



MONASH University

Probing cosmic history with merging
compact binaries

Jeff Riley

PhD, MAppSci(IT), GradDip(KBS)

A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

in the

School of Physics and Astronomy

August 10, 2023

Copyright notice

©Jeff Riley (2023)

This thesis must be used only under the normal conditions of “fair dealing” under the Copyright Act. It should not be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner’s permission.

The only real valuable thing is intuition.
- Albert Einstein (attrib.)

Abstract

The observation of the merger of compact binaries can help constrain the uncertain parameters that describe the underlying processes in the evolution of stars and binary systems. Since the detection of gravitational waves by the advanced Laser Interferometer Gravitational-wave Observatory (aLIGO) in 2015, the catalogue of compact binary mergers has grown to ninety events, and is expected to grow to hundreds of events during the fourth observing run of the LIGO, Virgo, and KAGRA (LVK) collaboration, which began on May 24, 2023. We can use the growing catalogue of compact binary merger detections to probe cosmic history by discovering constraints on the parameters that describe processes that help form the Universe.

Population synthesis is a powerful technique used in astrophysics to model the properties, evolution, and interaction of populations of astronomical objects such as isolated stars, binary stars, multiple star systems (typically triples, but higher-order systems are possible), star clusters, galaxies, and supernovae.

In this thesis, we enhance and extend the Compact Object Mergers: Population Astrophysics and Statistics (COMPAS) stellar and binary population synthesis computer model, then we demonstrate that properties of synthesised populations of merging compact binaries – namely, detection rates, chirp masses, and redshifts – can be used to measure cosmological parameters describing the redshift-dependent star formation rate and metallicity distribution.

We release the COMPAS toolset publicly as open-source software, and we present, as a proof-of-concept, a method that uses artificial neural networks to emulate a binary population synthesis computer model, and construct a fast, flexible, parallelisable surrogate model that can be used for inference. By using the surrogate model to rapidly synthesise populations of merging compact binaries, we can compare properties of the synthesised populations to the catalogue of gravitational-wave events in order to constrain the underlying process of star formation and evolution of stars and binary systems.

Declaration of Authorship

I, Jeff Riley, hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

This thesis includes (3) original papers published in peer reviewed journals and (1) accepted for publication but not yet published. The core theme of the thesis is probing cosmic history by exploring the population properties of simulated compact binary systems. The ideas, development and writing up of all the papers in the thesis were the principal responsibility of myself, the student, working within the School of Physics and Astronomy under the supervision of Professor Ilya Mandel.

My contribution to the work presented in this thesis is shown in the table below. Co-authors marked with an asterisk (*) were at the time of submission/publication enrolled as students at Monash University. The inclusion of co-authors reflects the fact that the work came from active collaboration between researchers and acknowledges input into team-based research.

I have not renumbered sections of submitted or published papers in order to generate a consistent presentation within the thesis.

Student: Jeff Riley

Date: August 10, 2023

I, Ilya Mandel, hereby certify that the above declaration correctly reflects the nature and extent of the student's and co-authors' contributions to this work. In instances where I am not the responsible author I have consulted with the responsible author to agree on the respective contributions of the authors.

Primary Supervisor: Ilya Mandel

Date: August 10, 2023

Published works.

Ch	Title	Status	Contribution	
			Student	Co-author
2	Compact Object Mergers: Population Astrophysics and Statistics (COMPAS)	Published	60% Code design and development; data generation, collection, and analysis; writing of manuscript.	10%, Simon Stevenson. Initial concept and code development; input into manuscript. 7.5% Ilya Mandel. Initial concept; code development; input into manuscript. 5% Alejandro Vigna Gómez. Code development; input into manuscript. 5% Floor Broekgaarden. Code development; input into manuscript. 5% Reinhold Willcox*. Code development; input into manuscript. 7.5% remaining authors*. Code development; input into manuscript.
2	COMPAS: A rapid binary population synthesis suite	Published	60% Code design and development; writing of manuscript.	10%, Simon Stevenson. Code development; input into manuscript. 7.5% Ilya Mandel. Code development; input into manuscript. 5% Alejandro Vigna Gómez. Code development; input into manuscript. 5% Floor Broekgaarden. Code development; input into manuscript. 5% Reinhold Willcox*. Code development; input into manuscript. 7.5% remaining authors*. Code development; input into manuscript.

Published works (cont'd).

Ch	Title	Status	Contribution	
			Student	Co-author
3	Chemically Homogeneous Evolution: A rapid population synthesis approach	Published	75% Code design and development; data generation, collection, and analysis; writing of manuscript.	10%, Ilya Mandel. Concept; input into manuscript. 5% Pablo Marchant. Data generation and collection; input into manuscript. 5% Alejandro Vigna Gómez. Input into manuscript. 5% remaining authors. Prior work; data generation.
4	Surrogate forward models for population inference on compact binary mergers	Accepted for publication in The Astrophysical Journal (ApJ)	90% Surrogate model design; code design and development; data generation and collection; data analysis; writing of manuscript.	10% Ilya Mandel. Concept; likelihood formula; input into manuscript.

Acknowledgements

I won't pretend that I can thank all of the very many people who have helped me along the way as I journeyed through life – one doesn't get to be my age, much less undertake a second PhD at this age, without the help and support of a lot of good and kind people – but I'll do my best.

Ilya, your willingness to take on someone with a passion for, but no formal education in, physics, and your readiness at every step to enable me to do this research with your patient mentoring and guidance, has broadened my horizon and enabled me to pursue my passion for physics. I won't forget that. To say thank you is not sufficient - I owe you a huge debt of gratitude.

A long time ago I took a position with one of the world's leading computer manufacturers. I approached that job with typical youthful confidence, comfortable in the belief that I was pretty good at what I did, and remember most vividly being thoroughly humbled by the calibre of the people I met after starting. To have that happen a second time, much later in life, is in some ways even more humbling. To be sure, I was somewhat less confident starting this PhD (and a lot less youthful), but I was again amazed by the calibre of the people I met after starting. It has been a privilege to work with so many extremely talented and remarkable people over the duration of this PhD. I have learned much, and for that I am indebted to those who took the time and effort. I hope I have also contributed in some way.

To the COMPAS group, COMPAS friends, family, and collaborators, I can't imagine what I would have done without you. I appreciate, and am grateful for, your ready acceptance of someone who was not typical of what you might have been used to as a peer. The local COMPAS group, at various times Reinhold, Coen, Serena, Simon, Mike, Ryo, Avi, and Evgeni – each of you enriched my life with your unique talents and personalities, and I learned a great deal from our interactions. Less local (now), but still connected, Alejandro, Floor, Isobel, and Jade – so talented and impressive, you each have my sincere admiration. All of you, and your friendships, are the enduring memories I'll take away from this time. Many of you have scattered to far-flung places to take up postdoctoral positions, or simply to go home – I miss you all. Thanks for the memories.

And finally, to my immediate family, your support, encouragement, and patience, made this journey not only possible, but all-the-more enjoyable.

My parents always encouraged learning, and would be thoroughly pleased with this circumstance.

Jen, your encouragement of my indulgence in my passions over many years is an inspiration. For that, and many other reasons, you will always have my love and gratitude.

My kids, Tim, Tina, and Nick - you were all born while I was undertaking my first PhD, and now, during my second, you have all recently graduated with university degrees yourselves. I am intensely proud of all of you, and your achievements. You are each so unique, and I encourage you to continue to be yourselves and show the world your individuality. In the words of Maya Angelou:

“If you are always trying to be normal, you will never know how amazing you can be.”

This research was supported in equal parts by the Dean’s scholarship from the School of Physics and Astronomy, the Australian Research Council (ARC) via Future Fellowship FT190100574 (CI Mandel), and the ARC Centre of Excellence for Gravitational-wave Discovery (OzGrav), project CE17010004.

Contents

Copyright notice	i
Abstract	iii
Declaration of Authorship	iv
Acknowledgements	vii
1 Introduction	1
1.1 Astronomical compact objects	1
1.1.1 Compact binaries	3
1.1.1.1 Isolated binary evolution	4
1.1.1.2 Binary evolution through dynamical interactions	4
1.1.1.3 Merging compact binaries	5
1.1.1.4 Gravitational wave detections	6
1.2 Population synthesis	7
1.2.1 Compact Object Mergers: Population Astrophysics and Statistics (COMPAS)	8
1.3 Thesis outline	9
2 Compact Object Mergers: Population Astrophysics and Statistics (COMPAS)	10
3 Chemically Homogeneous Evolution: A rapid population synthesis approach	46
4 Surrogate forward models for population inference on compact binary mergers	62
5 Summary and discussion	83
Bibliography	88

Chapter 1

Introduction

Rapidly growing catalogues of compact binary mergers from advanced gravitational-wave detectors allow us to explore the astrophysics of massive stellar binaries. Merger observations can constrain the uncertain parameters that describe the underlying processes in the evolution of stars and binary systems in population models. In this thesis, we demonstrate that properties of populations of merging compact binaries – namely, detection rates, chirp masses, and redshifts – can be used to measure cosmological parameters describing the redshift-dependent star formation rate and metallicity distribution. We present, as a proof-of-concept, a method that uses artificial neural networks (ANNs) to emulate a binary population synthesis computer model, and construct a fast, flexible, parallelisable surrogate model that is used for inference.

We use the COMPAS population synthesis computer model for this work [1]. A key part of the work done for this thesis was to enhance and extend COMPAS with the aim of releasing it publicly as an open-source population synthesis toolset (described in Chapter 2), and further work to extend the COMPAS model to include a new evolutionary track (chemically homogeneous evolution (CHE), described in Chapter 3).

In the remainder of this chapter we present an overview of some of the concepts and tools underpinning the work done for this thesis. An outline of the thesis is given in Section 1.3.

1.1 Astronomical compact objects

Astronomical compact objects are formed from the remnants of massive stars that have undergone a supernova explosion, leaving behind a collapsed core. These compact objects

are characterised by their extremely high density and small size. They include white dwarfs, neutron stars, and black holes.

White dwarfs are the least massive of the compact objects. They are formed from the cores of low-mass stars that have exhausted their nuclear fuel and shed their outer layers, leaving behind a hot, dense, core. White dwarfs typically have masses similar to that of the sun, but a radius that is only about the size of the Earth. They gradually cool over billions of years, eventually becoming dark and cold.

Neutron stars, formed when the core of a massive star collapses inwards, compressing protons and electrons together to form neutrons, are the most common type of compact object. This collapse to a neutron star results in an incredibly dense object, with a mass that is typically ~ 1.4 times that of the sun but a radius of only around 10km. Neutron stars have extremely strong magnetic fields, and initially spin rapidly, emitting intense beams of radiation that can be observed as pulsars. As a neutron star loses rotational energy it radiates less energy, and eventually is no longer considered a pulsar.

Black holes are formed when the core of a massive star collapses to a singularity - the singularity is no longer part of regular spacetime, and its properties can't be described with current physical theories. The extreme mass of the collapsed star causes such extreme curvature of spacetime that the singularity is surrounded by an event horizon from which nothing can escape. Black holes can have masses ranging from a few times that of the sun to billions of solar masses, and are believed to be present at the centre of most galaxies. Black holes can't be seen directly (optically), but can be observed indirectly through their effects on surrounding matter. When matter falls into a black hole, it is accelerated to very high speeds, producing intense radiation in the form of X-rays and gamma rays, and these emissions can be detected by telescopes and used to study the properties of black holes.

Astronomical compact objects play a crucial role in our understanding of the Universe. They are natural laboratories for testing the laws of physics under extreme conditions, such as strong gravity, high magnetic fields, and intense radiation. They also provide clues about the evolution of stars and galaxies, the formation of elements, and the origin of cosmic rays. Studying compact objects requires a range of observational techniques, including X-ray, gamma-ray, optical, radio, and gravitational-wave astronomy, as well as theoretical modelling and simulations.

1.1.1 Compact binaries

Astronomical compact binaries refer to pairs of stars that orbit around their common centre of mass, where at least one of the stars is a compact object such as a white dwarf, a neutron star, or a black hole. These systems can provide valuable insights into the nature of compact objects, the processes of stellar evolution, and the physics of extreme environments.

Compact binaries can be observed through a variety of observational techniques, such as detecting their X-ray and gamma-ray emissions, measuring their gravitational wave signals, and studying their optical and radio emissions. Different types of compact binaries exhibit distinct observational properties. For example, X-ray binaries are characterised by the presence of an accretion disk around the compact object, which emits X-rays as the infalling material heats up and releases energy.

Binary pulsars, which consist of a pulsar and a companion star (often a white dwarf or neutron star) in a close orbit, are an important, and much-studied, type of compact binary. Binary pulsars emit regular pulses of radio waves as the pulsar rotates, allowing astronomers to accurately measure the properties of the system, such as the masses of the two stars, and the eccentricity of their orbit. Binary pulsars are also important for testing theories of gravity, as they provide a strong testbed for the predictions of general relativity.

The study of compact binaries helps further our understanding of a wide range of astrophysical phenomena, from the formation and evolution of stars to the behaviour of matter under extreme conditions. Ongoing and future observations of these systems will continue to provide new insights into the nature of the Universe and its constituents.

The formation of compact binaries typically occurs through two main channels (see, e.g., [2, 3], for reviews):

- isolated binary evolution, in which two stars may interact through tides and mass transfer, but are dynamically decoupled from other stars (e.g., [4–6]). This often results in common envelope evolution, where a more massive star expands and transfers material to its companion, causing the two stars to spiral closer together.
- dynamical interactions in dense stellar environments, such as globular clusters, where three or more stars interact gravitationally to form a compact binary (e.g., [7–13]).

1.1.1.1 Isolated binary evolution

We know from gravitational wave observations that compact binaries exist with orbital separations of just a few tens of solar radii. To be the source of gravitational waves detected at advanced Laser Interferometer Gravitational-wave Observatory (aLIGO), compact objects orbiting each other must spiral in as they lose energy through the emission of gravitational waves. Orbital energy loss through gravitational wave emission is not efficient at wide separations, and the timescale for gravitational wave emission to drive a binary to merger scales as the fourth power of the orbital separation [14]. In order for two (say) $30 M_{\odot}$ black holes to merge within the age of the Universe, their initial separation must be below $\lesssim 50 R_{\odot}$. This challenges our understanding of binary star evolution.

During main-sequence evolution, typical slowly rotating massive stars can reach a radial extent of several thousand solar radii - much greater than the orbital separation required for a binary black hole (BBH) to merge within the age of the Universe. Close binary stars (those that should be able to merge within the age of the Universe by spiralling-in due to orbital energy loss through gravitational wave emission) that contain one or more stars that expand to even hundreds of solar radii should merge before they become compact binaries. One possible solution is CHE, which we address in Chapter 3, but we also know that a close binary star undergoes several physical processes during its life [15], possibly including mass transfer, Roche lobe overflow, common envelope evolution, tides, and supernovae with accompanying natal kicks [16], any of which can completely change the final fate of the system.

1.1.1.2 Binary evolution through dynamical interactions

In dense stellar environments, such as globular clusters, the orbits of stars, and multiple star systems, are very often perturbed by dynamical encounters with other stars or multiple star systems. These interactions affect the formation and the evolution of compact binaries such as BBHs in multiple ways [17].

If a binary star interacts with a third star in a three-body interaction - a common interaction in a dense cluster - the binary star will exchange some fraction of its internal energy with the third star. If the orbital separation of the binary is small, such an encounter will increase the kinetic energy of the ejected star at the expense of the orbital energy of the binary, and reduce the orbital separation even further, which, in the case of a BBH might drive the separation of the BBH into the regime where orbital decay by gravitational waves becomes efficient, and so hasten the merger [3].

Due to the short dynamical friction timescale, mergers of massive stars are common in dense young star clusters [18]. Such mergers can lead to the formation of black holes with mass in the pair-instability mass gap [19], which can then acquire a companion by dynamical exchanges, leading to the formation of BBHs in the pair-instability mass gap.

1.1.1.3 Merging compact binaries

When two compact objects, such as neutron stars or black holes, orbit each other, the emission of gravitational waves due to their acceleration causes them to lose energy and spiral inward. As they get closer and the curvature of spacetime becomes tighter, they orbit faster and emit more gravitational waves - this continues until they eventually merge. During the final stages of the merger, the objects are moving at a significant fraction of the speed of light, and the gravitational wave amplitude reaches its peak. As the objects merge, they release a burst of energy in the form of gravitational waves, creating ripples in the curvature of spacetime that can be detected by gravitational wave observatories on Earth.

The predictions of Michell [20] and Laplace [21, 22] notwithstanding, the first hints that black holes (as we now understand them) may exist emanate from the work of Karl Schwarzschild in 1916 [23, 24]. Schwarzschild's solution to the Einstein field equations [25] includes a parameter - the Schwarzschild radius - that defines the event horizon of a massive, non-rotating, spherical object with no charge (what we now know as the Schwarzschild black hole). The Schwarzschild radius is a characteristic radius associated with any object that has mass, but for objects much less massive than black holes it is negligible. The event horizon of a black hole is loosely considered to be its "surface".

In the hundred years since Schwarzschild's work, black holes proved to be elusive - while there were hints and circumstantial evidence, there was no unequivocal evidence confirming the existence of black holes. Then, on February 11th, 2016, the detection of a gravitational wave signature of two merging stellar-mass black holes of around 30 solar masses each, occurring about 1.3 billion light-years away, was announced by the LIGO Scientific Collaboration and Virgo Collaboration [26]. The detection, known as GW150914, was made on September 14th, 2015, just over a hundred years after Einstein's field equations were first published, and just short of the 100-year anniversary of the publication of Schwarzschild's work.

1.1.1.4 Gravitational wave detections

Since that first detection, the field of gravitational-wave astronomy has grown rapidly. The LVK collaboration has undertaken three observing runs since 2015 (runs O1, O2, and O3, with O3 being in two parts: O3a and O3b). Observing run O1 ran from September 12, 2015 to January 19, 2016; O2 from November 30, 2016 to August 25, 2017; O3a from April 1, 2019 to October 1, 2019; and O3b from November 1, 2019 to March 27, 2020. Observing run O4 is currently underway, having begun on May 24, 2023.

To date the LVK collaboration has, through the release of Gravitational-wave Transient Catalogs (GWTCs), reported ninety events with an astrophysical probability > 0.5 , with eighty-five merging BBHs, two double neutron stars (DNSs), and three black hole-neutron stars (BHNSs) detected [27–30]. Planned upgrades in the near future to aLIGO, including LIGO A+ [31] and LIGO Voyager [32], are expected to enable the detection of many more events, while third generation gravitational wave (GW) detectors such as the Einstein Telescope [33, 34] and the Cosmic Explorer [35, 36] are expected to detect several tens of thousands of BBH and DNS mergers per year.

In 2017 August, both aLIGO and Virgo detected the DNS merger GW170817 [37], coincident with the detection of the gamma ray burst (GRB) 170817A [38], and that, with subsequent detections of the resulting kilonova and afterglow [39, 40], was the first known event to enable GW multimessenger astronomy. GW170817, and the later BBH merger detected during LVK observing run O3, GW190521, which was identified as potentially having an electromagnetic (EM) counterpart, are the only GW events thought to have EM counterparts [41].

We have learned much from the detection of gravitational waves, and will continue to learn much more as we observe more. One of the most significant things to emerge from the detection of gravitational waves is the proven existence of black holes. While we have postulated for a very long time that black holes must exist, and we knew there was something very massive at the centre of many galaxies (ours amongst them), until the first gravitational wave detection in 2015 we had never directly observed a black hole. Now we have - GW150914, and the many detections since, proved that black holes really do exist. Moreover, we also now know for sure that black holes can, and do, orbit one another, and merge into single, larger, black holes.

Another very significant thing we learned from a gravitational-wave detection is what short-duration GRBs are. The DNS merger GW170817 proved that when neutron stars merge, a beam of gamma radiation is ejected - as was detected by the Fermi space telescope mere seconds after the gravitational wave from the merger was detected by

aLIGO. Such gamma ray bursts had been seen in the past, and it was postulated that they were the result of merging neutron stars, but until aLIGO detected GW170817, and the Fermi space telescope detected GRB 170817A, we had no direct proof.

1.2 Population synthesis

Population synthesis is a powerful technique used in astrophysics to model the properties, evolution, and interaction of populations of astronomical objects such as isolated stars, binary stars, multiple star systems (typically triples, but higher-order systems are possible), star clusters, galaxies, and supernovae. The technique involves combining the laws of physics and astrophysics with statistical methods to simulate the evolution and interaction of a population of objects over time. By generating synthetic populations, astronomers can explore the properties of these objects, such as their size (radius and mass), luminosity, colour, age, and chemical composition, etc., and compare them to observations, including gravitational wave observations from the LVK collaboration [31].

Population synthesis models are used to answer a wide range of questions in astrophysics, including how stars and galaxies form, how they evolve over time, and how they interact with their environment. The models take into account the physical processes that drive the evolution of astronomical objects, such as nuclear fusion, mass loss through winds, and interactions between objects in binary and higher-order systems, such as the effect of tides, mass transfer, common envelope evolution, mergers, etc. Population synthesis models can also be used to predict the properties of astronomical objects that have not yet been observed, allowing astronomers to design observational campaigns to test their predictions. Population synthesis has become an essential tool for modern astrophysics and has led to many important discoveries in the field.

By simulating the formation and evolution of a large population of stars, astronomers can study how the properties of those stars change over time. Moreover, population synthesis models can be used to study the properties of stars at different stages of their evolution, such as main-sequence stars, red giants, white dwarfs, including their (possible) explosive deaths as supernovae, and the remnants they leave (white dwarfs, neutron stars, black holes).

In recent years, the development of advanced computer simulations and observational data have made population synthesis models even more powerful. With these tools, astronomers can refine their models to make more accurate predictions and test them against increasingly detailed observations. Population synthesis models continue to be a

critical tool for understanding the properties and evolution of astronomical objects and for advancing our knowledge of the Universe.

1.2.1 Compact Object Mergers: Population Astrophysics and Statistics (COMPAS)

COMPAS is an open-source integrated suite of software tools combining a robust, rapid, and flexible population synthesis application for both single star evolution and binary star evolution, with tools for deployment on a range of platforms, including high-performance computing platforms, and a set of Python file management, population analysis, and plotting scripts.

The heart of the COMPAS suite is the single star evolution (SSE) and binary star evolution (BSE) simulation code, developed in the C++ programming language [42]. C++ is a cross-platform, object-oriented programming language that gives a clear structure to programs and allows code to be more easily understood and maintained - it is one of the world's most popular programming languages, and has an established history of being used to create high-performance applications. The COMPAS C++ code is modular, object-oriented code, designed to be easily understood and extended.

COMPAS version 1 was developed from 2015 to 2019 [43–49], and follows the Hurley SSE and BSE algorithms [16, 50], and the StarTrack code [51, 52]. The use of polynomial fitting formulae for stellar attributes, as a function of only the initial mass, metallicity, and age of a star, allows population synthesis applications based on the Hurley SSE and BSE algorithms to evolve systems rapidly, and with as much detail as possible without using the detailed evolutionary algorithms used in applications such as Modules for Experiments in Stellar Astrophysics (MESA) [53–55].

Other population synthesis tools exist, each with their own *raison d'être*. These include applications such as `binary_c` [56–58], Scenario Machine [59, 60], SeBa [61, 62], COSMIC [63], Mobse [64], and BPASS [65]. COMPAS was designed to be fast, allowing users to rapidly simulate large populations of either single stars or binary stars, and flexible, providing users with the means to easily configure the many tunable attributes of the stars and binaries simulated. For a comparison of the data products for some of the population synthesis tools listed here, interested readers should monitor the currently-underway Synthetic UCB (Ultra-compact binaries) Catalogs Project being conducted by the LISA Science Group [66].

The Hurley SSE and BSE algorithms used polynomial fits to stellar tracks from Pols et al. [67]. Recently a number of new population synthesis applications have been developed

that integrate stellar evolution by interpolation over a grid of stellar evolution tracks rather than using polynomial fits to a single set of tracks. These include ComBinE [68], METISSE [69, 70], SEVN [71], and POSYDON [72]. The benefit of this approach is that polynomial fits to new and updated stellar tracks are not required, however the quality of interpolation depends heavily on the density and completeness of the input tracks. The performance and memory requirements of the approach are proportional to the number of tracks used for interpolation - as the number of tracks increases, memory consumption rises and execution times increase.

1.3 Thesis outline

The remaining sections of the thesis are organised as follows:

- Chapter 2 presents a description of the development and public release of COMPAS version 2.
- Chapter 3 describes the extension of COMPAS version 2 to include the CHE track.
- Chapter 4 describes the construction and validation of the surrogate model.
- Chapter 5 concludes with a summary of the work presented in the thesis, and some remarks on possible future work.

Chapter 2

Compact Object Mergers: Population Astrophysics and Statistics (COMPAS)

As noted earlier, COMPAS version 1 was developed from 2015 to 2019 [43–49], and was an effective tool for simulating populations of binary stars, but while an object-oriented architecture was conceived of and the code was written in C++, over time through the iterative nature of research, with several contributors to the C++ code, the code base drifted somewhat from the object-oriented programming paradigm, and became difficult to understand (at least for users new to the code), and not easily extensible and maintainable.

We undertook the task of refactoring the existing COMPAS C++ code (COMPAS version 1) into a well-structured, object-oriented, and modular application that would be more easily extended and maintained, with a view to making it more accessible to a broader user-base by releasing it publicly as open-source code. COMPAS version 2 is the result.

Source code refactoring is the process of restructuring existing source code without changing its external behaviour. The motivation to refactor existing code is to improve the design, structure, and/or implementation of the software, while preserving its functionality. Typically, refactoring applies a series of small changes to the source code that either preserve the behaviour of the software, or at least do not modify its conformance to functional requirements. In this case, however, the entire code-base was redesigned and rewritten, while preserving the astrophysical algorithms that underpin the simulation of single star and binary evolution.

The architecture of the application now conforms to the object-oriented programming paradigm, making it much easier to extend and maintain. A new, comprehensive logging subsystem was added, as was flexible and extensible functionality for users to specify a grid of systems to evolve (specifying initial attributes of each system, rather than just a population of randomly sampled systems). Code versioning, change logs, and documentation in the form of an online user and developer guide were added. COMPAS was made available as open-source software in March 2020 [73], and has since developed an active, rapidly growing, user and developer base, and through them COMPAS continues to improve and expand.

The following two papers, Riley et al. [1, 74], describe the COMPAS version 2 code and the physics implemented therein. COMPAS version 2 is the product of the entire COMPAS collaboration over many years.

More information about COMPAS is available on the [COMPAS web page](#)¹.

The COMPAS code is available in the [COMPAS public repository](#)² on github [75].

[COMPAS documentation](#) is available online³.

¹<https://compas.science/index.html>

²<https://github.com/TeamCOMPAS/COMPAS>

³<https://compas.readthedocs.io/en/latest/>



Rapid Stellar and Binary Population Synthesis with COMPAS

Team COMPAS¹⁶,Jeff Riley^{1,2}, Poojan Agrawal^{2,3}, Jim W. Barrett⁴, Kristan N. K. Boyett⁵, Floor S. Broekgaarden⁶, Debatri Chattopadhyay^{2,3}, Sebastian M. Gaebel⁷, Fabian Gittins⁸, Ryosuke Hirai^{1,2}, George Howitt^{2,9}, Stephen Justham^{10,11,12}, Lokesh Khandelwal¹⁰, Floris Kummer¹⁰, Mike Y. M. Lau^{1,2}, Ilya Mandel^{1,2,4}, Selma E. de Mink^{6,10,12}, Coenraad Neijssel^{2,4}, Tim Riley^{1,2}, Lieke van Son^{6,10,12}, Simon Stevenson^{2,3},Alejandro Vigna-Gómez^{13,14}, Serena Vinciguerra¹⁰, Tom Wagg^{6,12,15}, and Reinhold Willcox^{1,2}¹ School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia; jeff.riley@monash.edu, ilya.mandel@monash.edu² OzGrav, Australian Research Council Centre of Excellence for Gravitational Wave Discovery, Australia; spstevenson@swin.edu.au³ Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia⁴ Institute of Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK⁵ Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK⁶ Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA; floor.broekgaarden@cfa.harvard.edu⁷ Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstrasse 38, D-30167 Hannover, Germany⁸ Mathematical Sciences and STAG Research Centre, University of Southampton, Southampton SO17 1BJ, UK⁹ School of Physics, University of Melbourne, Parkville, Victoria 3010, Australia¹⁰ Anton Pannekoek Institute of Astronomy and GRAPPA, Science Park 904, University of Amsterdam, 1098XH Amsterdam, The Netherlands¹¹ School of Astronomy & Space Science, University of the Chinese Academy of Sciences, Beijing 100012, People's Republic of China¹² Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany¹³ DARK, Niels Bohr Institute, University of Copenhagen, Jagtvej 128, DK-2200 Copenhagen, Denmark¹⁴ Niels Bohr International Academy, The Niels Bohr Institute, Blegdamsvej 17, DK-2100 Copenhagen, Denmark¹⁵ Department of Astronomy, University of Washington, Seattle, WA 98195, USA

Received 2021 October 1; revised 2021 November 24; accepted 2021 December 7; published 2022 February 2

Abstract

Compact Object Mergers: Population Astrophysics and Statistics (COMPAS; <https://compas.science>) is a public rapid binary population synthesis code. COMPAS generates populations of isolated stellar binaries under a set of parameterized assumptions in order to allow comparisons against observational data sets, such as those coming from gravitational-wave observations of merging compact remnants. It includes a number of tools for population processing in addition to the core binary evolution components. COMPAS is publicly available via the GitHub repository <https://github.com/TeamCOMPAS/COMPAS/>, and is designed to allow for flexible modifications as evolutionary models improve. This paper describes the methodology and implementation of COMPAS. It is a living document that will be updated as new features are added to COMPAS; the current document describes COMPAS v02.21.00.

Unified Astronomy Thesaurus concepts: [Stellar populations \(1622\)](#); [Binary stars \(154\)](#); [Neutron stars \(1108\)](#); [Black holes \(162\)](#)

Software reviewed by the [Journal of Open Source Software](#) JOSS

1. Introduction

The majority of massive stars are born in a stellar binary- or multiple-star system with other stellar companions (e.g., Chini et al. 2012; Sana et al. 2012, 2014; Kobulnicky et al. 2014; Almeida et al. 2017; Moe & Di Stefano 2017). The subsequent evolution of massive stellar binaries plays a critical role in many fields of astronomy. Massive binaries are thought to play key roles in the reionization of the universe (e.g., Conroy & Kratter 2012; Ma et al. 2016b; Eldridge et al. 2017; Rosdahl et al. 2018; Göteborg et al. 2019, 2020), nucleosynthesis (e.g., Dray & Tout 2003; Izzard et al. 2006; Woosley & Heger 2007; Langer 2012), and the diversity of observed supernovae (e.g., Podsiadlowski et al. 1992; Eldridge et al. 2013, 2018, 2019a;

Tauris et al. 2015; Moriya et al. 2017; Yoon et al. 2017, 2019; Zapartas et al. 2019, 2021). Some massive binaries will evolve into systems containing one or two compact objects, which can be observed as X-ray binaries (e.g., Remillard & McClintock 2006), double neutron stars (e.g., Tauris et al. 2017), short gamma-ray bursts (e.g., Woosley & Bloom 2006; Berger 2014), and gravitational-wave transients (e.g., Mandel & Farmer 2018; Abbott et al. 2019; Mandel & Broekgaarden 2021).

Many physical processes in the evolution of a binary system are uncertain. The uncertainties in stellar wind mass loss, mass transfer, common envelope physics, supernova remnant masses, and natal kicks, among others, can be constrained by comparing the observed populations listed above against theoretical predictions under a range of assumptions. The essence of binary population synthesis (BPS) simulations is to enable modeling of large populations by combining prescriptions for single-star evolution (SSE) and binary star evolution (BSE), modeling the evolution of stars from zero-age main sequence (ZAMS) until their final states.

In this paper we present the publicly available BPS suite Compact Object Mergers: Population Astrophysics and Statistics (COMPAS). The core of COMPAS is a BPS code that models

¹⁶ The public COMPAS code is a product of work by the entire COMPAS collaboration over many years; we therefore kindly request that, in recognition of this team effort, the paper is cited as “Team COMPAS: J. Riley et al.”

the evolution of a population of binary stars using a set of simplified prescriptions or recipes for stellar and binary evolution. By doing so, COMPAS can compute the full evolution of a typical binary system in around 10 ms on a modern laptop, and compute the evolution of a million binaries in a few CPU hours. The general approach is similar to other BPS codes including the Scenario Machine (Lipunov et al. 1996a, 1996b, 2009), IBiS (Tutukov & Yungelson 1996), SeBa (Portegies Zwart & Verbunt 1996; Portegies Zwart & Yungelson 1998; Nelemans et al. 2001; Toonen et al. 2012), BSE (Hurley et al. 2002), StarTrack (Belczynski et al. 2002, 2008, 2020b), `binary_c` (Izzard et al. 2004, 2006, 2009), MOBSE (Giacobbo & Mapelli 2018; Giacobbo et al. 2018), and COSMIC (Breivik et al. 2020).

In addition to the core BPS code, COMPAS also provides several other publicly available tools, including rapid single stellar evolution, postprocessing tools to study the evolution of populations over cosmic time (Neijssel et al. 2019), post-processing scripts to model the detectability of double compact objects, or DCOs [comprising double neutron stars (DNS), binary black holes (BBH), and black hole–neutron star (BH–NS) mergers] by ground-based gravitational-wave observatories (e.g., Barrett et al. 2018), a statistical sampling framework to optimize the computational cost of BPS (Broekgaarden et al. 2019), and models for specific evolutionary phases such as X-ray binaries and pulsars (e.g., Chattopadhyay et al. 2020; Vinciguerra et al. 2020).

COMPAS was developed with a primary focus on the study of compact object mergers that serve as sources of gravitational waves. It has been used extensively to investigate the properties of compact binaries containing neutron stars and black holes. Stevenson et al. (2017) studied the formation history of the first three BBHs detected via gravitational waves. Barrett et al. (2018) explored how future gravitational-wave observations will allow us to determine the physics of massive binary evolution. Stevenson et al. (2019) included the impact of (pulsational) pair-instability supernovae. Neijssel et al. (2019) investigated the consequences of uncertain metallicity-specific star formation history on the rate and properties of DCO mergers. Lau et al. (2020) predicted the number of DNSs detectable with LISA and the inference these observations will enable. Bavera et al. (2020) used COMPAS to predict the spin distribution of merging BBHs. van Son et al. (2020) investigated the robustness of the predicted pair-instability mass gap to uncertainties in the accretion efficiency for black holes. Riley et al. (2021) studied chemically homogeneous evolution as a pathway to BBH formation. Mandel et al. (2021) considered the consequences of an alternative, stochastic recipe for compact remnant masses. Vigna-Gómez et al. (2021) explored sequential BBH mergers during triple evolution. Broekgaarden et al. (2021a) and Broekgaarden & Berger (2021) focused on the formation of black hole–neutron star binaries. van Son et al. (2021) investigated the relative contributions of dynamically stable and unstable mass transfer to BBH mergers over cosmic history. Broekgaarden et al. (2021b) investigated the relative impact from uncertain metallicity-specific star formation history to that from uncertain stellar evolution on the rate and properties of DCO mergers.

A number of other consequences of massive binary evolution have been explored with COMPAS. Vigna-Gómez et al. (2018) used observations of Galactic DNSs to constrain evolutionary physics. Chattopadhyay et al. (2020) used the same population to constrain birth distributions of pulsar spin periods and magnetic

field strengths, as well as magnetic field decay scales. Schröder et al. (2020) investigated optical counterparts to common envelope (CE) events which fail to eject the envelope and result in a merger. Howitt et al. (2020) explored luminous red novae: red optical transients associated with CE events. Vigna-Gómez et al. (2020) cataloged the CE events en route to DNS formation. Mandel & Müller (2020) took advantage of COMPAS SSE modules in formulating a stochastic compact remnant mass prescription and momentum-conserving natal kick prescription. Vinciguerra et al. (2020) studied the population of Be X-ray binaries to establish constraints for mass accretion efficiency. Chattopadhyay et al. (2021) explored the prospects for observing neutron star–black hole binaries in future radio pulsar surveys. Miller-Jones et al. (2021) and Neijssel et al. (2021) constrained the massive stellar wind mass loss with observations of the black hole (BH) high-mass X-ray binary Cygnus X-1. Willcox et al. (2021) compared COMPAS models to observed pulsar velocities in order to study neutron star natal kicks.

COMPAS has also played an important role in developing more efficient sampling, inference, and model emulation techniques. Barrett et al. (2017) attempted to generate computationally efficient surrogate models of binary population synthesis through Gaussian process emulation. Taylor & Gerosa (2018) used COMPAS public data as an example in their hierarchical inference study. Broekgaarden et al. (2019) applied importance sampling to enable computationally efficient DCO simulations. Lin et al. (2021) developed new tools for classification and emulation based on (local) Gaussian process models and the COMPAS suite.

In this paper we introduce the COMPAS code and discuss its methodology and implementation. The paper is organized as follows. Section 2 presents the main features of COMPAS. Section 3 discusses the implementation of single stellar evolution in COMPAS. Section 4 discusses binary stellar evolution prescriptions. In Section 5 we describe how COMPAS can be used to model populations of binaries. Section 6 describes the postprocessing tools available within COMPAS. Section 7 illustrates a few applications of COMPAS. We conclude in Section 8.

2. COMPAS

2.1. Overview

COMPAS¹⁷ is an open-source integrated suite of software tools combining a robust, rapid, and flexible population synthesis application for both single-star evolution and binary star evolution, with tools for deployment on a range of platforms, including high-performance computing platforms, and a set of Python postprocessing analysis and plotting scripts.

The heart of the COMPAS suite is the SSE and BSE simulation code, developed in the C++ programming language. C++ is a cross platform, object-oriented programming (OOP) language that gives a clear structure to programs and allows code to be easily understood and maintained—it is one of the world’s most popular programming languages, and has an established history of being used to create high-performance applications. The COMPAS C++ code is a modular, object-oriented code, designed to be easily understood and extended.

Results produced with COMPAS are publicly available.¹⁸

¹⁷ compas.science. Code available at <https://github.com/TeamCOMPAS/COMPAS>.

¹⁸ <https://zenodo.org/communities/compas/>

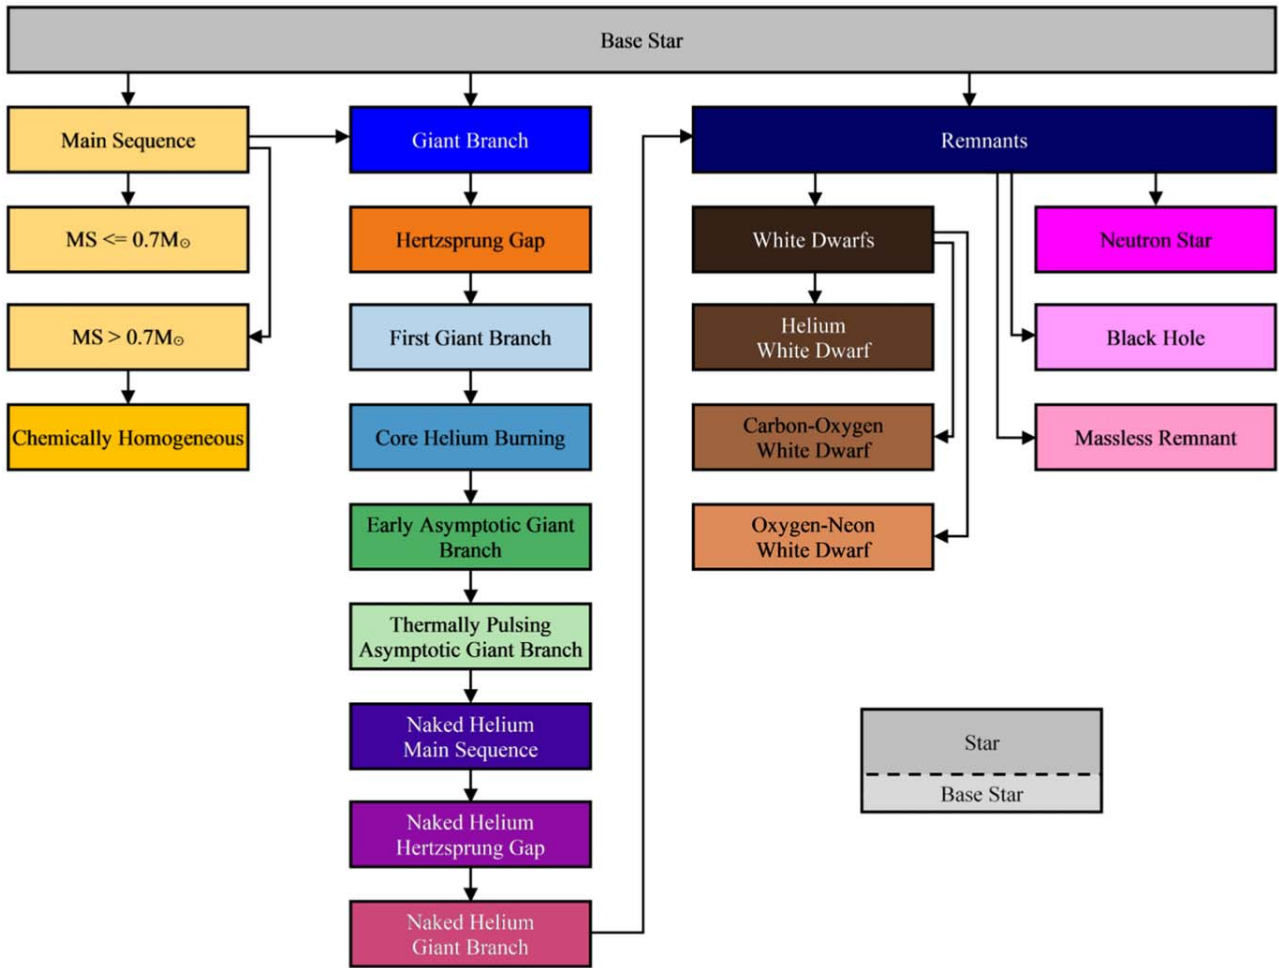


Figure 1. SSE class and container diagram (arrows indicate inheritance).

2.2. Mode: Single Star versus Binary Star

COMPAS operates in either one of two selectable modes: SSE or BSE. As the names suggest, in SSE mode COMPAS evolves single stars, and in BSE mode COMPAS evolves binary stars.

The SSE algorithm covers all evolution phases from the ZAMS up to and including the remnant stages. The allowed range of ZAMS masses, $M \in [0.1, 150] M_{\odot}$, extends the range of Hurley et al. (2000) models by extrapolation; the allowed range of metallicities $Z \in [10^{-4}, 0.03]$ follows these models (see Section 3 for a full description).

In BSE mode, the SSE code provides the stellar attributes (e.g., luminosity, radius, temperature, etc.) for each of the component stars as they evolve. COMPAS evolves binary stars until a DCO is formed, the component stars merge, or, optionally, the binary is disrupted.

In either mode (SSE or BSE), users can specify a maximum evolution time and/or maximum number of evolutionary steps, after which evolution is halted.

2.3. Architecture

The SSE and BSE simulation code at the heart of COMPAS is written in C++ using OOP concepts. The architecture is based on stellar type, with each stellar type being described by a separate C++ class. Figure 1 shows the SSE class and

container diagram, where the arrows indicate inheritance (the OOP mechanism used to base one class upon another: the inheriting class inherits the implementation of the inherited class). The COMPAS C++ code is implemented using multiple inheritance, and all stellar classes also inherit directly from the *BaseStar* class (arrows not shown in Figure 1 for clarity). Each of the stellar classes encapsulates data structures and algorithms specific to the evolutionary phase corresponding to the class.¹⁹

The *Star* class shown in Figure 1 is a container class for the stellar classes. An instance of the *Star* class is a single star being evolved by COMPAS, and contains an object that is created as a *BaseStar* object, and evolves, over time, through various SSE classes shown in Figure 1.

Figure 2 shows the BSE class and container diagram. The main class for binary star evolution is the *BinaryStar* class. The *BinaryStar* class is a wrapper, containing a *BaseBinaryStar* class object, and abstracts away the details of the binary star and the evolution. The *BaseBinaryStar* class is a container class for the objects that represent the component stars of a binary system. An instance of the *BaseBinaryStar* class is a binary system being evolved by COMPAS, and contains a *BinaryConstituentStar* class object for each of the component stars

¹⁹ The class names shown in Figure 1 do not match the class names used in the COMPAS C++ code, which are abbreviated into single words or acronyms.




Figure 2. BSE class and container diagram.

(i.e., the primary and secondary stars), as well as data structures and algorithms specific to the evolution of a binary system. The *BinaryConstituentStar* class inherits from the SSE *Star* class, so objects instantiated from the *BinaryConstituentStar* class inherit the characteristics of the SSE *Star* class, particularly the stellar evolution model. The *BinaryConstituentStar* class defines additional data structures and algorithms (to the data structures and algorithms provided by the SSE classes) required to support the evolution of a binary system component star.

2.4. Evolutionary Models

2.4.1. SSE Model

The SSE model implemented in COMPAS (see Section 3 for more details) follows Hurley et al. (2000), using their analytical fits to the models of Pols et al. (1998). After the creation of the star according to the initial conditions (specified or sampled), the evolution of a single star proceeds by integrating the attributes of the star over its lifetime, and stops when the star evolves to a remnant, or the maximum time, or maximum number of time steps is reached.

Figure 3 shows a high-level overview of the code flow for SSE.

2.4.2. BSE Model

The binary evolution model implemented in COMPAS is broadly similar to the BSE population synthesis application (Hurley et al. 2002), and other population synthesis applications derived from it, such as *binary_c* (Izzard et al. 2004, 2006, 2009; De Minkde Mink et al. 2013) and *StarTrack* (Belczynski et al. 2002, 2008).

After the creation of the binary system according to the initial conditions (specified or sampled), the evolution of a binary system proceeds by integrating the attributes of the system over its lifetime and stops if the component stars merge, when the system forms a DCO, is disrupted, or the maximum time or maximum number of time steps is reached.

Figure 4 shows a high-level overview of the code flow for BSE.

2.5. Time Stepping

The initial estimate of the time step used in SSE follows Hurley et al. (2000), where the time step varies depending upon the evolutionary phase of the star (see Section 3 for a full list of the stellar types). COMPAS then checks whether the time step produces excessive change, defined as either

1. mass loss greater than 1%, or
2. radial change greater than 10%

over the time step, and limits the time step accordingly.

For nuclear timescale evolution, we limit the time step to a minimum of 100 yr, and we impose an overall minimum time

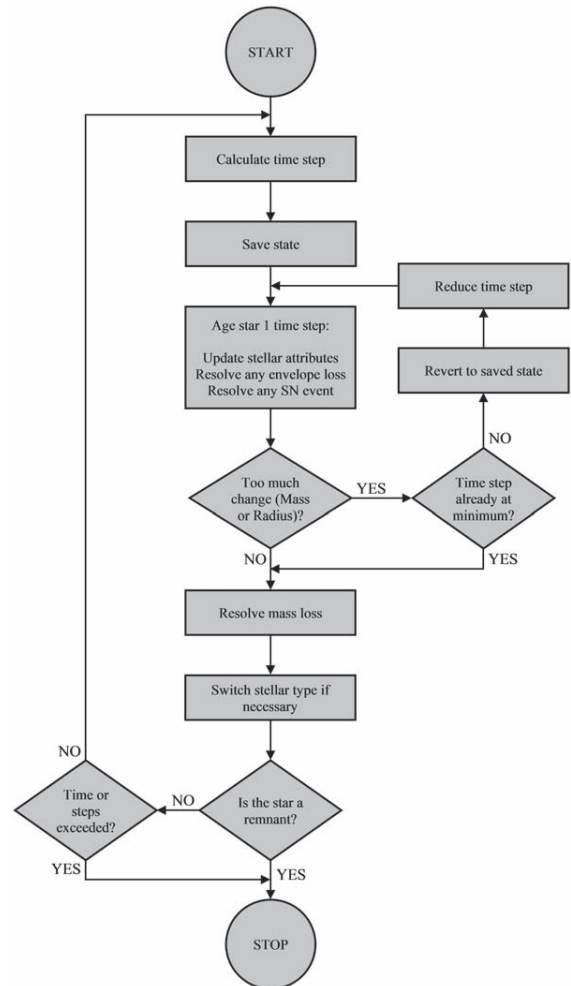


Figure 3. High-level SSE evolution.

step of 100 s, including for dynamical timescale evolution (see Section 3.3).

For the BSE time step, COMPAS uses the minimum of the binary constituent stars' SSE time steps: this allows the constituent stars to evolve using time steps that do not produce excessive change. Changes in binary properties are not separately considered when calculating the time step since large changes in binary properties would be accompanied by similarly large changes in constituent star properties.

COMPAS provides a mechanism for the user to scale the calculated time step by a positive scaling factor. Scaling is performed prior to limiting of the time step.

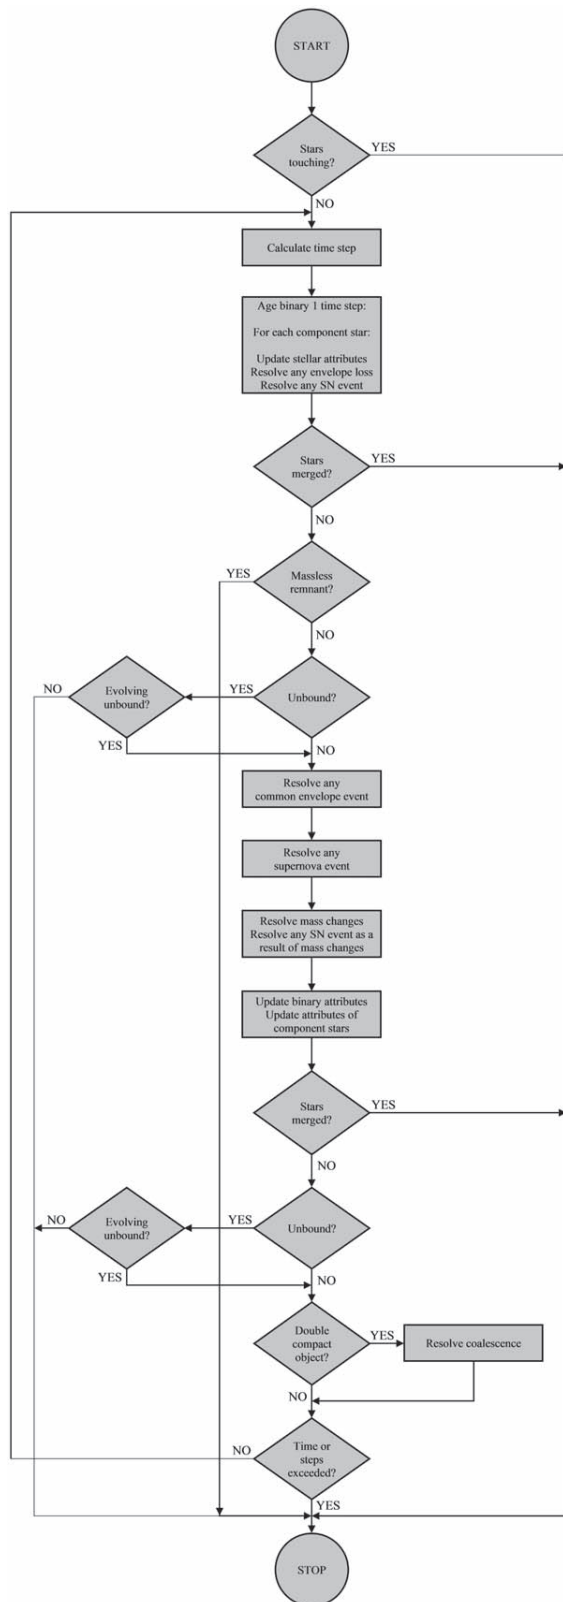


Figure 4. High-level BSE evolution.

2.6. Input and Configuration

COMPAS provides wide-ranging functionality and affords users much flexibility in determining how the synthesis and evolution of stars (single or binary) is conducted. Users

configure COMPAS’S functionality and provide initial conditions via the use of program options and input grid files. The full list and description of program options and grid files can be found in the COMPAS online documentation.²⁰

2.7. Default Model

Investigating massive star and binary evolution requires modeling many complicated astrophysical processes. In BPS codes, this is often done through the use of simple analytic prescriptions, which are calibrated either to theoretical predictions or to results from observations (see Sections 3 and 4 for more details in the context of SSE and BSE). In this paper, we refer to our default set of modeling assumptions, including which prescriptions are used, as our *Default* (or “default”) model. This model is summarized in Table 1. This model has been calibrated to match a range of observations (see Section 1 for a brief overview). There are, however, large uncertainties in the prescriptions used. COMPAS is flexible, and in many cases, we provide additional options to allow users to easily vary their choices from our defaults.

2.7.1. Program Options

COMPAS provides a rich set of configuration parameters via program options, allowing users to vary many parameters that define the initial attributes and/or affect the evolution of single and binary stars being evolved. Furthermore, COMPAS allows some parameters to be specified as ranges or sets of values via the program options, allowing users to specify a grid of parameter values on the command line. Combining command line program options, particularly ranges and sets, with a grid file allows users great flexibility in specifying more complex combinations of parameter values.

2.7.2. Grid Files

A grid file allows users to specify, in plain text, initial values and physics assumptions for multiple systems for both SSE and BSE. Each line of a grid file is used by COMPAS to set the initial values of an individual single star (SSE) or an individual binary system (BSE), and the physics assumptions to be used to evolve the star or system.

2.8. Output

COMPAS provides real-time status information during the evolution of systems. Detailed and summary information about the star or system being evolved is written to log files as the evolution proceeds.

A number of COMPAS log files may be produced depending upon the simulation type (SSE or BSE) and user specifications. These log files record, for each star or system being evolved:

1. summary information at the completion of evolution,
2. detailed information at each time step,
3. detailed information at the time of each stellar type switch,
4. summary information for all supernovae (SNe) events,
5. summary information for all CE events during BSE,
6. detailed information for all Roche-lobe overflow (RLOF) events during BSE,
7. summary information for all DCOs formed during BSE, and

²⁰ <https://compas.science/docs/>

Table 1
Initial Values and Default Settings for Binary Population Synthesis Simulation with COMPAS

Description and Name	Value/Range	Note/Setting
Initial conditions		
Initial mass $M_{1,i}$	[5, 150] M_{\odot}	Kroupa (2001) IMF $\propto M_{1,i}^{-\alpha_{\text{IMF}}}$ with $\alpha_{\text{IMF}} = 2.3$ for stars in this mass range
Initial mass ratio $q_i = M_{2,i}/M_{1,i}$	[0.01, 1]	We assume a flat mass ratio distribution $p(q_i) \propto 1$ with $M_{2,i} \geq 0.1 M_{\odot}$
Initial semimajor axis a_i	[0.01, 1000] au	Distributed flat-in-log $p(a_i) \propto 1/a_i$
Initial metallicity Z_i	[0.0001, 0.03]	Distributed flat-in-log $p(Z_i) \propto 1/Z_i$
Initial orbital eccentricity e_i	0	All binaries are assumed to be circular at birth
Fiducial parameter settings		
Chemically homogeneous evolution	Enabled	Following Riley et al. (2021), ‘‘pessimistic’’ version checking for threshold throughout evolution (Section 3.4).
Stellar winds for hydrogen-rich stars	Belczynski et al. (2010)	Based on Vink et al. (2000, 2001), including LBV wind mass loss with $f_{\text{LBV}} = 1.5$.
Stellar winds for hydrogen-poor helium stars	Belczynski et al. (2010)	Based on Hamann & Koesterke (1998) and Vink & de Koter (2005).
Mass-transfer stability criteria	ζ -prescription	Based on Vigna-Gómez et al. (2018) and references therein
Mass-transfer accretion rate	Thermal timescale	Limited by thermal timescale for stars Vigna-Gómez et al. (2018), Vinciguerra et al. (2020)
	Eddington-limited	Accretion rate is Eddington-limited for compact objects
Nonconservative mass loss	Isotropic re-emission	Massevitch & Yungelson (1975), Bhattacharya & van den Heuvel (1991) Soberman et al. (1997), Tauris & van den Heuvel (2006)
Case BB mass-transfer stability	Always stable	Based on Tauris et al. (2015, 2017), Vigna-Gómez et al. (2018) (Section 4.2)
Circularization at the onset of RLOF	On	Instantly circularized to periastris (Section 4.2)
CE prescription	$\alpha - \lambda$	Based on Webbink (1984), de Kool (1990)
CE efficiency α -parameter	1.0	(Section 4.2.4)
CE λ -parameter	λ_{Nanjing}	Based on Xu & Li (2010a, 2010b) and Dominik et al. (2012)
Hertzprung gap (HG) donor in CE	Pessimistic	Defined in Dominik et al. (2012): HG donors do not survive a CE phase
SN natal-kick magnitude for white dwarfs	0	We assume WDs do not receive natal kicks (Section 3.7.1)
SN natal-kick magnitude v_k for NS	$[0, \infty)$ km s $^{-1}$	Drawn from Maxwellian distribution with standard deviation $\sigma_{\text{rms}}^{\text{1D}}$
SN natal-kick magnitude for BH	$[0, \infty)$ km s $^{-1}$	Reduced relative to NS kicks by the fallback fraction (Fryer et al. 2012), see Section 3.9
SN natal-kick polar angle θ_k	$[0, \pi]$	$p(\theta_k) = \sin(\theta_k)/2$
SN natal-kick azimuthal angle ϕ_k	$[0, 2\pi]$	Uniform $p(\phi) = 1/(2\pi)$
SN mean anomaly of the orbit	$[0, 2\pi]$	Uniformly distributed
Core-collapse SN remnant mass prescription	Delayed	From Fryer et al. (2012), which has no lower BH mass gap
USSN remnant mass prescription	Delayed	From Fryer et al. (2012)
ECSN remnant mass prescription	$m_f = 1.26 M_{\odot}$	Based on Equation (8) in Timmes et al. (1996)
Core-collapse SN velocity dispersion $\sigma_{\text{rms}}^{\text{1D}}$	265 km s $^{-1}$	1D rms value based on Hobbs et al. (2005)
USSN and ECSN velocity dispersion $\sigma_{\text{rms}}^{\text{1D}}$	30 km s $^{-1}$	1D rms value based on e.g., Pfahl et al. (2002a), Podsiadlowski et al. (2004)
PISN/PPISN remnant mass prescription	Marchant et al. (2019)	As implemented in Stevenson et al. (2019)
Maximum NS mass	$\text{max}_{\text{NS}}=2.5 M_{\odot}$	Mass division between NS and BH (Section 3.7.2)
Tides and rotation		No tides and/or rotation except chemically homogeneous evolution
Simulation settings		
Binary fraction	$f_{\text{bin}} = 1$	Corrected factor to be consistent with, e.g., Sana (2017)
Solar metallicity Z_{\odot}	$Z_{\odot} = 0.0142$	Based on Asplund et al. (2009)

Note. COMPAS users who wish to provide a similar table accompanying their publication can find a template at <https://github.com/FloorBroekgaarden/templateForTableBPSsettings>.

8. detailed pulsar evolution information.

COMPAS log files are created and written as hierarchical data format, version 5 (HDF5) files,²¹ comma separated values (CSV) files, tab separated values (TSV) files, or plain text files, as specified by the user.

The COMPAS software suite includes a Python postprocessing script to combine all COMPAS output HDF5, CSV, or TSV files into a single HDF5 file, which is especially useful if a single large experiment is spread over several virtual machines.

3. Single Stellar Evolution

As stars evolve, they experience nuclear fusion while balancing gravity with pressure and radiating away excess energy.

²¹ <https://www.hdfgroup.org/>

Consequently, their composition, radius, temperature, and luminosity all change, they may lose mass in stellar winds, and sufficiently massive stars may explode in SNe at the end of their lives.

COMPAS currently relies on rapid algorithms that provide estimates for how various fundamental stellar properties, such as their radii and luminosities, change as a star evolves through different evolutionary phases. The algorithms that capture how stars evolve—the SSE library—are at the core of the COMPAS code. These routines govern the evolution of single stars, as the name suggests, but they are also used to capture how a star evolves under the external influence of a binary companion.

In this section we begin by giving an overview of the evolutionary algorithms implemented in COMPAS that govern the evolution of the main properties of stars (Section 3.1). We then discuss some of the key evolutionary stages that stars evolve through (Section 3.2), and highlight several important evolutionary timescales in stellar evolution (Section 3.3). We proceed to describe the effects of rapid rotation that are included in COMPAS (Section 3.4), and describe the prescriptions available in COMPAS for incorporating mass loss in stellar winds (Section 3.5). We end this section by discussing the properties of stellar remnants: white dwarfs (WDs), neutron stars (NSs), and BHs.

3.1. Evolutionary Algorithms

COMPAS computes the properties of a star (such as luminosity, radius, or core mass) as functions of a star’s mass, metallicity, and age using analytic formulae, fit to detailed stellar models. The ZAMS radius R_{ZAMS} and luminosity L_{ZAMS} are calculated as functions of mass and metallicity using the analytic formulae from Tout et al. (1996). During the evolution we use the formulae from Hurley et al. (2000), developed to match the detailed stellar models of Pols et al. (1998; see also Eggleton et al. 1989 and Tout et al. 1997 for the basis of this approach).

The Pols et al. (1998) stellar models are for nonrotating stars and span ZAMS masses between 0.1–50 M_{\odot} . The original models do not include any mass loss. We incorporate mass loss following Hurley et al. (2000) as described in Section 3.5. Since the Hurley et al. (2000) formulae are polynomials in M_{ZAMS} , they can easily be extrapolated to higher masses by evaluating them outside of this range. We find that the fits extrapolate smoothly to at least 150 M_{\odot} . We note that this approach is not ideal, but significantly improving upon this is far from trivial. The evolution of high-mass stars is still very uncertain (Maeder & Meynet 2000; Langer 2012; Agrawal et al. 2020; Bowman 2020; Belczynski et al. 2021). This is particularly true for the later, faster evolutionary phases, where observations are scarce, and for internal properties such as core masses, which we cannot probe directly. Although grids of detailed models exist, they vary widely in their predictions. Given the limitations in the accuracy of massive star models and the absence of robust rapid prescriptions that clearly provide a significant improvement, we opt for simply extrapolating the fits in the present version of COMPAS.

The models are applicable for metallicities between $Z = 10^{-4}$ and $Z = 0.03$. We *do not* extrapolate in metallicity, as we find that some of the fitting formulae are not well behaved outside of this range. We discuss the implications for population studies in Section 6.4.

3.2. Evolutionary Stages

To introduce some of the evolutionary stages captured by the SSE formulae, we briefly summarize the main evolutionary stages of single stars. Our extremely brief sketch of stellar evolution generally follows the characteristic behavior of massive stars, and should not be taken as fully general.

We follow the lives of stars from the ZAMS (we do not include pre-main-sequence evolution in COMPAS). Stars spend most of their lives on the main sequence, fusing hydrogen in their cores. Following Hurley et al. (2000), we distinguish between low-mass ($M_{\text{ZAMS}} \leq 0.7 M_{\odot}$) main-sequence stars (stellar type 0) which are expected to be fully convective, and more massive main-sequence stars ($M_{\text{ZAMS}} > 0.7 M_{\odot}$, stellar type 1). During main-sequence evolution, most stars increase their luminosity and decrease their effective temperature T_{eff} . Following the main sequence, sufficiently massive stars experience a rapid thermal timescale (see Section 3.3) phase of expansion, evolving to lower effective temperatures at near constant luminosity. This phase of evolution is sometimes known as the Hertzsprung gap (HG), as the short timescale leads to a paucity of stars observed in this phase. Stars then begin a longer-lived phase of core helium burning (CHeB), evolving onto and up the giant branch. We note that at low metallicity, massive stars may not reach the giant branch before beginning CHeB. Helium-shell-burning stars evolve along the asymptotic giant branch (AGB). We choose to use the name early asymptotic giant branch (EAGB) for H-rich massive stars with a C/O core, in addition to its usual meaning. Stars that lose their outer hydrogen envelopes, either through stellar winds or binary mass transfer, become helium main-sequence (HeMS) stars, which then evolve analogously to hydrogen-rich stars through to the helium giant branch (HeGB). Finally, most stars end their lives as a stellar remnant, either a WD, NS, or BH depending on their initial mass, while some stars leave behind no remnant.

We show the evolution of single stars of particular masses in the mass range 0.5–150 M_{\odot} at solar ($Z = 0.0142$, Asplund et al. 2009) and low ($Z = 0.001$) metallicity in the Hertzsprung–Russell diagram according to COMPAS in Figure 5, with the different evolutionary phases identified for each track.

A complete list of stellar phases is given in Table 2, following Hurley et al. (2000). See Figure 1 for a schematic of the corresponding classes in the COMPAS code.

Figure 6 shows the maximum radial extent of a star during each evolutionary phase for a star with ZAMS mass between 0.5–150 M_{\odot} at solar ($Z = 0.0142$) and low ($Z = 0.001$) metallicity. Phases during which a star expands are the phases when binary interactions are most likely to occur. We discuss binary interactions further in Section 4.

3.3. Evolutionary Timescales

Three key timescales in single and binary stellar evolution—dynamical, thermal, and nuclear—often create a very convenient timescale hierarchy. The separation of timescales allows many approximations to be made.

The shortest timescale is almost always the dynamical (or freefall) timescale, defined as

$$\tau_{\text{dyn}} = \sqrt{\frac{R^3}{GM}} \approx 1600 \text{ s} \left(\frac{R}{R_{\odot}} \right)^{3/2} \left(\frac{M}{M_{\odot}} \right)^{-1/2}, \quad (1)$$

where R is the radius of a star and M is its mass. The dynamical timescale is used as the minimum timescale in COMPAS (with a minimum cutoff of 100 s).

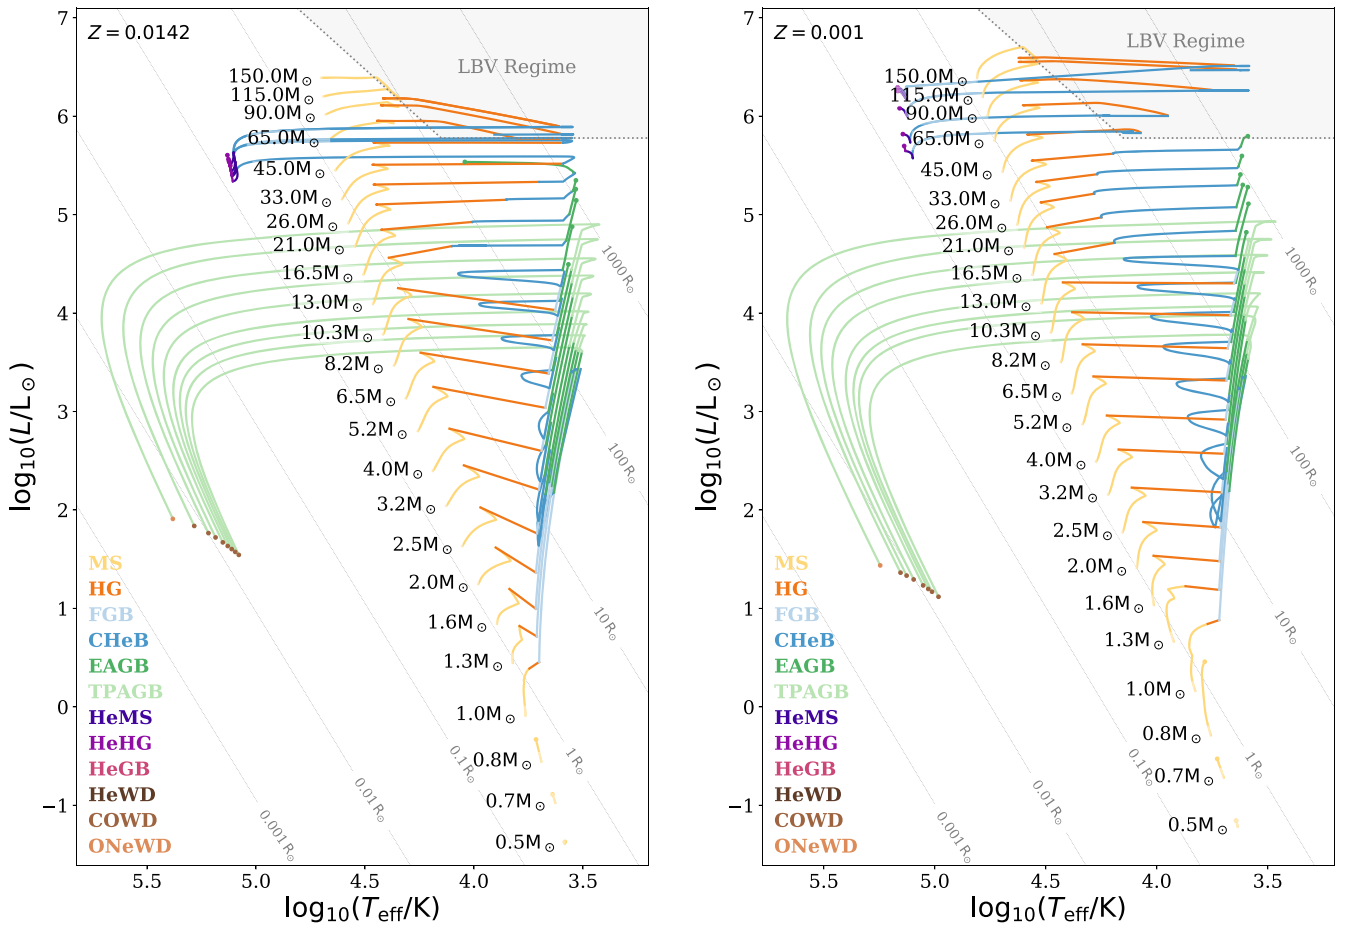


Figure 5. Evolutionary tracks in the Hertzsprung–Russell diagrams for single stars with ZAMS masses between 0.5–150 M_{\odot} at solar ($Z = 0.0142$) and low ($Z = 0.001$) metallicity using COMPAS default settings, limited to a maximum evolution time of 14 Gyr. The shaded region (bounded by the Humphreys–Davidson limit and a minimum luminosity) illustrates the regime in which luminous blue-variable (LBV) mass loss dominates. Diagonal contours show lines of constant radii.

The Kelvin–Helmholtz (or thermal) timescale is the time required for a star’s internal energy E_{int} to be radiated at its current luminosity L , and is given by $\tau_{\text{KH}} = E_{\text{int}}/L$. We estimate this as

$$\tau_{\text{KH}} \approx 3.0 \times 10^7 \text{ yr} \left(\frac{M}{M_{\odot}} \right) \left(\frac{M_{\text{s/env}}}{M_{\odot}} \right) \left(\frac{R}{R_{\odot}} \right)^{-1} \left(\frac{L}{L_{\odot}} \right)^{-1}, \quad (2)$$

where L is the luminosity of the star and $M_{\text{s/env}}$ is either the total mass of the star M for stellar types without a clearly defined envelope, or the mass of the envelope M_{env} for stars with a clearly defined envelope (see Hurley et al. 2002, for more details). This is used, e.g., when calculating thermal timescale mass transfer (see Section 4.2.3).

The nuclear timescale is relevant when nuclear fusion is setting the timescale at which the star evolves. This applies, e.g., to main-sequence stars that burn hydrogen in their centers, and stars that are undergoing central helium burning. The nuclear timescale is approximately

$$\tau_{\text{nuc}} \approx \phi f_{\text{nuc}} \frac{Mc^2}{L} \approx 10^{10} \text{ yr} \left(\frac{M}{M_{\odot}} \right) \left(\frac{L}{L_{\odot}} \right)^{-1}, \quad (3)$$

where c is the speed of light, ϕ is the efficiency which the rest mass of the relevant reacting nuclei is converted into energy

(for hydrogen, $\phi = 0.007$), and f_{nuc} is the fraction of the stellar mass that can serve as nuclear fuel. We do not use this timescale in any calculations; it is provided here, and optionally in COMPAS outputs, for reference only. Instead, we use the fitting formulae for stellar lifetimes in Hurley et al. (2000). The main-sequence lifetimes given by these fits are shown in Figure 7. At masses of $\gtrsim 50 M_{\odot}$ the Hurley et al. (2000) formulae, which extrapolate the Pols et al. (1998) model grid in this mass range, overestimate the stellar lifetimes by approximately 30% compared to the detailed stellar models considered in Agrawal et al. (2020).

3.4. Rotation

Most stars rotate sufficiently slowly (e.g., Dufton et al. 2013; Ramírez-Agudelo et al. 2013) for the impact of rotation to be modest. Rotation is expected to enhance mass-loss rates (e.g., Chiosi & Maeder 1986; Friend & Abbott 1986; see also Section 3.5), slightly increase main-sequence luminosities and lifetimes (e.g., Talon et al. 1997; Maeder & Meynet 2000), and lead to increased core masses (e.g., Maeder 1987; Langer 1992; Heger et al. 2000).

However, very rapid rotation, especially coupled with significant tidal effects in very close binaries, can have a dramatic impact on the evolution of a star. Sufficiently rapid rotation can lead to

enhanced mixing within a star, which may lead to chemically homogeneous evolution (Maeder 1987), where a star can burn almost all of its hydrogen into helium.

The stellar tracks used in COMPAS are based on nonrotating stellar models (Pols et al. 1998). We do not account for the effects of mild rotation, but we do implement chemically homogeneous evolution (CHE) following the recipes of Riley et al. (2021). COMPAS implements a metallicity-dependent rotational frequency threshold to determine whether a star is evolving chemically homogeneously. If a star is rotating faster than the threshold given by Riley et al. (2021) at ZAMS, we consider it to be evolving chemically homogeneously. We neglect the very limited radial evolution of a CHE star and fix its main-sequence radius equal to the ZAMS radius of a nonrotating star of the same mass and metallicity. The main-sequence evolution of a CHE star thus follows the Hurley et al. (2000) model of main-sequence stars, albeit with a fixed radius. COMPAS can be configured by the user to check, or not, the rotational frequency of the star against the CHE threshold at every time step on the main sequence. In the default model, the rotational frequency check is enabled, and if the rotational frequency drops below the threshold value for CHE (e.g., due to the orbit of the binary widening as a consequence of mass loss through winds), the star is thereafter evolved as a regular main-sequence star (i.e., it immediately jumps to the track of a regular main-sequence star of the same mass). If the rotational frequency check is disabled by the user, the star evolves chemically homogeneously through its main-sequence lifetime once it satisfies the CHE threshold at ZAMS. Finally, we assume that if a star evolves chemically homogeneously through the main sequence, it contracts directly into a naked helium star at the end of the main sequence, retaining its full mass at that point. Evolution then follows the Hurley et al. (2000) models of helium stars.

3.5. Wind Mass Loss (Single Star)

Stars lose mass through stellar winds. This impacts their evolution, and affects what remnants they form. Stars lose mass throughout their lives through several mechanisms. Hot stars lose mass through steady-state line-driven winds (Castor et al. 1975; Vink et al. 2001), while the mechanism through which cool stars like red supergiants (RSGs) lose mass is less well understood theoretically, leading to most mass-loss prescriptions for these stars being empirically derived (e.g., de Jager et al. 1988; Beasor & Davies 2018). In addition, stars close to the Humphreys–Davidson limit (Humphreys & Davidson 1994) are known to experience eruptive mass loss. For a review of mass loss from massive stars, see Smith (2014). Low-mass stars generally experience weaker stellar winds than their high-mass counterparts. However, they can experience strong wind mass loss during the later stages of their evolution on the AGB (see, e.g., Höfner & Olofsson 2018, for a recent review).

Wind mass-loss rates are highly uncertain, and can have a substantial impact on the evolution of a star (e.g., Renzo et al. 2017; Belczynski et al. 2020a). There is recent observational evidence that mass-loss rates may be overestimated for certain evolutionary phases (e.g., Smith 2014; Beasor & Davies 2018; Sander et al. 2020; Miller-Jones et al. 2021; Neijssel et al. 2021).

COMPAS currently includes two simple analytic prescriptions for wind mass loss based on a combination of theoretical simulations and observational measurements: the original wind

Table 2
Stellar Phases and Abbreviations Used in COMPAS

Stellar Phase	Abbreviation	Number
Main sequence, $M < 0.7 M_{\odot}$	MS	0
Main sequence, $M > 0.7 M_{\odot}$	MS	1
Hertzsprung gap	HG	2
First giant branch	FGB	3
Core helium burning	CHeB	4
Early asymptotic giant branch ^a	EAGB	5
Thermally pulsing asymptotic giant branch	TPAGB	6
Helium main sequence	HeMS	7
Helium Hertzsprung gap	HeHG	8
Helium giant branch	HeGB	9
Helium white dwarf	HeWD	10
Carbon–oxygen white dwarf	COWD	11
Oxygen–neon white dwarf	ONeWD	12
Neutron star	NS	13
Black hole	BH	14
Massless remnant	MR	15
Chemically homogeneously evolving	CHE	16
...	NONE ^b	19

Notes. Types 0–15 are from Hurley et al. (2000), while CHE stars (type 16) are an addition in COMPAS (Riley et al. 2021).

^a We apply the term EAGB to all H-rich stars with an inert C/O core that are primarily powered by He shell fusion, regardless of mass or location in the Hertzsprung–Russell diagram.

^b Stellar type 19 (NONE) can sometimes appear as a temporary placeholder in COMPAS outputs before the stellar type is set.

prescription from Hurley et al. (2000) and an updated prescription from Belczynski et al. (2010), which is our default.

3.5.1. Hurley Model

Hurley et al. (2000) define the total mass-loss rate as the dominant mass-loss rate during each stellar phase of the star (see Table 2), with a possible addition of LBV-like mass loss if the star is an luminous blue variables (LBV). This can be summarized as

$$\dot{M}_H = \begin{cases} \dot{M}_{\text{NJ}}, & \text{MS} \\ \max\{\dot{M}_{\text{NJ}}, \dot{M}_{\text{KR}}, \dot{M}_{\text{WR}}\} + \dot{M}_{\text{LBV}}, & \text{if HG-CHeB} \\ \max\{\dot{M}_{\text{NJ}}, \dot{M}_{\text{KR}}, \dot{M}_{\text{WR}}, \dot{M}_{\text{VW}}\}, & \text{AGB} \\ \max\{\dot{M}_{\text{NJ}}, \dot{M}_{\text{KR}}, \dot{M}_{\text{WR}}(\mu = 0)\}, & \text{HeMS-HeGB} \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where \dot{M}_{NJ} , \dot{M}_{KR} , \dot{M}_{WR} , \dot{M}_{VW} , \dot{M}_{LBV} are defined as follows.

For stars across the whole Hertzsprung–Russell (HR) diagram, Hurley et al. (2000) apply

$$\begin{aligned} \dot{M}_{\text{NJ}} = 9.6 \times 10^{-15} \max\left(0, \min\left(1, \frac{(L/L_{\odot}) - 4000}{500}\right)\right) \\ \times \left(\frac{Z}{0.02}\right)^{1/2} \left(\frac{R}{R_{\odot}}\right)^{0.81} \\ \times \left(\frac{L}{L_{\odot}}\right)^{1.24} \left(\frac{M}{M_{\odot}}\right)^{0.16} M_{\odot} \text{ yr}^{-1}, \end{aligned} \quad (5)$$

which is the mass-loss rate from Nieuwenhuijzen & de Jager (1990), modified by the metallicity scaling $Z^{1/2}$ (Kudritzki et al.




Figure 6. Maximum radial extent of stars with ZAMS masses between $0.5\text{--}150 M_{\odot}$ during different phases of stellar evolution, from the main sequence to the helium giant branch, using default settings. Solid lines show the maximum extent of stars during a given evolutionary phase, whose colors have the same meaning as in Figure 5. We additionally show R_{ZAMS} in gray to indicate how much the star expands during the main sequence. Note that the FGB curve is invisible when it overlaps with the CHeB curve. Shading indicates regions in which different types of mass transfer can occur due to the expansion of the donor; on the main sequence (yellow), after hydrogen exhaustion (orange) and after helium exhaustion (blue). The dashed lines indicate important transition masses in the Hurley et al. (2000) fitting formulae, where M_{HT} is the minimum mass of a star that will complete its main sequence in a Hubble time, M_{HeF} is the maximum initial mass for which helium ignites degenerately in a helium flash, M_{SN} is the mass above which a star collapses into an NS or BH, possibly with a supernova explosion, M_{FGB} is the maximum initial mass for which helium ignites on the first giant branch and M_{WR} is the minimum mass at which a star will self-strip to become a Wolf–Rayet star.




Figure 7. Main-sequence lifetimes of stars between $0.5\text{--}150 M_{\odot}$ using Hurley et al. (2000) fitting formulae across the range of metallicities at which these formulae can be used.

1989). This is only non-zero for luminous massive stars with $L > 4000 L_{\odot}$.

For stars on the giant branch and beyond, this model adopts the results of Kudritzki & Reimers (1978)

$$\dot{M}_{\text{KR}} = \eta \times 4 \times 10^{-13} M_{\odot} \text{ yr}^{-1} \left(\frac{L}{L_{\odot}} \right) \left(\frac{R}{R_{\odot}} \right) \left(\frac{M}{M_{\odot}} \right)^{-1}, \quad (6)$$

where η is a phenomenological scaling parameter of order unity. By default, we adopt $\eta = 0.5$ following Hurley et al. (2000).

For stars on the asymptotic giant branch they use the results of Vassiliadis & Wood (1993)

$$\dot{M}_{\text{VW}} = \min \left(10^{-11.4+0.0125(P_0-P_1)}, 1.36 \times 10^{-9} \frac{L}{L_{\odot}} \right) M_{\odot} \text{ yr}^{-1}, \quad (7)$$

where P_0 , the Mira pulsation period, is given by

$$P_0/\text{days} = \min \left[1.995 \times 10^3, 8.51 \times 10^{-3} \left(\frac{M}{M_{\odot}} \right)^{-0.9} \left(\frac{R}{R_{\odot}} \right)^{1.94} \right], \quad (8)$$

and $P_1/\text{days} = 100 \times \max(M/M_{\odot} - 2.5, 0)$.

Wolf–Rayet like wind mass loss is included for small hydrogen-envelope mass, $\mu < 1.0$, stars according to

$$\dot{M}_{\text{WR}} = 10^{-13} \left(\frac{L}{L_{\odot}} \right)^{1.5} (1 - \mu) M_{\odot} \text{ yr}^{-1}. \quad (9)$$

where the parameter μ describes the ratio of the envelope mass to the total mass, so that $\mu = 1$ on the main sequence, $0 < \mu < 1$ for stars with a developed core, and $\mu = 0$ for stripped stars with no hydrogen envelope. The full expressions for μ are given in Equation (97) in Hurley et al. (2000).

We designate post-main-sequence stars with $L > 6 \times 10^5 L_{\odot}$ and $\text{HD} \equiv 10^{-5} (R/R_{\odot}) (L/L_{\odot})^{0.5} > 1.0$ as LBV, following Humphreys & Davidson (1994). For these LBV stars Hurley et al. (2000) add an LBV-like mass loss, intended to account

for eruptive mass loss in an averaged sense

$$\dot{M}_{\text{LBV}} = 0.1(\text{HD} - 1.0)^3 \left(\frac{L/L_{\odot}}{6 \times 10^5} - 1.0 \right) M_{\odot} \text{ yr}^{-1}. \quad (10)$$

3.5.2. Belczynski Model

Belczynski et al. (2010) use a model for stellar winds based on results from Monte Carlo radiative transfer simulations of Vink et al. (2000, 2001). For stars that are not LBVs or helium stars, they define the mass-loss rate as

$$\dot{M}_{\text{B}} = \begin{cases} \dot{M}_{\text{H}}, & T/\text{K} < 12500 \\ \dot{M}_{\text{Vink,BBJ}}, & 12500 \leq T/\text{K} < 25000 \\ \dot{M}_{\text{Vink,ABJ}}, & 25000 \leq T/\text{K}, \end{cases} \quad (11)$$

where \dot{M}_{H} is the Hurley mass-loss rate defined in Section 3.5.1 and the Vink mass-loss rates are defined as

$$\begin{aligned} \log \left(\frac{\dot{M}_{\text{Vink,BBJ}}}{M_{\odot} \text{ yr}^{-1}} \right) &= -6.688 + 2.210 \log(L/10^5 L_{\odot}) \\ &- 1.339 \log(M/30 M_{\odot}) - 1.601 \log(V/2.0) + \\ &0.85 \log(Z/Z_{\odot}) + 1.07 \log(T/20000 \text{ K}), \end{aligned} \quad (12)$$

below the bistability jump, where the ratio of the wind speed at infinity to the star's escape velocity is $V = v_{\infty}/v_{\text{esc}} = 1.3$, and as

$$\begin{aligned} \log \left(\frac{\dot{M}_{\text{Vink,ABJ}}}{M_{\odot} \text{ yr}^{-1}} \right) &= -6.697 + 2.194 \log(L/10^5 L_{\odot}) \\ &- 1.313 \log(M/30 M_{\odot}) - 1.226 \log(V/2.0) \\ &+ 0.85 \log(Z/Z_{\odot}) + 0.933 \log(T/40000 \text{ K}) \\ &- 10.92 [\log(T/40000 \text{ K})]^2, \end{aligned} \quad (13)$$

above the bistability jump, where the ratio of the wind speed at infinity to the stars escape velocity is $V = v_{\infty}/v_{\text{esc}} = 2.6$. The mass-loss rate in Equations (12) and (13) above scales with metallicity as $\dot{M} \propto Z^{0.85}$, in agreement with observationally determined scaling of the mass-loss rates of O and B stars with metallicity in the Milky Way and the Magellanic Clouds (Mokiem et al. 2007).

For helium stars, Belczynski et al. (2010) assume a mass-loss rate

$$\dot{M}_{\text{B,He}} = f_{\text{WR}} \times 10^{-13} \left(\frac{L}{L_{\odot}} \right)^{1.5} \left(\frac{Z}{Z_{\odot}} \right)^m M_{\odot} \text{ yr}^{-1}, \quad (14)$$

from Hamann & Koesterke (1998), with $m = 0.86$, as given by Vink & de Koter (2005). We have introduced the phenomenological scaling parameter f_{WR} to allow the strength of WR winds to be varied (Barrett et al. 2018). Our default choice is $f_{\text{WR}} = 1$.

LBV stars have high mass-loss rates due to both line-driven winds and eruptive mass loss. The uncertainty in LBV mass-loss rates is parameterized with a scaling parameter f_{LBV} with a default value of 1.5 (Belczynski et al. 2010):

$$\dot{M}_{\text{B,LBV}} = f_{\text{LBV}} \times 10^{-4} M_{\odot} \text{ yr}^{-1}. \quad (15)$$




Figure 8. Top: the relation between the initial (ZAMS) mass of single stars at solar metallicity ($Z = 0.0142$) and either the WD remnant mass or the CO core mass at the moment of compact object formation for NSs and BHs, assuming the COMPAS default model, where colors denote the remnant type. Middle: the same relation at $Z = 0.01$ (solid), $Z = 0.001$ (dashed) and $Z = 0.0001$ (dotted) to show the metallicity dependence. The inset panel shows the same quantities as the main panel, but with linear axes, focusing on BH masses. Bottom: the initial–compact object mass relation at Z_{\odot} plotted for different remnant mass models (see Section 3.8.2): the *delayed* (dotted line) and *rapid* (dashed line) prescriptions from Fryer et al. (2012), and the stochastic model from Mandel & Müller (2020) (points converging to a solid line).

3.6. Treatment of the Impact of Mass Loss and Gain

The stellar tracks computed by Pols et al. (1998) assumed no mass loss. When a star loses mass during a core-burning phase

(either through stellar winds, as described in Section 3.5, or through mass transfer), its luminosity will decrease. This will extend the remaining lifetime. Meanwhile, mass gain through mass transfer (COMPAS does not account for accretion of winds from the companion) can *rejuvenate* the star. After mass gain such a star will appear younger than a single star of the same mass and age (a *blue straggler*).

For a star losing or gaining mass through either mass transfer or winds (mass loss only) on the main sequence, we follow Hurley et al. (2000; see also Tout et al. 1997) in modifying the lifetime according to

$$t' = f_{\text{rej}} \frac{t'_{\text{MS}}}{t_{\text{MS}}} t, \quad (16)$$

where t and t_{MS} are the effective age and main-sequence lifetime prior to a small change in mass, and t' and t'_{MS} are the age and main-sequence lifetime of the star after a small amount of mass loss/gain. Mass changes for helium main-sequence stars are treated by analogy with the above equation, with the main-sequence lifetimes t_{MS} and t'_{MS} replaced with t_{HeMS} and t'_{HeMS} , respectively (Hurley et al. 2000).

The prefactor f_{rej} is unity for all mass-losing stars (and for low-mass $M \leq 0.7 M_{\odot}$ main-sequence stars, stellar type 0), while it is taken to be the ratio of the mass before/after mass gain for main-sequence stars initially more massive than $0.7 M_{\odot}$ (Belczynski et al. 2008) and helium main-sequence stars. The ages of Hertzsprung gap stars are updated following mass changes as described in Hurley et al. (2000). For giants with clearly decoupled cores and envelopes, we assume that changes to the envelope (such as mass loss/gain) do not affect the remaining lifetime. See Schneider et al. (2015) and references therein for more details and possible improvements.

3.7. Stellar Remnants

Standard stellar evolution theory predicts that low and intermediate mass stars (with initial masses $\lesssim 8 M_{\odot}$) typically end their lives as WDs (Section 3.7.1), while more massive stars end their lives by collapsing into NSs and BHs (Section 3.7.2). We show the relation between initial masses, core masses and remnant masses of stars evolved under default COMPAS assumptions in Figure 8.

3.7.1. White Dwarfs

We distinguish between three different types of WDs based on their mass and composition, following Hurley et al. (2000). Stars that lose their envelopes prior to helium ignition leave behind helium white dwarfs (HeWDs), while those that lose their envelopes after core-He exhaustion leave behind either carbon–oxygen white dwarfs (COWDs) or oxygen–neon white dwarfs (ONeWDs) depending on their core masses upon reaching the base of the asymptotic giant branch. Larger core masses are associated with higher temperatures that allow carbon to fuse, forming oxygen–neon or oxygen–neon–magnesium cores. We assume that helium core masses below $1.6 M_{\odot}$ at the base of the asymptotic giant branch lead to COWD formation, while core masses above that lead to ONeWD formation. For our default model, we find that single stars with initial masses $\lesssim 0.8 M_{\odot}$ form HeWDs (though the evolutionary timescale for these stars to evolve is longer than a Hubble time), while heavier stars with initial masses up to $\approx 7 M_{\odot}$ form COWDs. Only stars in a narrow mass range

of $\approx 7\text{--}8 M_{\odot}$ form ONeWDs (see upper panel of Figure 8). The boundary between stars that form WDs and those that form NSs/BHs is uncertain (see, for example, Doherty et al. 2017, for further discussion). WDs have masses $\lesssim 1.4 M_{\odot}$ (the Chandrasekhar mass) in our model. We determine the radius of WDs following Tout et al. (1997) and Hurley et al. (2000). WDs cool as they age (Mestel 1952). We model the luminosity of WDs using the Mestel cooling track given in Equation (90) in Hurley et al. (2000). By default we assume that WDs do not receive any kick during formation (but see El-Badry & Rix 2018 for a discussion of evidence for small $\sim 1 \text{ km s}^{-1}$ recoil velocities associated with WD formation).

3.7.2. Neutron Stars and Black Holes

Stars with initial masses more than $\sim 8 M_{\odot}$ collapse into an NS or a BH at the end of their lives. This collapse may be accompanied by a supernova explosion. Asymmetry in the supernova may give the remnant a substantial momentum boost, commonly known as a *natal kick* (see Section 3.9). The mass of the supernova remnant (and by extension its stellar type), as well as the natal kick all depend on which kind of supernova it undergoes. In COMPAS we distinguish between several different types of supernovae, which are discussed in Section 3.8.

COMPAS distinguishes NSs from BHs by the remnant mass; the default value of the maximum NS mass is $2.5 M_{\odot}$, following Fryer et al. (2012).

Figure 8 shows the core mass and remnant mass as a function of initial mass for single stars in the mass range $0.1\text{--}150 M_{\odot}$ in the COMPAS default model. The top panel of Figure 8 shows the mass ranges in which each type of stellar remnant is formed at solar metallicity. The middle panel of Figure 8 shows how the initial mass-core mass relation varies with metallicity, while the bottom panel shows how the relation depends on which remnant mass prescription is used for solar metallicity. This figure can be compared to results from Belczynski et al. (2010) and Banerjee et al. (2020).

By default, in COMPAS we follow Hurley et al. (2000) in assuming all NSs have a radius of 10 km. We also include the possibility of defining the NS radius as a function of its mass $R_{\text{NS}}(M_{\text{NS}})$, as expected from the equation of state. Currently, the NS equation of state from Akmal et al. (1998) is implemented in COMPAS, which gives NS radii of 11–12 km for NSs in the astrophysically relevant mass range $1\text{--}2.4 M_{\odot}$. The maximum NS mass predicted by this equation of state is $2.4 M_{\odot}$.

The NS moment of inertia is assumed to follow the “universal” (equation of state insensitive) relation given by Lattimer & Schutz (2005). The luminosity of NSs is estimated following Hurley et al. (2000).

For output purposes only, the “radius” of a BH is given by its Schwarzschild radius

$$R_{\text{BH}} = \frac{2GM}{c^2} \approx 4.24 \times 10^{-6} R_{\odot} \frac{M}{M_{\odot}} \quad (17)$$

and we follow Hurley et al. (2000) in arbitrarily setting the BH luminosity to $10^{-10} L_{\odot}$.

Pulsar evolution is implemented in COMPAS as described in Chattopadhyay et al. (2020, 2021). We assign a spin period and an initial magnetic field to each newly born NS. We provide several options for the initial distribution of pulsar spin periods and magnetic fields, see Chattopadhyay et al. (2020) for details.

The evolution of the pulsar spin period and spindown rate are followed as a function of time, assuming the canonical magnetic dipole model for a pulsar. In this model, pulsars spindown over time due to magnetic braking, but may be spun up again (or recycled) through mass accretion. Our approach closely follows the methodology of Osłowski et al. (2011), Kiel et al. (2008), and Faucher-Giguere & Kaspi (2006; see also Ye et al. 2019). Other pulsar properties, such as the pulsar luminosity or beaming fraction, are not directly computed by COMPAS, but can easily be modeled in postprocessing (see Chattopadhyay et al. 2020, 2021, for details). Pulsar evolution is optional in COMPAS and is disabled in the default model.

BHs are assumed to be nonspinning. However, the spins of BHs have been modeled by postprocessing COMPAS data in a number of papers (Bavera et al. 2020; Chattopadhyay et al. 2021).

3.8. Supernova Types

In COMPAS we distinguish between several different types of supernovae. In the following, we describe the types of supernovae we model and the conditions under which each type of supernova is assumed to occur. COMPAS records whether the supernova progenitor star has a hydrogen-rich envelope, allowing for a crude estimation of whether it would appear observationally as a type I or II supernova.

3.8.1. Electron-capture Supernovae

In COMPAS we assume that a star undergoes an electron-capture supernova (ECSN; Miyaji et al. 1980; Nomoto 1984, 1987; Ivanova et al. 2008) if it has a helium core mass in the range $1.6\text{--}2.25 M_{\odot}$ (Hurley et al. 2002) at the base of the asymptotic giant branch, and the carbon–oxygen core mass reaches a threshold of $1.38 M_{\odot}$. In our default model, this corresponds to a ZAMS mass range of $7.5\text{--}8.1 M_{\odot}$ for single stars at Z_{\odot} .

The mass range of stars that undergo ECSNe is somewhat uncertain and model dependent. For example, Podsiadlowski et al. (2004) argued that a more realistic range of helium core masses leading to ECSNe is $1.4\text{--}2.5 M_{\odot}$. Andrews et al. (2015) and Vinciguerra et al. (2020) use different core mass ranges for ECSNe ($2\text{--}2.5 M_{\odot}$ and $1.83\text{--}2.25 M_{\odot}$, respectively, where the latter is based on Fryer et al. (2012), which could better reproduce observations of NSs). Willcox et al. (2021) argued that the (ZAMS) mass range of (effectively) single stars undergoing ECSNe cannot be wider than $0.2 M_{\odot}$ to avoid overproducing low velocity pulsars, but ECSNe could be more common in binaries.

If a star undergoes an electron-capture supernova, we set its remnant to be an NS with a mass of $1.26 M_{\odot}$, as an approximation to the solution of Equation (8) in Timmes et al. (1996), assuming a baryonic mass of $1.38 M_{\odot}$.

Massive oxygen–neon WDs (ONeWDs) close to the Chandrasekhar mass can accrete enough mass to undergo an accretion-induced collapse (AIC) to an NS (Nomoto & Kondo 1991). While AIC is nominally possible in COMPAS when the mass of a ONeWD reaches $1.38 M_{\odot}$, when the ECSN prescription is followed for the remnant mass and natal kick, this is based on a very simplistic implementation of accretion onto WDs.

3.8.2. Core-collapse Supernovae

Stars with helium core masses greater than $2.25 M_{\odot}$ at the base of the asymptotic giant branch undergo core-collapse

supernovae (CCSNe) in COMPAS, if/when their carbon–oxygen core mass reaches the threshold given by Hurley et al. (2000), where we replace their Chandrasekhar mass threshold with $1.38 M_{\odot}$.

There is still a great deal of uncertainty regarding the mechanism of CCSNe, including which stars explode, how they explode, and what the properties of their remnants are (see Müller 2020 and Burrows & Vartanyan 2021 for recent reviews). Some recent supernova simulations predict that whether a supernova is successful and leads to an explosion, or fails and leads to an implosion, is a nonmonotonic function of a star’s initial mass (Ugliano et al. 2012; Sukhbold & Woosley 2014; Nakamura et al. 2015; Ertl et al. 2016; Sukhbold & Adams 2020). COMPAS uses simple parameterized models (described below) to relate the properties of supernova remnants to their progenitor stars. We attempt to parameterize some of this uncertainty by including several different models.

By default, COMPAS uses the *delayed* supernova remnant mass prescription from Fryer et al. (2012) to map the carbon–oxygen core masses of stars to compact object remnant masses during core-collapse supernova events. An alternative *rapid* prescription assumes supernova explosions occur within 250 ms (compared to longer timescales assumed for the delayed model) and reproduces, by construction, a mass gap between NSs and BHs. Mandel & Müller (2020) proposed a model for compact object masses and kicks which is stochastic, with both NS and BH formation possible in certain regions of parameter space; this model was implemented in COMPAS in Mandel et al. (2021). Schneider et al. (2021) predict that the history of mass transfer impacts the remnant mass, and we include their remnant mass prescription. We also include the slightly older remnant mass prescriptions from Hurley et al. (2000) and Belczynski et al. (2008), which define the remnant mass as a piecewise function or linear function, respectively, of the progenitor carbon–oxygen core mass; these models are provided for historical consistency reasons, but are somewhat outdated.

We convert the baryonic mass of the remnant to a gravitational mass using Equation (13) of Fryer et al. (2012) for NSs. For BHs, the gravitational mass is assumed to be $0.1 M_{\odot}$ less than the baryonic mass to account for mass lost in neutrinos in the default prescription (Stevenson et al. 2019). This amount can be adjusted, or, alternatively, a fixed fraction of the mass can be lost in neutrinos during BH formation as in Equation (14) of Fryer et al. (2012).

We show the difference in the initial-final remnant mass relation between the two models from Fryer et al. (2012) and the stochastic model from Mandel & Müller (2020) in the bottom panel of Figure 8.

3.8.3. Ultra-stripped Supernovae

Mass transfer from a helium star that re-expands after core helium burning (so-called “case BB mass transfer”; Delgado & Thomas 1981; Tauris et al. 2015) in short-period binaries leads to severe stripping of the donor, leaving behind a helium envelope with mass $\lesssim 0.1 M_{\odot}$. If the remaining stellar core is sufficiently massive to undergo core collapse, then we say that it undergoes an ultra-stripped supernova (USSN; Tauris et al. 2013, 2015). Due to the lack of envelope, USSNe remnants are NSs with characteristically lower mass, and may receive smaller natal kicks than typical NSs (Suwa et al. 2015; Moriya

et al. 2017; Müller et al. 2018). We discuss kicks further in Section 3.9. By default in COMPAS we assume that case BB mass transfer is always stable (see Vigna-Gómez et al. 2018) and removes the entire helium envelope but none of the underlying carbon–oxygen core.

3.8.4. Pair-instability Supernovae

Stars with helium cores in the mass range $\sim 50\text{--}150 M_{\odot}$ are believed to become unstable as a result of electron–positron pair production (Fowler & Hoyle 1964; Barkat et al. 1967; Woosley 2017; Farmer et al. 2019). This causes the radiation pressure support in the core to drop, causing the core to contract. As it contracts, the temperature increases, triggering explosive oxygen burning. This may reverse the contraction and completely unbind the star in a pair-instability supernovae (PISNe) explosion, leaving no remnant behind (Fowler & Hoyle 1964; Barkat et al. 1967; Fraley 1968; Kozyreva et al. 2014a, 2014b, 2017; Takahashi 2018; Leung et al. 2019; Woosley 2019). BH formation is expected again above a helium core mass of $\sim 150 M_{\odot}$ (Woosley et al. 2002; Woosley 2019). In addition, it is theoretically well established that stars with helium core masses in the range $\sim 30\text{--}60 M_{\odot}$ lead to pulsational PISNe (e.g., Yoshida et al. 2016; Spera & Mapelli 2017; Woosley 2017; Takahashi 2018; Farmer et al. 2019; Marchant et al. 2019; Renzo et al. 2020), where material from the star is ejected in several supernova-like pulses, while the star returns to equilibrium between each pulse, and eventually undergoes an iron core collapse to form a BH. Together, these effects lead to a dearth of BHs in the mass range $\sim 45\text{--}130 M_{\odot}$ and a possible excess in the $\sim 35\text{--}45 M_{\odot}$ mass range with the current COMPAS default prescription.

The implementation of PISNe in COMPAS is discussed in detail in Stevenson et al. (2019). In brief, by default stars with helium core masses in the range between 35 and 60 M_{\odot} lose mass through pulsational PISNe prior to collapse, while those with helium core masses between 60 and 135 M_{\odot} explode in PISNe and leave no remnants behind. We use fits based on results from Marchant et al. (2019) in the COMPAS default model for the relation between the helium core mass and the final presupernova mass in the pulsational PISNe regime. In addition, COMPAS also provides the simple model from Belczynski et al. (2016a) to allow for comparison with StarTrack models, along with additional models based on detailed results from Woosley (2017) and Farmer et al. (2019).

3.9. Supernova Natal Kicks

Galactic pulsars are observed to have large proper motions, from which a distribution of their velocities is inferred (e.g., Gunn & Ostriker 1970; Lyne & Lorimer 1994; Hansen & Phinney 1997; Arzoumanian et al. 2002; Hobbs et al. 2005; Beniamini & Piran 2016; Verbunt et al. 2017). The high velocities are attributed to significant asymmetries in the supernova explosions which “kick” the pulsars, though the exact mechanism is uncertain and may be attributed to either hydrodynamic effects (e.g., Janka & Müller 1994; Burrows & Hayes 1996; Wongwathanarat et al. 2013) or neutrino emission (Woosley 1987; Bisnovatyi-Kogan 1993; Socrates et al. 2005; Nagakura et al. 2019). See Lai et al. (2001) for a broad overview of natal kicks.

By default, for NSs we draw the the natal-kick magnitudes from a Maxwell–Boltzmann distribution

$$p(v_{\text{kick}}|\sigma_{\text{kick}}) = \sqrt{\frac{2}{\pi}} \frac{v_{\text{kick}}^2}{\sigma_{\text{kick}}^3} \exp\left(-\frac{v_{\text{kick}}^2}{2\sigma_{\text{kick}}^2}\right), \quad (18)$$

with root-mean-square 1D velocity σ_{kick} (i.e., root-mean-square speed of $\sqrt{3}\sigma_{\text{kick}}$). For CCSNe, we assume that $\sigma_{\text{kick}} = \sigma_{\text{CCSN}} = 265 \text{ km s}^{-1}$ (Hobbs et al. 2005). ECSNe and USSNe are expected to have smaller kicks than standard iron core-collapse supernovae (e.g., Suwa et al. 2015; Gessner & Janka 2018; Müller et al. 2019). By default, we assume $\sigma_{\text{ECSN}} = \sigma_{\text{USSN}} = 30 \text{ km s}^{-1}$ (Vigna-Gómez et al. 2018). These lower natal kicks for ECSN and USSN follow Pfahl et al. (2002a) and Podsiadlowski et al. (2004), which are motivated by the subset of DNSs and NS–binary systems with low velocities and small eccentricities (Briskin et al. 2002; Schwab et al. 2010; Beniamini & Piran 2016; Tauris et al. 2017), as well as NS retention fractions in globular clusters (e.g., Pfahl et al. 2002b).

Several authors have proposed that NS kicks should be proportional to the amount of ejecta, and inversely proportional to the remnant mass (Bray & Eldridge 2016, 2018; Giacobbo & Mapelli 2020; Mandel & Müller 2020). A scaling with ejecta mass would naturally account for reduced kicks in USSNe with low ejecta mass. In COMPAS we have implemented the fits of this form from Bray & Eldridge (2018) and Mandel & Müller (2020), where the latter model self-consistently predicts both the remnant mass and natal kick.

COMPAS also includes a model for remnant masses and kicks based on the 1D parameterized supernova simulations of Müller et al. (2016; see Vigna-Gómez et al. 2018, for details).

We further provide the option to use a uniform distribution of kick velocities up to some maximum $v_{\text{kick}}^{\text{max}}$, as well as kicks fixed at a specific value.

Whether BHs also receive natal kicks, and what their magnitudes are is an open astrophysical question. There is some evidence, both theoretical and observational, that BHs receive smaller kicks than NSs (see, e.g., Janka 2013; Mandel 2016; Repetto et al. 2017; Atri et al. 2019).

COMPAS currently includes four different models for BH natal kicks: “full,” “reduced,” our default model “fallback,” and “zero.” In the “full” model, we assume that BHs receive the full kick drawn from Equation (18), where by default we assume $\sigma_{\text{kick}} = \sigma_{\text{CCSN}}$ for BHs. In both the “reduced” and the “fallback” models, we calculate the intensity of the kick velocity imparted to newly born BHs from the kick that a NS would have received according to Equation (18) with $\sigma_{\text{kick}} = \sigma_{\text{CCSN}}$. In the “reduced” model, we assume that NSs and BHs receive the same *momentum* during the explosion, such that the kick velocity of a BH should be rescaled according to

$$v_{\text{BH}} = \frac{M_{\text{NS}}}{M_{\text{BH}}} v_{\text{NS}}, \quad (19)$$

where M_{NS} is taken to be $1.4 M_{\odot}$. In the “fallback” model, we scale the BH kick by the fraction of mass falling back onto the proto-NS f_{fb} (Fryer et al. 2012)

$$v_{\text{BH}} = v_{\text{NS}}(1 - f_{\text{fb}}). \quad (20)$$

Note that unlike Fryer et al. (2012), we apply Equation (20) even if the amount of mass falling back is less than $M = 0.2 M_{\odot}$. In COMPAS, by default we apply Equation (20) to all compact objects when using this prescription. In the “zero” kicks model, BHs receive no kick during their collapse.

By default, we assume the supernova kick angle is drawn isotropically from the unit sphere in the rest frame of the supernova progenitor. However, there is some tentative observational evidence for spin-kick alignment in pulsars (e.g., Lai et al. 2001; Johnston et al. 2005; Noutsos et al. 2012; Yao et al. 2021), which may indicate that kicks are preferentially aligned to the spin axis of the progenitor. We provide various alternative models for preferential kick directionality, including along the progenitor spin axis (or in a finite cone around it), or within the progenitor spin plane (or in a wedge around it). We describe the impact of the supernova on the orbit of a binary in Section 4.3.

4. Binary Stellar Evolution

Interacting binaries are the core case of study for COMPAS. Mass transfer is arguably the most important process in interacting binaries, modifying the component stars and the orbital properties (Podsiadlowski et al. 1992; Podsiadlowski 2010). Binary evolution can lead to stellar mergers, disruption, or DCO formation.

A binary in COMPAS is parameterized by the orbital properties: the component masses of the primary M_1 and secondary M_2 , the semimajor axis a , and eccentricity e . The orbital angular momentum of a binary with nonrotating stars is

$$J_{\text{orb}} = \mu \sqrt{GM_{\text{tot}} a (1 - e^2)}, \quad (21)$$

where $M_{\text{tot}} = M_1 + M_2$ and $\mu = M_1 M_2 / M_{\text{tot}}$ is the reduced mass. The effects of binary interactions on the stellar components, such as mass transfer and rejuvenation, are considered within our SSE framework (Section 3). We currently only account for stellar rotation in the context of CHE systems (Section 3.4). In the absence of rotation and tides, the binary can be treated as a two-body problem in the point mass approximation.

In this Section we present the details of our treatment of BSE. In Section 4.1 we discuss stellar winds in the context of binaries. In Section 4.2 we present our implementation of mass transfer. In Section 4.3 we describe the effect of supernovae on the orbit of the binary. In Section 4.4 we describe our implementation of gravitational radiation. Finally, in Section 4.5 we briefly discuss the main caveats in our implementation of binary evolution.

4.1. Wind Mass Loss (Binary System)

In Section 3.5 we presented the prescriptions and parameterization of stellar wind mass-loss rates in COMPAS. Here we present the effect of wind mass loss on the orbit of the binary system.

We consider the case of gradual mass loss ($M/\dot{M} \gg P_{\text{orb}}$) in which material rapidly leaves the system with the specific angular momentum of the mass-losing star (wind velocities are fast compared to the orbital velocity). We assume that the fast winds are emitted spherically symmetrically from the star’s surface (Huang 1956). This is known as the *Jeans mode* or *fast winds*

mode (Huang 1963). In the case, the orbit widens according to

$$\frac{\dot{a}}{a} = -\frac{\dot{M}_{\text{tot}}}{M_{\text{tot}}}, \quad (22)$$

which can be simplified as $aM_{\text{tot}} = \text{constant}$, while eccentricity is unchanged (Dosopoulou & Kalogera 2016). Winds can thus be considered as an extreme case of nonconservative mass transfer (Section 4.2).

Currently we do not consider wind accretion (e.g., Bondi & Hoyle 1944), wind RLOF (Mohamed & Podsiadlowski 2007; Hirai & Mandel 2021), or wind interaction with the companion (Brookshaw & Tavani 1993; Schröder et al. 2021).

4.2. Mass Transfer

The physics and timescales involved in mass exchange are broad, complex, and their parameterizations can be convoluted. In COMPAS, we use a simplified approach to this complicated problem. Our approach is similar to that of Belczynski et al. (2002, 2008), Hurley et al. (2002), and Postnov & Yungelson (2014), among others. We distinguish between a donor (subscript d) and an accretor (subscript a). The donor is the star that transfers (and loses) mass, while the accretor gains mass. The accretor can be a star or a compact object, and can fully retain the transferred mass (conservative mass transfer) or only a fraction of it (nonconservative mass transfer). If the mass transfer is nonconservative, i.e., some mass is lost from the binary, there is a change in the total angular momentum of the binary. We follow the orbital evolution of the binary through a mass transfer episode by taking the time derivative of Equation (21) and rearranging it as

$$\frac{\dot{a}}{a} = 2 \frac{\dot{J}_{\text{orb}}}{J_{\text{orb}}} - 2 \frac{\dot{M}_d}{M_d} - 2 \frac{\dot{M}_a}{M_a} + \frac{\dot{M}_d + \dot{M}_a}{M_d + M_a} + \frac{2e\dot{e}}{1 - e^2}. \quad (23)$$

The default COMPAS assumption is that the binary is instantly circularized to periastris, $r_p = a(1 - e)$, at the onset of RLOF. COMPAS options allow circularization with angular momentum conservation (at a separation of the semilatus rectum, $a(1 - e^2)$) or mass transfer with unchanged eccentricity.

We parameterize the fraction of mass lost by the donor which is accreted by the accretor with a factor β :

$$\dot{M}_a = -\beta \dot{M}_d, \quad (24)$$

with $0 \leq \beta \leq 1$. We parameterize the change in angular momentum by assuming the nonaccreted matter leaves the system with γ times the specific orbital angular momentum, i.e., $\dot{J}_{\text{orb}}/\dot{M} = \gamma J_{\text{orb}}/M$. Following these assumptions we can rewrite Equation (23) as²²

$$\frac{\dot{a}}{a} = -2 \frac{\dot{M}_d}{M_d} \left[1 - \beta \frac{M_d}{M_a} - (1 - \beta) \left(\gamma + \frac{1}{2} \right) \frac{M_d}{M_d + M_a} \right], \quad (25)$$

which is the equation we use to solve for the orbital evolution during a mass transfer episode. Note that fully nonconservative mass transfer ($\beta = 0$), where mass is lost with the specific

²² See the educational lecture notes on binary evolution by Onno Pols: http://www.astro.ru.nl/~onnop/education/binaries_utrecht_notes/.

angular momentum of the donor star ($\gamma = M_a/M_d$), is equivalent to fast wind mass loss from Section 4.1.

Before solving for the orbital evolution during a mass transfer phase, we first need to determine if mass transfer will occur (Section 4.2.1) and, if so, whether the mass transfer episode will be dynamically stable (Sections 4.2.2 and 4.2.3). If the mass transfer episode is dynamically unstable, it will lead to a common envelope phase (Section 4.2.4).

4.2.1. Roche-lobe Overflow

A mass transfer phase can be initiated either by the radial expansion of a star as a consequence of stellar evolution or by a decrease in the binary separation. The star overflows its Roche lobe and the surface material is transferred from the donor to the companion through the first Lagrangian point. The first Lagrangian point is the juncture between the two Roche lobes, which are the regions that contain the gravitationally bound material around each star. In COMPAS we follow Eggleton (1983) and approximate the radius of the donor's Roche lobe normalized by the separation as

$$r_{\text{RL}} = R_{\text{RL}}/a = 0.49 \frac{q_{\text{RL}}^{2/3}}{0.6q_{\text{RL}}^{2/3} + \ln(1 + q_{\text{RL}}^{1/3})}, \quad (26)$$

with $q_{\text{RL}} = M_d/M_a$. Equation (26) assumes point-like masses in a circular orbit. The Roche lobe is generally shaped like a teardrop, but here the Roche lobe radius is defined as the radius of a sphere which has the same volume as the Roche lobe. The condition for RLOF, and therefore for initiating mass transfer, is when the radius of the star is larger than the Roche radius at periaapsis, i.e., when $R > R_{\text{RL}}(1 - e)$.

4.2.2. Stability Criteria

When the condition for RLOF is satisfied, we need to determine if the mass transfer episode will be dynamically stable or lead to a common envelope event. In COMPAS, the stability of the mass transfer phase is determined based on approximations to the mass–radius relationships $\zeta \equiv d \ln R/d \ln M$ (see, e.g., Soberman et al. 1997). Namely, the response of the radius of the donor star to mass loss $\zeta_* \equiv d \ln R_*/d \ln M$ is compared to the response of the Roche-lobe radius $\zeta_{\text{RL}} \equiv d \ln R_{\text{RL}}/d \ln M$ to mass transfer (Paczynski & Sienkiewicz 1972; Hjellming & Webbink 1987; Soberman et al. 1997). If $\zeta_* \geq \zeta_{\text{RL}}$, then the mass transfer episode is assumed to be stable, otherwise, mass transfer is assumed to lead to a common envelope phase.

We approximate the value of ζ_* based on stellar types (Section 3.2). We assume $\zeta_* = 2$ for main-sequence and HeMS stars and $\zeta_* = 6.5$ for HG stars, as implemented in Vigna-Gómez et al. (2018), based on typical values from Ge et al. (2015). These ζ_* values translate into critical mass ratios for stable mass transfer of $M_d/M_a < 1.72$ ($M_d/M_a < 2.25$) for fully conservative (fully nonconservative) mass transfer from main-sequence donors and $M_d/M_a < 3.83$ ($M_d/M_a < 4.58$) for HG donors. For stellar types HG, FGB, CHeB, EAGB, and TPAGB (where we use these stellar types, defined in Section 3.2, as loose proxies for having a convective envelope, but see, e.g., Klencki et al. 2021) we follow Soberman et al. (1997) in the

form

$$\zeta_{\text{SPH}} = \frac{2}{3} \left(\frac{m_{\text{SPH}}}{1 - m_{\text{SPH}}} \right) - \frac{1}{3} \left(\frac{1 - m_{\text{SPH}}}{1 + 2m_{\text{SPH}}} \right) - 0.03m_{\text{SPH}} + 0.2 \left[\frac{m_{\text{SPH}}}{1 + (1 - m_{\text{SPH}})^{-6}} \right], \quad (27)$$

where $m_{\text{SPH}} = M_{\text{core}}/M$ and M_{core} is the core mass as defined by Hurley et al. (2000). In COMPAS, we consider stripped post-helium-burning stars as a special case and assume by default that any mass transfer episode from HeHG and HeGB stellar types is always stable. This is in agreement with the expected outcome of mass transfer episodes from stripped stars onto NSs or BHs as suggested by Tauris et al. (2013, 2015). We currently do not model mass transfer from white dwarf donors (HeWD, COWD, and ONeWD).

The Roche-lobe mass–radius exponent ζ_{RL} depends on the accreted mass fraction β and the specific angular momentum that nonaccreted mass removes from the system (see Section 4.2.3 to see how we determine these). We follow Soberman et al. (1997) and Woods et al. (2012) in rewriting

$$\zeta_{\text{RL}} = \frac{d \ln R_{\text{RL}}}{d \ln M} = \frac{\partial \ln a}{\partial \ln M_d} + \frac{\partial \ln r_{\text{RL}}}{\partial \ln q_{\text{RL}}} \frac{\partial \ln q_{\text{RL}}}{\partial \ln M_d}, \quad (28)$$

where the terms are taken from Equation (25) and the partial derivative of Equation (26).

4.2.3. Stable Mass Transfer

The amount of mass transferred during a dynamically stable mass transfer episode is calculated depending on the stellar type of the donor (Section 3.2). Broadly, we distinguish between stars that have a clear core/envelope separation (HG, FGB, CHeB, EAGB, TPAGB, HeHG, and HeGB) and those which do not (MS and HeMS). We do not consider other donor types.

For donor stars with a clear core/envelope separation we transfer their entire envelope. For donor stars without a clear core/envelope separation, we calculate and remove the minimal mass necessary in order to have the donor fit within its Roche lobe. Numerically, we accomplish this by using a root-finder with specified tolerances. We compute the accreted mass fraction β once, at the start of the mass transfer phase, as follows.

We compute the donor and accretor mass transfer rates. For the donor, we assume a thermal timescale mass transfer rate

$$\dot{M}_d = M_d/\tau_{\text{KH},d}. \quad (29)$$

For the accreting object, the maximum accretion rate is limited depending on its stellar type

$$\dot{M}_a = \begin{cases} \dot{M}_{\text{KH},a} & \text{if stellar accretor,} \\ \dot{M}_{\text{edd}} & \text{if compact object accretor.} \end{cases} \quad (30)$$

In Equation (30), the thermal mass accretion rate is given by

$$\dot{M}_{\text{KH},a} = C \cdot M_a/\tau_{\text{KH},a}, \quad (31)$$

where the factor $C = 10$ (default value) is assumed to take into account the expansion of the star due to mass transfer, following Paczyński & Sienkiewicz (1972), Neo et al. (1977), Hurley et al. (2002) and Schneider et al. (2015). Meanwhile,

the Eddington-limited accretion rate is given by

$$\dot{M}_{\text{edd}} = 1.5 \times 10^{-8} \frac{R_*}{10 \text{ km}} \frac{M_{\odot}}{\text{yr}}, \quad (32)$$

where for a black hole the radius R_* is given by the Schwarzschild radius as in Equation (17). The Eddington accretion rate assumptions in COMPAS are flexible and can be change to a user-specified function (see van Son et al. 2020 for more details).

By comparing the donor and accretor mass transfer rates, we can determine how conservative the mass transfer episode is. Only for main-sequence (MS and HeMS) and Hertzsprung gap (HG and HeHG) stars, mass accretion leads to stellar rejuvenation (Section 3.6) as the accretor star transitions to a more massive and less evolved stellar track (Tout et al. 1997; Hurley et al. 2002; Belczynski et al. 2008). For fully conservative mass transfer, i.e., $\beta = 1$, there is no angular momentum drained from the binary and Equation (25) simplifies to $M_d^2 M_a^2 a = \text{constant}$. For nonconservative mass transfer, i.e., $\beta < 1$, the nonaccreted mass is lost, by default, through *isotropic re-emission* from the vicinity of the accreting star (e.g., Masevitch & Yungelson 1975; Bhattacharya & van den Heuvel 1991; Soberman et al. 1997). Isotropic re-emission corresponds to $\gamma = M_d/M_a$.

We briefly point out some alternatives from the default mass transfer model which the user can choose in COMPAS. The user can specify a fixed value for β , which enforces that a fraction β of the mass transferred by the donor is accreted, except for Eddington-limited accretion. The value of γ can also be changed. We have mentioned the Jeans ($\gamma = M_a/M_d$) mass-loss mode in the context of winds, but we also include the user-defined possibility of choosing these mass-loss modes in semiconservative and fully nonconservative mass transfer episodes. COMPAS also includes a mode of nonconservative mass transfer that represents mass loss via a circumbinary ring as an option (see, e.g., Soberman et al. 1997). In this case, the nonaccreted mass carries a specific angular momentum $\gamma = \sqrt{a_{\text{ring}}/a} M/\mu = \sqrt{2} M/\mu$ for a ring at radius $a_{\text{ring}} = 2a$.

4.2.4. Common Envelope Phase (Dynamically Unstable Mass Transfer)

In COMPAS, dynamically unstable mass transfer occurs if $\zeta_* < \zeta_{\text{RL}}$ or if both stars simultaneously experience RLOF (see Section 4.2.6 for an exception for CHE binaries). Dynamically unstable mass transfer leads to a CE phase. In this phase, the binary is engulfed in a shared envelope and experiences gas drag, which causes a dynamical timescale inspiral (Paczynski 1976; Podsiadlowski 2001; Ivanova et al. 2013). CE events are thought to be especially relevant for the formation of DCOs in tight orbits (van den Heuvel 1976; Ivanova et al. 2013).

In COMPAS we follow the energy formalism in the form of the $\alpha_{\text{CE}}\text{-}\lambda$ prescription to estimate the post-CE orbital separation (Webbink 1984; de Kool 1990). In the $\alpha_{\text{CE}}\text{-}\lambda$ prescription the initial (pre-CE) binding energy of the donor is equated to the orbital energy reservoir

$$E_{\text{bind}} = \alpha_{\text{CE}} \times \Delta E_{\text{orb}}, \quad (33)$$

where ΔE_{orb} is the difference between the binary orbital energies before and after the CE phase and α_{CE} is a user-specified efficiency factor which parameterizes the fraction of

the orbital energy that is used to unbind the CE. The default value is $\alpha_{\text{CE}} = 1.0$.

The value of the binding energy of the envelope depends on the location of the envelope’s inner boundary and the sources of energy considered. We follow de Kool (1990) and express it in terms of a structure parameter λ :

$$E_{\text{bind}} = -\frac{GMM_{\text{env}}}{\lambda R}. \quad (34)$$

By default, COMPAS calculates λ using the “Nanjing lambda” prescription described by Xu & Li (2010a, 2010b), who provide fitting formulae for λ . Our implementation is identical to that of StarTrack (Dominik et al. 2012), including the several improvements they have made to the λ fits. This λ was computed in two different ways; one using only the gravitational binding energy (λ_g) and one that also includes the contribution of the full internal energy²³ (λ_b). COMPAS allows the user to specify a linear combination of the two parameters, $\lambda = \alpha_{\text{th}}\lambda_b + (1 - \alpha_{\text{th}})\lambda_g$. The default is $\alpha_{\text{th}} = 1$, i.e., including the full internal energy.

Alternative COMPAS options to estimate λ include prescriptions for calculating the envelope binding energy from Loveridge et al. (2011), fitting formulae to results from Kruckow et al. (2016) as implemented in Vigna-Gómez et al. (2018), and the option of using a fixed constant value.

We generalize Equations (33) and (34) to include the potential case of a double-core CE (Brown 1995) in the form

$$\begin{aligned} & -\frac{GM_1 M_{1,\text{env}}}{\lambda_1 R_1} - \frac{GM_2 M_{2,\text{env}}}{\lambda_2 R_2} \\ & = \alpha_{\text{CE}} \left(\frac{GM_1 M_2}{2a_{\text{pre-CE}}} - \frac{GM_{1,\text{core}} M_{2,\text{core}}}{2a_{\text{post-CE}}} \right), \end{aligned} \quad (35)$$

where $a_{\text{pre-CE}}$ and $a_{\text{post-CE}}$ are the separation before and after the CE phase, respectively; if one of the $M_{\text{env}} = 0$, then Equation (35) simplifies to the classic single-core energy formalism. Given the simplicity of our parameterization and the short (dynamical) timescales involved in a CE episode, we assume the phase is instantaneous. Equation (35) is used to predict the post-CE orbital separation given the pre-CE binary parameters and the post-CE component masses. The criterion for successful envelope ejection is $a_{\text{post-CE}} > R_1 + R_2$.

After a successful envelope ejection we always assume a circular orbit (see, e.g., Ivanova et al. 2013, for discussion of this). Besides a successful envelope ejection, which leads to a close binary, the CE phase can lead to immediate RLOF or a stellar merger (Section 4.2.5). Immediate RLOF implies that the stripped star or its companion is filling the respective post-CE Roche lobe. By default we allow these systems to engage in a mass transfer episode again, but we flag them so they can be considered as mergers in postprocessing.

A key uncertainty in CE evolution is the fate of Hertzsprung Gap donors (Belczynski et al. 2007). Such stars are not expected to have developed a steep density gradient between core and envelope (Taam & Sandquist 2000; Ivanova & Taam 2004) making it challenging to successfully eject the envelope. It is not clear whether Hertzsprung Gap donor stars can survive CE evolution, or whether instead, this would lead

²³ Note that this includes the thermal energy, radiation energy, ionization energy, and the dissociation energy of molecular hydrogen.

to a merger. In order to account for this uncertainty, we adopt two extreme models following Dominik et al. (2012). In the “optimistic” model, CE events involving a Hertzsprung Gap donor are treated in the same way as a more evolved star, determining the fate of the binary according to the energy budget (Equation (35)). In the “pessimistic” model, it is assumed that all CE events involving a Hertzsprung Gap donor result in a stellar merger. COMPAS keeps track of systems that experience an “optimistic” CE event, allowing the user to remove them in postprocessing. For the results presented in this paper we use the pessimistic model as the default assumption.

CE events initiated by RLOF from main-sequence donors always lead to a stellar merger. CE events with main-sequence accretors are treated as all other CE events by default, but may optionally always lead to stellar mergers.

There is no mass accretion onto the companion during a CE phase in the default COMPAS model. However, different optional CE accretion rate prescriptions exist for the case of NS accretors, including a user-defined fixed value, or following prescriptions from Osłowski et al. (2011) and MacLeod & Ramirez-Ruiz (2015), as described by Chattopadhyay et al. (2020). Mass accretion during a CE phase involving a BH companion has been investigated with COMPAS by van Son et al. (2020), who considered both accretion of a user-defined fixed fraction of the envelope and the Hoyle–Lyttleton (Hoyle & Lyttleton 1939) accretion rate within a CE following prescriptions from MacLeod & Ramirez-Ruiz (2015) based on Chevalier (1993).

4.2.5. Stellar Mergers

In COMPAS, two stars are assumed to merge when $a \leq (R_1 + R_2)$. This can occur following runaway stable mass transfer, if the envelope fails to be ejected during a CE phase, or if the direction of the supernova kick drastically shrinks the orbit. Currently, stellar mergers are flagged and the calculation is stopped, without subsequent evolution of the merger product.

4.2.6. Massive Overcontact Binaries

COMPAS includes prescriptions for CHE stars (Section 3.4), particularly in the context of CHE binaries (Riley et al. 2021). CHE binaries may arise from massive overcontact binaries, in which stars overflow their Roche lobes and share mass during the main sequence (see, e.g., Marchant et al. 2016). We therefore make an exception for them and relax our criteria for RLOF and mergers (see Riley et al. 2021). We consider that CHE stars in massive overcontact binaries can be filling their Roche lobe throughout the main sequence as long as they do not overflow the second (outer) Lagrangian point (Marchant et al. 2016; Riley et al. 2021). In our model, CHE binaries filling the second (outer) Lagrangian point lead to an imminent stellar merger.

4.3. Supernovae (Binary)

In COMPAS, supernovae lead to NS or BH formation²⁴ (the nomenclature is not necessarily associated with their observational signature). Supernovae occur on timescales much shorter than the timescales we resolve in COMPAS; we therefore assume that they are instantaneous events and that they could occur with uniform probability over all orbital phases.

Supernovae affect the orbit of a binary via instantaneous mass loss and natal kicks (Blaauw 1961; Hills 1983; Brandt &

Podsiadlowski 1995; Kalogera 1996; Tauris & Takens 1998; Hurley et al. 2002). We follow Appendix B of Pfahl et al. (2002b) to solve for the response of the orbital elements to a supernova. This prescription accounts for the natal kick, mass ejection, interaction with the companion during the supernova, and modification of the center-of-mass velocity. If the postsupernova eccentricity of the binary exceeds one, we label the binary as gravitationally unbound and cease further calculations (with an option to continue stellar evolution of the noncompact-object companion, if any).

4.4. Gravitational Radiation

Gravitational radiation releases energy and angular momentum from a binary, reducing both the orbital separation and the eccentricity. In COMPAS we only consider gravitational waves (GWs) after DCO formation and follow Peters (1964) to calculate the time to coalescence as

$$\tau_{\text{GW}} = \frac{15}{304} \frac{a_{\text{DCO}}^4 c^5}{G^3 M_1 M_2 M_{\text{tot}}} \kappa(e_{\text{DCO}}), \quad (36)$$

where a_{DCO} and e_{DCO} are the separation and eccentricity of the binary after the second supernova and $\kappa(e_{\text{DCO}})$ is a function of the eccentricity given by

$$\begin{aligned} \kappa(e_{\text{DCO}}) = & \left[\frac{(1 - e_{\text{DCO}}^2)}{e_{\text{DCO}}^{12/19}} \left(1 + \frac{121}{304} e_{\text{DCO}}^2 \right)^{-\frac{870}{2299}} \right]^4 \\ & \times \int_0^{e_{\text{DCO}}} \frac{e^{29/19} \left(1 + \frac{121}{304} e^2 \right)^{\frac{1181}{2299}}}{(1 - e^2)^3} de, \end{aligned} \quad (37)$$

where the integral is calculated as a Riemann sum over 10,000 linearly spaced eccentricity bins. For almost circular ($e < 0.01$) or very eccentric ($e > 0.99$) binaries, we use the approximations as presented following Equation (5.14) in Peters (1964). If τ_{GW} is less than a Hubble time, computed with $H_0 = 67.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Planck Collaboration et al. 2016), we classify the DCO as a merger candidate (but see Section 6.4 to see how we account for cosmological evolution).

4.5. Caveats to Binary Evolution in COMPAS

We currently do not include tidal evolution in COMPAS, although tides are believed to play a non-negligible role in massive binary evolution.

We also do not include magnetic braking evolution in COMPAS. Most massive stars are non-magnetic (Donati & Landstreet 2009). Moreover, unlike for low-mass stars, the magnetic braking assumption of negligible mass loss does not hold for radiatively driven stellar winds. We do not expect magnetic braking to play a significant role in the formation of NS and BH binaries.

COMPAS has not been adequately tested in the context of systems hosting white dwarfs.

5. Evolving a Population

Typically users wish to study particular outcomes of stellar or binary system evolution: BBHs that merge within the age of the universe, DNSs, X-ray binaries, etc. The initial attributes of these systems are usually not known a priori, so a population of

²⁴ With the exception of PISNes as discussed in Section 3.8.4.

systems is evolved with the expectation that some of the systems will evolve into systems of interest.

Each single star or binary system in a COMPAS simulation is described at ZAMS by the initial values of the salient attributes: mass and metallicity for single stars; component star masses, separation, eccentricity, and metallicity for binary systems. Users can specify values to be used for each of the initial attributes, or allow values to be drawn from specified distributions.

5.1. Sampling

The population of objects (stars or binary systems) synthesized by COMPAS is intended to be a representative sample from the full population of stars or binary systems in the universe. Sampling allows us to infer information about the full population based on results from the sampled subset—the sampled collection is used as a proxy for the actual population.

COMPAS provides functionality that allows users to sample the initial attributes of systems outside the COMPAS application and provide those initial attribute values to COMPAS via input grid files (Section 2.7.2). This makes it possible to interface with importance sampling via STROOPWAFEL (Broekgaarden et al. 2019), and in principle, allows for interfaces with other sampling tools such as Dartboard (Andrews et al. 2018) or emulation packages (e.g., Barrett et al. 2017).

COMPAS also includes a basic set of initial condition distributions within the main code. This enables Monte Carlo sampling in which users specify, for each star or binary system to be evolved, fixed values for initial attributes that should not be sampled, and COMPAS will sample the remainder using the distributions described below. Each star or binary system is then evolved to its final state, and the results recorded in output files (see Section 2.8).

Here we briefly describe the basic initial parameter distributions available in the core COMPAS code. For simplicity, we assume that the overall distribution of initial parameters is an outer product over independent parameter distributions, despite evidence to the contrary (Abt et al. 1990; Duchêne & Kraus 2013; Moe & Di Stefano 2017; Klencki et al. 2018).

5.2. Single Star/Primary Mass

The initial mass of a single star, or the primary star (the more massive star at ZAMS) in a binary system, $M_{1,i}$, is, unless specified by the user, determined by the initial mass function (IMF) being used. By default, COMPAS uses the Kroupa (2001) IMF, the distribution function of which is given by

$$p(M_{1,i}) \propto M_{1,i}^{-2.3} \quad (38)$$

above $0.5 M_{\odot}$, and $M_{1,i} \in [5.0, 150.0] M_{\odot}$. Other IMF functions (e.g., Salpeter 1955) are available as configurable options. For simplicity, we assume that the IMF is the same for all metallicities.

5.3. Mass Ratio and Secondary Mass

The mass of the secondary star (less massive at ZAMS), $M_{2,i}$, in a binary system being evolved by COMPAS is, unless specified by the user, determined by the mass ratio

$$q_i \equiv M_{2,i} / M_{1,i}, \quad (39)$$

where $q_i \in [0.01, 1.0]$.

In the default COMPAS model, the mass ratio is drawn from a flat distribution (Sana et al. 2012; Kobulnicky et al. 2014). Other distributions for mass ratio provided in COMPAS as options are the distributions described by Duquennoy & Mayor (1991a) and Sana et al. (2012) (Table S3). A minimum value for $M_{2,i}$ can be specified by the user.

5.4. Metallicity

The metallicity, Z_i , of a single star, or both component stars of a binary system, is, unless otherwise specified by the user, given by $Z_i = Z_{\odot} = 0.0142$ (Asplund et al. 2009).

Users can specify that metallicity be sampled for each single or binary star (for binary stars, both component stars use the same, sampled, value for metallicity), using a log-uniform distribution:

$$p(Z_i) \propto \frac{1}{Z_i}, \quad (40)$$

where $Z_i \in [0.0001, 0.03]$.

5.5. Semimajor Axis

The initial semimajor axis of a binary star, a_i , is, unless specified by the user, sampled independently of the masses using a log-uniform distribution:

$$p(a_i) \propto \frac{1}{a_i}, \quad (41)$$

where $a_i \in [0.01, 1000]$ au (Öpik 1924; Abt 1983).

Other distributions for semimajor axis provided in COMPAS as options are those described by Duquennoy & Mayor (1991b) and Sana et al. (2012). A custom method allows the user to specify parameters of the distribution.

5.6. Orbital Period

The orbital period can be specified for binary stars instead of the semimajor axis. The user can either specify a value for the orbital period, P_i , or that the value be sampled. If the value is sampled, a log-uniform distribution is used:

$$p(P_i) \propto \frac{1}{P_i}, \quad (42)$$

where $P_i \in [1.1, 1000]$ days.

5.7. Orbital Eccentricity

Unless otherwise specified by the user, the COMPAS default model assumes all binary stars are circular at birth (i.e., initial eccentricity $e_i = 0$). Other distributions for eccentricity provided in COMPAS as options are:

1. a flat distribution, $p(e_i) = 1$,
2. a thermal eccentricity distribution $p(e_i) = 2e_i$ (Heggie 1975),
3. the M35 distribution described by Geller et al. (2013),
4. the distribution described by Duquennoy & Mayor (1991a), and
5. the distribution described by Sana et al. (2012, Table S3).




Figure 9. Schematic display of the different times in the formation and evolution of a binary system that impact the time t_m at which a DCO (here a BHNS) system will merge in the history of the universe. The relevant timescales are: the moment the binary is formed at ZAMS from a gas cloud, t_{form} , the moment of DCO formation, t_{DCO} , the time at which the merger takes place, t_m , the time it takes the binary to evolve from ZAMS to the DCO system, $t_{\text{evolve}} = t_{\text{DCO}} - t_{\text{form}}$, the inspiral time, $t_{\text{inspiral}} = t_m - t_{\text{DCO}}$ and the time between binary formation at ZAMS and merger, $t_{\text{delay}} = t_m - t_{\text{form}}$. All mergers that occur within the horizon distance of GW detectors are potentially detectable today. Arrows indicate that those times that vary per binary system. Figure from Broekgaarden et al. (2021a) available at https://github.com/FloorBroekgaarden/BlackHole-NeutronStar/tree/main/plottingCode/Fig_1/variations.

5.8. Supernova Kicks

Stars that experience a supernova event may experience a momentum boost as a result of the explosion (see Section 3.7). Unless otherwise specified by the user, the attributes of these so-called natal kicks (i.e., magnitude and direction) are drawn from default distributions, depending upon the supernova type and expected remnant. The COMPAS model for natal kicks is described in detail in Section 3.9.

5.9. Stellar Rotation

Unless otherwise specified by the user, the COMPAS default model assumes that all stars are nonrotating at birth (i.e., the initial rotational velocity $v_{\text{rot}} = 0 \text{ km s}^{-1}$). Other rotation distributions provided in COMPAS as options are:

1. the initial, individual rotational velocities for the two stars can be provided separately,
2. the rotational velocity distribution from Hurley et al. (2000),
3. a rotational velocity distribution for O and B stars based on results from the VLT-FLAMES survey (Ramírez-Agudelo et al. 2013).

We describe the limited aspects of rotation currently modeled by COMPAS in Section 3.4. For stars in binaries, if CHE is enabled (default), COMPAS overwrites the drawn or zero rotational velocities of the binary components with the orbital velocity of the binary (i.e., tidal locking is assumed).

6. Postprocessing

Using binary population synthesis simulations to make predictions and calculations for astrophysical populations requires converting the data from the simulation into meaningful astrophysical quantities. The foremost example, for

compact objects, is calculating DCO formation and merger rates as a function of redshift and/or component masses. The strengths of the COMPAS suite include its publicly available postprocessing scripts where these calculations are performed. This section describes the methodology behind the main postprocessing scripts that are publicly available in COMPAS.²⁵

6.1. Recording Properties

As described in Section 2.8, during the simulation, COMPAS calculates and records properties of the stars and/or binary systems such as the ages, masses, stellar radii, effective temperatures, velocities, eccentricities, and separations. The user can specify which properties are recorded and when during the simulation they are reported. Examples include the option to save the properties at every time step (*detailed output*, an example is given in Section 7.1) or only printing the properties of the binary at important evolutionary stages of the binary such as CE episodes and SNe. A detailed description of how the output can be specified is given in the code documentation.²⁰

6.2. Selecting Binary Systems of Interest

During a simulation, COMPAS returns the recorded properties of all of the simulated binary systems. Further subselection of binaries of specific interest, such as BBHs that merge in a Hubble time or systems that experienced a CE event, is done by means of postprocessing. The “optimistic” and “pessimistic” CE selection (as mentioned in Section 4.2.4) is also performed in postprocessing. The subselection of systems of interest is typically done in COMPAS by “slicing” or “masking” the data, which is described in the publicly available jupyter notebooks.

²⁵ <https://github.com/TeamCOMPAS/COMPAS>

6.3. Converting to Yields per Star-forming Mass

A COMPAS simulation is typically performed by modeling only a fraction of the underlying stellar population by, for example, not simulating single stars and/or not drawing the simulated binaries from their full initial birth distributions (e.g., by only simulating stars with masses $\geq 5 M_\odot$). To obtain meaningful estimates for formation rates of the binaries of interest, the population synthesis simulation is typically re-normalized to a formation yield per unit star-forming mass. In this section we will often write this DCO formation yield as a function of: birth metallicity Z_i ,²⁶ delay time t_{delay} (i.e., the time between the formation and the merger of a binary, see Figure 9), and compact object masses M_1 , M_2 . Doing so, the DCO formation yield for a binary with metallicity Z_i , delay time t_{delay} and final compact object masses M_1 and M_2 can be written as

$$\frac{d^4 N_{\text{form}}}{dM_{\text{SFR}} dt_{\text{delay}} dM_1 dM_2}(Z_i, t_{\text{delay}}, M_1, M_2),$$

where N_{form} is the number of systems of interest that form and M_{SFR} is a unit of star-forming mass. To get the total yield $dN_{\text{form}}/dM_{\text{SFR}}$ this is marginalized over t_{delay} , M_1 and M_2 . This yield is typically computed with a Monte Carlo approach through COMPAS simulations. The subsequent conversion to merger and detection rates is then done in postprocessing, as described below.

6.4. Calculating Astrophysical Rates over the Cosmic History of Our Universe: The Case of Double Compact Object Mergers

These COMPAS suite scripts to calculate cosmological formation and merger rates of astrophysical events, often referred to as ‘‘cosmological integration,’’ are based on the work presented in Neijssel et al. (2019). We describe the method behind these postprocessing scripts for the example of calculating DCO merger rates over redshift below, similarly to Neijssel et al. (2019) and Broekgaarden et al. (2021a), but the idea can be easily generalized to other phenomena.

6.4.1. Double Compact Object Merger Rates for Ground-based GW Detectors

The DCO merger rate measured by a comoving observer in the source frame of the merger at a given merger time t_m (measured since the Big Bang) is given by

$$\begin{aligned} \mathcal{R}_m(t_m, M_1, M_2) &\equiv \frac{d^4 N_{\text{merger}}}{dt_m dV_c dM_1 dM_2}(t_m, M_1, M_2) \\ &= \int dZ_i \int_0^{t_m} dt_{\text{delay}} \text{SFRD}(Z_i, z(t_{\text{form}} = t_m - t_{\text{delay}})) \\ &\quad \times \frac{d^4 N_{\text{form}}}{dM_{\text{SFR}} dt_{\text{delay}} dM_1 dM_2}(Z_i, t_{\text{delay}}, M_1, M_2), \end{aligned} \quad (43)$$

where we convolve the yield with the the metallicity-specific star formation rate density $\text{SFRD}(Z_i, z(t_{\text{form}}))$, which is a function of birth metallicity Z_i and redshift z .²⁷ In Equation (43) V_c is the comoving volume and the relevant star formation rate is computed at a redshift corresponding to the formation time

$t_{\text{form}} = t_m - t_{\text{delay}}$. Delay times and metallicities are integrated over. We describe the merger rate in Equation (43) as an explicit function of M_1 and M_2 because the probability of a DCO merger detection depends on their values (see below), but the total merger rate is often computed by marginalizing over the DCO component masses.

We obtain the SFRD (Z_i, z) by multiplying the total star formation rate density (SFRD) with a metallicity probability density function

$$\begin{aligned} \text{SFRD}(Z_i, z_{\text{form}}) &\equiv \frac{d^3 M_{\text{SFR}}}{dt_s dV_c dZ_i}(z_{\text{form}}) \\ &= \underbrace{\frac{d^2 M_{\text{SFR}}}{dt_s dV_c}(z_{\text{form}})}_{\text{SFRD}} \times \underbrace{\frac{dP}{dZ_i}(z_{\text{form}})}_{\text{GSMF+MZR}}, \end{aligned} \quad (44)$$

where we used the short hand notation $z_{\text{form}} = z(t_{\text{form}})$. In the available COMPAS postprocessing scripts, the metallicity distribution function, dP/dZ_i , is typically described as a convolution between a galaxy mass function, the number density of galaxies per logarithmic mass bin (GSMF) and the mass–metallicity relation (MZR). This is discussed in more detail in the following sections and schematically shown in Figure 10.

In practice, the integral in Equation (43) is approximated by a Monte Carlo estimate, integrating over simulated metallicities. The metallicities in the integral limit in Equation (43) that fall outside of the simulated metallicity range can be either included in the edge bins (see, e.g., Broekgaarden et al. (2021a) for more details) or conservatively ignored (effectively curtailing dP/dZ_i to zero outside the range of simulated metallicities).

The DCO merger rate in Equation (43) is then converted to a local detection rate \mathcal{R}_{det} by integrating over the comoving volume and taking into account the probability P_{det} of detecting a gravitational-wave source (Section 6.6) using

$$\begin{aligned} \mathcal{R}_{\text{det}}(t_{\text{det}}, M_1, M_2) &\equiv \frac{d^3 N_{\text{det}}}{dt_{\text{det}} dM_1 dM_2} \\ &= \int dz \frac{dV_c}{dz} \frac{dt_m}{dt_{\text{det}}} \mathcal{R}_m(t_m) P_{\text{det}}(M_1, M_2, z(t_m)), \end{aligned} \quad (45)$$

where t_{det} is the time in the detector (i.e., the observer) frame, $P_{\text{det}}(M_1, M_2, z)$ is the probability of detecting a gravitational-wave signal from a binary with component masses M_1 , M_2 merging at redshift z , and $\frac{dt_m}{dt_{\text{det}}} = \frac{1}{(1+z)}$ and $\frac{dV_c}{dz}$ are given by, e.g., Hogg (1999). In practice, this integral over redshift can be computed as a Riemann sum over discrete redshift bins; see, e.g., previous work by Dominik et al. (2013, 2015), Belczynski et al. (2016b), Mandel & de Mink (2016), Eldridge et al. (2019b), Baibhav et al. (2019), Bavera et al. (2020), Chruslinska et al. (2019). The total detection rate per unit time can be obtained by further integrating over DCO masses.

6.5. Metallicity-specific Star Formation Rate Density Prescriptions

In COMPAS, the current publicly available postprocessing scripts use convolutions between analytical prescriptions for the SFRD, GSMF, and MZR (Equation (44)). The analytical equations are based on observations and simulations and can be flexibly adapted to the user’s preferences. We present below

²⁶ Before we used Z instead of Z_i for metallicity. In the remaining section we use Z_i to be consistent with the notation for other birth parameters.

²⁷ We use SFRD for the star formation rate density, and SFRD (Z_i, z) for the metallicity-specific star formation rate density.

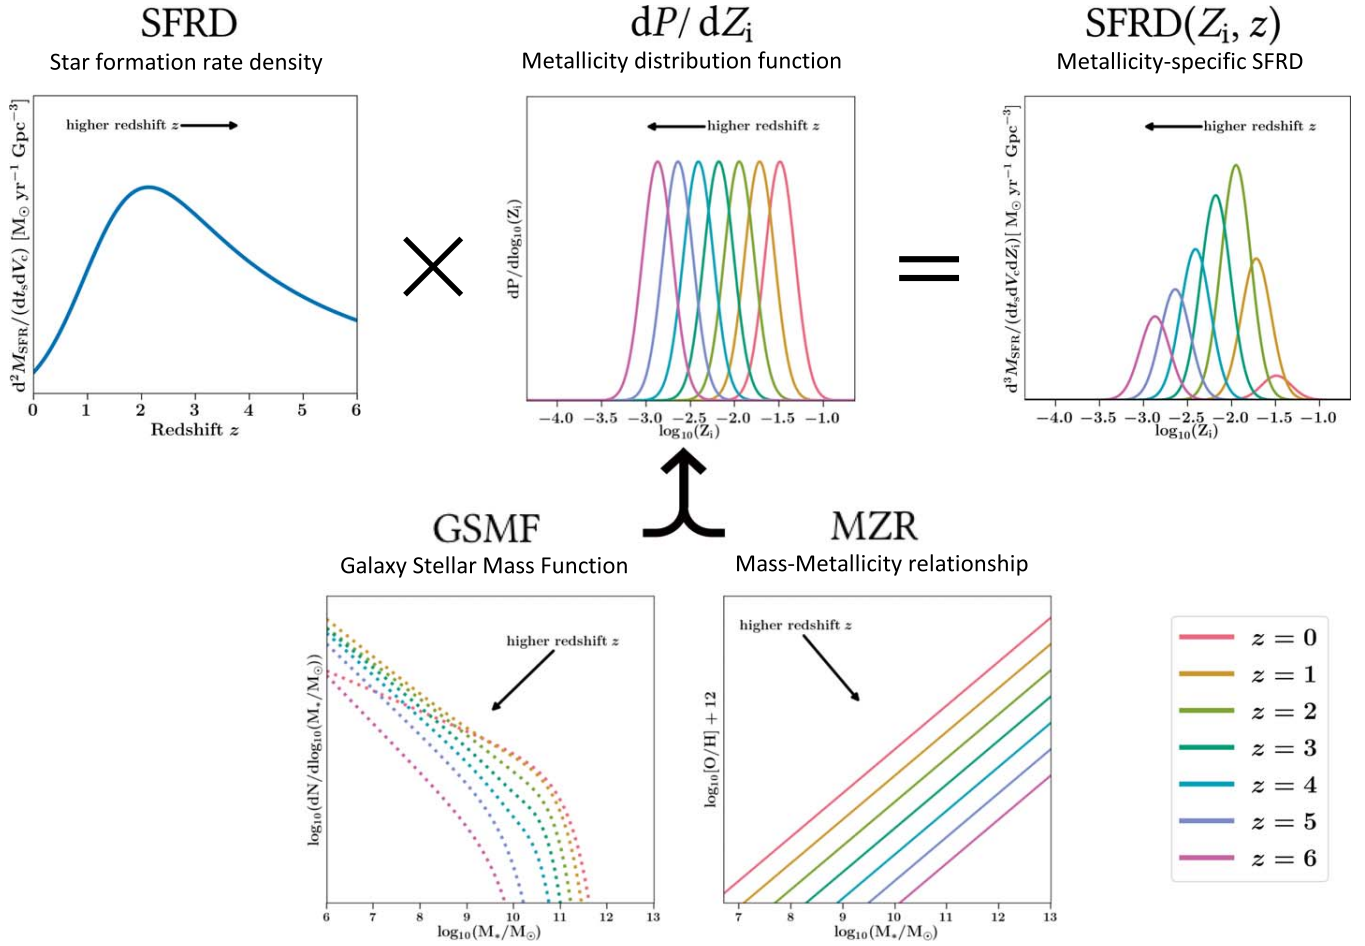


Figure 10. Schematic depiction of how models for the metallicity-specific star formation rate density, SFRD (Z_i, z), can be created in COMPAS by multiplying a star formation rate density with a metallicity probability distribution function, dP/dZ_i . The metallicity distribution function is typically constructed in the available COMPAS postprocessing scripts by convolving a galaxy stellar mass function with a mass–metallicity relationship. An exception is the “preferred” model from Neijssel et al. (2019), which uses a phenomenological model directly for the metallicity distribution function. The arrows in each subfigure indicate in which direction the distribution moves as redshift increases. Original figure from Broekgaarden et al. (2021a) available at https://github.com/FloorBroekgaarden/BlackHole-NeutronStar/tree/main/plottingCode/Fig_2.

the prescriptions existing in COMPAS postprocessing based on the work from Neijssel et al. (2019). An overview of the default available options for those prescriptions is given in Table 3.

6.5.1. Star Formation Rate Density Prescriptions

The COMPAS cosmic integration postprocessing scripts have different options for the SFRD. Several examples are given in the second column of Table 3 and shown in Figure 11. One of the SFRD prescriptions is based on the “preferred” model from Neijssel et al. (2019), which is calibrated to match the GW detections from the first two observing runs of LIGO and Virgo. Two other options are the Madau & Dickinson (2014) and Madau & Fragos (2017) SFRDs. All three SFRDs are given by the functional form

$$\frac{d^2 M_{\text{SFR}}}{dt_s dV_c}(z) = b_1 \frac{(1+z)^{b_2}}{1 + [(1+z)/b_3]^{b_4}} M_{\odot} \text{ yr}^{-1} \text{ Mpc}^{-3}. \quad (46)$$

where z is the redshift. The parameters $b_1, b_2, b_3,$ and b_4 in the equation for the SFRD are defined by: $b_1 = 0.01, b_2 = 2.77, b_3 = 2.9$ and $b_4 = 4.7$ for the preferred model in Neijssel et al. (2019); $b_1 = 0.015, b_2 = 2.7, b_3 = 2.9$ and $b_4 = 5.6$ for the Madau & Dickinson (2014) prescription (see their Equation

(15)) and $b_1 = 0.01, b_2 = 2.6, b_3 = 3.2$ and $b_4 = 6.2$ for the Madau & Fragos (2017) prescription (see their Equation (1)). Another available SFRD is the one from Strolger et al. (2004, see their Equation (1)), which uses an extinction-corrected model for SFRD described as a function of the universe’s age:

$$\begin{aligned} & \frac{d^2 M_{\text{SFR}}}{dt_s dV_c}(t_{\text{form}}) \\ &= c_1 \left[t_{\text{form}}^{c_2} \exp\left(-\frac{t_{\text{form}}}{c_3}\right) + c_4 \exp\left(c_4 \frac{t_{\text{form}} - t_0}{c_3}\right) \right] \frac{M_{\odot}}{\text{yr Mpc}^3}, \end{aligned} \quad (47)$$

with $c_1 = 0.182, c_2 = 1.26, c_3 = 1.865, c_4 = 0.071,$ and $t_0 = 13.47 \text{ Gyr}$.

6.5.2. Metallicity Distribution Function over Redshift (GSMF + MZR)

The chemical evolution of star formation in our universe is described in the COMPAS postprocessing scripts by the metallicity probability density, dP/dZ_i , which is a function of redshift. The cosmic integration postprocessing scripts offer several options for the metallicity density function, several examples are shown in Figure 12.

Table 3

Examples of the By-default Available Options to Construct a Metallicity-specific Star Formation Density Model SFRD(Z_i, z) in the Publicly Available Cosmic Integration Postprocessing Scripts of COMPAS

xyz index	SFRD [x]	GSMF [y]	MZR [z]
0 (default)	“Preferred” phenomenological model from Neijssel et al. (2019)		
1	Madau & Dickinson (2014)	Panter et al. (2004)	Langer & Norman (2006)
2	Strolger et al. (2004)	Furlong et al. (2015) single Schechter	Langer & Norman (2006) + offset
3	Madau & Fragos (2017)	Furlong et al. (2015) double Schechter	Ma et al. (2016a)

Note. A SFRD (Z_i, z) model can be obtained by combining a star formation rate (SFR) with a galaxy stellar mass function (GSMF) and mass–metallicity relation (MZR). See Sections 6.4 for more details. The labels 0, 1, 2, 3 are solely used in the figures to refer to these models. The code is flexible to easily adopt a user-specified prescription for the SFRD, GSMF, or MZR.

The preferred model from Neijssel et al. (2019) ($yz = 00$ in Table 3) uses a phenomenological model, which defines the metallicity density function as a symmetric log-normal distribution:

$$\frac{dP}{dZ_i}(z) = \frac{1}{Z_i \sigma_1 \sqrt{2\pi}} \exp\left(-\frac{(\ln(Z_i) - \mu_1(z))^2}{2\sigma_1^2}\right), \quad (48)$$

where σ_1 is the standard deviation in $\ln(Z)$ -space and $\mu_1(z)$ is the redshift-dependent mean in $\ln(Z)$ -space. Neijssel et al. (2019) use a redshift independent $\sigma_1 = 0.39$ and mean $\mu_1 = \langle \ln(Z) \rangle$ defined by

$$\langle Z(z) \rangle = \exp(\mu_1 + \sigma_1^2/2) = Z_0 10^{e_1 z}, \quad (49)$$

where Z_0 is the mean metallicity at redshift 0. In the preferred phenomenological model, $Z_0 = 0.035$ and $e_1 = -0.23$.

This parameterization of the mean (Equation (49)) follows the work by Langer & Norman (2006). Observational evidence suggests the metallicity distributions are likely not symmetric in log-metallicity (e.g., Langer & Norman 2006; Chruslinska et al. 2019; Boco et al. 2021). The other prescriptions for the metallicity distribution function, which are available in the COMPAS postprocessing, are therefore asymmetric convolutions of a GSMF and an MZR.

6.5.3. Galaxy Stellar-mass Function Prescriptions

The available GSMFs in the COMPAS postprocessing scripts are all based on observations of luminosity distributions of galaxies, which are converted to galaxy mass distributions based on a luminosity–mass relation. The options are listed in the third column of Table 3 and in Figure 13. More details are provided in Appendix A3 of Neijssel et al. (2019).

The GSMFs 1 and 2 use a functional form of a single Schechter function given by

$$\Phi_{M_*}(z) dM_* = \phi_1(z) \left(\frac{M_*}{M_c(z)}\right)^{-\psi_1(z)} \exp\left(\frac{-M_*}{M_c(z)}\right) dM_*, \quad (50)$$

where M_* is the galaxy stellar mass, ϕ_1 is the overall normalization, ψ_1 is the parameter for the slope of the GSMF for $M_* \leq M_c$ and M_c is the cutoff where the GSMF moves from a power-law into an exponential drop off (see Figure 13). Panter et al. (2004, Equation (1)) use a z -independent single Schechter function with $\phi_1 = 7.8 \times 10^{-3} \text{ Mpc}^{-3}$, $\psi_1 = +1.16$ and $M_c = 7.64 \times 10^{10} M_\odot$. The GSMF option “Furlong et al. (2015) single” ($y = 2$ in Table 3), uses a linear fit by Neijssel et al. (2019) to the tabulated redshift-dependent values for ϕ_1 , ψ_1 , and M_c . The GSMF option “Furlong et al. (2015) double”

($y = 3$ in Table 3) is instead based on a double Schechter function given by

$$\Phi_{M_*}(z) dM_* = \exp\left(\frac{-M_*}{M_c(z)}\right) \times \left[\phi_1(z) \left(\frac{M_*}{M_c(z)}\right)^{-\psi_1(z)} + \phi_2(z) \left(\frac{M_*}{M_c(z)}\right)^{-\psi_2(z)} \right] dM_*, \quad (51)$$

which is fitted in a similar way based on tabulated data.

6.5.4. Mass–Metallicity Relation Prescriptions

The default MZR in the COMPAS postprocessing scripts are analytical formulas for the mapping between galaxy stellar mass and the metallicity of star formation. Outside of the phenomenological model by Neijssel et al. (2019), there are three different MZR available by default in COMPAS.

The MZR 1 and 2 in Table 3 are based on Langer & Norman (2006), who derive a MZR based on Savaglio et al. (2005) observations of 56 galaxies in the Gemini Deep Deep and Canada–France Redshift Survey. They obtain a bisector fit, which is then simplified by Langer & Norman (2006) resulting in the MZR

$$\frac{M}{M_*} = \left(\frac{Z_i}{Z_\odot}\right)^2, \quad (52)$$

with $M_* = 7.64 \times 10^{10} M_\odot$ from Panter et al. (2004), and a mean metallicity that scales with redshift as Equation (49) with the COMPAS default values of $Z_0 = 0.035$, $\alpha = 0.3$. Combining these two relations results in the MZR by Langer & Norman (2006) given in Table 3 with the label $z = 1$. However, since this MZR has an offset from the Savaglio et al. (2005) bisector fit, Neijssel et al. (2019) also added a second default MZR prescription based on Langer & Norman (2006) with a fixed offset in metallicity to match the results by Savaglio et al. (2005).

The MZR from Ma et al. (2016a) is based on theoretical models of cosmological simulations combined with population synthesis simulations. Their MZR is given by

$$\begin{aligned} & \log_{10}(Z_{\text{gas}}/Z_\odot) \\ &= 0.35 \left[\log_{10}\left(\frac{M_*}{M_\odot}\right) - 10 \right] + 0.93 \exp(-0.43z) - 1.05, \end{aligned} \quad (53)$$

with Z_{gas} the metallicity of the star-forming gas (here Z_i).




Figure 11. The star formation rate density as a function of redshift for the four default available options in the cosmic integration postprocessing scripts of COMPAS.




Figure 12. Three examples of metallicity probability distribution functions (dP/dZ_i) available in the postprocessing scripts of COMPAS. For each of the prescriptions the metallicity distribution is shown for four different redshifts. Gray areas show Z_i values that fall outside of the metallicity range of SSE tracks from Hurley et al. (2000).

The three MZR options available by default in COMPAS are given in Figure 14.

6.6. Gravitational-wave Selection Effects

Whether a DCO merger is detectable by a GW interferometer depends on its distance (i.e., redshift), orientation, inclination, and source component masses M_1 and M_2 . The detectability by a network is approximated by checking whether the source signal-to-noise ratio (S/N) in a single detector surpasses a predefined threshold. COMPAS contains

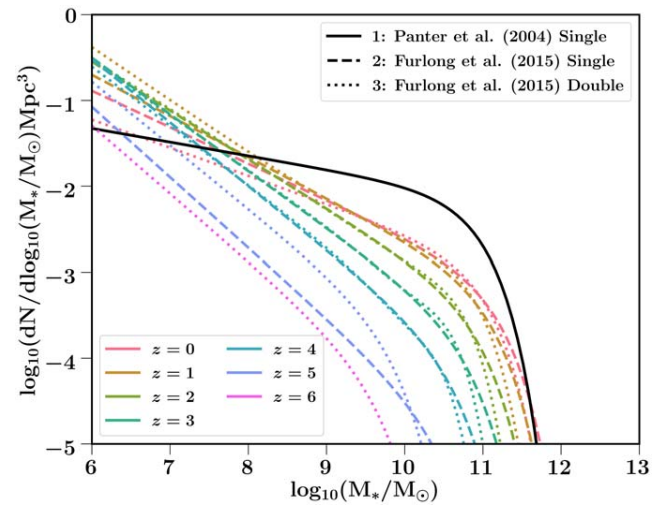


Figure 13. The galaxy stellar mass functions (GSMFs) for the three default available options in the cosmic integration postprocessing scripts of COMPAS (see also Table 3). The GSMFs are shown at seven different redshift values, except for the Panter et al. GSMF, which is not redshift-dependent.




Figure 14. The mass–metallicity relations for the three default available options in the cosmic integration postprocessing scripts of COMPAS. The MZRs are shown at different redshift values.

publicly available postprocessing scripts that calculate the detection probability of a gravitational-wave source, based on the method described in Barrett et al. (2018). Source component spins and any residual in-band eccentricity are not currently included in these detection probability estimates.

Typically a user sets in the COMPAS postprocessing scripts a S/N threshold of $S/N=8$ for a single ground-based GW detector (Finn & Chernoff 1993), such that sources with a higher S/N are detectable, but this can be changed to the user-specified value. The S/Ns of the DCO mergers are calculated by computing the source waveforms using a user-chosen waveform model included in the LAL software suite (LIGO Scientific Collaboration 2018), such as IMRPHEMOPV2 (Hannam et al. 2014; Husa et al. 2016; Khan et al. 2016) and SEOBNRV3 (Pan et al. 2014; Babak et al. 2017). We marginalize over the sky localization and source orientation

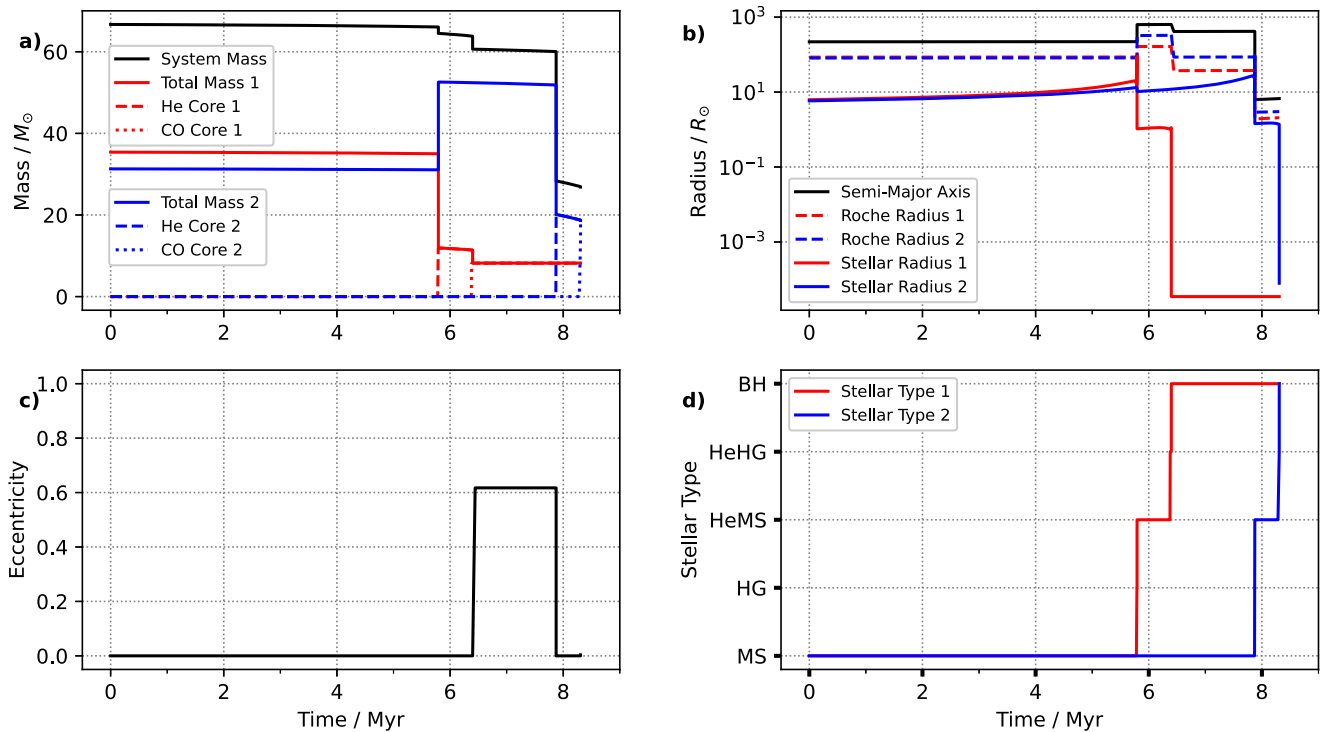


Figure 15. Plots produced using detailed COMPAS output of the evolution of a with initial masses $M_{\text{ZAMS},1} = 35.4 M_{\odot}$, $M_{\text{ZAMS},2} = 29.3 M_{\odot}$ on a circular orbit with semimajor axis $a = 1.02 \text{ AU}$ at metallicity $Z = 0.001$. These initial parameters produce a GW151226-like BBH with component masses 15.9 and $8.2 M_{\odot}$ and a coalescence time of 280 Myr (GW151226 had measured component masses of $14.2^{+8.3}_{-3.7}$ and $7.5^{+2.3}_{-2.3} M_{\odot}$, Abbott et al. 2016b). Panels (a)–(d) plot respectively the time evolution of the mass parameters (masses of both stars and their He and CO cores, and the total mass of the binary), the radius parameters (stellar radii and Roche radii for both stars, and the semimajor axis), eccentricity, and stellar types (as described in Table 2). Red, blue, and black lines denote the primary star, secondary star, and the binary, respectively.)

of the binary using the antenna pattern function from Finn & Chernoff (1993). The detector sensitivity can be chosen by the user. Available options include the sensitivity of a LIGO instrument at design sensitivity and O1, O2, and O3 configurations (Aasi et al. 2015; Abbott et al. 2016a), as well as the third-generation Einstein Telescope detector (Hild et al. 2011).

7. Usage Examples

In this section, we show a few practical examples of COMPAS usage. Although COMPAS has mainly been applied to large population studies (see examples in Section 1), it is often useful to visualize the full evolutionary path of a given isolated binary, especially when trying to reproduce a specific system. We show an example in Section 7.1, with the caveat that the approximate treatment of rapid population synthesis cannot match the precision of detailed stellar codes for individual systems. In Section 7.2 we show a more typical application of COMPAS postprocessing tools to predict a distribution over a population, in this case, the chirp mass distribution of detectable BBH mergers.

7.1. Detailed Evolution of a Binary

We provide an example plot of the COMPAS detailed output in Figure 15, which records the detailed evolution of the progenitor to a GW151226-like BBH (Abbott et al. 2016b; Stevenson et al. 2017). The code to reproduce this binary and the detailed output of any binary is available at <https://github.com/TeamCOMPAS/COMPAS>.

We now describe the evolution of this example binary system. Descriptions of the stellar types referenced below can be found in Table 2.

1. Figure 15(b) shows that the primary star (red line) exceeds its Roche lobe at 5.8 Myr and initiates mass transfer as it expands rapidly once it evolves off the main sequence. This dynamically stable mass transfer episode is nonconservative, as reflected by the concurrent dip in total mass (black curve in Figure 15(a)). This causes the semimajor axis to nearly triple, despite the larger initial mass of the donor (see Equation (25)). Figure 15(a) shows that the primary loses $23 M_{\odot}$ (corresponding to its hydrogen envelope), of which $21 M_{\odot}$ is accreted by the secondary. The primary emerges as a stripped Helium star (HeMS), as shown in Figure 15(d), with total mass equal to its He-core mass (Figure 15(a)).
2. After another 0.6 Myr , the primary collapses into a BH in a core-collapse supernovae (CCSNe), ejecting $3 M_{\odot}$ in the process (Figure 15(a)). Its natal kick induces an orbital eccentricity of 0.62 (Figure 15(c)), but due to a fortuitous combination of kick magnitude and direction, the orbital semimajor axis shrinks by $\sim 1/3$ (Figure 15(b)).
3. The secondary evolves off the main sequence at 7.8 Myr (Figure 15(d)). It too expands and exceeds its Roche lobe shortly thereafter, triggering dynamically unstable mass transfer back onto the primary (now a BH, Figure 15(b)). The binary enters a CE, characterized by orbital tightening by several orders of magnitude (Figure 15(b)) on the dynamical timescale of the donor. In the default COMPAS




Figure 16. Distribution of the detection rate of BBH mergers over chirp mass, assuming LIGO O3 sensitivity, an S/N threshold of 8, and the default metallicity-specific star formation history model from Neijssel et al. (2019). The gray bars show a binned histogram. The thick gray smooth solid line is a kernel density estimate of the distribution, and the shaded region shows the one and two σ sampling uncertainties estimated through bootstrapping.

model, CE phases are assumed to completely circularise the binary, removing the eccentricity imparted by the first SNe (Figure 15(c)).

4. The secondary is stripped by the CE episode and continues its evolution as a stripped helium star (HeMS) (Figure 15(d)), collapsing into a BH at 8.3 Myr. The complete fallback in the second supernova results in no natal kick and a consequently small post-SNe eccentricity (Figure 15(c)).

7.2. Chirp Mass Distribution of LIGO BBHs

We show an example application of COMPAS population postprocessing tools to predict the chirp mass distribution of BBHs detected by gravitational-wave detectors.

We evolve ten million binaries with COMPAS assuming the default model. Metallicities are sampled log-uniformly in the range $\log(Z) \in [0.0001, 0.03]$. A Monte Carlo integral over a smooth distribution of metallicities (see Section 6.4) avoids the metallicity binning artefacts encountered when using a discrete metallicity grid, as discussed, e.g., by Dominik et al. (2015). We compute the merger rate distribution under the assumption of a phenomenological, metallicity-specific star formation history (see Section 6.5) detailed in Neijssel et al. 2019. We apply gravitational-wave selection effects (see Section 6.6) for a LIGO detector network operating at O3 sensitivity and a S/N detection threshold of 8. Figure 16 shows the chirp mass distribution of detectable BBHs.

The script used to make this plot, which makes use of COMPAS’s postprocessing classes, can be found at <https://github.com/TeamCOMPAS/COMPAS>.

8. Conclusions

We described COMPAS, a public rapid population synthesis code. COMPAS v 02.21.00 evolves ~ 100 binaries per second on a modern laptop computer. Given its parallel structure (jobs can be readily split across multiple cores with no need for communication until results are re-combined), a population of a billion binaries can be evolved in 24 hours on a modest 128 core cluster. COMPAS is designed to make it easy to specify desired parameterized prescriptions or introduce new models for various stages of stellar and binary evolution. Together with its postprocessing toolbox, COMPAS is thus well suited for inference studies on observed stellar binary populations (e.g., Barrett et al. 2018).

The following are some of the planned enhancements to COMPAS that we hope to include in future versions:

1. COMPAS currently relies on Hurley et al. (2000) models for single stellar evolution. In order to evaluate the impact of uncertainties in these models, we plan to incorporate single stellar evolution tracks interpolated from other stellar evolution codes with METISSE (Method of Interpolation for Single Star Evolution, Agrawal et al. 2020).
2. Extend COMPAS to more accurately treat low-mass stars. While the code can evolve binaries with low-mass components, a number of features, ranging from white dwarf novae to magnetic braking, are either not included or insufficiently tested.
3. Include a proper treatment of stellar mergers, allowing the future evolution of merger products to be tracked.
4. Update and re-activate the treatment of tidal interactions, including tidal synchronization and circularization, to include the latest models (e.g., Vick & Lai 2020).

COMPAS is a public code, and we encourage the community to use it and, should they wish, to become involved in its development. In particular, any defects or enhancement requests can be brought to our attention via the github issue tracker²⁸ or by e-mail, compas-user@googlegroups.com. COMPAS will be published in the *Journal of Open Source Software* (Team Compas et al. 2021).

The authors thank Ben Bradnick, Isobel Romero-Shaw, and Rajath Sathyaprakash for past contributions to the code, and Simone Bavera, Chris Belczynski, Christopher Berry, Jan Eldridge, Tassos Fragos, David Hendriks, Jarrod Hurley, Vicky Kalogera, Morgan MacLeod, Pablo Marchant, Javier Morán Fraile, Philipp Podsiadlowski, Carl Rodriguez, Dorottya Szécsi, and Michael Zevin for discussions and advice. Multiple authors are supported by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), through project No. CE170100004. Multiple authors were funded in part by the National Science Foundation under grant No. (NSF grant No. 2009131), the Netherlands Organization for Scientific Research (NWO) as part of the Vidi research program BinWaves with project No. 639.042.728 and by the European Union’s Horizon 2020 research and innovation program from the European Research Council (ERC, grant agreement No. 715063). F.S.B. is supported in part by the Prins Bernard Cultuurfonds studiebeurs. I.M. is a recipient of an Australian Research Council Future Fellowship

²⁸ <https://github.com/TeamCOMPAS/COMPAS/issues>

(FT190100574). A.V.G. acknowledges funding support by the Danish National Research Foundation (DNRF132).

This research has made use of NASA's Astrophysics Data System Bibliographic Services.²⁹













Facilities: Some of the results in this manuscript were obtained using the following computing facilities: FAS Research Computing, Harvard University, and the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government.

Software: COMPAS is written in C++ and we acknowledge the use of the GNU C++ compiler, GNU scientific library (gsl), the BOOST C++ library, and the HDF5 C++ library from <http://www.gnu.org/software/gsl/> (Galassi et al. 2002). The COMPAS suite makes use of Python from the Python Software Foundation. Python Language Reference available at <http://www.python.org> (van Rossum 1995). In addition, the COMPAS suite makes use of the python packages Astropy³⁰ (Astropy Collaboration et al. 2013, 2018), hdf5³¹ (Collette 2013), the IPython,³² and Jupyter notebook package³³ (Pérez & Granger 2007; Kluyver et al. 2016), Matplotlib³⁴ (Hunter 2007), NumPy³⁵ (Harris et al. 2020), SciPy³⁶ (Virtanen et al. 2020), Seaborn (Waskom et al. 2020). The COMPAS postprocessing code for detection probability currently makes use of precomputed results from the LALSuite toolkit (LIGO Scientific Collaboration 2018), such as the IMRPHENOMPV2 waveform (Hannam et al. 2014; Husa et al. 2016; Khan et al. 2016).

Data Availability

We encourage the community to make results obtained with COMPAS publicly available at <https://zenodo.org/communities/compas/>. The scripts to reproduce the data and plots for all figures in this manuscript using v02.21.00 of COMPAS are provided in the corresponding directories at <https://github.com/Team-COMPAS/COMPAS>.

ORCID iDs

Jeff Riley  <https://orcid.org/0000-0003-1530-2557>
 Poojan Agrawal  <https://orcid.org/0000-0002-1135-984X>
 Jim W. Barrett  <https://orcid.org/0000-0002-0151-0753>
 Kristan N. K. Boyett  <https://orcid.org/0000-0003-4109-304X>
 Floor S. Broekgaarden  <https://orcid.org/0000-0002-4421-4962>
 Debatri Chattopadhyay  <https://orcid.org/0000-0001-5867-5033>
 Fabian Gittins  <https://orcid.org/0000-0002-9439-7701>
 Ryosuke Hirai  <https://orcid.org/0000-0002-8032-8174>
 George Howitt  <https://orcid.org/0000-0002-2593-0189>
 Stephen Justham  <https://orcid.org/0000-0001-7969-1569>
 Floris Kummer  <https://orcid.org/0000-0002-6056-3070>
 Mike Y. M. Lau  <https://orcid.org/0000-0002-6592-2036>

²⁹ <http://adsabs.harvard.edu/>

³⁰ <http://www.astropy.org>

³¹ <https://docs.h5py.org/en/stable/>



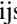
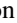





³² <http://ipython.org>

³³ <https://jupyter.org/>

³⁴ <http://www.matplotlib.org>

³⁵ <http://www.NumPy.org/>

³⁶ <https://www.scipy.org>

Ilya Mandel  <https://orcid.org/0000-0002-6134-8946>
 Selma E. de Mink  <https://orcid.org/0000-0001-9336-2825>
 Coenraad Neijssel  <https://orcid.org/0000-0001-9318-5259>
 Lieke van Son  <https://orcid.org/0000-0001-5484-4987>
 Simon Stevenson  <https://orcid.org/0000-0002-6100-537X>
 Alejandro Vigna-Gómez  <https://orcid.org/0000-0003-1817-3586>
 Serena Vinciguerra  <https://orcid.org/0000-0003-3068-6974>
 Tom Wagg  <https://orcid.org/0000-0001-6147-5761>
 Reinhold Willcox  <https://orcid.org/0000-0003-0674-9453>

References

- Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, *CQGra*, 32, 074001
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *LRR*, 19, 1
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *PhRvL*, 116, 241103
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, *PhRvX*, 9, 031040
 Abt, H. A. 1983, *ARA&A*, 21, 343
 Abt, H. A., Gomez, A. E., & Levy, S. G. 1990, *ApJS*, 74, 551
 Agrawal, P., Hurley, J., Stevenson, S., Szécsi, D., & Flynn, C. 2020, *MNRAS*, 497, 4549
 Akmal, A., Pandharipande, V., & Ravenhall, D. 1998, *PhRvC*, 58, 1804
 Almeida, L., Sana, H., Taylor, W., et al. 2017, *A&A*, 598, A84
 Andrews, J. J., Farr, W. M., Kalogera, V., & Willems, B. 2015, *ApJ*, 801, 32
 Andrews, J. J., Zezas, A., & Fragos, T. 2018, *ApJS*, 237, 1
 Arzoumanian, Z., Chernoff, D. F., & Cordes, J. M. 2002, *ApJ*, 568, 289
 Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, *ARA&A*, 47, 481
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
 Atri, P., Miller-Jones, J. C. A., Bahramian, A., et al. 2019, *MNRAS*, 489, 3116
 Babak, S., Taracchini, A., & Buonanno, A. 2017, *PhRvD*, 95, 024010
 Baibhav, V., Berti, E., Gerosa, D., et al. 2019, *PhRvD*, 100, 064060
 Banerjee, S., Belczynski, K., Fryer, C. L., et al. 2020, *A&A*, 639, A41
 Barkat, Z., Rakavy, G., & Sack, N. 1967, *PhRvL*, 18, 379
 Barrett, J. W., Gaebel, S. M., Neijssel, C. J., et al. 2018, *MNRAS*, 477, 4685
 Barrett, J. W., Mandel, I., Neijssel, C. J., Stevenson, S., & Vigna-Gomez, A. 2017, in IAU Symp. 325, *Astroinformatics* (Cambridge: Cambridge Univ. Press), 46
 Bavera, S. S., Fragos, T., Qin, Y., et al. 2020, *A&A*, 635, A97
 Beasar, E. R., & Davies, B. 2018, *MNRAS*, 475, 55
 Belczynski, K., Bulik, T., Fryer, C. L., et al. 2010, *ApJ*, 714, 1217
 Belczynski, K., Heger, A., Gladysz, W., et al. 2016a, *A&A*, 594, A97
 Belczynski, K., Hirschi, R., Kaiser, E. A., et al. 2020a, *ApJ*, 890, 113
 Belczynski, K., Kalogera, V., & Bulik, T. 2002, *ApJ*, 572, 407
 Belczynski, K., Kalogera, V., Rasio, F. A., et al. 2008, *ApJS*, 174, 223
 Belczynski, K., Kalogera, V., Rasio, F. A., Taam, R. E., & Bulik, T. 2007, *ApJ*, 662, 504
 Belczynski, K., Klencki, J., Fields, C. E., et al. 2020b, *A&A*, 636, A104
 Belczynski, K., Repetto, S., Holz, D. E., et al. 2016b, *ApJ*, 819, 108
 Belczynski, K., Romagnolo, A., Olejak, A., et al. 2021, arXiv:2108.10885
 Beniamini, P., & Piran, T. 2016, *MNRAS*, 456, 4089
 Berger, E. 2014, *ARA&A*, 52, 43
 Bhattacharya, D., & van den Heuvel, E. P. J. 1991, *PhR*, 203, 1
 Bisnovatyi-Kogan, G. S. 1993, *A&AT*, 3, 287
 Blaauw, A. 1961, *BAN*, 15, 265
 Boco, L., Lapi, A., Chruslinska, M., et al. 2021, *ApJ*, 907, 110
 Bondi, H., & Hoyle, F. 1944, *MNRAS*, 104, 273
 Bowman, D. M. 2020, *FRASS*, 7, 70
 Brandt, N., & Podsiadlowski, P. 1995, *MNRAS*, 274, 461
 Bray, J. C., & Eldridge, J. J. 2016, *MNRAS*, 461, 3747
 Bray, J. C., & Eldridge, J. J. 2018, *MNRAS*, 480, 5657
 Breivik, K., Coughlin, S., Zevin, M., et al. 2020, *ApJ*, 898, 71
 Briskin, W. F., Benson, J. M., Goss, W. M., & Thorsett, S. E. 2002, *ApJ*, 571, 906
 Broekgaarden, F. S., & Berger, E. 2021, *ApJL*, 920, L13
 Broekgaarden, F. S., Berger, E., Neijssel, C. J., et al. 2021a, *MNRAS*, 508, 5028
 Broekgaarden, F. S., Berger, E., Stevenson, S., et al. 2021b, arXiv:2112.05763
 Broekgaarden, F. S., Justham, S., de Mink, S. E., et al. 2019, *MNRAS*, 490, 5228
 Brookshaw, L., & Tavani, M. 1993, *ApJ*, 410, 719
 Brown, G. E. 1995, *ApJ*, 440, 270

- Burrows, A., & Hayes, J. 1996, *PhRvL*, **76**, 352
- Burrows, A., & Vartanyan, D. 2021, *Natur*, **589**, 29
- Castor, J. I., Abbott, D. C., & Klein, R. I. 1975, *ApJ*, **195**, 157
- Chattopadhyay, D., Stevenson, S., Hurley, J. R., Bailes, M., & Broekgaarden, F. 2021, *MNRAS*, **504**, 3682
- Chattopadhyay, D., Stevenson, S., Hurley, J. R., Rossi, L. J., & Flynn, C. 2020, *MNRAS*, **494**, 1587
- Chevalier, R. A. 1993, *ApJL*, **411**, L33
- Chini, R., Hoffmeister, V., Nasserri, A., Stahl, O., & Zinnecker, H. 2012, *MNRAS*, **424**, 1925
- Chiosi, C., & Maeder, A. 1986, *ARA&A*, **24**, 329
- Chruslinska, M., Nelemans, G., & Belczynski, K. 2019, *MNRAS*, **482**, 5012
- Collette, A. 2013, Python and HDF5 (Sebastopol, CA: O'Reilly & Associates)
- Conroy, C., & Kratter, K. M. 2012, *ApJ*, **755**, 123
- de Jager, C., Nieuwenhuijzen, H., & van der Hucht, K. A. 1988, *A&AS*, **72**, 259
- de Kool, M. 1990, *ApJ*, **358**, 189
- De Minkde Mink, S., Langer, N., Izzard, R., Sana, H., & de Koter, A. 2013, *ApJ*, **764**, 166
- Delgado, A. J., & Thomas, H. C. 1981, *A&A*, **96**, 142
- Doherty, C. L., Gil-Pons, P., Siess, L., & Lattanzio, J. C. 2017, *PASA*, **34**, e056
- Dominik, M., Belczynski, K., Fryer, C., et al. 2012, *ApJ*, **759**, 52
- Dominik, M., Belczynski, K., Fryer, C., et al. 2013, *ApJ*, **779**, 72
- Dominik, M., Berti, E., O'Shaughnessy, R., et al. 2015, *ApJ*, **806**, 263
- Donati, J. F., & Landstreet, J. D. 2009, *ARA&A*, **47**, 333
- Dosopoulou, F., & Kalogera, V. 2016, *ApJ*, **825**, 71
- Dray, L. M., & Tout, C. A. 2003, *MNRAS*, **341**, 299
- Duchêne, G., & Kraus, A. 2013, *ARA&A*, **51**, 269
- Dufton, P. L., Langer, N., Dunstall, P. R., et al. 2013, *A&A*, **550**, A109
- Duquennoy, A., & Mayor, M. 1991a, *A&A*, **500**, 337
- Duquennoy, A., & Mayor, M. 1991b, *A&A*, **248**, 485
- Eggleton, P. P. 1983, *ApJ*, **268**, 368
- Eggleton, P. P., Fitchett, M. J., & Tout, C. A. 1989, *ApJ*, **347**, 998
- El-Badry, K., & Rix, H.-W. 2018, *MNRAS*, **480**, 4884
- Eldridge, J. J., Fraser, M., Smartt, S. J., Maund, J. R., & Crockett, R. M. 2013, *MNRAS*, **436**, 774
- Eldridge, J. J., Guo, N. Y., Rodrigues, N., Stanway, E. R., & Xiao, L. 2019a, *PASA*, **36**, e041
- Eldridge, J. J., Stanway, E. R., & Tang, P. N. 2019b, *MNRAS*, **482**, 870
- Eldridge, J. J., Stanway, E. R., Xiao, L., et al. 2017, *PASA*, **34**, e058
- Eldridge, J. J., Xiao, L., Stanway, E. R., Rodrigues, N., & Guo, N. Y. 2018, *PASA*, **35**, 49
- Ertl, T., Janka, H. T., Woosley, S. E., Sukhbold, T., & Ugliano, M. 2016, *ApJ*, **818**, 124
- Farmer, R., Renzo, M., de Mink, S. E., Marchant, P., & Justham, S. 2019, *ApJ*, **887**, 53
- Faucher-Giguere, C.-A., & Kaspi, V. M. 2006, *ApJ*, **643**, 332
- Finn, L. S., & Chernoff, D. F. 1993, *PhRvD*, **47**, 2198
- Fowler, W. A., & Hoyle, F. 1964, *ApJS*, **9**, 201
- Fraley, G. S. 1968, *Ap&SS*, **2**, 96
- Friend, D. B., & Abbott, D. C. 1986, *ApJ*, **311**, 701
- Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, *ApJ*, **749**, 91
- Furlong, M., Bower, R. G., Theuns, T., et al. 2015, *MNRAS*, **450**, 4486
- Galassi, M., Davies, J., Theiler, J., et al. 2002, GNU Scientific Library (Godalming: Network Theory Limited)
- Ge, H., Webbink, R. F., Chen, X., & Han, Z. 2015, *ApJ*, **812**, 40
- Geller, A. M., Hurley, J. R., & Mathieu, R. D. 2013, *AJ*, **145**, 8
- Gessner, A., & Janka, H.-T. 2018, *ApJ*, **865**, 61
- Giacobbo, N., & Mapelli, M. 2018, *MNRAS*, **480**, 2011
- Giacobbo, N., & Mapelli, M. 2020, *ApJ*, **891**, 141
- Giacobbo, N., Mapelli, M., & Spera, M. 2018, *MNRAS*, **474**, 2959
- Götberg, Y., de Mink, S. E., Groh, J. H., Leitherer, C., & Norman, C. 2019, *A&A*, **629**, A134
- Götberg, Y., de Mink, S. E., McQuinn, M., et al. 2020, *A&A*, **634**, A134
- Gunn, J. E., & Ostriker, J. P. 1970, *ApJ*, **160**, 979
- Hamann, W.-R., & Koesterke, L. 1998, *A&A*, **335**, 1003
- Hannam, M., Schmidt, P., Bohé, A., et al. 2014, *PhRvL*, **113**, 151101
- Hansen, B. M. S., & Phinney, E. S. 1997, *MNRAS*, **291**, 569
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Heger, A., Langer, N., & Woosley, S. E. 2000, *ApJ*, **528**, 368
- Heggie, D. C. 1975, *MNRAS*, **173**, 729
- Hild, S., Abernathy, M., Acernese, F., et al. 2011, *CQGra*, **28**, 094013
- Hills, J. G. 1983, *ApJ*, **267**, 322
- Hirai, R., & Mandel, I. 2021, *PASA*, **38**, e056
- Hjellming, M. S., & Webbink, R. F. 1987, *ApJ*, **318**, 794
- Hobbs, G., Lorimer, D. R., Lyne, A. G., & Kramer, M. 2005, *MNRAS*, **360**, 974
- Höfner, S., & Olofsson, H. 2018, *A&ARv*, **26**, 1
- Hogg, D. W. 1999, arXiv:astro-ph/9905116
- Howitt, G., Stevenson, S., Vigna-Gómez, A. r., et al. 2020, *MNRAS*, **492**, 3229
- Hoyle, F., & Lyttleton, R. A. 1939, *PCPS*, **35**, 405
- Huang, S. S. 1956, *AJ*, **61**, 49
- Huang, S.-S. 1963, *ApJ*, **138**, 471
- Humphreys, R. M., & Davidson, K. 1994, *PASP*, **106**, 1025
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, *MNRAS*, **315**, 543
- Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, *MNRAS*, **329**, 897
- Husa, S., Khan, S., Hannam, M., et al. 2016, *PhRvD*, **93**, 044006
- Ivanova, N., Heinke, C. O., Rasio, F. A., Belczynski, K., & Fregeau, J. M. 2008, *MNRAS*, **386**, 553
- Ivanova, N., Justham, S., Chen, X., et al. 2013, *A&ARv*, **21**, 59
- Ivanova, N., & Taam, R. E. 2004, *ApJ*, **601**, 1058
- Izzard, R. G., Dray, L. M., Karakas, A. I., Lugaro, M., & Tout, C. A. 2006, *A&A*, **460**, 565
- Izzard, R. G., Glebbeek, E., Stancliffe, R. J., & Pols, O. R. 2009, *A&A*, **508**, 1359
- Izzard, R. G., Tout, C. A., Karakas, A. I., & Pols, O. R. 2004, *MNRAS*, **350**, 407
- Janka, H. T. 2013, *MNRAS*, **434**, 1355
- Janka, H.-T., & Müller, E. 1994, *A&A*, **290**, 496
- Johnston, S., Hobbs, G., Vigeland, S., et al. 2005, *MNRAS*, **364**, 1397
- Kalogera, V. 1996, *ApJ*, **471**, 352
- Khan, S., Husa, S., Hannam, M., et al. 2016, *PhRvD*, **93**, 044007
- Kiel, P., Hurley, J., Bailes, M., & Murray, J. 2008, *MNRAS*, **388**, 393
- Klencki, J., Moe, M., Gladysz, W., et al. 2018, *A&A*, **619**, A77
- Klencki, J., Nelemans, G., Istrate, A. G., & Chruslinska, M. 2021, *A&A*, **645**, A54
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, Positioning and Power in Academic Publishing: Players, Agents and Agendas (Amsterdam: IOS Press), 87
- Kobulnicky, H. A., Kiminki, D. C., Lundquist, M. J., et al. 2014, *ApJS*, **213**, 34
- Kozyreva, A., Blinnikov, S., Langer, N., & Yoon, S. C. 2014a, *A&A*, **565**, A70
- Kozyreva, A., Gilmer, M., Hirschi, R., et al. 2017, *MNRAS*, **464**, 2854
- Kozyreva, A., Yoon, S. C., & Langer, N. 2014b, *A&A*, **566**, A146
- Kroupa, P. 2001, *MNRAS*, **322**, 231
- Kruckow, M. U., Tauris, T. M., Langer, N., et al. 2016, *A&A*, **596**, A58
- Kudritzki, R. P., Pauldrach, A., Puls, J., & Abbott, D. C. 1989, *A&A*, **219**, 205
- Kudritzki, R. P., & Reimers, D. 1978, *A&A*, **70**, 227
- Lai, D., Chernoff, D. F., & Cordes, J. M. 2001, *ApJ*, **549**, 1111
- Langer, N. 1992, *A&A*, **265**, L17
- Langer, N. 2012, *ARA&A*, **50**, 107
- Langer, N., & Norman, C. A. 2006, *ApJL*, **638**, L63
- Lattimer, J. M., & Schutz, B. F. 2005, *ApJ*, **629**, 979
- Lau, M. Y. M., Mandel, I., Vigna-Gómez, A., et al. 2020, *MNRAS*, **492**, 3061
- Leung, S.-C., Nomoto, K., & Blinnikov, S. 2019, *ApJ*, **887**, 72
- LIGO Scientific Collaboration 2018, LIGO Algorithm Library—LALSuite, Free Software (GPL), doi:10.7935/GT1W-FZ16
- Lin, L., Bingham, D., Broekgaarden, F., & Mandel, I. 2021, arXiv:2106.01552
- Lipunov, V. M., Ozernoy, L. M., Popov, S. B., Postnov, K. A., & Prokhorov, M. E. 1996a, *ApJ*, **466**, 234
- Lipunov, V. M., Postnov, K. A., & Prokhorov, M. E. 1996b, *A&A*, **310**, 489
- Lipunov, V. M., Postnov, K. A., Prokhorov, M. E., & Bogomazov, A. I. 2009, *ARep*, **53**, 915
- Loveridge, A. J., van der Sluys, M. V., & Kalogera, V. 2011, *ApJ*, **743**, 49
- Lyne, A. G., & Lorimer, D. R. 1994, *Natur*, **369**, 127
- Ma, X., Hopkins, P. F., Faucher-Giguère, C.-A., et al. 2016a, *MNRAS*, **456**, 2140
- Ma, X., Hopkins, P. F., Kasen, D., et al. 2016b, *MNRAS*, **459**, 3614
- MacLeod, M., & Ramirez-Ruiz, E. 2015, *ApJL*, **798**, L19
- Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415
- Madau, P., & Fragos, T. 2017, *ApJ*, **840**, 39
- Maeder, A. 1987, *A&A*, **178**, 159
- Maeder, A., & Meynet, G. 2000, *ARA&A*, **38**, 143
- Mandel, I. 2016, *MNRAS*, **456**, 578
- Mandel, I., & Broekgaarden, F. S. 2021, arXiv:2107.14239
- Mandel, I., & de Mink, S. E. 2016, *MNRAS*, **458**, 2634
- Mandel, I., & Farmer, A. 2018, arXiv:1806.05820
- Mandel, I., & Müller, B. 2020, *MNRAS*, **499**, 3214
- Mandel, I., Müller, B., Riley, J., et al. 2021, *MNRAS*, **500**, 1380

- Marchant, P., Langer, N., Podsiadlowski, P., Tauris, T. M., & Moriya, T. J. 2016, *A&A*, **588**, A50
- Marchant, P., Renzo, M., Farmer, R., et al. 2019, *ApJ*, **882**, 36
- Massevitch, A., & Yungelson, L. 1975, *MmSAI*, **46**, 217
- Mestel, L. 1952, *MNRAS*, **112**, 583
- Miller-Jones, J. C. A., Bahramian, A., Orosz, J. A., et al. 2021, *Sci*, **371**, 1046
- Miyaji, S., Nomoto, K., Yokoi, K., & Sugimoto, D. 1980, *PASJ*, **32**, 303
- Moe, M., & Di Stefano, R. 2017, *ApJS*, **230**, 15
- Mohamed, S., & Podsiadlowski, P. 2007, in ASP Conf. Ser. 372, 15th European Workshop on White Dwarfs, ed. R. Napiwotzki & M. R. Burleigh (San Francisco, CA: ASP), 397
- Mokiem, M. R., de Koter, A., Vink, J. S., et al. 2007, *A&A*, **473**, 603
- Moriya, T. J., Mazzali, P. A., Tominaga, N., et al. 2017, *MNRAS*, **466**, 2085
- Müller, B. 2020, *LRCA*, **6**, 3
- Müller, B., Gay, D. W., Heger, A., Tauris, T. M., & Sim, S. A. 2018, *MNRAS*, **479**, 3675
- Müller, B., Heger, A., Liptai, D., & Cameron, J. B. 2016, *MNRAS*, **460**, 742
- Müller, B., Tauris, T. M., Heger, A., et al. 2019, *MNRAS*, **484**, 3307
- Nagakura, H., Sumiyoshi, K., & Yamada, S. 2019, *ApJL*, **880**, L28
- Nakamura, K., Takiwaki, T., Kuroda, T., & Kotake, K. 2015, *PASJ*, **67**, 107
- Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., et al. 2019, *MNRAS*, **490**, 3740
- Neijssel, C. J., Vinciguerra, S., Vigna-Gómez, A., et al. 2021, *ApJ*, **908**, 118
- Nelemans, G., Yungelson, L. R., Portegies Zwart, S. F., & Verbunt, F. 2001, *A&A*, **365**, 491
- Neo, S., Miyaji, S., Nomoto, K., & Sugimoto, D. 1977, *PASJ*, **29**, 249
- Nieuwenhuijzen, H., & de Jager, C. 1990, *A&A*, **231**, 134
- Nomoto, K. 1984, *ApJ*, **277**, 791
- Nomoto, K. 1987, *ApJ*, **322**, 206
- Nomoto, K., & Kondo, Y. 1991, *ApJL*, **367**, L19
- Noutsos, A., Kramer, M., Carr, P., & Johnston, S. 2012, *MNRAS*, **423**, 2736
- Öpik, E. 1924, *PTarO*, **25**, 1
- Oslowski, S., Bulik, T., Gondek-Rosińska, D., & Belczyński, K. 2011, *MNRAS*, **413**, 461
- Paczynski, B. 1976, in IAU Symp. 73, Structure and Evolution of Close Binary Systems, ed. P. Eggleton, S. Mitton, & J. Whelan (Dordrecht: D. Reidel), 75
- Paczynski, B., & Sienkiewicz, R. 1972, *AcA*, **22**, 73
- Pan, Y., Buonanno, A., Taracchini, A., et al. 2014, *PhRvD*, **89**, 084006
- Panther, B., Heavens, A. F., & Jimenez, R. 2004, *MNRAS*, **355**, 764
- Pérez, F., & Granger, B. E. 2007, *CSE*, **9**, 21
- Peters, P. C. 1964, *PhRv*, **136**, 1224
- Pfahl, E., Rappaport, S., & Podsiadlowski, P. 2002a, *ApJL*, **571**, L37
- Pfahl, E., Rappaport, S., & Podsiadlowski, P. 2002b, *ApJ*, **573**, 283
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, **594**, A13
- Podsiadlowski, P. 2001, in ASP Conf. Ser. 229, Common-envelope Evolution and Stellar Mergers, ed. P. Podsiadlowski et al. (San Francisco, CA: ASP), 239
- Podsiadlowski, P. 2010, *NewAR*, **54**, 39
- Podsiadlowski, P., Joss, P. C., & Hsu, J. J. L. 1992, *ApJ*, **391**, 246
- Podsiadlowski, P., Langer, N., Poelarends, A. J. T., et al. 2004, *ApJ*, **612**, 1044
- Pols, O. R., Schröder, K.-P., Hurley, J. R., Tout, C. A., & Eggleton, P. P. 1998, *MNRAS*, **298**, 525
- Portegies Zwart, S. F., & Verbunt, F. 1996, *A&A*, **309**, 179
- Portegies Zwart, S. F., & Yungelson, L. R. 1998, *A&A*, **332**, 173
- Postnov, K. A., & Yungelson, L. R. 2014, *LRR*, **17**, 3
- Ramírez-Agudelo, O. H., Simón-Díaz, S., Sana, H., et al. 2013, *A&A*, **560**, A29
- Remillard, R. A., & McClintock, J. E. 2006, *ARA&A*, **44**, 49
- Renzo, M., Farmer, R., Justham, S., et al. 2020, *A&A*, **640**, A56
- Renzo, M., Ott, C. D., Shore, S. N., & de Mink, S. E. 2017, *A&A*, **603**, A118
- Repetto, S., Igoshev, A. P., & Nelemans, G. 2017, *MNRAS*, **467**, 298
- Riley, J., Mandel, I., Marchant, P., et al. 2021, *MNRAS*, **505**, 663
- Rosdahl, J., Katz, H., Blaizot, J., et al. 2018, *MNRAS*, **479**, 994
- Salpeter, E. E. 1955, *ApJ*, **121**, 161
- Sana, H. 2017, in IAU Symp. 329, The Lives and Death-Throes of Massive Stars, ed. J. J. Eldridge et al. (Cambridge: Cambridge Univ. Press), 110
- Sana, H., de Mink, S. E., de Koter, A., et al. 2012, *Sci*, **337**, 444
- Sana, H., Le Bouquin, J. B., Lacour, S., et al. 2014, *ApJS*, **215**, 15
- Sander, A. A. C., Vink, J. S., & Hamann, W. R. 2020, *MNRAS*, **491**, 4406
- Savaglio, S., Glazebrook, K., Le Borgne, D., et al. 2005, *ApJ*, **635**, 260
- Schneider, F. R. N., Izzard, R. G., Langer, N., & de Mink, S. E. 2015, *ApJ*, **805**, 20
- Schneider, F. R. N., Podsiadlowski, P., & Müller, B. 2021, *A&A*, **645**, A5
- Schröder, S. L., MacLeod, M., Loeb, A., Vigna-Gómez, A., & Mandel, I. 2020, *ApJ*, **892**, 13
- Schröder, S. L., MacLeod, M., Ramirez-Ruiz, E., et al. 2021, arXiv:2107.09675
- Schwab, J., Podsiadlowski, P., & Rappaport, S. 2010, *ApJ*, **719**, 722
- Smith, N. 2014, *ARA&A*, **52**, 487
- Soberman, G. E., Phinney, E. S., & van den Heuvel, E. P. J. 1997, *A&A*, **327**, 620
- Socrates, A., Blaes, O., Hungerford, A., & Fryer, C. L. 2005, *ApJ*, **632**, 531
- Spera, M., & Mapelli, M. 2017, *MNRAS*, **470**, 4739
- Stevenson, S., Sampson, M., Powell, J., et al. 2019, *ApJ*, **882**, 121
- Stevenson, S., Vigna-Gómez, A., Mandel, I., et al. 2017, *NatCo*, **8**, 14906
- Strolger, L.-G., Riess, A. G., Dahlen, T., et al. 2004, *ApJ*, **613**, 200
- Sukhbold, T., & Adams, S. 2020, *MNRAS*, **492**, 2578
- Sukhbold, T., & Woosley, S. 2014, *ApJ*, **783**, 10
- Suwa, Y., Yoshida, T., Shibata, M., Umeda, H., & Takahashi, K. 2015, *MNRAS*, **454**, 3073
- Taam, R. E., & Sandquist, E. L. 2000, *ARA&A*, **38**, 113
- Takahashi, K. 2018, *ApJ*, **863**, 153
- Talon, S., Zahn, J.-P., Maeder, A., & Meynet, G. 1997, *A&A*, **322**, 209
- Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, *ApJ*, **846**, 170
- Tauris, T. M., Langer, N., Moriya, T. J., et al. 2013, *ApJL*, **778**, L23
- Tauris, T. M., Langer, N., & Podsiadlowski, P. 2015, *MNRAS*, **451**, 2123
- Tauris, T. M., & Takens, R. J. 1998, *A&A*, **330**, 1047
- Tauris, T. M., & van den Heuvel, E. P. J. 2006, in Compact Stellar X-ray Sources, ed. W. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 623
- Taylor, S. R., & Gerosa, D. 2018, *PhRvD*, **98**, 083017
- Team COMPAS 2022, *JOSS*, **7**, 3838
- Timmes, F. X., Woosley, S. E., & Weaver, T. A. 1996, *ApJ*, **457**, 834
- Toonen, S., Nelemans, G., & Portegies Zwart, S. 2012, *A&A*, **546**, A70
- Tout, C. A., Aarseth, S. J., Pols, O. R., & Eggleton, P. P. 1997, *MNRAS*, **291**, 732
- Tout, C. A., Pols, O. R., Eggleton, P. P., & Han, Z. 1996, *MNRAS*, **281**, 257
- Tutukov, A., & Yungelson, L. 1996, *MNRAS*, **280**, 1035
- Ugliko, M., Janka, H.-T., Marek, A., & Arcones, A. 2012, *ApJ*, **757**, 69
- van den Heuvel, E. P. J. 1976, in IAU Symp. 73, Structure and Evolution of Close Binary Systems, ed. P. Eggleton, S. Mitton, & J. Whelan (Cambridge: Cambridge Univ. Press), 35
- van Rossum, G. 1995, Python Tutorial, Tech. Rep., CS-R9526 (Amsterdam: Centrum voor Wiskunde en Informatica)
- van Son, L., de Mink, S., Broekgaarden, F., et al. 2020, *ApJ*, **897**, 100
- van Son, L. A. C., de Mink, S. E., Callister, T., et al. 2021, arXiv:2110.01634
- Vassiliadis, E., & Wood, P. R. 1993, *ApJ*, **413**, 641
- Verbunt, F., Igoshev, A., & Cator, E. 2017, *A&A*, **608**, A57
- Vick, M., & Lai, D. 2020, *MNRAS*, **496**, 3767
- Vigna-Gómez, A., MacLeod, M., Neijssel, C. J., et al. 2020, *PASA*, **37**, e038
- Vigna-Gómez, A., Neijssel, C. J., Stevenson, S., et al. 2018, *MNRAS*, **481**, 4009
- Vigna-Gómez, A., Toonen, S., Ramirez-Ruiz, E., et al. 2021, *ApJL*, **907**, L19
- Vinciguerra, S., Neijssel, C. J., Vigna-Gómez, A., et al. 2020, *MNRAS*, **498**, 4705
- Vink, J. S., & de Koter, A. 2005, *A&A*, **442**, 587
- Vink, J. S., de Koter, A., & Lamers, H. J. G. L. M. 2000, *A&A*, **362**, 295
- Vink, J. S., de Koter, A., & Lamers, H. J. G. L. M. 2001, *A&A*, **369**, 574
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Waskom, M., Gelbart, M., Botvinnik, O., et al. 2020, mwaskom/seaborn: v0.11.2 (August 2021), Zenodo, doi:10.5281/zenodo.592845
- Webbink, R. F. 1984, *ApJ*, **277**, 355
- Willcox, R., Mandel, I., Thrane, E., et al. 2021, *ApJL*, **920**, L37
- Wongwathanarat, A., Janka, H. T., & Mueller, E. 2013, *A&A*, **552**, A126
- Woods, T. E., Ivanova, N., van der Sluis, M. V., & Chaichenets, S. 2012, *ApJ*, **744**, 12
- Woosley, S., & Heger, A. 2007, *PhR*, **442**, 269
- Woosley, S. E. 1987, in IAU Symp. 125, The Origin and Evolution of Neutron Stars, ed. D. J. Helfand & J. H. Huang (Dordrecht: D. Reidel), 255
- Woosley, S. E. 2017, *ApJ*, **836**, 244
- Woosley, S. E. 2019, *ApJ*, **878**, 49
- Woosley, S. E., & Bloom, J. S. 2006, *ARA&A*, **44**, 507
- Woosley, S. E., Heger, A., & Weaver, T. A. 2002, *RvMP*, **74**, 1015
- Xu, X.-J., & Li, X.-D. 2010a, *ApJ*, **716**, 114
- Xu, X.-J., & Li, X.-D. 2010b, *ApJ*, **722**, 1985

Yao, J., Zhu, W., Manchester, R. N., et al. 2021, [NatAs](#), **5**, 788

Ye, C. S., Kremer, K., Chatterjee, S., Rodriguez, C. L., & Rasio, F. A. 2019, [ApJ](#), **877**, 122

Yoon, S.-C., Chun, W., Tolstov, A., Blinnikov, S., & Dessart, L. 2019, [ApJ](#), **872**, 174

Yoon, S.-C., Dessart, L., & Clocchiatti, A. 2017, [ApJ](#), **840**, 10

Yoshida, T., Umeda, H., Maeda, K., & Ishii, T. 2016, [MNRAS](#), **457**, 351

Zapartas, E., de Mink, S. E., Justham, S., et al. 2019, [A&A](#), **631**, A5

Zapartas, E., de Mink, S. E., Justham, S., et al. 2021, [A&A](#), **645**, A6

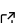
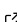
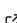
COMPAS: A rapid binary population synthesis suite

Team COMPAS¹, Jeff Riley^{2,3}, Poojan Agrawal^{4,3}, Jim W. Barrett⁵, Kristian N. K. Boyett⁶, Floor S. Broekgaarden⁷, Debatri Chattopadhyay^{8,4,3}, Sebastian M. Gaebel⁹, Fabian Gittins¹⁰, Ryosuke Hirai^{2,3}, George Howitt¹¹, Stephen Justham^{12,13,14}, Lokesh Khandelwal¹², Floris Kummer¹², Mike Y. M. Lau^{2,3}, Ilya Mandel^{2,3,5}, Selma E. de Mink^{14,12,7}, Coenraad Neijssel^{5,3}, Tim Riley^{2,3}, Lieke van Son^{7,12,14}, Simon Stevenson^{4,3}, Alejandro Vigna-Gómez^{15,16}, Serena Vinciguerra¹², Tom Wagg^{7,14}, and Reinhold Willcox^{2,3}

1 The public COMPAS code is a product of work by the entire COMPAS collaboration over many years; we therefore kindly request that, in recognition of this team effort, the paper is cited as Team COMPAS - J. Riley et al. 2 School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia 3 OzGrav, Australian Research Council Centre of Excellence for Gravitational Wave Discovery, Australia 4 Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia 5 Institute of Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Birmingham, B15 2TT 6 Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK 7 Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA 8 School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA, United Kingdom 9 Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstrasse 38, D-30167 Hannover, Germany 10 Mathematical Sciences and STAG Research Centre, University of Southampton, Southampton SO17 1BJ, UK 11 School of Physics, University of Melbourne, Parkville, Victoria, 3010, Australia 12 Anton Pannekoek Institute of Astronomy and GRAPPA, Science Park 904, University of Amsterdam, 1098XH Amsterdam, The Netherlands 13 School of Astronomy & Space Science, University of the Chinese Academy of Sciences, Beijing 100012, China 14 Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany 15 DARK, Niels Bohr Institute, University of Copenhagen, Jagtvej 128, 2200, Copenhagen, Denmark 16 Niels Bohr International Academy, The Niels Bohr Institute, Blegdamsvej 17, 2100 Copenhagen, Denmark

DOI: [10.21105/joss.03838](https://doi.org/10.21105/joss.03838)

Software

- [Review](#) 
- [Repository](#) 
- [Archive](#) 

Editor: [Christina Hedges](#) 

Reviewers:

- [@katiebreivik](#)
- [@HeloiseS](#)


Submitted: 12 October 2021

Published: 21 January 2022

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)).

In partnership with



Summary

Most massive stars—those with initial masses greater than $8 M_{\odot}$ —are born with another massive star as a companion (Moe & Di Stefano, 2017; Sana et al., 2012). Massive binary stars are responsible for producing many exotic astrophysical phenomena, such as the observed diversity of supernovae, binary pulsars, X-ray binaries and merging compact objects. The latter are now regularly observed by the ground-based gravitational wave observatories Advanced LIGO and Virgo (B. P. Abbott et al., 2016; R. Abbott et al., 2021). Population models of massive binary evolution make it possible to interpret existing observations and to make predictions for future observing campaigns.

Statement of need

Binary population synthesis generates population models of isolated stellar binaries under a set of parametrized assumptions. These models permit comparisons against observational data sets, such as X-ray binaries of gravitational-wave mergers.

In particular, rapid binary population synthesis is needed in order to efficiently explore a broad parameter space of uncertain assumptions about the physics of stellar and binary evolution, including supernova remnant masses and natal kicks, mass transfer efficiency and stability, and the outcome of common-envelope events.

A range of binary population synthesis codes have been developed over the last three decades. These include the Scenario Machine (Lipunov et al., 1996), IBiS (Tutukov & Yungelson, 1996), SeBa (Portegies Zwart & Verbunt, 1996), BSE (Hurley et al., 2002), StarTrack (Belczynski et al., 2008), binary_c (Izzard et al., 2004), MOBSE (Giacobbo et al., 2018) and COSMIC (Breivik et al., 2020). These codes range from private to semi-public to fully public, and differ in the range of available tools, computational complexity, and speed of execution.

COMPAS is a rapid binary population synthesis suite. It parametrizes complex astrophysical processes with prescriptions calibrated to detailed models. COMPAS is designed to allow for flexible modifications as evolutionary models improve. All code is fully public and, including pre-processing and post-processing tools. COMPAS is computationally efficient, with a focus on the statistical analysis of large populations, particularly but not exclusively in the context of gravitational-wave astronomy.

Details

The core engine of COMPAS—responsible for calculating the evolution of single (Hurley et al., 2000) and binary (Hurley et al., 2002) stars—is written in object oriented C++ for speed and flexibility. COMPAS is able to simulate the evolution of a typical binary over 10 Gyr in approximately 10 milliseconds.

A detailed description of the implementation of the COMPAS suite can be found in [Team COMPAS: Riley et al. \(2021\)](#).

In addition to the core stellar and binary evolution engine, we provide Python scripts for both pre- and post-processing COMPAS outputs. Post-processing can account for integrating populations formed throughout cosmic history (Neijssel et al., 2019) and methods to account for gravitational-wave selection effects (Barrett et al., 2018). A set of examples is also provided.

COMPAS is *embarrassingly* parallel and can be trivially run on high performance computers and distributed on cloud computing.

COMPAS was initially designed to focus on studies of merging binaries containing neutron stars and black holes that are being observed through gravitational waves (Stevenson et al., 2017; Vigna-Gómez et al., 2018). In recent years, the scope of systems investigated with COMPAS has expanded to incorporate, e.g., Be X-ray binaries (Vinciguerra et al., 2020) and luminous red novae (Howitt et al., 2020) (see [Team COMPAS: Riley et al. \(2021\)](#) or [the COMPAS collaboration website](#) for a summary of COMPAS publications to date.)

COMPAS development happens on [Github](#). We maintain a [Zenodo community](#) where data from many publications using COMPAS is publicly available.

Acknowledgements

Multiple authors are supported by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), through project number CE170100004. Multiple authors were funded in part by the National Science Foundation under Grant No. (NSF grant number 2009131), the Netherlands Organization for Scientific Research (NWO) as part of the Vidi research program BinWaves with project number 639.042.728 and by the European

Union's Horizon 2020 research and innovation program from the European Research Council (ERC, Grant agreement No. 715063). FSB is supported in part by the Prins Bernard Cultuurfonds studiebeurs. IM is a recipient of an Australian Research Council Future Fellowship (FT190100574). AVG acknowledges funding support by the Danish National Research Foundation (DNRF132)

References

- Abbott, B. P., Abbott, R., Abbott, T. D., Abernathy, M. R., Acernese, F., Ackley, K., Adams, C., Adams, T., Addesso, P., Adhikari, R. X., & al., et. (2016). Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett*, *116*(6), 061102. <https://doi.org/10.1103/PhysRevLett.116.061102>
- Abbott, R., Abbott, T. D., Acernese, F., Ackley, K., Adams, C., Adhikari, N., & others. (2021). GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. *arXiv e-Prints*. <http://arxiv.org/abs/2111.03606>
- Barrett, J. W., Gaebel, S. M., Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., Berry, C. P. L., Farr, W. M., & Mandel, I. (2018). Accuracy of inference on the physics of binary evolution from gravitational-wave observations. *MNRAS*, *477*, 4685–4695. <https://doi.org/10.1093/mnras/sty908>
- Belczynski, K., Kalogera, V., Rasio, F. A., Taam, R. E., Zezas, A., Bulik, T., Maccarone, T. J., & Ivanova, N. (2008). Compact Object Modeling with the StarTrack Population Synthesis Code. *ApJS*, *174*(1), 223–260. <https://doi.org/10.1086/521026>
- Breivik, K., Coughlin, S., Zevin, M., Rodriguez, C. L., Kremer, K., Ye, C. S., Andrews, J. J., Kurkowski, M., Digman, M. C., Larson, S. L., & Rasio, F. A. (2020). COSMIC Variance in Binary Population Synthesis. *Astrophysical Journal*, *898*(1), 71. <https://doi.org/10.3847/1538-4357/ab9d85>
- Giacobbo, N., Mapelli, M., & Spera, M. (2018). Merging black hole binaries: the effects of progenitor's metallicity, mass-loss rate and Eddington factor. *MNRAS*, *474*, 2959–2974. <https://doi.org/10.1093/mnras/stx2933>
- Howitt, G., Stevenson, S., Vigna-Gómez, A., Justham, S., Ivanova, N., Woods, T. E., Neijssel, C. J., & Mandel, I. (2020). Luminous Red Novae: population models and future prospects. *MNRAS*, *492*(3), 3229–3240. <https://doi.org/10.1093/mnras/stz3542>
- Hurley, J. R., Pols, O. R., & Tout, C. A. (2000). Comprehensive analytic formulae for stellar evolution as a function of mass and metallicity. *MNRAS*, *315*(3), 543–569. <https://doi.org/10.1046/j.1365-8711.2000.03426.x>
- Hurley, J. R., Tout, C. A., & Pols, O. R. (2002). Evolution of binary stars and the effect of tides on binary populations. *MNRAS*, *329*(4), 897–928. <https://doi.org/10.1046/j.1365-8711.2002.05038.x>
- Izzard, R. G., Tout, C. A., Karakas, A. I., & Pols, O. R. (2004). A new synthetic model for asymptotic giant branch stars. *MNRAS*, *350*, 407–426. <https://doi.org/10.1111/j.1365-2966.2004.07446.x>
- Lipunov, V. M., Postnov, K. A., & Prokhorov, M. E. (1996). The Scenario Machine: restrictions on key parameters of binary evolution. *A&A*, *310*, 489–507.
- Moe, M., & Di Stefano, R. (2017). Mind Your Ps and Qs: The Interrelation between Period (P) and Mass-ratio (Q) Distributions of Binary Stars. *ApJS*, *230*(2), 15. <https://doi.org/10.3847/1538-4365/aa6fb6>
- Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., Barrett, J. W., Gaebel, S. M., Broekgaarden, F. S., de Mink, S. E., Szécsi, D., Vinciguerra, S., & Mandel, I. (2019). The effect of the

- metallicity-specific star formation history on double compact object mergers. *MNRAS*, 490(3), 3740–3759. <https://doi.org/10.1093/mnras/stz2840>
- Portegies Zwart, S. F., & Verbunt, F. (1996). Population synthesis of high-mass binaries. *A&A*, 309, 179–196.
- Sana, H., de Mink, S. E., de Koter, A., Langer, N., Evans, C. J., Gieles, M., Gosset, E., Izzard, R. G., Le Bouquin, J.-B., & Schneider, F. R. N. (2012). Binary Interaction Dominates the Evolution of Massive Stars. *Science*, 337, 444. <https://doi.org/10.1126/science.1223344>
- Stevenson, S., Vigna-Gómez, A., Mandel, I., Barrett, J. W., Neijssel, C. J., Perkins, D., & de Mink, S. E. (2017). Formation of the first three gravitational-wave observations through isolated binary evolution. *Nature Communications*, 8, 14906. <https://doi.org/10.1038/ncomms14906>
- Team COMPAS: Riley, J., Agrawal, P., Barrett, J. W., Boyett, K. N. K., Broekgaarden, F. S., Chattopadhyay, D., Gaebel, S. M., Gittins, F., Hirai, R., Howitt, G., Justham, S., Khandelwal, L., Kummer, F., Lau, M. Y. M., Mandel, I., de Mink, S. E., Neijssel, C., Riley, T., van Son, L., ... Willcox, R. (2021). Rapid stellar and binary population synthesis with COMPAS. *arXiv e-Prints*, arXiv:2109.10352. <http://arxiv.org/abs/2109.10352>
- Tutukov, A., & Yungelson, L. (1996). Double-degenerate semidetached binaries with helium secondaries: cataclysmic variables, supersoft X-ray sources, supernovae and accretion-induced collapses. *MNRAS*, 280(4), 1035–1045. <https://doi.org/10.1093/mnras/280.4.1035>
- Vigna-Gómez, A., Neijssel, C. J., Stevenson, S., Barrett, J. W., Belczynski, K., Justham, S., de Mink, S. E., Müller, B., Podsiadlowski, P., Renzo, M., Szécsi, D., & Mandel, I. (2018). On the formation history of Galactic double neutron stars. *MNRAS*, 481, 4009–4029. <https://doi.org/10.1093/mnras/sty2463>
- Vinciguerra, S., Neijssel, C. J., Vigna-Gómez, A., Mandel, I., Podsiadlowski, P., Maccarone, T. J., Nicholl, M., Kingdon, S., Perry, A., & Salemi, F. (2020). Be X-ray binaries in the SMC as indicators of mass-transfer efficiency. *MNRAS*, 498(4), 4705–4720. <https://doi.org/10.1093/mnras/staa2177>

Chapter 3

Chemically Homogeneous Evolution: A rapid population synthesis approach

Rotationally-induced chemical mixing in massive stars can prevent the establishment of a strong chemical gradient, and so significantly affect the star's evolution [76, 77]. In such cases, as long as the star continues to rotate at a sufficiently high rate, it will remain quasi-chemically homogeneous [78, 79]. Contrary to the core-envelope structure exhibited by more slowly rotating stars, and the characteristic expansion of the envelope as the core contracts, the radius of chemically homogeneous (CH) stars will shrink or remain constant as they become hotter, brighter, and bluer [80].

de Mink et al. [81] argued that the conditions for chemical homogeneity can be achieved in very close massive binary systems, and proposed CHE as a viable formation channel for binary black holes that can coalesce within the age of the Universe and that would otherwise merge before, or soon after, the completion of hydrogen burning due to the expansion of the stars. VFTS 352 [82] and HD 5980 [83] are examples of observed binary systems thought to have undergone CHE [84].

Marchant et al. [85] conducted detailed simulations of the CHE formation channel using the MESA code [53–55], which were further developed and investigated by du Buisson et al. [86]. Mandel and de Mink [80] conducted simulations in which they explore a broad parameter space using a set of simplifying assumptions that allow rapid population synthesis.

COMPAS (Chapter 2) implemented a number of models for stellar and binary evolution, but lacked a model for CHE. We implemented a model for CHE in COMPAS, and

extended the work of Marchant et al. [85] and [80]. The following paper, Riley et al. [87], describes the COMPAS implementation of CHE.

The CHE model implemented in COMPAS is a fairly naïve treatment, and some simplifying assumptions were made (detailed in the following paper). This work, and the work described in Chapter 2, enables work to be done to improve the COMPAS CHE model. Stevenson and Houlden [88] have recently extended the CHE implementation in COMPAS to include better treatments of the lifetime and luminosity of CHE stars relative to non-rotating main-sequence stars, the effects of both metallicity and rotation on mass loss, and some work to determine the final spins of black holes formed through CHE.

Chemically homogeneous evolution: a rapid population synthesis approach

Jeff Riley¹,^{1,2}★ Ilya Mandel¹,^{1,2,3,4} Pablo Marchant,⁵ Ellen Butler,³ Kaila Nathaniel,⁶
Coenraad Neijssel,^{1,3} Spencer Shortt⁷ and Alejandro Vigna-Gómez⁸

¹*School of Physics and Astronomy, Monash University, Clayton, VIC 3800, Australia*

²*ARC Centre of Excellence for Gravitational Wave Discovery – OzGrav, Hawthorn, VIC 3122, Australia*

³*Birmingham Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, B15 2TT Birmingham, UK*

⁴*ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions – ASTRO 3D, Canberra, ACT 2601, Australia*

⁵*Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, B-3001 Leuven, Belgium*

⁶*Argelander-Institut für Astronomie, Universität Bonn, Auf dem Hügel 71, D-53121 Bonn, Germany*

⁷*Department of Mathematics, University of Colorado Boulder, Boulder, CO 80309, USA*

⁸*DARK, Niels Bohr Institute, University of Copenhagen, Jagtvej 128, DK-2200 Copenhagen, Denmark*

Accepted 2021 April 30. Received 2021 April 14; in original form 2020 September 29

ABSTRACT

We explore chemically homogeneous evolution (CHE) as a formation channel for massive merging binary black holes (BBHs). We develop methods to include CHE in a rapid binary population synthesis code, Compact Object Mergers: Population Astrophysics and Statistics (COMPAS), which combines realistic models of binary evolution with cosmological models of the star formation history of the Universe. For the first time, we simultaneously explore conventional isolated binary star evolution under the same set of assumptions. This approach allows us to constrain population properties and make simultaneous predictions about the gravitational-wave detection rates of BBH mergers for the CHE and conventional formation channels. The overall mass distribution of detectable BBHs is consistent with existing gravitational-wave observations. We find that the CHE channel may yield up to ~ 70 per cent of all gravitational-wave detections of BBH mergers coming from isolated binary evolution.

Key words: gravitational waves – stars: evolution – stars: massive – black hole mergers – binaries: close.

1 INTRODUCTION

On 2015 September 14, the first direct observation of gravitational waves was made by the Advanced Laser Interferometer Gravitational-wave Observatory (aLIGO; Abbott et al. 2016). The detected signal, now known as GW150914, was also the first observation of two black holes merging, thus confirming the existence of binary stellar-mass black hole systems and providing evidence that they can merge within the current age of the Universe. Based on 10 binary black hole (BBH) detections during the first two observing runs of aLIGO and advanced Virgo, Abbott et al. (2019b) estimate a local BBH merger rate of $25\text{--}109\text{ Gpc}^{-3}\text{ yr}^{-1}$ at a 90 per cent confidence.

How the BBH sources of these gravitational-wave signals form remains an open question. To be the source of gravitational waves detected at aLIGO, which is sensitive to signals with frequencies of tens to hundreds of Hz, compact objects orbiting each other must spiral in as they lose energy through the emission of gravitational waves. Orbital energy loss through gravitational-wave emission is not efficient at wide separations, and the time-scale for gravitational-wave emission to drive a binary to merger scales as the fourth power of the orbital separation (Peters 1964). In order for two $30M_{\odot}$ black holes to merge within ≈ 14 Gyr, the current age of the Universe, their initial separation must be below $\lesssim 50R_{\odot}$, and

therein lies a problem, as this is smaller than the radial extent reached by typical slowly rotating massive stars during their evolution. The different astrophysical channels proposed for forming merging BBHs generally fall into the following two categories (see e.g. Mandel & Farmer 2018; Mapelli 2018, for reviews):

(i) isolated binary evolution, in which two stars may interact through tides and mass transfer, but are dynamically decoupled from other stars (e.g. Tutukov & Yungelson 1973, 1993; van den Heuvel 1976).

(ii) dynamical formation, where dynamical interactions in a dense environment and/or a hierarchical triple system play a key role in forming and hardening a compact BBH (Sigurdsson & Hernquist 1993; Miller & Lauburg 2009; Ziosi et al. 2014; Rodriguez et al. 2015; Antonini et al. 2016; Bartos et al. 2017; Stone, Metzger & Haiman 2017).

A variant of the isolated binary evolution channel relies on rotationally induced chemical mixing in massive stars to prevent the establishment of a strong chemical gradient (Maeder 1987; Heger, Langer & Woosley 2000; Maeder & Meynet 2000). As long as the star continues to rotate at a sufficiently high rate, it will remain quasi-chemically homogeneous (CH; Maeder 1987; Langer 1992). Contrary to the core-envelope structure exhibited by conventional, more slowly rotating stars, and the characteristic expansion of the envelope as the core contracts, the radius of quasi-CH stars will shrink or remain constant as they become hotter and brighter (Yoon, Langer & Norman 2006; Mandel & de Mink 2016).

* E-mail: jeff.riley@monash.edu

As we discuss below, previous work on the chemically homogeneous evolution (CHE) channel for BBH formation (de Mink & Mandel 2016; Mandel & de Mink 2016; Marchant et al. 2016; du Buisson et al. 2020) explored this channel independently of the usual isolated binary evolution channel. In this paper, we present our rapid population synthesis model for the CHE of binary systems, allowing for a direct comparison of the rates and properties of CHE and non-CHE BBHs under the same set of assumptions. Our CHE model is implemented in the rapid binary population synthesis code Compact Object Mergers: Population Astrophysics and Statistics (COMPAS; Stevenson et al. 2017; Vigna-Gómez et al. 2018), with thresholds on CHE evolution computed using rotating stellar models in the Modules for Experiments in Astrophysics (MESA) code (Paxton et al. 2011, 2013, 2015).

The remainder of this paper is organized as follows. Section 2 is a brief outline of CHE and previous work on the formation of BBHs through this channel. Section 3 presents a description of our CHE model and the implementation of the model in COMPAS. We present our results in Section 4. We provide some concluding remarks in Section 5.

2 CHEMICALLY HOMOGENEOUS EVOLUTION

Stars evolving on the main sequence (MS) typically develop increasingly helium-rich cores and hydrogen-rich envelopes as radial mixing is inefficient. However, von Zeipel (1924) showed that rotating stars cannot simultaneously be in hydrostatic and thermal equilibrium if the rotational velocity is a function of only radius, which has been argued to result in meridional currents in the radiative layers of a rotating star (Eddington 1929; Sweet 1950). In massive rapidly rotating stars in low-metallicity environments, these currents can mix material from the convective core throughout the radiative envelope, leading to CHE for rapidly rotating stars (Maeder 1987).

Due to strong chemical mixing, CH stars do not maintain a hydrogen-rich envelope – thus avoiding the dramatic expansion exhibited during the post-MS phase by non-CH stars. The radius of a CH star remains stable, or shrinks slowly, as the star becomes increasingly helium rich over the course of the MS, with the star contracting to a massive naked helium star post-MS. CH components of a very close binary system can thus avoid overfilling their Roche lobes, mass transfer, and probable merger.

de Mink et al. (2009) modelled the evolution of rotating massive stars using the hydrodynamic stellar evolution code described by Yoon et al. (2006) and Petrovic et al. (2005), which includes the effects of rotation on the stellar structure and the transport of angular momentum via rotationally induced hydrodynamic instabilities (Heger et al. 2000). The binary models developed by de Mink et al. (2009) and Song et al. (2016) show that constituent stars in very tight binary systems can achieve rotational frequencies sufficient to induce CHE. de Mink et al. (2009) proposed CHE as a viable formation channel for high-mass black hole X-ray binaries. VFTS 352 (Almeida et al. 2015) and HD 5980 (Koenigsberger et al. 2014) are examples of observed binary systems thought to have undergone CHE (de Mink & Mandel 2016).

Mandel & de Mink (2016) and Marchant et al. (2016) introduced and investigated CHE as a channel for forming merging BBHs. They concluded that for sufficiently high masses and sufficiently low metallicities, a narrow range of initial orbital periods (short enough to allow rapid rotation necessary for CHE, but not so short that the binary would immediately merge) could allow this channel to produce merging BBHs.

Mandel & de Mink (2016) and de Mink & Mandel (2016) used approximate thresholds for CHE based on the models of Yoon et al. (2006) to investigate the rates and properties of BBHs formed through the CHE channel. They estimated a merger rate of $\sim 10 \text{ Gpc}^{-3} \text{ yr}^{-1}$ in the local Universe for this channel, subject to a number of evolutionary uncertainties, which they explored in a population-synthesis-style study.

Marchant et al. (2016) used the MESA code to conduct detailed simulations of the CHE channel, which were followed until the BBH stage. The simulations were conducted for close binaries with component masses above $\sim 20 M_{\odot}$, and included the overcontact phase in a majority of CHE BBH progenitors. Marchant et al. (2016) suggested that as long as material does not overflow the L2 point in overcontact binaries, co-rotation can be maintained, and a spiral-in due to viscous drag can be avoided. In this scenario, close binary systems typically enter the overcontact phase in the early stages of core hydrogen burning, and then equilibrate their masses through mass transfer between the constituent stars. du Buisson et al. (2020) extended the results of the MESA simulations performed by Marchant et al. (2016) and combined them with the cosmological simulations of the chemical and star formation history of the universe by Taylor & Kobayashi (2015). Their population synthesis study investigated the population properties, cosmological rates, and aLIGO detection rates of BBHs, including the dependence on the early-Universe star formation rate (SFR), which they find to be mild for moderate variations in the high-redshift SFR.

3 METHODS

In this section, we describe the implementation of CHE within the COMPAS rapid binary population synthesis code. Using COMPAS allows us to rapidly evolve a large synthetic population of binaries, which includes binaries whose component stars evolve conventionally [i.e. along a redwards track on the Hertzsprung–Russell diagram (HRD)], and others whose components evolve via CHE (i.e. along a bluewards track on the HRD), thus providing data for both pathways that can be compared directly. Below, we summarize the key physics implemented in COMPAS, starting with our approximate model of quasi-CHE based on MESA experiments, as well as the choices made for the metallicity-specific star formation history.

3.1 Physics implemented in COMPAS

The basics of stellar and binary evolution and BBH population modelling in COMPAS are described by Stevenson et al. (2017), Vigna-Gómez et al. (2018), and Neijssel et al. (2019). Here, we provide a brief summary and describe differences from previous COMPAS studies.

3.1.1 CHE in COMPAS

We used a set of MESA models of single stars with a fixed rotational frequency and no mass-loss to determine the minimal angular frequency ω necessary for CHE as a function of mass and metallicity. Our fits to these angular frequency thresholds are provided in Appendix A. COMPAS uses these fits to determine whether a star is evolving chemically homogeneously.

Stellar evolution in COMPAS follows the analytical fits of Hurley, Pols & Tout (2000) to the stellar models from Pols et al. (1998). In order to address CHE, we introduce a new CH stellar type to the Hurley et al. (2000) collection of stellar types. In our simplified

model, we neglect the very limited radial evolution of a CH star and set its radius equal to the zero-age MS (ZAMS) radius of a non-rotating star of the same mass and metallicity (see Appendix A). We compute the mass-loss rate for CH stars in the same way as for regular MS stars, but with this fixed rather than evolving radius. As a consequence, the total mass lost over the MS by CH stars in our COMPAS models is generally within $\lesssim 10$ per cent of that lost by non-CH stars of the same ZAMS mass and metallicity, except for the most massive stars in our simulations, with initial masses above $100 M_{\odot}$, where the absence of radial expansion leads to significantly reduced MS mass-loss estimates for CH stars. Finally, we assume that if a star evolves chemically homogeneously through the MS, it contracts directly into a naked helium star at the end of the MS, retaining its full mass at that point. Future evolution follows the Hurley et al. (2000) models of helium stars.

Tides are very efficient at ensuring circularization and synchronization in very close binaries through tidal locking (e.g. Hut 1981). We therefore assume that all potential candidates for CHE are tidally synchronized at birth, so that their rotational angular frequency equals the orbital angular frequency. We check this angular frequency at birth to determine whether a star belongs to the CH type and continue to check it at every time-step on the MS. If the angular frequency ever drops below the threshold value for CHE, e.g. because of binary widening as a consequence of mass-loss through winds, the star is henceforth evolved as a regular MS star (in our simplified treatment, it immediately jumps to the track of a regular MS star of the same mass). We assume that once a chemical gradient is formed, it is very challenging to overcome and ensure efficient mixing, so in our model, a star that is not evolving chemically homogeneously cannot become a CH star (cf. BPASS models, which allow quasi-CHE through accretion-induced spin-up; Eldridge et al. 2017). Although we assume perfect tidal synchronization for CH stars, we disregard the angular momentum stored in the stellar rotation when considering binary evolution with mass-loss.

3.1.2 Initial conditions

Each binary system in a COMPAS simulation is described at birth (i.e. at ZAMS) by its initial conditions: constituent star masses, separation, eccentricity, and metallicity. Initial conditions for our experiments were chosen using statistical distribution functions from the literature that were themselves based on observations. We describe the most important of these, and some important parameters that affect the evolution of the constituent stars as well as the binary system, in the following paragraphs.

The mass of the primary star in the binary system (the more massive star at ZAMS) $m_{1,i}$ is described by the Kroupa (2001) initial mass function (IMF), the distribution function of which is given by

$$p(m_{1,i}) \propto m_{1,i}^{-\alpha}, \quad (1)$$

where $\alpha = 2.3$ for the simulated range of primary masses $m_{1,i} \in [5, 150] M_{\odot}$. We assume that the IMF is the same for all metallicities.

The mass of the secondary star (less massive at ZAMS) $m_{2,i}$ is determined by drawing a mass ratio between the constituent stars $q_i \equiv m_{2,i}/m_{1,i}$ that follows a flat distribution, $p(q_i) = 1$ (Sana et al. 2012; Kobulnicky et al. 2014). Since we are interested in BBH formation, we explore only $m_{2,i} \geq 3.0 M_{\odot}$ here. However, for both the primary and secondary masses, we consider the full mass range to normalize the simulation results to a given star-forming mass or SFR (e.g. Neijssel et al. 2019).

The initial separation is drawn from a flat-in-log distribution independently of the masses (see Moe & Di Stefano 2017, for coupled

initial conditions):

$$p(a_i) \propto \frac{1}{a_i}, \quad (2)$$

where $a_i \in [0.01, 1000]$ au (Öpik 1924; Abt 1983).

We assume that all binaries are circular at birth (i.e. $e_i = 0$); see Vigna-Gómez et al. (2020) for further discussion. Close binaries are tidally circularized at birth, so this has no impact on potential CHE systems.

We simulate 30 different metallicities spaced uniformly in the logarithm across the range $-4 \leq \log_{10} Z \leq -1.825$.

3.1.3 Wind-driven mass-loss

We use the mass-loss rates as prescribed by Hurley et al. (2000) and Hurley, Tout & Pols (2002) and references therein for cooler stars with temperatures of 12 500 K and below. For stars hotter than 12 500 K, we use the wind mass-loss rates from Vink, de Koter & Lamers (2001), as implemented in Belczynski et al. (2010).

The luminous blue variable (LBV) stars (Maeder 1989; Pasquali et al. 1997), located close to the Humphreys–Davidson limit in the HRD (Humphreys & Davidson 1979), are treated differently. For these stars, we use the LBV wind mass-loss rate prescribed by Belczynski et al. (2010):

$$\frac{dM}{dt} = f_{\text{LBV}} \times 10^{-4} M_{\odot} \text{ yr}^{-1}, \quad (3)$$

where $f_{\text{LBV}} = 1.5$.

For massive, hot and bright naked helium stars, we use a metallicity-dependent Wolf–Rayet (WR) wind mass-loss rate (Vink & de Koter 2005). We parametrize the rate of mass-loss by following Belczynski et al. (2010):

$$\frac{dM}{dt} = f_{\text{WR}} \times 10^{-13} L^{1.5} \left(\frac{Z}{Z_{\odot}} \right)^m M_{\odot} \text{ yr}^{-1}, \quad (4)$$

where L is the luminosity and $m = 0.86$ (Vink & de Koter 2005); we take $Z_{\odot} = 0.014$ (Asplund et al. 2009) and $f_{\text{WR}} = 1.0$ in our default model.

In our model, all stars that remain CH on the MS convert their entire mass into helium at the end of their MS lifetime, so all such stars evolve into naked helium stars. Lower mass-loss rates would promote the formation of black holes as the end-products of the evolution of these massive stars, so we consider four different values of f_{WR} in order to study the impact of WR mass-loss on CHE: $f_{\text{WR}} \in \{0.0, 0.2, 0.6, 1.0\}$.

All mass lost in winds is assumed to promptly depart the binary without further interaction with the companion in so-called ‘Jeans mode’ mass-loss, carrying away the specific angular momentum of the donor.

Our mass-loss rate models do not include the impact of stellar rotation, which is likely to drive additional mass-loss. In particular, the WR star formed when the CH star contracts and spins up at the end of its MS is likely to be critically rotating, and the same process may repeat during core contraction after core helium depletion. The star’s angular momentum can be reduced to subcritical levels with a small amount of mass-loss, and therefore does not significantly impact the overall mass budget; however, this does affect the remnant spin (Marchant & Moriya 2020).

3.1.4 Mass transfer and overcontact systems

We use the prescriptions described in Vigna-Gómez et al. (2018) and Neijssel et al. (2019) to determine the dynamical stability of mass transfer through Roche lobe overflow (RLOF), the fraction of mass accreted on to the companion and the specific angular momentum carried away by non-conservative dynamically stable mass transfer, and the outcome of common-envelope evolution. For non-CHE binaries that go through a common-envelope phase, we assume that Hertzsprung-gap donors do not survive (the ‘pessimistic’ prescription of Belczynski et al. 2007; Neijssel et al. 2019) and we assume that immediate post-common-envelope RLOF indicates a merger.

We deviate from previous COMPAS models in the treatment of binaries that experience RLOF at ZAMS. Unlike previous work, we now allow such binaries to equilibrate their masses. The new separation of the equal-mass binary with a conserved total mass is determined by angular momentum conservation. Binary components are allowed to overfill their Roche lobes, creating overcontact systems. However, if the components extend past the L2 Lagrange points after equilibration, we assume that the binary loses co-rotation and promptly merges (Marchant et al. 2016). For equal-mass circular binaries, the volume-equivalent radius for half of the volume within the L2 equipotential surface equals half the orbital separation. Therefore, our criterion for avoiding a prompt merger is equivalent to demanding that the sum of the unperturbed stellar radii is smaller than the orbital separation a .

3.1.5 Pair-instability supernovae (SNe)

Stellar evolution models predict that single stars with helium cores in the range of ~ 60 – $130 M_{\odot}$ can become unstable due to electron–positron pair production, leading to pair-instability SNe (PISNe) that disrupt the star, leaving no remnant behind (e.g. Fowler & Hoyle 1964; Barkat, Rakavy & Sack 1967; Fraley 1968; Woosley, Heger & Weaver 2002; Farmer et al. 2019; Woosley 2019). Stars with helium cores more massive than $130 M_{\odot}$ also experience a rapid collapse driven by pair production, but in these stars photodisintegration prevents a subsequent explosion; such stars may again produce merging BBHs (e.g. Marchant et al. 2016; du Buisson et al. 2020), but they are not explored in our models, which have maximum initial stellar masses of $150 M_{\odot}$. Meanwhile, stars with somewhat lower helium core masses, between ~ 35 and $\sim 60 M_{\odot}$, are predicted to eject significant fractions of their total mass over several episodes (e.g. Yoshida et al. 2016; Woosley 2017; Marchant et al. 2019; Renzo et al. 2020). Such pulsational pair-instability SNe (PPISNe) leave behind a black hole remnant, albeit with a reduced mass. PISNe and PPISNe are expected to produce a PISNe mass gap in the distribution of remnant masses from single stellar evolution – a dearth of black holes with masses between ~ 45 and $\sim 130 M_{\odot}$.

Some superluminous SNe have been identified as PISN candidates (Gal-Yam 2012, and references therein), while iPTF2014hls has been identified as a PPISN candidate (Arcavi et al. 2017). Furthermore, the distribution of masses of gravitational-wave observations appeared consistent with a cut-off due to (P)PISNe (Abbott et al. 2019a), though GW190521 is a BBH merger with at least one component in the predicted PISN mass gap (Abbott et al. 2020).

Here, we follow the Stevenson et al. (2019) fit to the Marchant et al. (2019) models for predicting the range of PISN masses and the PPISN remnant masses from the masses of the progenitor helium cores. We apply the entire PPISN mass-loss in one time-step. Moreover, in our




Figure 1. The total SFR as a function of redshift (red) and subdivided into different ranges of metallicity, following the preferred model of Neijssel et al. (2019). The dark blue and green curves are most relevant for BBH formation.

treatment both SNe happen in one time-step for equal-mass stars. This overestimates the post-SN period and eccentricity of binaries whose components lose significant mass in a PPISN.

We use the ‘Delayed’ prescription of Fryer et al. (2012) for compact object remnant masses and modulate the natal kicks by fallback for regular core-collapse SNe, with reduced kicks for electron-capture and ultra-stripped SNe as in Vigna-Gómez et al. (2018).

3.2 Star Formation Rate

The local merger rate of BBHs depends on their formation rate at higher redshifts due to the possibly significant time delays between formation and merger, and is therefore sensitive to the SFR as a function of redshift. Furthermore, the yield of BBHs per unit star-forming mass, the BBH mass distribution, and the distribution of delay times between formation and merger are all sensitive functions of the metallicity of progenitor stars, both for CHE (e.g. Marchant et al. 2016) and non-CHE (Chruslinska, Nelemans & Belczynski 2019; Neijssel et al. 2019) systems. We must therefore specify a metallicity-specific SFR (MSSFR) in order to estimate the merger rate and properties of BBHs. We use the preferred model of Neijssel et al. (2019) for the MSSFR. Fig. 1 shows the contribution of different ranges of star formation metallicities to the total SFR. This model has higher star formation metallicities in the local Universe than the Taylor & Kobayashi (2015) model used by du Buisson et al. (2020) (cf. their fig. 2).

4 RESULTS AND DISCUSSION

We evolved a total of 12 million binaries as described in Section 3. These were equally divided into 30 metallicity bins and 4 choices of the WR mass-loss rate multipliers f_{WR} , for a total of 100 000 binaries for each of 120 combinations of Z and f_{WR} . Binaries are evolved until a double compact object is formed, or until an event happens that makes this outcome impossible (e.g. the stars merge or the binary becomes unbound), or the system reaches 14 Gyr in age.

Our simulations are based on a Monte Carlo sampling of binaries. We estimate the sampling uncertainty on all derived quantities via

Table 1. Population statistics.

	Population				
	$f_{WR} = 0$	$f_{WR} = 0.2$	$f_{WR} = 0.6$	$f_{WR} = 1.0$	Total
Number of binaries evolved	3000 000	3000 000	3000 000	3000 000	12 000 000
L2 overflow at ZAMS	243 717	243 638	243 419	243 930	974 704
Surviving binaries	2756 283	2756 362	2756 581	2756 070	11 025 296
	Surviving population				
At least one star experiencing RLOF at ZAMS	75 530	75 410	75 328	75 707	301 975
Both stars in binary CH at ZAMS	4193	4281	4201	4216	16 891
Primary only CH at ZAMS	2607	2615	2593	2604	10 419
Secondary only CH at ZAMS	0	0	0	0	0
Post-ZAMS Merger	618 001	616 749	618 425	620 250	2473 425
BBHs formed	68 231	67 200	66 016	60 294	261 741
BBHs merging in 14 Gyr	11 004	11 048	10 926	10 647	43 625
	Both stars CH at ZAMS				
At least one star experiencing RLOF at ZAMS	3661	3761	3715	3715	14 852
Both stars remained CH on MS	3444	3461	3379	3360	13 644
Primary only remained CH on MS	43	89	116	160	408
Secondary only remained CH on MS	0	0	0	0	0
Neither star remained CH on MS	706	731	706	696	2839
BBHs formed	2152	2370	2527	2621	9670
BBHs merging in 14 Gyr	2057	2322	2377	2306	9062
	Primary only CH at ZAMS				
At least one star experiencing RLOF at ZAMS	0	0	0	0	0
Primary remained CH on MS	1405	1353	1341	1337	5436
BBHs formed	0	2	3	7	12
BBHs merging in 14 Gyr	0	0	0	0	0

bootstrapping: We uniformly resample, with replacement, a new population of 12 million binaries from the original, evolved, population of 12 million. Error bars on plots, where shown, correspond to the 5th and 95th percentiles from bootstrapping.

4.1 Population statistics

The population statistics are shown in Table 1. From a population of 11 025 296 binaries that survived beyond ZAMS (i.e. did not merge at ZAMS), 16 891 were composed of two CH stars at ZAMS, with a further 10 419 composed of one CH star and one MS star at ZAMS. Furthermore, in all of the binaries with only one CH star at ZAMS it was, as we would expect, the primary, more massive, star that was CH. A total of 261 741 BBHs were formed in the simulation, but only 43 625 of these were close enough to merge within 14 Gyr, the current age of the Universe. Among the 13 644 simulated binaries that evolved chemically homogeneously throughout the MS, 9670 went on to form BBHs, the vast majority of which, 9062, merged within 14 Gyr (the few non-merging ones are those that lost significant mass in PPISNe).

4.2 Evolved system properties

Fig. 2 presents a visual summary of the evolutionary outcomes for each of the 12 million binary systems synthesized, with each point on the plot representing a single binary system, and the colour indicating the initial parameters and the outcome of the evolution (per the legend). We are particularly interested in systems for which both stars evolve chemically homogeneously and eventually collapse to form a BBH, so we have agglomerated some of the less interesting progenitor types and outcomes into groups so that the plot is not overly busy.

Because Fig. 2 is a summary over the entire grid of metallicities and WR mass-loss rate multipliers synthesized, it allows us to see on a broad scale the evolutionary outcomes for both CHE systems and non-CHE systems. The COMPAS models for the formation of non-CHE BBHs have been discussed by Neijssel et al. (2019), so we will focus our discussions hereafter on the CHE channel.

Fig. 3 shows the parameter space in which CHE is expected to occur in synchronously rotating binaries according to our CHE threshold. The darkest grey area in the lower part of the diagram indicates the region in which L2 overflow occurs and the stellar components merge; the lighter grey area in the upper part indicates the region in which the stellar components do not rotate rapidly enough to induce CHE. The central, lighter, area of the diagram indicates the region in which we expect CHE to occur, with the darker, lower, part of the central area indicating the important region of binaries whose components overflow their Roche lobes but avoid L2 overflow, occupied by the overcontact systems described by Marchant et al. (2016). This overcontact region is responsible for much of the BBH formation through CHE (cf. Fig. 2).

As expected (given our PPISN and PISN mass limits, see section 3.1.5), we see BBHs from PPISNe begin to appear at a total ZAMS mass of $\gtrsim 70 M_{\odot}$ while PISN events appear at a total ZAMS mass of $\gtrsim 120 M_{\odot}$. A few unbound CHE systems correspond to simultaneous PISNe that instantaneously removed more than half the mass of the binary in our treatment (see Section 3.1.5). In practice, such systems will undergo a series of pulsations leading to non-simultaneous mass-loss and may survive, but at separations too large to merge within the current age of the Universe. The horizontal band of PISNe just above the CH binaries in Fig. 2 are hybrid systems comprised of a CH star and an MS star, whereas the vertical band of PISNe at the upper right of the plot are systems comprised of two MS stars.




Figure 2. Initial parameters and final outcomes for each of the binary systems synthesized, showing the initial orbital period T_{ZAMS} (in days) versus the initial total mass (in M_{\odot}). The population represents a grid of 30 metallicities evenly spaced over the range $-4 \leq \log_{10}(Z) \leq -1.825$, and 4 WR mass-loss multipliers, $f_{\text{WR}} \in \{0.0, 0.2, 0.6, 1.0\}$. Regions shaded in black represent all systems that experienced L2 overflow at ZAMS; pale green represents systems that did not form BBHs. Systems for which both stars were CH at ZAMS and remained so throughout their MS lifetime are represented by regions shaded cyan if they formed BBHs via regular core-collapse SNe, blue if they formed BBHs after undergoing PPISNe, and magenta if they exploded as PISNe. Systems in which at least one of the stars did not evolve chemically homogeneously for its entire MS lifetime are represented by areas shaded light green if they formed BBHs via core-collapse SNe, dark green if they formed BBHs following PPISNe, and maroon if either star exploded as a PISN.

4.3 Population synthesis

The initial system total masses and orbital periods of CHE systems that go on to form BBHs merging within 14 Gyr are shown in Fig. 4. We show binaries evolved with the WR mass-loss multiplier $f_{\text{WR}} = 1$. Each point on the plot represents a simulated binary shaded according to its metallicity. Higher metallicity binaries are shifted towards the top of the plot. This is consistent with Fig. 3, which shows that higher metallicity stars have greater stellar radii and hence greater minimal separation, as well as lower CHE threshold rotational frequency.

Binaries with reduced WR winds have similar initial distributions, but show a clear-cut maximum total mass of $\approx 120 M_{\odot}$, which matches the mass threshold of $60 M_{\odot}$ for individual He star masses beyond which PISNe occur and leave no remnants. At higher f_{WR} , high-metallicity systems can lose a significant fraction of their mass, so binaries with initial total masses above $120 M_{\odot}$ can avoid PISNe.

To illustrate this, we plot the mass lost by a CH star with a ZAMS mass of $40.5 M_{\odot}$ over the naked helium phase in Fig. 5, for a range of WR mass-loss multipliers and metallicities. At $f_{\text{WR}} = 1$ and $Z = Z_{\odot}$, this star loses nearly half of its mass in WR winds. Meanwhile, at low metallicities, which are typical for high formation redshifts, the total mass lost in WR winds is very low, except at artificially enhanced f_{WR} values of 5 and 10, which disagree with observational constraints and are not considered in this study. Consequently, we do not expect to see a significant impact of f_{WR} on low-metallicity BBH formation, which matches our findings as discussed below.

Table 1 shows that, across all simulated metallicities and WR mass-loss multipliers, ~ 80 per cent of binaries composed of two CH stars at ZAMS retain two CH stars at the end of the MS. For binaries composed of one CH star and one MS star at ZAMS, the CH star will remain CH by the end of the MS in only ~ 50 per cent of simulations. Since we assume tidal locking in the CHE model implemented in COMPAS, as a binary widens due to mass-loss and the orbital frequency of the binary slows, the rotational frequency of the constituent CH stars slows commensurably. Binaries in which only the primary is CH at ZAMS avoided RLOF and are typically wider, so further widening through winds is more likely to spin-down the primary sufficiently to evolve off the CHE track.

Fig. 6 shows the distribution of the BBHs total masses and orbital periods just after BBH formation for systems evolving through the CHE channel. As in Fig. 4, we select only BBHs that will merge in 14 Gyr and shade binaries by metallicity. On this plot, we select $f_{\text{WR}} = 0.2$. This allows us to show not only the sharp disappearance of BBHs with total masses above $\approx 80 M_{\odot}$ due to PPISN mass-loss and complete disruption in PISNe, but also their reappearance at masses above $\approx 250 M_{\odot}$, on the other side of the ‘PISN mass gap’. There are only very few such high-mass binaries in our simulations because, with our ZAMS mass upper limit of $150 M_{\odot}$, they require very low mass-loss. Consequently, there are no such binaries in our $f_{\text{WR}} = 1.0$ simulations because their progenitors lose too much mass to remain above the PISN threshold.




Figure 3. Parameter space for equal-mass binary systems with the indicated companion mass at which CHE is expected to occur at ZAMS. Solid lines show the thresholds for CHE implemented in COMPAS (see Appendix A), dotted lines are RLOF thresholds, and dashed lines are L2 overflow thresholds. Colours differentiate metallicities. Shading corresponds to $Z = 0.001$: The dark colour at the bottom indicates L2 overflow at ZAMS, grey at the top indicates periods too low for CHE at ZAMS, and the region below the solid line and above the dashed lined indicates the possible range for CHE, with the parameter space for overcontact binaries that undergo RLOF at ZAMS shaded more darkly.




Figure 4. Initial total masses and orbital periods for CHE systems that go on to form BBHs that will merge within 14 Gyr. Each point represents one simulated binary, evolved with WR mass-loss multiplier $f_{\text{WR}} = 1$, shaded according to its metallicity.

The shortest post-BBH formation periods, and thus the shortest delay times, are seen for the lowest metallicity systems. This is due to the combined effects of their lower period at ZAMS as seen in Fig. 4 and the reduced orbital widening due to reduced mass-loss at low metallicities. However, some low-metallicity binaries lose sufficient mass in PPISNe to create wider, more eccentric binaries found towards the top of Fig. 6.




Figure 5. Total mass lost by a WR star with a ZAMS mass of $40.5 M_{\odot}$ as a function of metallicity. Line colour indicates the WR mass-loss rate multiplier (solid lines). Also shown are the mass (on the same scale as the mass-loss curves) and luminosity at the start of the WR phase as a function of metallicity (dashed lines).

4.4 Binary Black Holes

4.4.1 Formation rates

Fig. 7 shows the merging BBHs yield: the formation rate per unit star-forming mass as a function of metallicity for systems that will merge within 14 Gyr. The solid lines are the rates for the entire population – both CHE and non-CHE binaries – while the dashed




Figure 6. Total masses and orbital periods immediately after BBH formation for CHE systems that will merge within 14 Gyr. Each point represents a simulated binary, evolved with WR mass-loss multiplier $f_{\text{WR}} = 0.2$, shaded according to its metallicity. The empty area between ≈ 80 and $\approx 250 M_{\odot}$ is a consequence of systems that lost mass as PPISNe or left no remnants after exploding as PISNe.




Figure 7. Yield of BBHs that will merge within 14 Gyr per unit star-forming mass as a function of metallicity. The solid lines are the rates for the entire population – both CHE and non-CHE binaries – while the dashed lines are the rates for only the CHE binaries. Colours indicate WR mass-loss rate multipliers. Error bars indicate 90 per cent confidence intervals from sampling uncertainty.

lines are the rates for only CHE binaries. WR mass-loss multipliers are differentiated by the colour of the lines.

The overall yield of merging BBHs is quantitatively similar to the simulations of Neijssel et al. (2019), who predicted a yield of ~ 6 , 4, and 1 merging BBHs per $10^5 M_{\odot}$ of star formation at $Z = 0.01$, 0.1, and $0.3 Z_{\odot}$, respectively. The small differences are due partly to the inclusion of the CHE channel as well as PISNe and PPISNe in this work, which were not included in Neijssel et al. (2019).




Figure 8. BBH formation rate per Gpc^3 of comoving volume per year as a function of redshift for BBHs that will merge within 14 Gyr. Error bars indicate sampling uncertainty.

Meanwhile, the low-metallicity CHE channel yield of slightly less than 1 merging BBHs per $10^5 M_{\odot}$ of star formation is similar to both the Mandel & de Mink (2016) back-of-the-envelope estimate and the Marchant et al. (2016) detailed models that indicate ~ 0.7 merging BBHs below the PISN mass gap per 1000 core-collapse SNe or per $10^5 M_{\odot}$ of star formation at $Z = 0.02 Z_{\odot}$.

The paucity of CHE BBHs at high metallicity, $Z \gtrsim 0.3 Z_{\odot}$, is due primarily to a combination of the upward shifting of the allowed initial periods at higher metallicities (see Figs 3 and 4) and greater orbital widening by stronger high-metallicity winds. The increase in orbital period at BBH formation increases the delay times, preventing the BBHs from merging within 14 Gyr. The widening by mass-loss is ameliorated by reduced WR mass-loss rates. However, the WR mass-loss multipliers have negligible effect at low metallicities because the total mass-loss rate is too low even for $f_{\text{WR}} = 1$ (see Fig. 5 and associated discussion). Neijssel et al. (2019) discuss the impact of metallicity on the non-CHE BBH yield, highlighting the contributions of wind-driven widening and stellar evolutionary stage at mass transfer.

Fig. 8 shows the BBH formation rate per unit comoving volume per unit source time as a function of redshift. The formation rate for CHE BBHs peaks at $z \approx 4.25$ for $f_{\text{WR}} = 1.0$, and at $z \approx 3.5$ for other WR mass-loss multipliers for the chosen MSSFR history. The BBH formation rate for both CHE and non-CHE channels peaks at higher redshifts than the assumed SFR because both have higher yields per unit star formation at lower metallicities, which are prevalent at higher redshifts.

4.4.2 Merger delay times

Fig. 9 indicates the distribution of delay times between star formation and BBH mergers. This figure combines all metallicities with equal weights, without considering their contribution to the observable systems, so should be viewed as an indicative sketch.

Non-CHE binaries in Fig. 9 have a very broad distribution of delay times. Some are very short, less than 10 Myr, due to significant hardening during mass transfer episodes, including through dynamically unstable mass transfer and common-




Figure 9. Distribution of delay times between formation and merger for BBHs. All metallicities from the simulation are combined with equal weights and arbitrary counts per uniform bins in log delay time are shown. Error bars indicate sampling uncertainty.

envelope ejection, as well as fortuitously directed SN natal kicks. Meanwhile, there is an almost flat tail of long delay times on this logarithmic plot, corresponding to a $p(\tau_{\text{delay}}) \sim 1/\tau_{\text{delay}}$ distribution.

On the other hand, binaries formed through CHE are seen to have a more strongly clustered delay time distribution, with typical delay times of between 100 Myr and 1 Gyr. There are no ultra-short delay times because, with the exception of RLOF at ZAMS, such binaries do not undergo mass transfer that could harden the binary. Moreover, the high masses of CHE stars imply that they do not experience asymmetric SNe and associated natal kicks in the COMPAS model.

The smallest time delay between formation and merger for CHE systems in our simulations ranges from ~ 0.025 Gyr for $f_{\text{WR}} = 0.0$ to ~ 0.033 Gyr for $f_{\text{WR}} = 1.0$. The combination of lower metallicities and reduced mass-loss rates yields the shortest delay times, allowing binaries to start evolution from closer separations while avoiding L2 overflow and to avoid subsequent widening through mass-loss. This is consistent with the minimal delay times found in other studies. Mandel & de Mink (2016), who consider only $Z = 0.004 Z_{\odot}$, estimate minimum delay times of ~ 3.5 Gyr. Marchant et al. (2016) find minimal delay times of ~ 0.4 Gyr and point out the metallicity dependence. du Buisson et al. (2020) consider the lowest metallicities among these studies, $Z = 10^{-5}$, and find the shortest delay times, ~ 0.02 Gyr.

Some CHE binaries will be significantly widened by mass-loss, potentially losing up to a factor of ~ 2 in mass during the WR phase (see Fig. 5) and thereby increasing their separation by the same factor. The gravitational-wave-driven coalescence time scales as $a^4 M^{-3}$ (Peters 1964), so a factor of 2 each in mass decrease and semimajor axis increase would yield a factor of $2^7 \sim 100$ increase in the delay time. This explains the long delay time tail of the CHE BBH distribution, as well as the decrease in the prominence of this tail as the WR wind mass-loss multiplier is reduced. Even when $f_{\text{WR}} = 0$, some CHE BBHs will have long delay times due to the mass lost in PPISNe.




Figure 10. BBH merger rate per Gpc^3 of comoving volume per year of source time as a function of redshift. Error bars indicate sampling uncertainty.




Figure 11. The merger rate of BBHs detectable by aLIGO at final design sensitivity, as a function of merger redshift. Error bars indicate sampling uncertainty.

4.4.3 Merger rates

Fig. 10 shows the BBH merger rate per Gpc^3 of comoving volume per year of source time as a function of redshift. The merger rate for CHE BBHs peaks at $z \approx 4$ for $f_{\text{WR}} = 1.0$, and at $z \approx 3$ for other WR mass-loss multipliers. Both CHE and total BBH merger rates peak at higher redshifts than the SFR, which peaks at $z \approx 2$ (see Fig. 1), because both CHE and non-CHE channels have higher yields at lower metallicity (see Fig. 7). The relatively small difference between the peak formation and merger rates is explained by the short delay times for the CHE systems (see Figs 8 and 9). The delay times are particularly short for $f_{\text{WR}} = 0$ CHE BBHs, which explains their suppressed merger rate in the local Universe.

The merger rates of BBHs that could be observed by aLIGO operating at final design sensitivity merger rates are shown in Fig. 11. Binaries formed through CHE have higher average masses than non-CHE binaries, which increases the range within which they




Figure 12. Cumulative BBH detections as a function of merger redshift, per year of observing at aLIGO O1 sensitivity. Error bars indicate sampling uncertainty.




Figure 13. Cumulative BBH detections as a function of merger redshift, per year of observing at aLIGO final design sensitivity. Error bars indicate sampling uncertainty.

are detectable by aLIGO. Therefore, CHE BBHs make up a higher fraction of all detections at greater redshifts.

4.4.4 aLIGO detection rates

Figs 12 and 13 show the predicted cumulative detection rates per year of observing time as a function of redshift for aLIGO O1 and final design sensitivities, respectively.

Fig. 12 shows that the total expected detection rate at O1 sensitivity is 38–55 detections per year, depending on the assumed value of the WR mass-loss rate multiplier. This would correspond to 17–25 detections over the 166 d of coincident data over the first two advanced detector observing runs, namely O1 and O2. In fact, only 10 BBHs were observed during this time (Abbott et al. 2019b).

The increased detection rate relative to the preferred MSSFR model of Neijssel et al. (2019), who predicted 22 detections per

year, in agreement with the O1 and O2 observations, is due to the contribution of CHE BBHs. CHE BBHs may constitute up to ~70 per cent of all BBH detections at both the O1 sensitivity and the final design sensitivity of aLIGO.

The star formation history model of Neijssel et al. (2019) was tuned to the gravitational-wave observations, and explaining the relatively high masses of observed BBHs required significant high-redshift, low-metallicity star formation. The inclusion of CHE BBHs naturally yields a population of high-mass sources, allowing for the high-mass SFR to be reduced in line with the Madau & Dickinson (2014) and Madau & Fragos (2017) models. This would naturally bring rate predictions in line with the O1 and O2 observations and correspondingly reduce predicted detection rates for future detectors.

Using the preferred cosmic metallicity star formation model of Neijssel et al. (2019), as we do here, and assuming $f_{\text{WR}} = 1$, we predict a total BBH detection rate of ≈ 660 per year at aLIGO design sensitivity (versus ≈ 37 at O1 sensitivity), with ≈ 470 (≈ 27) of these coming from the CHE channel. The CHE BBH detection rates are a factor of ~ 2 larger than those estimated by du Buisson et al. (2020), who found that ≈ 250 (≈ 13) CHE BBHs per year may be detected at aLIGO design (O1) sensitivity. The differences in the assumed MSSFRs in these studies are responsible for much of this difference.

We note that in both Figs 12 and 13 the order of the lines with respect to the number of detections does not match the order of the WR mass-loss multipliers. This is due to the interplay between the formation rate of BBHs and their delay times as a function of f_{WR} , which are described in Figs 8 and 9 and associated discussion. For example, in the absence of WR winds ($f_{\text{WR}} = 0$), reduced delay times due to a lack of binary widening relative to simulations with higher WR mass-loss rates mean that very few CHE BBHs, which predominantly form at lower metallicities and thus higher redshifts, merge in the local Universe, where they would be detectable.

4.4.5 Mass distribution of detectable BBH mergers

The cumulative distribution functions for the modelled chirp mass distribution of detectable BBH mergers are shown in Fig. 14. The dark blue lines indicate the chirp mass distribution of all BBHs while the light blue lines indicate the chirp mass distribution of CHE BBHs. In both cases, results for the WR mass-loss multiplier $f_{\text{WR}} = 1.0$ are reported, and the O1 aLIGO sensitivity, which is similar to that of the second observing run, is used. We show cumulative distribution functions for sets of 10 randomly selected samples from the COMPAS models – the number of BBHs detected during O1 and O2 – in order to indicate the variation due to sampling fluctuations. To avoid granularity due to the discreteness of the metallicity grid in COMPAS models (see Dominik et al. 2015; Neijssel et al. 2019, for a discussion), we used continuous sampling in metallicity to construct the model predictions for this plot.

As mentioned previously, CHE BBHs are more massive than typical non-CHE BBHs. The initial masses of CHE BBHs must be high to allow for CHE (see Fig. 3). Moreover, CHE in our model allows stars to convert all of their mass to helium, whereas non-CH massive stars typically have $\gtrsim 50$ per cent of their mass in hydrogen-rich envelopes, which they lose prior to collapse into black holes in the course of binary evolution. This is highlighted in Fig. 15, which indicates the fraction of all BBHs detectable at aLIGO O1 sensitivity that formed through the CHE channel, as a function of chirp mass. The CHE channel dominates the production of BBHs at high chirp masses, particularly for reduced WR mass-loss models, when it yields increasingly large chirp masses ($\gtrsim 30 M_{\odot}$ in the absence of WR winds).




Figure 14. Chirp mass posteriors for the 10 BBH mergers detected during the first and second aLIGO observing runs (Abbott et al. 2019a) are shown in colour at the bottom, with labels at top. These are randomly sampled to construct the cumulative density functions shown in lavender (each curve corresponds to a cumulative distribution through 10 samples, one from each posterior). Cumulative density functions for COMPAS chirp mass predictions based on $f_{\text{WR}} = 1.0$ models are also based on 10 samples from either the full population (dark blue lines) or CHE systems only (light blue lines).




Figure 15. The fraction of BBHs formed through the CHE channel among all BBHs detectable at aLIGO O1 sensitivity, plotted as a function of chirp mass.

Fig. 14 allows for a direct comparison between the modelled chirp mass distribution and the aLIGO observations from the first two observing runs. The individual posterior samples from the 10 aLIGO BBH detections during those observing runs are plotted at the bottom of the plot. Randomly sampled cumulative distribution functions of the chirp mass of observed events are constructed by taking 10 random samples, one from each of the 10 aLIGO observation posteriors and displayed as light lavender curves. The overlap of the lavender and dark blue lines in Fig. 14 shows that the COMPAS model

of BBH formation, which includes the contribution of CHE, yields a chirp mass distribution of detectable BBH mergers that is consistent with detections during the first two aLIGO observing runs.

5 CONCLUDING REMARKS

We described the model of CHE that we implemented in the rapid population synthesis code COMPAS. We used MESA models to determine the critical rotation thresholds for CHE, and provided fits that can be used in other rapid binary population synthesis applications. We synthesized 12 million binary systems over a range of metallicities (30 metallicities evenly spaced across the range $-4 \leq \log_{10} Z \leq -1.825$) and WR wind mass-loss multipliers ($f_{\text{WR}} \in \{0.0, 0.2, 0.6, 1.0\}$). We confirmed that our simplified models match detailed binary evolution simulations (Marchant et al. 2016; du Buisson et al. 2020) well.

We investigated the contribution of CHE and non-CHE channels to BBH formation under a single set of assumptions. We found that the CHE channel may contribute more than half, and perhaps as much as three quarters, of all aLIGO BBH detections arising from isolated binary evolution. CHE BBHs may represent $\gtrsim 80$ per cent of detectable sources with the highest chirp masses of $\gtrsim 30 M_{\odot}$. A comparison between our model population and the population of detected binaries from the first two advanced detector observing runs indicates that the current model overpredicts the total number of sources by a factor of ~ 2 , but matches the observed chirp mass distribution.

We made a number of simplifying assumptions in this study that can be investigated and improved on in the future. We generally erred on the side of being conservative about CHE predictions:

(i) We used Hurley et al. (2000) non-rotating MS models to set the radii and mass-loss rates of CH MS stars. The imperfect radius model for rapidly rotating CH stars in turn leads to differences in the orbital separation boundary for avoiding merger through L2 overflow between the COMPAS and MESA models (see Appendix A).

(ii) We used simplified tidal interaction assumptions under which CH stars are immediately tidally synchronized, yet do not store angular momentum. Accounting for the angular momentum stored in stars – and the additional angular momentum carried away by winds from a rotating star – impacts the response of the binary’s orbit to mass-loss, and reduces the amount of orbital widening by wind mass-loss in close binaries.

(iii) Contrary to our simplified assumptions, winds may interact with binary companions. This is particularly true in close binaries, when the wind speeds are comparable to the orbital speeds, and wind interactions may produce additional drag and reduce the amount of orbital widening (e.g. Brookshaw & Tavani 1993; MacLeod & Loeb 2020).

(iv) We ignored the possibility of initially non-CH stars switching to CHE in response to mass accretion.

(v) We assumed that all mass-loss in PPISNe happens instantaneously, rather than over several pulsations (although the first pulsation is likely to be dominant, so this approximation may not be especially problematic).

Our predicted BBH merger rate at redshift zero of $50 \text{ Gpc}^{-3} \text{ yr}^{-1}$ (including $20 \text{ Gpc}^{-3} \text{ yr}^{-1}$ from the CHE channel) for the default WR mass-loss rate $f_{\text{WR}} = 1.0$ overestimates the number of BBH detections during the first two observing runs of gravitational-wave detectors. This is at least partly due to our using an MSSFR prescription from Neijssel et al. (2019) that was designed to reproduce gravitational-wave observations without accounting for CHE. A resolution may

involve reducing the high-redshift SFR back to levels more closely matching the models of Madau & Dickinson (2014) and Madau & Fragos (2017).

Two other observational constraints on CHE BBH formation come from the spins of observed BBH mergers and from potential electromagnetic observations of their progenitors. At first glance, the effective spins of BBHs observed to date (Abbott et al. 2019b) do not match the large reservoirs of angular momentum in CH stars. However, WR winds can carry away much of angular momentum. Marchant et al. (2016) argued that typical dimensionless effective spins of CHE sources should be ~ 0.4 , much lower than the supercritical spins expected at WR star formation. The fraction of stellar angular momentum lost in winds during the WR phase can be estimated as

$$\frac{\Delta L}{L} \sim \frac{2}{3} \left(\frac{R_{\text{WR}}}{R_{\text{WR,g}}} \right)^2 \frac{\Delta M}{M}, \quad (5)$$

where the ratio of the radius of the WR star to its gyration radius is $R_{\text{WR}}/R_{\text{WR,g}} \sim 10$. Thus, WR winds could lose the overwhelming bulk of the angular momentum that CHE stars have, as long as $\Delta M/M > 0.01$, which is true even at $Z = 0.01 Z_{\odot}$ if WR mass-loss is not suppressed (see Fig. 5). Unlike binaries that are hardened during the common-envelope phase to the point where tides can efficiently spin-up the WR companion (Kushnir et al. 2016; Bavera et al. 2020; Belczynski et al. 2020), binaries evolving through the CHE channel will tidally decouple during the WR phase. For example, a typical binary from our simulations with component ZAMS masses of $60 M_{\odot}$ at metallicity $Z = 0.00089$ and an initial orbital period of just over 1 d will evolve through the CHE channel and form a merging BBH with individual black hole masses of $38 M_{\odot}$. By the end of the WR phase, following mass-loss with $f_{\text{WR}} = 1.0$, the components will have masses of $52 M_{\odot}$ and an orbital separation of $25 R_{\odot}$. The WR radius at this time is $< 2 R_{\odot}$ (e.g. Yoon, Dierks & Langer 2012), so the tidal synchronization time-scale will be several hundred Myr. This is much longer than the duration of the WR phase, so once spun down by winds, these stars cannot be spun up again by tides.

CHE BBH progenitors could yield interesting observational candidates. Systems such as WR20a (Rauw et al. 2004) and BAT99-32 (Shenar et al. 2019) may belong in this category. The metallicity of the Galaxy is too high to allow for merging CHE BBHs according to our models, but we expect them to be formed at a rate of $\sim 3 \times 10^{-6}$ per year in the Magellanic clouds. Given the typical MS and WR phase lifetimes of 3×10^6 and 3×10^5 yr, respectively, we may hope to detect ~ 10 MS CH binaries and ~ 1 binary composed of two naked helium stars formed through CHE and en route to BBH formation in the Magellanic clouds today.

The joint model for the classical and CHE isolated binary evolution channels developed here will enable simultaneous inference on binary evolution model parameters and the metallicity-specific star formation history once the full trove of observations from the third gravitational-wave observing run is available. Ultimately, the relatively short delay times of CHE BBHs make them ideal probes of high-redshift star formation history, while their high masses make them perfect targets for third-generation gravitational-wave detectors with good low-frequency sensitivity, such as the Einstein Telescope (Punturo et al. 2010) or the Cosmic Explorer (Abbott et al. 2017).

ACKNOWLEDGEMENTS

Simulations in this paper made use of the COMPAS rapid population synthesis code, which is freely available at <http://github.com/TeamCOMPAS/COMPAS>. The version of COMPAS used for these simulations was v02.11.01a, built specifically for these simulations;

functionality in this release was integrated into the public COMPAS code base in v02.11.04.

The authors thank Selma de Mink and other colleagues in Team COMPAS, as well as Morgan MacLeod, for helpful discussions. We also thank Tim Riley for assistance in running COMPAS simulations. IM is a recipient of the Australian Research Council Future Fellowship FT190100574. AVG acknowledges funding support by the Danish National Research Foundation (DNRF132).

DATA AVAILABILITY

The data underlying this article will be available via <https://zenodo.org/communities/compas/>.

REFERENCES

- Abbott B. P. et al., 2016, *Phys. Rev. Lett.*, 116, 061102
 Abbott B. P. et al., 2017, *Class. Quantum Gravity*, 34, 044001
 Abbott B. P. et al., 2019a, *Phys. Rev. X*, 9, 031040
 Abbott B. P. et al., 2019b, *ApJ*, 882, L24
 Abbott R. et al., 2020, *Phys. Rev. Lett.*, 125, 101102
 Abt H. A., 1983, *ARA&A*, 21, 343
 Almeida L. A. et al., 2015, *ApJ*, 812, 102
 Angulo C. et al., 1999, *Nucl. Phys. A*, 656, 3
 Antonini F., Chatterjee S., Rodriguez C. L., Morscher M., Pattabiraman B., Kalogera V., Rasio F. A., 2016, *ApJ*, 816, 65
 Arcavi I. et al., 2017, *Nature*, 551, 210
 Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481
 Barkat Z., Rakavy G., Sack N., 1967, *Phys. Rev. Lett.*, 18, 379
 Bartos I., Kocsis B., Haiman Z., Márka S., 2017, *ApJ*, 835, 165
 Bavera S. S. et al., 2020, *A&A*, 647, A153
 Belczynski K., Bulik T., Fryer C. L., Ruiter A., Valsecchi F., Vink J. S., Hurley J. R., 2010, *ApJ*, 714, 1217
 Belczynski K., Taam R. E., Kalogera V., Rasio F. A., Bulik T., 2007, *ApJ*, 662, 504
 Belczynski K. et al., 2020, *A&A*, 636, A104
 Brookshaw L., Tavani M., 1993, *ApJ*, 410, 719
 Brott I. et al., 2011, *A&A*, 530, A115
 Caughlan G. R., Fowler W. A., 1988, *At. Data Nucl. Data Tables*, 40, 283
 Chaboyer B., Zahn J. P., 1992, *A&A*, 253, 173
 Chruslinska M., Nelemans G., Belczynski K., 2019, *MNRAS*, 482, 5012
 de Mink S. E., Cantiello M., Langer N., Pols O. R., Brott I., Yoon S. C., 2009, *A&A*, 497, 243
 de Mink S. E., Mandel I., 2016, *MNRAS*, 460, 3545
 Dominik M. et al., 2015, *ApJ*, 806, 263
 du Buisson L. et al., 2020, *MNRAS*, 499, 5941
 Eddington A. S., 1929, *MNRAS*, 90, 54
 Eldridge J. J., Stanway E. R., Xiao L., McClelland L. A. S., Taylor G., Ng M., Greis S. M. L., Bray J. C., 2017, *Publ. Astron. Soc. Aust.*, 34, e058
 Farmer R., Renzo M., de Mink S. E., Marchant P., Justham S., 2019, *ApJ*, 887, 53
 Fowler W. A., Hoyle F., 1964, *ApJS*, 9, 201
 Fraley G. S., 1968, *Ap&SS*, 2, 96
 Fryer C. L., Belczynski K., Wiktorowicz G., Dominik M., Kalogera V., Holz D. E., 2012, *ApJ*, 749, 91
 Gal-Yam A., 2012, *Science*, 337, 927
 Grevesse N., Sauval A. J., 1998, *Space Sci. Rev.*, 85, 161
 Heger A., Langer N., Woosley S. E., 2000, *ApJ*, 528, 368
 Humphreys R. M., Davidson K., 1979, *ApJ*, 232, 409
 Hurley J. R., Pols O. R., Tout C. A., 2000, *MNRAS*, 315, 543
 Hurley J. R., Tout C. A., Pols O. R., 2002, *MNRAS*, 329, 897
 Hut P., 1981, *A&A*, 99, 126
 Iglesias C. A., Rogers F. J., 1996, *ApJ*, 464, 943
 Kippenhahn R., 1974, in Tayler R. J., Hesser J. E., eds, *Proc. IAU Symp.* 66, Late Stages of Stellar Evolution. Reidel, Dordrecht, p. 20
 Kobulnicky H. A. et al., 2014, *ApJS*, 213, 34

- Koenigsberger G., Morrell N., Hillier D. J., Gamen R., Schneider F. R. N., González-Jiménez N., Langer N., Barbá R., 2014, *AJ*, 148, 62
- Kroupa P., 2001, *MNRAS*, 322, 231
- Kushnir D., Zaldarriaga M., Kollmeier J. A., Waldman R., 2016, *MNRAS*, 462, 844
- Langer N., 1992, *A&A*, 265, L17
- MacLeod M., Loeb A., 2020, *ApJ*, 902, 85
- Madau P., Dickinson M., 2014, *ARA&A*, 52, 415
- Madau P., Fragos T., 2017, *ApJ*, 840, 39
- Maeder A., 1987, *A&A*, 178, 159
- Maeder A., 1989, in Davidson K., Moffat A. F. J., Lamers H. J. G. L., eds, *Astrophysics and Space Science Library*, Vol. 157, *On the Evolutionary Status and Instability Mechanism of the Luminous Blue Variables*. Springer-Verlag, Berlin, p. 15
- Maeder A., Meynet G., 2000, *ARA&A*, 38, 143
- Mandel I., de Mink S. E., 2016, *MNRAS*, 458, 2634
- Mandel I., Farmer A., 2018, preprint (arXiv:1806.05820)
- Mapelli M., 2018, preprint (arXiv:1809.09130)
- Marchant P., Langer N., Podsiadlowski P., Tauris T. M., Moriya T. J., 2016, *A&A*, 588, A50
- Marchant P., Moriya T. J., 2020, *A&A*, 640, L18
- Marchant P., Renzo M., Farmer R., Pappas K. M. W., Taam R. E., de Mink S. E., Kalogera V., 2019, *ApJ*, 882, 36
- Miller M. C., Lauburg V. M., 2009, *ApJ*, 692, 7
- Moe M., Di Stefano R., 2017, *ApJS*, 230, 15
- Neijssel C. J. et al., 2019, *MNRAS*, 490, 3740
- Öpik E., 1924, *Publ. Tartu Astrofizika Obs.*, 25, 1
- Pasquali A., Langer N., Schmutz W., Leitherer C., Nota A., Hubeny I., Moffat A. F. J., 1997, *ApJ*, 478, 340
- Paxton B., Bildsten L., Dotter A., Herwig F., Lesaffre P., Timmes F., 2011, *ApJS*, 192, 3
- Paxton B. et al., 2013, *ApJS*, 208, 4
- Paxton B. et al., 2015, *ApJS*, 220, 15
- Peters P. C., 1964, *Phys. Rev.*, 136, B1224
- Petrovic J., Langer N., Yoon S. C., Heger A., 2005, *A&A*, 435, 247
- Polis O. R., Schröder K.-P., Hurley J. R., Tout C. A., Eggleton P. P., 1998, *MNRAS*, 298, 525
- Potekhin A. Y., Chabrier G., 2010, *Contrib. Plasma Phys.*, 50, 82
- Punturo M. et al., 2010, *Class. Quantum Gravity*, 27, 084007
- Rauw G. et al., 2004, *A&A*, 420, L9
- Renzo M., Farmer R., Justham S., Götberg Y., de Mink S. E., Zapartas E., Marchant P., Smith N., 2020, *A&A*, 640, A56
- Rodríguez C. L., Morscher M., Pattabiraman B., Chatterjee S., Haster C.-J., Rasio F. A., 2015, *Phys. Rev. Lett.*, 115, 051101
- Rogers F. J., Nayfonov A., 2002, *ApJ*, 576, 1064
- Sana H. et al., 2012, *Science*, 337, 444
- Saumon D., Chabrier G., van Horn H. M., 1995, *ApJS*, 99, 713
- Shenar T. et al., 2019, *A&A*, 627, A151
- Sigurdsson S., Hernquist L., 1993, *Nature*, 364, 423
- Song H. F., Meynet G., Maeder A., Ekström S., Eggenberger P., 2016, *A&A*, 585, A120
- Stevenson S., Sampson M., Powell J., Vigna-Gómez A., Neijssel C. J., Szécsi D., Mandel I., 2019, *ApJ*, 882, 121
- Stevenson S., Vigna-Gómez A., Mandel I., Barrett J. W., Neijssel C. J., Perkins D., de Mink S. E., 2017, *Nat. Commun.*, 8, 14906
- Stone N. C., Metzger B. D., Haiman Z., 2017, *MNRAS*, 464, 946
- Sweet P. A., 1950, *MNRAS*, 110, 548
- Taylor P., Kobayashi C., 2015, *MNRAS*, 448, 1835
- Timmes F. X., Swesty F. D., 2000, *ApJS*, 126, 501
- Tutukov A., Yungelson L., 1973, *Nauchnye Inf.*, 27, 70
- Tutukov A. V., Yungelson L. R., 1993, *MNRAS*, 260, 675
- van den Heuvel E. P. J., 1976, in Eggleton P., Mitton S., Whelan J., eds, *Proc. IAU Symp. 73, Structure and Evolution of Close Binary Systems*. Kluwer, Dordrecht, p. 35
- Vigna-Gómez A., MacLeod M., Neijssel C. J., Broekgaarden F. S., Justham S., Howitt G., de Mink S. E., Mandel I., 2020, *Publ. Astron. Soc. Aust.*, 37, e038
- Vigna-Gómez A. et al., 2018, *MNRAS*, 481, 4009
- Vink J. S., de Koter A., 2005, *A&A*, 442, 587
- Vink J. S., de Koter A., Lamers H. J. G. L. M., 2001, *A&A*, 369, 574
- von Zeipel H., 1924, *MNRAS*, 84, 665
- Woosley S. E., 2017, *ApJ*, 836, 244
- Woosley S. E., 2019, *ApJ*, 878, 49
- Woosley S. E., Heger A., Weaver T. A., 2002, *Rev. Mod. Phys.*, 74, 1015
- Yoon S.-C., Langer N., Norman C., 2006, *A&A*, 460, 199
- Yoon S. C., Dierks A., Langer N., 2012, *A&A*, 542, A113
- Yoshida T., Umeda H., Maeda K., Ishii T., 2016, *MNRAS*, 457, 351
- Ziosi B. M., Mapelli M., Branchesi M., Tormen G., 2014, *MNRAS*, 441, 3703

APPENDIX A: CHE THRESHOLDS

We evolved single stars over a range of masses, metallicities, and rotational frequencies with version 10108 of the MESA code in order to find the boundary between CHE and regular non-CH stellar evolution.¹ Simulations were performed until the end of the MS without mass-loss, while enforcing solid body rotation at a constant angular frequency.

Opacities are computed using tables from the OPAL project (Iglesias & Rogers 1996) with solar-scaled metal mass fractions as given by Grevesse & Sauval (1998). The equation of state is a combination of the OPAL (Rogers & Nayfonov 2002), HELM (Timmes & Swesty 2000), PC (Potekhin & Chabrier 2010), and SCVH (Saumon, Chabrier & van Horn 1995) equations of state. Nuclear reaction rates are taken from Caughlan & Fowler (1988) and Angulo et al. (1999) with preference for the latter when available.

Our choices for overshooting and rotational mixing processes follow those of Brott et al. (2011). Namely, overshooting from convective hydrogen-burning cores is modelled as step overshooting, increasing the size of the convective core by $0.335H_p$, where H_p is the pressure scale height at the edge of the convective boundary. As we consider solid body rotation, the only significant mixing process included in our simulations is the effect of Eddington–Sweet circulations as described by Kippenhahn (1974), with an efficiency factor of 1/30 (Chaboyer & Zahn 1992; Heger, Langer & Woosley 2000). We also include the inhibiting effect of composition gradients in rotational mixing as described by Heger et al. (2000), given by the dimensionless parameter $f_\mu = 0.1$ (Yoon, Langer & Norman 2006). The star was considered to evolve chemically homogeneously if the difference between the helium fraction across the star did not exceed 0.2.

Fig. A1 shows the maximum rotational frequency at which the star remains non-CH (upward triangles), and the minimum rotational frequency at which the star becomes CH (downward triangles), for a grid of masses ranging from 10 to $150 M_\odot$ and three metallicities, $Z = 0.01, 0.001, \text{ and } 0.0001$.

The following fits for the threshold angular frequency for CHE are implemented in COMPAS and shown in Fig. A1:

$$\omega_{M,Z} = \frac{\omega_{M,Z,0.004}}{0.09 \ln \left(\frac{Z}{0.004} \right) + 1}, \quad (\text{A1})$$

where

$$\omega_{M,Z,0.004} = \begin{cases} \sum_{i=0}^5 a_i \frac{M^i}{M^{0.4}} & \text{rad s}^{-1}, M \leq 100 M_\odot \\ \sum_{i=0}^5 a_i \frac{100^i}{M^{0.4}} & \text{rad s}^{-1}, M > 100 M_\odot \end{cases} \quad (\text{A2})$$

and

¹The complete set of MESA input files necessary to reproduce these simulations will be made available after acceptance of the manuscript.




Figure A1. Rotational frequency threshold for CHE as a function of mass and metallicity. Downward and upward triangles represent the slowest rotating CHE model and fastest rotating non-CHE model at the given mass and metallicity, respectively. The curves indicate the fits of equations (A1) and (A2).

$$a_0 = 5.7914 \times 10^{-4}$$

$$a_1 = -1.9196 \times 10^{-6}$$

$$a_2 = -4.0602 \times 10^{-7}$$

$$a_3 = 1.0150 \times 10^{-8}$$

$$a_4 = -9.1792 \times 10^{-11}$$

$$a_5 = 2.9051 \times 10^{-13}.$$

We expect these fits to be valid over the range where they are constructed ($10 M_{\odot} \leq M \leq 150 M_{\odot}$, $10^{-4} \leq Z \leq 0.01$) but caution should be exercised if the fits are extrapolated significantly beyond these boundaries.

Fig. A2 shows the range of binaries in mass–orbital period space that leads to CHE at $Z = 0.001$. Red points indicate the population of binaries that remain CH through the MS according to the model described here and implemented in COMPAS. For comparison, the background colours show the outcomes from detailed MESA binary models of du Buisson et al. (2020), with cyan, purple, and blue denoting binaries that undergo CHE on the MS without merging. In general, there is good agreement between the two sets of models, particularly at higher orbital periods. For the tightest orbits, du Buisson et al. (2020) models predict mergers through L2 overflow when our COMPAS fits suggest that the overcontact binary may still survive with CHE stars. This is likely due to the combination of rotational deformation and tidal deformation, as well as some mild expansion early during CHE evolution, that are not accounted for in the COMPAS models.




Figure A2. Comparison between the COMPAS mass and orbital period range leading to CHE, implemented as described in this appendix (red dots), and the detailed MESA binary models of du Buisson et al. (2020) (background colour shading of cyan, blue, and purple indicates binaries that survive CHE on the MS) at $Z = 0.001$.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.

Chapter 4

Surrogate forward models for population inference on compact binary mergers

As previously noted, observations of mergers of compact binaries can help us constrain the uncertain parameters that describe the underlying processes in the evolution of stars and binary systems. A naïve, brute-force, method to determine constraints on these parameters is to use a population synthesis computer model to generate a synthetic state space by integrating over all possible values of initial conditions and evolutionary parameters, then search that synthetic state space for states that correspond to observed data - but even an indicative selection of that parameter space is very large, and generating such a large state space with current modelling tools (e.g. COMPAS) is computationally intractable.

The goal of this proof-of-concept study was to produce a surrogate model - an interpolant - that worked, not necessarily the most efficient or accurate interpolant. "Good enough" accuracy was all that we required - our goal was to prove that a surrogate model could be constructed that is able to interpolate over the range of values for which it was trained, and that is reasonably accurate and a good deal faster than the population synthesis model it was built to emulate. Optimisation, with respect to both accuracy and speed, would come later.

For this study we focussed on just four of the metallicity specific star formation rate (MSSFR) parameters used in the post-processing of COMPAS simulations ($(\alpha, \lambda, a_{SF}, d_{SF})$) - see Neijssel et al. [48] for details).

We developed a surrogate model that is able to quickly (~ 200 faster than COMPAS), and with fairly good accuracy, simulate a state space large enough to allow us to infer constraints for some of the parameters of interest. We constructed the surrogate model using ANNs to predict binned merger detection rates for double compact objects (DCOs) over a range of chirp mass and redshift values.

We trained the interpolant by using COMPAS to generate a smaller number of states, using a targeted selection of initial conditions and evolutionary parameters, and then used that simulated state space as training data to teach the constituent ANNs how to map initial conditions and evolutionary parameters to the final state, thus avoiding the computational time of calculating the intermediate steps.

Using our trained surrogate model we generated a very large simulated state space and searched that state space for states that matched observed gravitational-wave data.

The following paper, Riley and Mandel [89], describes the development of our surrogate model.

Next steps for this work include optimising the performance of the surrogate model, with respect to both accuracy and speed, and expanding the surrogate model to other astrophysical parameters that govern stellar and binary evolution in the population synthesis code itself (e.g. COMPAS), rather than just the post-processing code.

Surrogate forward models for population inference on compact binary mergers

JEFF RILEY^{1,2} AND ILYA MANDEL^{1,2}

¹*School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia*

²*ARC Centre of Excellence for Gravitational Wave Discovery – OzGrav, Australia*

ABSTRACT

Rapidly growing catalogs of compact binary mergers from advanced gravitational-wave detectors allow us to explore the astrophysics of massive stellar binaries. Merger observations can constrain the uncertain parameters that describe the underlying processes in the evolution of stars and binary systems in population models. In this paper, we demonstrate that binary black hole populations – namely, detection rates, chirp masses, and redshifts – can be used to measure cosmological parameters describing the redshift-dependent star formation rate and metallicity distribution. We present a method that uses artificial neural networks to emulate binary population synthesis computer models, and construct a fast, flexible, parallelisable surrogate model that we use for inference.

Keywords: surrogate model – inference – binaries – stars: evolution – gravitational waves – machine learning – artificial neural networks

1. INTRODUCTION

Rapid binary population synthesis software packages (e.g., BSE (Hurley et al. 2002), StarTrack (Belczynski et al. 2008), binary_c (Izzard et al. 2018), SEVN (Iorio et al. 2022), COMPAS (Team COMPAS: Riley, J. et al. 2022a), cosmic (Breivik et al. 2020)) have a number of free parameters that allow modellers to:

- set the initial conditions of the simulation (e.g. initial stellar mass, metallicity, separation (for binary stars), etc.), and
- determine the nature of some of the (simulated) physical processes that a star, single or binary, undergoes during its evolution through time (e.g. wind mass loss rate, supernova remnant masses and kicks, mass transfer and common envelope efficiency for binary stars).

Many of these free parameters and the physical states and processes they represent are so far not well constrained by observation.

Our ultimate goal is to develop a method to determine constraints for (some of) the physical states and processes modelled by various modelling and simulation

software packages, thus not only learning more about the physical nature of these states and processes, but also improving our ability to model them. A naïve, brute-force, method to determine such constraints is to generate a synthetic state space by integrating over all possible values of initial conditions and evolutionary parameters, then search that synthetic state space for states that correspond to observed data. However, even an indicative selection of that parameter space is very large (e.g., Broekgaarden et al. (2022) considered 560 distinct model variations), and generating such a large state space with current modelling tools is computationally intractable.

As a first step towards achieving our goal, we develop a tool that can quickly simulate a state space large enough to allow us to infer astrophysical constraints. We can do that by generating a smaller number of states using a targeted selection of initial conditions and evolutionary parameters, and using that simulated state space as a set of training examples to teach an interpolant how to map initial conditions and evolutionary parameters to the final state, thus avoiding the computational time of calculating the intermediate steps. Once we have a working tool we can then generate a very large simulated state space and search that state space for states that match observed data. Here, as proof of the concept, we describe the construction of the tool for a reduced space of parameters: four parameters governing the cosmic

metallicity-specific star formation rate (MSSFR) in the model of [Neijssel et al. \(2019\)](#).

The remainder of this paper is organised as follows. Section 2 presents a description of the tool we construct for this proof-of-concept study, and the method used to train the tool. We present and discuss our results in Section 3. We provide some concluding remarks in Section 4.

2. METHOD

2.1. Overview

The COMPAS population synthesis software suite ([Stevenson et al. 2017](#); [Vigna-Gómez et al. 2018](#); [Team COMPAS: Riley, J. et al. 2022a](#)) allows users to generate synthetic populations of binary systems. The COMPAS suite includes post-processing tools that allow users to, amongst other things, calculate merger rates, detection rates of mergers, etc., for double compact objects (DCOs) that will merge within the age of the universe ([Neijssel et al. 2019](#); [Barrett et al. 2018](#); [Team COMPAS: Riley, J. et al. 2022b](#)).

For this proof-of-concept study we develop a surrogate model for COMPAS that provides fast predictions of detection rates of merging DCOs as a function of chirp mass and merger redshift. Surrogate models have been used to emulate various aspects of the simulated evolution of binary systems and to infer constraints on simulation parameters, including the evolution of individual binary systems (e.g. [Lin et al. \(2021\)](#), using a Gaussian Process emulator ([Rasmussen & Williams 2005](#)) and classifier) and populations of binaries (e.g. [Barrett et al. \(2016\)](#), [Taylor & Gerosa \(2018\)](#), [Wong & Gerosa \(2019\)](#), [Cheung et al. \(2022\)](#) using Gaussian Process emulators or normalising flows, and [Wong et al. \(2021\)](#) using deep-learning enhanced hierarchical Bayesian analysis).

We chose detection rates because these are calculated by the COMPAS post-processing tools after the evolution of the population of binary systems by convolving binary evolution outcomes with the MSSFR. This makes generating example (training) data for our surrogate model by varying MSSFR parameters much less onerous than it would be if we needed to evolve new populations of binaries whenever our astrophysical parameters are varied.

We construct an interpolant that, given a set of input values, predicts a detection rate matrix for the merger of DCOs. Input to the interpolant is a set of values for four MSSFR parameters chosen for this study (see Section 2.2). The interpolant outputs detection rates binned into a matrix of 113 chirp mass bins by 15 redshift bins (see Section 2.3). The detection rates gener-

ated by the interpolant are combined rate predictions for all DCOs: binary black holes, binary neutron stars, and mixed neutron star – black hole binaries. The interpolant we construct is comprised of a matrix of simple feed-forward artificial neural networks (ANNs), with each ANN trained to predict the detection rate for a single cell of the resultant detection rate matrix.

We first simulate the evolution of 512 million binary systems using COMPAS, resulting in a population of approximately 1.3 million DCOs that merge within the age of the Universe, ~ 13.8 Gyr.

The COMPAS modelling tool has many configurable options that allow the user to change initial conditions and evolutionary parameters. The simulation for this study was performed using the COMPAS fiducial values for all configurable options (see Appendix B, Table 4; described in detail in [Team COMPAS: Riley, J. et al. \(2022a\)](#)), and for our purposes we consider the population of binary systems simulated by COMPAS to be indicative of the population of binary systems in the Universe.

The COMPAS suite includes tools to take a population of binary systems produced by COMPAS, integrate over the known initial condition distribution and the cosmological star formation history, and produce a population of merging DCOs with individual masses and merger redshifts, to which observational selection effects can be applied, converting the population into a prediction for the observable merger rate over the population of DCOs.

We construct a grid of six different values for each of the four MSSFR parameters under study (see Section 2.2), and for each entry in the grid we use the COMPAS tools described above to calculate a detection rate matrix for the population of merging DCOs, resulting in 1,296 detection rate matrices, which we use to create the training data for the ANNs that comprise our interpolant. The calculated detection rate matrices contain, for each bin in the chirp mass – redshift space, the expected detection rate (detections per year) for an interferometer network with an indicative sensitivity of the LIGO O3 run ([Abbott et al. 2020](#)).

For each cell in the detection rate matrix we construct and train an ANN that predicts the detection rate for that cell, given values for the four MSSFR parameters – this matrix of ANNs comprises our interpolant¹. We chose a matrix of simple ANNs rather than a single, large ANN for several reasons:

¹ Some cells in the detection matrix may not require the use of an ANN – see Section 2.4.

- (i) Training a large ANN to learn > 1500 relationships is difficult and time-consuming, whereas training a small ANN to learn a single relationship is easier and faster.
- (ii) A large ANN is likely to be less accurate because the network will try to get most relationships close, whereas the single, small ANN just needs to get one relationship right.
- (iii) Over-fitting in a large ANN learning multiple relationships is more likely than a single, small ANN learning a single relationship. It is likely that some of the > 1500 different relationships to be learned will be easier to learn than others. Training a large ANN to learn multiple relationships could result in the relationships that are easier to learn being over-fitted because the ANN needs to be trained for longer to learn the more difficult relationships.
- (iv) ANNs for different cells can have different architectures - some relationships might require more, or fewer, nodes and/or layers.
- (v) While other global surrogate models could provide computational cost savings by utilising the regularity (smoothness) of the output over the input parameter space, assumptions of regularity are risky when some astrophysical parameters could lead to bifurcations in the outputs.
- (vi) We can replace the ANN for some cells with a different model without affecting all cells - some relationships may not need to be modelled by an ANN: a simple function fit (possibly even a constant value) might be sufficient.
- (vii) We can retrain the ANNs for individual cells as necessary, without the need to retrain all cells.
- (viii) We can improve performance by running individual ANNs in parallel on different CPUs.

2.2. MSSFR parameters

The parameters we chose to vary for this study are four free parameters for the calculation of the MSSFR in the phenomenological model of Neijssel et al. (2019). In this model, the MSSFR is split into two parts, the star formation rate (SFR) and the metallicity distribution:

$$\frac{d^3 M_{\text{SFR}}}{dt_s dV_c dZ}(z) = \frac{d^2 M_{\text{SFR}}}{dt_s dV_c}(z) \times \frac{dP}{dZ}(z), \quad (1)$$

where the SFR is given by:

$$\frac{d^2 M_{\text{SFR}}}{dt_s dV_c} = a \frac{(1+z)^b}{1 + \left(\frac{1+z}{c}\right)^d} M_{\odot} \text{ yr}^{-1} \text{ Mpc}^{-3}, \quad (2)$$

and we use a log-normal distribution in metallicity at each redshift (cf. the skewed log-normal model proposed by van Son et al. (2022)):

$$\frac{dP}{dZ}(z) = \frac{1}{Z\sigma\sqrt{2\pi}} e^{-\frac{(\ln(Z) - \ln(\langle Z \rangle) + \sigma^2/2))^2}{2\sigma^2}}, \quad (3)$$

with a redshift-independent standard deviation σ in $\ln(Z)$ space around a redshift-dependent mean μ of $\ln(Z)$ given by:

$$\langle Z \rangle = e^{(\mu + \frac{\sigma^2}{2})}, \quad (4)$$

with mean metallicity parametrised as in Langer & Norman (2006):

$$\langle Z(z) \rangle = Z_0 10^{\alpha z} \quad (5)$$

We vary SFR parameters a and d , and the metallicity distribution parameters α and σ , in this study, while fixing $b = 2.77$, $c = 2.9$, and $Z_0 = 0.035$. Neijssel et al. (2019) demonstrated that a range of star-formation history models can be mimicked by varying only this subset of MSSFR parameters. We therefore choose to vary the same parameters for consistency with their analysis. Moreover, reducing the number of parameters allows us to limit computational complexity for this proof-of-principle study.

2.3. Binned rates

As noted in Section 2.1, the detection rates generated by our interpolant are binned into a matrix of 113 chirp mass bins by 15 redshift bins.

The redshift bins are constructed as fixed-width bins, with the lower edge of the first bin at redshift 0.0 and each bin having a width of 0.1 - the upper edge of the final (15th) bin is at redshift 1.5.

The chirp mass bins are constructed as variable-width bins, with the lower edge of the first bin at $L_1 = 0.0 M_{\odot}$ and the upper edge at $L_2 = 0.5 M_{\odot}$. The remaining chirp mass bin edges fall at

$$L_{i+1} = \frac{41}{39} L_i, \quad 2 \leq i \leq 112. \quad (6)$$

Thus, the width of the 2nd to 112th bins is equal to 5% of the median chirp mass for the bin. The final (113th) bin extends to infinity.

2.4. The Interpolant

As noted in Section 2.1, the interpolant we construct is comprised of a matrix of ANNs. For this study, all ANNs in the matrix have identical architecture, though




Figure 1. The Interpolant

that is not a requirement of the method: each ANN is independent of the others, and the architecture of individual ANNs could be varied. Nominally, there will be a corresponding ANN for each cell in the detection rate matrix, but there may be some cells for which an ANN is not required. For example, if a given cell in the detection rate matrix is known to be a constant value (e.g. 0), training an ANN is both not useful and not necessary - we can instead replace the ANN for that cell with (e.g.) a simple equation and skip training an ANN. A schematic diagram of the interpolant is shown in Figure 1. In Figure 1, ANNs are trained to predict values only for the unshaded cells in the detection rate matrix: values for the shaded cells will be predicted using some other model, as discussed above.

2.4.1. The Artificial Neural Networks

Each ANN is a fully-connected feed-forward network of artificial neurons (nodes), with an input layer, three hidden layers, and an output layer. The input layer consists of four input nodes (one for each MSSFR parameter), the output layer consisting of a single node (to indicate the predicted detection rate for the corresponding matrix element), and each of the three hidden layers contains 64 nodes. The hidden and output layers have an associated bias node that provides each node in respective layer with a bias weight. A schematic diagram of a representative ANN is shown in Figure 2.

While there are some generally accepted rules of thumb, the selection of ANN architecture is very often arbitrary. In this case the architecture was chosen to produce relatively robust results quickly – no attempt was made to optimise the architecture for the best performance, either with respect to accuracy or speed. The input layer and output layer are fixed (the four MSSFR parameters are the inputs, and since the ANN is a regressor, not a classifier, we have a single, continuous, output node), and the choice of three hidden layers, each with 64 nodes, was a brute-force rather than finessed




Figure 2. ANN architecture

approach. Per the Universal Approximation Theorem², a shallow neural network (with a single, wide, hidden layer) can approximate any function – after a little experimentation with different architectures the architecture described above was found to produce good enough results reasonably quickly. While there is no universally accepted rule as to how many layers constitute a deep ANN (how many grains make a heap?³), the architecture implemented here probably qualifies as a deep network (because the networks have more than two hidden layers) - but it is essentially a simple feed-forward ANN (albeit, given the number of nodes on each layer, a large simple feed-forward ANN).

2.4.2. Training

Training an ANN requires many examples of the relationship(s) to be learned (e.g. example outputs for given inputs) - anywhere from hundreds to millions depending upon the complexity of the relationship(s) to be learned. By designing the interpolant as a matrix of independent, small ANNs, we have minimised the number, and complexity, of the relationships to be learned by a single ANN, and as a result the size of the training data set required is significantly reduced.

The data used to train the ANNs was constructed by using the COMPAS post-processing tools to calculate detection rate matrices for 1,296 combinations of the MSSFR parameters under study - 6 different values for each of the 4 parameters, collectively labeled λ , as shown in Table 1. The trained interpolant is expected to be able to interpolate between the bounds of each of the constituent parameters of λ .

We chose the ranges of values based on the preferred values from Neijssel et al. (2019), which are the default

² https://en.wikipedia.org/wiki/Universal_approximation_theorem

³ https://en.wikipedia.org/wiki/Sorites_paradox

Table 1. λ training values

λ_x	Values
α	[-0.500, -0.400, -0.300, -0.200, -0.100, -0.001]
σ	[0.100, 0.200, 0.300, 0.400, 0.500, 0.600]
a_{SF}	[0.005, 0.007, 0.009, 0.011, 0.013, 0.015]
d_{SF}	[4.200, 4.400, 4.600, 4.800, 5.000, 5.200]

COMPAS values for each of the parameters - the ranges are a spread around the preferred values.

The question of how many training examples are needed to train an ANN effectively has no simple answer. We need a sample of data that is representative of the problem we are trying to solve (in our case, constructing an interpolant for DCO merger detection rates) and, in general, the examples in the sample should be independent and identically distributed. We are training ANNs which we hope will learn to correctly map input data to output data, and moreover will learn to interpolate and map input data that they have not seen during training - that is, we hope our ANNs will learn the relationships between the input data and output data. We need sufficient training data to reasonably capture the relationships that (we expect) exist both between input features, and between input features and output features.

The dataset described above contains a single example of the 1,296 mappings (relationships) that we will present to the ANNs. We know that a single example is almost certainly not sufficient to correctly train the ANNs - with too few training examples we risk overfitting the ANNs, and possibly compromising their ability to generalise beyond the training dataset (but also see Appendix A). In fact, we know that the training data are not perfectly accurate because the COMPAS binary population represents a Monte Carlo sampler over initial conditions. Therefore, training data suffer from statistical sampling uncertainty. To account for this, we create two new training datasets which we will use to train and test the networks. The results will give us an indication of the number of training examples required. For each new training dataset, we use bootstrap sampling of COMPAS outputs to create a number of new detection rate matrices.

The first sampled dataset contains 10 different bootstrapped examples for each of the 1,296 mappings, resulting in a training dataset with 12,960 matrices, and the second contains 100 different bootstrapped examples for each of the 1,296 mappings, resulting in a training dataset of 129,600 matrices.

As discussed previously, we don't need to train ANNs for cells in the detection rate matrix that we know to

contain constant rates. For each training dataset we inspect all training examples and determine which cells in the matrix are constant across all examples - ANNs are not required for those cells. We found that for the 1,695 cells in the detection rate matrix we only needed to train 292 ANNs for the smaller training dataset, and 293 ANNs for the larger training dataset - for all other cells the detection rate across all training examples was below our threshold for zero (0.001 detections per year).

Training of the interpolant was conducted using the Keras API (Chollet et al. 2015) for Python (Van Rossum & Drake 2009). Keras is a high-level neural networks API that uses the Tensorflow machine learning platform (Abadi et al. 2015).

We used a fairly naïve training methodology. Just as no attempt was made to optimise the architecture for the best performance, no real attempt was made to optimise the training for the best accuracy of the networks - "good enough" accuracy was all we were looking for in this proof-of-concept study. We divided our collection of models with different choices of parameters λ and corresponding bootstrapped examples into a training set (80% of all choices of λ), and a validation set (the remaining 20% choices of λ). Networks were trained using the set of training examples, and tested against the validation set - examples that the network undergoing training had not seen during training. k -fold cross-validation was not performed - some non-exhaustive tests were performed using k -fold cross-validation with no significant improvement in accuracy, so in the interest of reduced training time we chose not to implement k -fold cross-validation.

Each network was initialised with random connection weights and biases, and trained for a maximum of 4,500 epochs (an epoch is the presentation of the entire training set to the network), after which the network was tested on the validation set. Early stopping was enabled - if no improvement in accuracy (on the training data) was evident after a specified number of epochs (since the last improvement), training was halted.

The network attributes (weights and biases) were adjusted during training using the Keras Adam optimiser (a stochastic gradient descent optimisation algorithm), with a custom metric for measuring network performance (⁴).

The metric used for measuring the performance of the network was half the Poisson uncertainty of the detection count associated with the bin being predicted assuming 1 year of observations: where the expected (tar-

⁴ <https://keras.io/api/optimizers/adam/>

get) detection count for a given bin is t , a prediction, p , is considered correct if $(t - \frac{\sqrt{t}}{2}) \leq p \leq (t + \frac{\sqrt{t}}{2})$. We used sufficiently large COMPAS binary populations for this study to ensure that the bootstrapping uncertainty is typically small relative to \sqrt{t} : fewer than 0.01% of training and validation samples fall outside a range of width \sqrt{t} . The majority of these outliers were found in specific bins in chirp mass – redshift space with low DCO counts (see Section 3.1.2) and correspondingly high bootstrapping uncertainty. The accuracy requirement could have been loosened for such bins, but given the low outlier count, we were comfortable with using this definition of performance.

If the network did not achieve 100% accuracy on the validation set, it was trained again with a different set of (random) initial weights and biases. If the network did not achieve 100% accuracy on the validation set after 10 tries at training, the network that achieved the highest accuracy during training was accepted.

2.4.3. Performance

There are three aspects of surrogate model performance that interest us: speed, accuracy, and the ability to generalise (specifically, in our case, interpolate). If we are to use the interpolant to create large synthetic state spaces that we can explore, we need it to be very fast, at least by comparison with existing methods (e.g. the COMPAS tools described in Section 2.1), to be accurate, and to generalise from the training examples.

The accuracy of the interpolant, with respect to generating a correct detection rate matrix for a given set of MSSFR parameters, is measured during training (see Section 2.4.2). Since the validation data are technically used as part of training – in determining whether training can be stopped or can continue to another iteration – we conducted several ad-hoc spot checks of the interpolant against results generated using COMPAS, and in all cases the interpolant matched (within acceptable error bounds) the COMPAS results, thus confirming the ability to generalise/interpolate.

Measuring the speed of the interpolant is straightforward. For this study, since the interpolant consists only of ANNs, the speed of the interpolant is dependent mainly upon the execution (forward propagation) speed of its constituent ANNs, and how many ANNs comprise the interpolant - with a small amount of input/output (IO) overhead to read the input data (i.e., the values of λ), present it to the network, and write the output data.

All speed-related performance tests were conducted on a Hewlett Packard ProLiant ML350p Gen8 with two Intel E5-2650v2 processors running at 2.6GHz, providing

32 threads (with hyper-threading) and 96GB of RAM - performance figures quoted are for that system.

The total elapsed time to create both training datasets was ~ 142 hours.

Since in this study the ANNs that comprise our interpolant have identical architectures, if we ignore the IO overhead, the speed of the interpolant is directly proportional to the execution speed of a single ANN. Training of the ANNs that comprise the interpolant is independent, so can be done in parallel. We measured the average execution speed of a single ANN over 1,000,000 executions to be 2.19×10^{-5} seconds per execution. If we consider the interpolant to be the serial execution of its (for this study) 293 constituent ANNs (292 for the smaller dataset), the execution speed of the interpolant (ignoring IO overhead) is $2.19 \times 10^{-5} \times 293 = 6.42 \times 10^{-3}$ seconds. However, since the execution of the constituent ANNs is independent, the architecture of the interpolant lends itself to parallel execution (it is, in fact, embarrassingly parallel), so the speed of execution of the interpolant could be reduced significantly by spreading the execution of the constituent ANNs across multiple CPUs/cores.

The total elapsed time to train all ANNs was ~ 1.5 hours for the smaller training dataset, and ~ 12 hours for the larger training dataset.

As noted earlier, our interpolant is only useful if it is significantly faster than existing methods. We measured the average execution speed of the COMPAS toolset to create a merger detection rate matrix from a population of DCOs over 50 executions to be 111.97 seconds. The COMPAS post-processing toolset (to create the detection rate matrix) is not parallelisable.

Our interpolant, if run in serial fashion, is $\sim 17,500$ times faster than the COMPAS toolset, and, naively, could be up to ~ 5 million times faster if run in a fully parallel fashion (though there would be some overhead in running in parallel). We should note that the COMPAS post-processing tools are written in Python (but it is unlikely that any optimisation would result in improvements of more than tens of percent).

2.5. Searching the state space

With the interpolant constructed, we shift our focus to showing that the interpolant enables us to infer constraints on our chosen MSSFR parameters. To do this we search the state space over the range of values of the parameters we used to train our interpolant – the interpolant is not guaranteed to be valid outside those ranges – and seek out models that match the observed data.

First, to validate the method and test the interpolant, we create a mock data set that, for this test, we consider

to be representative of the Universe. We use COMPAS as our forward model to create a population of DCOs representative of the Universe, and, as noted earlier, we use the COMPAS fiducial values for all configurable options.

We use the COMPAS post-processing tools to create mock “true” data, \mathcal{D} , as follows. We specify a detector noise spectrum (equivalent to an approximate LIGO O3 sensitivity) and choose “true” values of the MSSFR parameters λ . We use the COMPAS tools to generate a DCO merger detection rate matrix that represents the universe defined by the COMPAS fiducial model. We multiply the detection rate matrix by a fixed run duration to produce a matrix of expected detections, and then independently sample mock observations from each chirp mass – redshift in this matrix following Poisson statistics.

We then use our surrogate model to generate a detection rate matrix at various points in the state space (characterised by varying λ), multiply this by the same run duration, and calculate the likelihood of observing \mathcal{D} given λ at those points as described below. This produces a “likelihood landscape” that we can search. Finding the maximum likelihood on the landscape tells us the value of λ (the interpolant considers is) most likely to produce our “true” data. The likelihood surface (or, in Bayesian language, the posterior on the parameters λ under the assumption of flat priors over the range considered) contains information about uncertainty in the MSSFR parameter inference and any correlations between these parameters.

Results of the validation of the method are presented in Section 3.2.1.

2.5.1. Likelihood calculation

We define the likelihood \mathcal{L} that we will see the data \mathcal{D} for a particular set of parameters λ as

$$\log \mathcal{L}(\mathcal{D}|\lambda) = \log \mathcal{L}(N_{\text{obs}}|\lambda) + \sum_{i=1}^{N_{\text{obs}}} \log p(D_i|\lambda), \quad (7)$$

where N_{obs} is the number of observations in the data $\mathcal{D} \equiv \{D_1, D_2, \dots, D_{N_{\text{obs}}}\}$.

The first term of Equation 7 is a Poisson likelihood on there being N_{obs} detections,

$$\log \mathcal{L}(N_{\text{obs}}|\lambda) = N_{\text{obs}} \log(\mu(\lambda)) - \mu(\lambda), \quad (8)$$

where $\mu(\lambda)$ is the expected number of observations and we omit terms that depend on the data only and therefore disappear on normalisation, such as $\log(N_{\text{obs}}!)$ and permutation coefficients.

The second term of Equation 7 is comprised of a product of the probabilities of individual detections,

$$p(D_i|\lambda) = p(z = z_i, \mathcal{M} = \mathcal{M}_i|\lambda), \quad (9)$$

where $p(z = z_i, \mathcal{M} = \mathcal{M}_i|\lambda)$ is the entry in the probability distribution matrix $p(z, \mathcal{M}|\lambda)$ in the z and \mathcal{M} bin of the observed i^{th} event. The probability distribution matrix $p(z, \mathcal{M}|\lambda)$ is constructed by dividing the entries in the matrix of expected detections for a given λ , generated by multiplying the surrogate model detection rate matrix by the run duration, by the sum of the entries in the matrix, μ .

2.6. Using gravitational wave observations

The catalogue of LIGO, Virgo, and KAGRA (LVK) gravitational wave observations (LIGO Scientific and Virgo Collaborations (2022), LIGO Scientific, Virgo and KAGRA Collaborations (2021)) provides us with a number of merger detections from a population of DCOs. For each LVK event we know, among other things, the chirp mass of the merging binary and the merger redshift, up to a measurement uncertainty. We only include confident detections with a minimum astrophysical probability $p_{\text{astro}} \geq 0.95$; although less confident events could be included (e.g. Farr et al. 2015), they will typically contribute little information due to greater measurement uncertainties, so we omit them for this proof-of-concept study.

Once we validate our method, we can use our surrogate model and the LVK data to infer real physical constraints on the parameters under study, with the important caveat that we have fixed the astrophysical parameters, and any errors in those could bias the inferred MSSFR model values.

2.6.1. LVK likelihood calculation

In practice, for LVK events, we may not know the values z_i, \mathcal{M}_i and hence the probability $p(z = z_i, \mathcal{M} = \mathcal{M}_i|\lambda)$ for the observed i^{th} event in Equation 9 perfectly, but may only have K samples from a posterior $p(z_i, M_{c,i}|D_i)$. For LVK data, we substitute the following for Equation 9:

$$p(D_i|\lambda) = \frac{1}{K} \sum_{k=1}^K \frac{p(z = z_i^k, M_c = M_{c,i}^k|\lambda)}{\pi(z_i^k, M_{c,i}^k)}, \quad (10)$$

where the subscript k refers to the k^{th} posterior sample among K available samples for event i and $\pi(z, M_c)$ is the prior used by LVK data analysts in the original inference on the observed events (see, e.g. Mandel et al. 2019).

3. RESULTS AND DISCUSSION

3.1. Interpolant training results

The detection rate matrix generated by our interpolant consists of 1695 cells: 113 chirp mass bins x 15 redshift bins. Of the 1695 cells in the matrix, only 17% (292 for the smaller training dataset, and 293 for the larger training dataset) require an ANN to predict the rate for that cell: all other cells have zero rates for all training examples. The time to train the interpolant, and the accuracy achieved during training, are shown below.

3.1.1. Training times

Details of the training times for the ANNs trained are presented in Table 2. As noted in Section 2.4.3, the total elapsed time to train all ANNs was ~ 1.5 hours for the smaller training dataset, and ~ 12 hours for the larger training dataset.

Table 2. ANN training statistics

Training Statistic	#examples	
	10	100
Lowest train time (min)	0.26	1.37
Highest train time (min)	42.98	301.37
Average train time (min)	9.85	74.1
Average tries to train	5.76	7.18
Fewest tries to train	1	1
Most tries to train	10	10
Number of networks trained	292	293

3.1.2. Interpolant accuracy

As discussed in Section 2.4.2, each of the constituent ANNs of the interpolant was trained until an accuracy threshold was reached - with accuracy being measured as the percentage of the validation data set the ANN was able to correctly predict (the meaning of “correct” is explained in Section 2.4.2).

Details of the interpolant accuracy achieved during training are presented in Table 3.

The accuracy figures in Table 3 show that for both training datasets, the average accuracy achieved by the ANNs trained was better than 99.4%, with more than 90% of the ANNs achieving an accuracy of 99% or better.

Figure 3 shows in detail the accuracy achieved for each of the ANNs - the grid shown corresponds to the detection matrix, with each cell representing the ANN corresponding to a $\{z, \mathcal{M}\}$ bin.

Table 3. ANN training accuracy

	#examples	
	10	100
% correct		
Highest	100.0	100.0
Lowest	54.16	83.0
Average	99.44	99.58
Stddev	3.23	1.84
Counts		
% correct = 100	139	97
$99 \leq \% \text{ correct} < 100$	124	172
$95 \leq \% \text{ correct} < 99$	26	20
$90 \leq \% \text{ correct} < 95$	0	0
$85 \leq \% \text{ correct} < 90$	0	1
$80 \leq \% \text{ correct} < 85$	1	3
% correct < 80	2	0
Number of networks evaluated	292	293

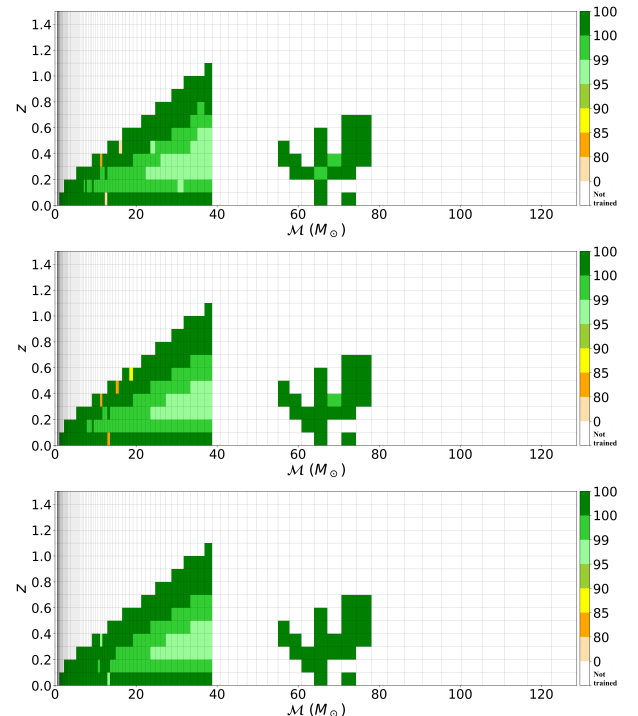


Figure 3. Interpolant training accuracy: 10 training examples (upper panel); 100 training examples (middle panel); 100 training examples with correctness threshold of $|t - p| \leq \max(\sqrt{t}/2, 0.5)$ (bottom panel). White cells represent ANNs that did not require training (zero rates); the meaning of other colours is shown on the colour bar.

From the top two panels of Figure 3 we can see that most of the ANNs (for both training datasets) were

trained to a high degree of accuracy, with only a few ANNs not achieving 95% accuracy or better. We can inspect the data and the training records for those few ANNs, and because the architecture of the interpolant allows us to easily replace individual constituent ANNs (or whatever model we have used for a particular cell), we can retrain them to achieve a better accuracy, change the ANN architecture, etc. For example, we find that increasing the number of hidden layers from 3 to 5 allows us to achieve 100% accuracy across all cells. However, there is a risk of attempting to memorise rather than learn the model with over-trained ANNs (see Appendix A).

The shape and location of the green areas in Figure 3 is interesting. Recall that the white cells in Figure 3 are $\{z, \mathcal{M}\}$ bins for which the detection rate was below the threshold for zero (0.001 detections per year). The corollary is that the green areas are $\{z, \mathcal{M}\}$ bins that contain DCOs that are detectable at LIGO O3 sensitivity. The triangle on the left is delineated by lower horizon distances for smaller masses along the diagonal line at the top, and by the lower edge of the pair-instability mass gap on the right (Heger & Woosley (2002), Marchant & Moriya (2020)).

Many of the bins along the top edge of the triangle (where the detection probability decreases toward zero at the horizon distance) and the bottom edge (only a small comoving volume is contained between redshifts of $z = 0$ and $z = 0.1$) have low detection rates. These bins also contain the apparently poorly performing ANNs in the top two panels. However, this is largely an artefact of our choice of a relative tolerance threshold for correctness, which becomes increasingly stringent as the detection rate decreases. We can instead consider a combination of relative and absolute tolerances, which better accounts for the fact that a detection in a bin is unlikely if the predicted detection rate is low, making excessive precision in the surrogate model unnecessary. If we replace the threshold for correctness of surrogate model prediction p given expected target count detection count t with $|t - p| \leq \max(\sqrt{t}/2, 0.5)$ (cf. the default threshold of $|t - p| \leq \sqrt{t}/2$ from Section 2.4.2), we find that all ANNs achieve an accuracy above 95% (see bottom panel of Figure 3).

We see a green cactus-shaped area to the right of the triangle (that is, to the right of the lower edge of the pair-instability mass gap). To understand this we need to understand the distribution of DCOs over the chirp mass range - this is shown in Figure 4. Figure 4 shows that almost all of the ≈ 1.3 million merging DCOs in the training data are located in the bins where $\mathcal{M} \lesssim 40M_{\odot}$ (the green triangle in Figure 3), with a few bins in the \approx




Figure 4. DCO counts for chirp mass \mathcal{M} bins.

$50 - 80M_{\odot}$ range (the green cactus-shaped area) with 10 or fewer DCOs per bin, and then two bins at $\approx 120M_{\odot}$ with a total of ≈ 190 DCOs (not seen in Figure 3 - the detection rates for these bins were below the zero threshold).

Further inspection of the data reveals that of the ≈ 1.3 million merging DCOs, only 50 are located in the $\approx 50 - 80M_{\odot}$ chirp mass range, and a further 193 are in the two chirp mass bins at $\approx 120M_{\odot}$, and all of these binaries were BBHs comprised of two massive stars that began their evolution on the main sequence as chemically homogeneous stars (Riley et al. 2021). Of the 50 binaries in the $\approx 50 - 80M_{\odot}$ chirp mass range, 23 had both stars evolve as chemically homogeneous stars for the entirety of their lifetime on the main sequence, while for the remaining 27 the secondary star did not evolve as a chemically homogeneous star for its main sequence lifetime. Of the 193 binaries in the $\approx 120M_{\odot}$ chirp mass range, all but one had both stars evolve as chemically homogeneous stars for the entirety of their lifetime on the main sequence - the secondary star of the remaining binary did not evolve as a chemically homogeneous star for its main sequence lifetime.

With only 243 of the ≈ 1.3 million DCOs produced from a population of 512 million binaries above the lower edge of the pair-instability mass gap, it is clear that these DCOs are very rare (Marchant et al. (2016), Riley et al. (2021)). The cactus shape in Figure 3 is thus determined by the delay times of chemically homogeneously evolving DCOs and convolution with the MSSFR, with the detailed features likely an artefact of limited sampling.

3.2. Inference with surrogate model

3.2.1. Method validation: inference on perfect measurements

In Section 2.5 we outlined how we validate our method. Here we present the results of that validation.

Using COMPAS post-processing tools and the synthesised population of 512 million binaries, we created a detection rate matrix, binned as described in Section 2.3, for a known value of λ : $\lambda(\alpha, \sigma, a_{SF}, d_{SF}) = (-0.325, 0.213, 0.012, 4.253)$. Using that detection rate

matrix and assuming an observing time of either 0.1 year or 1 year, we created mock data sets \mathcal{D} with perfect parameter measurement accuracy for each of the “detected” sources (58 in the 0.1 year data set, 578 in the 1 year data set).

We carried out a Markov Chain Monte Carlo (MCMC) search (e.g. [Andrieu et al. \(2003\)](#)), using the *emcee* Python package ([Foreman-Mackey et al. 2013](#)), over the $\lambda = \{\alpha, \sigma, a_{SF}, d_{SF}\}$ parameter space, assuming flat priors on λ . This search calculates the likelihood $\mathcal{L}(\mathcal{D}|\lambda_i)$ at each λ_i visited by the MCMC algorithm using the surrogate model (see Section 2.5.1 for the likelihood calculation). To confirm that our inference step was performing adequately, we also performed a fine-grained grid search consisting of 51 equidistant values for each of α , σ , and a_{SF} , and 101 equidistant values for d_{SF} - for a total of $51 \times 51 \times 51 \times 101 \approx 13.4$ million points evaluated, followed by a naïve hill-climbing search (e.g. [Russell & Norvig \(1995\)](#)), using the results of the grid search as a starting point. We confirmed that the searches found the same maximum likelihoods (within expected sampling variations).

The MCMC search posteriors are shown in Figure 5. The true value is found within 68% credible intervals for most MSSFR parameters under study, and within 95% credible intervals for all parameters. As expected, the larger data set produces more precise inference on MSSFR parameters, with the ratio of posterior width for α approaching a factor of $\sqrt{10}$ narrower on the data set with 10 times more data (right panel), though the improvement is much smaller for poorly measured parameters such as d_{SF} whose posteriors are prior-dominated.

We considered surrogate models trained on data sets with 10 bootstrapped samples per λ value and 100 bootstrapped samples and found that they are not significantly different, which is consistent with the accuracy of the surrogate model shown in Figure 3. We show results for surrogate models trained with 100 bootstrapped samples per λ from here on.

3.2.2. Method validation: inference on uncertain measurements

We used the same detection rate matrix as in the previous subsection, corresponding to $\lambda(\alpha, \sigma, a_{SF}, d_{SF}) = (-0.325, 0.213, 0.012, 4.253)$, to validate inference on mock observations with measurement uncertainties. Using that detection rate matrix, assuming an observing time of one year, we created a dataset of mock LVK data containing 578 events. We then randomly sampled 58 events from that dataset, creating a new dataset with an observing time of 0.1 year.

We then replaced each of the chosen events with samples from an associated mock posterior. To create these

samples, we used a mock model of the LVK prior. We built the source-frame chirp mass prior by assuming that the component masses $m_1 > m_2$ are uniformly drawn from the range $[1, 1000]M_\odot$, with additional cuts $m_2 \in [0.05m_1, m_1]$ and $\mathcal{M} \equiv m_1^{0.6}m_2^{0.6}(m_1 + m_2)^{-0.2} \in [1, 200]M_\odot$. For the mock redshift prior, we used $\pi(z) \propto z^2$ on $z \in [0.01, 1.5]$. We then weighed mock chirp mass and redshift samples taken from

$$\begin{aligned} \mathcal{M} &\sim \mathcal{M}^T (1 + 0.03 \frac{12}{\rho} (r_0 + r)); \\ z &\sim z^T (1 + 0.3 \frac{12}{\rho} (r_0 + r)) \end{aligned} \quad (11)$$

by these priors. In Eq. 11, which follows [Powell et al. \(2019\)](#), the superscript T denotes the true value, ρ is the signal-to-noise ratio sampled from $p(\rho) \propto \rho^{-4}$ with a minimum of $\rho \geq 12$, r_0 is a normal random variable which stochastically shifts the peak of the posterior away from the truth, and r is a vector of normal random variables which provides the spread of the posterior. Thus, the small and large data sets are composed of 58 and 578 events, each represented by a set of mock posterior samples that contain mock measurement uncertainties on the observed parameters \mathcal{M} and z .

Using each of those datasets as “true” data, we used our surrogate model to infer posteriors on λ assuming flat priors. The MCMC search statistics are shown in Figure 6.

Figure 6 shows that, as expected, inference provides consistent credible intervals, with most “true” values falling within the 1- σ credible interval, and the remainder within the 2- σ credible interval. As before, the larger data set produces more precise inference on MSSFR parameters. Comparing Figure 6 with Figure 5, we see that mock measurement uncertainties have limited impact on MSSFR inference for the smaller data set, where inference is predominantly limited by the total number of events. For the larger data set, incorporating measurement uncertainty does lead to a moderate deterioration in the accuracy of inference.

3.2.3. Inference from real LVK observations

Since the goal of this work is to (eventually) develop a method to allow us to infer constraints for (some of) the astrophysical states and processes, as an illustration of the ability of this methodology to assist in inferring constraints on cosmological and astrophysical parameters from real data, we performed an MCMC search (as described in Section 3.2.1) on actual LVK observations from observing runs 3a and 3b (run duration 275.3 days), where $p_{astro} \geq 0.95$ (52 events).

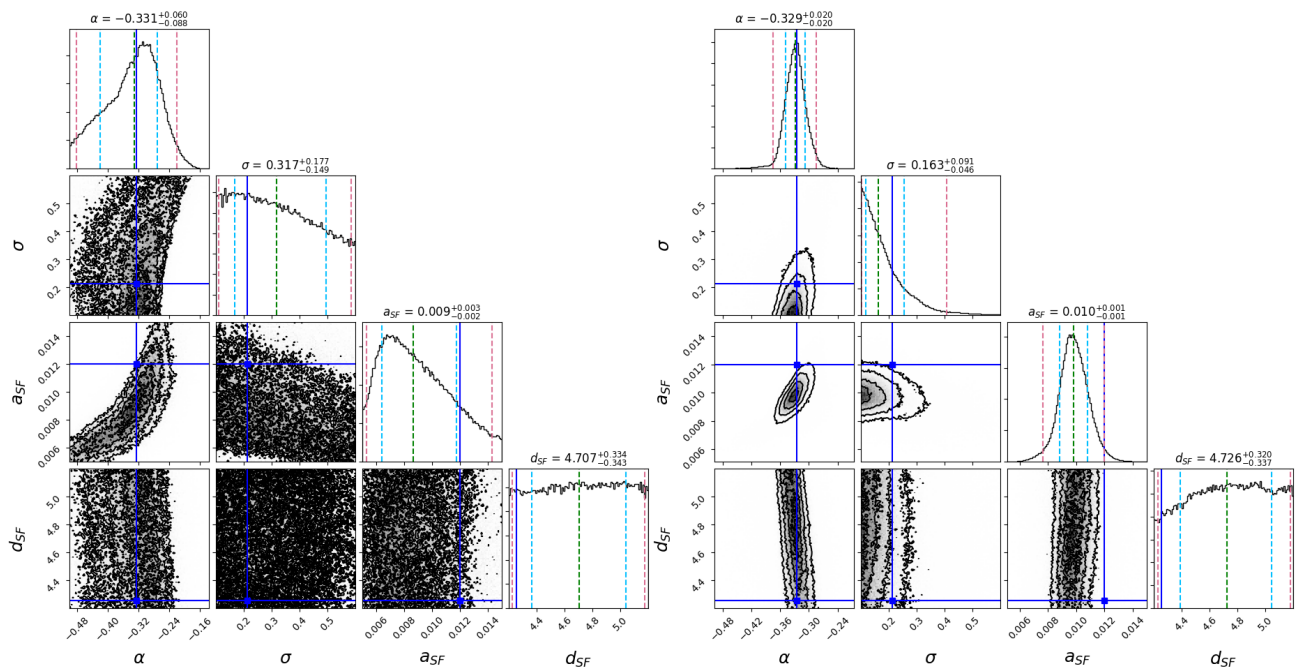


Figure 5. Corner plots showing MSSFR parameter inference on mock observations without measurement uncertainties. The marginalised 1-d histograms for each MSSFR parameter appear in the panels along the diagonal, and the marginalised 2-d projections of the posterior probability distributions in the off-diagonal panels. The left panel is for the smaller dataset (58 events); the right panel for the larger dataset (578 events). The blue solid line indicates the best-fit value, the green dashed line shows the median, and the pale-blue and pale-red dashed lines show the 68% and 95% credible intervals, respectively. Labels on the diagonals show the mean and 68% credible interval.

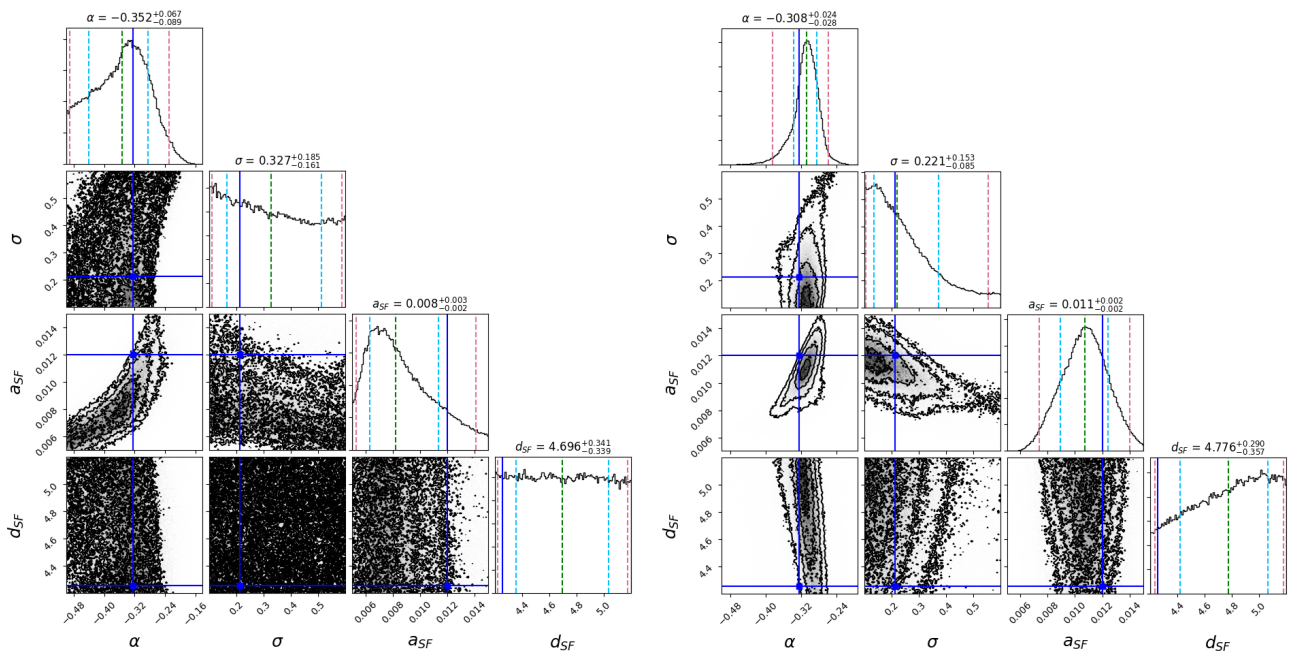


Figure 6. Corner plots for MSSFR parameter inference on mock observations with measurement uncertainty. The left panel is for the smaller dataset (58 events); the right panel for the larger dataset (578 events). The blue solid line indicates the “true” value, the green dashed line shows the median, and the pale-blue and pale-red dashed lines show the 1σ and 2σ credible intervals respectively.

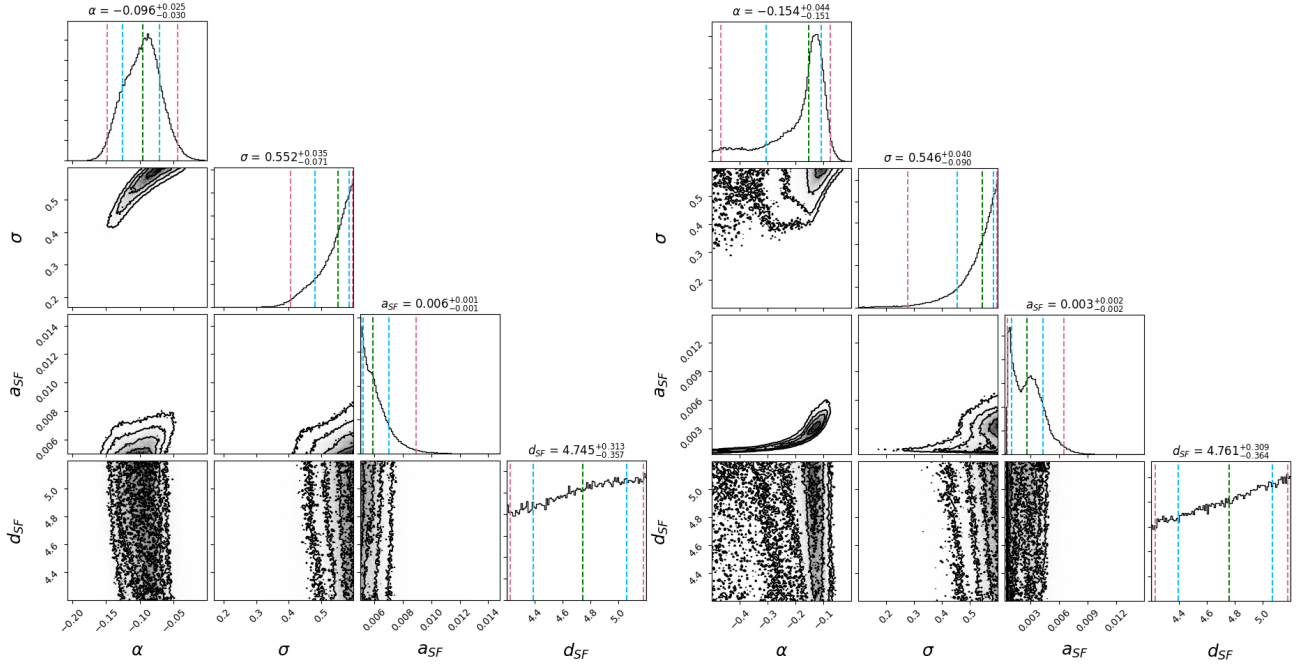


Figure 7. Corner plot for MSSFR parameter inference on LVK data with $p_{astro} \geq 0.95$ using the original priors (left panel) and using the extended prior on a_{SF} (right panel).




Figure 8. Cumulative Density Functions for chirp mass (left panel) and redshift (right panel). In each panel the 100 orange curves are randomly sampled cumulative density functions from the posterior samples (with LVK priors) for each of the DCO mergers detected during the third observing run ($p_{astro} \geq 0.95$), and indicates the spread due to measurement uncertainty. The 100 blue curves are cumulative density functions constructed by sampling 52 events from the predicted distribution of detectable DCO mergers from the best-fit λ and indicate the sample variance.

Here, we present the results of this analysis. However, a few disclaimers are warranted. We only varied parameters in a MSSFR model, while fixing the assumptions and parameters describing stellar and binary evolution (mass transfer, winds, supernovae, etc.) to fiducial, but likely incorrect, values. We also used a simple approximation for the selection function imposed by LVK observations, picking a sample noise spectrum rather than a variable one and using a single-detector signal-to-noise ratio threshold of 8, independent of the source parameters, as a proxy for detectability by the network. A more realistic selection function would be based on an actual injection campaign into the detector noise as accumulated over the third LVK observing run.

The posteriors on the MSSFR parameters on the LVK O3 data are shown in the left panel of Figure 7. The highest-likelihood point found by the MCMC chain is at $\lambda(\alpha, \sigma, a_{SF}, d_{SF}) = (-0.085, 0.598, 0.005, 5.042)$. It is clear that the posteriors rail against the prior boundary on a_{SF} .

We therefore repeated the analysis by extending the uniform prior on a_{SF} to incorporate the range $[0, 0.015]$. Fortunately, a_{SF} is a scaling factor only, so this did not require re-training the neural network. We used the surrogate model prediction at the mid-point of the a_{SF} range over which it was trained to predict the detection rate matrix at values of a_{SF} outside this range. Specifically, we assumed that the surrogate model for the detection rate matrix at $a_{SF} < 0.005$ was equal to $(a_{SF}/0.01)$ times the prediction of the surrogate model at $\{a_{SF} = 0.01, \lambda'\}$, where $\lambda' \equiv \lambda/a_{SF}$ refer to all other parameters. We show the posteriors on MSSFR parameters with this extended prior in the right panel of Figure 7. The highest-likelihood point found by the MCMC chain with extended priors is at $\lambda(\alpha, \sigma, a_{SF}, d_{SF}) = (-0.115, 0.599, 0.003, 5.158)$. These posteriors rail against the prior on σ , suggesting that an even broader distribution of metallicities is preferred under this MSSFR model. On the other hand, the preferred value of the normalisation of the star-formation rate a_{SF} is a factor of ~ 5 lower than the estimate of Madau & Dickinson (2014).

In Figure 8 we show the cumulative density functions (CDFs) for chirp mass (left panel) and redshift (right panel), for both LVK observations and for our best-fit λ (from the analysis with the extended prior on a_{SF}). This provides a visual posterior predictive check. In this case, our preferred MSSFR model fails this check. The paucity of higher-redshift observations in the model prediction likely indicates an overly conservative detectability estimate, probably through the use of a pessimistic threshold on detectable signal-to-noise ratio and/or a

pessimistic noise power spectral density estimate. While this problem is relatively easy to fix, the failure to reproduce the observed features of the chirp mass distribution, such as the sharp, narrow peak near $\mathcal{M} \approx 10 M_{\odot}$, speaks to a flawed underlying binary evolution model. As noted above, we did not consider variations in the stellar and binary evolution assumptions and model parameters, as this was only intended as a proof-of-concept exercise.

Moreover, some of the events, particularly the most massive ones, have been conjectured to arise from channels other than isolation binary evolution, which we did not consider here (see, e.g., Mandel & Farmer 2022 for a review). Other channels may be particularly relevant for higher-mass binaries, including hierarchical dynamical mergers (e.g., Gerosa & Fishbach 2021) or mergers in AGN disks in which black holes could grow by accretion (e.g., Tagawa et al. 2021). We therefore consider an alternative data set in which we remove the five LVK binaries whose median source-frame chirp mass exceed $40M_{\odot}$ and reweigh the posteriors for remaining binaries by imposing a strict chirp mass prior upper limit of $40M_{\odot}$. We also remove any binaries with chirp mass exceeding $40M_{\odot}$ from the forward models. The resulting posteriors on MSSFR parameters (using the extended prior on a_{SF}) and the cumulative posterior predictive check on the chirp mass distribution are shown in Figure 9. The posteriors on MSSFR parameters shift within one standard deviation relative to Figure 7, allowing for a more consistent chirp mass posterior predictive check than in Figure 8, although our model still fails to reproduce the narrow peak near $\mathcal{M} \approx 10 M_{\odot}$. However, this analysis is based on a single model of binary evolution and only considers the impact of MSSFR parameters; complete inference must simultaneously incorporate the parameters governing stellar and binary evolution along with MSSFR models.

3.2.4. Inference performance

All MCMC searches performed for this study used an ensemble of 10 walkers (single-threaded), and were allowed to run for 100,000 steps per walker, resulting in 1,000,000 steps per search. Each step involves the creation of a new detection rate matrix (for the values of λ visited at that step). We used the ANNs trained on the 100-sample dataset, so our surrogate model was comprised of 293 ANNs, with an execution time of $\sim 6.42 \times 10^{-3}$ seconds (Section 2.4.3).

The elapsed time for each MCMC search using our surrogate model is $\sim 6.42 \times 10^{-3} \times 1,000,000 = \sim 6,420$ seconds (~ 1.8 hours), ignoring MCMC algorithm overheads, which are small compared to the execution time

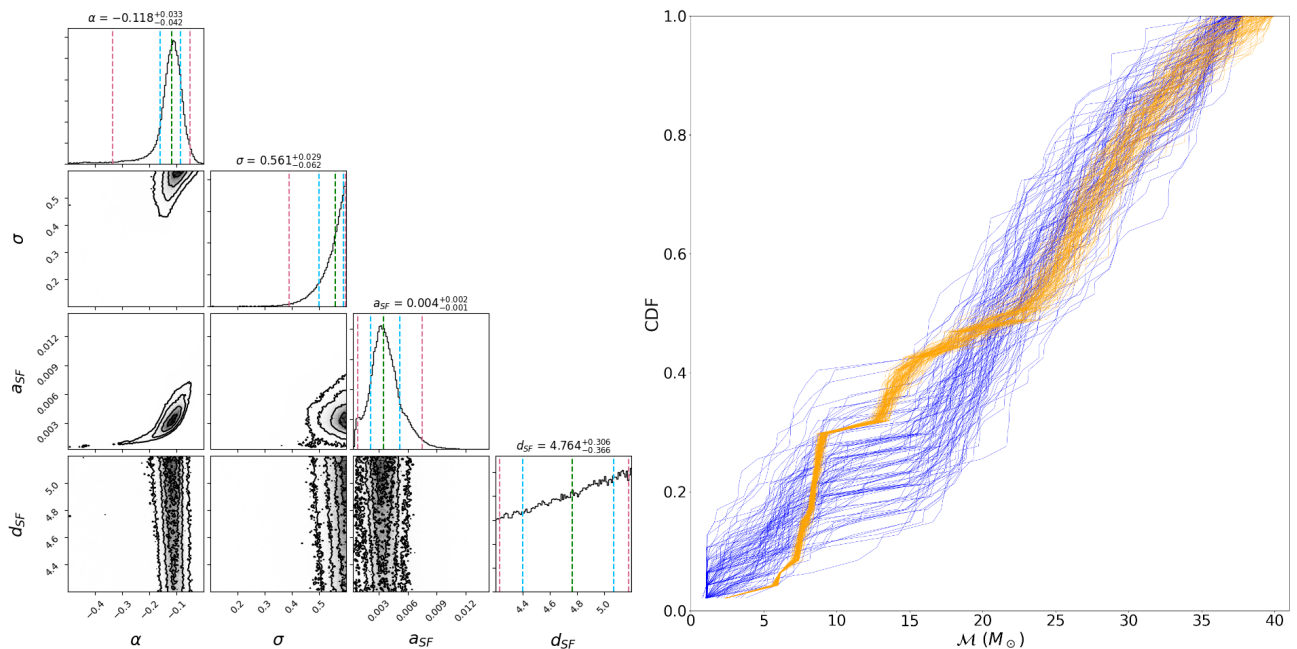


Figure 9. Inference on LVK data after removing the 5 most massive binaries from the data set (see section 3.2.3). Left panel: Corner plot for MSSFR parameter inference on LVK data, as in Figure 7; right panel: posterior predictive check on the chirp mass distribution for the associated best-fit parameters, as in Figure 8.

of the ANNs. The time to create all training data sets was 142 hours, and the time to train the ANNs was 12 hours, for a total time 156 hours.

The COMPAS post-processing tools execution time to create a single detection rate matrix is ~ 112 seconds (Section 2.4.3). The elapsed time for equivalent MCMC searches using the COMPAS post-processing tools to create the detection rate matrices is $\sim 112 \times 1,000,000 = \sim 112,000,000$ seconds (~ 3.5 years), ignoring MCMC algorithm overheads, which are small compared to the execution time of the COMPAS post-processing tools. Thus, the use of surrogate models reduced the total inference cost by a factor of 200 after accounting for the time needed to create training data and train these models.

4. CONCLUDING REMARKS

We conducted a proof-of-concept study that showed that it is possible to construct an interpolant that can calculate DCO merger detection rates for a set of astrophysical parameters, with fairly good accuracy, and in milliseconds (several orders of magnitude faster than methods currently available). This allows us to create a large synthetic state-space in a very reasonable timeframe, which can then be used to conduct inference about the parameters.

The method we developed is fast, flexible, highly parallelisable, and robust - the constituent ANNs can be re-trained, or replaced, individually and as necessary. We

achieved a total analysis cost reduction by a factor of 200 after accounting for surrogate model training.

We found that our first attempt at analysis on LVK data pushed our inference to rail against the prior boundaries. In this case this was likely due to poor choices in the fiducial model describing astrophysical evolution, exacerbated by an overly simplistic observational selection function. We were able to partially remedy this issue because one of the parameters required only a simple rescaling. However, to avoid similar problems in future analyses, it would be wise to use active learning and grow the training data set in regions of parameter space which provide models with the best match to the data.

For this study we focussed on just the four MSSFR parameters used in the post-processing of COMPAS simulations, but we are not limited to those parameters, or to the post-processing code. The next step is to expand this work to other astrophysical parameters that govern stellar and binary evolution in the population synthesis code itself (e.g. COMPAS), rather than just the post-processing code, thereby reducing the need for time-consuming simulations.

ACKNOWLEDGEMENTS

Simulations in this paper made use of the COMPAS rapid population synthesis code, which is freely available at <http://github.com/TeamCOMPAS/COMPAS>. The

version of COMPAS used for these simulations was v02.31.00.

IM is a recipient of the Australian Research Council Future Fellowship FT190100574.

DATA AVAILABILITY

The data underlying this article is available from the corresponding author upon request.

REFERENCES

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/>
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020, Living Rev. Relativ., 23, 3, doi: [10.1007/s41114-020-00026-9](https://doi.org/10.1007/s41114-020-00026-9)
- Andrieu, C., de Freitas, N., Doucet, A., & Jordan, M. 2003, Machine Learning, 50, 5, doi: [10.1023/A:1020281327116](https://doi.org/10.1023/A:1020281327116)
- Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, ARA&A, 47, 481, doi: [10.1146/annurev.astro.46.060407.145222](https://doi.org/10.1146/annurev.astro.46.060407.145222)
- Barrett, J. W., Gaebel, S. M., Neijssel, C. J., et al. 2018, Monthly Notices of the Royal Astronomical Society, 477, 4685, doi: [10.1093/mnras/sty908](https://doi.org/10.1093/mnras/sty908)
- Barrett, J. W., Mandel, I., Neijssel, C. J., Stevenson, S., & Vigna-Gómez, A. 2016, Proceedings of the International Astronomical Union, 12, 46, doi: [10.1017/s1743921317000059](https://doi.org/10.1017/s1743921317000059)
- Belczynski, K., Bulik, T., Fryer, C. L., et al. 2010, ApJ, 714, 1217, doi: [10.1088/0004-637X/714/2/1217](https://doi.org/10.1088/0004-637X/714/2/1217)
- Belczynski, K., Kalogera, V., Rasio, F. A., et al. 2008, ApJSupplement, 174, 223, doi: [10.1086/521026](https://doi.org/10.1086/521026)
- Bhattacharya, D., & van den Heuvel, E. P. J. 1991, PhR, 203, 1, doi: [10.1016/0370-1573\(91\)90064-S](https://doi.org/10.1016/0370-1573(91)90064-S)
- Breivik, K., Coughlin, S., Zevin, M., et al. 2020, ApJ, 898, 71, doi: [10.3847/1538-4357/ab9d85](https://doi.org/10.3847/1538-4357/ab9d85)
- Broekgaarden, F. S., Justham, S., de Mink, S. E., et al. 2019, MNRAS, 490, 5228, doi: [10.1093/mnras/stz2558](https://doi.org/10.1093/mnras/stz2558)
- Broekgaarden, F. S., Berger, E., Stevenson, S., et al. 2022, MNRAS, 516, 5737, doi: [10.1093/mnras/stac1677](https://doi.org/10.1093/mnras/stac1677)
- Cheung, D. H. T., Wong, K. W. K., Hannuksela, O. A., Li, T. G. F., & Ho, S. 2022, PhRvD, 106, 083014, doi: [10.1103/PhysRevD.106.083014](https://doi.org/10.1103/PhysRevD.106.083014)
- Chizat, L., Oyallon, E., & Bach, F. 2019, in Proceedings of the 33rd International Conference on Neural Information Processing Systems (Red Hook, NY, USA: Curran Associates Inc.)
- Chollet, F., et al. 2015, Keras, <https://github.com/fchollet/keras>, GitHub
- Dasgupta, S., & Tosh, C. 2020, Expressivity of expand-and-sparsify representations, arXiv, doi: [10.48550/ARXIV.2006.03741](https://doi.org/10.48550/ARXIV.2006.03741)
- de Kool, M. 1990, ApJ, 358, 189, doi: [10.1086/168974](https://doi.org/10.1086/168974)
- Dominik, M., Belczynski, K., Fryer, C., et al. 2012, ApJ, 759, 52, doi: [10.1088/0004-637X/759/1/52](https://doi.org/10.1088/0004-637X/759/1/52)
- Farr, W. M., Gair, J. R., Mandel, I., & Cutler, C. 2015, PhRvD, 91, 023005, doi: [10.1103/PhysRevD.91.023005](https://doi.org/10.1103/PhysRevD.91.023005)
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306, doi: [10.1086/670067](https://doi.org/10.1086/670067)
- Frankle, J., & Carbin, M. 2019, in International Conference on Learning Representations. <https://openreview.net/forum?id=rJl-b3RcF7>
- Frankle, J., Schwab, D. J., & Morcos, A. S. 2020, Training BatchNorm and Only BatchNorm: On the Expressive Power of Random Features in CNNs, arXiv, doi: [10.48550/ARXIV.2003.00152](https://doi.org/10.48550/ARXIV.2003.00152)
- Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, ApJ, 749, 91, doi: [10.1088/0004-637X/749/1/91](https://doi.org/10.1088/0004-637X/749/1/91)
- Gerosa, D., & Fishbach, M. 2021, Nat. Astron., 5, 749, doi: [10.1038/s41550-021-01398-w](https://doi.org/10.1038/s41550-021-01398-w)
- Hamann, W. R., & Koesterke, L. 1998, A&A, 335, 1003
- Heger, A., & Woosley, S. E. 2002, ApJ, 567, 532, doi: [10.1086/338487](https://doi.org/10.1086/338487)
- Hobbs, G., Lorimer, D. R., Lyne, A. G., & Kramer, M. 2005, MNRAS, 360, 974, doi: [10.1111/j.1365-2966.2005.09087.x](https://doi.org/10.1111/j.1365-2966.2005.09087.x)
- Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, MNRAS, 329, 897, doi: [10.1046/j.1365-8711.2002.05038.x](https://doi.org/10.1046/j.1365-8711.2002.05038.x)
- Iorio, G., Costa, G., Mapelli, M., et al. 2022, arXiv e-prints, doi: [10.48550/arXiv.2211.11774](https://doi.org/10.48550/arXiv.2211.11774)
- Izzard, R. G., Preece, H., Jofre, P., et al. 2018, MNRAS, 473, 2984, doi: [10.1093/mnras/stx2355](https://doi.org/10.1093/mnras/stx2355)
- Kroupa, P. 2001, MNRAS, 322, 231, doi: [10.1046/j.1365-8711.2001.04022.x](https://doi.org/10.1046/j.1365-8711.2001.04022.x)
- Langer, N., & Norman, C. A. 2006, The Astrophysical Journal, 638, L63, doi: [10.1086/500363](https://doi.org/10.1086/500363)

- LIGO Scientific and Virgo Collaborations. 2022, GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run - Parameter Estimation Data Release, v2, Zenodo, doi: [10.5281/zenodo.6513631](https://doi.org/10.5281/zenodo.6513631)
- LIGO Scientific, Virgo and KAGRA Collaborations. 2021, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run – Parameter estimation data release, Zenodo, doi: [10.5281/zenodo.5546663](https://doi.org/10.5281/zenodo.5546663)
- Lin, L., Bingham, D., Broekgaarden, F., & Mandel, I. 2021, *The Annals of Applied Statistics*, 15, 1604 , doi: [10.1214/21-AOAS1484](https://doi.org/10.1214/21-AOAS1484)
- Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415, doi: [10.1146/annurev-astro-081811-125615](https://doi.org/10.1146/annurev-astro-081811-125615)
- Mandel, I., & Farmer, A. 2022, *Physics Reports*, 955, 1, doi: [10.1016/j.physrep.2022.01.003](https://doi.org/10.1016/j.physrep.2022.01.003)
- Mandel, I., Farr, W. M., & Gair, J. R. 2019, *MNRAS*, 486, 1086, doi: [10.1093/mnras/stz896](https://doi.org/10.1093/mnras/stz896)
- Marchant, P., Langer, Norbert, Podsiadlowski, Philipp, Tauris, Thomas M., & Moriya, Takashi J. 2016, *A&A*, 588, A50, doi: [10.1051/0004-6361/201628133](https://doi.org/10.1051/0004-6361/201628133)
- Marchant, P., & Moriya, T. J. 2020, *A&A*, 640, L18, doi: [10.1051/0004-6361/202038902](https://doi.org/10.1051/0004-6361/202038902)
- Marchant, P., Renzo, M., Farmer, R., et al. 2019, *ApJ*, 882, 36, doi: [10.3847/1538-4357/ab3426](https://doi.org/10.3847/1538-4357/ab3426)
- Massevitch, A., & Yungelson, L. 1975, *Mem. Soc. Astron. Italiana*, 46, 217
- Mukherjee, S., & Huberman, B. A. 2022, *Why Neural Networks Work*, arXiv, doi: [10.48550/ARXIV.2211.14632](https://doi.org/10.48550/ARXIV.2211.14632)
- Neijssel, C. J., Vigna-Gómez, A., Stevenson, S., et al. 2019, *MNRAS*, 490, 3740, doi: [10.1093/mnras/stz2840](https://doi.org/10.1093/mnras/stz2840)
- Pfahl, E., Rappaport, S., & Podsiadlowski, P. 2002, *ApJ*, 573, 283, doi: [10.1086/340494](https://doi.org/10.1086/340494)
- Podsiadlowski, P., Langer, N., Poelarends, A. J. T., et al. 2004, *ApJ*, 612, 1044, doi: [10.1086/421713](https://doi.org/10.1086/421713)
- Powell, J., Stevenson, S., Mandel, I., & TiÅo, P. 2019, *MNRAS*, 488, 3810, doi: [10.1093/mnras/stz1938](https://doi.org/10.1093/mnras/stz1938)
- Rasmussen, C. E., & Williams, C. K. I. 2005, *Gaussian Processes for Machine Learning* (The MIT Press), doi: [10.7551/mitpress/3206.001.0001](https://doi.org/10.7551/mitpress/3206.001.0001)
- Riley, J., Mandel, I., Marchant, P., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 663, doi: [10.1093/mnras/stab1291](https://doi.org/10.1093/mnras/stab1291)
- Russell, S. J., & Norvig, P. 1995, *Artificial intelligence: a modern approach* (Englewood Cliffs, N.J.: Prentice Hall), 111–113
- Soberman, G. E., Phinney, E. S., & van den Heuvel, E. P. J. 1997, *A&A*, 327, 620
- Stevenson, S., Sampson, M., Powell, J., et al. 2019, *ApJ*, 882, 121, doi: [10.3847/1538-4357/ab3981](https://doi.org/10.3847/1538-4357/ab3981)
- Stevenson, S., Vigna-Gómez, A., Mandel, I., et al. 2017, *Nature Communications*, 8, 14906, doi: [10.1038/ncomms14906](https://doi.org/10.1038/ncomms14906)
- Tagawa, H., Kocsis, B., Haiman, Z., et al. 2021, *ApJ*, 908, 194, doi: [10.3847/1538-4357/abd555](https://doi.org/10.3847/1538-4357/abd555)
- Tauris, T. M., Langer, N., & Podsiadlowski, P. 2015, *MNRAS*, 451, 2123, doi: [10.1093/mnras/stv990](https://doi.org/10.1093/mnras/stv990)
- Tauris, T. M., & van den Heuvel, E. P. J. 2006, *Formation and evolution of compact stellar X-ray sources*, Vol. 39, 623–665
- Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, *ApJ*, 846, 170, doi: [10.3847/1538-4357/aa7e89](https://doi.org/10.3847/1538-4357/aa7e89)
- Taylor, S. R., & Gerosa, D. 2018, *Phys. Rev. D*, 98, 083017, doi: [10.1103/PhysRevD.98.083017](https://doi.org/10.1103/PhysRevD.98.083017)
- Team COMPAS: Riley, J., Agrawal, P., Barrett, J. W., et al. 2022a, *ApJS*, 258, 34, doi: [10.3847/1538-4365/ac416c](https://doi.org/10.3847/1538-4365/ac416c)
- . 2022b, *Journal of Open Source Software*, 7, 3838, doi: [10.21105/joss.03838](https://doi.org/10.21105/joss.03838)
- Timmes, F. X., Woosley, S. E., & Weaver, T. A. 1996, *ApJ*, 457, 834, doi: [10.1086/176778](https://doi.org/10.1086/176778)
- Van Rossum, G., & Drake, F. L. 2009, *Python 3 Reference Manual* (Scotts Valley, CA: CreateSpace)
- van Son, L. A. C., de Mink, S. E., Chruslinska, M., et al. 2022, arXiv e-prints, arXiv:2209.03385. <https://arxiv.org/abs/2209.03385>
- Vigna-Gómez, A., Neijssel, C. J., Stevenson, S., et al. 2018, *MNRAS*, 481, 4009, doi: [10.1093/mnras/sty2463](https://doi.org/10.1093/mnras/sty2463)
- Vinciguerra, S., Neijssel, C. J., Vigna-Gómez, A., et al. 2020, *MNRAS*, 498, 4705, doi: [10.1093/mnras/staa2177](https://doi.org/10.1093/mnras/staa2177)
- Vink, J. S., & de Koter, A. 2005, *A&A*, 442, 587, doi: [10.1051/0004-6361:20052862](https://doi.org/10.1051/0004-6361:20052862)
- Vink, J. S., de Koter, A., & Lamers, H. J. G. L. M. 2000, *A&A*, 362, 295
- . 2001, *A&A*, 369, 574, doi: [10.1051/0004-6361:20010127](https://doi.org/10.1051/0004-6361:20010127)
- Webbink, R. F. 1984, *ApJ*, 277, 355, doi: [10.1086/161701](https://doi.org/10.1086/161701)
- Wong, K. W. K., Breivik, K., Kremer, K., & Callister, T. 2021, *Phys. Rev. D*, 103, 083021, doi: [10.1103/PhysRevD.103.083021](https://doi.org/10.1103/PhysRevD.103.083021)
- Wong, K. W. K., & Gerosa, D. 2019, *Phys. Rev. D*, 100, 083015, doi: [10.1103/PhysRevD.100.083015](https://doi.org/10.1103/PhysRevD.100.083015)
- Xu, X.-J., & Li, X.-D. 2010a, *ApJ*, 716, 114, doi: [10.1088/0004-637X/716/1/114](https://doi.org/10.1088/0004-637X/716/1/114)
- . 2010b, *ApJ*, 722, 1985, doi: [10.1088/0004-637X/722/2/1985](https://doi.org/10.1088/0004-637X/722/2/1985)

Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O.
2017, in International Conference on Learning
Representations.
<https://openreview.net/forum?id=Sy8gdB9xx>

Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O.
2021, Commun. ACM, 64, 107–115, doi: [10.1145/3446776](https://doi.org/10.1145/3446776)
Zhang, C., Bengio, S., & Singer, Y. 2022, Journal of
Machine Learning Research, 23, 1.
<http://jmlr.org/papers/v23/20-069.html>

APPENDIX

A. ANN TRAINING AND PERFORMANCE

We did not use extremely large datasets to train the ANNs that comprise the interpolant in this work. Even so, we achieved very good results with the ANNs we trained - most achieved 100% accuracy when tested on the training and validation data. These results might lead us to wonder whether the ANNs have been over-fitted during training (though we did test that they are capable of generalising).

How many training examples are required to properly train an ANN? The answer depends on the complexity of the problem, and the ANN being used to solve it. It is effectively unknowable *a priori*, and must be discovered through empirical investigation. In other words, we can tell whether the training data were sufficient by testing that the ANN gets the right answer for the examples it was trained on, and whether it predict correct answers for inputs it was not trained / validated on. But were the data necessary?

The conventional wisdom is that at least thousands of training examples are required, certainly no fewer than hundreds. Tens, or hundreds, of thousands for an "average" modelling problem; millions, or tens of millions, for a "hard" problem. This basically translates to "Get as much data as possible and use it" - not an unreasonable strategy in a situation where a better answer is not immediately apparent.

We know that when a sufficiently large ANN tries to learn from a small dataset it will tend to memorise the entire dataset - an example of over-fitting. But does this compromise the ability of the network to generalise beyond the the training dataset? In the past we might have thought so, but recent work suggests that this is not necessarily the case (Zhang et al. (2017), Zhang et al. (2021)).

Mukherjee & Huberman (2022) claim that the properties of a particular mathematical operation, *expand-and-sparsify* (Dasgupta & Tosh 2020), explain the ability of an over-parameterised ANN (i.e. the number of parameters in the model exceeds the size of the training dataset) to both memorise the entire training dataset *and* generalise beyond the training dataset (by learning the underlying structure to the training data).

As noted above, we achieved very good results with the ANNs we trained, yet in most cases the number of epochs over which the networks were trained, and so the time to train, was not particularly onerous. This, especially when coupled with the performance achieved by the ANNs when trained on a relatively small dataset, raises the question of why it wasn't more difficult and time-consuming to find a solution to such a difficult problem. The answer is two-fold.

First, the phenomenon of untrained networks with randomly-initialised weights performing surprisingly well has been reported in the past (e.g. Zhang et al. (2022), Frankle et al. (2020)). This is often explained by referring to the Lottery Ticket Hypothesis (Frankle & Carbin 2019), which posits that ANNs are really running large lotteries where sub-networks whose weights have been serendipitously randomly-initialised to values that produce good results are scattered throughout the network. Thus, in a sufficiently large, untrained, and randomly-initialised ANN, a sub-network with random weights that performs as well as a trained network can be found, and the remainder of the network is ignored.

Next, as noted in Section 2.4.2, we trained the ANNs using the Keras Adam optimiser, which uses an enhanced stochastic gradient descent algorithm. We know that the standard stochastic gradient descent algorithm is usually sufficient to train an ANN - it is not often that other methods of optimisation for the network parameters are required. Chizat et al. (2019) note that often during training, the parameters of an ANN change very little from their initialised values. Mukherjee and Huberman claim that since very often the randomly-initialised weights are "good enough" (as described above), they require hardly any fine-tuning for better performance, so an optimiser such as stochastic gradient descent needs to do very little work to train a network to achieve reasonably good performance. Indeed, Mukherjee & Huberman (2022) contend that the training of an ANN is the third most important component of a successful ANN solution, behind architecture selection and parameter initialisation.

B. COMPAS CONFIGURATION FIDUCIAL VALUES

Table 4. Initial values and default settings for the COMPAS fiducial model.

Description and name	Value/range	Note / setting
Initial conditions		
Initial primary mass $m_{1,i}$	[5, 150] M_{\odot}	Kroupa (2001) IMF $\propto m_{1,i}^{-\alpha}$ with $\alpha_{\text{IMF}} = 2.3$ for stars above $5 M_{\odot}$
Initial mass ratio $q_i = m_{2,i}/m_{1,i}$	[0, 1]	We assume a flat mass ratio distribution $p(q_i) \propto 1$ with $m_{2,i} \geq 0.1 M_{\odot}$
Initial semi-major axis a_i	[0.01, 1000] AU	Distributed flat-in-log $p(a_i) \propto 1/a_i$
Initial metallicity Z_i	[0.0001, 0.03]	Distributed uniform-in-log
Initial orbital eccentricity e_i	0	All binaries are assumed to be circular at birth
Fiducial parameter settings:		
Stellar winds for hydrogen rich stars	Belczynski et al. (2010)	Based on Vink et al. (2000, 2001), including LBV wind mass loss with $f_{\text{LBV}} = 1.5$
Stellar winds for hydrogen-poor helium stars	Belczynski et al. (2010)	Based on Hamann & Koesterke (1998) and Vink & de Koter 2005
Max transfer stability criteria	ζ -prescription	Based on Vigna-Gómez et al. (2018) and references therein
Mass transfer accretion rate	thermal timescale	Limited by thermal timescale for stars Hurley et al. (2002); Vinciguerra et al. (2020)
	Eddington-limited	Accretion rate is Eddington-limit for compact objects
Non-conservative mass loss	isotropic re-emission	Massevitch & Yungelson (1975); Bhattacharya & van den Heuvel (1991); Soberman et al. (1997) Tauris & van den Heuvel (2006)
Case BB mass transfer stability	always stable	Based on Tauris et al. (2015, 2017); Vigna-Gómez et al. (2018)
CE prescription	$\alpha - \lambda$	Based on Webbink (1984); de Kool (1990)
CE efficiency α -parameter	1.0	
CE λ -parameter	λ_{Nanjing}	Based on Xu & Li (2010a,b) and Dominik et al. (2012)
Hertzsprung gap (HG) donor in CE	pessimistic	Defined in Dominik et al. (2012): HG donors don't survive a CE phase
SN natal kick magnitude v_k	$[0, \infty)$ km s^{-1}	Drawn from Maxwellian distribution with standard deviation $\sigma_{\text{rms}}^{\text{1D}}$
SN natal kick polar angle θ_k	$[0, \pi]$	$p(\theta_k) = \sin(\theta_k)/2$
SN natal kick azimuthal angle ϕ_k	$[0, 2\pi]$	Uniform $p(\phi) = 1/(2\pi)$
SN mean anomaly of the orbit	$[0, 2\pi]$	Uniformly distributed
Core-collapse SN remnant mass prescription	delayed	From Fryer et al. (2012), which has no lower BH mass gap
USSN remnant mass prescription	delayed	From Fryer et al. (2012)
ECSN remnant mass prescription	$m_f = 1.26 M_{\odot}$	Based on Equation 8 in Timmes et al. (1996)
Core-collapse SN velocity dispersion $\sigma_{\text{rms}}^{\text{1D}}$	265 km s^{-1}	1D rms value based on Hobbs et al. (2005)
USSN and ECSN velocity dispersion $\sigma_{\text{rms}}^{\text{1D}}$	30 km s^{-1}	1D rms value based on e.g., Pfahl et al. (2002); Podsiadlowski et al. (2004)
PISN / PPISN remnant mass prescription	Marchant et al. (2019)	As implemented in Stevenson et al. (2019)
Maximum NS mass	$\text{max}_{\text{NS}} = 2.5 M_{\odot}$	Following Fryer et al. (2012)
Tides and rotation		We do not include tides and/or rotation in this study
Binary fraction	$f_{\text{bin}} = 0.7$	
Solar metallicity Z_{\odot}	$Z_{\odot} = 0.0142$	based on Asplund et al. (2009)
Simulation settings		
Binary population synthesis code	COMPAS	Stevenson et al. (2017); Barrett et al. (2018); Vigna-Gómez et al. (2018); Neijssel et al. (2019) Broekgaarden et al. (2019); Team COMPAS: Riley, J. et al. (2022a).

Chapter 5

Summary and discussion

The major outcomes of the work conducted for this thesis are:

- Major rewrite and public release of the COMPAS population synthesis toolset.
We took COMPAS version 1, a stellar and binary population syntheses computer model, and by refactoring the source code into a well-structured, object-oriented, and modular application, we made it more robust and efficient, as well as more easily extensible and maintainable. We made it more accessible to a broader user and developer base by releasing COMPAS version 2 as publicly available open-source code.
- Study of the chemically homogeneous evolution (CHE) evolutionary track.
We extended COMPAS version 2 to include the CHE track to provide an alternative method to synthesise populations of DCOs that could merge within the age of the Universe - significant because BBHs formed through CHE may have unique properties that provide a signature of the formation pathway that could be evident in the gravitational waves emitted as they merge. We found that the CHE channel may contribute more than half, and perhaps as much as three quarters, of all aLIGO BBH detections arising from isolated binary evolution.
- Proof-of-concept development of a fast surrogate model for COMPAS.
We used COMPAS version 2 to generate data that we then used to train a surrogate model for COMPAS itself that we constructed using artificial neural networks (ANNs). The surrogate model allows us to rapidly explore the parameter space of processes underlying the evolution of stars and binary systems, and infer constraints on those parameters.

In Chapter 2 we described the work done to improve and extend the COMPAS stellar and binary population synthesis suite of tools. The architecture of the C++ core of the toolset was improved, and the code base rewritten, to reflect the modern object-oriented programming paradigm. COMPAS code is designed to make it easy to specify desired parameterised prescriptions, and introduce new models for various stages of stellar and binary evolution. COMPAS version 2 is much easier to extend and maintain.

We improved the robustness of the code, and the output it produces. We improved efficiency to the extent that COMPAS version 2 is at least an order of magnitude faster than COMPAS version 1: COMPAS version 2 is capable of evolving ~ 100 binaries per second on a modern laptop computer. Given its parallel structure (jobs can be readily split across multiple cores with no need for communication until results are combined), a population of a billion binaries can be evolved in ~ 24 hours on a modest 128-core cluster.

The user interface was rewritten: a new, comprehensive logging subsystem was added to the C++ code, as was flexible and extensible functionality for users to specify a grid of systems to evolve (specifying initial attributes of each system, rather than just a population of randomly sampled systems).

Source and quality control were improved: code versioning, change logs, and documentation in the form of an online user and developer guide were added.

In [1] we noted that COMPAS has been used extensively to investigate the properties of compact binaries (we refer the reader to Riley et al. [1] for details of a selection of those publications).

Since its release as open-source software in March 2020 [73], COMPAS has continued to develop an active, rapidly growing, user and developer base, and through them COMPAS continues to improve and expand. A representative selection of work published after the public release of COMPAS, and so not detailed in [1], for which COMPAS was used as a tool to develop the research follows. The list presented here, and the list presented in Riley et al. [1], is a good indication of the increasing impact that COMPAS is having in the field.

van Son et al. [90] used COMPAS to investigate the redshift evolution of the BBH merger rate. Mandel and Farmer [2] studied formation channels of merging stellar-mass black-hole binaries. Broekgaarden et al. [91] studied the impact of massive binary star and cosmic evolution on gravitational wave observations. van Son et al. [92] explored the stable mass transfer channel for the formation of GW sources. Stevenson et al. [93] used COMPAS to model wide binary pulsars from electron-capture supernovae. Wagg

et al. [94] present new simulations of the populations of binary black holes, black hole-neutron stars, and double neutron stars that will be detectable by the planned space-based gravitational-wave detector, Laser Interferometer Space Antenna (LISA) [95, 96]. Broekgaarden et al. [97] study signatures of mass ratio reversal in gravitational waves from merging binary black holes. D’Orazio et al. [98] used synthetic populations created by COMPAS in a multi-messenger analysis of BHNS formation scenarios and mergers, and the physics of BHNS magnetospheric interactions. Pejcha et al. [99] use COMPAS in a study of the attributes of the double-eclipsing binary CzeV343. Raveh et al. [100] used synthetic populations of BBHs created by COMPAS to study the effect of flyby perturbations on the production of GW sources. Yarza et al. [101] used data created by COMPAS in a study to develop a revised energy formalism for common envelope evolution (CEE). van Son et al. [102] used observations and predictions from COMPAS to study, and propose a simple analytical function for, the metallicity-dependent cosmic star formation history. Moreno Méndez et al. [103] used COMPAS to evolve the inner massive binary in hierarchical triples to study the effects of hypercritical accretion during common envelopes in triples, highlighting the possibility of BBHs forming in the pair-instability supernova (PISN) mass gap. Gupta et al. [104] explore the BHNS detection and measurement capabilities of proposed new GW detectors, as well as sensitivity improvements to current detectors, and use COMPAS to create their fiducial model for BHNS systems that form through the isolated binary formation channel. Rauf et al. [105] explore the connection between the GW merger rates of stellar-mass BBHs and galaxy properties by generating populations of stars using COMPAS and evolving them in galaxies using the semi-analytic galaxy formation model Shark [106]. Grichener [107] presents a population synthesis study of mergers of neutron stars and black holes with cores of giant stars during CEE. Vigna-Gómez and Ramirez-Ruiz [108] explore the landscape of compact objects in the context of microlensing targets using binary population synthesis models created by COMPAS. Agudo et al. [109] present the results from multi-wavelength observations of a transient discovered during a follow-up campaign of a GW event reported as a possible binary neutron star merger, and use ultra-stripped supernovae (USSNe) rates calculated by COMPAS as a comparison against observations. de Sá et al. [110] use populations created by COMPAS to study the effects of their implementation of a metallicity- and redshift-dependent initial mass function (IMF) on population properties, particularly merger rates. Mosbech et al. [111] use properties of a synthetic population created by COMPAS to investigate the merger rate of stellar-mass BBHs throughout cosmic time as a new observational probe of dark matter microphysics.

In Chapter 3 we described the model of chemically homogeneous evolution (CHE) that we implemented in COMPAS. We used MESA models to determine the critical rotation thresholds for CHE, and provided fits that can be used in other rapid binary population

synthesis applications. We confirmed, through the analysis of a synthesised population of 12 million binary systems over a range of metallicities and Wolf-Rayet (WR) wind mass-loss rates that our simplified models match detailed binary evolution simulations well [85, 86].

We investigated the contribution of CHE and non-CHE channels to BBH formation under a single set of assumptions. As noted above, we found that the CHE channel may contribute more than half, and perhaps as much as three quarters, of all aLIGO BBH detections arising from isolated binary evolution. CHE BBHs may represent $\gtrsim 80\%$ of detectable sources with the highest chirp masses of $\gtrsim 30 M_{\odot}$. A comparison between our model population and the population of detected binaries from the first two advanced detector observing runs indicates that the current model over-predicts the total number of sources by a factor of ~ 2 , but matches the observed chirp mass distribution.

As noted in Chapter 3, work is already underway to improve and extend the implementation of CHE in COMPAS [88], which may alleviate some of the issues we see with the current implementation over-producing BBHs. Stevenson and Houlden [88] in turn suggest further extensions to the model.

In Chapter 4 we described a proof-of-concept study that showed that it is possible to construct a surrogate model that can predict DCO merger detection rates for a set of astrophysical parameters, with fairly good accuracy, and in milliseconds (several orders of magnitude faster than methods currently available). This allows us to create a large synthetic state-space in a very reasonable timeframe, which can then be used to conduct inference about the parameters.

The surrogate model we constructed consists of a number of ANNs, and is able to interpolate across the entire range of parameter values for which it was trained. The method we developed is fast, flexible, highly parallelisable, and robust - the constituent ANNs can be retrained, or replaced, individually and as necessary. We achieved a total analysis cost reduction by a factor of ~ 200 after accounting for the surrogate model training time.

For this study we focussed on just the four MSSFR parameters used in the post-processing of COMPAS simulations ($(\alpha, \lambda, a_{SF}, d_{SF})$ - see Neijssel et al. [48] for details). We found that our first attempt at analysis on LVK data pushed our inference on two of the parameters (α and a_{SF}) to rail against the prior boundaries. This was probably due to poor choices in the fiducial model describing astrophysical evolution, exacerbated by an overly simplistic observational selection function, and in the case of a_{SF} , also likely as a result of the over-production of BBHs by our CHE implementation in COMPAS. We were able to partially remedy this issue because a_{SF} required only a simple rescaling.

However, to avoid similar problems in future analyses, it would be wise to use active learning and grow the training data set in regions of parameter space which provide models with the best match to the data.

We are not limited to the four MSSFR parameters we chose to focus on for this study, or indeed to the COMPAS post-processing code. The next step is to expand this work to other astrophysical parameters that govern stellar and binary evolution in the population synthesis code itself (e.g. COMPAS), rather than just the post-processing code, thereby reducing the need for time-consuming simulations.

Bibliography

- [1] Jeff Riley, Poojan Agrawal, Jim W. Barrett, Kristan N. K. Boyett, Floor S. Broekgaarden, Debatri Chattopadhyay, Sebastian M. Gaebel, Fabian Gittins, Ryosuke Hirai, George Howitt, Stephen Justham, Lokesh Khandelwal, Floris Kummer, Mike Y. M. Lau, Ilya Mandel, Selma E. de Mink, Coenraad Neijssel, Tim Riley, Lieke van Son, Simon Stevenson, Alejandro Vigna-Gómez, Serena Vinciguerra, Tom Wagg, Reinhold Willcox, and Team Compas. Rapid Stellar and Binary Population Synthesis with COMPAS. *ApJS*, 258(2):34, February 2022. doi: 10.3847/1538-4365/ac416c.
- [2] Ilya Mandel and Alison Farmer. Merging stellar-mass binary black holes. *Physics Reports*, 955:1–24, 2022.
- [3] Michela Mapelli. *Astrophysics of stellar black holes*, 2018.
- [4] A. Tutukov and L. Yungelson. Evolution of massive close binaries. *Nauchnye Informatsii*, 27:70, 1973.
- [5] A. V. Tutukov and L. R. Yungelson. The merger rate of neutron star and black hole binaries. *Monthly Notices of the Royal Astronomical Society*, 260:675–678, February 1993. doi: 10.1093/mnras/260.3.675.
- [6] E. P. J. van den Heuvel. Late Stages of Close Binary Systems. In P. Eggleton, S. Mitton, and J. Whelan, editors, *Structure and Evolution of Close Binary Systems*, volume 73 of *IAU Symposium*, page 35, 1976.
- [7] Steinn Sigurdsson and Lars Hernquist. Primordial black holes in globular clusters. *Nature*, 364:423–425, 1993. doi: 10.1038/364423a0.
- [8] M. Coleman Miller and Vanessa M. Lauburg. Mergers of stellar-mass black holes in nuclear star clusters. *The Astrophysical Journal*, 692(1):7, Feb 2009. ISSN 1538-4357. doi: 10.1088/0004-637X/692/1/917. URL <https://doi.org/10.1088/0004-637X/692/1/917>.

- [9] B. M. Ziosi, M. Mapelli, M. Branchesi, and G. Tormen. Dynamics of stellar black holes in young star clusters with different metallicities - II. Black hole-black hole binaries. *Monthly Notices of the Royal Astronomical Society*, 441:3703–3717, July 2014. doi: 10.1093/mnras/stu824.
- [10] Carl L. Rodriguez, Meagan Morscher, Bharath Pattabiraman, Sourav Chatterjee, Carl-Johan Haster, and Frederic A. Rasio. Binary black hole mergers from globular clusters: Implications for advanced ligo. *Phys. Rev. Lett.*, 115:051101, Jul 2015. doi: 10.1103/PhysRevLett.115.051101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.115.051101>.
- [11] Fabio Antonini, Sourav Chatterjee, Carl L. Rodriguez, Meagan Morscher, Bharath Pattabiraman, Vicky Kalogera, and Frederic A. Rasio. Black hole mergers and blue stragglers from hierarchical triples formed in globular clusters. *The Astrophysical Journal*, 816(2):65, Jan 2016. ISSN 1538-4357. doi: 10.3847/0004-637x/816/2/65. URL <http://dx.doi.org/10.3847/0004-637X/816/2/65>.
- [12] I. Bartos, B. Kocsis, Z. Haiman, and S. Márka. Rapid and Bright Stellar-mass Binary Black Hole Mergers in Active Galactic Nuclei. *The Astrophysical Journal*, 835:165, February 2017. doi: 10.3847/1538-4357/835/2/165.
- [13] N. C. Stone, B. D. Metzger, and Z. Haiman. Assisted inspirals of stellar mass black holes embedded in AGN discs: solving the ‘final au problem’. *Monthly Notices of the Royal Astronomical Society*, 464:946–954, January 2017. doi: 10.1093/mnras/stw2260.
- [14] Philip Carl Peters. Gravitational radiation and the motion of two point masses. *Physical Review*, 136(4B):B1224, 1964.
- [15] Peter Eggleton. *Evolutionary Processes in Binary and Multiple Stars*. Cambridge University Press, 2006.
- [16] Jarrod R. Hurley, Christopher A. Tout, and Onno R. Pols. Evolution of binary stars and the effect of tides on binary populations. *MNRAS*, 329(4):897–928, Feb 2002. doi: 10.1046/j.1365-8711.2002.05038.x.
- [17] Simon F. Portegies Zwart and Stephen L. W. McMillan. Black hole mergers in the universe. *The Astrophysical Journal*, 528(1):L17, dec 1999. doi: 10.1086/312422. URL <https://dx.doi.org/10.1086/312422>.
- [18] Simon F. Portegies Zwart, Stephen L.W. McMillan, and Mark Gieles. Young massive star clusters. *Annual Review of Astronomy and Astrophysics*, 48(1):431–493, 2010. doi: 10.1146/annurev-astro-081309-130834. URL <https://doi.org/10.1146/annurev-astro-081309-130834>.

- [19] Ugo N Di Carlo, Nicola Giacobbo, Michela Mapelli, Mario Pasquato, Mario Spera, Long Wang, and Francesco Haardt. Merging black holes in young star clusters. *Monthly Notices of the Royal Astronomical Society*, 487(2):2947–2960, 05 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz1453. URL <https://doi.org/10.1093/mnras/stz1453>.
- [20] J. Michell. On the Means of Discovering the Distance, Magnitude, &c. of the Fixed Stars, in Consequence of the Diminution of the Velocity of Their Light, in Case Such a Diminution Should be Found to Take Place in any of Them, and Such Other Data Should be Procured from Observations, as Would be Farther Necessary for That Purpose. By the Rev. John Michell, B. D. F. R. S. In a Letter to Henry Cavendish, Esq. F. R. S. and A. S., January 1784. URL <https://doi.org/10.1098/rstl.1784.0008>.
- [21] P.S. de Laplace. *Exposition du syst eme du monde*,. Exposition du syst eme du monde. De l’Imprimerie du Cercle-Social, 1796.
- [22] Pierre Simon Laplace. Beweis des Satzes, dass die anziehende Kraft bey einem Weltk orper so gro s seyn k onne, dass das Licht davon nicht ausstr omen kann. *Allgemeine Geographische Ephemeriden*, 4(1):1–6, July 1799.
- [23] Karl Schwarzschild.  ber das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie. *Sitzungsberichte der K oniglich Preussischen Akademie der Wissenschaften*, pages 189–196, January 1916.
- [24] Karl Schwarzschild.  ber das Gravitationsfeld einer Kugel aus inkompressibler Fl ussigkeit nach der Einsteinschen Theorie. In *Sitzungsberichte der K oniglich Preussischen Akademie der Wissenschaften zu Berlin*, pages 424–434, March 1916.
- [25] Albert Einstein. Die Feldgleichungen der Gravitation. *Sitzungsberichte der K oniglich Preussischen Akademie der Wissenschaften*, pages 844–847, January 1915.
- [26] Abbott, B. P. et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016. doi: 10.1103/PhysRevLett.116.061102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>.
- [27] The LIGO Scientific Collaboration and Virgo Collaboration, et al. Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs. *Phys. Rev. X*, 9:031040, Sep 2019. doi: 10.1103/PhysRevX.9.031040. URL <https://link.aps.org/doi/10.1103/PhysRevX.9.031040>.

- [28] The LIGO Scientific Collaboration and Virgo Collaboration, et al. Gwtc-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *Phys. Rev. X*, 11:021053, Jun 2021. doi: 10.1103/PhysRevX.11.021053. URL <https://link.aps.org/doi/10.1103/PhysRevX.11.021053>.
- [29] The LIGO Scientific Collaboration and Virgo Collaboration, et al. GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. *arXiv e-prints*, art. arXiv:2108.01045, August 2021. doi: 10.48550/arXiv.2108.01045.
- [30] The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. *arXiv e-prints*, art. arXiv:2111.03606, November 2021. doi: 10.48550/arXiv.2111.03606.
- [31] Abbott, B. P. et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Reviews in Relativity*, 23(1):3, September 2020. doi: 10.1007/s41114-020-00026-9.
- [32] R X Adhikari et al. A cryogenic silicon interferometer for gravitational-wave detection. *Classical and Quantum Gravity*, 37(16):165003, jul 2020. doi: 10.1088/1361-6382/ab9143. URL <https://dx.doi.org/10.1088/1361-6382/ab9143>.
- [33] M Punturo et al. The einstein telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity*, 27(19):194002, sep 2010. doi: 10.1088/0264-9381/27/19/194002. URL <https://dx.doi.org/10.1088/0264-9381/27/19/194002>.
- [34] S Hild et al. Sensitivity studies for third-generation gravitational wave observatories. *Classical and Quantum Gravity*, 28(9):094013, apr 2011. doi: 10.1088/0264-9381/28/9/094013. URL <https://dx.doi.org/10.1088/0264-9381/28/9/094013>.
- [35] B P Abbott et al. Exploring the sensitivity of next generation gravitational wave detectors. *Classical and Quantum Gravity*, 34(4):044001, jan 2017. doi: 10.1088/1361-6382/aa51f4. URL <https://dx.doi.org/10.1088/1361-6382/aa51f4>.
- [36] David Reitze, Rana X Adhikari, Stefan Ballmer, Barry Barish, Lisa Barsotti, GariLynn Billingsley, Duncan A. Brown, Yanbei Chen, Dennis Coyne, Robert Eisenstein, Matthew Evans, Peter Fritschel, Evan D. Hall, Albert Lazzarini, Geoffrey Lovelace, Jocelyn Read, B. S. Sathyaprakash, David Shoemaker, Joshua Smith, Calum Torrie, Salvatore Vitale, Rainer Weiss, Christopher Wipf, and

- Michael Zucker. Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO. *Bulletin of the AAS*, 51(7), sep 30 2019. <https://baas.aas.org/pub/2020n7i035>.
- [37] B. P. et al. Abbott. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett*, 119(16):161101, October 2017. doi: 10.1103/PhysRevLett.119.161101.
- [38] A. Goldstein, P. Veres, E. Burns, M. S. Briggs, R. Hamburg, D. Kocevski, C. A. Wilson-Hodge, R. D. Preece, S. Poolakkil, O. J. Roberts, C. M. Hui, V. Connaughton, J. Racusin, A. von Kienlin, T. Dal Canton, N. Christensen, T. Littenberg, K. Siellez, L. Blackburn, J. Broida, E. Bissaldi, W. H. Cleveland, M. H. Gibby, M. M. Giles, R. M. Kippen, S. McBreen, J. McEnery, C. A. Meegan, W. S. Paciesas, and M. Stanbro. An ordinary short gamma-ray burst with extraordinary implications: Fermi-gbm detection of grb 170817a. *The Astrophysical Journal Letters*, 848(2):L14, oct 2017. doi: 10.3847/2041-8213/aa8f41. URL <https://dx.doi.org/10.3847/2041-8213/aa8f41>.
- [39] B. P. et al. Abbott. Multi-messenger Observations of a Binary Neutron Star Merger. *The Astrophysical Journal Letters*, 848(2):L12, October 2017. doi: 10.3847/2041-8213/aa91c9.
- [40] Riccardo Ciolfi. Binary neutron star mergers after GW170817. *Frontiers in Astronomy and Space Sciences*, 7:27, June 2020. doi: 10.3389/fspas.2020.00027.
- [41] Geoffrey Mo, Rahul Jayaraman, Michael Fausnaugh, Erik Katsavounidis, George R. Ricker, and Roland Vanderspek. Searching for gravitational-wave counterparts using the transiting exoplanet survey satellite. *The Astrophysical Journal Letters*, 948(1):L3, apr 2023. doi: 10.3847/2041-8213/acca70. URL <https://dx.doi.org/10.3847/2041-8213/acca70>.
- [42] ISO. *ISO/IEC 14882:2011 Information technology — Programming languages — C++*. International Organization for Standardization, Geneva, Switzerland, February 2012. URL http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=50372.
- [43] Simon Stevenson, Alejandro Vigna-Gómez, Ilya Mandel, Jim W. Barrett, Coenraad J. Neijssel, David Perkins, and Selma E. de Mink. Formation of the first three gravitational-wave observations through isolated binary evolution. *Nature Communications*, 8:14906, April 2017. doi: 10.1038/ncomms14906.
- [44] Jim W. Barrett, Sebastian M. Gaebel, Coenraad J. Neijssel, Alejandro Vigna-Gomez, Simon Stevenson, Christopher P. L. Berry, Will M. Farr, and Ilya Mandel.

- Accuracy of inference on the physics of binary evolution from gravitational-wave observations. *Mon. Not. Roy. Astron. Soc.*, 477(4):4685–4695, 2018. doi: 10.1093/mnras/sty908.
- [45] Alejandro Vigna-Gómez, Coenraad J. Neijssel, Simon Stevenson, Jim W. Barrett, Krzysztof Belczynski, Stephen Justham, Selma E. de Mink, Bernhard Müller, Philipp Podsiadlowski, Mathieu Renzo, Dorottya Szécsi, and Ilya Mandel. On the formation history of Galactic double neutron stars. *MNRAS*, 481:4009–4029, December 2018. doi: 10.1093/mnras/sty2463.
- [46] Simon Stevenson, Matthew Sampson, Jade Powell, Alejandro Vigna-Gómez, Coenraad J. Neijssel, Dorottya Szécsi, and Ilya Mandel. The Impact of Pair-instability Mass Loss on the Binary Black Hole Mass Distribution. *ApJ*, 882(2):121, September 2019. doi: 10.3847/1538-4357/ab3981.
- [47] Floor S. Broekgaarden, Stephen Justham, Selma E. de Mink, Jonathan Gair, Ilya Mandel, Simon Stevenson, Jim W. Barrett, Alejandro Vigna-Gómez, and Coenraad J. Neijssel. STROOPWAFEL: simulating rare outcomes from astrophysical populations, with application to gravitational-wave sources. *MNRAS*, 490(4):5228–5248, Dec 2019. doi: 10.1093/mnras/stz2558.
- [48] Coenraad J. Neijssel, Alejandro Vigna-Gómez, Simon Stevenson, Jim W. Barrett, Sebastian M. Gaebel, Floor S. Broekgaarden, Selma E. de Mink, Dorottya Szécsi, Serena Vinciguerra, and Ilya Mandel. The effect of the metallicity-specific star formation history on double compact object mergers. *MNRAS*, 490(3):3740–3759, Dec 2019. doi: 10.1093/mnras/stz2840.
- [49] Debatri Chattopadhyay, Simon Stevenson, Jarrod R. Hurley, Luca J. Rossi, and Chris Flynn. Modelling Double Neutron Stars: Radio and Gravitational Waves. *Mon. Not. Roy. Astron. Soc.*, 494(2):1587–1610, 2020. doi: 10.1093/mnras/staa756.
- [50] Jarrod R. Hurley, Christopher A. Tout, Sverre J. Aarseth, and Onno R. Pols. Direct n-body modeling of stellar populations: blue stragglers in m67. *Mon. Not. Roy. Astron. Soc.*, 323:630, 2001. doi: 10.1046/j.1365-8711.2001.04220.x.
- [51] Krzysztof Belczynski, Vassiliki Kalogera, and Tomasz Bulik. A Comprehensive Study of Binary Compact Objects as Gravitational Wave Sources: Evolutionary Channels, Rates, and Physical Properties. *ApJ*, 572(1):407–431, Jun 2002. doi: 10.1086/340304.
- [52] Krzysztof Belczynski, Vassiliki Kalogera, Frederic A. Rasio, Ronald E. Taam, Andreas Zezas, Tomasz Bulik, Thomas J. Maccarone, and Natalia Ivanova. Compact

- Object Modeling with the StarTrack Population Synthesis Code. *ApJS*, 174(1): 223–260, Jan 2008. doi: 10.1086/521026.
- [53] Bill Paxton, Pablo Marchant, Josiah Schwab, Evan B. Bauer, Lars Bildsten, Matteo Cantiello, Luc Dessart, R. Farmer, H. Hu, N. Langer, R. H. D. Townsend, Dean M. Townsley, and F. X. Timmes. MODULES FOR EXPERIMENTS IN STELLAR ASTROPHYSICS (MESA): BINARIES, PULSATIONS, AND EXPLOSIONS. *The Astrophysical Journal Supplement Series*, 220(1):15, September 2015. doi: 10.1088/0067-0049/220/1/15. URL <https://doi.org/10.1088/2F0067-0049%2F220%2F1%2F15>.
- [54] Bill Paxton, Matteo Cantiello, Phil Arras, Lars Bildsten, Edward F. Brown, Aaron Dotter, Christopher Mankovich, M. H. Montgomery, Dennis Stello, F. X. Timmes, and Richard Townsend. Modules for Experiments in Stellar Astrophysics (MESA): Planets, Oscillations, Rotation, and Massive Stars. *ApJS*, 208(1):4, September 2013. doi: 10.1088/0067-0049/208/1/4.
- [55] Bill Paxton, Lars Bildsten, Aaron Dotter, Falk Herwig, Pierre Lesaffre, and Frank Timmes. Modules for Experiments in Stellar Astrophysics (MESA). *ApJS*, 192(1): 3, January 2011. doi: 10.1088/0067-0049/192/1/3.
- [56] R. G. Izzard, C. A. Tout, A. I. Karakas, and O. R. Pols. A new synthetic model for asymptotic giant branch stars. *MNRAS*, 350:407–426, May 2004. doi: 10.1111/j.1365-2966.2004.07446.x.
- [57] R. G. Izzard, L. M. Dray, A. I. Karakas, M. Lugaro, and C. A. Tout. Population nucleosynthesis in single and binary stars. I. Model. *A&A*, 460(2):565–572, Dec 2006. doi: 10.1051/0004-6361:20066129.
- [58] R. G. Izzard, E. Glebbeek, R. J. Stancliffe, and O. R. Pols. Population synthesis of binary carbon-enhanced metal-poor stars. *A&A*, 508(3):1359–1374, Dec 2009. doi: 10.1051/0004-6361/200912827.
- [59] V. M. Lipunov, K. A. Postnov, and M. E. Prokhorov. The Scenario Machine: restrictions on key parameters of binary evolution. *A&A*, 310:489–507, Jun 1996.
- [60] V. M. Lipunov, K. A. Postnov, M. E. Prokhorov, and A. I. Bogomazov. Description of the “Scenario Machine”. *Astronomy Reports*, 53(10):915–940, Oct 2009. doi: 10.1134/S1063772909100047.
- [61] S. F. Portegies Zwart and F. Verbunt. Population synthesis of high-mass binaries. *A&A*, 309:179–196, May 1996.

- [62] S. Toonen, G. Nelemans, and S. Portegies Zwart. Supernova Type Ia progenitors from merging double white dwarfs. Using a new population synthesis model. *A&A*, 546:A70, Oct 2012. doi: 10.1051/0004-6361/201218966.
- [63] Katelyn Breivik, Scott Coughlin, Michael Zevin, Carl L. Rodriguez, Kyle Kremer, Claire S. Ye, Jeff J. Andrews, Michael Kurkowski, Matthew C. Digman, Shane L. Larson, and Frederic A. Rasio. COSMIC Variance in Binary Population Synthesis. *ApJ*, 898(1):71, July 2020. doi: 10.3847/1538-4357/ab9d85.
- [64] Nicola Giacobbo, Michela Mapelli, and Mario Spera. Merging black hole binaries: the effects of progenitor’s metallicity, mass-loss rate and Eddington factor. *MNRAS*, 474(3):2959–2974, March 2018. doi: 10.1093/mnras/stx2933.
- [65] E. R. Stanway and J. J. Eldridge. Re-evaluating old stellar populations. *MNRAS*, 479(1):75–93, September 2018. doi: 10.1093/mnras/sty1353.
- [66] LISA Science Group. Lisa synthetic ucb catalogs project, 2023. URL https://lisa.pages.in2p3.fr/consortium-userguide/g_lsg.html.
- [67] O. R. Pols, K.-P. Schröder, J. R. Hurley, C. A. Tout, and P. P. Eggleton. Stellar evolution models for $Z = 0.0001$ to 0.03. *MNRAS*, 298:525–536, August 1998. doi: 10.1046/j.1365-8711.1998.01658.x.
- [68] Matthias U. Kruckow, Thomas M. Tauris, Norbert Langer, Michael Kramer, and Robert G. Izzard. Progenitors of gravitational wave mergers: binary evolution with the stellar grid-based code COMBINE. *MNRAS*, 481(2):1908–1949, Dec 2018. doi: 10.1093/mnras/sty2190.
- [69] Poojan Agrawal, Jarrod Hurley, Simon Stevenson, Dorottya Szécsi, and Chris Flynn. The fates of massive stars: exploring uncertainties in stellar evolution with metisse. *Monthly Notices of the Royal Astronomical Society*, 497(4): 4549–4564, 08 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa2264. URL <https://doi.org/10.1093/mnras/staa2264>.
- [70] Poojan Agrawal, Jarrod Hurley, Simon Stevenson, Carl L. Rodriguez, Dorottya Szecsi, and Alex Kemp. Modelling stellar evolution in mass-transferring binaries and gravitational-wave progenitors with metisse, 2023. URL <https://arxiv.org/abs/2303.10187>.
- [71] Mario Spera, Michela Mapelli, Nicola Giacobbo, Alessandro A Trani, Alessandro Bressan, and Guglielmo Costa. Merging black hole binaries with the SEVN code. *MNRAS*, 485(1):889–907, 02 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz359. URL <https://doi.org/10.1093/mnras/stz359>.

- [72] Tassos Fragos, Jeff J. Andrews, Simone S. Bavera, Christopher P. L. Berry, Scott Coughlin, Aaron Dotter, Prabin Giri, Vicky Kalogera, Aggelos Katsaggelos, Konstantinos Kovelakas, Shamal Lalvani, Devina Misra, Philipp M. Srivastava, Ying Qin, Kyle A. Rocha, Jaime Román-Garza, Juan Gabriel Serra, Petter Stahle, Meng Sun, Xu Teng, Goce Trajceviski, Nam Hai Tran, Zepei Xing, Emmanouil Zapartas, and Michael Zevin. Posydon: A general-purpose population synthesis code with detailed binary-evolution simulations. *The Astrophysical Journal Supplement Series*, 264(2):45, feb 2023. doi: 10.3847/1538-4365/ac90c1. URL <https://dx.doi.org/10.3847/1538-4365/ac90c1>.
- [73] ARC Centre of Excellence for Gravitational Wave Discovery. Unravelling the mystery of black holes: Scientists release stellar code to the public, 3 2020. URL <https://phys.org/news/2020-03-unravelling-mystery-black-holes-scientists.html>.
- [74] Jeff Riley, Poojan Agrawal, Jim W. Barrett, Kristan N. K. Boyett, Floor S. Broekgaarden, Debatri Chattopadhyay, Sebastian M. Gaebel, Fabian Gittins, Ryosuke Hirai, George Howitt, Stephen Justham, Lokesh Khandelwal, Floris Kummer, Mike Y. M. Lau, Ilya Mandel, Selma E. de Mink, Coenraad Neijssel, Tim Riley, Lieke van Son, Simon Stevenson, Alejandro Vigna-Gómez, Serena Vinciguerra, Tom Wagg, Reinhold Willcox, and Team Compas. Compas: A rapid binary population synthesis suite. *Journal of Open Source Software*, 7(69):3838, 2022. doi: 10.21105/joss.03838. URL <https://doi.org/10.21105/joss.03838>.
- [75] github. Github, 2020. URL <https://github.com/>.
- [76] André Maeder and Georges Meynet. The evolution of rotating stars. *Annual Review of Astronomy and Astrophysics*, 38(1):143–190, Sep 2000. ISSN 1545-4282. doi: 10.1146/annurev.astro.38.1.143. URL <http://dx.doi.org/10.1146/annurev.astro.38.1.143>.
- [77] A. Heger, N. Langer, and S. E. Woosley. Presupernova Evolution of Rotating Massive Stars. I. Numerical Method and Evolution of the Internal Stellar Structure. *ApJ*, 528(1):368–396, January 2000. doi: 10.1086/308158.
- [78] A. Maeder. Evidences for a bifurcation in massive star evolution. The ON-blue stragglers. *A&A*, 178:159–169, May 1987.
- [79] N. Langer. Helium enrichment in massive early type stars. *A&A*, 265:L17–L20, November 1992.
- [80] Ilya Mandel and Selma E. de Mink. Merging binary black holes formed through chemically homogeneous evolution in short-period stellar binaries. *Monthly Notices*

- of the Royal Astronomical Society*, 458(3):2634–2647, Feb 2016. ISSN 1365-2966. doi: 10.1093/mnras/stw379. URL <http://dx.doi.org/10.1093/mnras/stw379>.
- [81] S. E. de Mink, M. Cantiello, N. Langer, O. R. Pols, I. Brott, and S. Ch. Yoon. Rotational mixing in massive binaries. Detached short-period systems. *A&A*, 497(1):243–253, April 2009. doi: 10.1051/0004-6361/200811439.
- [82] L. A. Almeida, H. Sana, S. E. de Mink, F. Tramper, I. Soszyn ´ski, N. Langer, R. H. Barbá, M. Cantiello, A. Daminieli, A. de Koter, and et al. Discovery of the massive overcontact binary vfts 352: Evidence for enhanced internal mixing. *The Astrophysical Journal*, 812(2):102, Oct 2015. ISSN 1538-4357. doi: 10.1088/0004-637x/812/2/102. URL <http://dx.doi.org/10.1088/0004-637x/812/2/102>.
- [83] Gloria Koenigsberger, Nidia Morrell, D. John Hillier, Roberto Gamen, Fabian R. N. Schneider, Nicolás González-Jiménez, Norbert Langer, and Rodolfo Barbá. The hd 5980 multiple system: Masses and evolutionary status. *The Astronomical Journal*, 148(4):62, Sep 2014. ISSN 1538-3881. doi: 10.1088/0004-6256/148/4/62. URL <http://dx.doi.org/10.1088/0004-6256/148/4/62>.
- [84] S. E. de Mink and I. Mandel. The chemically homogeneous evolutionary channel for binary black hole mergers: rates and properties of gravitational-wave events detectable by advanced ligo. *Monthly Notices of the Royal Astronomical Society*, 460(4):3545–3553, May 2016. ISSN 1365-2966. doi: 10.1093/mnras/stw1219. URL <http://dx.doi.org/10.1093/mnras/stw1219>.
- [85] Pablo Marchant, Norbert Langer, Philipp Podsiadlowski, Thomas M. Tauris, and Takashi J. Moriya. A new route towards merging massive black holes. *Astron. Astrophys.*, 588:A50, 2016. doi: 10.1051/0004-6361/201628133.
- [86] L du Buisson, P Marchant, Ph Podsiadlowski, C Kobayashi, F B Abdalla, P Taylor, I Mandel, S E de Mink, T J Moriya, and N Langer. Cosmic rates of black hole mergers and pair-instability supernovae from chemically homogeneous binary evolution. *Monthly Notices of the Royal Astronomical Society*, 499(4):5941–5959, 10 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa3225. URL <https://doi.org/10.1093/mnras/staa3225>.
- [87] Jeff Riley, Ilya Mandel, Pablo Marchant, Ellen Butler, Kaila Nathaniel, Coenraad Neijssel, Spencer Shortt, and Alejandro Vigna-Gomez. Chemically homogeneous evolution: a rapid population synthesis approach. *Mon. Not. Roy. Astron. Soc.*, 505(1):663–676, 2021. doi: 10.1093/mnras/stab1291.

- [88] Simon Stevenson and Blaze Houlden. An improved population synthesis model for binary black hole formation via chemically homogeneous evolution. In preparation, 2023.
- [89] Jeff Riley and Ilya Mandel. Surrogate forward models for population inference on compact binary mergers, 2023. URL <https://doi.org/10.48550/arXiv.2303.00508>.
- [90] LAC van Son, SE de Mink, T Callister, S Justham, M Renzo, T Wagg, FS Broekgaarden, F Kummer, R Pakmor, and I Mandel. The redshift evolution of the binary black hole merger rate: a weighty matter. *The Astrophysical Journal*, 931(1):17, 2022.
- [91] Floor S Broekgaarden, Edo Berger, Simon Stevenson, Stephen Justham, Ilya Mandel, Martyna Chruślińska, Lieke AC van Son, Tom Wagg, Alejandro Vigna-Gómez, Selma E de Mink, et al. Impact of massive binary star and cosmic evolution on gravitational wave observations—ii. double compact object rates and properties. *Monthly Notices of the Royal Astronomical Society*, 516(4):5737–5761, 2022.
- [92] LAC van Son, SE de Mink, M Renzo, S Justham, E Zapartas, K Breivik, T Callister, WM Farr, and C Conroy. No peaks without valleys: The stable mass transfer channel for gravitational-wave sources in light of the neutron star–black hole mass gap. *The Astrophysical Journal*, 940(2):184, 2022.
- [93] Simon Stevenson, Reinhold Willcox, Alejandro Vigna-Gómez, and Floor Broekgaarden. Wide binary pulsars from electron-capture supernovae. *Monthly Notices of the Royal Astronomical Society*, 513(4):6105–6110, 05 2022. ISSN 0035-8711. doi: 10.1093/mnras/stac1322. URL <https://doi.org/10.1093/mnras/stac1322>.
- [94] Tom Wagg, Floor S Broekgaarden, Selma E de Mink, N Frankel, LAC van Son, and Stephen Justham. Gravitational wave sources in our galactic backyard: Predictions for bbbh, bhns, and nsns binaries detectable with lisa. *The Astrophysical Journal*, 937(2):118, 2022.
- [95] Pau Amaro-Seoane, Heather Audley, Stanislav Babak, John Baker, Enrico Barausse, Peter Bender, Emanuele Berti, Pierre Binétruy, Michael Born, Daniele Bortoluzzi, Jordan Camp, Chiara Caprini, Vitor Cardoso, Monica Colpi, John Conklin, Neil Cornish, Curt Cutler, Karsten Danzmann, Rita Dolesi, Luigi Ferraioli, Valerio Ferroni, Ewan Fitzsimons, Jonathan Gair, Lluís Gesa Bote, Domenico Giardini, Ferran Gibert, Catia Grimani, Hubert Halloin, Gerhard Heinzl, Thomas Hertog, Martin Hewitson, Kelly Holley-Bockelmann, Daniel Hollington, Mauro Hueller, Henri Inchauspe, Philippe Jetzer, Nikos Karnesis, Christian Killow, Antoine Klein, Bill Klipstein, Natalia Korsakova, Shane L Larson, Jeffrey Livas, Ivan

- Lloro, Nary Man, Davor Mance, Joseph Martino, Ignacio Mateos, Kirk McKenzie, Sean T McWilliams, Cole Miller, Guido Mueller, Germano Nardini, Gijs Nelemans, Miquel Nofrarias, Antoine Petiteau, Paolo Pivato, Eric Plagnol, Ed Porter, Jens Reiche, David Robertson, Norna Robertson, Elena Rossi, Giuliana Russano, Bernard Schutz, Alberto Sesana, David Shoemaker, Jacob Slutsky, Carlos F. Sopuerta, Tim Sumner, Nicola Tamanini, Ira Thorpe, Michael Troebs, Michele Vallisneri, Alberto Vecchio, Daniele Vetrugno, Stefano Vitale, Marta Volonteri, Gudrun Wanner, Harry Ward, Peter Wass, William Weber, John Ziemer, and Peter Zweifel. Laser interferometer space antenna, 2017.
- [96] Monica Colpi, Kelly Holley-Bockelmann, Tamara Bogdanovic, Priya Natarajan, Jillian Bellovary, Alberto Sesana, Michael Tremmel, Jeremy Schnittman, Julia Comerford, Enrico Barausse, Emanuele Berti, Marta Volonteri, Fazeel Khan, Sean McWilliams, Sarah Burke-Spolaor, Jeff Hazboun, John Conklin, Guido Mueller, and Shane Larson. Astro2020 science white paper: The gravitational wave view of massive black holes, 2019.
- [97] Floor S Broekgaarden, Simon Stevenson, and Eric Thrane. Signatures of mass ratio reversal in gravitational waves from merging binary black holes. *The Astrophysical Journal*, 938(1):45, 2022.
- [98] Daniel J D’Orazio, Zoltan Haiman, Janna Levin, Johan Samsing, and Alejandro Vigna-Gómez. Multimessenger constraints on magnetic fields in merging black hole–neutron star binaries. *The Astrophysical Journal*, 927(1):56, 2022.
- [99] Ondřej Pejcha, Pavel Cagaš, Camille Landri, Michael M Fausnaugh, Gisella De Rosa, Jose L Prieto, Zbyněk Henzl, and Milan Pešta. The complex dynamical past and future of double eclipsing binary czev343: Misaligned orbits and period resonance. *Astronomy & Astrophysics*, 667:A53, 2022.
- [100] Yael Raveh, Erez Michaely, and Hagai B Perets. Detailed properties of gravitational-wave mergers from flyby perturbations of wide binary black holes in the field. *Monthly Notices of the Royal Astronomical Society*, 514(3):4246–4258, 2022.
- [101] Ricardo Yarza, Rosa Wallace Everson, and Enrico Ramirez-Ruiz. A revised energy formalism for common-envelope evolution: repercussions for planetary engulfment and the formation of neutron star binaries. *arXiv preprint arXiv:2210.00010*, 2022.
- [102] LAC van Son, SE de Mink, M Chruślińska, C Conroy, R Pakmor, and L Hernquist. The locations of features in the mass distribution of merging binary black holes are robust against uncertainties in the metallicity-dependent cosmic star formation history. *The Astrophysical Journal*, 948(2):105, 2023.

- [103] Enrique Moreno Méndez, Fabio De Colle, Diego López-Cámara, and Alejandro Vigna-Gómez. Hypercritical accretion during common envelopes in triples leading to binary black holes in the pair-instability-supernova mass gap. *Monthly Notices of the Royal Astronomical Society*, 522(2):1686–1696, 2023.
- [104] Ish Gupta, Ssohrab Borhanian, Arnab Dhani, Debatri Chattopadhyay, Rahul Kashyap, V Ashley Villar, and BS Sathyaprakash. Neutron star-black hole mergers in next generation gravitational-wave observatories. *Physical Review D*, 107(12):124007, 2023.
- [105] Liana Rauf, Cullan Howlett, Tamara M Davis, and Claudia DP Lagos. Exploring binary black hole mergers and host galaxies with shark and compas. *Monthly Notices of the Royal Astronomical Society*, page stad1757, 2023.
- [106] Claudia del P Lagos, Rodrigo J Tobar, Aaron S G Robotham, Danail Obreschkow, Peter D Mitchell, Chris Power, and Pascal J Elahi. Shark: introducing an open source, free, and flexible semi-analytic model of galaxy formation. *Monthly Notices of the Royal Astronomical Society*, 481(3):3573–3603, 09 2018. ISSN 0035-8711. doi: 10.1093/mnras/sty2440. URL <https://doi.org/10.1093/mnras/sty2440>.
- [107] Aldana Grichener. Mergers of neutron stars and black holes with cores of giant stars: a population synthesis study. *Monthly Notices of the Royal Astronomical Society*, 523(1):221–232, 2023.
- [108] Alejandro Vigna-Gómez and Enrico Ramirez-Ruiz. A binary origin for the first isolated stellar-mass black hole detected with astrometric microlensing. *The Astrophysical Journal Letters*, 946(1):L2, 2023.
- [109] I Agudo, L Amati, T An, FE Bauer, S Benetti, MG Bernardini, R Beswick, K Bhimbhakdi, T de Boer, M Branchesi, et al. Panning for gold, but finding helium: Discovery of the ultra-stripped supernova sn 2019wxt from gravitational-wave follow-up observations. *Astronomy & Astrophysics*, 675:A201, 2023.
- [110] Lucas M de Sá, Antônio Bernardo, Riis RA Bachega, Livia S Rocha, and Jorge E Horvath. Effects of a non-universal imf and binary parameter correlations on compact binary mergers. *Astronomische Nachrichten*, 344(1-2):e220089, 2023.
- [111] Markus R. Mosbech, Alexander C. Jenkins, Sownak Bose, Celine Boehm, Mairi Sakellariadou, and Yvonne Y. Y. Wong. Gravitational-wave event rates as a new probe for dark matter microphysics. *Phys. Rev. D*, 108:043512, Aug 2023. doi: 10.1103/PhysRevD.108.043512. URL <https://link.aps.org/doi/10.1103/PhysRevD.108.043512>.