The Architecture of the LISA Science Analysis

study leads Saul Teukolsky Michele Vallisneri



The Architecture of the LISA Science Analysis

A W. M. Keck Institute for Space Studies workshop, January 16-19, 2018

This work had its inception at the Architecture of the LISA Science Analysis study funded by the W. M. Keck Institute for Space Studies. The study organizers and all study participants are deeply grateful to KISS Executive Director Michele Judd and to her amazing staff for creating the perfect environment for insight, creativity, and collegiality. Deirdre Shoemaker acknowledges support from NSF award PHY1806580. Michele Vallisneri, Alvin Chua, Curt Cutler, and Chad Galley acknowledge support from the JPL RTD program. Part of this work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Government sponsorship acknowledged.

Editing and Formatting: Meg Rosenburg

Cover Image: Chuck Carter/Keck Institute for Space Studies (KISS)
Header images: SXS project (http://www.black-holes.org); R. Hurt, Caltech-JPL; AEI/MM/exozet;
ESO/L. Calçada/spaceengine.org; NASA/ESA; iStock MF3d; NASA.
© September 2019. All rights reserved.

Authors and Study Participants

Rachel Akeson California Institute of Technology

Anne Archibald Universiteit van Amsterdam

Stanislav Babak Université Paris Diderot

Katelyn Breivik Canadian Institute for Theoretical Astrophysics

C. Titus Brown University of California, Davis

Alvin Chua (scientific secretary) Jet Propulsion Laboratory

Neil Cornish Montana State University

Curt Cutler Jet Propulsion Laboratory

Scott Davidoff Jet Propulsion Laboratory

Francois Foucart University of New Hampshire

Chad Galley Jet Propulsion Laboratory

Lawrence Kidder Cornell University

Prayush Kumar Cornell University

Astrid Lamberts California Institute of Technology Geoffrey Lovelace California State University, Fullerton

Ashish Mahabal California Institute of Technology

Christine Corbett Moran California Institute of Technology

Laura Nuttall University of Portsmouth

Maria Okounkova California Institute of Technology

Travis Robson Montana State University

Mark Scheel California Institute of Technology

Deirdre Shoemaker Georgia Institute of Technology

Stephen Taylor California Institute of Technology

Saul Teukolsky (co-organizer) California Institute of Technology

Massimo Tinto University of California, San Diego

Michele Vallisneri (co-organizer) Jet Propulsion Laboratory

Vijay Varma California Institute of Technology

Alberto Vecchio University of Birmingham



	List of Figures 5
1	Overview
1.1	Our Goals
1.2	Our Study 8
1.3	This Report
2	The Complexity and Richness of LISA Science 11
2.1	LISA Needs Waveforms! 11
2.2	LISA Hearts Astronomy! 16
3	Hard Problems
3.1	The Global Fit
3.2	Designing the LISA Catalog
3.3	Preparing for the Unexpected

4	New Tools	31
4.1	Machine Learning	31
4.2	Toward an Open-Science Model with Computational Notebooks	34
	Bibliography	36



2.1	Reduced-order modeling	12
2.2	Parameter-space coverage of current waveform catalogs	14
2.3	EMRI surrogate waveforms	15
2.4	Astronomical observatories that will generate LISA-relevant science	21
3.1 3.2	The "full enchilada" dataset from the Mock LISA Data Challenges	23 25
4.1	Receiver operating characteristic curves for a neural-network classifier	32



1.1 Our Goals

The space-based gravitational-wave (*henceforth, GW*!) observatory **LISA** (Amaro-Seoane et al., 2017) will offer **unparalleled science returns**, including a view of massive black hole mergers up to high redshifts, precision tests of general relativity and black hole structure, a census of thousands of compact binaries in the Galaxy, and the possibility of detecting stochastic signals from the early Universe.

While the Mock LISA Data Challenges (2006–2011) gave us confidence that LISA will be able to fulfill its scientific potential, we still have a rather incomplete idea of what the **end-to-end LISA** science analysis should look like. The task at hand is substantial:

- Our data reduction process needs to ensure the phase coherence of GW measurements across data gaps and instrument glitches over multiple years.
- Our waveform models need to reach part-in-10⁵ accuracy to maximize the science payoff of the mission, with sufficient computational efficiency to sample parameter space broadly.
- Our algorithms need to resolve thousands of individual sources of different types and strengths, all of them superimposed on the same multi-year dataset, while simultaneously characterizing the underlying noise-like stochastic background.
- Our catalogs need to represent the complex and high-dimensional joint distributions of estimated source parameters for all sources.

It is tempting to assume that current algorithms and prototype codes will scale up to this challenge, thanks to the greatly increased computational power that will become available by the time of LISA's launch in the early 2030s. In reality, harnessing that power will require very different **methods, adapted to future high-performance computational architectures** that we can only glimpse now. Thus, we need to begin our exploration at this time, seeking inspiration from other disciplines (e.g., big data processing, computational biology, the most advanced applications in astroinformatics) and learning to pose the same physical questions in different, future-proof ways—or even daring to imagine questions that will be tractable only with future machines.

The broad objective of this study was to imagine how evolved or rethought data analysis algorithms and source-modeling codes will perform the LISA science analysis on the computers of the future, with the hope of guiding LISA science and data analysis research and development in the years to come.

1.2 Our Study

The study began with a short course open to the Caltech community, with video available online,¹ featuring sections on LISA science (by Curt Cutler), the numerical modeling of LISA sources (by Deirdre Shoemaker), and LISA data analysis (by Stas Babak).

The study participants then convened for three sessions consisting of tutorials, discussions, and breakout groups:

- The session on data analysis problems and approaches included talks on stochastic techniques, machine/deep learning, the LISA Data Processing Center being prototyped in Paris, modern astronomical archives, as well as a panel on social aspects of writing scientific software. Breakout groups convened to discuss the global fit problem, data products, and the detection of unexpected signals.
- The session on numerical modeling included talks on new computational paradigms, parameter estimation, hydrodynamics in numerical relativity, and reduced-order modeling waveforms, as well as a panel on EMRI waveforms, which are especially challenging analytically and numerically. Breakout groups convened to discuss new ideas for waveform accuracy estimates, searches and parameter estimation, and the astronomy community interface.
- The session on **computation** included talks on data-intensive architectures, the scientific use of GPUs, "live" data products, and advanced visualization techniques.

¹http://kiss.caltech.edu/short_courses/LISA.html

The study participants then discussed follow-up work and the organization of report writing. Slides from all presentations are available on the KISS website.²

1.3 This Report

This report does not aim at providing an exhaustive retelling of the many interesting discussions that took place over the course of the study; rather, it focuses on the *stories and questions* that we consider likely to generate the most interesting and urgent directions for research.

The report includes three main sections. First, **The Complexity and Richness of LISA Science** covers the requirements of efficiency and accuracy imposed on waveforms by LISA's remarkable sensitivity, as well as the many astronomical use cases for the LISA data, which will need to be fulfilled with *ad hoc* data products. **Hard Problems** focuses on the difficulty of detecting and characterizing many thousands of superimposed sources; of compiling and making available a source catalog that can fully represent the LISA findings and enable insightful science; and of building confidence in the detection of previously unknown sources that inhabit LISA's proverbial *discovery space*. Finally, **New Tools** is concerned with the mostly untapped power of deep learning for GW analysis applications, and with emerging tools (such as computational notebooks) that promise agile, efficient, and reproducible access to astronomical archives and to all computations that can be performed with them.

An executive summary, such as would be customary for a report like this one, could not do justice to the exploratory nature of this study, and to the many strands that we touched upon briefly and that yet may figure into the ultimate solution to the LISA science analysis. We offer instead an **unordered list of takeaways** that one of us (M.V.) compiled shortly after the study. As you will see, some of these points are represented more generously than others in the rest of this report, so if you do read through, you will likely come up with your own different list. Nevertheless:

- There are many very deep research questions that stand between us and the solution to the LISA science analysis. It will take time to answer them, and we will need to think about them carefully and purposefully. However, since we need to get started on a practical level, we can begin by providing infrastructure for experimentation that will later become infrastructure for implementation. We should be especially careful to remain flexible while we are getting oriented.
- As for numerical modelers, they need practical performance and accuracy benchmarks as soon as possible, so that they can strategize to make LISA-grade waveforms across all source classes. As Saul Teukolsky points out, "This is a billion dollar experiment, and we do not want to sacrifice even the last vestiges of science that can be extracted from it."

²http://kiss.caltech.edu/workshops/LISA/LISA.html#Presentations

- The awe-inducing demonstrations of deep learning (*hereafter DL*) in many domains (and the first tentative applications in GW data analysis) show the potential of this technique, which will only become more important as it continues to drive processor development; thus, our wish to continue riding Moore's law will only be fulfilled if we embrace DL as a data analysis workhorse, or at least if we recast standard calculations using DL's mathematical language.
- Speaking of Moore's law, much of its ascent in recent times has come in the shape of ever more powerful graphical processing units, with larger and larger collections of elementary compute cores. GPU programming used to be a perilous and abstruse land, but thankfully the right tools (especially compilers and mathematical libraries) are now available to wield all of the Gigaflops we have been promised.
- We need to look outside the GW community and provide diverse venues to engage scientists of different kinds who will make the most of the LISA data, as well as untapped talent more generally. We can take advantage of "citizen" programmers, but we need virtuous policies and modern practices to organize scientific software development and to distribute credit fairly.
- The high-level LISA data products will need to serve a diverse community ranging from cosmographers to population modelers, from seekers of counterparts to experts in alternative theories of gravity. We can take inspiration from recent astronomical archives (Fermi, Kepler, SDSS) in tailoring and diversifying our data products. However, we should remember that the best internal data representation (for source characterization, data cleaning, and to handle the continuous influx of data) may not be same as is needed for archiving and for external users. Thus we need help from database professionals.
- New powerful collaborative tools are fulfilling a vision of open science and bringing computations to otherwise unwieldy datasets *in the cloud*. In addition, bespoke visualizations can be extremely insightful. While LISA will fly with onboard technology that is necessarily frozen at the beginning of its arc,³ the same need not be true for its science segment, if we remain flexible and observant.

³At least for hardware; it would be great to be able to perform small updates to software if they result in large functional improvements.

2. The Complexity and Richness of LISA Science

2.1 LISA Needs Waveforms!

In discussing **waveform accuracy**, it is useful to distinguish between the intrinsic error in the underlying *models* of physics (e.g., a numerical relativity simulation or a post-Newtonian expansion of Einstein's field equations; Buonanno and Sathyaprakash, 2015; Choptuik et al., 2015) and the errors incurred in making these models available as *templates*: *hybridization* error when marrying analytical and numerical waveforms; *approximation* error when fitting phenomenological models to simulations; *representation* error when casting waveforms onto reduced-order modeling bases. (One speaks generally of *surrogates* to encompass waveforms that can be evaluated efficiently and that approximate more careful and expensive computations.)

Physical model errors are addressed by theoretical work; template errors by more careful and potentially costly computation. Waveform accuracy requirements affect both the theoretical inputs and their data-analytic representations, and they will depend strongly on the sources under consideration. Ideally, we want the error budget to be dominated by instrument noise, and not by waveform limitations; this criterion remains the same for astrophysical parameter estimation, for anomaly identification in tests of general relativity, and for minimizing residuals after the subtraction of the loudest signals. For current (second-generation) ground-based GW observatories, the signal-to-noise ratio (SNR) for black hole–binary mergers will be ~100 at design sensitivity. For LISA, the comparable number will be ~10,000. Since the required phase accuracy for matched filtering scales as 1/SNR, we need about a factor of 100 improvement in waveform accuracy from today.¹ No current numerical relativity code can achieve this accuracy,

¹A sufficient condition for the systematic error from mismodeling to be negligible compared to the random effects of noise is for the *mismatch* to be smaller than $D/(2 \text{ SNR}^2)$, where D is the number of parameters in the waveform. See Chatziioannou et al. (2017).

even extrapolating beyond exascale computers. For EMRIs and Galactic binaries the SNR is lower, and the accuracy requirement scales accordingly. (By contrast, detection rates would be damaged only marginally by losses in accuracy of a few percent.)

How do we define **waveform accuracy requirements** formally for LISA? It is unclear that we have all of the mathematical tools and criteria to pose this question profitably (although we have well-established foundations; Lindblom et al., 2008). In space-project jargon, **what is the modeling and templating error budget for LISA?** Once we have this answer, we can ask questions such as:

- Are post-Newtonian (PN), effective-one-body (EOB), and numerical relativistic (NR) binary
 models sufficiently accurate to extract all science from loud massive black hole binary
 inspirals? E.g, are numerical relativity simulations sufficient to capture the fine details of
 mergers? How many cycles do they need to cover so that the "handoff" from PN/EOB to
 numerical models is sufficiently accurate?
- What about the representations of waveform models as phenomenological (*Phenom*) or reduced-order modeling (ROM) templates? Can these be improved at the cost of more computation, whether *online* during parameter estimation or during a preliminary stage?



Figure 2.1: The reduced-order modeling approach involves the offline construction and interpolation of a reduced basis for the signal space, in order to greatly accelerate the online use of templates for detection and parameter estimation. Credit: C. Galley & A. Chua.

The other important consideration that we should concentrate on is the specific **region of source parameter space to which LISA will be sensitive**: for instance, significant values of binary eccentricity and precession are challenging to numerical simulations, and they will be more relevant to LISA than they are to ground-based detections.

In terms of work that is happening within the LISA Consortium at the time of writing, it is relevant to ask if the **LISA Data Challenges** (LDC) (*LISA Data Challenges* 2018) could be used to answer these questions empirically. (However, the template families currently used in the LDC need to be updated to the state of the art.) Specifically, we would address the level of waveform error that is tolerable for different parameter estimation tasks. To do so, we need to designate a waveform family *as truth*, then perform parameter estimation with a different family.

For instance, if the reference waveforms are provided by long(er) numerical relativity simulations and parameter estimation is performed with *Phenom* surrogates, we would be testing the accuracy of the *Phenom* template implementation (which can also be done directly; Kumar et al., 2016). Conversely, we can use a surrogate family as reference and perform parameter estimation with a modified surrogate that includes *plausible* phasing/amplitude errors across parameter space (e.g., as obtained by comparing numerical grid refinements): doing so would exemplify the parameter estimation consequences of mismodeling error and potentially lead to quantitative accuracy criteria for models. In doing such tests, we would:

- where possible, seed numerical investigations with analytical results, such as Cutler and Vallisneri's analysis of parameter estimation bias (Cutler and Vallisneri, 2007), formulated for PN waveforms;
- make sure the analysis is sufficiently general—we should be able to apply the same methods to all waveform families using compatible definitions of intrinsic errors;
- perform the study up to the highest expected signal-to-noise ratios (for binary black-hole coalescences, 10⁴).

Once accuracy requirements are understood, the modeling community can set out to fulfill them. In the case of **numerical relativity**, it is unclear that current paradigms (e.g., domain-based parallelism) will suffice, even when coupled with the continued growth of computing power. Indeed, the movement toward sharing power among many compute cores raises many problems that limit performance, including communication bandwidth and latency, memory contention, and synchronization bottlenecks.

Recent efforts to reformulate the numerical solution of Einstein's equation in terms of **task-based parallelism** (Kidder et al., 2017) aim to address these limits, dividing computation into tasks within a dependency tree. The tasks correspond (roughly) to evaluating the functions required to advance the evolution equations; the parameters needed by the functions define their mutual



Figure 2.2: The current waveform catalogs for modern NR codes sample parameter space in different ways, but always sparsely. Credit: P. Kumar.

dependencies. A smart scheduler assigns the tasks that are ready to the compute cores, managing asynchronous information transfer, and (crucially) handling load balancing and *task stealing*.

It may be possible to improve the *status quo* for other aspects of NR modeling; for instance, by applying machine learning to automation in the "hand holding" of simulation runs, and possibly replacing certain computational bottlenecks with appropriately trained neural networks (an approach used in particle physics for expensive Monte Carlo likelihoods).

We would hope that the eventual modeling and representation errors would be so small that they can simply be ignored in data analysis. If this is not the case, we must formulate a robust method to propagate a characterization of the errors into data analysis pipelines, where the corresponding degrees of freedom can be treated as nuisance parameters and marginalized over (Moore et al., 2016).

Beyond the general schemes outlined above, there are specific considerations to be made for each class of LISA sources. For massive and stellar-mass black hole binaries with comparable mass ratios, we already have banks of numerical relativity simulations that can provide fiducial injections, although we need more simulations of eccentric binaries; PN waveforms also need to include eccentricity organically. To orient effort in both domains, we need to understand astrophysical expectations for eccentric systems. NR simulations are currently able to comfortably handle eccentricity and spin precession for mass ratios of up to 1/20 and spin magnitudes of up to 0.8. An extension to higher spins would require improved initial data, while higher mass ratios will be addressed better by black hole perturbation theory.



Figure 2.3: A recently developed EMRI surrogate is able to generate fast self-forced inspiral trajectories in milliseconds, but is presently implemented only for the first-order self-force on Schwarzschild. Credit: M. van de Meent & N. Warburton.

Intermediate-mass-ratio black hole binaries (IMRIs) are very problematic already at the modeling stage. They would require pushing perturbative analytical expansions to third order in the mass ratio—very difficult—although Le Tiec's "q to v" trick may help (Le Tiec et al., 2013). At worst, we may need to investigate new perturbative schemes for the self-force. Looking at currently available models and simulations, we could test EOB waveforms at high mass ratio with numerical relativity runs, which however would be very expensive.

The modeling of extreme-mass-ratio inspirals (EMRIs) for parameter estimation requires up to second-order self-force corrections to a Kerr geodesic orbit, or specifically their average dissipative effects. (The complete second-order contribution would be useful for IMRIs.) Waveforms produced directly from self-force calculations would be too computationally expensive for use in parameter estimation algorithms, and must be approximated by a surrogate with very minimal loss of accuracy. The fast near-identity-transform method of van de Meent and Warburton (van de Meent and Warburton, 2018) is promising for this purpose, although it depends on the construction of an accurate interpolant for the self-force when the calculations become available. For EMRI detection, fast kludge waveforms (Babak et al., 2007; Barack and Cutler, 2004; Chua and Gair, 2015) that utilize various approximations have been shown to be adequate (Chua et al., 2017), but current implementations need to be improved in speed and accuracy.

2.2 LISA Hearts Astronomy!

The LISA observations will seed many studies, investigations, and follow-ups across the astronomical community. Because the LISA archive will be the first of its kind (notwithstanding many similarities with source catalogs for mature ground-based GW detectors), it is important to **reach out to astronomers** to teach them the peculiarities of LISA data, and to learn about their requirements and use cases. KISS study participants outlined several avenues to do so, including holding regular workshops at professional conferences; providing worked-out data analysis examples (e.g., as computational notebooks, see below); and structuring the LISA archive to accept contributions from outside the project. This outreach could extend to the public through a carefully designed *citizen science project* similar to Gravity Spy (*Gravity Spy* 2018).

Thus, the LISA data archive will need to fulfill **many different use cases** in the astronomical community. The following are five examples (certainly not exhaustive) that we chose to probe the range of possible data products, and their connections to other astronomical catalogs.

- A census of massive black holes throughout the Universe. LISA will observe massive black hole (MBH) mergers throughout the entire history of the Universe, and EMRI events out to *z*~1. With several tens of events per year and a 5–10-year mission duration, LISA will build up a catalog of potentially hundreds of these events that can be used to form a census of MBH properties over cosmic time. In particular, LISA will be sensitive to MBH mergers with constituent BHs in the range 10⁴–10⁷ M_☉, which will still harbor a signature of seed formation at high redshifts. This catalog will also shed light on the impact of metallicity on MBH formation, MBH accretion efficiencies, and MBH accretion geometries. Models of hierarchical structure formation that have been constructed for consistency with observed AGN properties can agree with a wide variety of BH seed conditions (Gair et al., 2010, 2011; Sesana et al., 2011). Hence, LISA offers an unparalleled opportunity to peer into seed formation and accretion models of BHs across and throughout the history of the Universe. Required data products: Catalog of events, with independent posterior samples for each. The recovered event parameters should include luminosity distance, component spins, and component masses.
- Mapping Galactic structure and formation using white dwarf binaries. The population of white dwarfs (WDs) residing in the Milky Way is an important tracer of the Galactic gravitational potential and formation history (Kalirai, 2012). While the most interesting probes into Galactic structure and evolution history reside with the distant populations in the Milky Way halo, the Gaia mission² will observe the WD population with 100% completeness out to only 100 pc (Carrasco et al., 2014). This is largely due to the combined restrictions dictated by Gaia's limiting magnitude (20) and the inherent dim nature of WDs. Studies have shown, however, that LISA will observe gravitational

²http://sci.esa.int/gaia

waves from hundreds to thousands of resolved double WDs (DWDs) in the outer regions of the Galaxy with high signal to noise. Included in this large population is the subset of several dozens of detached and accreting DWDs observable by *both* LISA and Gaia (Korol et al., 2017; Kremer et al., 2017). The combined observations of WDs from Gaia in the regions out to 100 pc and LISA in the regions beyond could provide an unprecedented view into the structure and formation of the Milky Way. **Required data products**: Catalog of parameters of resolved WDs, including three-dimensional positions and masses.

- Understanding unknown GW signals. Although there are a number of sources LISA will definitely observe (such as MBH mergers and WD binaries), the mission may also discover completely unexpected signals. (See also 3.3 below.) In order to rule out instrumental origins, there will have to be a standard (and, if possible, automated) checklist of investigations to conduct, such as checking the time delay between the excitations felt at each spacecraft, as well as auxiliary sensor data. Furthermore, Robson and colleagues (Robson and Cornish, 2019) showed that differences in phasing in the LISA observables may allow us to distinguish unmodeled GW bursts from instrument glitches. Once an astrophysical origin can be established, as much information as possible concerning the unknown signal will need to be appear in the LISA catalog, for follow-up study and comparisons with electromagnetic catalogs. This is an exciting opportunity for scientists across many disciplines to work together to understand a new source of GWs. Required data products: Strain data containing the signal; as much information about the signal as possible, such as sky localization, length of the signal, peak strain amplitude, etc.
- Episodic accretion in white dwarf binary populations. Cataclysmic variables (CVs) and AM Canum Venaticorum stars (AM CVns) will comprise a large subset of the resolved compact binaries observed by LISA in the Milky Way (Nelemans, 2005; Solheim, 2010). Based on EM observations, several of these systems (verification binaries) are known to emit GWs detectable by LISA (Stroeer and Nelemans, 2009). Observing AM CVns in outburst leads to a wide range of opportunities to witness astrophysical processes of great interest, including accretion disk spiral density waves, guasi-period oscillations, and the launching of jets (Bonnet-Bidaud et al., 2015; Coppejans et al., 2016; Groot, 2001; Steeghs et al., 1997). AM CVn outbursts have been shown to have orbital period-dependent brightness variations, recurrence times, and durations, with shorter orbital periods leading to smaller brightness variations and shorter durations and recurrence times (Levitan et al., 2015). While duration times are 10-40 days over orbital periods of 22-36 minutes, the recurrence times are 30–500 days over the same orbital period range, highlighting the necessity for the rapid release of GW data from accreting WD binaries. Required data products: Catalog of parameters for accreting WDs, especially mass, chirps, and positions, made available on timescales of days to weeks.

• Multi-messenger observations of massive black hole mergers. Combined detections of the GW and EM signals powered by merging MBH binaries can significantly enhance the scientific return of LISA. EM signals provide improved localization and distance estimates, as well as information about the environment of the merging MBHs, thus opening new ways to use LISA detections in cosmology, astrophysics, and for tests of general relativity. Given the large uncertainties in the sky localization of LISA events, this will only be possible if the EM signals are sufficiently well understood to limit the rate of false positives in the LISA detection volume. The most promising EM counterparts to merging MBHs are powered by the accretion of gas onto the MBHs, both before and after merger. Our theoretical understanding of that process remains limited, in part due to the difficulty involved in simulating the many different time scales relevant to this problem, and in part due to large uncertainties regarding the feeding mechanism of MBH binaries and the large scale structure of the accretion flow. Required data products: Prompt alerts provided by LISA to EM observers, with tight sky localization.

As shown by these use cases, archives are a crucial part of astronomy! They preserve data gathered at tremendous expense of both treasure and talent, and they enable impactful archival science long after the end of a mission or survey. Study participant Rachel Akeson emphasized that easily searchable archival datasets facilitate innovative uses, especially where databases combine information from multiple missions and facilities (*NASA/IPAC Extragalactic Database* 2018). Archives contain:

- Raw and processed data from a mission. For LISA, data products have been tentatively assigned to four levels: 0 raw interferometric measurements and auxiliary data; 1 time-delay interferometry observables (TDI—the closest LISA comes to GW strain) and data quality indicators; 2 individual source parameter estimates and partially source-subtracted observables; 3 full catalogs. Transient event alerts (e.g., for imminent MBH mergers) would fall between 2 and 3.
- Other mission data. These might include spacecraft data that may not seem immediately relevant,³ metadata characterizing pipeline output for completeness and reliability, or simulated data of various kinds.
- High-level derived products and associations from the archive. For LISA, premiere
 products would include cross-identifications with EM objects from other missions and
 observatories.
- High-level data products from the community. For LISA, these would include reanalyses of individual sources or recalculation of the global solution, as well as independent

³For instance, Spitzer developed an exoplanet transit mode eight years after launch, using pointing data not originally used in processing pipeline. See Beichman and Deming (2017).

associations with other databases, for which good interfaces will be crucial. Standardization and reliability of community inputs is often an issue.

The level of archive support varies greatly, and for large space missions it normally includes dedicated user interfaces, bespoke tools and visualizations, and helpdesk support. Datasets can be updated continuously or episodically, when costly or delicate processing is needed. While astronomy is not (or not yet) "Big Data" by Silicon Valley standards, volume still poses issues for retrieval, searching, and processing. The standard mantra is now *bring the computation to the data*, with the usual qualifier, *in the cloud*. Akeson emphasized **careful planning and organization** early on, especially for large datasets that will then be hard to move and change, or for capabilities that may be precluded by cost if pursued later. She praised task-specific optimization and reduction of overheads, and warned against complexity, since "things that are complicated when they are small explode on you when they grow big."

Ultimately, the design of the LISA archive will reflect the peculiarities of a space-based GW observatory (see 3.2 below) and the use cases collected from the astronomical community. Nevertheless, it should reuse established solutions where available, and it can benefit from the experience (with regards to serving the scientific community and to engaging *citizen scientists*) of ground-based detectors and of survey space missions such as Fermi and Kepler. The rest of this section summarizes our discussion of such lessons learned.

The LIGO–Virgo Collaboration currently releases prompt event alerts, limited data stretches (4,096 s) around events published in refereed journals, and full strain datasets two years after the end of an observing run (although this delay may be shortened in the future). The Gravitational Wave Open Science Center⁴ (formerly the LIGO Open Science Center) hosts public datasets along with tutorials and tools to aid data use. It recently released the first *catalog of detected GW transients* (The LIGO Scientific Collaboration and the Virgo Collaboration, 2018b)⁵, as well as *posterior samples* for selected events. Some lessons applicable to LISA are that:

- we should carefully consider the content and format of data products ahead of time (and possibly release mock examples), so that future users have time to prepare downstream tools;
- we should think carefully about giving guidance as to the interpretation of the data. In the case of LIGO-Virgo, strain data releases were never meant to allow the reproduction of GW searches, but were nevertheless used in that way, resulting in almost certainly spurious claims (Creswell et al., 2017).

⁴https://www.gw-openscience.org ⁵http://www.gw-openscience.org/catalog

NASA's **Fermi**⁶ is a sky-scanning telescope with public data (*photons*!) and software. Photons are posted as they are detected, but then are revised with much better characterization in successive data releases. These releases include updated calibration data and software tools. In addition, Fermi publishes catalogs describing confirmed point sources with spectral model fits, plus a few diffuse sources. Typical users who wish to study a source use all "nearby" photons together with calibration and catalogs, and perform a **multisource spectral fit** (with a Fermi-provided tool) to account for the possibility that each photon may come from the target or from the wings of the point spread function of other nearby sources. This kind of entanglement has obvious parallels to LISA's confusion problem. Lessons include:

- Even simple tasks, such as computing light curves, requires reprocessing mission data with sophisticated mission-provided software. Furthermore, users need to exploit higher-level products (the catalogs) to make sense of lower-level ones (the photons). This will definitely be the case for certain LISA investigations, such as tests of general relativity on loud MBH inspirals, where fine waveform detail will be correlated with the parameters of other overlapping signals.
- For LISA, a (figurative) spectrum of increasingly informative catalogs may address such correlation problems. Some users would be happy with a basic catalog listing maximumlikelihood source parameters plus uncertainties for sources that pass some quantitative falsealarm criterion. At the next level, sources could list parameter correlations in multivariate normal form. Next, catalog entries could list substantial correlations with other sources, which may be shaky for the weakest signals. Finally, more elaborate single-source and joint posteriors would be specified wherever useful, or possibly on the basis of user requests.

NASA's **Kepler**⁷ employed the *transit method* to survey our Galactic neighborhood and discover hundreds of Earth-size and smaller planets in or near the habitable zone, characterizing the fraction of stars that might have such planets. The statistical nature of Kepler's goals, its increase in sensitivity at longer baselines, and its resurrection in a different operating mode after a hardware failure resulted in **multiple versions and multiple kinds** of astronomical data products and catalogs⁸ covering many different use cases. Some lessons applicable to LISA are that:

- the astronomical community can deal with multiple purposes of object catalogs; with iterative catalog versions; and with objects changing status between catalogs (or even appearing and disappearing)—but consistent versioning is paramount;
- the full range of products should be defined early to allow for good interfaces and data flow in processing and archiving;

⁶https://fermi.gsfc.nasa.gov

⁷https://www.nasa.gov/mission_pages/kepler

⁸https://keplerscience.arc.nasa.gov/data-products.html



Figure 2.4: Astronomical observatories that will generate LISA-relevant science before LISA flies. The search for EM counterparts will largely rely on existing catalogs, so the LISA archive will need interfaces and interoperability with several of them.

• project definitions and thresholds should be enforced strictly to avoid confusion (e.g., we should avoid changes in SNR thresholds or object status definitions).



3.1 The Global Fit

It was realized early (Prince et al., 2005) in the development of the LISA mission that the problem of combining the identification and characterization of thousands of individual sources would be a difficult one, especially in the presence of data gaps and artifacts, or of poorly measured, nonstationary/non-Gaussian noise.

Many reasons for worry. The "confusion" problem in LISA data analysis lies in the fact that multiple signals (if *overlapping* geometrically in signal space) interfere in the likelihood function. Typical overlaps are small, but small groups of sources will be in close proximity, and will thus be poorly resolvable. Moreover, with so many signals the confusion adds up and unresolved signals (including many millions of Galactic binaries) act as an effective non-Gaussian and nonstationary noise term (Cutler and Harms, 2006; Robson and Cornish, 2017).

Furthermore, even the number of sources is unknown; this calls for *transdimensional* likelihood exploration algorithms, where goodness of fit (which will always increase with the number of model parameters) is tempered by a regularization term (which happens naturally in a Bayesian framework; Laumann, 2018), or ideally by a strong astrophysical prior.

The basic building block of GW analysis is the likelihood/sampling distribution of instrument noise, which is evaluated most efficiently in a diagonal basis. For stationary Gaussian noise, the frequency domain provides such a representation. However LISA instrument noise will be nonstationary because of hardware effects, and also because of the seasonally modulated confusion background from the Galaxy.



Figure 3.1: The "full enchilada" dataset from the Mock LISA Data Challenges exemplifies the LISA confusion problem, albeit with great understatement in the actual number of sources. Credit: M. Vallisneri

Exploiting randomness. Stochastic methods such as Markov-Chain Monte Carlo (MCMC) (Sharma, 2017) have long been a workhorse in GW analysis. Applying them to specific families of GW signals is more in the realm of art (or, at least, craft) than than science, especially in the choice of efficient proposals that provide a reliable guide for stochastic exploration. In addition to the basic strategy of local posterior approximation with Fisher-matrix covariance, GW work has successfully employed *ad hoc* steps based on physical symmetries (e.g., jumps between multiple EMRI harmonics), precomputed (and necessarily approximated) global likelihood maps (e.g., an *F*-statistic map for global binaries), *differential evolution* (i.e., using past stochastic history to suggest jumps), and more. Approaches that improve on the basic Metropolis–Hastings exploration algorithm include *parallel tempering* (Swendsen and Wang, 1986; Vousden et al., 2016), Hamiltonian Monte Carlo (Neal, 2011), several variants of nested sampling, and more.

Other GW-specific tricks involve the analytical maximization or marginalization of the likelihood, whenever possible.

Glimpses of solutions. Cornish suggests that a wavelet decomposition similar to that used by LIGO's WaveBurst may allow the efficient representation of slowly varying nonstationary noise (Dahlhaus, 1997; Nason et al., 2000), as well as an expedient treatment of gaps. Such a technique creates a need for efficient waveform generation in the wavelet domain.

In the LIGO context, *BayesWave* (Cornish and Littenberg, 2015) represents a sophisticated and efficient approach to modeling signals, instrument glitches, and noise out of basic components (wavelets for signals and glitches) assembled in transdimensional fashion. Its approach to signal and noise priors is pragmatic rather than fundamental (what "true" prior probability can, after all, be assigned to a given wavelet?), and it is validated by the standard strategy of injecting and recovering simulated signals in actual and synthesized instrument noise.

To tackle the LISA global fit problem, Cornish suggests a strategy of "solving for anything," using a Bayesian model with dimension ~500,000, possibly updating blocks of sources and parameters in turn. A smaller-scale application of this approach successfully resolved tens of thousands Galactic binaries, and was demonstrated in the first round of Mock LISA Data Challenges (Crowder and Cornish, 2007). The problem can be parallelized by creating many equivalent chains simultaneously (although each will need to deal with all parameters), or many temperatures in a *parallel-tempering* sampler.

After surveying this landscape, study participants identified a list of key needs and challenges:

- Waveform overlaps need to be mapped extensively, both among different waveform families, and among different parameter regions for the same type of waveforms (e.g., for EMRIs, which have very complex waveforms with multiple harmonics that may lead to partial "matches" between distant regions of parameter space).
- Priors for the **number of detectable sources** are poorly known; overly strong assumptions may bias astrophysical inference from catalogs down the line.
- Extreme efficiency will be needed in generating waveforms and in producing the corresponding LISA responses (not a trivial operation for LISA).
- Bias from systematic/unmodeled effects (e.g., white-dwarf triples; Robson et al., 2018) needs to be excluded conclusively or accounted for explicitly. Likewise, instrument glitches will need to be recognized, then fit or excised.
- The **implementation** of the global fit is still defined very vaguely, and it is likely to be challenging technically. Issues to be clarified include:



Figure 3.2: Schematic representation of a possible global fit algorithm. Source-specific codes update the number and parameters of sources in each class, then pass on the resulting residuals (an example of *blocked Gibbs sampling*, see, e.g., Jensen et al., 1995). Once convergence is achieved, this iteration would sample the global parameter distribution. Credit: N. Cornish and C. Cutler

- Computational parallelism, which may be based on tasks, populations, or source classes. Can we have multiple pipelines that fit different clusters of "similar" signals?
- The logistics of iterative exploration (possibly proceeding from the strongest to the weakest sources). If iteration proceeds by subtracting progressively weaker signals, then the imperfect fit of the louder signals may leave significant residuals that confuse the fit of the weaker signals.
- The representation of very highly dimensional posteriors.
- The global fit solution will **evolve with time**: it will become better defined but more complex as we accrete more and more data, with subthreshold events transitioning into detectability.
- The quality of data and prevalence of noise will also vary naturally across time. Should our analysis proceed from the quietest times?
- How do we **evaluate and visualize** the **performance** (and correctness!) of likelihood exploration? (E.g., how do we know that we are exploring all relevant modes?)

Beyond these well-specified areas of investigations, study participants brainstormed several more open ideas to explore:

- Are there any **new methods** that could be applied to the problem of finding global maxima in high-dimensional search problems?
 - Is MCMC truly the only appropriate paradigm?
 - Are there alternatives to reversible-jump MCMC to handle transdimensionality?
 - Study participants mentioned the Hilbert–Huang transform, empirical mode decomposition, intrinsic time-scale decomposition.
- Could we use subsets of data to **soften likelihood surfaces**, and subsets of parameters to reduce the dimensionality of the problem? (E.g., with only a week of data, all binaries would be well-described by two parameters.) These fits/posteriors can then be evolved incrementally into more complete and complex ones.
- It may be possible to take sub-threshold EMRI fits on short data spans and use **pattern** recognition to assemble them into a consistent EMRI solution.
- Can machine learning help with the parallel allocation of stochastic chains (possibly for subsets of sources)? Likewise, can we use machine learning to tune other search parameters?
- How can we accelerate likelihood evaluations (e.g., with GPUs, *tensor processing units*, precomputation/reuse of computation) and where can we usefully approximate them?
- Can we fit confusion backgrounds of subthreshold sources? (E.g., EMRIs, by somewhat keying on their distinctive features?)

3.2 Designing the LISA Catalog

When Hogg and Lang state forcefully that "Telescopes don't make catalogs" (Hogg and Lang, 2010), they mean that the direct output of telescopes are images (arrays of pixels), whereas the desired scientific outcome of observations and surveys is *statistical inference* about astrophysical properties or astronomical populations. Standard catalogs, which typically describe best-fit parameters and error bars for all detected sources, are neither; they represent a historically entrenched intermediate data product, which may nevertheless diminish science payoffs if information is lost in compiling them.

There are several reasons to believe that a "traditional" LISA catalog may suffer mightily from this problem:

- Source parameter correlations. Most of LISA's sources are long-lived (radiating GWs in LISA's detection band for years or more), so tens of thousands of them will be present simultaneously in the LISA time series, and will need to be disentangled. If two detectable sources have strong overlap *in signal space*, then the parameter errors of one will be strongly correlated with the parameter errors of the other. Such correlations would need to be represented in the LISA catalog. As discussed above, even the *number* of sources may be uncertain: e.g., three WD binaries with similar sky positions and frequencies may be fit reasonably well by a set of two binaries. It is not obvious how such alternative explanations of the data may even be represented in a catalog. That such things will occur is to some extent an unavoidable fact of this mission, and the catalog must be designed with this in mind. Hints of understanding these sorts of "overlaps" appeared in the literature (e.g., Racine and Cutler (2007) use large-deviation theory to show that it is extremely improbable that a large number of low-SNR WD binaries could mimic either an EMRI signal or MBHB signal), but much remains murky.
- Quasi-detections. The catalog should include a list of sub-threshold sources that come near the statistical threshold of confident detection, as well as preliminary estimates of their parameters (and possibly estimates of when the detector threshold would be reached, for the purpose of potential EM follow-up).¹ Such sources contain information that is significant for astrophysical rate and population studies; furthermore, they play an important role at the boundary between confirmed individual detections and diffuse signal components. Subthreshold sources will suffer even more from inter-source parameter correlations.
- Non-Gaussian posteriors. Even for clearly detected sources, the inferred posterior probability distributions of parameters may have complicated, non-normal shapes, and possibly multiple local maxima with comparable probability weights. (The Mock LISA Data Challenges indicated that this is the case for binaries of rapidly spinning MBHs, as well as EMRIs.) The most benign catalog presentation for these sources would be a Gaussian mixture (a weighted sum of normal distributions) in parameter space, but a more expensive description as a Monte Carlo sampling of parameter values may be needed. Inter-source correlations then become even more problematic: for each source, we would want distributions that are marginalized with respect to all other sources and systematics; but we would also need a way to represent the correlations.

A vision for the LISA catalog. Looking beyond problems to aspirations and solutions, we may adopt Hogg and Lang's goal for a maximally performing catalog to allow statistical inference *just as if we were analyzing the raw data directly*. In mathematical terms, this means that we need access to the likelihood of the observed data as a function of a model (i.e., the parameters of all sources, plus parameters describing instrument noise and other effects). Hogg and Lang suggest three possible implementation steps of increasing technical complexity and computational cost:

¹Other missions, such as WISE, have already done this (see http://bit.ly/2NfECoJ).

- A traditional catalog, with sufficient information (parameter correlations, systematic effects) such that we are able reconstruct the likelihood of the data given variations with respect to nominal parameter estimates. The Gaia catalog² will come close to meeting these criteria. Because of the problems outlined above, it seems difficult for a LISA catalog to do the same, except perhaps for some special subsets of detected sources.
- 2. A collection (indeed, a sampling) of traditional catalogs, which collectively approximate the joint posterior distribution of all source parameters. A LISA data analysis pipeline that culminates in a global fit would in principle be able to produce a sampling of catalogs, but we would need to ensure that the sampling is actually representative of the underlying joint posterior. Hybrid solutions may be expedient. For the purpose of astrophysical population inference, it would be especially useful to commission a service (an application or web server) that could return equal-weight (sub-)catalog samples for a user-selected category of sources.
- 3. A self-contained software application or web service that returns the likelihood of the data given an input catalog. Computing one such response would amount to computing waveforms for all postulated sources and cross-correlating their sum with the LISA data—obviously a very expensive computational task, which may nevertheless be possible on large dedicated clusters or with clever distributed-data algorithms. Queries could exploit a reservoir of pre-computed waveforms, waveform basis decompositions, partially evaluated likelihoods, or other approximations.

The analytical and computational feasibility of these steps needs to be evaluated. In addition, the LISA catalog should embody a set of mission-specific features, which again pose intriguing data analytical or data science challenges.

- Internal representation. The global fit procedure discussed above can be seen, effectively, as creating and updating the most data-rich embodiment of the LISA catalog. Nevertheless, the most effective representation for the iterative fit is probably not what we will release for external use. The raw output of stochastic algorithms consists of very large numbers of variable-length parameter strings, each a sample of the posterior distribution. Are there lower-dimensional alternatives (somewhere between multivariate normals and posterior samples) that would be more expedient? Should we attach waveforms to posterior samples, or some sort of probabilistic description of subtraction residuals? What about waveform overlaps, or a waveform *Bayesogram* (Neil Cornish's coinage), consisting of a histogram of waveform values at each time?
- Updates. The signal-to-noise ratios for most detected sources will grow continuously throughout the mission, so frequent catalog updates will be useful for the broader scientific

²https://gea.esac.esa.int/archive

community and, crucially, to organize EM follow-ups and request *protected observations* at the expected time of binary mergers. Indeed, the updates would both refine parameter estimates and add new sources that have passed the detection threshold.

- A living catalog. The LISA catalog should be capable of archiving and presenting the results of EM observations, whether searches for EM transient counterparts for MBH mergers, or optical observations of Galactic binaries (as discussed in 2.2 above).
- Powerful visualization. Visualization tools will be especially useful for many tasks. For
 instance, mission scientists will want to perform spot checks of the largely automated
 catalogs, by zooming in on sources that have recently transitioned from *subthreshold* to *detected*, and/or sources whose parameter estimates are changing rapidly. The relatively
 high dimensionality of posterior distributions makes their visualization challenging.

3.3 Preparing for the Unexpected

In one of the breakout sessions, study participants were asked to discuss and document a chain of events that leads to detecting a GW signal that cannot be mapped to any known astrophysical source. How would LISA scientists convince the world that the signal is real? Ground-based observatories may face the same problem for signals that are only detected in *unmodeled* search pipelines, but the LISA case is complicated by the presence of many concomitant signals, and by the relative paucity of auxiliary channels monitoring instrument performance.

The participants discussed several scenarios:

- signals with clearly attributable EM counterparts (the easiest case);
- · signals with weird or mismatched EM counterparts;
- signals found by black box machine learning algorithms with strange parameter attributions;
- signals that match alternative-gravity dynamics better than general relativity;
- signals found in open data releases by nonstandard third-party pipelines;
- a set of similar low-significance "triggers" (according to standard pipelines) that were reanalyzed at the end of the mission;
- finally, the absence of a GW signal where one is expected (e.g., for a verification binary).

All of these scenarios have very interesting implications (both scientific and sociological), and may be instructive examples for roleplaying exercises in future mock data challenges. The study participants outlined a strong candidate, **GW350914**, which they characterized as follows:

• very strong, with SNR = 200, and high Bayes factor (above 5σ);

- not periodic, lasting 1.2 days, with peak frequency of 3 mHz;
- localized to within 25 square degrees in the sky, but without apparent EM counterparts;
- with no associated features in auxiliary channels (i.e., "green" data quality);
- and a very good way to celebrate the 20th anniversary of GW detection!

The participants further developed a *real-vs-not-real* checklist (Gawande, 2010), which starts out by questioning the state of the instrument, then transitions to data analysis. Such a process is formalized well for LIGO (The LIGO Scientific Collaboration and the Virgo Collaboration, 2016, 2018a):

- 1. All spacecraft systems are nominal.
- 2. Nothing (too) suspicious is seen in the auxiliary channels.
- 3. The rate of test mass decharging was not increased.
- 4. The spacecraft antennae were not moving or settling.
- 5. The spacecraft gas tanks were not experiencing any unusual phenomena.
- 6. Space weather was not significant.
- 7. Consistent time delays are observed between the excitations seen at the three LISA spacecraft (best assessed if all six LISA links are available).
- 8. The signal cannot be explained as the residual from the imperfect subtraction of other sources.
- 9. Nothing is seen in the null/breathing mode channel (if the frequency is sufficiently low that the null channel is actually GW-insensitive).
- 10. Signal consistency tests pass (though it is unclear what these would be for an unmodeled signal).
- 11. The signal can be consistently attributed to one region in the sky.
- 12. The distance estimate is sensible (i.e., not in our planetary backyard).
- 13. A reasonable combination of search pipelines detects the signal.
- The possibility of a hoax (through either hardware of software injection) can be investigated an excluded. This was done very carefully for the first LIGO detection (Landry and O'Reilly, 2015).



4.1 Machine Learning

Modern astronomical surveys and time domain astronomy programs produce huge amounts of data, often characterized by non-ideal properties (gaps, irregular sampling, heteroskedastic errors) that are ignored by many standard statistical methods. To tackle this onslaught, astronomers have been increasingly turning to *machine learning*, which largely eschews the principled application of statistical theory in favor of a **training-testing-application** sequence. Machine learning algorithms draw inferences *directly from the data*, either by comparing their evolving guesses with the "right" answers provided by the user (supervised learning), or by discovering hidden structure or groupings from first principles (unsupervised learning).

Traditional machine learning relies on designing and deriving relatively low-dimensional **feature vectors** from homogeneous data, and then feeding the features to classification (guessing the discrete class of an input) or regression (estimating continuous parameters) algorithms. For the same type of data, many different features can be defined (and many end up being very correlated in practice), and the choice of features can drastically alter performance. Furthermore, "naturally" defined features can carry physical meaning that is lost when they are chosen from a more mathematically amorphous set, as in Principal Component Analysis (Donalek et al., 2013). When features or classes are obvious to human eyes but elusive for mathematical definition, citizen science¹ can provide training "labels" for supervised learning.

Deep learning (Schmidhuber, 2015), which generally refers to stochastically trained *deep* neural networks (DNNs) with many layers, works directly from "raw" data, but its box is rather black.

¹https://www.zooniverse.org



Figure 4.1: Receiver operating characteristic curves for a neural-network classifier as compared to standard matched-filtering, in the context of LIGO-type detection. Credit: H. Gabbard, et al.

In its most impressive application so far, Google DeepMind's *AlphaZero* application learned to play chess, shogi, and Go at superhuman levels by unsupervised reinforcement learning—that is, by *playing itself* on the basis of the rules of the game alone (Silver et al., 2018). Among other techniques, AlphaZero relies on **convolutional neural networks** (CNNs), which apply multiple linear convolutions to input vectors and arrays, and which are employed broadly in computer vision and other image processing. Trained convolutional networks are opaque in how they achieve their tasks, but the technique of *DeconvNets* (deconvolution networks) can be used to project the "feature" activations, such as the notion of "cat" for an animal image classifier, back to the input pixel space (Olah et al., 2018; Zeiler and Fergus, 2014). In fact, studying activation patterns can identify useful features, learned directly from the data.

Study participant Ashish Mahabal described the application of an encoder-decoder CNN to the task of image differencing in transient detection (Sedaghat and Mahabal, 2018). The approach "encapsulates all of the steps of a traditional image-subtraction pipeline—image registration, background subtraction, noise removal, PSF matching and subtraction—in a single real-time convolutional network." Datasets characterized by an "arrow of time" can be tackled by a

deep-learning technique known as *recurrent neural networks* (RNNs; Karpathy, 2015), which allow for loops, or equivalently for the propagation of information across multiple iterations of the network. *Long Short Term Memory* networks (LSTMs; Olah, 2015) are RNNs that are capable of learning long-term dependencies, and can be used for light curves and other time series.

Convolutional networks have been applied to the detection and localization of GW signals from LIGO black hole mergers, in both the time domain (Gebhard and Kilbertus, 2017) and the frequency domain (Gabbard et al., 2018; George and Huerta, 2018b). The analysis spanned simulated LIGO-type signals (Fan et al., 2019; Nakano et al., 2019; Rebei et al., 2018), as well as real LIGO data (George and Huerta, 2018a). CNNs were also tasked with the classification of transient noise events ("glitches") in LIGO data (George et al., 2018; Razzano and Cuoco, 2018), and with the denoising of LIGO signals within *recurrent autoencoders* (Shen et al., 2017; Wei and Huerta, 2019). Finally, study participants Alvin Chua, Chad Galley, and Michele Vallisneri have demonstrated the use of DNNs as fast waveform interpolants in a ROM representation, providing benefits such as native parameter estimation and analytic waveform derivatives (Chua et al., 2019).

The KISS discussion of machine learning (and DL especially), raised a number of questions (Q) and (wishful?) suggestions (S). These are some of them:

- Q: Can DNNs learn the ideal proposal distribution for MCMC? The ideal Bayesian proposal is, in fact, the posterior itself; it may be possible, even if it seems very ambitious, to train DNNs to return approximations to Bayesian posteriors given the data by exposing them to a large variety of noisy datasets. More in line with the spirit of the question, DNNs may learn to offer "smart" proposal moves within MCMC runs, since for each proposed move we know immediately whether it is accepted; however the Metropolis–Hastings rule requires that for each proposal x → y we know the probability of the inverse proposal y → x. This may be difficult to arrange.
- Q: How are DNN architecture (e.g., depth, connectivity...) and hyperparameters (e.g., learning rate) decided? It's a **black art**. Some groups use Bayesian methods.
- S: While DL is opaque (and by extension not statistically "pure"), it could provide a rough map of where to look for signals in GW searches. The parallel with stochastic methods is interesting: these are used to solve problems that we can specify analytically, but cannot solve easily; supervised DNNs excel where the desired outcome is clear, but we cannot even formulate a way to get the solution. However, we may still need some reframing of GW problems before we can make headway.
- S: Training sets containing GW signals are generated quickly with simulation; generative networks may be more appropriate for glitch modeling, where examples in data are limited.

- S: Instead of using machine learning to replace classical solutions to the big problems of GW data analysis (detection, parameter estimation, data quality), we should go over the components of pipelines, and figure out which could be made faster by replacing them with fast (but expensive to train) black boxes. For instance, likelihoods (given the data) can be learned and *emulated* (Graff et al., 2012); the maximum-likelihood parameters could be learned given loads of noisy datasets; as suggested above, the control systems that govern the solution of NR could benefit from a steady automated hand.
- S: DL may be applicable to the multi-messenger problem of correlating LISA catalog entries with other EM catalogs.
- S: We should tap into the large and enthusiastic user/scholar community flocking to DL. If we describe problems well (e.g., on "competition" portals such as Kaggle,² or in the LISA Data Challenges themselves) and if we abstract them as much as possible from GW jargon, talented people may be interested in applying their own methods.
- S (based on a presentation by an Nvidia GPU software engineer): DL is increasingly driving the development and capabilities of new GPUs; conversely, the continued increase in computing power in modern computers relies increasingly on the teraflops unleashed by massively parallel arrays of GPU compute cores. It follows that if GW data analysis is to continue benefiting from Moore's law, it must reformulate its tasks as DL problems, or at least as algorithms that use the same mathematical primitives. (In fact, Google, Apple, and other large players are now building very fast and efficient *tensor processing units* that are strictly designed for DNNs; see Sato, 2018).

4.2 Toward an Open-Science Model with Computational Notebooks

Computational notebooks are akin to laboratory notebooks for scientific computing. They let users mix and match code (and the output of running it), data, text, equations, images, and more. They allow for an exploratory style of computing that can be compared to a conversation between data and the analyst. The conversation is recorded and can be replayed, modified, adapted, and (crucially) shared with other scientists, who are able to inspect, reproduce, and build on the annotated computations of their colleagues. While computational notebooks appeared 30 years ago in Wolfram's *Mathematica* (which later brought them to an unmatched level of sophistication), **their adoption has grown exponentially** in the last few years thanks to the *Jupyter* project (Perkel, 2018), which brings together several freely available components (notably the Python language, JavaScript, and modern web browsers) into a usable, portable, and overall powerful experience.

²https://www.kaggle.com

The next enabling step toward open³ and seamless collaboration in computational and data science projects are technologies that allow **packaging data with code**, with all of the elements and libraries that it needs to run. To continue the laboratory metaphor, packaging allows researchers (whether experienced or novice) and even the public to instantly replicate a copy of a laboratory (with equipment and samples) and immediately start research, either from scratch or by following a completed experiment. Examples of packaging tools include *JupyterHub*,⁴ which makes it possible to spawn Jupyter instances running on a server and accessible remotely, as well as *Binder*,⁵ which turns a *GitHub* repository containing notebooks into an executable environment Binder leverages *Docker* containers (Silver, 2017) to fully reproduce software environments for complete reproducibility⁶. Major *cloud* providers such as Google and Microsoft likewise offer computational notebook services, providing access (both free and paid) to significant computing resources.

Study participant Titus Brown noted eloquently that live data products (not just notebooks, but whatever they will evolve into) can "sucker scientists into doing better software development by creating a community of practice around open science," and offered several suggestions for how a **mission like LISA could benefit**:

- Notebooks can be engineered as a primary mode of access to making and using the LISA data products, and they can sit next to the data (and to all kinds of derived data products) to substantially reduce data transfer and synchronization. (As a matter of fact, modern astronomical data archives such as NOAO's DataLab⁷ already allow this today.)
- Notebooks can be generated and published automatically when new raw data arrives, documenting data reduction and allowing "live" inspection of data quality, new and improved detections, and so on. (NANOGrav has already moved most of its data reduction to notebooks that are run with varying degrees of supervision.)
- Notebooks can form the basis for a live paper of the future (*Envisioning the Scientific Paper of the Future* 2017), although other components of the scientific process (such as review and annotation⁸ need to find an appropriate digital format.
- Tutorials in the form of notebooks simplify the onboarding of new people, support greater participation by a broader community (leading to concrete measures of impact), and

³*Open Science* is the practice of science in such a way that others can collaborate and contribute, where research data, lab notes and other research processes are freely available, under terms that enable reuse, redistribution and reproduction of the research and its underlying data and methods (https://www.fosteropenscience.eu).

⁴https://jupyterhub.readthedocs.io

⁵https://mybinder.org

⁶...but probably not for long, as Titus Brown argues in http://ivory.idyll.org/blog/ 2017-pof-software-archivability.html.

⁷https://datalab.noao.edu

⁸https://web.hypothes.is

significantly ease the training of students and users. (Indeed the LISA Data Challenges have already offered tutorials based on notebooks hosted on Google's *Colaboratory* service.)

The purpose of this section is not to glorify one tool as the solution for the ills and needs of the world (a well rehearsed digital pitfall). Rather, we are awed and encouraged by the community effort that made the tool possible, and by the philosophy that it embodies: an embrace of **openness**, **collaboration**, **community**, **and sustainability**. Furthermore, as emphasized during the study's panel on the social elements of coding, transformative tools must be accompanied by changes in the practice and culture of science and scientific institutions:

- Academia and space agencies must recognize (for hiring and promotion) code as a scientific artifact on par with technical articles. In short, incentives must be in place to reward virtuous behavior.
- Likewise, we should **nurture the authors of scientific software** as much as we do instrument builders.
- Even if we treat LISA software (or parts of it) as an open-source project, we can still rely on professional software engineers, and hire them for key tasks (including code reviews!).

Finally, we should honor the selfless commitment of so many talented individuals to LISA science by living up to the highest ideals of our endeavor—fostering diversity through inclusion, and striving to realize an **empowering**, **embracing**, **and egalitarian environment**. Only a happy and thriving LISA community will create the strongest and most impactful science.



- Amaro-Seoane, P. et al. (2017). "Laser Interferometer Space Antenna". In: *arXiv e-prints*. arXiv: 1702.00786 [astro-ph.IM].
- Babak, S. et al. (2007). ""Kludge" gravitational waveforms for a test-body orbiting a Kerr black hole". In: *Phys. Rev. D* 75, p. 024005. DOI: 10.1103/PhysRevD.75.024005.
- Barack, L. and C. Cutler (2004). "LISA capture sources: Approximate waveforms, signal-tonoise ratios, and parameter estimation accuracy". In: *Phys. Rev. D* 69, p. 082005. DOI: 10.1103/PhysRevD.69.082005.
- Beichman, C. A. and D. Deming (2017). "Observing Exoplanets with the Spitzer Space Telescope".
 In: Handbook of Exoplanets. Ed. by Hans J. Deeg and Juan Antonio Belmonte. Springer, pp. 1–25. DOI: 10.1007/978-3-319-30648-3_78-1.
- Bonnet-Bidaud, J. M. et al. (2015). "Quasi-periodic oscillations in accreting magnetic white dwarfs. I. Observational constraints in X-ray and optical". In: Astronomy and Astrophysics 579, A24. DOI: 10.1051/0004-6361/201425482. arXiv: 1506.05891 [astro-ph.HE].
- Buonanno, A. and B. S. Sathyaprakash (2015). "Sources of Gravitational Waves: Theory and Observations". In: *General Relativity and Gravitation: A Centennial Perspective*. Ed. by Abhay Ashtekar et al. Cambridge University Press, pp. 287–346. DOI: 10.1017/CB09781139583961. 009. Also at https://arxiv.org/pdf/1410.7832.pdf.
- Carrasco, J. M. et al. (2014). "Gaia photometry for white dwarfs". In: *Astronomy and Astrophysics* 565, A11. DOI: 10.1051/0004-6361/201220596. arXiv: 1403.6045 [astro-ph.IM].

- Chatziioannou, K. et al. (2017). "Analytic Gravitational Waveforms for Generic Precessing Binary Inspirals". In: *Phys. Rev. Lett.* 118.5, p. 051101. DOI: 10.1103/PhysRevLett.118.051101. arXiv: 1606.03117 [gr-qc].
- Choptuik, M. W., L. Lehner, and F. Pretorius (2015). "Probing Strong-Field Gravity Through Numerical Simulations". In: *General Relativity and Gravitation: A Centennial Perspective*.
 Ed. by Abhay Ashtekar et al. Cambridge University Press, pp. 361–411. DOI: 10.1017/ CB09781139583961.011. Also at https://arxiv.org/pdf/1502.06853.pdf.
- Chua, A. J. K. and J. R. Gair (2015). "Improved analytic extreme-mass-ratio inspiral model for scoping out eLISA data analysis". In: *Class. Quantum Grav.* 32, p. 232002. DOI: 10.1088/0264– 9381/32/23/232002.
- Chua, A. J. K., C. R. Galley, and M. Vallisneri (2019). "Reduced-Order Modeling with Artificial Neurons for Gravitational-Wave Inference". In: *Physical Review Letters* 122.21, 211101, p. 211101. DOI: 10.1103/PhysRevLett.122.211101. arXiv: 1811.05491 [astro-ph.IM].
- Chua, A. J. K., C. J. Moore, and J. R. Gair (2017). "Augmented kludge waveforms for detecting extreme-mass-ratio inspirals". In: *Phys. Rev. D* 96, p. 044005. DOI: 10.1103/PhysRevD.96. 044005.
- Coppejans, D. L. et al. (2016). "Dwarf nova-type cataclysmic variable stars are significant radio emitters". In: *MNRAS* 463, pp. 2229–2241. DOI: 10.1093/mnras/stw2133. arXiv: 1608.06295 [astro-ph.SR].
- Cornish, N. J. and T. B. Littenberg (2015). "BayesWave: Bayesian inference for gravitational wave bursts and instrument glitches". In: *Class. Quantum Grav.* 32, p. 135012. DOI: 10.1088/0264-9381/32/13/135012. arXiv: 1410.3835 [gr-qc].
- Creswell, J. et al. (2017). "On the time lags of the LIGO signals". In: Journal of Cosmology and Astro-Particle Physics 2017, p. 013. DOI: 10.1088/1475-7516/2017/08/013. arXiv: 1706.04191 [astro-ph.IM].
- Crowder, J. and N. Cornish (2007). "A Solution to the Galactic Foreground Problem for LISA". In: *Phys. Rev. D* 75, p. 043008. DOI: 10.1103/PhysRevD.75.043008. arXiv: astro-ph/0611546 [astro-ph].
- Cutler, C. and J. Harms (2006). "Big Bang Observer and the neutron-star-binary subtraction problem". In: *Phys. Rev. D* 73, p. 042001. DOI: 10.1103/PhysRevD.73.042001. eprint: gr-qc/0511092.
- Cutler, C. and M. Vallisneri (2007). "LISA detections of massive black hole inspirals: Parameter extraction errors due to inaccurate template waveforms". In: *Phys. Rev. D* 76, p. 104018. DOI: 10.1103/PhysRevD.76.104018. arXiv: 0707.2982 [gr-qc].

- Dahlhaus, R. (1997). "Fitting time series models to nonstationary processes". In: Ann. Statist. 25, pp. 1–37. DOI: 10.1214/aos/1034276620.
- Donalek, C. et al. (2013). "Feature selection strategies for classifying high dimensional astronomical data sets". In: 2013 IEEE International Conference on Big Data, pp. 35–41. URL: https://doi.org/10.1109/BigData.2013.6691731.
- Envisioning the Scientific Paper of the Future (2017). Caltech symposium. URL: http://bit. ly/2WyOKNB.
- Fan, X. et al. (2019). "Applying deep neural networks to the detection and space parameter estimation of compact binary coalescence with a network of gravitational wave detectors". In: *Science China Physics, Mechanics, and Astronomy* 62, p. 969512. DOI: 10.1007/s11433-018-9321-7. arXiv: 1811.01380 [astro-ph.IM].
- Gabbard, H. et al. (2018). "Matching Matched Filtering with Deep Networks for Gravitational-Wave Astronomy". In: *Phys. Rev. Lett.* 120, p. 141103. DOI: 10.1103/PhysRevLett.120.141103. arXiv: 1712.06041 [astro-ph.IM].
- Gair, J. R., C. Tang, and M. Volonteri (2010). "LISA extreme-mass-ratio inspiral events as probes of the black hole mass function". In: *Phys. Rev. D* 81, 104014, p. 104014. DOI: 10.1103/PhysRevD.81.104014. arXiv: 1004.1921 [astro-ph.GA].
- Gair, J. R. et al. (2011). "Constraining properties of the black hole population using LISA". In: *Class. Quantum Grav.* 28, 094018, p. 094018. DOI: 10.1088/0264-9381/28/9/094018. arXiv: 1009.6172 [gr-qc].
- Gawande, A. (2010). The Checklist Manifesto: How to Get Things Right. Henry Holt and Company. URL: https://books.google.com/books?id=x3IcNujwHxcC.
- Gebhard, T. and N. Kilbertus (2017). ConvWave: Searching for Gravitational Waves with Fully Convolutional Neural Nets. URL: https://github.com/nikikilbertus/convwave.
- George, D. and E. A. Huerta (2018a). "Deep Learning for real-time gravitational wave detection and parameter estimation: Results with Advanced LIGO data". In: *Phys. Lett. B* 778, pp. 64–70. DOI: 10.1016/j.physletb.2017.12.053. arXiv: 1711.03121 [gr-qc].
- (2018b). "Deep neural networks to enable real-time multimessenger astrophysics". In: *Phys. Rev. D* 97, p. 044039. DOI: 10.1103/PhysRevD.97.044039.
- George, D., H. Shen, and E. A. Huerta (2018). "Classification and unsupervised clustering of LIGO data with Deep Transfer Learning". In: *Phys. Rev. D* 97, 101501, p. 101501. DOI: 10.1103/PhysRevD.97.101501. arXiv: 1706.07446 [gr-qc].

- Graff, P. et al. (2012). "BAMBI: blind accelerated multimodal Bayesian inference". In: *MN-RAS* 421, pp. 169–180. DOI: 10.1111/j.1365-2966.2011.20288.x. arXiv: 1110.2997 [astro-ph.IM].
- Gravity Spy (2018). URL: https://www.zooniverse.org/projects/zooniverse/gravityspy/about/results.
- Groot, P. J. (2001). "Evolution of Spiral Shocks in U Geminorum during Outburst". In: Astrophys. J. Lett. 551, pp. L89–L92. DOI: 10.1086/319826. eprint: astro-ph/0102416.
- Hogg, D. W. and D. Lang (2010). "Telescopes don't make catalogues!" In: GAIA: At the Frontiers of Astrometry. Ed. by C. Turon, F. Meynadier, and F. Arenou. Vol. 45. EAS Publications Series, pp. 351–358. DOI: 10.1051/eas/1045059. arXiv: 1008.0738 [astro-ph.IM].
- Jensen, C. S., U. Kjærulff, and A. Kong (1995). "Blocking Gibbs sampling in very large probabilistic expert systems". In: *International Journal of Human-Computer Studies* 42, pp. 647–666.
- Kalirai, J. S. (2012). "The age of the Milky Way inner halo". In: *Nature* 486, pp. 90–92. DOI: 10.1038/nature11062. arXiv: 1205.6802.
- Karpathy, K. (2015). The Unreasonable Effectiveness of Recurrent Neural Networks. blog post. URL: https://karpathy.github.io/2015/05/21/rnn-effectiveness.
- Kidder, L. E. et al. (2017). "SpECTRE: A task-based discontinuous Galerkin code for relativistic astrophysics". In: J. Comp. Phys. 335, pp. 84–114. DOI: 10.1016/j.jcp.2016.12.059. arXiv: 1609.00098 [astro-ph.HE].
- Korol, V. et al. (2017). "Prospects for detection of detached double white dwarf binaries with Gaia, LSST and LISA". In: *MNRAS* 470, pp. 1894–1910. DOI: 10.1093/mnras/stx1285. arXiv: 1703.02555 [astro-ph.HE].
- Kremer, K. et al. (2017). "Accreting Double White Dwarf Binaries: Implications for LISA". In: Astrophys. J. 846, p. 95. DOI: 10.3847/1538-4357/aa8557. arXiv: 1707.01104 [astro-ph.HE].
- Kumar, P. et al. (2016). "Accuracy of binary black hole waveform models for aligned-spin binaries".
 In: *Phys. Rev. D* 93, p. 104050. DOI: 10.1103/PhysRevD.93.104050. arXiv: 1601.05396 [gr-qc].
- Landry, M. and B. O'Reilly (2015). *LLO, LHO site status for GW150914*. LIGO technical note. URL: https://dcc.ligo.org/LIGO-L1500138/public.
- Laumann, Felix (2018). How Bayesian methods embody Occam's razor (blog post). URL: https://medium.com/neuralspace/how-bayesian-methods-embody-occams-razor-43f3d0253137.

- Le Tiec, A. et al. (2013). "Periastron advance in spinning black hole binaries: Gravitational self-force from numerical relativity". In: *Phys. Rev. D* 88, p. 124027. DOI: 10.1103/PhysRevD. 88.124027.
- Levitan, D. et al. (2015). "Long-term photometric behaviour of outbursting AM CVn systems". In: MNRAS 446, pp. 391-410. DOI: 10.1093/mnras/stu2105. arXiv: 1410.6987 [astro-ph.SR].
- Lindblom, L., B. J. Owen, and D. A. Brown (2008). "Model waveform accuracy standards for gravitational wave data analysis". In: *Phys. Rev. D* 78, p. 124020. DOI: 10.1103/PhysRevD. 78.124020. arXiv: 0809.3844 [gr-qc].
- LISA Data Challenges (2018). URL: https://lisa-ldc.lal.in2p3.fr/ldc.
- Moore, C. J. et al. (2016). "Improving gravitational-wave parameter estimation using Gaussian process regression". In: *Phys. Rev. D* 93, p. 064001. DOI: 10.1103/PhysRevD.93.064001.
- Nakano, H. et al. (2019). "Comparison of various methods to extract ringdown frequency from gravitational wave data". In: *Phys. Rev. D.* 99.12, p. 124032. DOI: 10.1103/PhysRevD.99. 124032. arXiv: 1811.06443 [gr-qc].
- NASA/IPAC Extragalactic Database (2018). URL: https://ned.ipac.caltech.edu.
- Nason, G. P., R. Von Sachs, and G. Kroisandt (2000). "Wavelet processes and adaptive estimation of the evolutionary wavelet spectrum". In: *J. Royal Stat. Soc. B* 62, pp. 271–292. URL: https://www.jstor.org/stable/3088859.
- Neal, R. M. et al. (2011). "MCMC using Hamiltonian dynamics". In: Handbook of Markov Chain Monte Carlo. Ed. by Steve Brooks et al., pp. 113–162.
- Nelemans, G. (2005). "AM CVn stars". In: The Astrophysics of Cataclysmic Variables and Related Objects. Ed. by J.-M. Hameury and J.-P. Lasota. Vol. 330. Astronomical Society of the Pacific Conference Series, p. 27. eprint: astro-ph/0409676. URL: http://aspbooks.org/custom/ publications/paper/330-0027.html.
- Olah, C. (2015). Understanding LSTM Networks. blog post. URL: https://colah.github.io/ posts/2015-08-Understanding-LSTMs.
- Olah, C. et al. (2018). "The Building Blocks of Interpretability". In: *Distill*. URL: https://distill.pub/2018/building-blocks.
- Perkel, J. M. (2018). "Why Jupyter is data scientists' computational notebook of choice". In: *Nature* 563, pp. 145–146. DOI: 10.1038/d41586-018-07196-1.

- Prince, T., B. Schumaker, and M. Vallisneri (2005). Analysis Methods for Interferometric Gravitational-wave Observations from Space (AMIGOS). URL: http://www.vallis.org/ publications/amigos-2005.pdf.
- Racine, É. and C. Cutler (2007). "Gaussianity of LISA's confusion backgrounds". In: *Phys. Rev. D* 76, p. 124033. DOI: 10.1103/PhysRevD.76.124033. arXiv: 0708.4242 [gr-qc].
- Razzano, M. and E. Cuoco (2018). "Image-based deep learning for classification of noise transients in gravitational wave detectors". In: *Class. Quantum Grav.* 35, p. 095016. DOI: 10.1088/1361-6382/aab793. arXiv: 1803.09933 [gr-qc].
- Rebei, A. et al. (2018). In: arXiv e-prints. arXiv: 1807.09787 [gr-qc].
- Robson, T. and N. Cornish (2017). "Impact of galactic foreground characterization on a global analysis for the LISA gravitational wave observatory". In: *Class. Quantum Grav.* 34, p. 244002. DOI: 10.1088/1361-6382/aa9601. arXiv: 1705.09421 [gr-qc].
- Robson, T. and N. J. Cornish (2019). "Detecting Gravitational Wave Bursts with LISA in the presence of Instrumental Glitches". In: *Phys. Rev. D* 99.2, p. 024019. DOI: 10.1103/PhysRevD. 99.024019. arXiv: 1811.04490 [gr-qc].
- Robson, T. et al. (2018). "Detecting hierarchical stellar systems with LISA". In: *Phys. Rev. D* 98, p. 064012. DOI: 10.1103/PhysRevD.98.064012. arXiv: 1806.00500 [gr-qc].
- Sato, K. (2018). What makes TPUs fine-tuned for deep learning? Google Cloud blog. URL: https://bit.ly/2Wt78HL.
- Schmidhuber, J. (2015). "Deep learning in neural networks: An overview". In: *Neural Networks* 61, pp. 85–117. ISSN: 0893-6080. DOI: 10.1016/j.neunet.2014.09.003.
- Sedaghat, N. and A. Mahabal (2018). "Effective image differencing with convolutional neural networks for real-time transient hunting". In: MNRAS 476, pp. 5365–5376. DOI: 10.1093/ mnras/sty613. arXiv: 1710.01422 [astro-ph.IM].
- Sesana, A. et al. (2011). "Reconstructing the massive black hole cosmic history through gravitational waves". In: *Phys. Rev. D* 83, p. 044036. DOI: 10.1103/PhysRevD.83.044036. arXiv: 1011.5893 [astro-ph.CO].
- Sharma, S. (2017). "Markov Chain Monte Carlo Methods for Bayesian Data Analysis in Astronomy". In: Annual Reviews of Astronomy and Astrophysics 55, pp. 213–259. DOI: 10.1146/annurevastro-082214-122339. arXiv: 1706.01629 [astro-ph.IM].
- Shen, H. et al. (2017). "Denoising Gravitational Waves using Deep Learning with Recurrent Denoising Autoencoders". In: *arXiv e-prints*. arXiv: 1711.09919 [gr-qc].
- Silver, A. (2017). "Software simplified". In: Nature 546, pp. 173-174. DOI: 10.1038/546173a.

- Silver, D. et al. (2018). "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play". In: *Science* 362, pp. 1140–1144. ISSN: 0036-8075. DOI: 10.1126/science.aar6404.
- Solheim, J.-E. (2010). "Update on AM CVn stars". In: ed. by K. Werner and T. Rauch. Vol. 1273. American Institute of Physics Conference Series, pp. 299–304. DOI: 10.1063/1.3527826.
- Steeghs, D., E. T. Harlaftis, and K. Horne (1997). "Spiral structure in the accretion disc of the binary IP Pegasi". In: MNRAS 290, pp. L28–L32. DOI: 10.1093/mnras/290.2.L28. eprint: astro-ph/9708005.
- Stroeer, A. and G. Nelemans (2009). "The influence of short-term variations in AM CVn systems on LISA measurements". In: MNRAS 400, pp. L24–L28. DOI: 10.1111/j.1745-3933.2009. 00754.x. arXiv: 0909.1796 [gr-qc].
- Swendsen, R. H. and J.-S. Wang (1986). "Replica Monte Carlo simulation of spin glasses". In: *Phys. Rev. Lett.* 57, pp. 2607–2609. DOI: 10.1103/PhysRevLett.57.2607.
- The LIGO Scientific Collaboration and the Virgo Collaboration (2016). "Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914". In: *Class. Quantum Grav.* 33, p. 134001. DOI: 10.1088/0264-9381/33/13/134001. arXiv: 1602.03844 [gr-qc].
- (2018a). "Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO's first observing run". In: *Class. Quantum Grav.* 35, p. 065010. DOI: 10.1088/1361-6382/aaaafa. arXiv: 1710.02185 [gr-qc].
- (2018b). "GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs". In: *arXiv e-prints*. arXiv: 1811.12907 [astro-ph.HE].
- Van de Meent, M. and N. Warburton (2018). "Fast self-forced inspirals". In: *Class. Quantum Grav.* 35, p. 144003. DOI: 10.1088/1361-6382/aac8ce.
- Vousden, W. D., W. M. Farr, and I. Mandel (2016). "Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations". In: MNRAS 455, pp. 1919–1937. DOI: 10.1093/mnras/stv2422. arXiv: 1501.05823 [astro-ph.IM].
- Wei, W. and E. A. Huerta (2019). "Gravitational Wave Denoising of Binary Black Hole Mergers with Deep Learning". In: *arXiv e-prints*. arXiv: 1901.00869 [gr-qc].
- Zeiler, M. D. and Rob Fergus (2014). "Visualizing and Understanding Convolutional Networks".
 In: Computer Vision ECCV 2014. Ed. by David Fleet et al. Springer, pp. 818–833. URL: https://doi.org/10.1007/978-3-319-10590-1_53.