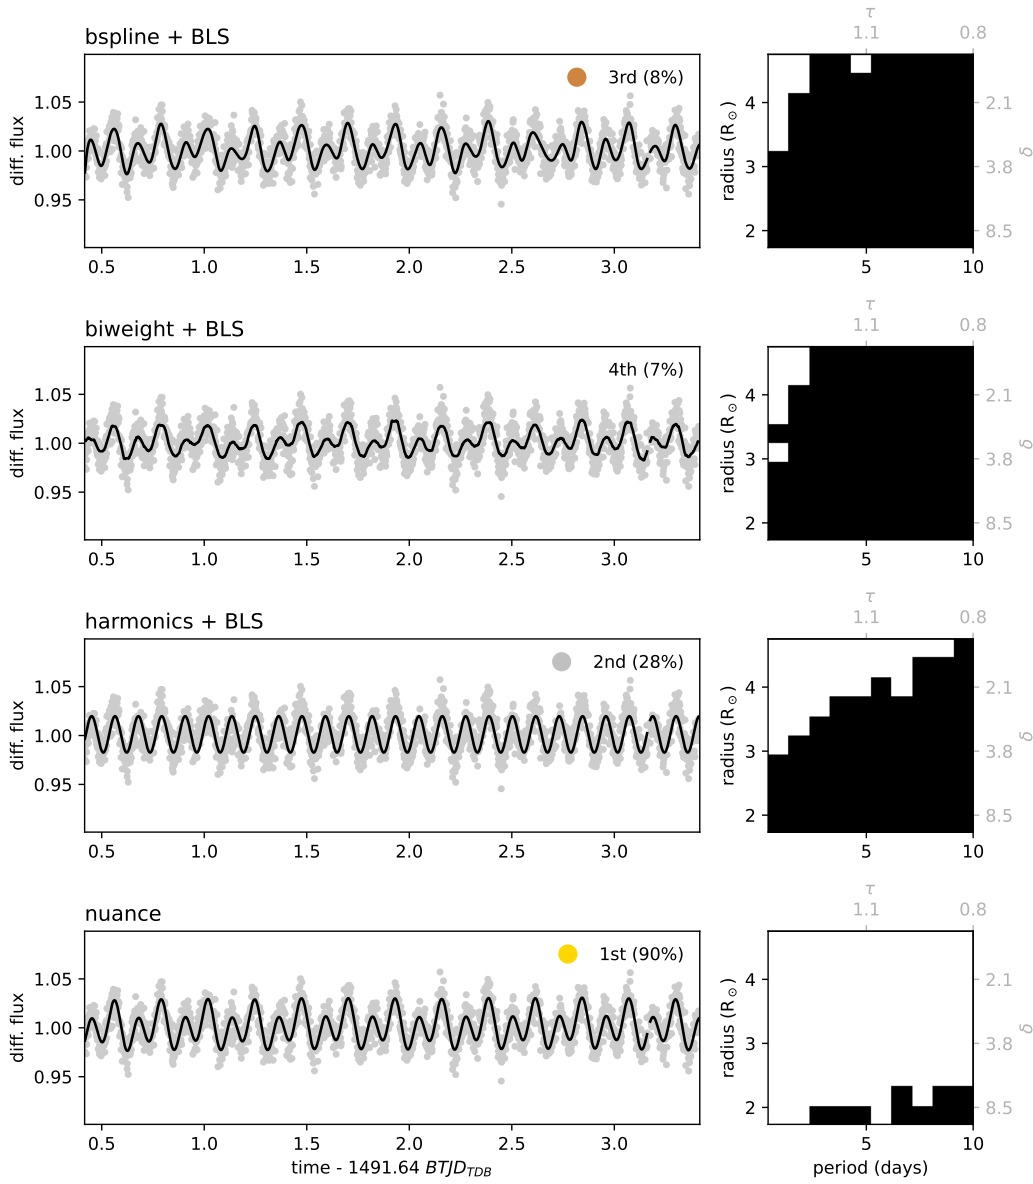


Figure 3.14: Results of the transits injection-recovery on TIC 153193474 single-sector light curve. See [section 3.5.1](#) and legend of [Figure 3.12](#) for a detailed description of this figure. The **harmonics+TLS** method is the best after **nuance**, with 62% less transits detected in total, indicating that all presented methods are unable to detrend without removing many of the injected transit signals (which is not surprising when considering the trend being built, except in the case of **bspline+TLS**, but being apparently too flexible).

TIC 306331621 - sector 6

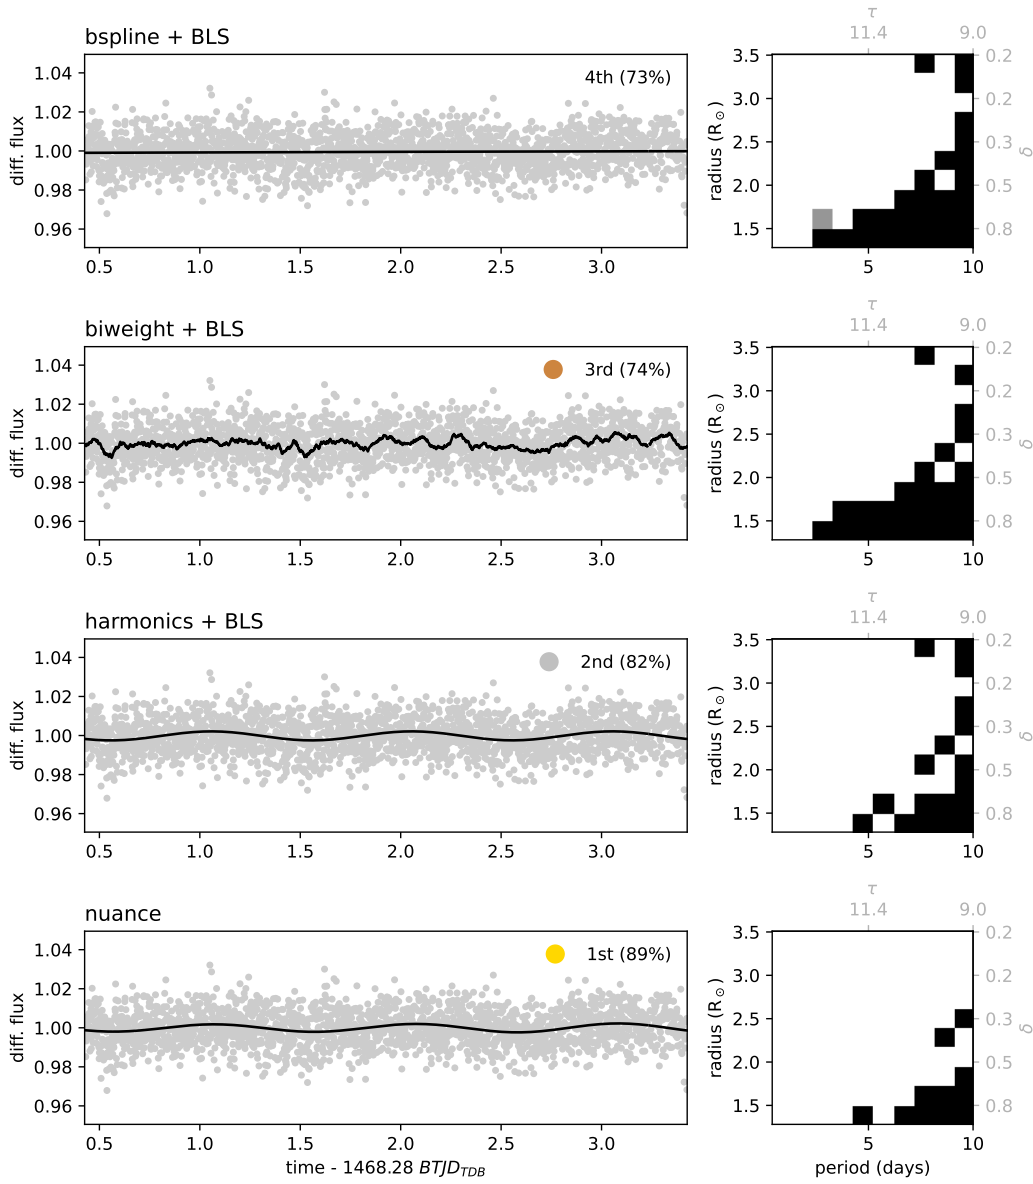


Figure 3.15: Results of the transits injection-recovery on TIC 306331621 single-sector light curve. See [section 3.5.1](#) and legend of [Figure 3.12](#) for a detailed description of this figure. The **harmonics+TLS** method is the best after nuance, with 7% less transits detected in total that could, again, be due to the sole difference between nuance and BLS discussed in [section 3.3](#).

3.6 Discussion

The results from [section 3.5](#) show that `nuance` always outperforms other methods when searching for transit signals in light curves containing correlated noise. While performances with other methods can be very comparable, `nuance` alleviates the need to choose for a specific detrending technique and rather exploit the complete modeling of the light curve using a physically interpretable GP.

In few cases, `nuance` offers a clear advantage in comparison with other methods, in particular for low SNR transit signals in light curves containing fast or non-sinusoidal stellar variability. In these cases, `nuance` is able to robustly detect transits, and provide a unique option in comparison with techniques that likely degrade the searched signals through detrending.

Processing time

As of June 2023, `nuance` is very much a work in progress and a full assessment of its performances is still ongoing. [Figure 3.16](#) is a preliminary study of the computational cost of `nuance` against BLS. For reasons described in the next section, the linear search is first considered separately from the periodic search.




Figure 3.16: Preliminary study of `nuance` processing time against BLS. The lighter blue curve shows the performance of `nuance` when applied to chunks of 10 000 points continuous observations, instead of considering these observations all together (the only option when using BLS).

In [Figure 3.16](#), the processing time of the linear search from `nuance` is recorded against the number of points in the light curve. One advantage of `nuance` is that the linear search can be performed separately on different continuous observations, and then combined in the periodic search. Hence, if searching for transits in separate observations with similar durations (such as different TESS sectors or different ground-based nightly observations), the computational cost of `nuance` grows linearly with the number of observations. Such a parallel cannot be made with BLS, as the algorithm considers all observations as a single one. This way, after assuming that 10 000 points observations can be combined into the periodic search (roughly corresponding to the number of data points in continuous TESS observations for example), `nuance` will constantly perform a single order of magnitude slower than BLS, only considering

the linear search. For this reason, **nuance** must not be employed in the general case, but rather used when light curves contain correlated noise with characteristics described in this chapter. Proper guidelines and metrics to assess the presence of such noise, which would indicate if **nuance** is required, are currently being developed.

For reference, the complete search (including the periodic search) of a single TESS sector for TIC 140212114 (15992 points partly shown in [Figure 3.12](#)) took a total of 30 seconds, executed on a 10 core Apple M1 Pro chip. The latter processing time would grow linearly with the number of sectors being processed.

Periodic search

The periodic search of **nuance** can be done separately from the linear search. For a given light curve, a single linear search can be performed then masked and combined iteratively to search for multiple transiting exoplanets. Hence, [Figure 3.16](#) only shows the processing time of the linear search, assuming that a single planet is searched for. As of June 2023, the periodic search in the current implementation of **nuance** is the least optimized part. But recent developments indicate that the periodic search can be performed in a time equal to the complete processing time of the BLS algorithm. Hence, as the BLS algorithm grows linearly with the number of planets being searched (because reran every time on masked light curves), the increased computing time for **nuance** would be equal, bridging the gap between the computing time of BLS and **nuance** when searching for multiplanetary transiting systems.

Design matrix

In this chapter, starting from [subsection 3.1.2](#), it was assumed that systematic signals (as opposed to astrophysical signals) were completely corrected for without affecting the searched transit signals. For this reason, all analyses in this chapter were done with a single column design matrix \mathbf{X} , corresponding to the mean, constant, differential flux (ideally 1) of a light curve. In practice, **nuance** has been developed to linearly model systematic signals through more complex design matrices, in addition with its capability to model correlated noise while searching for transits. This aspect is completely unexploited in the previous comparisons, and would require to search for transits in completely raw light curves, which comes with few challenges. Given a light curve, this approach involves choosing a design matrix containing specific explanatory measurements, hence requiring a proper model comparison. In addition, large design matrices are expected to increase **nuance** processing time, something that needs to be accounted for in its performance assessment.

Used in the context of the SPECULOOS survey, **nuance** is currently being tested on light curves containing systematic signals that can be modeled linearly using polynomials of common explanatory variables (such as airmass, background flux, pointing error or PSF full-width at half-maximum). On light curves spanning only few hours, with less than a thousand photometric measurements, the use of a linear model to detrend the light curve ignoring the presence of transits is highly susceptible to remove the signals of interest, giving a clear advantage to a full-fledge method like **nuance**. The same applies for filter-based detrending algorithms. Proving the later empirical statement is out of the scope of this chapter and will be treated in a following publication. However, **nuance** might be particularly well suited to search for transits in ground-based observations, where the detrending of sparse observations

is believed to remove the lowest SNR transit signals.

Overall, `nuance` is a powerful tool to search for transits in light curves containing systematic signals and correlated noise, characteristic of late M-dwarfs light curves. Due to several advances in computational methods, as well as its analytical formalism, `nuance` remains highly tractable, and may already be applied to space-based continuous observations. `nuance` is currently being tested on sparse observations, offering a unique tool to search for transiting exoplanet in the light curves of SPECULOOS and other ground-based surveys.

Acknowledgements

I would like to thank Julien De Wit and Prajwal Niraula for meaningful discussions at the beginning of this project, Germain Garcia for his useful insights about numerical optimization, and Michaël Gillon for his overall support. I would like to thank co-author Daniel Foreman-Mackey for his invitation to the Flatiron Institute (Simons Foundation), leading to discussions that shaped the mathematical formalism of the method, as well as introducing the use of `JAX` for its tractable implementation. Co-author Francisco J. Pozuelos helped in comparing the transit search of `nuance` on the light curves of the existing planetary system TOI-540 (positive results that are not presented in this chapter, and that will be featured in a following publication).

Chapter 4

Atmospheric reconnaissance of TRAPPIST-1h

A chapter based on the work published in Garcia et al. 2022a with Sarah E. Moran, Benjamin V. Rackham, Hannah R. Wakeford, Michaël Gillon, Julien de Wit and Nikole K. Lewis. In this chapter, my personal contributions have been distinctly emphasized and clearly differentiated from those made by co-authors.

Starting in 2011, the TRAPPIST-South telescope monitored the fluxes of 50 of the nearest and brightest southern very-late M-dwarfs, each for a total of 50 hours in the search for transiting exoplanets. Doing so, its mission was to assess the typical photometric precision that can be reached for these cool objects, and to estimate the feasibility of the fully dedicated and optimized SPECULOOS transit survey (Gillon et al., 2013). On September 16th 2015, the TRAPPIST telescope, focusing on target *2MASS J23062928-0502285*, was interrupted at the beginning of the night by a magnitude 6 earthquake¹, but was luckily restarted by a conscientious remote operator². Following this restart was the first ever observation of one of the seven earth-sized exoplanets transiting TRAPPIST-1, planet c.

Following up on this first transit observation, the TRAPPIST-South telescope monitored the star for more than 250 hours, unveiling transits from the inner planets TRAPPIST-1b, c and d (Gillon et al., 2016). In a continuous follow-up effort, observations from other ground-based facilities unveiled four additional planets (Gillon et al., 2017), with the orbit of the seventh planet fully characterized using the Kepler Telescope during the K2 mission (Luger et al., 2017). TRAPPIST-1 planets have radii ranging from 0.76 to 1.13 R_{\oplus} , and masses ranging from 0.3 to 1.3 M_{\oplus} (Agol et al., 2021). They are all rocky and likely to be tidally locked, with equilibrium temperatures (assuming null albedo) ranging from 173 to 400 K (Gillon et al., 2017).

TRAPPIST-1 is an M8-type star with an effective temperature of 2566 ± 26 K, and a radius only 12% that of the Sun (Agol et al., 2021). In addition to its high infrared luminosity (magnitude of 10.3 in the K band), these properties makes TRAPPIST-1 a truly

¹Magnitude 8.3 at epicenter, one of the strongest earthquake in the last 100 years in the region (<https://earthquake.usgs.gov/earthquakes/eventpage/us20003k7a/executive>).

²Laetitia Delrez, a member of the TRAPPIST team located in Europe.

unique candidate for the detailed atmospheric characterization of rocky exoplanets with the James Webb Space Telescope (JWST) (e.g. [Morley et al. 2017](#); [Lustig-Yaeger et al. 2019](#)). As of 2022, spectroscopic observations from the Hubble Space Telescope Wide Field Camera 3 (HST WFC3) have been able to rule out the presence of clear primary hydrogen-dominated atmospheres for all planets ([de Wit et al., 2016, 2018](#); [Wakeford et al., 2019](#)) except for the outermost planet TRAPPIST-1h, the most likely to have retained such an extended atmosphere ([Bourrier et al., 2017](#)).

In this chapter, I present the transmission spectrum of TRAPPIST-1h obtained from HST WFC3 observations in order to rule out the presence of an extended hydrogen-rich atmosphere.

4.1 Observations

Three transits of TRAPPIST-1h were observed with the Hubble Space Telescope (HST) Wide Field Camera 3 (WFC3) as part of HST GO program 15304 (PI: Julien de Wit) on UT 2018 July 19 (visit 1), 2019 September 24 (visit 2), and 2020 July 20 (visit 3)³. Each of the three transits was observed over a five-hour window, each requiring four HST orbits, all consisting in approximately 45 minutes of observation separated by 45 minutes gaps due to Earth occultation. Time-series spectroscopy in the 1.12–1.65 μm wavelength band were obtained using the G141 grism in scanning mode, spreading the stellar spectrum perpendicularly to its dispersion axis. The scan rate was set to 0.02 $'' \cdot s^{-1}$ with an exposure time of 112s, resulting in 17-pixel-wide scans acquired in the forward direction only. Each scan is composed of six non-destructive readouts (including the zeroth-read) which were used in the reduction process to remove part of the accumulated background over the complete exposures. A direct image of the target was acquired using the F139M narrow-band filter, at the beginning of each orbit in order to perform G141 wavelength calibration.

Data reduction

I extracted the time-resolved spectra from the three visits presented above following the method presented in [Kreidberg et al. \(2014\)](#), implemented into a `prose` pipeline ([Garcia et al., 2022b](#)) (see [Chapter 2](#)). As in [Kreidberg et al. \(2014\)](#), I started the reduction from the `ima` pre-calibrated images, each processed through the same steps. During a full exposure, nondestructive readouts are recorded such that a specific subexposure can be constructed by subtracting its readout from the previous one, hence removing part of the accumulated background. The pipeline started by forming subexposures out of nondestructive readouts. For each of these subexposures:

1. bad pixels identified by the `calcwfc3` pipeline are masked.
2. The wavelength trace is computed using a direct image of the source and the solution from [Pirzkal et al. \(2016\)](#).
3. A wavelength-dependent flat-field calibration is applied, using [Kuntzschner et al. \(2011\)](#).
4. All image rows are interpolated to the wavelength solution of the direct image so that all values found in a given column correspond to photons from a common wavelength bin.

³These three visits are respectively denoted 03, 02 and 04 in the observing plan

5. The subexposure background is estimated as the median pixel value within a spectrum-free portion of the image, directly subtracted from the subexposure.
6. The spectrum is cut out of the image using a 40-pixel-tall aperture, centered in the scanning direction on the spectrum center of light. This large aperture (compared to the 17-pixel-tall spectrum) is manually set in order to include the tail of the WFC3/IR point spread function.
7. The 1D spectrum is finally built from this cutout using an optimal spectrum extraction algorithm (Horne, 1986). By using this technique, the fact that more background pixels are contained in the wide aperture does not affect the noise of the extracted signal, as these are optimally weighted.

Once these seven steps completed for all subexposures of an image, I interpolated and summed each subexposure spectra over a common wavelength axis, yielding the final 1D spectrum of the image.

To compare my results with results obtained from an independent pipeline, I performed an additional processing of the data using *iraclis*⁴, an open-source Python package presented in Tsiaras et al. (2016). This comparative reduction started from the raw observation products and went through the calibration steps of the *calwfc3* pipeline⁵. Then, for every image, the wavelength-dependent photon trajectories along the scanned spectrum were computed, so that fluxes could be extracted within accurately placed polygonal boxes along the wavelength trace (see Tsiaras et al. (2016), Figure 6). By doing so, the method used in *iraclis* properly accounts for the off-axis nature of the G141 slitless spectra and refines the wavelength solution of every exposure.

I performed the spectra extraction with *iraclis* using the same wavelength bins as the ones used in the *prose* pipeline: 12 bins from 1.1 μm to 1.67 μm , leading to bins of $\sim 0.04 \mu m$ in width. This sampling was chosen so that wavelength bins included entire pixels (in opposition with fractional ones) in the dispersion axis.

The raw data obtained this way featured ramp-like signals characteristic of WFC3 observations. In order to consistently model these signals over all orbits, I discarded the first orbit of each visit, containing a higher-amplitude ramp, as well as the first and second exposures of each orbit. The resulting white-light curves from the three visits are shown in Figure 4.1 (raw data are in light gray), obtained by summing the spectra over all wavelength bins.

⁴<https://github.com/ucl-exoplanets/Iraclis>

⁵<https://hst-docs.stsci.edu/wfc3dhb/chapter-3-wfc3-data-calibration/3-3-ir-data-calibration-steps>




Figure 4.1: White-light curves of visit 1, 2 and 3 of GO 15304 observations. The first orbit is plotted as purple crosses and is not used in the analysis.

4.2 Light-curve modeling

4.2.1 Systematic models comparison

Spectroscopic light curves obtained with HST are known to display a high level of wavelength-, visit-, and orbit-dependent systematic signals (Wakeford et al. 2016 and references therein). As in previous studies, I modeled these signals using empirical functions of time t and HST orbital phase ϕ , which can take a variety of forms. For this study, I chose a model of the form:

$$F(t, \phi) = T(t) + P(t, \phi) + \epsilon(t), \quad (4.1)$$

where T is the transit light curve model, P is a polynomial systematic model of time t and HST orbital phase ϕ , and ϵ is a Gaussian noise. As in Wakeford et al. (2016), I selected the combination of polynomial orders of t and ϕ using model comparison, based on the minimization of the Akaike information criterion (AIC, Akaike 1974), such that the selected models were predictive enough while depending on a minimal number of parameters. By allowing polynomials of up to order three and excluding cross terms, a set of 16 models have been compared.

For a given visit and for each order combination, I inferred the best transit and systematic model parameters in a Bayesian framework using *exoplanet* (Foreman-Mackey et al., 2021)⁶, making use of the inference framework *PyMC3*⁷ (Salvatier et al., 2016) and the *starry* transiting light curve model (Luger et al. 2019, see subsection 1.3.2). The priors I used for the orbital parameters of TRAPPIST-1h are listed in Table 4.1, with T_0 having a uniform prior over the full time of the visit. This uninformative prior on the mid-transit time was used to encompass the wide transit time variations (TTVs) observed for the TRAPPIST-1 planets, reaching more than 1 hour for TRAPPIST-1h (Agol et al., 2021, Figure 2). Finally, I adopted a quadratic form to model the limb-darkening of the transit light curve, with parameters fixed to the values in Table 4.2, obtained from ExoCTK⁸ (Bourque et al., 2021)

⁶<https://docs.exoplanet.codes/en/latest/>

⁷<https://docs.pymc.io/>

⁸https://exoctk.stsci.edu/limb_darkening

using the PHOENIX ACES stellar atmosphere model (Husser et al., 2013).

parameter	unit	value
R_\star	R_\odot	$\mathcal{N}(0.1192, 0.0013)$
M_\star	M_\odot	$\mathcal{N}(0.0898, 0.0023)$
P	<i>days</i>	$\mathcal{N}(18.7672745, 0.00001876)$
b	-	$\mathcal{N}(0.448, 0.054)$
T_0	BJD tdb	$\mathcal{U}(t_0, t_f)$

Table 4.1: Priors on TRAPPIST-1 stellar parameters and TRAPPIST-1h orbital parameters. R_\star and M_\star are respectively the stellar radius and mass with prior distributions from Agol et al. (2021). P and b are TRAPPIST-1h orbital period and impact parameters, taken from Ducrot et al. (2020). T_0 is the transit mid-time, and the transit depth is set to an uninformative uniform prior. $\mathcal{U}(a, b)$ denotes a uniform distribution bounded by (a, b) , and $\mathcal{N}(\mu, \sigma)$ a normal distribution of variance σ^2 centered on μ . t_0 and t_f correspond to the times of the first and last exposures of a given visit.

Bandpass (μm)	u_1	u_2
1.101 - 1.119	0.167 ± 0.020	0.373 ± 0.029
1.140 - 1.159	0.168 ± 0.020	0.384 ± 0.030
1.180 - 1.200	0.175 ± 0.020	0.385 ± 0.029
1.222 - 1.242	0.150 ± 0.019	0.362 ± 0.028
1.265 - 1.286	0.146 ± 0.018	0.338 ± 0.026
1.310 - 1.332	0.177 ± 0.018	0.352 ± 0.026
1.355 - 1.379	0.230 ± 0.019	0.410 ± 0.028
1.403 - 1.428	0.326 ± 0.014	0.359 ± 0.021
1.453 - 1.479	0.289 ± 0.016	0.378 ± 0.023
1.504 - 1.530	0.224 ± 0.019	0.396 ± 0.027
1.558 - 1.585	0.178 ± 0.019	0.373 ± 0.027
1.612 - 1.641	0.135 ± 0.018	0.331 ± 0.026

Table 4.2: Quadratic limb darkening parameters obtained from ExoCTK (Bourque et al., 2021).

Given the data, I found the maximum a posteriori model parameters using the BFGS algorithm (Head & Zerner, 1985), as implemented in `scipy`⁹, and used the resulting likelihood to compute the AIC for each order combination. Finally, for each visit, I retained the model yielding the minimal AIC (Figure 4.2).

⁹<https://docs.scipy.org/doc/scipy/reference/optimize.minimize-bfgs.html>




Figure 4.2: Model comparison for visit 1, 2 and 3. The x-axis represents systematics models where 1 designates a constant, θ^n a polynomial order n of θ , and ϕ^m a polynomial order m of ϕ . The *best* model is circled. For one model to be considered *better* than another, we require that the AIC difference between it and the other must be greater than 20.

While each model was found by modeling the broadband white light curve, i.e. the integrated fluxes time-series over all wavelengths, the systematics model parameters in the rest of the analysis were allowed to vary from one wavelength bin to the other

4.2.2 Transmission spectra inference

I modeled the spectroscopic light curves from each visit using the systematic models found in [subsection 4.2.1](#), following the wavelength-dependent expression of the flux

$$F_\lambda(t, \phi) = T_\lambda(t) + P_\lambda(t, \phi) + \epsilon_\lambda(t), \quad (4.2)$$

using the same notation as in [Equation 4.1](#) with λ denoting the wavelength bin. Since the wavelength-dependent transit depth of TRAPPIST-1h is the parameter of interest, I adopted an uninformative prior for the wavelength-dependent apparent planet’s radius $R_{p,\lambda}$, following

$$R_{p,\lambda} \sim \mathcal{U}(0.5, 2)R_\oplus, \quad (4.3)$$

where $\mathcal{U}(a, b)$ is a uniform distribution bounded by a and b . Again, the orbital parameters and priors I used are listed in [Table 4.1](#). Like in [subsection 4.2.1](#), I chose to model stellar limb-darkening using a quadratic model, the coefficients of which were fixed to the values found in [Table 4.2](#) (a comparative analysis using free limb-darkening coefficients was performed and yielded comparable results). For each visit, I estimated the wavelength-dependent transit

depths and their uncertainties by sampling the data likelihood with a Hamiltonian Monte Carlo (see [subsection 1.4.2](#)), using the PyMC3 implementation ([Salvatier et al., 2016](#)).

Once this was done on individual visits (see the transmission spectra in the left plot of [Figure 4.3](#)), I performed a global analysis of all visits using the priors listed in [Table 4.1](#), yielding a global transmission spectrum shown in the right plot of [Figure 4.3](#) (associated to the spectroscopic light curves shown in [Figure 4.4](#)).




Figure 4.3: Transmission spectra obtained for each visit individually (left) as well as the one obtained through a joint analysis (right). The fact that error bars of the global spectrum (right) seems smaller than the combined errors from the individual transits (left) is due to the fact that the 3 visits combined allow to model a full transit in phase, while visits 2 and 3 do not contain the midpoint of the transit signal leading to larger uncertainties when being taken individually (see also [Figure 4.11](#))

I modeled each spectroscopic light curve with a separate systematic model (shown in [Table 4.3](#)), held to the polynomial form found in the previous section but with wavelength-dependent parameters.

visit	transit time	systematics model
1	2458319.4206 ± 0.0030	$1 + \phi$
2	2458751.0660 ± 0.0003	$1 + \theta + \phi^2$
3	2459051.3365 ± 0.0008	$1 + \theta + \phi$

Table 4.3: Transit times inferred individually from each visit as well as the selected systematic model.

As TTVs are important for TRAPPIST-1h, I performed a similar analysis with a non-periodic transit model, leading to the same transmission spectrum. The transit times obtained by considering each visit individually are reported in [Table 4.3](#).




Figure 4.4: Joint spectroscopic light curve models for the three visits. The raw data are represented by light points, and systematic models are shown by light lines. Solid points correspond to light curves corrected for the systematic models, and solid lines correspond to the inferred transit models per wavelength bands. Finally, histograms of the residuals between the data and models are plotted to the right of each spectroscopic light curve.

Finally, fluxes obtained from *iraclis* were modeled using the same approach, and led to similar transmission spectra. To further validate these results, I compared these transmission spectra to spectra obtained from an independent extraction and modeling provided by co-author Hannah Wakeford, following the approach from [Wakeford et al. \(2019\)](#). To conclude, the inferred transmission spectra were consistent across the three separate analyses, *prose*, *iraclis* and those produced by Hannah Wakeford.

[Figure 4.3](#) highlights variations of the transmission spectra from one visit to another, which seemed more pronounced at shorter wavelengths. To study the possible astrophysical origin of these variations, I conducted a careful assessment of stellar contamination, presented in the next section.

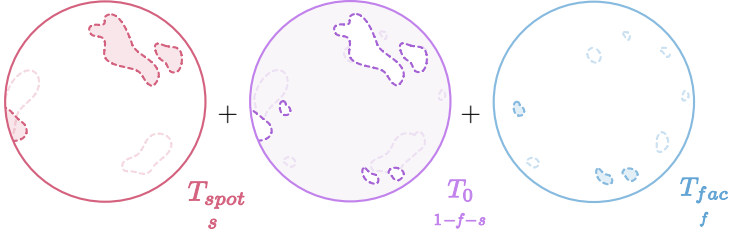
4.3 Stellar Contamination

4.3.1 Photosphere model

To properly interpret the global transmission spectrum inferred in [section 4.2](#), as well as understanding its variations between visits, the possible impact of stellar contamination must be studied. To do that, I employed an approach similar to [Wakeford et al. \(2019\)](#), which consists in modeling the median out-of-transit spectrum of TRAPPIST-1 accounting for a non-uniform surface, one potentially covered in spots and faculae. Having a model of the stellar disk alone then allows to correct the measured transmission spectrum from the stellar one, given which portion of the photosphere is occulted. Like in previous studies (e.g., [Rackham et al. 2018, 2019](#)), I modeled the complete photosphere as a combination of portions of cooler and hotter areas so that the spectrum of the star could be expressed as

$$F = (1 - s - f)F_{phot} + sF_{spot} + fF_{fac}, \quad (4.4)$$

where s and f are the covering fractions of spots and faculae (here simply denoting the cooler and hotter portions of the photosphere) and F_{phot} , F_{spot} , and F_{fac} are the intrinsic spectra of the quiescent photosphere, spots, and faculae, respectively,

i.e. $F =$ 

This simple model assumed that surface heterogeneities have spectra similar to that of a global photosphere with the same temperature. In this model, covering fractions f and s can be zero, so that the model reduces to only one or two temperature components. To model the individual components of the TRAPPIST-1 photosphere, I used the BT-Settl stellar atmosphere models ([Allard et al., 2012](#)), allowing to probe stellar components with temperatures down to 1500K.

Before inferring the parameters of this model against the data, one may wonder why the model described in [Equation 4.4](#) was used instead of the full stellar surface model provided by *starry* (presented in [subsection 1.3.2](#)). The main reason is simplicity. Because no spot occultation was featured in the data (or observed at a very low SNR), there was a very low advantage to model the full surface of the star, especially because of the higher number of parameters it required in comparison with the simpler model described in [Equation 4.4](#). The point of [Garcia et al. \(2022a\)](#) was to reproduce the approach of [Wakeford et al. \(2019\)](#) and show its limitations. Additional data, such as full rotational light curves of TRAPPIST-1 or spot occultations (if any), would be required to properly model the full stellar surface and may then justify the use of *starry* as a unified model to represent the surface. This point is further discussed in [Chapter 5](#).

4.3.2 Models and data preparation

Fitting [Equation 4.4](#) to the data required some preliminary steps. Indeed, BT-Settl theoretical spectra were provided on a grid of stellar metallicities, log-gravity, and effective

temperatures. From this grid, only spectra corresponding to $[Fe/H] = 0$ were kept, close to the inferred value of $[Fe/H] = 0.04$ (Van Grootel et al., 2018). I linearly interpolated these models along the log-gravity parameter, in order to produce spectra with $\log_{10}(g) = 5.22$ (Van Grootel et al., 2018). By fixing these values, only the effective temperature of the models was varied during their inference (except for the case discussed in subsection 4.3.3), ranging from 1500K to 5000K in steps of 100K. Then, I converted these models to flux density at Earth, accounting for TRAPPIST-1’s angular diameter, i.e., $\alpha = (R_*/d)^2$ using the stellar radius $R_* = 0.119 R_\odot$ (Agol et al., 2021) and the distance $d = 40.54$ ly (Gaia Collaboration et al., 2018, DR2). Finally, I convolved these higher resolution fluxes with the WFC3 point-spread-function¹⁰, before being sampled accordingly to the WFC3 IR detector (50\AA pixels).




Figure 4.5: Conversion of visit 2 out-of-transit stellar spectrum. I used the trimmed portion to perform the model fit.

The modeled dataset consisted in three out-of-transit spectra of TRAPPIST-1, one for each visit, that were fitted to a theoretical model. For each visit, I built a median spectrum over time, from which I corrected systematic signals (as they strongly affected the variations of the mean flux value modeled in subsection 4.2.1 and shown in Figure 4.4). To be compared with the models, I converted these spectra from $e^{-1} s^{-1} \text{\AA}^{-1}$ to $\text{erg s}^{-1} \text{cm}^{-2} \text{\AA}^{-1}$ by correcting for the WFC3 IR detector sensitivity¹¹ and accounting for the 112 s exposure time and wavelength passband. WFC3 sensitivity quickly drops to zero outside of the band it covers, so that correcting for it (through division) led to large errors on the tails of the spectra.

¹⁰<https://hst-docs.stsci.edu/wfc3ihb/chapter-7-ir-imaging-with-wfc3/7-6-ir-optical-performance>

¹¹<http://www.stsci.edu/hst/instrumentation/wfc3/documentation/grism-resources/wfc3-g141-calibrations>

While it could be accounted for in the fitting procedure, I obtained a better convergence by trimming the measured spectra to the 1.15–1.63 μm band (see Figure 4.5).

Finally, I added to the data the g, r, i, and z filter photometry from TRAPPIST-1 Pan-STARSS observations (Chambers et al., 2016), in order to better constrain the visible portion of the stellar spectrum. The addition of these data was motivated by the discrepancy observed when comparing the multi-component model from Wakeford et al. (2019) to broad band observations of TRAPPIST-1 in the visible (Figure 4.6)




Figure 4.6: Wakeford et al. (2019) three-component model against our HST data and Pan-STARSS g, r, i, and z wide-band photometry

4.3.3 Fitting the out-of-transit spectra

I inferred the parameters of Equation 4.4 from the dataset previously described, for each visit individually accounting for three distinct cases: a quiescent photosphere (1T), a photosphere with either spots or faculae (2T), and a photosphere with both spots and faculae (3T).

For each visit and for each model, I found the best maximum a posteriori model parameters following a brute force approach, that consisted in estimating the likelihood of the model in a grid of discrete temperatures combinations with 100 K steps (see Table 4.4).

model	components	min-max (K)
1T	F_{phot}	1500-5000
2T	F_{spot}	1500-5000
	F_{phot}	1500-5000
3T	F_{spot}	1500-2300
	F_{phot}	2300-2700
	F_{fac}	2700-5000

Table 4.4: Temperature grids where covering fractions and model likelihoods were estimated, sampled in steps of 100K. The covering fractions for all components are allowed to vary from 0 to 1.




Figure 4.7: 2T model likelihood sampled on the temperature grids described in [Table 4.4](#) and marginalized over the photospheric heterogeneity covering fractions. The symmetry of this distribution is simply due to the symmetry of our two-component model, with covering fractions f and $1 - f$.

For each point in the grid, i.e., for each temperature combination, I found the best covering fractions (between 0 and 1) using the BFGS optimization algorithm. This led to a complete sampling of the likelihood function over all possible temperatures, and allowed me to explore the multi-modal nature of the likelihood distribution over the components' temperatures. In this grid, I retained the maximum likelihood parameters, that I refined with a parallelized Markov Chain Monte Carlo using `emcee` ([Foreman-Mackey et al., 2013](#)), this time in a continuous parameter space, from which I estimated the parameters uncertainties. The 2T model likelihood shown in [Figure 4.7](#) has a well-defined maximum ([Figure 4.7](#)) while the 3T model is bi-modal (see [Figure 4.8](#)).

[Figure 4.9](#) shows the fitted out-of-transit spectrum for model 1T, 2T and 3T, obtained by following the brute force procedure on the visit 2 data. As the theoretical stellar spectra are produced in 100K steps, I used the step size as the uncertainty on the inferred temperatures for each component (see for example the temperature distributions in [Figure 4.8](#))




Figure 4.8: 3T model likelihood sampled on the temperature grids described in [Table 4.4](#) and marginalized over the photospheric heterogeneities' covering fractions, as well as each component's temperatures. For each visit, colored crosses correspond to the global maximum likelihood parameters. The black circle shows a local maximum of the likelihood distribution, highlighting the bi-modal nature of the model.

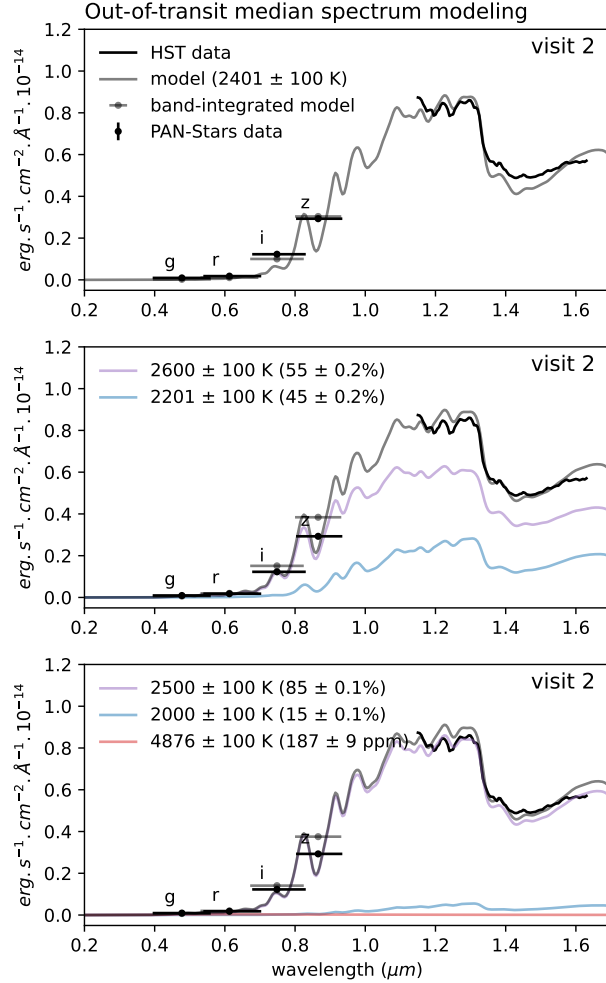


Figure 4.9: 1T, 2T, and 3T models of the visit 2 out-of-transit stellar spectrum.

Like in [Wakeford et al. \(2019\)](#), chi-squared statistics indicated that none of the multi-component models provided a good fit to the data from any visit. Moreover, the 1T model showed a larger discrepancy, also observed by [Wakeford et al. \(2019\)](#) when modeling HST/WFC3 TRAPPIST-1g out-of-transit spectrum against PHOENIX ACES models ([Husser et al., 2013](#)). In their study, this discrepancy was associated to a poor constraint on the stellar radius R_* . This led the authors to introduce a scaling parameter to the stellar radius with the overall motivation that it would account for theoretical models underestimating the radii of low-mass stars. However, introducing such a factor not only affects α but also the log surface gravity $\log_{10}(g) = GM_*/R_*^2$ involved in the stellar model parameters. Using such a radius scaling, varying the value of α and the log surface gravity, did provide a better fit to the data but led to $R_* = 0.1134 \pm 0.0013 R_\odot$, which was $>3\sigma$ discrepant from the radius inferred in [Agol et al. \(2021\)](#). For this reason, I preferred to adopt the unscaled radius found in [Agol et al. \(2021\)](#).

In order to understand the differences of the model parameters between visits, the inferred temperatures and covering fractions found for each visit are shown in [Figure 4.10](#).




Figure 4.10: Inferred temperatures and coverings fraction for the 1T, 2T, and 3T models on each visit.

Using the transmission spectrum correction described in the next section, none of these models were able to explain the differences observed between the transmission spectra of the three visits. However, it should be noted that only visit 1 contained the bottom of the transit signal and resulted in a transmission spectrum very similar to the one obtained through a joint analysis of all three visits. For this reason, the observed variability was probably not astrophysical but instead due to the partial coverage of two of the three transits observed, combined with the systematics effects due to the instrument (see Figure 4.11 for an empirical evidence).

With this conclusion, I assumed that the photospheric structure of TRAPPIST-1 was consistent over the three visits. The model from Equation 4.4 was then compared to the median visits-combined out-of-transit spectrum, leading to the temperature components reported in Table 4.5 for the 1T, 2T, and 3T models. The hotter component of model 3T ($\sim 5000\text{K}$) corresponded to the maximum effective temperature in the model's parameter space. Hence, this inferred temperature was considered as a lower limit to the hot-spot temperature.




Figure 4.11: Correlation between systematics model slope in time and transit depth. This figure has been produced by sampling the likelihood distribution of the data against the model using a Markov Chain Monte Carlo (see [section 1.4](#)). It showed that, apart from visit 1 where the transit mid-time was observed, the transit depth was correlated with the slope in time θ used to model visit-long systematics effects (a shallower slope is compensated by a deeper transit). As only the ingress and egress were observed in visit 2 and 3, it is assumed that the spectrum-to-spectrum differences in this study was partially due to this effect. However, as the visits-combined analysis contained a complete transit (in phase), the systematics slope should have been less correlated with the transit depth, as in visit 1, which was confirmed by the similarity observed between visit 1 and the global spectrum ([Figure 4.3](#))

model	covering fractions	temperatures	configuration
1T	100%	2400 ± 100 K	1T
2T	55.5 ± 0.1 %	2600 ± 100 K	2Tm
	44.5 ± 0.1 %	2200 ± 100 K	2Tc
3T	85.1 ± 0.1 %	2500 ± 100 K	3Tm
	14.9 ± 0.1 %	2000 ± 100 K	3Tc
	185.2 ± 8.8 ppm	4876 ± 100 K	3Th

Table 4.5: Multi-temperatures models parameters inferred from the visits-combined out-of-transit spectrum. For each of the $T1$, $T2$, and $T3$ models, we consider a configuration where the stellar disk is occulted across a unique temperature component. The name of the configuration associated to each component is reported in the last column.

It is interesting to note that the inferred 3T model, which includes bright spots covering a very small fraction of the stellar disk, was compatible (in order of magnitude) with the findings of [Morris et al. \(2018\)](#). In that study, the authors found that 32 ppm of 5300 ± 200 K bright spots were able to explain why the 3.3-day variability observed in the TRAPPIST-1 K2 light curve was undetectable in the *Spitzer* $4.5 \mu\text{m}$ band. To date, there is no direct evidence, nor theoretical support, for the presence of such hot spots on TRAPPIST-1 or any other M-dwarf.

In order to compare this analysis to the one of [Wakeford et al. \(2019\)](#), I repeated the same modeling procedure but using the PHOENIX ACES models (used in the former study) that led to a poorer fit to the data. While the inferred stellar components temperatures were different from the ones found using the PHOENIX BT-Settl models, a bi-modal likelihood

distribution was also observed with these alternative models. For this reason, I adopted the results based on PHOENIX BT-Settle theoretical spectra in the rest of the analysis.

4.3.4 Corrected transit depths

Using the models found in the previous section, I was able to correct the measured transmission spectra for the effect of stellar contamination. I made corrections for 5 different configurations, corresponding to the 1T, 2T, and 3T models with the transit chords either passing completely over the base (F_{phot}) or cooler (F_{spot}) portion of the photosphere. Due to its very small coverage, I discarded the configuration where only the hotter component (F_{fac}) is occulted. For each configuration and wavelength bin, the corrected transit depth was expressed as

$$\delta_c = \delta \times \frac{F}{F_{occ}}, \quad (4.5)$$

where δ is the measured transit depth, F_{occ} is the flux of the stellar disk occulted by the planet over the complete transit chord, and F is the overall flux of the stellar disk. As noted in Wakeford et al. (2019), the planet could pass on top of both components (with a light curve SNR low enough for the occultations not to be detectable), resulting in F_{occ} being a linear combination of both photosphere components' fluxes. However, my study focused on the cases where the occulted portion of the photosphere was made of a single temperature component. Applying this model to the 5 configurations outlined in Table 4.5 led to the corrected transmission spectra plotted in Figure 4.12. Under the assumptions described above, these spectra should be of planetary origin, the star's contribution being modeled and corrected out. 2Tm and 3Tm denote the two-component and three-component photosphere models with the transit chord being respectively over the 2600K and 2500K component. On the other hand, notations 2Tc and 3Tc correspond to transit configurations where the chord is over the cooler component, respectively 2200K and 2000K. Finally, 1T denotes a quiescent photosphere at 2400K for which no spectrum correction is required.

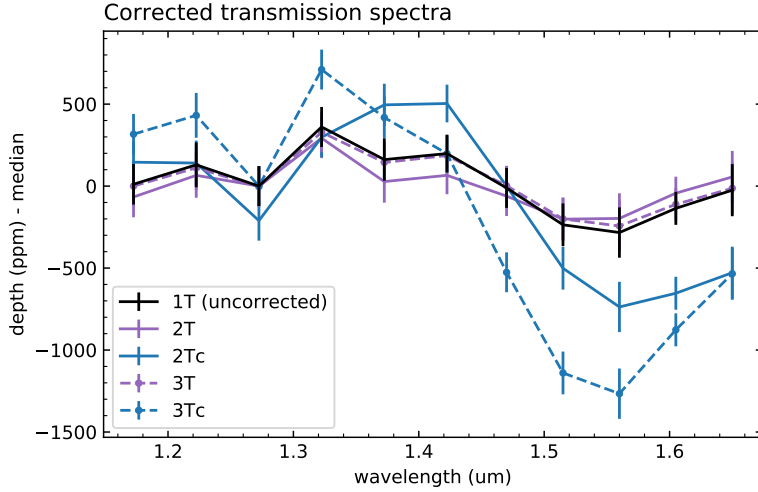


Figure 4.12: Corrected transmission spectra under the 5 configurations reported in Table 4.5. Colored solid lines correspond to the two-temperatures model corrected spectra (2T) and dotted lines denote the three-temperatures model corrected spectra (3T). Purple lines represent spectra of planet h transiting the quiescent component of the stellar photosphere (2T and 3T) while blue lines indicate spectra of planet h transiting the cooler component of the stellar photosphere (2Tc and 3Tc).

4.4 Planetary spectrum

While modeling planetary atmospheres is out of the scope of this dissertation, this section synthesizes the findings of Garcia et al. (2022a) concerning the interpretation of TRAPPIST-1h corrected transmission spectra (an analysis led by co-author Sarah Moran). As in Wakeford et al. (2019), the forward model framework CHIMERA (Line et al., 2013) was used, previously adapted to the inner TRAPPIST-1 planets (Batalha et al., 2018; Moran et al., 2018). These models included Rayleigh scattering due to H₂/He while any other form of scattering or absorption from clouds or hazes were neglected. Given the precision of the HST observations, simple aerosol-free atmospheric models were produced. For each model atmosphere, a planetary mass of 0.755 M_⊕ and a solid body radius of 0.326 R_⊕ was adopted (Agol et al., 2021).

The following planetary atmospheres were considered:

- A 1× solar metallicity atmosphere, which for an equilibrium temperature of 170 K corresponds to a scale height of 100 km. Opacities from water, methane, carbon monoxide, carbon dioxide, ammonia, molecular nitrogen, and H₂/He collision-induced-absorption (CIA) (Freedman et al., 2008, 2014) are included.
- A 10× solar metallicity atmosphere including only opacities from water, methane, and H₂/He-CIA, in order to produce the largest molecular features in the WFC3 G141 band-pass.
- A 200× solar metallicity CO₂-rich atmosphere, including opacities from CO₂, CH₄, and CO.

- A pure N₂ atmosphere without Rayleigh scattering, which in the near-infrared region of WFC3 G141 is also representative of an airless body, as well as potentially representing a heavily aerosol-laden atmosphere.

4.4.1 Model Results with HST

Accounting for all planetary, stellar atmosphere models and transit configurations, [Table 4.6](#) presents the scenarios compatible with the HST observations (see also [Figure 4.13](#)).

Photospheric configuration	Plausible?	Explanation
1T	Yes	Cannot be ruled out beyond 1.4σ with any atmospheric model
2Tm	Yes	Cannot be ruled out beyond 1σ with any atmospheric model
2Tc	Unlikely	Can be ruled out beyond 3.5σ with any H ₂ -rich model; Can be ruled out to $\sim 1\sigma$ with flat model
3Tm	Yes	Cannot be ruled out beyond 1.3σ with any atmospheric model
3Tc	No	Can be ruled out beyond 2.5σ with every atmospheric model
3Th	No	Hot spot coverage too small for planet to transit

Table 4.6: Summary of which stellar photospheric configurations remain plausible when including the interpretation of planetary atmospheric models.

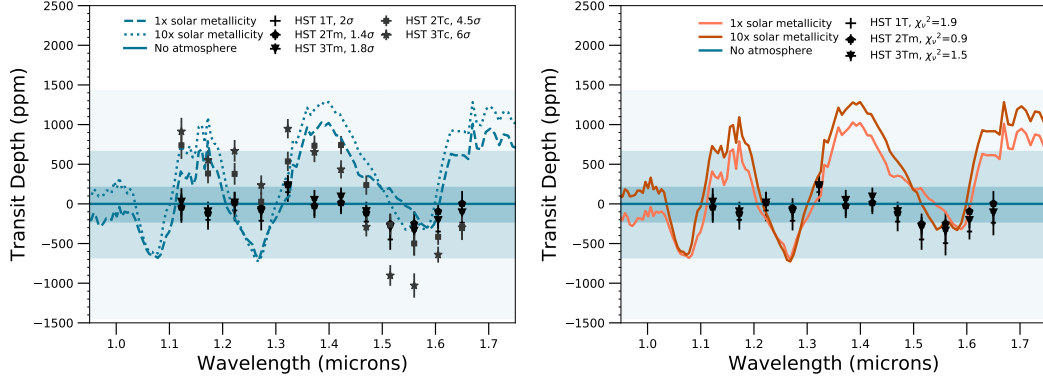


Figure 4.13: Model atmospheres produced using CHIMERA compared to the five potential planetary spectrum scenarios. Left: In addition to the mean-subtracted models (solid, dashed, and dotted lines) are included the 1σ , 2σ , and 3σ errors of the flat “no atmosphere” model as blue shaded regions. The confidence to which the solar metallicity scenario can be ruled out is reported after each scenario. The cold spot scenarios, 3Tc and 2Tc, was clearly ruled out at high confidence. Of the remaining models and scenarios, no model could be confidently excluded at $\geq 3\sigma$. However, a featureless infrared spectrum (solid line) either due to an airless body or an atmosphere without infrared features (e.g., a molecular nitrogen atmosphere) was statistically preferred over a $1\times$ solar metallicity atmosphere (dashed line) or a $10\times$ solar, water and methane-rich atmosphere (dotted line). Right: The scenarios with the best goodness of fit (1T, 2Tm, 3Tm) against the three model atmospheres ($1\times$ solar, light orange; $10\times$ solar, dark orange; no atmosphere, blue). The χ^2_ν of the flat line are reported after the three scenarios in the legend.

4.4.2 Model Results with Full Wavelength Coverage

Figure 4.14 shows the results of the previously modeled planetary atmospheres but including an extended wavelength coverage from $0.8 \mu\text{m}$ to $4.5 \mu\text{m}$ of all existing transits of TRAPPIST-1 h, using data from the space-based K2 campaign (Grimm et al., 2018), the SPECULOOS-South Observatory (SSO) and the Liverpool Telescope (LT) on the ground (Ducrot et al., 2018), and Spitzer/IRAC data at $3.6 \mu\text{m}$ and $4.5 \mu\text{m}$ (Ducrot et al., 2020).




Figure 4.14: As in Figure 4.13 (right) with the addition of K2/SSO/LT data points (Ducrot et al., 2018) and Spitzer IRAC channel 1 and channel 2 data (Ducrot et al., 2020), shown here corrected for stellar contamination scenario 3Tm, the “best fit” configuration. We also included a carbon dioxide-rich atmospheric model to compare against the data, though in all scenarios for the stellar photosphere, we minimized the χ^2 with the flat “high mean molecular weight/no atmosphere” model, which we reported as χ_ν^2 after each transit depth scenario in the legend. Note that we have applied offsets to the K2, SSO, LT, and Spitzer data, as discussed in more details in Garcia et al. (2022a).

Given these additional data, final implications for TRAPPIST-1h planetary atmosphere are described in Table 4.7

Atmosphere	Description	Plausible?	Explanation
1× solar	Chemical equilibrium Rayleigh scattering H ₂ , He, H ₂ O, CH ₄ , CO, CO ₂ , NH ₃ , N ₂	Strongly unlikely	Poor fit to data, though not ruled out to 1 σ
10× solar	Chemical equilibrium Rayleigh scattering H ₂ , H ₂ O, CH ₄	Strongly unlikely	Poor fit to data, but only ruled out to 1 σ
200× solar	Chemical equilibrium Rayleigh scattering CH ₄ , CO, CO ₂	Unlikely	Fits data with χ_ν^2 of ~ 4 or better
No atmosphere or Flat model or Aerosol-laden or N ₂	No Rayleigh scattering	Yes	Fits data with χ_ν^2 of ~ 1.4 or better

Table 4.7: Summary of atmospheric models considered and their plausibility given the data

4.5 Conclusion

I presented an analysis of the infrared transmission spectrum of TRAPPIST-1h, from 1.12 to 1.65 μm , obtained with the Hubble Space Telescope Wide Field Camera 3 and using the G141 grism. I extracted the spectra from the `ima` images with a pipeline based on `prose`, and modeled the resulting spectroscopic light curves within a Bayesian framework, with systematics model selected through the minimization of the Akaike Information Criteria. In order to address the effect of stellar contamination, I modeled the median out-of-transit spectrum of TRAPPIST-1 against a multi-temperature combination of the PHOENIX BT-Settl theoretical models, following the approach of Wakeford et al. (2019). While the retrieved transmission spectra from each visit treated individually varied, none of the single-, two-, nor three-component photospheric models (each component having a different temperature) was able to explain this variability. Despite that, the three-component model suggested the possibility that TRAPPIST-1 surface may be covered by $14.9\% \pm 0.1\%$ of cold spots (2000K \pm 100K) and by a very small fraction (185.2 ± 8.8 ppm) of hot spots (hotter than 5000K). This result was compatible with the results of Morris et al. (2018) that reported the varying amplitude of the photometric variability of TRAPPIST-1 observed between K2 and Spitzer.

With the potential stellar photospheric scenarios in hand, my colleague Sarah Moran studied a number of potential planetary atmospheres using the forward model CHIMERA, including various H₂-dominated atmospheres, a carbon dioxide-rich atmosphere, and a featureless “flat” line model. While the data quality combined with the cold, small (likely rocky) nature of the planet was not sufficient to exclude any atmospheric scenario to high statistical significance, this study concluded that TRAPPIST-1 h is not likely to possess a hydrogen-dominated atmosphere. The likeliest scenario for this planet was that it possesses a very high mean molecular weight ($\geq 1000\times$ solar metallicity) atmosphere, is enshrouded by an opaque aerosol layer, or is devoid of atmosphere entirely. Determination of the true nature of TRAPPIST-1 h’s atmosphere, or lack thereof, awaits upcoming measurements by JWST.

Overall, none of the photosphere models considered, including a homogeneous photosphere, provided a good fit to the data. While it might come from the accuracy of the PHOENIX models being used, this discrepancy is inherent to any M-dwarf photosphere model (as of early 2023, Iyer et al. 2023). Modeling the photosphere of late M-dwarfs during transit (essential to tackle the effect of stellar contamination) will require some prior knowledge about the structure of their photospheres. This prior knowledge is currently not available, but will certainly be essential for the atmospheric characterizations to come. The next chapter presents some avenues to achieve this goal, by developing methods and observations to better understand the surface of late M-dwarfs.

Chapter 5

Prospects in studying the photosphere of late M-dwarfs

In [Chapter 4](#), the lack of knowledge about the photosphere of TRAPPIST-1 affected what we could infer about the atmosphere of TRAPPIST-1b. While the latter transmission spectra were obtained at low SNR, this problem is expected to have similar impacts on the interpretation of higher SNR transmission spectra obtained with any instrument, including those of the James Webb Space Telescope ([Rackham et al., 2018](#)). As of June 2023, our incapacity to tackle the effect of stellar contamination hinders our plan to characterize the atmosphere of any transiting rocky exoplanet, a plan envisioned for the next two decades using major flagship missions.

To tackle the effect of stellar contamination, I propose to consider two different avenues: having a deterministic representation of the stellar surface during every transit, or having statistical information about the surface that can be combined with complementary measurements to constrain the photosphere of a star during any transit¹. Given current technologies, having a precise snapshot of the surface of a star as cold as TRAPPIST-1 will probably not be possible in the next decade, and even then, would only represent a single picture of a potentially dynamic object. On the other hand, statistical information about the surface of late M-dwarfs, grouped in populations of stars with similar properties, can be obtained, and can be used to constrain the photosphere of a given star during any transit.

In 2021, a study from one of NASA’s Exoplanet Exploration Program Analysis Group (SAG21) published 14 important findings ([Rackham et al., 2023](#)), *“contextualized statements of what we understand to be the current research needs that can be addressed to further our understanding of stellar photospheres and make the best use of precise space-based transmission spectra of exoplanets”*. In this chapter, I present some technical and observational avenues to better understand the surface of late M-dwarfs, and relate them to the findings of the SAG21.

¹A third avenue could consist in measuring transmission spectra at wavebands where active regions have a negligible effect.

5.1 Making every transit count

JWST already conducted secondary eclipse observations of TRAPPIST-1b (Greene et al., 2023) and more observations are expected to revolutionize our understanding of the rocky exoplanets orbiting TRAPPIST-1. But secondary eclipse spectroscopy may require a significantly greater JWST time commitment than transmission spectroscopy to achieve comparable constraints on the detection of the TRAPPIST-1 planet atmospheres (Lustig-Yaeger et al., 2019). Nonetheless, a detailed characterization of the atmosphere of any of the TRAPPIST-1 planets will require numerous transits to be observed (see Figure 5.1 from Lustig-Yaeger et al. 2019).

Detect Atmospheres in Transit with $\langle \text{SNR} \rangle = 5.0$
NIRSpec Prism sub512 ngroup6

Type of Atmosphere	b	c	d	e	f	g	v
1 bar H ₂ O	—	—	—	13	—	—	—
1 bar H ₂ O cloudy	—	—	—	23	—	—	—
10 bar CO ₂	2	4	2	7	7	7	7
92 bar CO ₂	2	4	2	8	7	7	7
10 bar Venus	—	18	15	30	12	9	8
92 bar Venus	—	22	24	31	12	11	8
10 bar O ₂ outgassing	2	3	2	10	9	10	9
100 bar O ₂ outgassing	2	4	2	7	5	4	4
10 bar O ₂ desiccated	2	3	2	8	6	6	5
100 bar O ₂ desiccated	2	4	2	11	9	8	6

TRAPPIST-1

Figure 5.1: Figure 10 from Lustig-Yaeger et al. (2019): Number of transits for each TRAPPIST-1 planet necessary to rule out a featureless spectrum with $\text{SNR} = 5$ for different self-consistent atmospheric compositions using JWST NIRSpec Prism.

For transmission spectroscopy to deliver on its promises, the effect of stellar contamination must be tackled. The less is known about the surface of late M dwarfs, the less powerful transmission spectroscopy become, independently of the number of transits observed.

5.2 Ensemble light curves analysis

Before the discovery of exoplanets, the impact of active regions on the study of celestial bodies was already a concern, especially present in the 1980s literature about the study of eclipsing binaries. As its origin, the hypothesis that starspots could cause stellar photometric variability dates back more than 300 years ago. In 1667, when talking about the variability of the Mira star (Omicron Ceti), Ismaël Boulliau wrote:

”[...] one part of the spherical mass of that star should be bright on one hemisphere and the other part should be dark on the other hemisphere, and the body of the star should revolve and rotate around its own center and around one axis, in such a way that a straight line drawn from the Earth to the star forms right angles with it.”

More than two centuries later, Heinrich Bruns (1882) formulated this hypothesis mathematically, highlighting a cumbersome issue: while appealing to the scientific community, the problem of inferring the surface map of a body based on an unresolved light curve is highly degenerate and involves an infinite number of solutions. In a more seminal study, [Russell \(1906\)](#) independently reached the same conclusion, by expanding the surface intensity of a spherical body as a sum of spherical harmonics (cf. Introduction)².

More recently, [Luger et al. \(2021b\)](#) highlighted that the surface inversion problem arises partly from uncertainties in stellar inclinations. To break these degeneracies, [Morris \(2020\)](#) and [Luger et al. \(2021a\)](#) employed ensemble analyses of light curves, and demonstrated both theoretically and empirically that unique constraints on the surface maps of stars with similar statistical properties can be obtained (related to the *finding 3.1* of the SAG21).




Figure 5.2: Number of observation hours per spectral type by the SPECULOOS survey as of June 2023. On top of each bar is the number of stars of a certain spectral type observed by the survey. This figure contains rough non-exhaustive estimates with missing targets.

As of June 2023, the SPECULOOS survey observed a considerable number of late M-dwarfs (see [Figure 5.2](#)), resulting in a unique sample of high-precision light curves featuring stellar variability (some examples given in [Figure 5.3](#)).

²An historical account of the starspot hypothesis can be found in [Kopal \(1982\)](#).




Figure 5.3: Example of SPECULOOS rotational light curves.

Linking these observations to populations of stars likely to exhibit similar surface properties (e.g. using their age and other activity indicators) is a challenge, but may open the door to a statistical understanding of the surface of late M-dwarfs. While this applies to TESS light-curves, the SPECULOOS survey is uniquely positioned thanks to the precision of its observations, allowing to probe the variability of the coolest M-dwarfs (e.g. [Figure 5.3](#)).

The distribution of surface properties obtained from such an ensemble analysis will be directly exploitable as priors for analyses similar to the one conducted in [Chapter 4](#), breaking several of the degeneracies observed when trying to account for the surface properties of the star in future atmospheric retrievals. The extent to which such priors will indeed allow to break the aforementioned degeneracies is still unknown, and should be carefully studied.

5.3 Spots occultations

A lot of the degeneracies inherent to the light curve inversion problem break when a spot is occulted by another celestial body. This phenomenon is denoted *spot occultation* and can be used to directly probe the physical and spectroscopic properties of an active region ([Silva, 2003](#)). Recently, a $6 R_{\oplus}$ companion to the M4-type star TOI-3884 was discovered to transit a large and static polar spot ([Almenara et al. 2022](#), [Libby-Roberts et al. 2023b](#)). Using the SPECULOOS telescopes, this repeated event was observed in several broad-band filters, allowing to constrain the temperature of the spot (albeit the choice and accuracy

of the stellar atmosphere model used). This intensive follow-up is shown in Figure 5.4 and illustrates the powerful capabilities of ground-based observations for such events, that can be conducted at smaller cost than space-based follow-up using low-resolution spectrographs




Figure 5.4: SPECULOOS broad-band observations of TOI-3884 polar spot occultation. Modeling of these multi-color observations employed the *starry* model (Luger et al., 2019) using spherical harmonics up to degree 12. Due to an absence of variability in the continuous TESS observations, the spot was assumed to be polar so that only the spot contrasts and size were free. The spot is modeled on the surface using a Gaussian function expanded over the spherical harmonics maps. The best-fit model is shown in solid lines and the inferred contrasts and spot radius are reported in Figure 5.5.

A full description of the modeling of TOI-3384 polar spot using SPECULOOS observations will be addressed in an upcoming publication. In the meantime, already published analyses raise several questions:

- In Almenara et al. (2022) and Libby-Roberts et al. (2023b), the spot occultations from TOI-3884 were modeled separately from the transit light curve, potentially due to the difficulty of sampling the posterior likelihood of a complete model. How much bias is introduced by this two-step modeling, both in terms of orbital parameters and spot properties?
- More generally, what can be theoretically inferred about spot properties from a single spot occultation?
- Beyond these observations, what can be theoretically inferred about spot properties from a combination of measurements (such as rotational light curves and spectroscopic observations). And must these measurements be modeled simultaneously? (related to *finding 2.1* and *2.2* of the SAG21)
- What are the biases introduced by the use of inaccurate stellar atmosphere models?




Figure 5.5: (Left) Characterization of TOI-3884 polar spot using BT-Sett models (Allard et al., 2012). The inferred spot temperature is 200 K lower than the 3192 K quiescent photosphere. (Top) Transmission of each filter. Error bars are not shown, as sampling the posterior distribution of the spot contrast was one of the important issue I encountered when modeling these observations using a complete modeling of this event (potentially leading to the two-step modeling employed in previous studies).

This questioning call for a theoretical study of the information content of such combined observations, and would be facilitated by the existence of a unified model. The formalism of *starry*, i.e. modeling the photometric and spectroscopic observations of a star with a surface represented by a sum of spherical harmonics, could act as such, and might provide a natural framework to address these questions (related to the *finding 2.3* of the SAG21).

5.4 Open-source tools

Something equally important to the modeling of transmission spectra, to the analysis of stellar variability and to the detection of transit signals in light curves of late M-dwarfs, is the availability, maintenance and reproducibility of data processing tools.

Reduction of spectroscopic images

As the signals of interest are obtained at low SNR, reduction of space-based transmission spectra from raw images are very sensitive to data processing choices. To avoid hidden biases, an option is to maintain good practices, and develop open-source and documented reduction

tools that foster the reproducibility of published results. Tools like Eureka³, Iraclis (Tsiaras et al., 2016) or PACMAN (Zieba & Kreidberg, 2022) are going in that direction, and should be carefully accompanied by an extensive documentation to avoid their use as black-boxes. While it is solely an engineering issue, it will be key to disentangle biosignatures from data processing or modeling artifacts. In that respect, the analysis presented in Chapter 4 failed at delivering an open-source `prose` pipeline to reduce HST/WFC3 images, but a similar project targeting both HST and JWST instruments images is ongoing (by external collaborators).

Image processing of ground-based observations

Ground-based photometric time-series are imprinted by numerous instrumental signals. While it is true for all stars, it is particularly true for late M-dwarfs observations conducted in the near-infrared, where thermal and atmospheric effects increase spurious signals. A way to tackle this effect is to record as much explanatory measurements as possible when processing raw images, and to allow pipelines to adapt to the specificities of each instrument. These capabilities are at the core of `prose` development, a tool that keep being maintained through an increasing number of contributors, and applied to an increasing number of projects. Future developments will include:

- Fast PSF photometry using JAX.
- Difference imaging
- Better background extraction techniques
- Image objects adapted to spectroscopic observations

Overall, `prose` entertains the idea that improvement of methods in a single field of astronomy can benefit all other fields, at no cost for the community.

Detection of transiting exoplanets around active stars

Further development of `nuance` (section 3.2) will enable its application to sparse ground-based observations. Whether from space or from the ground, detecting planets transiting active stars using `nuance` will not only allow a better estimate of the occurrence rate of planets around M-dwarfs (Dressing & Charbonneau, 2015; Ment & Charbonneau, 2023), but will uncover systems that likely yield spot occultations, a powerful way to understand the active regions and photosphere of late M-dwarfs. `nuance` is also maintained as an open-source Python package⁴, allowing the community to perform more informed transit searches.

³<https://github.com/kevin218/Eureka>

⁴<https://github.com/lgrcia/nuance>

Conclusion

Astrobiology is a relatively young field. Given current technologies, late M-dwarfs represent a unique opportunity to search for biosignatures on rocky exoplanets, that must be seized while more powerful instrumentation and techniques are being developed. However, although being the most numerous stars in our solar neighborhood, M-dwarfs are challenging to observe, resulting in unknowns that affect what we can learn about their transiting companions. During my PhD, I studied several challenges associated with this issue, and developed improved tools whose benefits extend the study of exoplanets around M-dwarfs.

I led the development of `prose`⁵, an instrument-agnostic Python package to process astronomical images in a modular way. Using this framework, I developed fast and robust software, among which is a photometric reduction pipeline widely used among my collaborators, as well as by external teams. With this tool, ground-based observations from the TRAPPIST telescopes are processed on a nearly-daily basis, in response to the major follow-up needs of the Transiting Exoplanet Survey Satellite (TESS) mission. In parallel of this development, I made several contributions to the SPECULOOS survey, by performing maintenance missions, developing a web-based interface to allow its consortium to better explore the survey's scientific products, and by performing remote operations of the Northern telescope Artemis.

Light curves of late M-dwarfs show enhanced stellar variability and correlated noise that act as nuisance signals when searching for planetary transits. As of June 2023, common techniques consist in detrending a light curve from correlated noise (including stellar variability) before searching for transits. During my PhD I have shown the impact of this approach on transit detectability and the extent to which it degrades the searched signals. Needing an alternative, I developed `nuance`, an algorithm using linear models and Gaussian processes to search for transits while modelling correlated noise in a tractable way⁶. On TESS light curves, I have shown that `nuance` always outperforms common two-step transit search strategies, doing so approximately one order of magnitude slower than commonly used techniques. Despite this figure, `nuance` remains highly tractable, with unique performances on light curves featuring short-scale or non-sinusoidal variability. This work is still ongoing, and good practices about how and when to use `nuance` will be presented in an upcoming publication.

In 2022, I led the atmospheric reconnaissance of TRAPPIST-1h, the outermost planet of the TRAPPIST-1 system. Using observations from the Hubble Space Telescope Wide Field Camera 3, I extracted spectroscopic light curves from three separate transits, and modeled

⁵<https://github.com/lgrcia/prose>

⁶Featured in the open source Python package of the same name (<https://github.com/lgrcia/nuance>).

the wavelength-dependent apparent radius of the planet to derive its transmission spectrum. As the transmission spectra of the three visits featured variations that could not be of planetary origin, I investigated the potential contribution of stellar contamination. Using the out-of-transit spectrum of the star, I showed that our lack of knowledge about TRAPPIST-1 photosphere lead to degeneracies in the interpretation of the resulting transmission spectra. By assuming five separate scenarios, I produced corrected spectra that were compared to forward models produced by colleague Sarah Moran, overall showing that TRAPPIST-1h is unlikely to possess a hydrogen-dominated atmosphere.

Finally, using the knowledge acquired during my PhD, I started building a multidisciplinary project that will help in better understanding the surface of late M-dwarfs, a key ingredient in tackling the effect of stellar contamination. This project will highly benefit from the SPECULOOS survey and its unique capabilities, as well as the expertise of a community that recently learned about the impact of stellar contamination.

To conclude, the various studies conducted during my PhD allowed me to gather a multifaceted expertise on a growing issue impacting the field of Astrobiology, and develop ideas pushing towards the detailed atmospheric characterization of terrestrial exoplanets around late M-dwarfs, in the search for life.

Bibliography

- Adami, C., & LaBar, T. 2015, arXiv e-prints, arXiv:1506.06988, doi: [10.48550/arXiv.1506.06988](https://doi.org/10.48550/arXiv.1506.06988)
- Agol, E., Luger, R., & Foreman-Mackey, D. 2020, *AJ*, 159, 123, doi: [10.3847/1538-3881/ab4fee](https://doi.org/10.3847/1538-3881/ab4fee)
- Agol, E., Dorn, C., Grimm, S. L., et al. 2021, Refining the transit timing and photometric analysis of TRAPPIST-1: Masses, radii, densities, dynamics, and ephemerides. <https://arxiv.org/abs/2010.01074>
- Aigrain, S., Parviainen, H., & Pope, B. J. S. 2016, *MNRAS*, 459, 2408, doi: [10.1093/mnras/stw706](https://doi.org/10.1093/mnras/stw706)
- Akaike, H. 1974, *IEEE Transactions on Automatic Control*, 19, 716
- Allard, F., Homeier, D., & Freytag, B. 2012, *Philosophical Transactions of the Royal Society of London Series A*, 370, 2765, doi: [10.1098/rsta.2011.0269](https://doi.org/10.1098/rsta.2011.0269)
- Almenara, J. M., Bonfils, X., Forveille, T., et al. 2022, *A&A*, 667, L11, doi: [10.1051/0004-6361/202244791](https://doi.org/10.1051/0004-6361/202244791)
- Artigau, É., Kouach, D., Donati, J.-F., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, *Ground-based and Airborne Instrumentation for Astronomy V*, ed. S. K. Ramsay, I. S. McLean, & H. Takami, 914715, doi: [10.1117/12.2055663](https://doi.org/10.1117/12.2055663)
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)
- Bailer-Jones, C. A. L., & Mundt, R. 2001, *A&A*, 367, 218, doi: [10.1051/0004-6361:20000416](https://doi.org/10.1051/0004-6361:20000416)
- Barnes, J. W. 2007, *PASP*, 119, 986, doi: [10.1086/522039](https://doi.org/10.1086/522039)
- Batalha, N. E., Lewis, N. K., Line, M. R., Valenti, J., & Stevenson, K. 2018, *ApJL*, 856, L34, doi: [10.3847/2041-8213/aab896](https://doi.org/10.3847/2041-8213/aab896)
- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393, doi: [10.1051/aas:1996164](https://doi.org/10.1051/aas:1996164)

- Bonfils, X., Almenara, J. M., Jocu, L., et al. 2015, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9605, Techniques and Instrumentation for Detection of Exoplanets VII, ed. S. Shaklan, 96051L, doi: [10.1117/12.2186999](https://doi.org/10.1117/12.2186999)
- Bourque, M., Espinoza, N., Filippazzo, J., et al. 2021, The Exoplanet Characterization Toolkit (ExoCTK), 1.0.0, Zenodo, doi: [10.5281/zenodo.4556063](https://doi.org/10.5281/zenodo.4556063)
- Bourrier, V., de Wit, J., Bolmont, E., et al. 2017, AJ, 154, 121, doi: [10.3847/1538-3881/aa859c](https://doi.org/10.3847/1538-3881/aa859c)
- Bradley, L., Sipőcz, B., Robitaille, T., et al. 2020, astropy/photutils: 1.0.0, 1.0.0, Zenodo, doi: [10.5281/zenodo.4044744](https://doi.org/10.5281/zenodo.4044744)
- Broeg, C., Fernández, M., & Neuhäuser, R. 2005, Astronomische Nachrichten, 326, 134, doi: [10.1002/asna.200410350](https://doi.org/10.1002/asna.200410350)
- Browning, M. K. 2008, ApJ, 676, 1262, doi: [10.1086/527432](https://doi.org/10.1086/527432)
- Burdanov, A. Y., de Wit, J., Gillon, M., et al. 2022, PASP, 134, 105001, doi: [10.1088/1538-3873/ac92a6](https://doi.org/10.1088/1538-3873/ac92a6)
- Butler, R. P., Marcy, G. W., Fischer, D. A., et al. 1999, ApJ, 526, 916, doi: [10.1086/308035](https://doi.org/10.1086/308035)
- Caldwell, D. A., Tenenbaum, P., Twicken, J. D., et al. 2020, Research Notes of the American Astronomical Society, 4, 201, doi: [10.3847/2515-5172/abc9b3](https://doi.org/10.3847/2515-5172/abc9b3)
- Chamberlain, J. W., & Hunten, D. M. 1987, Theory of planetary atmospheres. An introduction to their physics and chemistry., Vol. 36
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints, arXiv:1612.05560. <https://arxiv.org/abs/1612.05560>
- Charbonneau, D., Brown, T. M., Noyes, R. W., & Gilliland, R. L. 2002, ApJ, 568, 377, doi: [10.1086/338770](https://doi.org/10.1086/338770)
- Charbonneau, P. 2005, Living Reviews in Solar Physics, 2, 2, doi: [10.12942/lrsp-2005-2](https://doi.org/10.12942/lrsp-2005-2)
- Claret, A. 2000, A&A, 363, 1081
- Collins, K. A., Kielkopf, J. F., Stassun, K. G., & Hessman, F. V. 2017, AJ, 153, 77, doi: [10.3847/1538-3881/153/2/77](https://doi.org/10.3847/1538-3881/153/2/77)
- de Wit, J., Gillon, M., Demory, B. O., & Seager, S. 2012, A&A, 548, A128, doi: [10.1051/0004-6361/201219060](https://doi.org/10.1051/0004-6361/201219060)
- de Wit, J., Wakeford, H. R., Gillon, M., et al. 2016, Nature, 537, 69, doi: [10.1038/nature18641](https://doi.org/10.1038/nature18641)
- de Wit, J., Wakeford, H. R., Lewis, N. K., et al. 2018, Nature Astronomy, 2, 214, doi: [10.1038/s41550-017-0374-z](https://doi.org/10.1038/s41550-017-0374-z)
- Delrez, L., Gillon, M., Queloz, D., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10700, Ground-based and Airborne Telescopes VII, ed. H. K. Marshall & J. Spyromilio, 107001I, doi: [10.1117/12.2312475](https://doi.org/10.1117/12.2312475)

- Delrez, L., Murray, C. A., Pozuelos, F. J., et al. 2022a, *A&A*, 667, A59, doi: [10.1051/0004-6361/202244041](https://doi.org/10.1051/0004-6361/202244041)
- . 2022b, *A&A*, 667, A59, doi: [10.1051/0004-6361/202244041](https://doi.org/10.1051/0004-6361/202244041)
- Deming, D., Knutson, H., Kammer, J., et al. 2015, *ApJ*, 805, 132, doi: [10.1088/0004-637X/805/2/132](https://doi.org/10.1088/0004-637X/805/2/132)
- Dietrich, J., Apai, D., Schlecker, M., et al. 2023, *AJ*, 165, 149, doi: [10.3847/1538-3881/acba0b](https://doi.org/10.3847/1538-3881/acba0b)
- Dodd, M. S., Papineau, D., Grenne, T., et al. 2017, *Nature*, 543, 60, doi: [10.1038/nature21377](https://doi.org/10.1038/nature21377)
- Dransfield, G., Mékarnia, D., Triaud, A. H. M. J., et al. 2022, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12186, *Observatory Operations: Strategies, Processes, and Systems IX*, ed. D. S. Adler, R. L. Seaman, & C. R. Benn, 121861F, doi: [10.1117/12.2629920](https://doi.org/10.1117/12.2629920)
- Dressing, C. D., & Charbonneau, D. 2015, *ApJ*, 807, 45, doi: [10.1088/0004-637X/807/1/45](https://doi.org/10.1088/0004-637X/807/1/45)
- Ducrot, E., Sestovic, M., Morris, B. M., et al. 2018, *AJ*, 156, 218, doi: [10.3847/1538-3881/aade94](https://doi.org/10.3847/1538-3881/aade94)
- Ducrot, E., Gillon, M., Delrez, L., et al. 2020, *A&A*, 640, A112, doi: [10.1051/0004-6361/201937392](https://doi.org/10.1051/0004-6361/201937392)
- England, J. L. 2013, *The Journal of Chemical Physics*, 139, 121923, doi: [10.1063/1.4818538](https://doi.org/10.1063/1.4818538)
- Farmer, J. D. 2005, *Nature*, 436, 627, doi: [10.1038/436627a](https://doi.org/10.1038/436627a)
- Fletcher, R. 1987, *Practical methods of optimization: V. 1-2*, 2nd edn. (Chichester, England: John Wiley & Sons)
- Foreman-Mackey, D. 2016, *The Journal of Open Source Software*, 1, 24, doi: [10.21105/joss.00024](https://doi.org/10.21105/joss.00024)
- . 2018, *Research Notes of the American Astronomical Society*, 2, 31, doi: [10.3847/2515-5172/aaaf6c](https://doi.org/10.3847/2515-5172/aaaf6c)
- Foreman-Mackey, D., Agol, E., Ambikasaran, S., & Angus, R. 2017, *AJ*, 154, 220, doi: [10.3847/1538-3881/aa9332](https://doi.org/10.3847/1538-3881/aa9332)
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306, doi: [10.1086/670067](https://doi.org/10.1086/670067)
- Foreman-Mackey, D., Montet, B. T., Hogg, D. W., et al. 2015, *ApJ*, 806, 215, doi: [10.1088/0004-637X/806/2/215](https://doi.org/10.1088/0004-637X/806/2/215)
- Foreman-Mackey, D., Luger, R., Agol, E., et al. 2021, arXiv e-prints, arXiv:2105.01994. <https://arxiv.org/abs/2105.01994>
- Foster, J. J., Smolka, J., Nilsson, D.-E., & Dacke, M. 2018, *Proceedings of the Royal Society B: Biological Sciences*, 285, 20172322, doi: [10.1098/rspb.2017.2322](https://doi.org/10.1098/rspb.2017.2322)

- Freedman, R. S., Lustig-Yaeger, J., Fortney, J. J., et al. 2014, *ApJS*, 214, 25, doi: [10.1088/0067-0049/214/2/25](https://doi.org/10.1088/0067-0049/214/2/25)
- Freedman, R. S., Marley, M. S., & Lodders, K. 2008, *ApJS*, 174, 504, doi: [10.1086/521793](https://doi.org/10.1086/521793)
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, 616, A1, doi: [10.1051/0004-6361/201833051](https://doi.org/10.1051/0004-6361/201833051)
- Garcia, L. J., Moran, S. E., Rackham, B. V., et al. 2022a, *A&A*, 665, A19, doi: [10.1051/0004-6361/202142603](https://doi.org/10.1051/0004-6361/202142603)
- Garcia, L. J., Timmermans, M., Pozuelos, F. J., et al. 2022b, *MNRAS*, 509, 4817, doi: [10.1093/mnras/stab3113](https://doi.org/10.1093/mnras/stab3113)
- Gillon, M. 2018, *Nature Astronomy*, 2, 344, doi: [10.1038/s41550-018-0443-y](https://doi.org/10.1038/s41550-018-0443-y)
- Gillon, M., Jehin, E., Fumel, A., Magain, P., & Queloz, D. 2013, in *European Physical Journal Web of Conferences*, Vol. 47, *European Physical Journal Web of Conferences*, 03001, doi: [10.1051/epjconf/20134703001](https://doi.org/10.1051/epjconf/20134703001)
- Gillon, M., Jehin, E., Magain, P., et al. 2011, in *European Physical Journal Web of Conferences*, Vol. 11, *European Physical Journal Web of Conferences*, 06002, doi: [10.1051/epjconf/20101106002](https://doi.org/10.1051/epjconf/20101106002)
- Gillon, M., Jehin, E., Lederer, S. M., et al. 2016, *Nature*, 533, 221, doi: [10.1038/nature17448](https://doi.org/10.1038/nature17448)
- Gillon, M., Triaud, A. H. M. J., Demory, B.-O., et al. 2017, *Nature*, 542, 456, doi: [10.1038/nature21360](https://doi.org/10.1038/nature21360)
- Greene, T. P., Bell, T. J., Ducrot, E., et al. 2023, *arXiv e-prints*, arXiv:2303.14849, doi: [10.48550/arXiv.2303.14849](https://doi.org/10.48550/arXiv.2303.14849)
- Grimm, S. L., Demory, B.-O., Gillon, M., et al. 2018, *A&A*, 613, A68, doi: [10.1051/0004-6361/201732233](https://doi.org/10.1051/0004-6361/201732233)
- Hartmann, W. K. 2014, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372, 20130249, doi: [10.1098/rsta.2013.0249](https://doi.org/10.1098/rsta.2013.0249)
- Head, J. D., & Zerner, M. C. 1985, *Chemical Physics Letters*, 122, 264, doi: [https://doi.org/10.1016/0009-2614\(85\)80574-1](https://doi.org/10.1016/0009-2614(85)80574-1)
- Henry, T. J., Kirkpatrick, J. D., & Simons, D. A. 1994, *AJ*, 108, 1437, doi: [10.1086/117167](https://doi.org/10.1086/117167)
- Herbel, J., Kacprzak, T., Amara, A., Refregier, A., & Lucchi, A. 2018, *JCAP*, 2018, 054, doi: [10.1088/1475-7516/2018/07/054](https://doi.org/10.1088/1475-7516/2018/07/054)
- Hippke, M., David, T. J., Mulders, G. D., & Heller, R. 2019, *AJ*, 158, 143, doi: [10.3847/1538-3881/ab3984](https://doi.org/10.3847/1538-3881/ab3984)
- Hippke, M., & Heller, R. 2019, *A&A*, 623, A39, doi: [10.1051/0004-6361/201834672](https://doi.org/10.1051/0004-6361/201834672)
- Horne, K. 1986, *PASP*, 98, 609, doi: [10.1086/131801](https://doi.org/10.1086/131801)
- Howell, S. B. 2006, *Handbook of CCD Astronomy*, Vol. 5

- Huber, P. J. 1964, *The Annals of Mathematical Statistics*, 35, 73 , doi: [10.1214/aoms/1177703732](https://doi.org/10.1214/aoms/1177703732)
- Husser, T. O., Wende-von Berg, S., Dreizler, S., et al. 2013, *A&A*, 553, A6, doi: [10.1051/0004-6361/201219058](https://doi.org/10.1051/0004-6361/201219058)
- Iyer, A. R., Line, M. R., Muirhead, P. S., Fortney, J. J., & Gharib-Nezhad, E. 2023, *ApJ*, 944, 41, doi: [10.3847/1538-4357/acabc2](https://doi.org/10.3847/1538-4357/acabc2)
- Jarzynski, C. 2006, *Physics Today*, 59, 60, doi: [10.1063/1.2195318](https://doi.org/10.1063/1.2195318)
- Jenkins, J. M., Chandrasekaran, H., McCauliff, S. D., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7740, *Software and Cyberinfrastructure for Astronomy*, ed. N. M. Radziwill & A. Bridger, 77400D, doi: [10.1117/12.856764](https://doi.org/10.1117/12.856764)
- Kaltenegger, L., & Traub, W. A. 2009, *ApJ*, 698, 519, doi: [10.1088/0004-637X/698/1/519](https://doi.org/10.1088/0004-637X/698/1/519)
- Kingma, D. P., & Ba, J. 2017, *Adam: A Method for Stochastic Optimization*. <https://arxiv.org/abs/1412.6980>
- Kopal, Z. 1982, *Ap&SS*, 87, 149, doi: [10.1007/BF00648914](https://doi.org/10.1007/BF00648914)
- Kovács, G., Bakos, G., & Noyes, R. W. 2005, *MNRAS*, 356, 557, doi: [10.1111/j.1365-2966.2004.08479.x](https://doi.org/10.1111/j.1365-2966.2004.08479.x)
- Kovács, G., Hartman, J. D., & Bakos, G. Á. 2016, *A&A*, 585, A57, doi: [10.1051/0004-6361/201527124](https://doi.org/10.1051/0004-6361/201527124)
- Kovács, G., Zucker, S., & Mazeh, T. 2002, *A&A*, 391, 369, doi: [10.1051/0004-6361:2002080210.48550/arXiv.astro-ph/0206099](https://doi.org/10.1051/0004-6361:2002080210.48550/arXiv.astro-ph/0206099)
- Kreidberg, L., Bean, J. L., Désert, J.-M., et al. 2014, *Nature*, 505, 69, doi: [10.1038/nature12888](https://doi.org/10.1038/nature12888)
- Kreidberg, L., Line, M. R., Parmentier, V., et al. 2018, *AJ*, 156, 17, doi: [10.3847/1538-3881/aac3df](https://doi.org/10.3847/1538-3881/aac3df)
- Kuntschner, H., Kümmel, M., Walsh, J. R., & Bushouse, H. 2011, *Revised Flux Calibration of the WFC3 G102 and G141 grisms*, *Space Telescope WFC Instrument Science Report*
- Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. 2010, *The Astronomical Journal*, 139, 1782–1800, doi: [10.1088/0004-6256/139/5/1782](https://doi.org/10.1088/0004-6256/139/5/1782)
- Libby-Roberts, J. E., Schutte, M., Hebb, L., et al. 2023a, *arXiv e-prints*, arXiv:2302.04757, doi: [10.48550/arXiv.2302.04757](https://doi.org/10.48550/arXiv.2302.04757)
- . 2023b, *arXiv e-prints*, arXiv:2302.04757, doi: [10.48550/arXiv.2302.04757](https://doi.org/10.48550/arXiv.2302.04757)
- Line, M. R., Wolf, A. S., Zhang, X., et al. 2013, *ApJ*, 775, 137, doi: [10.1088/0004-637X/775/2/137](https://doi.org/10.1088/0004-637X/775/2/137)
- Lovelock, J. E. 1965, *Nature*, 207, 568, doi: [10.1038/207568a0](https://doi.org/10.1038/207568a0)
- Luengo, D., Martino, L., Bugallo, M., Elvira, V., & Särkkä, S. 2020, *EURASIP Journal on Applied Signal Processing*, 2020, 25, doi: [10.1186/s13634-020-00675-6](https://doi.org/10.1186/s13634-020-00675-6)

- Luger, R., Agol, E., Foreman-Mackey, D., et al. 2019, *The Astronomical Journal*, 157, 64, doi: [10.3847/1538-3881/aae8e5](https://doi.org/10.3847/1538-3881/aae8e5)
- Luger, R., Agol, E., Kruse, E., et al. 2016, *AJ*, 152, 100, doi: [10.3847/0004-6256/152/4/100](https://doi.org/10.3847/0004-6256/152/4/100)
- Luger, R., Foreman-Mackey, D., & Hedges, C. 2021a, arXiv e-prints, arXiv:2102.01697. <https://arxiv.org/abs/2102.01697>
- Luger, R., Foreman-Mackey, D., Hedges, C., & Hogg, D. W. 2021b, *AJ*, 162, 123, doi: [10.3847/1538-3881/abfdb8](https://doi.org/10.3847/1538-3881/abfdb8)
- Luger, R., Kruse, E., Foreman-Mackey, D., Agol, E., & Saunders, N. 2018, *AJ*, 156, 99, doi: [10.3847/1538-3881/aad230](https://doi.org/10.3847/1538-3881/aad230)
- Luger, R., Sestovic, M., Kruse, E., et al. 2017, *Nature Astronomy*, 1, 0129, doi: [10.1038/s41550-017-0129](https://doi.org/10.1038/s41550-017-0129)
- Lustig-Yaeger, J., Meadows, V. S., & Lincowski, A. P. 2019, *AJ*, 158, 27, doi: [10.3847/1538-3881/ab21e0](https://doi.org/10.3847/1538-3881/ab21e0)
- MacKay, D. J. C. 2003, *Information theory, inference and learning algorithms* (Cambridge, England: Cambridge University Press)
- Majeau, C., Agol, E., & Cowan, N. B. 2012, *ApJL*, 747, L20, doi: [10.1088/2041-8205/747/2/L20](https://doi.org/10.1088/2041-8205/747/2/L20)
- Mann, A. W., Feiden, G. A., Gaidos, E., Boyajian, T., & von Braun, K. 2015, *ApJ*, 804, 64, doi: [10.1088/0004-637X/804/1/64](https://doi.org/10.1088/0004-637X/804/1/64)
- Mayor, M., & Queloz, D. 1995, *Nature*, 378, 355, doi: [10.1038/378355a0](https://doi.org/10.1038/378355a0)
- Meadows, V. S., Reinhard, C. T., Arney, G. N., et al. 2018, *Astrobiology*, 18, 630, doi: [10.1089/ast.2017.1727](https://doi.org/10.1089/ast.2017.1727)
- Meierhenrich, U. J. 2013, *European Review*, 21, 190, doi: [10.1017/s106279871200035x](https://doi.org/10.1017/s106279871200035x)
- Ment, K., & Charbonneau, D. 2023, arXiv e-prints, arXiv:2302.04242, doi: [10.48550/arXiv.2302.04242](https://doi.org/10.48550/arXiv.2302.04242)
- Moffat, A. F. J. 1969, *A&A*, 3, 455
- Moran, S. E., Hörst, S. M., Batalha, N. E., Lewis, N. K., & Wakeford, H. R. 2018, *AJ*, 156, 252, doi: [10.3847/1538-3881/aae83a](https://doi.org/10.3847/1538-3881/aae83a)
- Morin, J., Donati, J. F., Petit, P., et al. 2010, *MNRAS*, 407, 2269, doi: [10.1111/j.1365-2966.2010.17101.x](https://doi.org/10.1111/j.1365-2966.2010.17101.x)
- . 2008, *MNRAS*, 390, 567, doi: [10.1111/j.1365-2966.2008.13809.x](https://doi.org/10.1111/j.1365-2966.2008.13809.x)
- Morley, C. V., Kreidberg, L., Rustamkulov, Z., Robinson, T., & Fortney, J. J. 2017, *ApJ*, 850, 121, doi: [10.3847/1538-4357/aa927b](https://doi.org/10.3847/1538-4357/aa927b)
- Morris, B. M. 2020, *ApJ*, 893, 67, doi: [10.3847/1538-4357/ab79a0](https://doi.org/10.3847/1538-4357/ab79a0)
- Morris, B. M., Agol, E., Davenport, J. R. A., & Hawley, S. L. 2018, *ApJ*, 857, 39, doi: [10.3847/1538-4357/aab6a5](https://doi.org/10.3847/1538-4357/aab6a5)

- Mosteller, F., & Tukey, J. W. 1977, *Data analysis and regression. A second course in statistics*
- Murray, C. A., Delrez, L., Pedersen, P. P., et al. 2020, *MNRAS*, 495, 2446, doi: [10.1093/mnras/staa1283](https://doi.org/10.1093/mnras/staa1283)
- Murray, C. D., & Correia, A. C. M. 2010a, in *Exoplanets*, ed. S. Seager, 15–23, doi: [10.48550/arXiv.1009.1738](https://doi.org/10.48550/arXiv.1009.1738)
- . 2010b, in *Exoplanets*, ed. S. Seager, 15–23, doi: [10.48550/arXiv.1009.1738](https://doi.org/10.48550/arXiv.1009.1738)
- Nutzman, P., & Charbonneau, D. 2008, *PASP*, 120, 317, doi: [10.1086/533420](https://doi.org/10.1086/533420)
- Pál, A. 2012, *MNRAS*, 420, 1630, doi: [10.1111/j.1365-2966.2011.20151.x](https://doi.org/10.1111/j.1365-2966.2011.20151.x)
- Pirzkal, N., Ryan, R., & Brammer, G. 2016, *Trace and Wavelength Calibrations of the WFC3 G102 and G141 IR Grisms*, Space Telescope WFC Instrument Science Report
- Pont, F., Zucker, S., & Queloz, D. 2006, *MNRAS*, 373, 231, doi: [10.1111/j.1365-2966.2006.11012.x](https://doi.org/10.1111/j.1365-2966.2006.11012.x)
- Protopapas, P., Jimenez, R., & Alcock, C. 2005, *MNRAS*, 362, 460, doi: [10.1111/j.1365-2966.2005.09305.x](https://doi.org/10.1111/j.1365-2966.2005.09305.x)
- Quirrenbach, A., Amado, P. J., Caballero, J. A., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, *Ground-based and Airborne Instrumentation for Astronomy V*, ed. S. K. Ramsay, I. S. McLean, & H. Takami, 91471F, doi: [10.1117/12.2056453](https://doi.org/10.1117/12.2056453)
- Rackham, B. V., Apai, D., & Giampapa, M. S. 2018, *ApJ*, 853, 122, doi: [10.3847/1538-4357/aaa08c](https://doi.org/10.3847/1538-4357/aaa08c)
- . 2019, *AJ*, 157, 96, doi: [10.3847/1538-3881/aaf892](https://doi.org/10.3847/1538-3881/aaf892)
- Rackham, B. V., Espinoza, N., Berdyugina, S. V., et al. 2023, *RAS Techniques and Instruments*, 2, 148, doi: [10.1093/rasti/rzad009](https://doi.org/10.1093/rasti/rzad009)
- Rajpurohit, A. S., Reylé, C., Allard, F., et al. 2013, *A&A*, 556, A15, doi: [10.1051/0004-6361/201321346](https://doi.org/10.1051/0004-6361/201321346)
- Ramsay, G., Doyle, J. G., & Doyle, L. 2020, *MNRAS*, 497, 2320, doi: [10.1093/mnras/staa2021](https://doi.org/10.1093/mnras/staa2021)
- Rasmussen, C. E., & Williams, C. K. I. 2005, *Gaussian processes for machine learning*, *Adaptive Computation and Machine Learning series* (London, England: MIT Press)
- Reiners, A., & Basri, G. 2009, *A&A*, 496, 787, doi: [10.1051/0004-6361:200811450](https://doi.org/10.1051/0004-6361:200811450)
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003, doi: [10.1117/1.JATIS.1.1.014003](https://doi.org/10.1117/1.JATIS.1.1.014003)
- Russell, H. N. 1906, *ApJ*, 24, 1, doi: [10.1086/141361](https://doi.org/10.1086/141361)
- Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. 2016, *PeerJ Computer Science*, 2, e55
- Schanche, N., Pozuelos, F. J., Günther, M. N., et al. 2022, *A&A*, 657, A45, doi: [10.1051/0004-6361/202142280](https://doi.org/10.1051/0004-6361/202142280)

- Sebastian, D., Gillon, M., Ducrot, E., et al. 2021, *A&A*, 645, A100, doi: [10.1051/0004-6361/202038827](https://doi.org/10.1051/0004-6361/202038827)
- Silva, A. V. R. 2003, *ApJL*, 585, L147, doi: [10.1086/374324](https://doi.org/10.1086/374324)
- Smith, E., & Morowitz, H. J. 2018, *The origin and nature of life on earth: The emergence of the Fourth geosphere* (Cambridge University Press)
- Solé, R. V., Munteanu, A., Rodriguez-Caso, C., & Macía, J. 2007, *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 1727, doi: [10.1098/rstb.2007.2065](https://doi.org/10.1098/rstb.2007.2065)
- Stetson, P. B. 1987, *PASP*, 99, 191, doi: [10.1086/131977](https://doi.org/10.1086/131977)
- Still, S., Sivak, D. A., Bell, A. J., & Crooks, G. E. 2012, *PhRvL*, 109, 120604, doi: [10.1103/PhysRevLett.109.120604](https://doi.org/10.1103/PhysRevLett.109.120604)
- Tamuz, O., Mazeh, T., & Zucker, S. 2005, *MNRAS*, 356, 1466, doi: [10.1111/j.1365-2966.2004.08585.x](https://doi.org/10.1111/j.1365-2966.2004.08585.x)
- Tody, D. 1986, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 627, *Instrumentation in astronomy VI*, ed. D. L. Crawford, 733, doi: [10.1117/12.968154](https://doi.org/10.1117/12.968154)
- Tsiaras, A., Waldmann, I. P., Rocchetto, M., et al. 2016, *ApJ*, 832, 202, doi: [10.3847/0004-637X/832/2/202](https://doi.org/10.3847/0004-637X/832/2/202)
- Van Grootel, V., Fernandes, C. S., Gillon, M., et al. 2018, *ApJ*, 853, 30, doi: [10.3847/1538-4357/aaa023](https://doi.org/10.3847/1538-4357/aaa023)
- Wakeford, H. R., Sing, D. K., Evans, T., Deming, D., & Mandell, A. 2016, *ApJ*, 819, 10, doi: [10.3847/0004-637X/819/1/10](https://doi.org/10.3847/0004-637X/819/1/10)
- Wakeford, H. R., Lewis, N. K., Fowler, J., et al. 2019, *AJ*, 157, 11, doi: [10.3847/1538-3881/aaf04d](https://doi.org/10.3847/1538-3881/aaf04d)
- Walker, S. I. 2017, *Physics Today*, 70, 58, doi: [10.1063/pt.3.3694](https://doi.org/10.1063/pt.3.3694)
- Walker, S. I., & Davies, P. C. W. 2013, *Journal of The Royal Society Interface*, 10, 20120869, doi: [10.1098/rsif.2012.0869](https://doi.org/10.1098/rsif.2012.0869)
- Wells, D. C., & Greisen, E. W. 1979, in *Image Processing in Astronomy*, ed. G. Sedmak, M. Capaccioli, & R. J. Allen, 445
- Wissner-Gross, A. D., & Freer, C. E. 2013, *Physical Review Letters*, 110, doi: [10.1103/physrevlett.110.168702](https://doi.org/10.1103/physrevlett.110.168702)
- Wolszczan, A. 1994, *Science*, 264, 538, doi: [10.1126/science.264.5158.538](https://doi.org/10.1126/science.264.5158.538)
- Zieba, S., & Kreidberg, L. 2022, *The Journal of Open Source Software*, 7, 4838, doi: [10.21105/joss.04838](https://doi.org/10.21105/joss.04838)