

Front page

Exam information

NFYK10020E - Physics Thesis 60 ECTS, Niels Bohr
Institute - Contract:124469 (Katrine Høier Frantzen)

Handed in by

Katrine Høier Frantzen
pjl124@alumni.ku.dk

Exam administrators

Eksamensteam, tel 35 33 64 57
eksamen@science.ku.dk

Assessors

Darach Jafar Watson
Examiner
darach@nbi.ku.dk
☎ +4535325994

vrs434 vrs434
Co-examiner
jerome@space.dtu.dk

Hand-in information

Title: Spectral analysis of the kilonova AT2017gfo using TARDIS and X-shooter data

Title, english: Spectral analysis of the kilonova AT2017gfo using TARDIS and X-shooter data

The sworn statement: Yes

Does the hand-in contain confidential material: No



SPECTRAL ANALYSIS OF THE KILONOVA AT2017GFO USING TARDIS AND X-SHOOTER DATA

MASTER'S THESIS

Written by *Katrine Høier Frantzen*

November 1, 2021

Supervised by

Darach Jafar Watson

UNIVERSITY OF COPENHAGEN



NAME OF INSTITUTE: Niels Bohr Institute

NAME OF DEPARTMENT: DAWN

AUTHORS: Katrine Høier Frantzen

EMAIL: pjl124@alumni.ku.dk

TITLE: Spectral analysis of the kilonova AT2017gfo using TARDIS and X-shooter data

SUPERVISORS: Darach Jafar Watson

HANDED IN: November 1, 2021

Abstract

On August 17 2017 the Advanced LIGO and Virgo interferometers measured the first gravitational waves signal from the merger of two neutron stars. The detection was followed by measurements of the associated electromagnetic radiation and thus became the first measurement of a so-called the kilonova. The kilonova was later found to originate from the coalescence of two neutron stars. These are thought to be a site of synthesis of heavy elements through rapid neutron capture, and for this reason this rare type of events has received increasing interest among astronomers during the last decades.

In this thesis we present a spectral analysis of the kilonova AT2017gfo based on synthetic spectra created using the radiative transfer code TARDIS and observed spectra measured by the X-shooter spectrograph. We use the Markov chain Monte Carlo method to compare the synthetic and observed spectra in order to place constraints on the distance, luminosity, density and the velocities of the photosphere and the outer edge. Of these parameters the distance is of great interest. We model the spectra measured during the first four epochs after the explosion and find that the first epoch spectrum gives the best and most reliable fit. As such, we find a best distance estimate to the transient of 38.58 ± 0.02 Mpc, which is consistent with distances found in the literature. We use the inferred distance estimates to calculate the Hubble constant and find a value between 71 and 77 $\text{km s}^{-1} \text{Mpc}^{-1}$. This is consistent with previous estimates based on measurements from the late universe, and this proves that the emission from neutron star mergers may be a useful tool for cosmology. Furthermore, we investigate the evolution of strontium (Sr) abundance and find that it decreases over time.

Despite leaving behind a number of improvements to be implemented in future work, the consistency between the distance estimate and previously determined estimates is striking considering the assumptions that were made throughout the analysis. This is promising for future work and proves that the method may become useful for spectral analysis of AT2017gfo and future kilonovae.

Acknowledgements

First of all, I would like to thank my supervisor Darach Watson for providing me with the opportunity to work on this topic. Thank you for the interesting and educational discussions, for all the great ideas and input, for challenging me and for the enthusiasm surrounding this work.

Furthermore, I would like to thank the people behind VALD for providing us with the full line list and ERDA Support for letting us use their facilities and for helping out with the parallelisation.

This research made use of TARDIS, a community-developed software package for spectral synthesis in supernovae ([Kerzendorf & Sim, 2014](#); [Kerzendorf et al., 2021](#)). The development of TARDIS received support from GitHub, the Google Summer of Code initiative, and from ESA's Summer of Code in Space program. TARDIS is a fiscally sponsored project of NumFOCUS. TARDIS makes extensive use of Astropy and Pyne.

Contents

1	Introduction	1
1.1	Supernovae and neutron stars	1
1.2	Kilonovae	4
1.3	Origin of the elements	7
1.4	The kilonova AT2017gfo	11
1.5	Thesis outline	14
2	Methods	15
2.1	Monte Carlo radiative transfer	15
2.2	TARDIS	15
2.3	Markov chain Monte Carlo	21
3	Creating synthetic spectra with TARDIS	24
3.1	Setting up the configuration files	24
3.2	Carsus and VALD	30
3.3	Parallelisation	34
4	Markov chain Monte Carlo	37
5	Results and discussion	41
5.1	Determining the Hubble constant, H_0	52
5.2	Constraining the strontium abundance	53
6	Future perspectives	58
7	Conclusion	63
	References	64
	Appendix A YAML example file	66
	Appendix B $D_L, L, \rho_0, v_{start}$ Corner plots Epoch 2-4	68
	Appendix C D_L, L, ρ_0, Sr Corner plots Epoch 2-4	71

1 Introduction

After the discovery of the first gravitational waves signal accompanied by electromagnetic emission, neutron star mergers have received increasing interest among astronomers. The interest lies in the fact that this rare type of events is a promising site of formation of some of the heaviest elements in the universe. Much research has been done to understand the nature of these events, but due to limited understanding a detailed study of the spectra of the kilonova still remains.

As a result of increasingly improved sensitivity of the gravitational waves detectors, more kilonova detections are expected in the future. For this reason a tool for detailed modeling of their spectra is of great interest. In this work we aim to investigate the properties of the kilonova AT2017gfo and to test and develop a new method for determining parameters of its spectra. To model the spectra we will use the radiative transfer code TARDIS to create synthetic spectra and compare these to the series of spectra measured by the X-shooter spectrograph. Modelling the spectra of AT2017gfo with TARDIS has previously been done with the aim of investigating the abundances of heavy elements (Smartt et al., 2017; Watson et al., 2019; Gillanders et al., 2021), and in these works it has proved itself as a useful tool. Motivated by these analyses we aim to develop a method using TARDIS in a more detailed and systematic manner than previously and to place constraints on the parameters of the kilonova.

The following section contains an introduction to the relevant theory necessary to understand the nature of kilonovae and how they relate to formation of heavy elements.

1.1 Supernovae and neutron stars

Kilonovae are transient events that take place when binaries of two massive objects, such as two neutron stars or a neutron star and a black hole, spiral towards each other and coalesce. In this work, we will investigate the inspiral and coalescence of two neutron stars, so to fully understand the nature of these events one first needs to understand the origin and nature of neutron stars.

A neutron star is one of the possible remnants after the death of a massive star. During the lifetime of a star it is said to be in hydrostatic equilibrium meaning that the pressure from nuclear fusion within its core balances out the gravitational pull. As a star enters the last stage of its evolution nuclear fusion stops and different scenarios can take place depending on the mass of the stellar core.

If the stellar core is below the Chandrasekhar mass, i.e. $1.4 M_{\odot}$ (Chandrasekhar, 1935), the star will collapse and form a white dwarf, which is a dense object supported against further collapse due to electron degeneracy pressure.

Stars with cores greater than the Chandrasekhar mass end their lives as core-collapse supernovae. At its last evolutionary stage such a star has an iron core and a layered structure of elements produced through nuclear fusion. Iron is a stable element which means that the binding energy is high, and this causes fusion to stop after the formation of iron leaving behind an iron-rich core. Fusion continues in the shells of the star, and as a result the iron core becomes progressively heavier. Since core fusion has stopped, there is no longer a pressure that supports the star from a gravitational collapse. The core is held by electron degeneracy, but as the mass increases to above the Chandrasekhar mass, the core starts to collapse further. During the collapse the temperature increases within a fraction of a second. Gamma-ray photons are released by the hot core, and these collide with the iron nuclei, breaking them apart into helium nuclei in a process called photodisintegration. The core contracts until the pressure becomes sufficiently high for the electrons to combine with the protons and produce neutrons and neutrinos:



This process is also known as inverse beta decay or electron capture. It takes place because the Fermi energy of the electrons increases, which means that the energy difference between the lowest and the highest energy level of the electrons within a small volume increases. This is a consequence of the Pauli exclusion principle, which will be described below. When the Fermi energy becomes greater than the mass energy difference between a proton and a neutron, the electrons will combine with protons and form neutrons and neutrinos. The neutrinos escape the contracting core, while the neutrons become packed at a very high density, and a neutron star is formed.

Neutron stars are some of the densest objects in the universe apart from black holes and hypothetical objects like quark stars, strange stars and white holes. They are relatively small with sizes of around ten kilometers (e.g. Capano et al., 2020). Unlike stars they do not produce energy through nuclear fusion, but instead they are supported from further gravitational collapse due to degenerate neutron pressure. This works in the same fashion as in the case of white dwarfs which are supported by electron degeneracy pressure. The phenomenon of degenerate matter is described by the Pauli exclusion principle (Pauli, 1925). In brief the principle states that in a closed system no fermions (e.g. electrons and neutrons) can have the same quantum number at the same time. This means that in a sufficiently small volume within a neutron star, all neutrons must occupy different energy states. As a consequence, when compressing more neutrons together within that

volume, some of the neutrons are forced to occupy high energy states. This causes the neutrons to move at higher speed, and at some point the fast moving neutrons provide a pressure against further compression.

Newly formed neutron stars are also known as pulsars due to their observational properties. As a star collapses, its angular momentum is conserved which means that the resulting neutron star has a short rotational period and spins up to 10^3 times per second (e.g. [Freedman et al., 2010](#)). The magnetic field of the progenitor star is conserved during the collapse, but since the size of the neutron star is much smaller than that of the star, the magnetic field becomes more concentrated. Charged particles are accelerated along the magnetic poles of the neutron star creating two beams of electromagnetic radiation along the magnetic axis. Due to the inclination between the magnetic axis and the rotational axis, the beams will sweep around the sky as the pulsar rotates. If the observer is located on the path of the beam, the signal will appear to pulse.

The upper bound to the mass of a cold non-rotating neutron star is called the Tolman-Oppenheimer-Volkoff limit ([Tolman, 1939](#); [Oppenheimer & Volkoff, 1939](#)). This is the mass that a neutron star can reach until the degenerate neutron pressure is no longer sufficient to support it against gravitational collapse. If the mass limit is exceeded, the neutron star will collapse to form a black hole or a denser object supported by degenerate matter (e.g. quarks ([Ivanenko & Kurdgelaidze, 1965](#))). Thus, the limit is analogous to the Chandrasekhar mass for white dwarfs, but due to the complexity of neutrons and their equation of state compared to that of electrons it is not as well constrained. [Bombaci \(1996\)](#) originally estimated the limit to lie between 1.5 and $3 M_{\odot}$ by investigating the equation of state for baryonic matter like neutrons. The limit was later estimated as 2.2 - $2.9 M_{\odot}$ ([Kalogera & Baym, 1996](#)) by reevaluating the equation of state. Motivated by the measurement of the gravitational waves event GW170817, which originated from the merger of two neutron stars (see [Section 1.4](#)), modern attempts have been made to constrain the maximum mass of a neutron star (e.g. [Margalit & Metzger, 2017](#); [Shibata et al., 2017](#); [Rezzolla et al., 2018](#); [Ruiz et al., 2018](#); [Shibata et al., 2019](#)). These estimates were based on the measurements of the gravitational waves signal ([Abbott et al., 2017a](#)), the resulting gamma-ray burst ([Abbott et al., 2017b](#)) and/or the electromagnetic counterpart observed by a wide range of telescopes. The results all lie in the range between 2.01 - $2.3 M_{\odot}$. This means that stellar cores of masses between the Chandrasekhar mass and the Tolman-Oppenheimer-Volkoff limit will collapse to form neutron stars, and with the newest constraints on the latter limit, it corresponds to cores of masses between 1.4 and $2.3 M_{\odot}$. Given the approximate mass of a neutron star of $2 M_{\odot}$ and a radius of around 10 km, one finds that the

density of a neutron star is close to that of an atomic nucleus ($2.3 \times 10^{17} \text{ kg/m}^3$):

$$\rho_{NS} = \frac{m_{NS}}{\frac{3}{4}\pi r_{NS}^3} \approx \frac{4 \times 10^{30} \text{kg}}{\frac{3}{4}\pi (10^4 \text{m})^3} \approx \frac{10^{30} \text{kg}}{10^{13} \text{m}^3} = 10^{17} \text{kg/m}^3 \quad (2)$$

While core-collapse supernovae remain the most popular explanation to the origin of neutron stars, another theory exists that suggests that neutron stars may also be formed from massive rapidly rotating white dwarfs (Schwab et al., 2017). Using the Zwicky Transient Facility, Caiazzo et al. (2021) discovered a white dwarf, dubbed ZTF J1901+1458, that exhibits the required properties. This massive white dwarf is likely to be the result of a merger of two white dwarfs with a total mass less than the Chandrasekhar mass. This made them unable to detonate as a Type 1a supernova, and they instead merged to form a massive white dwarf with a mass close to the Chandrasekhar limit, a small radius, a strong magnetic field and a short rotational period. The density of the white dwarf causes protons and electrons to combine inside the core, which will eventually lead to a thermonuclear collapse, since the density is further increased to a point where electron degeneracy can no longer support it. Whether this alternative explanation to the origin of neutron stars is true is still unknown and has not yet been proven. However, the relatively short distance to ZTF J1901+1458 ($\sim 41 \text{ pc}$) implies that it is not a rare object, and thus upon confirmation of this type of origin to neutron stars the total number of neutron stars in the galaxy would increase remarkably.

1.2 Kilonovae

As previously mentioned, kilonovae are transient events that take place when binaries of two compact objects, such as two neutron stars or a neutron star and a black hole, spiral towards each other and coalesce. They got their name due to the fact that their peak brightness is one thousand times greater than that of a regular nova (Metzger et al., 2010). However, they are still ten to one hundred times fainter than supernovae.

Due to their compactness, collisions between two neutron stars, two black holes and a neutron star and a black hole result in bursts of gravitational waves. Collisions of binaries of two black holes are not expected to produce any electromagnetic radiation, since black holes do not consist of matter. However, in the case of a collision involving a neutron star the gravitational waves signal is accompanied by electromagnetic radiation. This emission initially comes in the form of a short gamma-ray burst and later as a fainter and more long-lived transient called the kilonova. While the short gamma-ray bursts are known to last for $< 2\text{s}$ (e.g. Metzger, 2019), the kilonova is usually visible for days or weeks and is powered by the

radioactive decay of r-process elements synthesized during the merger (see Section 1.3).

The formation of a binary of compact objects starts with two stars with stellar cores heavier than the Chandrasekhar mass orbiting each other. As these stars enter their last evolutionary stage, they will detonate as core collapse supernovae, and provided that these supernovae do not disrupt the binary system and depending on the masses of the progenitor stars, the end result will be a binary of either two neutron stars, two black holes or a neutron star and a black hole. Another way of obtaining such a binary is if two of these objects associate in a dense globular cluster and start orbiting their mutual center of mass. Over time, the objects in a binary will inspiral towards each other while releasing energy as gravitational radiation until they merge and detonate. Since this thesis is an analysis of a binary of two neutron stars, focus will be on this type of binary for the rest of this section.

As mentioned above neutron star collisions are thought to be progenitors of short gamma-ray bursts (GRBs) (e.g. Paczynski, 1986; Eichler et al., 1989). These GRBs are thought to be the result of material falling into the black hole formed by the merger. This produces a relativistic outflow of material formed in a beam, and for that reason it depends on the viewing angle of the observer whether the GRB is observable or not (see Figure 1). This especially becomes clear when comparing the rate of short GRBs in the local universe of $\sim 5 \text{ Gpc}^{-3} \text{ yrs}^{-1}$ (Wanderman & Piran, 2015) to the much higher rate of neutron star mergers of $\sim 10^3 \text{ Gpc}^{-3} \text{ yrs}^{-1}$ (Abbott et al., 2019a).

The result of a binary neutron star merger is a black hole or a massive neutron star surrounded by a disk, and the type of remnant depends on the total mass of the progenitor neutron stars. If the total mass is above $\sim 2.6 - 3.9 M_{\odot}$ (Hotokezaka et al., 2011; Bauswein et al., 2013a) a black hole is formed immediately after the merger, and if the mass is less than this threshold the resulting remnant will be a rapidly rotating neutron star. Due to its rotation its mass can exceed the Tolman-Oppenheimer-Volkoff limit which only applies to non-rotating neutron stars. The neutron star can either be hypermassive, supramassive or stable depending on its mass and rotation. The hypermassive neutron star is supported against collapse by its differential rotation and thermal pressure, but due to cooling and loss of differential rotation, it will collapse into a black hole in less than a second (Baumgarte et al., 2000). A supramassive neutron star is less massive than a hypermassive neutron star, and it is supported by its solid body rotation (Cook et al., 1992). It slows down over time and will collapse into a black hole in 100-1000 seconds. If the total mass is below the Tolman-Oppenheimer-Volkoff limit the resulting remnant will be a neutron star that never collapses into a black hole. However, this scenario is unlikely to occur (e.g. Metzger, 2019).




Figure 1: Overview of the electromagnetic counterparts of a kilonova. Based on figure by O. Korobkin.

The emission of a kilonova can be considered as a combination of two components: A red component and a blue component (Metzger & Fernández, 2014). These are illustrated in Figure 1. The red component comes from lanthanide-rich ejecta ($Z = [57,71]$) heated through radioactive decay, it is visible at NIR wavelengths and lasts for weeks. It originates from matter dynamically ejected during the merger in the equatorial plane and from spherical outflows from the accretion disk surrounding the remnant. Here the electron fraction is low giving a neutron-rich matter capable of forming heavy lanthanides. The production of lanthanides leads to a larger opacity, which that gives the lanthanide-rich ejecta its red color.

The blue component comes from lanthanide-poor ejecta, it lasts for approximately a day and is observable at UV/visible wavelengths. The blue component has a high electron fraction and enough neutrons to synthesize r-process elements. However, the ratio of neutrons to lighter element nuclei is too small to reach the heavy lanthanide elements of $A > 140$. The blue component comes from outflows from the accretion disk (e.g. Miller et al., 2019), but it is especially pronounced when the merger remnant is a massive neutron star. Due to the high temperature of the neutron star, the production of electron-positron pairs is enhanced. These interact with the neutrons and form protons, neutrinos and anti-neutrinos, which are ejected spherically from the remnant as a neutrino wind. The neutrinos and anti-neutrinos interact with the protons and neutrons causing the before neutron-rich ejecta to consist of an equal amount of protons and neutrons. In many cases

the disk around the remnant blocks the spherical blue emission, causing it to be primarily visible in the polar direction (Kasen et al., 2015).

1.3 Origin of the elements

As previously mentioned, kilonovae have received an increasing interest and have become subject to extensive research during the last decades. One of the main reasons is the fact that neutron star mergers are thought to be the predominant site of formation of neutron-rich heavy elements synthesized through rapid neutron capture (r-process elements). Elements are synthesized through a wide range of processes, and this section contains a brief overview of the most important ones resulting in elements of different masses.

Elements lighter than iron are formed through nucleosynthesis in the cores of stars during their lifetimes. As stars are formed they consist primarily of hydrogen, but inside their cores the temperature is sufficiently high for hydrogen to fuse into helium through thermonuclear reactions. This process releases energy, which provides a pressure that supports the star against gravitational collapse. For massive stars nucleosynthesis continues into progressively heavier elements throughout their lifetimes. As a massive star has converted all of its hydrogen into helium, its core will contract due to the lack of pressure from core fusion, and this causes the temperature to increase until it becomes high enough for helium to fuse into carbon and oxygen. The structure of the star is now such that the core consists of carbon and oxygen and outside the core it has two layers of hydrogen and helium respectively. The process of nucleosynthesis continues for the most massive stars, and by the last evolutionary stage of such a star it will have an iron core and a layer-like structure of lighter elements. Since iron is a stable element, core fusion stops, and the star will eventually detonate as a supernova explosion. During this explosion, the elements are thrown into the interstellar medium and may become part of the gas that constitute a newly formed star in the proximity of the site of the supernova.

Elements heavier than iron are synthesized through neutron capture, where an atomic nucleus collides with a neutron to form a heavier nucleus. The most predominant mechanisms for neutron capture are the s-process and r-process (Figure 2), and these are thought to be responsible for the formation of the majority of the heaviest elements in the universe.

The s-process elements are formed through slow neutron capture. During this process a nucleus captures a neutron and forms a new isotope, and if the isotope is unstable it undergoes a beta decay. The result is an element with one higher atom number. The s-process has its name due to the fact that there is sufficient




Figure 2: The two predominant mechanisms for neutron capture and their timescales. Figure from 5W Infographics.

time for the decay to happen before it captures another neutron. This contrasts to elements formed through r-process, where neutrons are added to the atom so rapidly that even if the newly formed element is unstable it does not have time to undergo a beta decay before another neutron is captured. These two mechanisms are illustrated in Figure 2.

The two neutron capture processes require different environments to take place. s-process elements are formed where the neutron flux is low but sufficient for the capture of a neutron every $100-10^5$ year (Burbidge et al., 1957). Such an environment is present within stars and especially asymptotic giant branch (AGB) stars (e.g. Busso et al., 1999), where iron seeds left by a supernova and later captured during the formation of the star interact with neutrons and form heavier elements. Contrary to this the formation of r-process elements require an environment where the neutron flux is sufficiently high for an element to catch around $100-10^4$ neutrons per second (Burbidge et al., 1957). The exact site of r-process is still unknown, but one of the favoured candidates is the neutron-rich outflow that is produced during the rupture of a neutron star when it collides with another neutron star or a black hole (e.g. Lattimer & Schramm, 1974; Symbalisty & Schramm, 1982). It is due to the radioactive decay of these r-process elements that the kilonova produces its electromagnetic radiation (Li & Paczyński, 1998).

When studying r-process elements it is of great interest to know the expected distribution of the atoms formed in this manner. Such a distribution has been determined by considering the abundances of heavy elements in both the solar system and metal-poor stars, and all estimates show the same pattern (e.g. [Burbidge et al., 1957](#); [Snedden et al., 2003](#); [Kratz et al., 2007](#)) (Figure 3). The solar system contains elements formed through various nucleosynthesis processes taking place before the Sun was formed. The abundance of r-process elements has then been calculated as the total abundance of heavy elements with the s-process contribution subtracted. This contribution is known from theoretical calculations and was first done by [Burbidge et al. \(1957\)](#).

Furthermore, measurements of the abundances of heavy elements in metal-poor stars provide an estimate of the distribution of r-process elements ([Truran, 1981](#)). These stars have been formed at early times of the universe and do not contain elements formed through s-process. This is due to the fact that these elements are synthesized in AGB stars and for this reason the process had not yet kicked in and enriched the interstellar medium when the metal-poor stars were formed. The distributions of r-process elements obtained from the solar system and metal-poor stars show the same pattern, and Figure 3 shows the distribution of r-process elements obtained from the solar system. As seen in the figure the distribution has three peaks around atomic masses of $A \approx 80, 130$ and 196 , which corresponds to elements around selenium (Se), xenon (Xe) and platinum (Pt) respectively. For this reason the search for r-process elements are often concentrated around these peaks.




Figure 3: Distribution of solar r-process abundances calculated as the total abundance of heavy elements ([Lodders et al., 2009](#)) with the s-process contribution subtracted ([Bisterzo et al., 2014](#)).

While the site of formation of s-process elements is relatively well-known, the exact site(s) of r-process remains an unanswered question. Historically, numerous candidates have been suggested, but studies on these environments have ruled out the majority of them leaving two favoured candidates: Neutron star mergers and core-collapse supernovae.

Core-collapse supernovae were long thought to be the predominant site of r-process synthesis. However, only a fraction of these explosions have the right conditions for r-process synthesis (Podsiadlowski et al., 2004), and for that reason they cannot account for the total r-process abundance in a galaxy. To obtain the observed abundance of heavy r-process elements ($A > 140$) in a galaxy, the required production rate of this type of elements is $\sim 2 \times 10^{-7} M_{\odot} \text{yr}^{-1}$ (Metzger, 2019). The Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo collaboration found a volumetric rate of neutron star mergers of 110 - 3840 $\text{Gpc}^{-3} \text{yrs}^{-1}$ (Abbott et al., 2019a). This implies that such a merger occurs once every $\sim 6 - 120$ years. Given these numbers, one can infer that to explain the r-process abundance in a galaxy it requires that the mass of the ejected r-process elements per event is equal to (Metzger, 2019):

$$M_r \approx 10^{-2} M_{\odot} \left(\frac{\mathcal{R}_{NS-NS}}{10 \text{yr}^{-1}} \right)^{-1} \quad (3)$$

where \mathcal{R}_{NS-NS} is the rate of neutron star detections by The Advanced LIGO/Virgo interferometers. Numerical calculations (e.g. Bauswein et al., 2013b; Hotokezaka et al., 2013) and estimates based on the measurement of GW170817 (e.g. Kasen et al., 2017) show results consistent with this mass, meaning that neutron star mergers is likely the predominant source of r-process synthesis.

Despite being proved as an important site of formation of r-process elements neutron star mergers cannot explain the early r-process enrichment of the universe and the r-process abundances found in metal-poor stars. This is due to the fact that neutron star mergers happen on large timescales since they require the formation and inspiral of a binary, which means that the metal-poor stars were formed before the first neutron star merger had taken place. This indicates the existence of at least one additional source of r-process elements and has lead to the suggestion that energetic types of supernovae called magnetorotational hypernovae are important contributors (Yong et al., 2021). A magnetorotational hypernove is a type of energetic supernova originating from the death of a rapidly spinning and highly magnetized star with a mass above $25 M_{\odot}$, and these may have the required neutron-rich environment necessary to form r-process elements. By investigating the abundance pattern of an extremely metal-poor star Yong et al. (2021) found an overall low metallicity but a large enhancement of heavy elements. By comparing this measurement to theoretical r-process abundance values resulting from a mag-

netorotational hypernova they found the distribution patterns to be consistent. The abundance of certain heavy elements like thorium (Th) and uranium (U) was found to be higher than the values predicted from a neutron star merger, and the star contained lighter elements which cannot be produced through a merger. As such, they found that the abundance pattern from a magnetorotational hypernova reproduced the pattern found in the star better than the abundance pattern found from a neutron star merger. This suggests that the enrichment of the star was caused by a hypernova rather than a neutron star merger.

In summary, the required environment for r-process puts strong limits to its potential sites, but in spite of these limits the origin of the process remains a challenging and unsolved problem. Even though most recent research on the topic revolves around either neutron star mergers or extreme core-collapse supernovae, the contribution from each of these two events to the total r-process abundances has yet to be proved. This makes both core-collapse supernovae and neutron star mergers interesting in the search for r-process elements.

1.4 The kilonova AT2017gfo

The first simultaneous detection of a gravitational waves (GW) signal and wide-band electromagnetic radiation took place on August 17 2017 (Abbott et al., 2017a,b). Prior to the discovery a number of GW events had already been detected (Abbott et al., 2019a), but these events all originated from binaries of two black holes, and for that reason they were not expected to produce any electromagnetic radiation. Thus, the event in 2017, which was later associated with the merger of two neutron stars, became the beginning of an era for multi-messenger astronomy.

The gravitational waves signal resulting from the merger was measured by the Advanced LIGO and Virgo interferometers, and it was designated GW170817. Approximately 1.7 seconds after the merger, the signal was followed by a short gamma-ray burst GRB (GRB170817A) lasting for ~ 2 seconds and measured by the *Fermi* and *INTEGRAL* satellites (Abbott et al., 2017b). This detection together with the detection of another GRB (GRB150101B) in 2015, which was later suggested to be associated with the kilonova, supports the theory that kilonovae are the progenitors of short GRBs.

Less than 11 hours after the GW and GRB measurements, a transient (SSS 17a, later AT2017gfo) was discovered in the elliptical galaxy NGC 4993 in the area of 31 deg^2 indicated by the GW signal (Abbott et al., 2017b). It was measured by a wide range of observatories covering wavelengths from radio to X-rays. Among these was the X-shooter spectrograph mounted on the Very Large Telescope (VLT), which covered wavelengths from 3300\AA to $2.5\mu\text{m}$ (Pian et al., 2017; Smartt et al., 2017). Ten spectra were measured with one day intervals, and these are all plotted

in Figure 4 together with the time after the explosion. As seen in the figure, the transient started out blue in color with a brightness peak at UV wavelengths, and it later evolved to become more red and with a peak at NIR wavelengths.




Figure 4: Spectral data from the ten epochs measured by X-shooter. The data was downloaded from the ENGRAVE web page¹.

One aim of this thesis is to investigate the distance to the kilonova. Numerous attempts have already been made to try to determine the distance using different methods than the one that will be used in this work. Using the gravitational waves signal, [Abbott et al. \(2017a\)](#) initially found a luminosity distance to the source of 40.8^{+8}_{-14} Mpc considering all potential sky locations found by the GW signal. This distance was later revised assuming that the optical counterpart represents the true sky location, and it was then estimated as $43.8^{+2.9}_{-6.9}$ Mpc ([Abbott et al., 2017](#)). Prior to the event the I-band Tully-Fisher distance to the cluster ESO 508, in which the host galaxy NGC 4993, is located had been estimated as 41.1 ± 5.8 Mpc ([Sakai et al., 2000](#); [Freedman et al., 2001](#)). Using a distance calculated from Hubble’s law and the cosmological redshift of the galaxy (40.4 ± 3.4 Mpc) in combination

¹<http://www.engrave-eso.org/AT2017gfo-Data-Release>

with the fundamental plane (FP) distance (44.0 ± 7.5 Mpc), [Hjorth et al. \(2017\)](#) found a total distance to NGC 4993 of 41.0 ± 3.1 Mpc. Investigating the surface brightness fluctuations of the galaxy the distance has furthermore been estimated as $40.7 \pm 1.4 \pm 1.9$ Mpc, where the first error is the random error and the second one is the systematic error ([Cantiello et al., 2018](#)). Finally, considering the luminosity function of the globular clusters within NGC 4993 the distance has been estimated as 41.65 ± 3.00 Mpc ([Lee et al., 2018](#)), but given that the systematic error is likely underestimated, the uncertainty can be as high as ± 6 Mpc.

Attempts have been made to constrain the presence of r-process elements in the spectra of the transient. As seen in Figure 4, the first epoch spectrum shows two pronounced absorption features around 3500\AA and 8100\AA , and these were initially suggested to originate from absorption from CsI and TaI ([Smartt et al., 2017](#)). This was later ruled out by [Watson et al. \(2019\)](#), who argued that CsI and TaI are not expected to produce strong lines at 4800K, which is the approximate temperature of the kilonova determined by fitting the spectrum with a single temperature black body model. [Watson et al. \(2019\)](#) found that the features are more likely to originate from strontium (Sr), and that the feature around 8100\AA comes almost exclusively from three strong SrII lines at rest frame wavelengths of $\sim 1\mu\text{m}$. These three lines can be seen in the Sr level diagram in Figure 5. The discovery of Sr became the first direct evidence that r-process elements are formed in neutron star mergers and that neutron stars are comprised of neutron-rich matter.




Figure 5: Level diagram for Sr². The two lines at 674 nm and 687 nm are forbidden and have not been detected in the spectrum of AT2017gfo.

²<https://mike.depalatis.net/ionptable/strontium/>

1.5 Thesis outline

In this work we present a spectral analysis of the electromagnetic radiation arising from the kilonova AT2017gfo. We consider spectra measured by the X-shooter spectrograph and use the radiative transfer code TARDIS in combination with the Markov chain Monte Carlo method to place constraints on a number of parameters.

In Section 2 we present the methods used to perform the analysis including the radiative transfer code TARDIS and Markov chain Monte Carlo (MCMC). In Section 3 we describe the creation of synthetic spectra with TARDIS and in Section 4 we present the analysis of the synthetic spectra with MCMC. The results of the analysis are presented in Section 5 together with a discussion of these. Section 6 contains a discussion on improvements and future work on this method, and in Section 7 we summarize our results.

2 Methods

The following section describes the methods used to do the analysis. This includes a brief introduction to Monte Carlo radiative transfer (Section 2.1), which is a keystone to the radiative transfer code TARDIS used to create synthetic spectra (Section 2.2). Finally, Section 2.3 contains a description of the Markov chain Monte Carlo method which will be used to constrain parameters of the kilonova AT2017gfo.

2.1 Monte Carlo radiative transfer

Radiative transfer is the study of how light propagates through a medium. During its propagation the light is subject to light-matter interactions which causes the light to either scatter or be absorbed by the medium. Due to the large number of interactions, solving the radiative transfer problem analytically requires approximations and becomes challenging in many astronomical cases as the complexity of the problem increases. For this reason one often benefits from numerical tools and especially Monte Carlo methods.

The idea behind Monte Carlo radiative transfer methods is to simulate and trace the propagation of a number of test particles (or packets of energy) through a medium. As will be discussed below, these packets behave like real photons as they propagate through the medium. They are initialized according to the properties of the medium, and after they are launched they start interacting with the surroundings undergoing scattering and absorption. The frequency and type of interactions are likewise dictated by the properties of the medium. Using a sufficiently large number of energy packets, they can resemble the total number of emitted photons, and by examining the sample, one can put constraints on the properties of the entire radiation field.

2.2 TARDIS

TARDIS³ (**T**emperature **A**nd **R**adiative **D**iffusion **I**n **S**upernovae) is a Monte Carlo radiative transfer spectral synthesis code designed for 1D spectral modelling of supernova (SN) ejecta (Kerzendorf & Sim, 2014; Vogl et al., 2019). Hence, TARDIS is used to simulate and analyze the spectra of transient astronomical events such as supernovae and kilonovae, and by comparing these simulated models to observed spectra, one can put constraints on the composition, mass and explosion energy of the object.

³<https://tardis-sn.github.io/tardis/index.html>

As a consequence of modelling the explosion in one dimension, TARDIS assumes that the object expands homologously. This means that the explosion takes place at a starting point and expands radially outwards such that the velocity of the ejecta is proportional to the distance from the ejecta to the central point. TARDIS calculates the propagation of photons travelling from the photosphere (the inner boundary) to the edge of the ejecta (the outer boundary), and this is where the spectrum is measured. The computational domain can be seen in Figure 6. The velocity of the two edges are defined in the TARDIS configuration file, which will be explained in detail below, and the radii of the boundaries are defined by the time of explosion such that $r_{inner} = v_{inner}t_{exp}$ and $r_{outer} = v_{outer}t_{exp}$. The space between the inner and outer boundary is divided into shells in which the state of the plasma is constant. Furthermore, TARDIS assumes that there is no time dependence meaning that the generated spectrum is a snapshot at a given time and the radii of the shells are constant throughout the simulation.

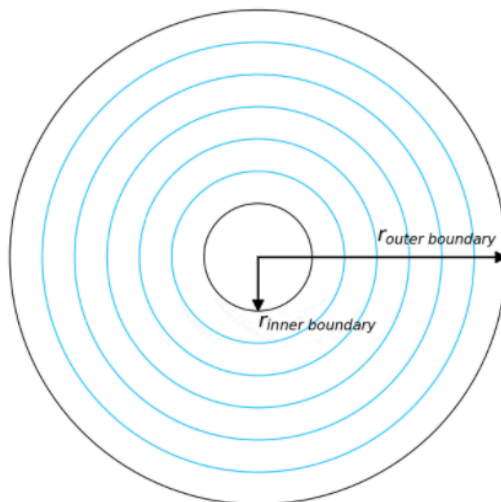


Figure 6: The computational domain of the TARDIS simulations showing the spherical structure with inner and outer radius and a number of shells. The figure was taken from the TARDIS documentation.⁴

To run a simulation in TARDIS, one first needs to set up a YAML configuration file containing all the parameters necessary to run the simulation. This file is divided into multiple sections with parameters pertaining to different parts of the simulation. The most important sections and the ones relevant for this analysis are the following:

⁴<https://tardis-sn.github.io/tardis/physics/setup/model.html>

Supernova configuration: Contains key information about the observational properties of the object, e.g. the requested luminosity and the time since explosion. TARDIS uses the requested luminosity to calculate the temperature of the explosion, and it will aim for the output spectrum to converge towards this luminosity.

Atomic data: Contains the name of the HDF5 file which has a list of the relevant atomic lines. The TARDIS built-in atomic database is the Kurucz line list (Kurucz, 2017), but for this analysis we use the Vienna Atomic Line Database (VALD3) (Piskunov et al., 1995; Ryabchikova et al., 2015). The implementation of this line list is described in Section 3.2.

Plasma configuration: Contains the relevant parameters used by TARDIS to calculate the state of the plasma. This includes information about the ionization, excitation and radiative rate treatment modes and the line interaction mode.

Model configuration: Section of parameters describing the geometry of the explosion and properties of the ejecta. This includes the number of shells that the ejecta should be divided into, the velocity shell structure, density and abundances.

Monte Carlo configuration: Contains the parameters determining the Monte Carlo transport and convergence criteria. This includes the number of energy packets used per iteration and the maximum number of iterations executed during the simulation before it ends. Furthermore, this section specifies the number of packets and virtual packets used for the last iteration, which is what determines the final spectrum.

Spectrum configuration: Contains the parameters needed to generate the spectrum. This includes the wavelength range of the output spectrum and the number of bins. The latter governs the resolution of the spectrum.

As previously mentioned, Monte Carlo radiative transfer codes like TARDIS usually use energy packets to represent light rather than real photons. These energy packets (also called real packets or r-packets) represent bundles of photons of the same frequency, and as they propagate through the ejecta they behave in the same manner as photons. The reason for using energy packets rather than individual photons is that tracking the propagation of these has some numerical advantages. In TARDIS the packets all have the same energy, and the total energy of the packets is 1 erg. This means that the energy of each packet is $1/N$, where N is the number of packets used for the simulation. Since the energies of the packets are conserved in the co-moving frame and the frequencies of the photons

that they represent are different, the packets will represent different numbers of photons throughout the simulation.

A TARDIS simulation works by launching packets at the inner boundary, i.e. the photosphere. The packets are initialized by assigning them a frequency of the constituent photons and a propagation direction. The frequencies are sampled from a black-body distribution at the temperature of the inner boundary, and if not specified in the configuration file, this temperature is calculated as:

$$T_{inner} = \left(\frac{L_{req}}{4\pi r_{inner}^2 \sigma_B} \right)^{1/4} \quad (4)$$

where L_{req} is the requested luminosity from the configuration file, r_{inner} is the radius of the photosphere as previously described, and σ_B is the Stefan–Boltzmann constant. Apart from a frequency each packets is assigned an initial direction. These are assigned to them to avoid limb-darkening, which is the effect where the central part of the object appear brighter than the edges.

After the initialization and launch of the packets, they start propagating along their assigned directions on straight trajectories. TARDIS calculates the distance from each packet to the next shell together with the distance that each packet has to travel before it interact with the surrounding material though either electron scattering or line absorption. The shortest distance determines whether the packet moves to the next shell or if it moves to the location of the interaction.

The optical depth of a packet (τ) determines the probability of a radiation-matter interaction. As each packet propagates through the medium it accumulates optical depth, and after accumulating an optical depth $\Delta\tau$ the probability of an interaction is $1 - e^{-\Delta\tau}$. To be able to model this, each packet is initially assigned an optical depth $\tau_{interaction} = -\ln(z)$, where z is a random number between 0 and 1. As their accumulated optical depth reaches this values, an interaction is performed.

There are two ways that a packet can interact with the matter of the ejecta: Through electron scattering or line interaction. TARDIS treats all electron scatterings as Thomson scatterings, where the kinetic energy of the particle and the frequency of the packet do not change. The optical depth that a packet accumulates over a distance l due to electron scattering is:

$$\Delta\tau = \sigma_T n_e l \quad (5)$$

where σ_T is the Thomson cross section and n_e is the free electron density.

An interaction with an atom can only take place if the frequency of the packet in the co-moving frame is equal to the energy difference required for a particular

transition between two energy levels. As each packet travels through the expanding ejecta, the frequency in the co-moving frame is either redshifted or blueshifted depending on the propagation direction of the packet. As such, if the packet is travelling towards the outer boundary, and hence in the same direction as the expansion, its frequency is redshifted due to the Doppler effect. On the other hand, the frequency is blueshifted if the packet is travelling towards the inner boundary and in opposite direction of the expansion. This also means that throughout the simulation the frequency of the packet may cause several transitions (i.e. it may come into resonance with many line transitions). The optical depth accumulated by a packet due to line interactions is calculated using the Sobolev approximation. In brief, this approximation applies to systems like supernova/kilonova ejecta, where the velocity gradient is sufficiently large for local velocity variations to be negligible compared to variations on larger scales. This means that the radiative transfer problem can be solved locally, which simplifies the problem, and as a consequence a packet can only interact with an atom if its co-moving frequency is exactly equal to the energy required for a particular transition to take place. TARDIS calculates the distance to the next so-called Sobolev point, which is the distance required for a packet to travel for its frequency to be red/blueshifted into the Sobolev resonance of the closest line. If a line interaction takes place, the packet instantaneously accumulates the optical depth $\tau_{Sobolev}$.

To determine whether a packet will move to the next shell or become subject to either an electron scattering or an atomic line interaction, TARDIS first calculates the distance that the packet needs to travel to reach the next Sobolev line resonance. As the packet travels towards the line resonance, it accumulates optical depth due to the possibility of electron scattering. If it reaches the Sobolev point, the Sobolev optical depth is added to the optical depth of the packet. This process is repeated to the next Sobolev point until the packet reaches the randomly assigned optical depth $\tau_{interaction}$. The point at which this happens determines the type of the next interaction. Meanwhile, the distance to the next shell is tracked. As previously mentioned, the state of the plasma is constant within each shell, which means that whenever a packet enters the next shell, quantities of the new plasma state is calculated, and this will have an impact on the radiation-matter interactions. If a packet enters the next shell before accumulating enough optical depth for an interaction to occur, the plasma state is re-evaluated with the parameters of the new shell and the accumulation of optical depth continues. When an interaction takes place, the packet moves to the location of the interaction, the interaction is performed and the packet is assigned a new direction and $\tau_{interaction}$. The propagation of a packet within a single shell is illustrated in the diagram in Figure 7. If the randomly assigned $\tau_{interaction}$ has one of the values in case (I), the optical depth will be reached between two Sobolev points due to electron scat-

tering. In this case the packet will therefore scatter and be re-assigned a new $\tau_{interaction}$ and direction. In case (II) the optical depth is reached at a Sobolev point, and hence a line interaction will take place. Case (III) is when the optical depth is not reached before the packet enters the next shell.

The propagation process is repeated until the packet reaches either the inner or the outer boundary. In the case of hitting the inner boundary the propagation of the packet terminates and the packet is discarded. In the case of reaching the outer boundary, the contribution of the packet to the emergent spectrum is measured and the propagation is likewise terminated.




Figure 7: Diagram of how a packet accumulates optical depth as it propagates through a shell. The randomly assigned optical depth determines if the packet will scatter (I), interact with an atom (II) or move to the next shell (III). The figure was taken from the TARDIS documentation⁵.

The spectra generated by TARDIS suffer from Monte Carlo noise meaning that the noise scales with the number of packets N as $1/\sqrt{N}$. This means that to obtain a spectrum with less noise one needs to increase the number of packets, and this will quickly become computationally expensive. To solve this problem, TARDIS has the option to create spectra using virtual packets (v-packets) instead of r-packets. Whenever an r-packet is launched or is subject to an event such as a line interaction or electron scattering, a number of v-packets are generated at the site of the event. These v-packets propagate through the atmosphere in the same fashion as the r-packets but differs from these by never changing their trajectories. Once the v-packets pass the outer boundary, their contribution to the final spectrum is

⁵<https://tardis-sn.github.io/tardis/physics/montecarlo/propagation.html>

measured and is weighted by the optical depth which has been accumulated by the packet through electron scattering and line interactions. Since the determination of the propagation path of the r-packets is computationally expensive, introducing v-packets that propagate along straight line decreases the runtime and allows for a larger signal-to-noise ratio of the final spectrum.

2.3 Markov chain Monte Carlo

Markov chain Monte Carlo (MCMC) is a statistical method used to sample from and put constraints on probability distributions whose parameters are unknown and difficult to examine analytically (e.g. [Hogg & Foreman-Mackey, 2018](#); [van Ravenzwaaij et al., 2018](#)).

MCMC is a combination of two statistical tools: The Monte Carlo method and Markov chains. The former is a method used to estimate properties of a distribution by drawing random samples from it. This is particularly useful in cases where the equations of the distribution are either complex or unknown, and if random samples can be easily drawn from it. When sampling from a distribution a number of values are drawn from it, and if that number is sufficiently large, creating a histogram of these values will closely resemble the original distribution up to some normalization constant. By examining this sample one can thereby estimate parameters of the original distribution. The Markov chain aspect of the MCMC introduces the idea that the samples are probabilistically related such that the probability of a newly drawn sample depends on the probability of the previous sample. The newest sample only depends on the previously drawn sample and not the samples drawn before that, and for that reason the Markov chains are said to have no memory. What makes Markov chains useful is the fact that after a sufficient number of steps/iterations they converge towards a stationary distribution independently of the starting points in the parameter space.

MCMC is a useful tool to estimate parameters of a model to describe some observed data. One important thing to note about MCMC is that the general purpose is not to optimize the parameters that describe the data, but rather to sample from the probability distribution around the optimum values. As such, MCMC is useful for estimating the uncertainties of the parameters.

The main purpose of MCMC is to numerically solve multi-dimensional integrals which are difficult to examine analytically. This is often useful in for example Bayesian statistics when estimating a posterior distribution, i.e. the distribution of probabilities assigned to random events taking observed data into account. Ac-

According to Bayes' theorem, the posterior distribution is given by:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (6)$$

Where $P(A|B)$ is the posterior distribution and hence the probability of observing the theoretical/model data A given the observed data B . $P(B|A)$ is the prior distribution, which contains information about what is already known about the system. This could for instance be constraints on the parameters of the model such that they do not become unphysical. $P(A)$ is the likelihood function, which is the probability of observing the data given that the model is true, and $P(B)$ is the marginal likelihood. The latter is a normalization constant and does not depend on the parameters of the model (A). However, it can be difficult to compute, and for that reason, Bayes' theorem is often approximated to:

$$\text{posterior} \propto \text{prior} \times \text{likelihood} \quad (7)$$

This approximation is particularly useful when dealing with a large number of parameters, since the computational cost otherwise increases drastically with the number of parameters.

MCMC works by initially assigning each chain a starting point in the parameter space. During each step or iteration a small random perturbation is added to the current state through Monte Carlo sampling creating what is called the proposal sample. The proposal sample is governed by the proposal function, which differs according to what MCMC method is being used. For instance, a simple proposal function is to take a Gaussian step in each dimension each with a standard deviation σ . The new state is evaluated, and depending on whether the probability of the new state is larger or smaller than that of the previous one, the new state is either accepted or rejected. In the latter case the first state is retained and the process is repeated.

As already mentioned, when running Markov chains, they will eventually converge towards a stationary distribution, i.e. the posterior distribution. The first steps of the chain before the convergence constitute the so-called burn-in sample, and this sample is usually discarded when determining the posterior distribution.

Many algorithms exist for constructing Markov chains, and these have different ways of adding random perturbations to the current states and also different approaches to accepting and rejecting the proposed sample. For this analysis we use the Python package `emcee` (Foreman-Mackey et al., 2013) to perform the MCMC. This is a so-called affine invariant MCMC ensemble sampler, and it is based on the algorithm described in Goodman & Weare (2010). The fact that the sampler is

affine invariant means that it is invariant to affine transformations of the parameter space, i.e. transformations that preserve points, straight lines and planes. This means that one can transform the parameter space such that the probability distribution becomes more simple to sample from, while the sampler will perform identically to what it did on the untransformed problem. By doing this one can minimize the computational cost. The emcee sampler works by throwing an ensemble of random walkers into the parameter space by randomly assigning each of them a starting point. The walkers start exploring the parameter space by taking random steps to points with different parameters, and after a number of steps all (or the majority of) the chains will converge towards the posterior distribution.

3 Creating synthetic spectra with TARDIS

The aim of this analysis is to put constraints on four parameters used to generate spectra with TARDIS: The requested luminosity (L), the initial density of the explosion (ρ_0), the inner boundary velocity (v_{start}) and the outer boundary velocity (v_{stop}). To do this we run a large number of simulations with ten different values of each of the four parameters. This means that we will run a total of 10 000 simulations that will later be used to construct a likelihood function to run the MCMC.

For this analysis we use TARDIS v3.0.dev3912 (installed on March 11th). To be able to run the simulations necessary for this analysis it requires setting up the configuration files for the simulations (Section 3.1), parallelisation of the simulations in order to minimize the computational cost (Section 3.3) and creating an atomic line list (Section 3.2).

3.1 Setting up the configuration files

For each TARDIS simulation we need to set up a YAML configuration file containing all parameters necessary to run the simulation. This means that when running 10 000 simulations we will need an equal number of configuration files. Most of the parameters inside these files are the same, but as mentioned above four of them will be different. This section contains a description of some of the important parameters chosen for our simulations, and an example of a YAML file is shown in Appendix A.

For this analysis we assume that the ejecta of the transient is in local thermodynamic equilibrium (LTE). This means that the temperature does not change significantly over the mean free path that a photon travels before it interacts with matter, and the transient can therefore be described by a single temperature. This assumption is justified by the fact that the spectrum of the first epoch can be reproduced by a single temperature blackbody profile (Watson et al., 2019). It also means that we can set both the ionization and the excitation parameters inside the YAML file to LTE.

Inside the YAML file we specify the abundances of the kilonova. For simplicity and similarly to what was done by Watson et al. (2019), the abundances are assumed to be uniformly distributed and equal to the solar r-process abundance ratios. These are calculated as the total solar abundances (Lodders et al., 2009, Table 6) with the s-process elements (Bisterzo et al., 2014, Table 1) subtracted. Since there are no measurements available for the s-process abundances of elements lighter than krypton (Kr), the total solar abundances of these elements were assumed to be equal to the r-process abundances.

We include elements up to strontium (Sr), since this is the element that causes the pronounced absorption feature around 6500Å to 9500Å, and for this reason Sr is important for the overall shape of the emergent spectrum. Elements lighter than arsenic (As) are excluded from the simulations too, since they are not expected to be produced through r-process, and the conditions for formation of other elements than r-process elements are not present during a neutron star merger. Light elements formed through nucleosynthesis and s-process within the progenitor stars of the neutron stars are neither expected to be present in the kilonova spectrum, since these are disrupted during the supernova explosion and broken down into neutrons within the neutron stars. For elements heavier than iron it is unclear exactly which ones are formed through r-process and how much of them is formed in this manner. However, measurements of the abundances of metal-poor stars may give an indication of what elements are produced through r-process. These stars were formed at an early time of the universe before the interstellar medium had been enriched with metals from supernovae. Since s-process elements are expected to be formed in AGN stars, these are not expected to be present within metal-poor stars. Therefore their abundances of heavy elements are thought to be dominated by elements formed through r-process. One example of such a star is CS 22892-052 (or the so-called Sneden-star), which is an Extremely Metal-Poor star (i.e. $[\text{Fe}/\text{H}] < -3$) (Snedden et al., 2003). Measurements of its abundances of heavy elements show that it has only a small amount of gallium (Ga) and germanium (Ge), which indicates that the major part of these elements are likely to be produced through other processes than r-process. For this reason we choose not to include these elements and let As be the lightest element included in our model.

Furthermore, elements heavier than Sr are not expected to have a significant impact on the overall shape of the spectrum and have therefore been excluded too. If all elements from As to uranium (U) are included in the TARDIS simulation and all are set to solar r-process abundance ratios, the resulting spectra will show significantly more absorption at short wavelengths compared to the observed X-shooter spectrum (Figure 8). This suggests that the abundance ratios of r-process elements within the kilonova are not equal to that of the Sun. We therefore exclude elements heavier than Sr, which also decreases the runtime of the TARDIS simulations significantly. The abundances used for this analysis have been normalised and are shown in Table 1. Since the actual values of the r-process abundances within a kilonova are yet unknown, the optimal way of running the simulations would be to leave the abundances as free parameters. However, due to the drastically increased runtime that this would require, we leave them at constant values. The possibility of an analysis of the abundance ratios in the kilonova is discussed in Section 6.




Figure 8: The observed X-shooter spectrum plotted together with a TARDIS spectrum resulting from a simulation including all elements from As to U and with the abundances set to solar r-process values.

Table 1: Normalized solar r-process abundances for arsenic to strontium calculated as the total solar abundances (Lodders et al., 2009, Table 6) with the s-process elements (Bisterzo et al., 2014, Table 1) subtracted.

Element	Abundance
Arsenic (As)	0.04194
Selenium (Se)	0.46412
Bromine (Br)	0.07357
Krypton (Kr)	0.33034
Rubidium (Rb)	0.03998
Strontium (Sr)	0.05004

Within the configuration file we furthermore specify the density of the transient, and we choose to adopt a power-law density profile as a function of velocity and time since explosion:

$$\rho(v, t) = \rho_0(t_0/t)^3(v/v_0)^n \quad (8)$$

where t_0 , ρ_0 and v_0 are the reference time, density and velocity of the ejecta respectively. For this analysis t_0 is considered to be the time of the explosion and is therefore set to 0. ρ_0 and v_0 are the density and velocity of the photosphere respectively. n is the density index, which controls the skewness of the lines. As investigated by [Kasen et al. \(2017\)](#) the exact functional form of the density profile and the value of the exponent does not have a big impact on the overall light curve. Thus, there is some freedom in the choice of exponent, and therefore we choose $n = -3$ for this analysis. This choice is consistent with previous work, where the value has been adopted for modelling AT2017gfo ([Tanaka et al., 2017](#); [Watson et al., 2019](#); [Gillanders et al., 2021](#)).

As previously mentioned, four of the parameters inside the configuration files are left as free parameters. These four and their impact on the shape of the spectrum are:

Luminosity, L : The luminosity required for the output spectrum. It determines the overall color of the spectrum (i.e. how blue/red the spectral shape is) and is related to the temperature of the blackbody.

Inner velocity, v_{start} : The velocity of the ejecta at its inner boundary (i.e. at the photosphere). It is important for the central velocity of the spectral lines due to Doppler shift of the radiation emitted from the photosphere. Since the inner velocity defines the size of the photosphere, it also controls to a large extent the brightness of the spectrum. In this analysis we set $v_0 = v_{\text{start}}$ and thus it also has an impact on the density profile (Eq. 8)

Outer velocity, v_{stop} : The velocity of the ejecta at its outer boundary (i.e. the edge of the ejecta), which controls the width of the lines due to Doppler broadening. Along line-of-sight the absorption is more blueshifted, so the faster the ejecta moves the broader the lines become.

Density, ρ_0 : The density at the edge of the photosphere. It drives the the density profile (Eq. 8) of the ejecta and thereby the depth of the absorption lines.

These four parameters are left free by defining a range for each of them within which their best fit values are believed to be. We then run a number of TARDIS simulations, where each parameter takes one of ten possible values within this range, so with four free parameters this gives a total number of simulations of 10^4 .

We aim to find the best fit models to the first four epochs which were measured at times 1.427, 2.42, 3.41 and 4.40 days after the explosion (Smartt et al., 2017). We first focus on modelling the spectrum of the first epoch and will come back to the remaining three epochs later in the analysis.

In order to find the ranges of the four parameters that we want to scan over, we start by hand fitting the first epoch spectrum to find the best fit by eye. The parameters that provide the best fit are shown in Table 2, and the spectrum created with these parameters is plotted in Figure 9. As the next step we change the values of the four parameters individually while keeping the remaining three at the best fit values. By plotting this together with the best fit and the residuals, we can determine the ranges of the four parameters that provide fits close to the best fit. Watson et al. (2019, Extended Data Figure 4) provides two estimates of the ejecta expansion velocity. One estimate is based on the width of the P-Cygni absorption lines, while the other is calculated from the blackbody radius, since it depends on the square-root of the luminosity, and can thereby be converted into an expansion velocity using the time of observation. These two independent estimates give comparable results, and they are used to give a first indication of the value of the v_{start} parameter.

Table 2: Best fit parameters evaluated by eye for the first epoch spectrum together with the ranges chosen to be scanned over.

Parameter	Best fit by eye	Minimum value	Maximum value
L [erg/s]	$3.55 \cdot 10^{41}$	$3.2 \cdot 10^{41}$	$3.9 \cdot 10^{41}$
v_{start} [m/s]	$0.87 \cdot 10^8$	$0.8 \cdot 10^8$	$0.93 \cdot 10^8$
v_{stop} [m/s]	$1.01 \cdot 10^8$	$0.96 \cdot 10^8$	$1.06 \cdot 10^8$
ρ_0 [g/cm ³]	$1.25 \cdot 10^{-16}$	$0.9 \cdot 10^{-16}$	$2 \cdot 10^{-16}$
D_L [Mpc]	38.6	35	41

Inside the configuration file we specify the number of bins for the resulting spectrum, which is what governs the resolution. Furthermore, we set the number of r-packets and v-packets for the last iteration, where the spectrum is calculated, and these numbers together set the signal to noise ratio (SNR) per wavelength unit (SNR/Å) of the final TARDIS spectrum.

The number of bins is chosen such that the resolution of the spectrum is comparable to the resolution of the X-shooter spectrum. To find the SNR of the TARDIS spectrum we first take the median filter of the simulated data using the command `scipy.signal.medfilt` and then calculate the root mean square deviation (RMS) between the median filtered data and the simulation. The SNR/Å is given by $1/\text{RMS}/\Delta\lambda$, where $\Delta\lambda$ is the resolution. To find the SNR of the X-




Figure 9: The best fit to the first epoch X-shooter spectrum obtained by eye.

shooter data, we take the flux spectrum and divide it by the error spectrum, and the $\text{SNR}/\text{\AA}$ is then given by $\text{SNR}/\Delta\lambda$.

To determine the number of r-packets used for our TARDIS simulations, we increase the number of packets until the $\text{SNR}/\text{\AA}$ is comparable to that of the X-shooter spectrum. The $\text{SNR}/\text{\AA}$ for both spectra are shown in Figure 10. As seen in the figure the $\text{SNR}/\text{\AA}$ of the TARDIS spectrum is greater than that of the X-shooter spectrum at most wavelengths. The $\text{SNR}/\text{\AA}$ should preferably be 3-5 times larger than that of the X-shooter spectrum in order to get a better comparison between model and data. However, this would require an larger number of packets and as a result the simulations would be computationally expensive. This is due to the fact that the noise of the TARDIS spectrum is governed by Monte Carlo noise, meaning that it scales as $1/\sqrt{N}$, where N is the number of r-packets. Thus, in order to decrease the noise by a factor of four and bring us to the optimal $\text{SNR}/\text{\AA}$, we would need 16 times more packets. This would likewise increase the runtime drastically by a factor of 16 and thereby take approximately a month to run all the simulations (see Section 3.3 for a discussion on runtime). Therefore we settle for a $\text{SNR}/\text{\AA}$ similar to that of the observed spectrum, which requires 10^8 r-packets and 5 v-packets for the last iteration.

As mentioned in Section 2.2, the plasma state is constant within each shell such that quantities like abundances and the density of the ejecta are the same within each shell. We want to adopt the power-law density profile as shown in Eq. 8,




Figure 10: Signal to noise ratios (SNR) for the X-shooter spectrum and the TARDIS simulations.

which means that we need the density to depend on the radius of the transient. In order to obtain such a density gradient, we need to run the simulations with several shells. We choose to use 10 shells for the TARDIS simulations. This number was determined by running the same simulation with the same parameters for different numbers of shells. The resulting spectra started to converge at 10 shells meaning that the spectra did no longer depend on the number of shells when running with more than 10. This convergence is shown in Figure 11 together with the runtimes of each simulation. As seen in the figure, the runtime of a simulation increases with the number of shells, and we therefore want to use as few shells as possible without having the dependency of the exact number.

3.2 Carsus and VALD

Previous works on spectral modelling of AT2017gfo have highlighted the importance a detailed line list when investigating the presence of heavy elements in the ejecta of kilonovae (e.g. [Smartt et al., 2017](#); [Gillanders et al., 2021](#); [Watson et al., 2019](#)). Both [Smartt et al. \(2017\)](#) and [Watson et al. \(2019\)](#) used the Kurucz line list ([Kurucz, 2017](#)) as the atomic database for their models. However, this line list is incomplete since it only contains lines for the lowest ionisation states of elements up to the first r-process peak (i.e. $A \lesssim 80$). This is primarily due to difficulties on constraining the weak NIR lines of the heaviest elements, and for this reason the majority of these lines are still unknown. The lack of complete atomic data makes




Figure 11: The resulting spectra from seven TARDIS simulations run with the same parameters, but with different numbers of shells. The left hand side shows the full spectrum, and the right hand side shows the Sr absorption feature, where it is clear to see how the spectrum depends on the number of shells. The spectra are plotted together with their runtimes.

any detailed modelling and conclusions on the presence of r-process elements in kilonovae difficult. For this reason Gillanders et al. (2021) used a more detailed atomic data set based on theoretical atomic data of gold (Au) and platinum (Pt). Despite not being able to prove the presence of these two elements, they highlight the importance of more detailed line lists to study the abundances of heavy elements in kilonovae.

The publicly available atomic data table for TARDIS is based primarily on the Kurucz line list. However, motivated by the importance of detailed atomic data when modelling kilonovae, we choose to run our simulations using the Vienna Atomic Line Database⁶ (VALD3) (Piskunov et al., 1995; Ryabchikova et al., 2015). This database provides a line list that is more complete and contains more lines than the Kurucz line list. We use all atomic lines from 0.1nm to 0.1mm excluding non-local thermodynamic equilibrium (NLTE) lines, hyperfine structure (HFS) lines and lines from elements with ionization states above three.

Despite running the simulations with wavelength ranges similar to that of the X-shooter spectra (i.e. 3250-22260 Å), we choose to include lines outside this range, since we do not know how important the transitions at these wavelengths are to the overall opacity. Furthermore, in the nebular phase, some fraction of the emission may emerge in long wavelength transitions. We do not include NLTE lines since the earliest epoch spectra can be well reproduced by a single temperature

⁶<https://www.astro.uu.se/valdwiki>

blackbody model suggesting that the transient is close to local thermal equilibrium (LTE) (Watson et al., 2019). For this reason we run our simulations assuming that the plasma is in LTE (see Section 3.1), so we do not need to implement NLTE lines in our line list. HFS lines are neither included since the synthetic spectra will not have a resolution that will allow us to pick up the HFS lines, and even if we did it is likely that broadening of the fine structure lines would cause them to overlap the HFS lines. We do not consider molecules for our line list, since we assume that the conditions within the explosion do not provide an environment for formation of these. One reason is that the timescale of the explosion is short, which means that the molecules would have to form rapidly. No solid surfaces are neither expected to be present during the merger, and these are often necessary for the chemistry to take place. Furthermore, only heavy elements are synthesized during the merger which means that the most common molecules in the universe containing elements like carbon (C), nitrogen (N) and oxygen (O) will not be present in the transient. Such molecules without these components are neither represented in the VALD line list. Lastly, we choose not to include lines from elements at ionization states above three, since heavy elements are only expected to be either neutral, singly, doubly or triply ionized in a kilonova more than one day after the explosion and when $T \lesssim 20000$ (Tanaka et al., 2017). With all these lines excluded from the line list, the resulting list has a total of ~ 156 million lines compared to Kurucz line list which has ~ 0.55 million atomic lines.

To run the simulations we need to set up the atomic line list and convert it into a format readable by TARDIS. To do so we use CARSUS⁷ (Pássaro et al., 2019), which is a Python package used to manage atomic datasets and output a file readable by radiative transfer codes like TARDIS.

CARSUS uses the Kurucz line list as its main source for energy levels and spectral lines, and these are typically supplemented with lines and levels from The Chianti Atomic Database and CMFGEN Atomic Data. Since we are interested in using VALD, we convert this line list into the same format as the Kurucz line list (i.e. with the same columns and units) such that we can use the Kurucz line list reader to parse the VALD data. We then use CARSUS to convert the VALD line list (excluding data from Chianti and CMFGEN) into a hierarchical data format (HDF5) file suitable for TARDIS. This type of file contains a number of data tables that include data on mass, ionization, levels and lines. Apart from extracting the data on lines and levels from VALD, CARSUS takes the atomic weights and ionisation energies from the National Institute of Standards and Technology (NIST)⁸, and Long & Knigge (2002) provides the ground state recombination fractions (ζ).

⁷<https://tardis-sn.github.io/carsus/>

⁸<https://www.nist.gov/pml/atomic-spectra-database>

For our line list we choose to include all elements from arsenic (As) to uranium (U). As such, we include elements present in all three r-process peaks. By default CARSUS has a threshold for the $\log gf$ value of the lines (i.e. the logarithm of the product of the statistical weight of the lower level and the oscillator strength for the transition) such that all weak lines where $\log gf < -3$ are excluded. However, a large number of weak lines from lanthanide elements may contribute substantially to the spectral shape of the kilonova. For this reason we change the threshold of the $\log gf$ value to -10 to ensure that no important lines are excluded. Including the weaker lines increases the size of the final line list and causes the runtime of the simulations to increase. This does not have a significant influence for our analysis, since we include only a few number of elements (see Section 3.1). If we were to run the simulations with a larger number of elements we would have to reconsider this value, since a threshold of -3 may be enough.

After running CARSUS the resulting VALD line list has $\sim 300\,000$ lines compared to Kurucz line list which has $\sim 23\,000$ lines from As to U. The distributions of these lines are shown in Figure 12, where the number of lines per element for each of the two line lists is plotted.




Figure 12: The number of lines for each element from As-U in the Kurucz and VALD line lists respectively.

3.3 Parallelisation

After setting up the configuration files and the atomic line list we are ready to run the 10 000 TARDIS simulations. With the setup described above every simulation takes approximately 25-30 minutes, so with 10 000 simulations it requires a total runtime of up to 200 days. To be able to run our simulations in a feasible amount of time, we will therefore benefit from parallelisation.

In order to parallelise our simulations we run them on the Electronic Research Data Archive (ERDA) at University of Copenhagen. This is an external server meant for storing, sharing, analyzing and archiving research data, and it is particularly useful when working with a large amount of data. Within ERDA we use the MPI Oriented Development and Investigation (MODI) notebook, which is a notebook configured as a SLURM (Simple Linux Utility for Resource Management) cluster⁹. This type of cluster is a workload manager that allows for a user to access resources (compute nodes), and it provides a framework to allocate and execute jobs on these nodes. This is particularly useful for parallelisation, since it gives the opportunity for batch jobs submitted to the cluster to be scaled to a parallel execution on several nodes. The cluster on ERDA has a total of eight compute nodes, and each of these has 64 CPUs and 256 Gb RAM.

When running a job in parallel on ERDA, the files that are necessary to execute the job should all be located within the `modi_mount` directory. This directory is a shared Network File System (NFS) mounted on every compute node, and thus all files should be placed here in order for the compute nodes to load them. To be able to run a TARDIS simulation on the cluster we also need to install the TARDIS package on every compute node. This is only possible by running the installation through the SLURM cluster, meaning that we have to submit the installation to the cluster as part of a SLURM job using the `sbatch` command. In order to keep it simple, the job is split into two scripts: A bash script and a SLURM script. As such, we create a bash script within the `modi_mount` folder, in which we run the installation of TARDIS and execute the application. The bash script is then run inside a SLURM script using the `srun` command. The contents of the bash script and the SLURM script are shown in Figure 13. This is an example of a job submitted to one node (in this case the `modi000` node) first installing TARDIS and afterwards running a Python script that executes a TARDIS simulation. As seen in the bash script we activate conda before the installation, and the directory to where packages are located is changed to a folder within the `modi_mount` directory. This is done in order to use a directory where we have permission to write new packages and to ensure that the nodes can access the TARDIS packages

⁹<https://slurm.schedmd.com/overview.html>

once the job is being executed. As shown in Figure 13, the bash script is run within the singularity image *hpc-notebook_latest.sif*. Utilizing a singularity image when executing a SLURM job allows us to use special libraries that are not pre-installed on the compute nodes directly. Thus, we run our job inside an environment that has the required dependencies.

```

pj1124_alumni_ku_dk@38dd534409c0:~/modi_mount/tardis/grid$ cat run_tardis.sh
#!/bin/bash

source $CONDA_DIR/etc/profile.d/conda.sh
export CONDA_PKGS_DIRS=~/.modi_mount/conda_dir

wget https://raw.githubusercontent.com/tardis-sn/tardis/master/tardis_env3.yml
conda env create -f tardis_env3.yml --force
conda activate tardis

git clone https://github.com/tardis-sn/tardis.git
python setup.py install --user

cd scripts
python tardis_script.py

```

(a) Bash script

```

pj1124_alumni_ku_dk@38dd534409c0:~/modi_mount/tardis/grid$ cat slurm_job.sh
#!/bin/bash

#SBATCH --job-name=tardis
#SBATCH --partition=modi_short
#SBATCH --nodes=1
#SBATCH --odelist=modi000

srun singularity exec ~/modi_images/hpc-notebook_latest.sif ~/modi_mount/tardis/grid/run_tardis.sh

```

(b) SLURM script

Figure 13: Example of the SLURM and bash scripts used to run the installation of TARDIS and submit a job to the cluster.

The SLURM cluster has four different partitions that can be utilized when running a job, and within the SLURM script we specify what partition should be used for the job. Each partition has a different priority and upper limit to the runtime of the job. The job is automatically terminated in case this limit is exceeded or if a higher priority job is submitted to the queue. As seen in Figure 13 we run our jobs utilizing the `modi_short` partition. This partition has the second highest priority and allows for jobs to run for up to two days. Finally, the SLURM script is submitted to the queue using the `sbatch` command, which is a standard SLURM command used to submit jobs.

By using the SLURM cluster to run our jobs, we can benefit from a total of 512 CPU cores and approximately 2 Tb RAM. In our case, the number of jobs that can be run simultaneously is limited by the memory usage of each simulation. Due to the large number of r-packets needed to obtain the required SNR and resolution, each simulation uses ~ 12 Gb of memory. This means that we can

run approximately 150 simulations simultaneously if we use all nodes. Thus, by parallelising our job, we can decrease the total runtime for all 10 000 simulations from ~ 200 days to less than two days.

4 Markov chain Monte Carlo

From the 10 000 TARDIS simulations run with different sets of parameters we aim to find the model with the set of parameters that best describes the observed X-shooter spectrum for the first epoch. To constrain these four parameters we use MCMC to sample from the posterior distribution of likelihoods in order to find the set of parameters with the maximum likelihood. To do this we need to define a likelihood function and a prior function according to Eq. 6.

The likelihood function is calculated from the output spectra from the TARDIS simulations. These are arranged in a four dimensional grid defined by the four free parameters L , ρ_0 , v_{start} and v_{stop} , and the likelihood function is found by first determining the likelihood of each of the spectra matching the observed X-shooter spectrum. However, the TARDIS spectra are given by the luminosity of the transient, whereas the X-shooter spectrum is given by the flux. This means that to be able to compare the simulated spectra to the observed data, the luminosity must be converted into a flux using the following equation:

$$F = \frac{L}{(1+z) \cdot 4\pi D_L^2} \quad (9)$$

where $(1+z)$ is a term that corrects for the Å bandpass, meaning that it corrects for the fact that the observed Å bandpass is stretched out compared to the rest frame. This is necessary when converting a luminosity in units of [erg/s/Å] to a flux in units of [erg/cm²/s/Å]. The redshift is the cosmological redshift, and for this source $z = 0.0099$ (Abbott et al., 2019b). The distance to the source D_L is not well known, and thus constraining this parameter is of great interest. For that reason we leave the distance as an additional free parameter ranging from 35 Mpc to 41 Mpc. This range was determined in the same way as for the remaining four parameters by varying the distance and only include the distances that gave reasonable fits. With this additional parameter the resulting grid is of five dimensions.

We want to compare the simulated spectra to observed data, and to do this we use data measured by the X-shooter spectrograph. This data set consists of ten spectra measured with one day intervals (Figure 4). It was originally published by Pian et al. (2017) and Smartt et al. (2017) and is publicly available through the ENGRAVE webpage¹⁰. As previously mentioned, we first concentrate on the first epoch spectrum and discuss the analysis of the three next epochs later. All

¹⁰<http://www.engage-eso.org/AT2017gfo-Data-Release/>

X-shooter spectra have been flux-calibrated, de-reddened, redshift corrected and smoothed.

Using each TARDIS spectrum as our proposed model for the data, we calculate the likelihood of measuring the X-shooter spectrum given the model. The likelihood is the product of the probabilities of measuring each data point given the TARDIS spectrum (Barlow, 1989):

$$\mathcal{L}(y_1, \dots, y_N | a) = \prod P(y_i | a) \quad (10)$$

where a is the set of free parameters (i.e. D_L , L , ρ_0 , v_{start} and v_{stop}).

The X-shooter spectrum consists of many independently measured data points meaning that according to the central limit theorem, the distribution of data points around their ideals are Gaussian. Thus, the probability of measuring y_i given x_i is:

$$P(y_i, a) = \frac{1}{\sigma_i \sqrt{2\pi}} e^{-(y_i - f(x_i, a))^2 / 2\sigma_i^2} \quad (11)$$

where y_i and x_i in this case are the observed flux and wavelength respectively, and $f(x_i, a)$ is the flux of the TARDIS spectrum at wavelength x_i . σ_i is the combined uncertainty of the uncertainty on the observed flux and the uncertainty on the TARDIS spectrum, i.e.:

$$\sigma_i = \sqrt{\sigma_{i,obs}^2 + \sigma_{i,sim}^2} \quad (12)$$

Where $\sigma_{i,obs}$ is given by the error spectrum and $\sigma_{i,sim}$ is calculated as the RMS of the TARDIS spectrum relative to the median filtered data. Given Eq. 10 and 11, the natural logarithm of the Gaussian likelihood is:

$$\ln(\mathcal{L}) = \sum \ln(P(y_i, a)) = -\frac{1}{2} \sum \left[\frac{(y_i - f(x_i, a))^2}{\sigma_i^2} + \ln(2\pi\sigma_i^2) \right] \quad (13)$$

When calculating the likelihood of each spectrum we choose to discard data points 4000 Å and above 13000 Å. The reason for discarding the NIR wavelengths is that as previously discussed, the weak lines from the heavy elements are still unknown and therefore not a part of our model. It is likely that the NIR emission of the kilonova is dominated by a large number of these weak lines, and for this reason we do not expect to be able to fit this part of the spectrum. The reason for excluding data at short wavelength is that, as shown in Figure 8, this part of the spectrum is likely to be dominated by absorption from heavy elements, and the majority of these have been excluded from the model.

From our grid of likelihoods we can obtain a likelihood function by interpolating the grid points. To do this we use `scipy.ndimage.map_coordinates`, which allows us to interpolate in five dimensions using a cubic spline. The likelihood function can then be used to run the MCMC and find the set of parameters that provide the spectrum with the highest likelihood. The priors are flat priors constraining the values allowed for each of the five parameters. These are chosen such that the only values the parameters can take are the ones within the ranges that were scanned over and thus the values for which the likelihood function is defined.

As mentioned in Section 2.3, we use the python package `emcee` (Foreman-Mackey et al., 2013) to run the MCMC. According to Foreman-Mackey et al. (2013) one should use a number of walkers of the order of 100, so we run the MCMC with 100 walkers, since this seems to reproduce the posterior distribution well. The number of iterations that the chains should run is determined from the fact that the chains should be approximately 50 times the auto-correlation time. This auto-correlation time is a measure of the number of evaluations of the posterior distribution necessary to obtain independent samples. By making a test run with a random number of iterations, the integrated auto-correlation times for the five parameters can be obtained, and for the first epoch spectrum these are:

$$D_l : 17.9 \quad L : 21.2 \quad \rho_0 : 26.9 \quad v_{start} : 33.3 \quad v_{stop} : 52.1$$

Considering these numbers, we choose to run the chains with 2700 steps, since this is approximately 50 times the longest auto-correlation time. 200 of these steps are discarded as the burn-in sample. The convergence of the walkers and the burn-in sample are shown in Figure 14.

The walkers are initialized such that they are placed in a tight 5-dimensional bundle close to the point that is expected to give the highest likelihood. This point is assumed to be at the set of parameters that give the best fit by eye (Table 2). As discussed by Foreman-Mackey et al. (2013), the most objective way to run the MCMC would be to initialize the walkers such that they are evenly distributed throughout the parameter space. However, this very rarely works in practice and causes the walkers not to converge towards one final likelihood but rather get stuck in local maxima. This is an effect that we see in our analysis too, and the walkers converge to a number of local maxima. Therefore we start the walkers close to the best fit by eye, which instead causes them to converge to the same maximum. Despite the seemingly subjective initialization, the walkers will explore the rest of the parameter space, and some of them should converge towards a point with a higher likelihood if one exists.




Figure 14: Convergence of the 100 walkers. The blue dotted line marks the end of the burn-in phase, which is after 200 iterations.

5 Results and discussion

The results for the five parameters for the first four epochs are listed in Table 3. The parameters for the last three epochs were obtained in a similar way to what was done for the first epoch, but the ratio between the two velocities was kept constant, and for this reason only v_{start} was left as a free parameter. The reason behind this approach will be discussed below.

Table 3: Results for the best fit parameters for the first four epochs. For the last three epochs v_{start} is 87.64% of v_{stop} , and therefore only one velocity was kept free. The uncertainties are the 2σ deviation from the mean.

Epoch No.	D_L [Mpc]	L [erg/s] $\times 10^{41}$	ρ_0 [g/cm ³] $\times 10^{-16}$	v_{start} [m/s] $\times 10^8$	v_{stop} [m/s] $\times 10^8$
1	38.58 ± 0.02	3.541 ± 0.004	1.501 ± 0.008	0.8696 ± 0.0010	0.9922 ± 0.0005
2	38.35 ± 0.08	2.426 ± 0.005	0.1036 ± 0.0016	0.7570 ± 0.0009	-
3	41.47 ± 0.10	2.096 ± 0.009	0.0952 ± 0.0012	0.6090 ± 0.0010	-
4	48.65 ± 0.09	2.362 ± 0.008	0.1280 ± 0.0023	0.6188 ± 0.0014	-

The spectra created from best fit parameters are plotted in Figure 15 together with the observed X-shooter spectra. As seen in the figure we have obtained a good fit for the first epoch spectrum, while the last three epochs suffer from remarkable deviations between model and data. More specifically we were unable to fit the emission line around $1.1 \mu\text{m}$, which is clear to see in the residual plots, and for this reason the peak was excluded from the analysis. The origin of the peak and why we were unable to fit it is discussed below. As such, the grey zones in Figure 15 mark the parts of the data that were excluded when calculating the likelihood. Short wavelengths were excluded from the first epoch spectra since we expect the spectrum to be dominated by absorption from r-process elements at these wavelengths, and these were not included in the model. Long wavelengths were excluded due to the fact that we expect a large number of weak r-process lines to emit at these wavelengths, and thus we did not expect our model to match the data in the NIR.

Figure 16 shows the corner plot from the MCMC for the first epoch spectrum. This plot allows us to visualize the multidimensional samples obtained from the MCMC in a scatterplot matrix. As such, each scatterplot contains a number of data points corresponding to the number of samples after excluding the burn-in sample. These samples have been plotted together with contours at 0.5σ , 1σ , 1.5σ , and 2σ . By creating a corner plot we obtain a scatterplot for every combination of the five parameters, which reveals any correlation between some of them. For instance we find from Figure 16 that there is a strong correlation between luminosity and distance to the object. This is as expected, since the more luminous the

object is the further away it would have to be in order to match the observed flux. Furthermore, the corner plot shows the final distributions of the parameters in five histograms. Considering the obtained distributions we find that all five of them look reasonable, which confirms that the walkers have converged to the posterior distribution. The corner plots for the last three epochs can be found in Appendix B.

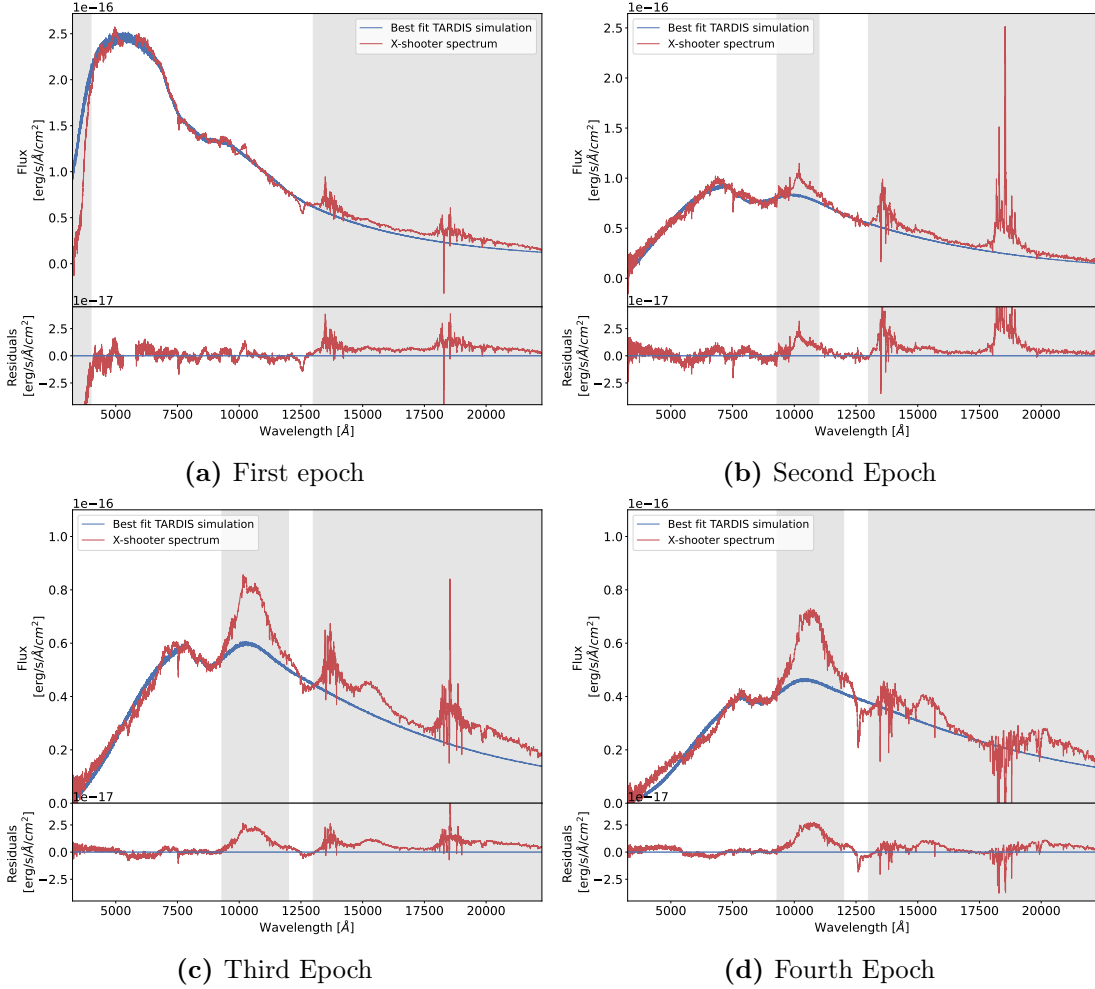


Figure 15: The best fit TARDIS spectra plotted together with the observed X-shooter spectra for the first four epochs. The grey zones were excluded from fitting for reasons described in the text.

One important thing to note about the observed X-shooter spectra is the two discrepancies in the measured flux where the UVB/VIS (5600 Å) and the VIS/NIR (10240 Å) arms meet. This is possibly due to calibration issues causing a mismatch between the fluxes measured by each arm. As a result all TARDIS spectra seem to

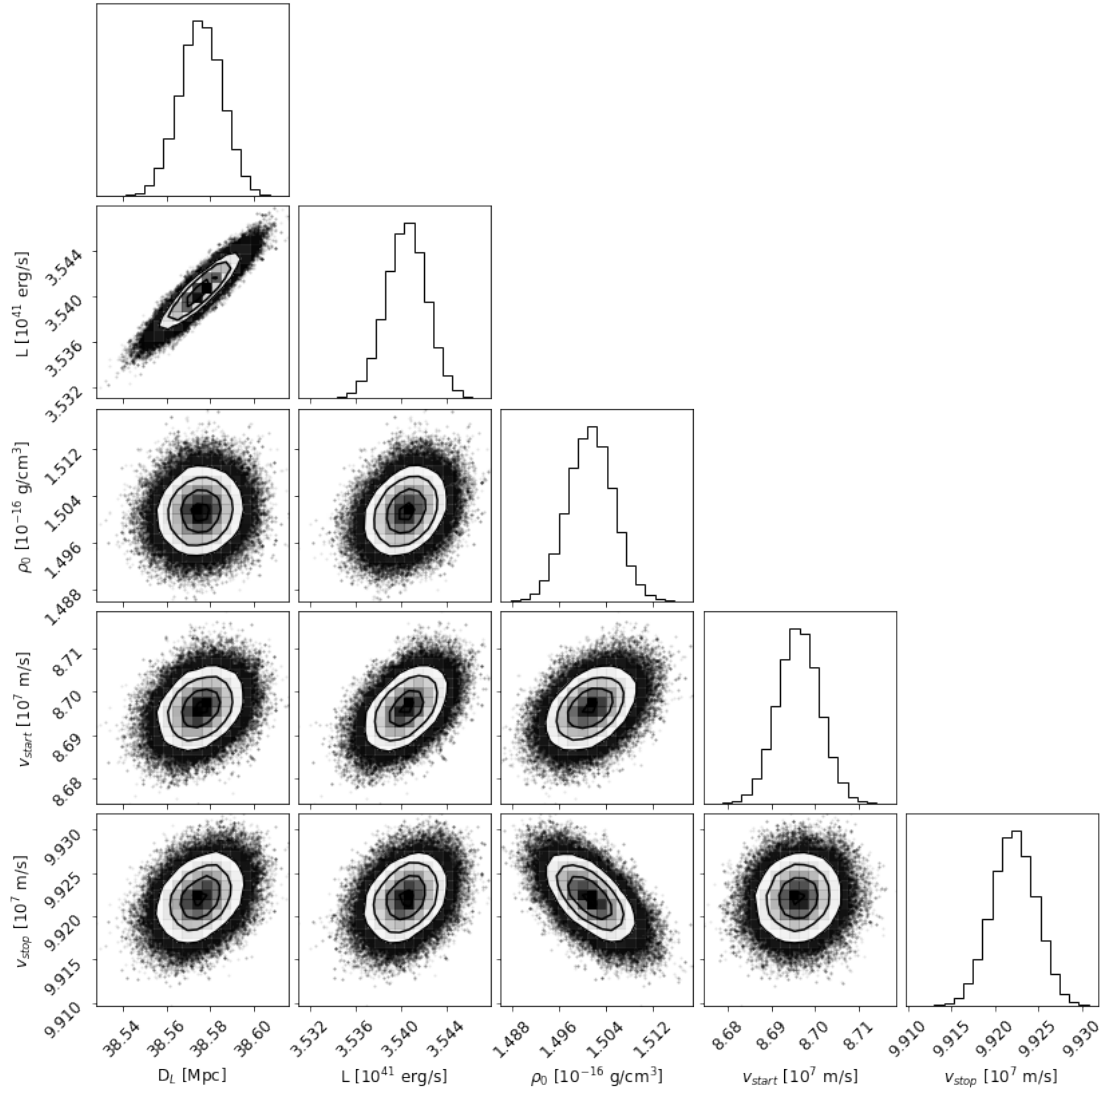


Figure 16: Corner plot for the first epoch spectrum.

fit poorly around these two miscalibrations.

As seen in Table 3 all parameters have remarkably small uncertainties. One way to visualize these is to plot a sample of spectra where all five parameters are within the 90% fractiles of the mean. To do so we take our 270 000 samples (100 walkers \times 2700 iterations) and remove the burn-in sample and the samples where at least one parameter has a value outside the 90% fractiles. This leaves us with approximately 200 000 samples, and by running TARDIS with every 300 of these sets of parameters we create 666 spectra. These are all plotted in Figure 17. Since the spectra have parameters close to each other they are expected to be very similar. However, as seen in the figure, seven of the spectra deviate remarkably from the majority, which is likely an artefact of TARDIS occasionally calculating spectra that are not normalized correctly. For this reason we calculate χ^2 for all the spectra relative to the observed X-shooter spectrum. Assuming that TARDIS calculates all spectra correctly, the probability of creating a spectrum with a χ^2 value 5σ away from the mean value is 6 in 10^7 . This means that it is very unlikely to find such a spectrum when running 666 simulations. Based on this, we discard any spectrum with a χ^2 value more than 5σ from the mean. This process of calculating the mean χ^2 and discarding the spectra more than 5σ away is continued until there are no more discards.




Figure 17: 666 TARDIS spectra with parameters within the 90% fractiles of the mean values. Seven spectra clearly deviate from the majority against expectation and thus indicate a bug within TARDIS.

The remaining spectra are all plotted in Figure 18 together with the observed spectrum. The discards of the spectra 5σ away from the mean caused the exclusion of the seven spectra as expected. As seen in the figure, the spectra cover a relatively small flux range confirming the initial suspicion that the uncertainties are very small. Reasonable uncertainties would have caused the flux range to cover the entire X-shooter spectrum or at least the majority of it. However, especially at long wavelengths we find that the flux range is so small that the observed spectrum exceeds the range at most wavelengths. The main reason for the small uncertainties is the fact that these are only statistical uncertainties, and the systematic ones have likely been underestimated. These systematic uncertainties for instance include the assumptions made by TARDIS on the geometry of the explosion, abundances etc.




Figure 18: The X-shooter spectrum plotted together with a sample of TARDIS spectra with all five parameters within the 90% fractiles of the mean values. The parts of the spectrum that were not included when fitting have likewise been excluded from the plot.

As previously mentioned the last three epochs were modelled differently from the first epoch. To model these epochs we assumed that the velocities at the inner

and outer boundaries are correlated such that the ratio between the two velocities is constant at all times. For the first epoch we find a ratio of $\frac{v_{inner}}{v_{outer}} \approx 0.8764$. Since $\frac{v_{inner}}{v_{outer}} = \frac{t_{exp} r_{inner}}{t_{exp} r_{outer}} = \frac{r_{inner}}{r_{outer}}$ the approximation of keeping this ratio constant means that the size of the photosphere compared to the total size of the transient is constant. Thus we leave v_{start} as a free parameter and make v_{stop} dependent of this value. By using this approximation, we can reduce the number of free parameters for the simulations by one and thereby decrease the runtime from ~ 2 days to 3-4 hours.

A different approach to decreasing the number of parameters would have been to set the inner boundary velocity to something small and constant and leave the outer boundary velocity as a free parameter. While the exact value of the outer boundary velocity is important for the width of the spectral lines, the value of the inner boundary velocity is mostly important to make sure that the atmosphere that the packets propagate through is thick enough for interactions to take place and thus to ensure that we do not look directly into the blackbody. However, we do not expect a much different result from using this approach, since the major issue with fitting the later epoch spectra does not seem to be the constant velocity ratio but rather the fact that TARDIS does not take nebular emission into account when running the simulations, and this is likely an important component in the later epoch spectra. Keeping the velocity ratio constant still gives us an indication of how the parameters evolve over time, and especially the strontium abundance, which will be discussed in Section 5.2. The constant velocity ratio is also a way of ensuring the thickness of the atmosphere.

As seen in Figure 15 the later epoch spectra all show a pronounced emission line centered around $1.1\mu\text{m}$, which TARDIS was unable to fit. It has been suggested by Watson et al. (2019) that this is the emission component of a SrII P-Cygni line. A P-Cygni line is a spectral line comprised by a redshifted emission line with an absorption line to its blue-ward side, and they are present in spectra from objects that have strong outflows of material. The absorption feature comes from radiation travelling directly towards the observer along the line-of-sight. This radiation passes through an expanding ejecta moving away from the object, which causes absorption features to be blueshifted. The gas surrounding the object radiates such that emission coming from parts of the gas moving away from the observer will be redshifted, while emission coming from parts of the gas moving towards the observer will be blueshifted. However, the latter emission is not visible in the spectrum due to the absorption, which means that the observed emission is all redshifted. A P-Cygni profile becomes more pronounced as the object expands, which is why the emission line in the spectra of AT2017gfo is very pronounced in the later epoch spectra. It does however become less prominent as the kilonova cools down and becomes less luminous. Since this feature is an effect of the three

dimensional structure of the object, its absence in the TARDIS spectra may be an artefact of modelling the kilonova in one dimension. Furthermore, TARDIS does not take nebular emission into account when running simulations, which may be another reason that we do not see the emission.

The distances to the kilonova determined from the first two epochs are relatively short compared to the distances that have previously been determined to the object or its host galaxy. For comparison, the distance found in this work for the first epoch spectrum has been plotted together with the distances found in the literature in Figure 19. We choose to compare the distances from the literature to the one inferred for the first epoch spectrum, since we consider this to be our most reliable distance estimate based on the quality of the four fits as shown in Figure 15. A description of how the distances from the literature were calculated can be found in Section 1.4. As seen in the figure, the distance estimate from this work is generally small compared to the rest. However, even though our measurement has a small uncertainty, it is still within the uncertainties of the rest.




Figure 19: The previously estimated distances to the kilonova or its host galaxy plotted together with the distance found in this work for the first epoch spectrum (orange dashed line). The methods used to find the distances are described in Section 1.4. The statistical and systematic uncertainties on the [Cantiello et al. \(2018\)](#) have been added in quadratures to find a combined uncertainty.

The short distance may be an artefact of running a one dimensional simulation and thereby assuming that the explosion resulting from a neutron star merger is spherically symmetric. The shorter distance suggests that the transient is oblate meaning that it is more extended radially compared to along line-of-sight. Such an explosion would appear more luminous to the observer than in the case of a spherically symmetric explosion, so to match the spectra of such two events one would need a shorter distance to the spherical explosion compared to the oblate one. However, if the explosion was spherical as assumed by TARDIS, one would have expected the observed explosion to be prolate and thus elongated along line-of-sight. This is due to the effect of light delay which is not taken into account when running TARDIS. Light delay means that the light emitted from the edges of the transient reaches the observer later than the light emitted from the center. The radiation emitted from the edges needs to travel an additional distance comparable to the radius of the object at the time of emission. For the first epoch, the velocity of the outer edge (v_{stop}) is approximately $0.3c$. Given the time of the explosion of ~ 1.5 day, the radius of the object is $v_{stop}t_{exp} \sim 0.5$ light days. This means that the light emitted from the edges of the object should arrive at the observer ~ 0.5 days after the light emitted from the center, which would make the transient appear prolate. Due to the high velocity of the ejecta, this effect is expected to be important in cases of kilonovae.

Investigating the shape of the spectral lines could potentially reveal the geometry of the explosion. If the explosion is prolate, absorption along line-of-sight will dominate the line causing the line to be dominated by blueshifted absorption and with a wing towards redder wavelengths. On the other hand, if the explosion is oblate it will have more absorption at red wavelengths. The shapes of the spectral lines could potentially be investigated by fitting the lines with a variety of spectral line functions (e.g. P-Cygni profiles) and in this way determine the skewness of the lines. However, this analysis is relatively complex, and we will therefore leave it for future work.

From the density found for the first epoch we can calculate the approximate mass of the ejecta inside the photosphere assuming that the radius is $R = v_{start}t_{exp}$ and the density is uniform:

$$M_{ej} \approx \frac{4}{3}\pi R^3 \rho_0 \approx 4 \cdot 10^{-4} M_{\odot} \quad (14)$$

This is approximately a factor of 10-100 smaller than the ejecta mass found in previous works of $\sim 10^{-4} - 10^{-2} M_{\odot}$ (e.g. [Hotokezaka et al., 2013](#); [Bauswein et al., 2013b](#)). ρ_0 is the density at the edge of the photosphere, and one would expect the density to increase towards the center. This means that the above estimate is a lower boundary to the ejecta mass, and since it is lower than previously inferred

ejecta masses, it implies that the density found for the kilonova is reasonable.

From the inferred densities in Table 3 we can calculate a density profile that depends on the radius at an instant in time. To do this we use the following relation between density and time:

$$\rho \propto \frac{m}{V} \propto \frac{m}{R^3} \propto \frac{m}{t^3} \quad (15)$$

where m is the total mass, V is the volume of the explosion, R is the radius and ρ and t are the density and time respectively. The density measured at time t_1 can then be converted into a density at time t_2 using the relation:

$$\frac{\rho_1}{\rho_2} \propto \left(\frac{t_1}{t_2}\right)^{-3} \Rightarrow \rho_2 \propto \rho_1 \left(\frac{t_2}{t_1}\right)^{-3} \quad (16)$$

By inserting the densities in Table 3 and their corresponding times since explosion we can infer the densities 1.427 days after the explosion. This gives the following:

$$\begin{array}{ll} \text{Epoch 1 : } 1.501 \cdot 10^{-16} \text{ g/cm}^3 & \text{Epoch 2 : } 0.505 \cdot 10^{-16} \text{ g/cm}^3 \\ \text{Epoch 3 : } 1.299 \cdot 10^{-16} \text{ g/cm}^3 & \text{Epoch 4 : } 3.752 \cdot 10^{-16} \text{ g/cm}^3 \end{array}$$

By plotting these numbers as a function of radius we can obtain a density profile. Since the radius depends on velocity and time as $R = vt$, where t is now the same for all densities, the density profile can instead be plotted as a function of velocity (v_{stop}). This is shown in Figure 20. As seen in the figure the four densities do not seem to show a clear tendency, but in spite of this we try to fit a power law profile to the data. To do the fitting we use the Python package `iminuit`, which gives us a power law index of $n = -2.4 \pm 0.8$. Even though the data points do not show a clear connection, the obtained power law index is consistent with predictions that the density of the inner part of the transient decreases gradually or is close to flat (Chevalier & Soker, 1989; Kasen et al., 2017).

The lack of a clear connection between the data points may be due to difficulties fitting the later epochs, especially due the emission line centered at $1.1\mu\text{m}$. As previously mentioned this caused us to exclude parts of the spectra when calculating the likelihood. What impact this approach had on the final parameters is not known, but it is likely that we would have found different and more reliable values if we had been able to fit the emission line.

Apart from difficulties when fitting and thus possibly obtaining unreliable values for the parameters, we do not know what kind of density profile to expect for the kilonova. If the profile can be described by a single power law, four data points may be enough to constrain the function. However, it has been suggested that the




Figure 20: The densities from Table 3 converted into densities at the same time after explosion and plotted as a function of v_{stop} . Since the radius is proportional to the velocity this gives a density profile of the kilonova. The data points have been fitted with a power law, which gives a power law index of $n = -2.4 \pm 0.8$.

density profile is more complex than this and instead should be described by a broken power law density profile such that the density declines gradually at small radii and drops steeply at the edge with a power law index of $n \sim -10$ (Chevalier & Soker, 1989; Kasen et al., 2017). In this case we would need more data points to get a better estimate of the profile.

In this work we have assumed that the transient is in LTE. This means that the density is high enough for collisions between atoms and ions/electrons to both excite and de-excite the atoms, and as a result the transient can be described by a single temperature. In a plasma in NLTE collisions excite the atoms, but they are de-excited through radiative emission. This gives long lifetimes and large densities of metastable states, which means that an environment in NLTE allows for a substantial number of forbidden transitions, and a spectrum from such an object will therefore contain forbidden lines. SrII has two forbidden line transition at rest frame wavelengths of 6738.392Å and 6868.171Å (Kramida et al., 2020), but these lines have not been identified in the observed spectrum of the kilonova, which means that it cannot be ruled out that the transient is in LTE.

Every spectral line has a critical density below which the radiative de-excitation rate exceeds the collisional de-excitation rate, and thus the population causing the line will be in NLTE. For this reason it is of great interest to calculate the electron

density of the plasma in order to be able to investigate whether a specific line is governed by collisions or radiative de-excitations. The electron density is used rather than the density of nuclei, since it is the number of free electrons that dictates the collisional recombinations in the plasma. This is due to the fact that it is the electrons that are responsible for de-exciting the atoms through collisions. The nuclei are namely too heavy, and as a result they move too slowly to have a considerable impact on the rate of collisions.

The electron density is calculated from the density of the first epoch using that the mass of a nucleon is $m = 1.66 \cdot 10^{-27}$ kg and assuming that the mean atomic mass of the nuclei inside the kilonova is 100. This number is obtained by integrating the r-process abundance distribution (as seen in Figure 3) and finding the mean value, which is approximately 100. Since the atoms are singly ionized on average, there will be one electron per nucleus and thus the electron density is equal to the density of nuclei. This gives:

$$\rho_e = \frac{1}{100} \frac{\rho_0}{m} = 10^6 \text{cm}^{-3} \quad (17)$$

One can then consider specific lines and calculate their critical densities based on their collision rate and their spontaneous decay rate. The former depends on the temperature and density of the plasma, while the latter is simply the Einstein A coefficient of the specific line. As such, the critical density is the density at which these two rates are equal, and by comparing this to the electron density one can tell whether the line is governed by collisional or spontaneous decay. Investigating the critical densities of a number of lines may thereby indicate whether the plasma is in LTE or if it may have NLTE effects.

The method used to determine distances in this work is similar to the Spectral-fitting Expanding Atmosphere Method (SEAM) (e.g. Baron et al., 1995, 1996; Baron et al., 2004; Baron et al., 2007) used to constrain distances to Type II-P SNe. SEAM works by comparing observed spectra from Type II-P SNe to synthetic spectra obtained from detailed NLTE radiative transfer codes. This gives a spectral energy distribution, and one can then calculate the absolute magnitude of the supernova in a specific photometric band. In combination with the apparent magnitude in the same band, this is used to find the distance modulus and hence the distance to the object. Despite using simplified models to create the synthetic spectra, the distances inferred to the SNe are consistent with the Cepheid distances to the host galaxies of the SNe (Baron et al., 2004). The method assumes that the explosion is spherically symmetric, which is the reason why Type II-P SNe is the chosen source for the analysis. Due to their intact hydrogen envelope they are namely the type of SNe that closest resembles a spherically symmetric explosion. The method have primarily been applied to SNe within the local universe, but its

independence from the distance ladder makes it a promising tool for cosmology.

5.1 Determining the Hubble constant, H_0

The inferred distances in Table 3 and 4 can be used to place constraints on the Hubble constant, H_0 . This constant describes the rate at which the universe is expanding, and there are commonly two ways of determining it. One uses Cepheid variable stars to calibrate the distances to Type Ia supernovae, and by measuring the redshift and the observed flux of these one can determine H_0 . This method is also known as the late universe method or the distance ladder method. The other method used to find H_0 fits the standard Λ CDM model to the observed temperature and anomalies of the cosmic microwave background (CMB). This method is known as the early universe method. The two different methods yield very different results for H_0 . Estimates based on the late universe method are generally higher (e.g. $73.24 \pm 1.74 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Riess et al., 2016)) than the estimates found by the early universe method (e.g. $67.4 \pm 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Aghanim et al., 2020)).

Using the distances and the cosmological redshift of the host galaxy NGC 4993 one can estimate H_0 . The cosmological redshift is the total redshift of the galaxy, where the redshift due to peculiar motions has been subtracted. According to Hjorth et al. (2017), the cosmological redshift of the galaxy is $z_{\text{cosmic}} = 0.00980 \pm 0.00079$. At small redshifts, i.e. $z \ll 1$, the luminosity distance to an object is given by (Ryden, 2003):

$$D_L \approx \frac{c}{H_0} z_{\text{cosmic}} \left[1 + \frac{1 - q_0}{2} z_{\text{cosmic}} \right] \quad (18)$$

where q_0 is called the deceleration parameter and is set to -0.53 for standard cosmological parameters. Inserting the luminosity found for the first epoch spectrum of $(38.58 \pm 0.002) \text{ Mpc}$, which is the smallest distance inferred, we find the following value for the Hubble constant:

$$H_0 = 76.77 \pm 6.24 \text{ km s}^{-1} \text{ Mpc}^{-1} \quad (19)$$

where the uncertainty is primarily governed by the uncertainty on the redshift. The parameters found for the fourth epoch do not seem to be consistent with the other epochs, especially in terms of the distance and the luminosity, and for this reason we do not consider the distance inferred from the fourth epoch spectrum to be reliable. Apart from this distance the largest one is that of $(41.47 \pm 0.10) \text{ Mpc}$. This distance yields a H_0 of:

$$H_0 = 71.42 \pm 5.80 \text{ km s}^{-1} \text{ Mpc}^{-1} \quad (20)$$

The calculations above show that the inferred distances to AT2017gfo imply that H_0 is likely somewhere between 71 and 77 $\text{km s}^{-1} \text{Mpc}^{-1}$. This means that our estimate is more consistent with the Hubble constant inferred from other astronomical objects in the late universe. In Figure 21 we have plotted a number of estimates of H_0 based on both methods together with the H_0 values obtained in this work. Despite the large uncertainties compared to previous estimates the two mean values seem to be more consistent with the late universe estimates as first assumed.




Figure 21: A number of estimates of the Hubble constant obtained using the early universe method and the late universe method¹¹. These are plotted together with the two estimates from this analysis.

The two estimates of H_0 and the fact that they are consistent with previous estimates also proves that the electromagnetic radiation resulting from a neutron star merger is useful for constraining H_0 and may become an important tool for cosmology in the future.

5.2 Constraining the strontium abundance

Apart from determining the four parameters shown in Table 3 it is of great interest to investigate the presence of r-process elements in the kilonova. Primarily due to the lack of a detailed line list and the computational cost of running a large number of simulations, we will not be able to make a thorough analysis of the

¹¹https://en.wikipedia.org/wiki/Hubble%27s_law

abundances of all the heavy elements. However, as found by [Watson et al. \(2019\)](#) the spectra show remarkable amounts of Sr. But even though it has been proved that Sr is present in the spectra of the kilonova, the abundance of the element remains unknown. For this reason we limit the analysis of r-process elements to Sr and look into the abundance of this element only. By investigating the amount of Sr present during each epoch, one can determine the evolution of the element and thus find an indication of how the abundances of light r-process elements change over time.

Table 3 shows that the uncertainties on the inner and outer velocities are small meaning that they are well constrained. The corner plot in Figure 16 additionally shows that there is no correlation between the two parameters. In order to reduce the number of free parameters when investigating the Sr abundance we therefore choose to keep the velocities constant at the values listed in Table 3. Next, we leave the Sr abundance as a free parameter instead of keeping it at a value equal to the solar r-process Sr abundance. We then re-run the simulations with the four free parameters: D_L , L , ρ_0 and Sr .

The results after running the simulations and the MCMC to constrain the best fit parameters are shown in Table 4, and the corresponding spectra are plotted in Figure 22 together with the X-shooter spectra and the spectra obtained with a solar r-process Sr abundance. The corner plot for the first epoch is shown in Figure 23, and one thing to note about the plot is the fact that the Sr abundance and the density are anti-correlated. This means that the smaller the abundance is, the larger the density is. This is likely because of the fact that the main absorption feature originates from three Sr lines, and thus the total density will strongly depend on the presence of Sr such that a small abundance will be compensated for with a large density. The corner plots for the last three epochs can be found in Appendix C.

As seen in the figure the simulated spectra deviate from the X-shooter spectra similarly to the spectra obtained when keeping the Sr abundance at a constant value. For this reason, the true values of the inferred parameters are likely to be different than the ones obtained here, but they are sufficient to get an idea of the evolution of the abundance. The values show that there is a tendency in the Sr abundance going from several times the solar r-process Sr abundance down to approximately solar. Hence we find that the Sr abundance decreases over time.

We can calculate the expected decrease in Sr abundance due to expansion of the transient. The abundance depends on the column density N of the ejecta, such that as the ejecta expands the column density decreases with radius as $N \propto R^{-2}$,

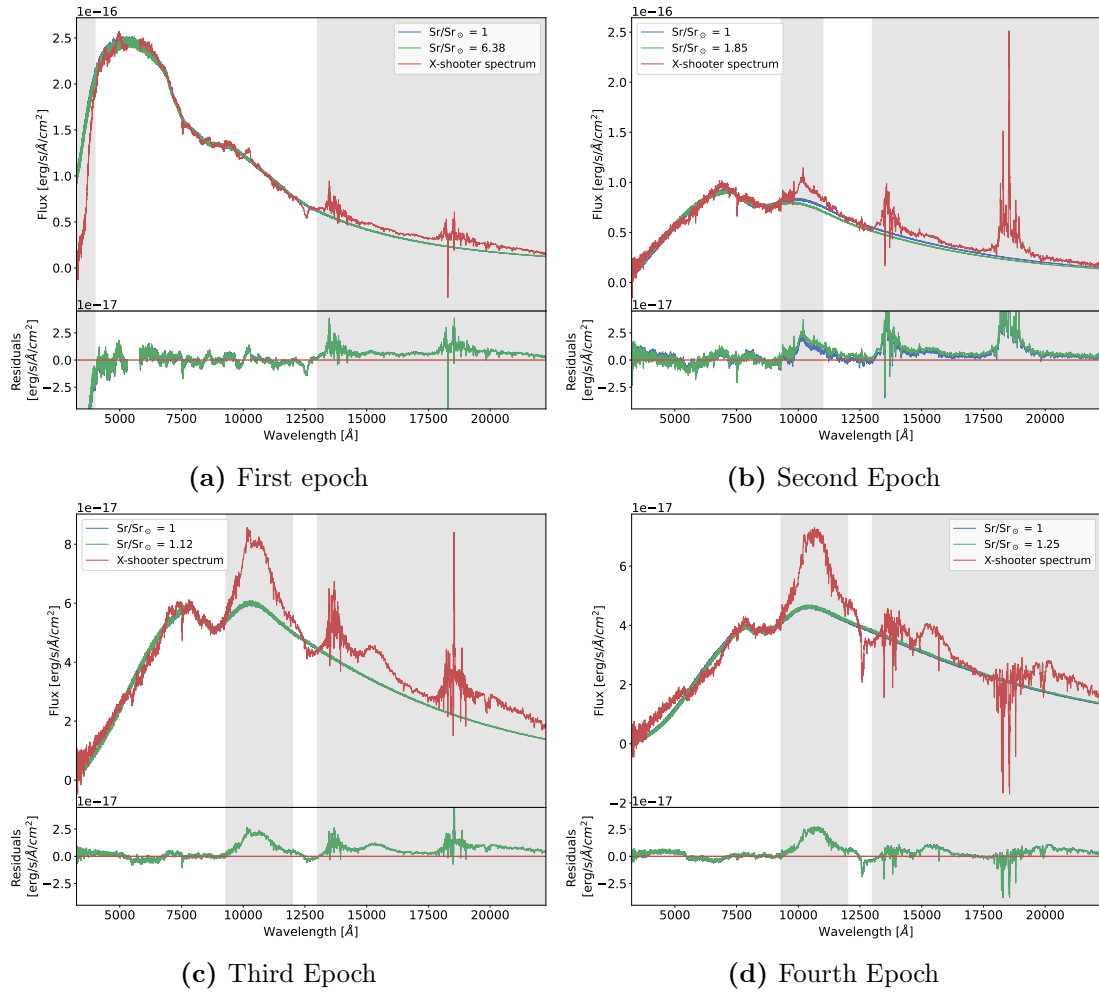


Figure 22: The best fit TARDIS spectra with the Sr abundance left as a free parameter plotted together with the observed X-shooter spectra and the best fit spectra with solar Sr abundances. The grey zones were excluded from fitting for reasons described in the text.

Table 4: Results for the best fit parameters for the first four epochs when keeping the velocities of the boundaries constant and leaving the abundance of Sr as a free parameter. The uncertainties are the 2σ deviation from the mean.

Epoch No.	D_L [Mpc]	L [erg/s] $\times 10^{41}$	ρ_0 [g/cm ³] $\times 10^{-16}$	Sr/Sr _⊙
1	38.257 ± 0.009	3.4853 ± 0.0012	0.571 ± 0.005	6.38 ± 0.11
2	41.13 ± 0.08	2.506 ± 0.004	0.0493 ± 0.0010	1.85 ± 0.04
3	41.38 ± 0.09	2.097 ± 0.004	0.0889 ± 0.0014	1.12 ± 0.02
4	47.66 ± 0.09	2.284 ± 0.004	0.1013 ± 0.0027	1.25 ± 0.04

and since $R = tv$ we find:

$$\frac{[Sr]_1}{[Sr]_2} \propto \frac{N_1}{N_2} \propto \left(\frac{R_1}{R_2}\right)^{-2} \propto \left(\frac{t_1 v_1}{t_2 v_2}\right)^{-2} \quad (21)$$

Inserting the velocities for the first and second epoch we find the expected Sr abundance ratio between the two epochs to be:

$$\frac{[Sr]_1}{[Sr]_2} \approx \left(\frac{1.427 \text{ days} \cdot 0.8696 \cdot 10^8 \text{ m/s}}{2.41 \text{ days} \cdot 0.7570 \cdot 10^8 \text{ m/s}}\right)^{-2} \approx 2.16 \quad (22)$$

As seen in Table 4 the inferred Sr abundances give a ratio of ~ 3.45 , which is somewhat larger than the ratio calculated above. The abundances are primarily due to SrII, since the pronounced absorption feature around $1\mu\text{m}$ originates from three SrII lines. Assuming that the values in Table 4 are correct, finding a lower Sr abundance for the second epoch than expected suggests that the ionization balance has changed such that SrII has been ionized into SrIII. This would weaken the absorption feature around $1\mu\text{m}$ and hence result in a lower Sr abundance. Furthermore, as the ejecta expands the optical depth increases, and since the Sr abundance decreases it suggests that there is not much Sr further within the transient.




Figure 23: Corner plot for the first epoch spectrum with the inner and outer velocities kept constant and the Sr abundance set as a free parameter.

6 Future perspectives

To be able to perform this analysis we have had to make a number of assumptions that are worth reviewing in hindsight. Furthermore, we have encountered some challenges that future work on this topic should focus on solving. This section contains a review of the method used for this analysis and comes with a few suggestions to changes and updates for future work.

One of the greatest limitations performing this analysis was runtime challenges. Ideally we would have scanned over more parameters, and especially for the later epochs we would have avoided keeping the ratio between the velocities constant. Furthermore, we would have preferred making a finer likelihood grid by running the simulations with more than ten different values for each parameter. However, implementing just one additional parameter would have made the simulations run for approximately 20 days, which would have been infeasible.

One way of decreasing the runtime would be to construct a so-called emulator (or surrogate model). An emulator is a neural network proxy trained on a large number of models with known parameters, and their purpose is to serve as a more computationally effective method than the original function. TARDIS can be considered as a function that takes a number of input parameters and transforms them into a spectral model, and this makes it a perfect example of a function that can be used to train an emulator. This was done by [Kerzendorf et al. \(2021\)](#), whose aim was to construct an emulator used for modelling of Type 1a supernovae. They ran a large number of TARDIS simulations varying the inner boundary velocity, temperature and nine abundances and leaving the remaining parameters constant. As such, the emulator was trained on a total of 68 600 TARDIS spectra using different neural network architectures. The five best performing neural network architectures were then combined to construct the final emulator. Finally, the spectra created by the emulator were compared to spectra created by TARDIS, and this comparison showed that the emulator could predict the TARDIS spectra down to a precision of 1%. Furthermore, the runtime of the emulator was $\sim 10^{-4}$ times the runtime of TARDIS. This method can be applied to other cases than Type 1a SNe, so training such an emulator could be useful for future work when modelling a large number of kilonova spectra.

When running the large number of TARDIS simulations we were limited by computational resources. The simulations required a lot of memory and caused the server to crash when pushing it to the limit. As mentioned above we would preferably have had more than ten values per parameter such that the likelihood grid was better resolved. One can also argue that we were a bit conservative with the ranges chosen to scan over, and that we could have picked smaller ones around

the best fit value. This could have solved the issue with the low resolution of the grid, but by doing so we would have run a risk of excluding parts of the parameter space with higher likelihood and thereby bias the results.

As mentioned above [Kerzendorf et al. \(2021\)](#) trained their emulator on 68 600 TARDIS spectra, and including the spectra used to test their emulator, they had to run a total of 98 000 simulations. The reason that they were able to run this many simulations in a reasonable amount of time is the fact that their output spectrum was the integrated spectrum rather than the v-packet spectrum, which was used in this analysis. In brief, the integrated spectrum is calculated by integrating the line source functions calculated during the simulation rather than measuring and binning the energy and frequency of the packets. As a result, the integrated spectrum is noise-free, so using this spectrum means that one can run the simulations with fewer packets. As such, [Kerzendorf et al. \(2021\)](#) ran their simulations with 10^5 packets for the last iteration, whereas we used 10^8 packets to get down to a reasonable SNR/Å. The memory required to run a simulation increases drastically with the number of packets, so using the integrated spectrum seems to be an obvious solution. However, we found that the integrated spectrum does not resemble the r- and v-packet spectra well. These three spectra have been plotted in [Figure 24](#) for the simulation of the first epoch with the best fit parameters. The figure shows that the integrated spectrum does not agree with the two other spectra, especially around the Sr absorption feature at $\sim 1\mu\text{m}$. For this reason we chose to work with the v-packet spectra, because we could not trust the quality of the integrated spectrum. This made the simulations computationally expensive, but future work could benefit from the integrated spectrum if one can find a way to calculate these without the discrepancies.

For this analysis we set the abundances of the r-process elements to that of the Sun. As previously mentioned this is very unlikely to be the case, since including all elements from As to U and letting the abundances be equal to that of the Sun results in a spectrum with too much absorption at blue wavelengths compared to the observed spectrum ([Figure 8](#)). Therefore the analysis can be expanded into leaving the abundance parameters as free parameters.

Furthermore, including more elements could potentially reveal information about other spectral features in the observed spectrum. As pointed out by [Watson et al. \(2019\)](#) there are two emission lines at NIR wavelengths that appear to increase in strength over time. These are centered at $\sim 1.6\mu\text{m}$ and $\sim 2.1\mu\text{m}$, and they can be seen in [Figure 25](#), where all X-shooter spectra are plotted and normalized in order to highlight the spectral features and their evolution. These peaks can also be seen in the residual plots of [Figure 15](#) and [22](#), and it is clear to see here too how they are more pronounced at later epochs compared to the first epoch. These emission lines have not yet been addressed, and thus future




Figure 24: The r-packet, v-packet and integrated spectra plotted for the same TARDIS simulation and showing inconsistencies between the two packet spectra and the integrated spectrum.

work focusing on the implementation of more elements in modelling the spectra could bring new information about the origin of these peaks. They likely contain important information about the plasma and its evolution, so constraining their origin is of great interest.

The method used in this work will be particularly useful for investigating r-process abundances when a more detailed line list is available and hence more features can be identified. However, investigating the abundances of the heavy elements will quickly increase the number of free parameters, so it will require a more effective method for calculating the spectra. For instance the simulation plotted in Figure 8 took approximately two days to run, which is not feasible when running a large number of simulations.




Figure 25: All X-shooter spectra normalized and with an offset.

Another way of improving this method would be to combine the MCMC and the TARDIS simulations such that instead of creating a grid of spectra and run a MCMC on the interpolated likelihood function, the TARDIS spectra are calculated for each walker for every iteration whenever they move to a new set of parameters. By running the simulations this way one would only have to calculate the necessary simulations instead of trying to cover an entire parameter space. This

method would neither suffer from the previously mentioned bug causing TARDIS to miscalculate spectra that are not normalized correctly. This is due to the fact that when running a Markov chain the next sample will only be accepted if the likelihood is greater than the current sample, so if a spectrum is miscalculated it will be rejected and the process will continue. This approach should then be combined with the use of an emulator in order to decrease the runtime of every MCMC iteration.

The analysis could furthermore benefit from an extension of TARDIS. As already discussed TARDIS assumes that the transient expands homologously which is not necessarily the case, and we find a distance to the kilonova that implies that the transient is elongated. Also, the code does not allow us to fit nebular emission, which becomes important at later epochs. TARDIS neither takes the effect of light delay into account, which may be important for kilonovae, where the ejecta moves at very high velocities. Thus, extending TARDIS to simulate in two or three dimensions and implement nebular emission and light delay would be useful to study the spectra of kilonovae.

Many assumptions were made throughout this analysis to get to the the results listed in Table 3 and 4. These assumptions mean that there are a number of systematic errors that have not been taken into account, which would have increased the uncertainties on the parameters. As previously mentioned, these assumptions include modelling the transient in one dimension, including only a few metals, not taking nebular emission and light delay into account and assuming that the environment is in LTE. Despite the numerous assumptions, the inferred distances are strikingly close to the the distances determined in previous works using other methods. These results either indicate that the assumptions have cancelled out with each other, or that they may not have had as big of an impact on the observed spectrum. This is promising for future work.

7 Conclusion

In this work, we used the radiative transfer code TARDIS to create a large number of synthetic spectra with the purpose of modelling the spectra of the first four epochs of the kilonova AT2017gfo. These spectra were compared to the observed X-shooter spectra in order to study the properties of the transient 1.5 to 4.5 days after the explosion. Using the Markov chain Monte Carlo method to compare the synthetic spectra to the observed spectra we were able to place constraints on five parameters: The distance (D_l), luminosity (L), density (ρ_0), velocity of the photosphere (v_{start}) and velocity of the outer edge (v_{stop}). Of these parameters the distance was of particular interest. Furthermore, we investigated the evolution of the strontium (Sr) abundance over the four epochs and used our distance estimates to calculate the Hubble constant, H_0 .

From our spectral analysis we found the best fit model for the first epoch spectrum and consider the parameters of this epoch to be the most reliable. The analysis revealed a distance to the kilonova of (38.58 ± 0.02) Mpc. In spite of the distance being shorter than the estimates found in the literature, it is still within the uncertainties of these. The inferred distance furthermore has a small uncertainty which is due to the fact that it only accounts for statistical errors, and the systematical ones have likely been underestimated.

From our distance estimates we calculated the Hubble constant and found a value between approximately 71 and 77 $\text{km s}^{-1} \text{Mpc}^{-1}$. This estimate is consistent with previous results based on late universe measurements, and since being independent from the distance ladder method it has been proved that kilonovae may become an important tool for cosmology.

Furthermore, we investigated the presence of Sr in the four epoch spectra, and found that the abundance decreases over time.

Many assumptions were made throughout this analysis, and a number of improvements are to be implemented in future work. This includes an extension of TARDIS in order to implement nebular emission and an additional dimension, a combination of the MCMC and running the TARDIS simulations and the training of an emulator using neural networks to find a less computationally expensive way of obtaining the spectra. Even though many assumptions were made we found a distance strikingly close to the previously inferred distances, which is promising for future work. Thus, with the implementation of the mentioned improvements the method may become a useful tool to investigate the presence of r-process elements in the ejecta of AT2017gfo and future kilonovae.

References

- Abbott, B., Abbott, R., Abbott, T., Abraham, S., & Acernese, F. 2019a, *Physical Review X*, 9
- Abbott, B., et al. 2017, *Nature*, 551, 85–88
- Abbott, B., et al. 2017a, *Physical Review Letters*, 119
- . 2017b, *The Astrophysical Journal*, 848, L12
- . 2019b, *Physical Review X*, 9
- Aghanim, N., et al. 2020, *Astronomy Astrophysics*, 641, A6
- Barlow, R. J. 1989, *Statistics: A Guide to the Use of Statistical Methods in the Physical Sciences* (Manchester Physics Series), reprint edn. (WileyBlackwell)
- Baron, E., Branch, D., & Hauschildt, P. H. 2007, *ApJ*, 662, 1148
- Baron, E., Hauschildt, P. H., Branch, D., Kirshner, R. P., & Filippenko, A. V. 1996, *MNRAS*, 279, 799
- Baron, E., et al. 1995, *ApJ*, 441, 170
- Baron, E. A., Nugent, P. E., Branch, D., & Hauschildt, P. H. 2004, *Astrophys. J. Lett.*, 616, L91
- Baumgarte, T. W., Shapiro, S. L., & Shibata, M. 2000, 528, L29
- Bauswein, A., Baumgarte, T. W., & Janka, H.-T. 2013a, *Physical Review Letters*, 111
- Bauswein, A., Goriely, S., & Janka, H.-T. 2013b, *The Astrophysical Journal*, 773, 78
- Bisterzo, S., Travaglio, C., Gallino, R., Wiescher, M., & Käppeler, F. 2014, *The Astrophysical Journal*, 787, 10
- Bombaci, I. 1996, *A&A*, 305, 871
- Burbidge, E. M., Burbidge, G. R., Fowler, W. A., & Hoyle, F. 1957, *Rev. Mod. Phys.*, 29, 547
- Busso, M., Gallino, R., & Wasserburg, G. J. 1999, *Annual Review of Astronomy and Astrophysics*, 37, 239
- Caiazzo, I., et al. 2021, *Nature*, 595, 39–42
- Cantiello, M., et al. 2018, *The Astrophysical Journal*, 854, L31
- Capano, C. D., et al. 2020, *Nature Astronomy*, 4, 625–632
- Chandrasekhar, S. 1935, *Monthly Notices of the Royal Astronomical Society*, 95, 207
- Chevalier, R. A., & Soker, N. 1989, *ApJ*, 341, 867
- Cook, G. B., Shapiro, S. L., & Teukolsky, S. A. 1992, *ApJ*, 398, 203
- Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, *Nature*, 340, 126
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306–312
- Freedman, R., Geller, R., & Kaufmann, W. 2010, *Universe* (W. H. Freeman)
- Freedman, W. L., et al. 2001, *The Astrophysical Journal*, 553, 47
- Gillanders, J. H., McCann, M., Sim, S. A., Smartt, S. J., & Ballance, C. P. 2021, *MNRAS*, 506, 3560
- Goodman, J., & Weare, J. 2010, *Communications in Applied Mathematics and Computational Science*, 5, 65
- Hjorth, J., et al. 2017, *The Astrophysical Journal*, 848, L31
- Hogg, D. W., & Foreman-Mackey, D. 2018, *The Astrophysical Journal Supplement Series*, 236, 11
- Hotokezaka, K., Kiuchi, K., Kyutoku, K., Okawa, H., Sekiguchi, Y.-i., Shibata, M., & Taniguchi, K. 2013, *Phys. Rev. D*, 87, 024001
- Hotokezaka, K., Kyutoku, K., Okawa, H., Shibata, M., & Kiuchi, K. 2011, *Physical Review D*, 83
- Ivanenko, D. D., & Kurdgelaidze, D. F. 1965, *Astrophysics*, 1, 251
- Kalogera, V., & Baym, G. 1996, *The Astrophysical Journal*, 470, L61
- Kasen, D., Fernández, R., & Metzger, B. D. 2015, *Monthly Notices of the Royal Astronomical Society*, 450, 1777–1786
- Kasen, D., Metzger, B., Barnes, J., Quataert, E., & Ramirez-Ruiz, E. 2017, *Nature*, 551, 80–84
- Kerzendorf, W., et al. 2021, *tardis-sn/tardis: TARDIS v2021.10.1*
- Kerzendorf, W. E., & Sim, S. A. 2014, *MNRAS*, 440, 387
- Kerzendorf, W. E., Vogl, C., Buchner, J., Contardo, G., Williamson, M., & van der Smagt, P. 2021, *ApJL*, 910, L23

- Kramida, A., Yu. Ralchenko, Reader, J., & and NIST ASD Team. 2020, NIST Atomic Spectra Database (ver. 5.8), [Online]. Available: <https://physics.nist.gov/asd> [2021, October 25]. National Institute of Standards and Technology, Gaithersburg, MD.
- Kratz, K., Farouqi, K., Pfeiffer, B., Truran, J. W., Sneden, C., & Cowan, J. J. 2007, *The Astrophysical Journal*, 662, 39–52
- Kurucz, R. L. 2017, *Canadian Journal of Physics*, 95, 825
- Lattimer, J. M., & Schramm, D. N. 1974, *ApJL*, 192, L145
- Lee, M. G., Kang, J., & Im, M. 2018, *The Astrophysical Journal*, 859, L6
- Li, L.-X., & Paczyński, B. 1998, *ApJL*, 507, L59
- Lodders, K., Palme, H., & Gail, H.-P. 2009, *Landolt-Börnstein - Group VI Astronomy and Astrophysics*, 712–770
- Long, K. S., & Knigge, C. 2002, *The Astrophysical Journal*, 579, 725–740
- Margalit, B., & Metzger, B. D. 2017, *The Astrophysical Journal*, 850, L19
- Metzger, B. D. 2019, *Living Reviews in Relativity*, 23
- Metzger, B. D., & Fernández, R. 2014, *Monthly Notices of the Royal Astronomical Society*, 441, 3444–3453
- Metzger, B. D., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 406, 2650–2662
- Miller, J. M., et al. 2019, *Physical Review D*, 100
- Oppenheimer, J. R., & Volkoff, G. M. 1939, *Phys. Rev.*, 55, 374
- Paczynski, B. 1986, *ApJL*, 308, L43
- Pauli, W. 1925, *Zeitschrift für Physik*, 31, 765
- Pian, E., et al. 2017, *Nature*, 551, 67–70
- Piskunov, N. E., Kupka, F., Ryabchikova, T. A., Weiss, W. W., & Jeffery, C. S. 1995, , 112, 525
- Podsiadlowski, P., Mazzali, P. A., Nomoto, K., Lazzati, D., & Cappellaro, E. 2004, *The Astrophysical Journal*, 607, L17–L20
- Pássaro, E. A., E., K. W., A., F., & C., V. 2019, *An Open Source Framework for Atomic Data*
- Rezzolla, L., Most, E. R., & Weih, L. R. 2018, *The Astrophysical Journal*, 852, L25
- Riess, A. G., et al. 2016, *The Astrophysical Journal*, 826, 56
- Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2018, *Physical Review D*, 97
- Ryabchikova, T., Piskunov, N., Kurucz, R. L., Stempels, H. C., Heiter, U., Pakhomov, Y., & Barklem, P. S. 2015, , 90, 054005
- Ryden, B. 2003, *Introduction to cosmology*
- Sakai, S., et al. 2000, *The Astrophysical Journal*, 529, 698
- Schwab, J., Bildsten, L., & Quataert, E. 2017, *Monthly Notices of the Royal Astronomical Society*, 472
- Shibata, M., Fujibayashi, S., Hotokezaka, K., Kiuchi, K., Kyutoku, K., Sekiguchi, Y., & Tanaka, M. 2017, *Phys. Rev. D*, 96, 123012
- Shibata, M., Zhou, E., Kiuchi, K., & Fujibayashi, S. 2019, *Physical Review D*, 100
- Smartt, S. J., et al. 2017, *Nature*, 551, 75–79
- Sneden, C., et al. 2003, *ApJ*, 591, 936
- Symbalisty, E., & Schramm, D. N. 1982, , 22, 143
- Tanaka, M., et al. 2017, *Publications of the Astronomical Society of Japan*, 69
- Tolman, R. C. 1939, *Phys. Rev.*, 55, 364
- Truran, J. W. 1981, *A&A*, 97, 391
- van Ravenzwaaij, D., Cassey, P., & Brown, S. D. 2018, *Psychonomic Bulletin & Review*, 25, 143
- Vogl, C., Sim, S. A., Noebauer, U. M., Kerzendorfer, W. E., & Hillebrandt, W. 2019, *Astronomy & Astrophysics*, 621, A29
- Wanderman, D., & Piran, T. 2015, *Monthly Notices of the Royal Astronomical Society*, 448, 3026–3037
- Watson, D., et al. 2019, *Nature*, 574, 497–500
- Yong, D., et al. 2021, *Nature*, 595, 223–226

A YAML example file

```
tardis_config_version: v1.0

supernova:
  luminosity_requested: 3.541e+41 erg/s
  time_explosion: 1.427 day

atom_data: VALD_U.h5

model:
  structure:
    type: specific
    velocity:
      start: 0.8696e+08 m/s
      stop: 0.9922e+08 m/s
      num: 10
    density:
      type: power_law
      rho_0: 1.501e-16 g/cm^3
      v_0: 0.8696e+08 m/s
      exponent: -3
  abundances:
    type: uniform
    As: 0.041942818873718424
    Se: 0.4641213563895073
    Br: 0.07357182982767003
    Kr: 0.3303430166237955
    Rb: 0.03998250778176807
    Sr: 0.05003847050354071

plasma:
  disable_electron_scattering: false
  ionization: lte
  excitation: lte
  radiative_rates_type: dilute-blackbody
  line_interaction_type: macroatom

montecarlo:
  seed: 23111963
```



```
no_of_packets: 1.e+5
iterations: 30
nthreads: 1
last_no_of_packets: 1.e+8
no_of_virtual_packets: 5
convergence_strategy:
  type: damped
  damping_constant: 1.0
  threshold: 0.05
  fraction: 0.8
  hold_iterations: 3
  t_inner:
    damping_constant: 1.0

spectrum:
  start: 3250 angstrom
  stop: 22260 angstrom
  num: 22000
```

B $D_L, L, \rho_0, v_{start}$ Corner plots Epoch 2-4




Figure B.1: Corner plot for the second epoch spectrum.

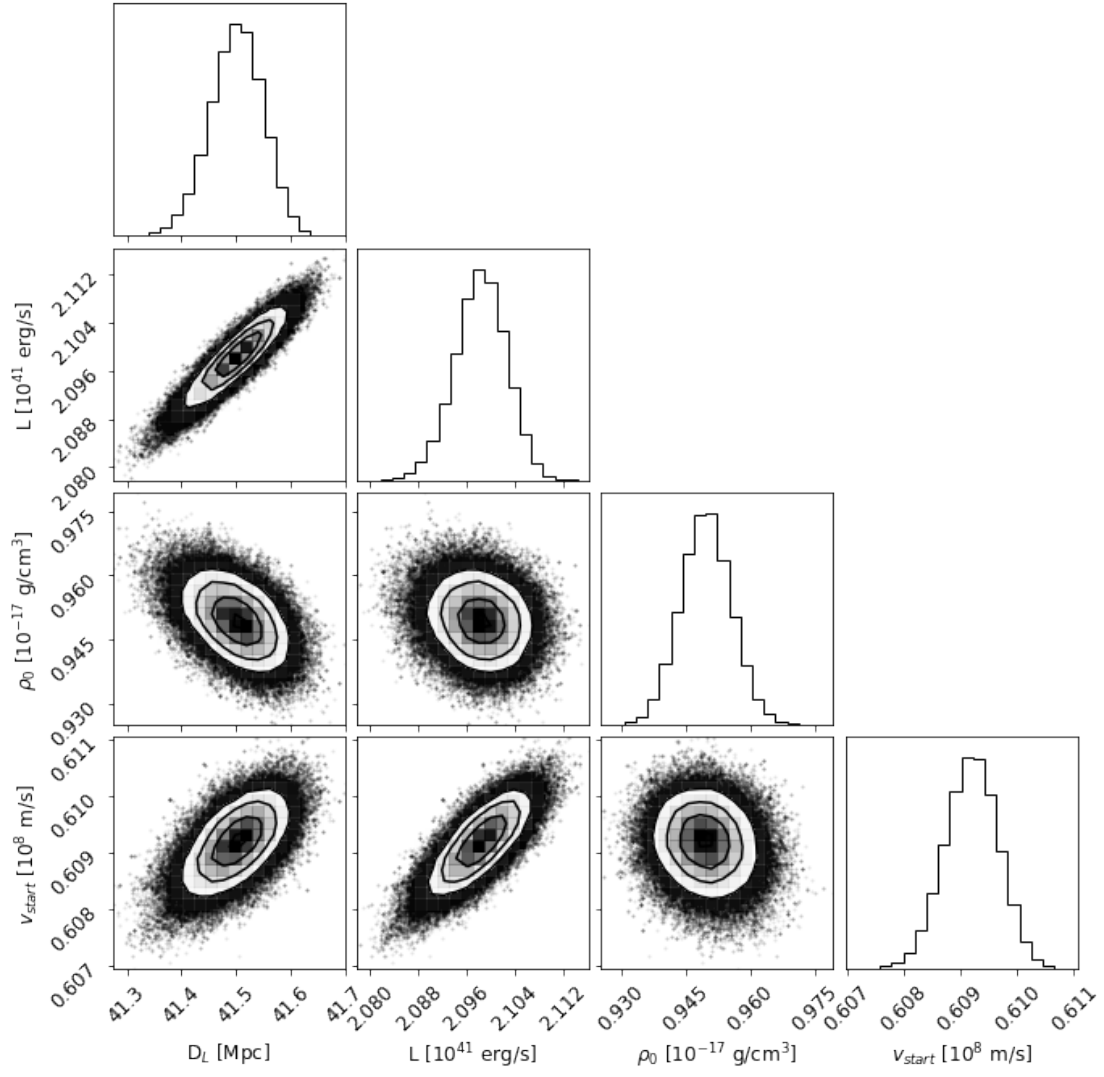


Figure B.2: Corner plot for the third epoch spectrum.

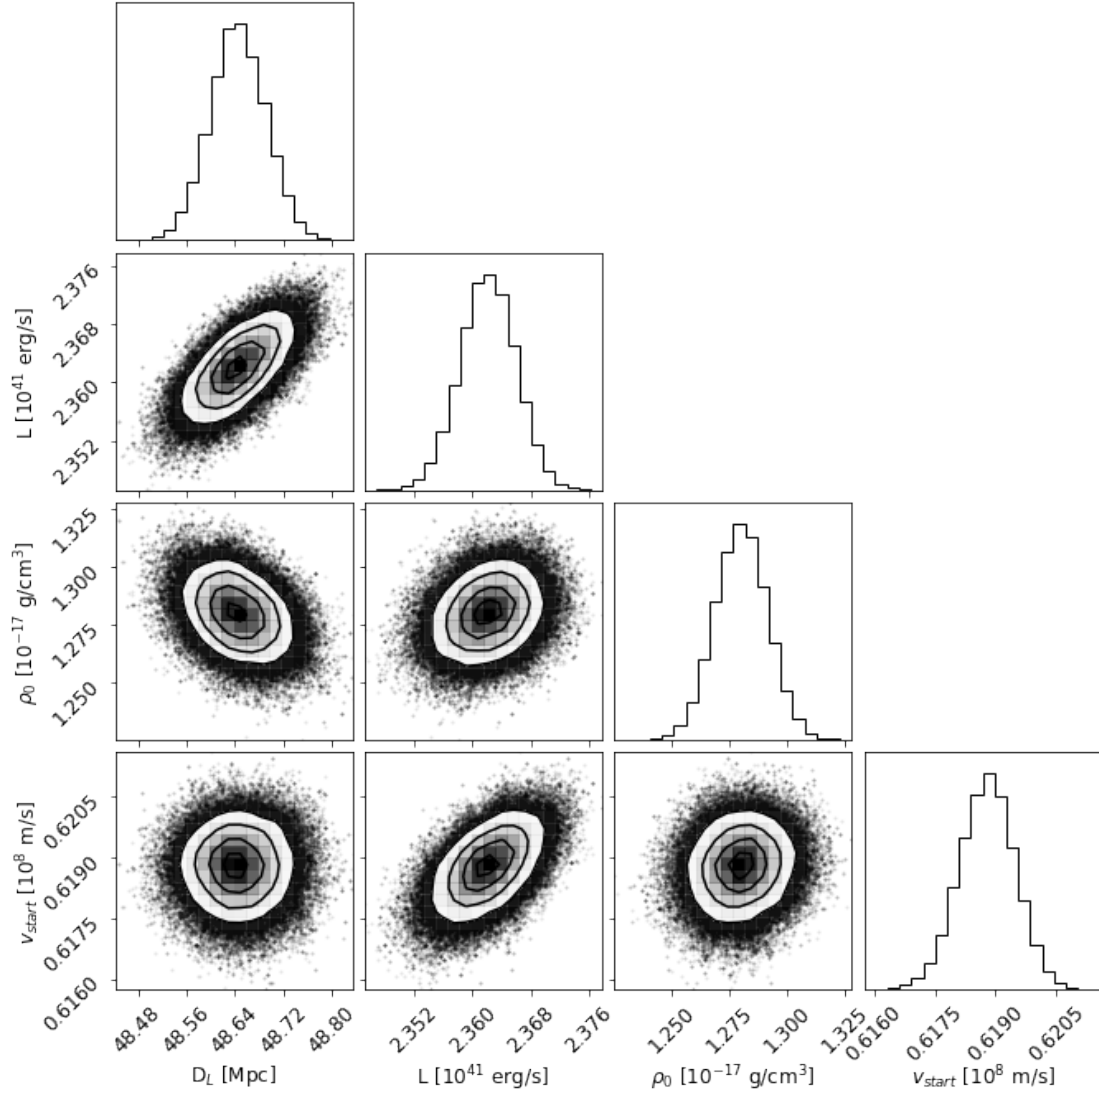


Figure B.3: Corner plot for the fourth epoch spectrum.

C D_L , L , ρ_0 , Sr Corner plots Epoch 2-4

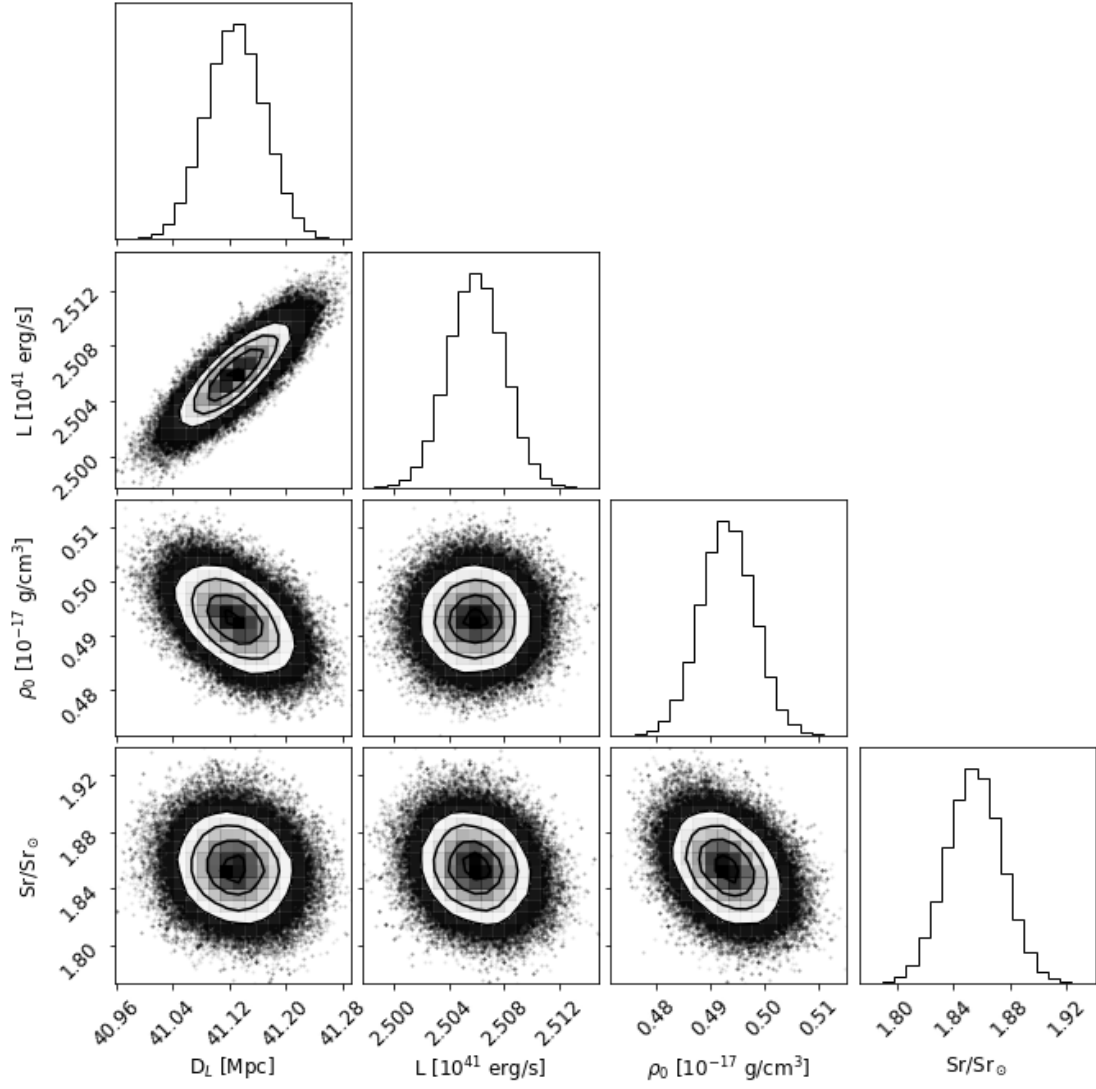


Figure C.1: Corner plot for the second epoch spectrum.




Figure C.2: Corner plot for the third epoch spectrum.

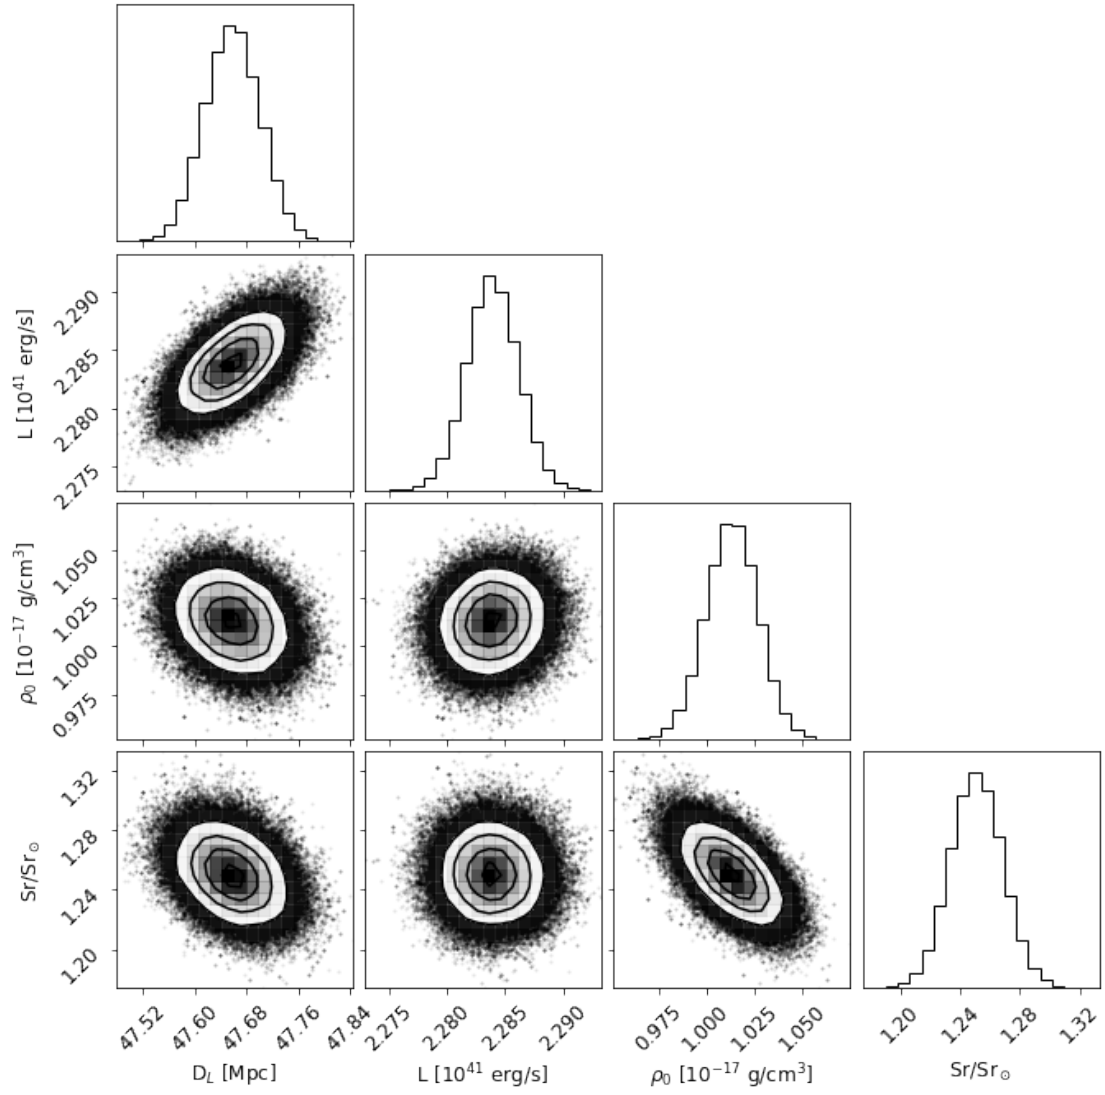


Figure C.3: Corner plot for the fourth epoch spectrum.