

Enabling real-time multi-messenger astrophysics discoveries with deep learning

E. A. Huerta¹, Gabrielle Allen¹, Igor Andreoni², Javier M. Antelis³, Etienne Bachelet⁴, G. Bruce Berriman², Federica B. Bianco⁵, Rahul Biswas⁶, Matias Carrasco Kind¹, Kyle Chard⁷, Minsik Cho⁸, Philip S. Cowperthwaite⁹, Zachariah B. Etienne¹⁰, Maya Fishbach⁷, Francisco Forster¹¹, Daniel George¹², Tom Gibbs¹³, Matthew Graham², William Gripp¹, Robert Gruendl¹, Anushri Gupta¹, Roland Haas¹, Sarah Habib¹, Elise Jennings¹⁴, Margaret W. G. Johnson¹, Erik Katsavounidis¹⁵, Daniel S. Katz¹, Asad Khan¹, Volodymyr Kindratenko¹, William T. C. Kramer¹, Xin Liu¹, Ashish Mahabal², Zsuzsa Marka¹⁶, Kenton McHenry¹, J. M. Miller¹⁷, Claudia Moreno¹⁸, M. S. Neubauer¹, Steve Oberlin¹³, Alexander R. Olivás Jr¹⁹, Donald Petravick¹, Adam Rebei¹, Shawn Rosofsky¹, Milton Ruiz¹, Aaron Saxton¹, Bernard F. Schutz²⁰, Alex Schwing¹, Ed Seidel¹, Stuart L. Shapiro¹, Hongyu Shen¹, Yue Shen¹, Leo P. Singer²¹, Brigitta M. Sipocz²², Lunan Sun¹, John Towns¹, Antonios Tsokaros¹, Wei Wei¹, Jack Wells²³, Timothy J. Williams¹⁴, Jinjun Xiong⁸ and Zhizhen Zhao¹*

Abstract | Multi-messenger astrophysics is a fast-growing, interdisciplinary field that combines data, which vary in volume and speed of data processing, from many different instruments that probe the Universe using different cosmic messengers: electromagnetic waves, cosmic rays, gravitational waves and neutrinos. In this Expert Recommendation, we review the key challenges of real-time observations of gravitational wave sources and their electromagnetic and astroparticle counterparts, and make a number of recommendations to maximize their potential for scientific discovery. These recommendations refer to the design of scalable and computationally efficient machine learning algorithms; the cyber-infrastructure to numerically simulate astrophysical sources, and to process and interpret multi-messenger astrophysics data; the management of gravitational wave detections to trigger real-time alerts for electromagnetic and astroparticle follow-ups; a vision to harness future developments of machine learning and cyber-infrastructure resources to cope with the big-data requirements; and the need to build a community of experts to realize the goals of multi-messenger astrophysics.

Within 2 years of operation (2015–2017), the advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo detectors made unprecedented discoveries, including several gravitational wave (GW) observations of binary black hole (BBH) mergers¹. These advances, recognized with the 2017 Nobel Prize in Physics, have firmly established GW astrophysics as a field, adding GWs to the existing arsenal of cosmic messengers, namely, neutrinos, cosmic rays and electromagnetic waves. The combination of these complementary signals, known as multi-messenger astrophysics

(MMA), has enabled groundbreaking discoveries, such as the observation of supernova 1987A with neutrinos and electromagnetic waves², the binary neutron star (BNS) merger GW170817 with gravitational and electromagnetic waves^{3,4}, and the observation of the blazar TXS 0506+056 with neutrinos and γ -rays⁵.

As the worldwide network of kilometre-scale GW detectors continues to expand, and each detector gradually reaches design sensitivity, the volume of space they probe will continue to grow, increasing the number of GW observations in upcoming observing runs to

*e-mail: elihu@illinois.edu
<https://doi.org/10.1038/s42254-019-0097-4>

110–3,840 Gpc⁻³ yr⁻¹ for BNS and 9.7–101 Gpc⁻³ yr⁻¹ for BBH, at the 90% confidence level. Upper limits at the 90% confidence level for the neutron star–black hole (NSBH) merger rate is 610 Gpc⁻³ yr⁻¹ (REF.¹). Furthermore, next-generation electromagnetic surveys will substantially increase the survey area, field of view, the number of alerts for the production of new and unexpected astronomical events, volumes of image data and catalogue sizes, and computational resources to process the image data and prepare annual data releases⁶. For instance, the Large Synoptic Survey Telescope (LSST)⁷ will collect about 20 TB of data within a 24 h period. Its alert latency will be 60 s, producing up to 10 million alerts per night. With a field of view of 9.6 deg², it will survey a total area of 18,000 deg² and will complete about 1,000 visits per night. By the end of the survey, the final raw image data will be 60 PB, requiring 0.4 EB of disk storage^{6,8}. The convergence of all-sky GW observations with deep, high-cadence electromagnetic observations with next-generation astronomical facilities will provide unique opportunities for new discoveries.

To realize the full potential of MMA, the computational challenges that currently limit the scope of existing multi-messenger searches need to be addressed. First, the computational cost of low-latency GW searches based on implementations of matched-filtering (an optimal, linear, signal-processing tool to detect a signal of known or expected shape in the presence of additive noise⁹) is such that they can only probe a 4D signal manifold (parametrized by the masses of the binary components (m_1, m_2) and a configuration in which the 3D spin vectors of the objects ($\mathbf{s}_1, \mathbf{s}_2$) are aligned or anti-aligned with the orbital angular momentum, \mathbf{L} , of the binary system, that is, the only non-zero components of the 3D spin vectors are (s_1^z, s_2^z)) out of the 9D parameter space available to GW detectors ($(m_1, m_2, \mathbf{s}_1, \mathbf{s}_2)$ and the orbital eccentricity e).

Extending these algorithms to sample higher-dimensional signal manifolds is computationally unfeasible, since this would require the use of datasets of modelled waveforms at least an order of magnitude larger than existing ones^{10,11}, which already require the use of supercomputers and the Open Science Grid for core GW data analyses^{12–14}. Furthermore, algorithms that search for time-correlated transients in multiple detectors, making minimal assumptions about the morphology of GW signals¹⁵, are tailored for burst-like GW sources, and they may miss second-long GW signals with moderate signal-to-noise ratios.

Second, despite considerable development in signal-processing algorithms to identify transients in telescope images, the problem remains complex and will be exacerbated by the extremely large zoo of astrophysical transients that is expected in next-generation surveys such as LSST, which will enable the observation of up to one million objects per image. In this big-data scenario, it is apparent that rudimentary applications of machine learning techniques that have complexity $N \log(N)$, where N represents the number of points in a d -dimensional parameter space, will be too labour-intensive and computational-intensive for the real-time validation of LSST detections, classifications and alerts^{7,8}. Machine learning techniques that fall into this category include nearest-neighbour methods¹⁶, whose computationally intensive nature has motivated their implementation in hardware architectures such as graphic processing units (GPUs) and field-programmable gate arrays¹⁷.

In addition to these points, the existing cyber-infrastructure resources are oversubscribed¹⁸, and MMA needs could only be accommodated by sacrificing other science applications in existing and planned supercomputer centres, as described in REF.¹⁹. Therefore, a fundamental change is needed to secure the potential of MMA. To address these pressing issues requires the involvement of the astrophysics, data science, high-performance computing (HPC) and cyber-infrastructure communities, and the creation of efficient communication channels between the GW astrophysics and time-domain astronomy communities to facilitate plans for joint, real-time observation campaigns. In this Expert Recommendation, we make suggestions to accelerate the adoption of innovative signal-processing algorithms and computing approaches, driven by the big-data revolution, that may help address computational challenges in MMA searches.

Identification of optical counterparts

The optical counterparts of BNS and NSBH mergers are known as kilonovae or macronovae²⁰. Their emission spans the ultraviolet, optical and near-infrared bands, and encodes key information about ejected material powered by radioactive decays of r-process nuclei²¹.

There are two issues related to the detection of such optical counterparts. First, the area, typically between a few tens to hundreds of square degrees²², within which LIGO-type detectors can localize a MMA trigger includes many unrelated optical transients (for example, supernovae) whose properties are similar to those of optical counterparts of compact binary mergers. Second, optical counterparts evolve rapidly. Their fast decay rates

Author addresses

¹University of Illinois at Urbana-Champaign, Urbana, IL, USA.

²California Institute of Technology, Pasadena, CA, USA.

³Tecnologico de Monterrey, School of Engineering and Sciences, Zapopan, Mexico.

⁴Las Cumbres Observatory, Goleta, CA, USA.

⁵University of Delaware, Newark, DE, USA.

⁶The Oskar Klein Centre for Cosmoparticle Physics, Stockholm University, AlbaNova, Stockholm, Sweden.

⁷The University of Chicago, Chicago, IL, USA.

⁸IBM T.J. Watson Research Center, Yorktown Heights, New York, NY, USA.

⁹Observatories of the Carnegie Institute for Science, Pasadena, CA, USA.

¹⁰West Virginia University, Morgantown, WV, USA.

¹¹Center for Mathematical Modelling, Santiago, Chile.

¹²Google X, Mountain View, CA, USA.

¹³NVIDIA, Santa Clara, CA, USA.

¹⁴Argonne National Laboratory, Leadership Computing Facility, Lemont, IL, USA.

¹⁵Massachusetts Institute of Technology, Cambridge, MA, USA.

¹⁶Columbia University in the City of New York, New York, NY, USA.

¹⁷Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM, USA.

¹⁸Universidad de Guadalajara, Centro Universitario de Ciencias Exactas e Ingenieria, Guadalajara, Mexico.

¹⁹University of Maryland, College Park, MD, USA.

²⁰Cardiff University, Cardiff, UK.

²¹NASA Goddard Space Flight Center, Code 661, Greenbelt, MD, USA.

²²University of Washington, Seattle, WA, USA.

²³Oak Ridge National Laboratory, Oak Ridge, TN, USA.

(for example, the initial optical emission of the counterpart of GW170817 faded more than 1 mag d^{-1} , followed by a longer-lived red transient²³) require identification within a few hours of the compact binary merger, with follow-up across the electromagnetic spectrum, especially with optical spectroscopy.

The location of a GW source within the optical counterpart's host galaxy and the follow-up observations can shed light on the evolution of the progenitor system, its dynamics in the host galaxy, the production of heavy elements, and diagnostics regarding the mass and composition of the ejected material, properties of the circumstellar medium and whether jets are generated during the event²⁴. Optical counterpart identification requires a prompt response to the GW trigger to initiate a multi-filter imaging campaign covering wide sky areas using large field-of-view, deep-imaging telescopes. Alternatively, smaller field-of-view telescopes can search for optical counterparts by imaging pre-selected galaxies identified as plausible hosts. In both cases, these campaigns use automated discovery pipelines that perform image subtraction and artefact rejection. Deep learning (DL) is already a standard component of such pipelines²⁵ and has a critical role in distinguishing astrophysical sources from noise. Rapid comparison with archive data will be valuable in pinpointing the optical counterparts. Furthermore, DL is now being used as a key discovery method in new pipelines, in the optical²⁶ and at other frequencies. Future analyses may also use photometric classification methods, including DL, for the identification of transients^{27,28}. The possibility of serendipitous kilonovae discovery using this approach in wide-field surveys such as LSST have been discussed in REFS^{29,30}.

GW sources and their optical counterparts may also be used to quantify how the Universe evolves. Assuming that Einstein's general relativity is the correct description of gravity, GW observations of compact binary mergers enable a direct measurement of the luminosity distance to their source³¹. This property promotes them as standard-siren indicators, and may be used in conjunction with a catalogue of potential host galaxies, which are accompanied by their associated redshifts, apparent magnitudes and sky localizations³², to establish a redshift–distance relationship and measure the Hubble constant. GW standard measurements of the Hubble constant have already been carried out by combining GW observations from the BNS merger GW170817 and its electromagnetic counterpart, which allowed an unambiguous identification of its host galaxy^{33,34}. Multi-messenger measurements of the Hubble constant are also possible with BBH mergers that have no electromagnetic counterparts so long as GW observations are combined with photometric redshift catalogues. This has been demonstrated in the case of the BBH merger GW170814, and redshift information from the Dark Energy Survey Year 3 data³². To ensure that galaxy catalogues are as complete as possible within the redshift range within which GW detectors observe BBH mergers³⁵, DL algorithms^{36,37} are being used to classify galaxies in the Sloan Digital Sky Survey³⁸ and the Dark Energy Survey³⁹ to construct galaxy catalogues at higher redshifts using the full 6-year data gathered by the Dark

Energy Survey. These activities are essential to address similar challenges in future LSST-type surveys.

As the number of GW standard-siren measurements of the Hubble constant continue to increase in the near future, they may help understand whether the current tension between local measurements of the Hubble constant using type Ia supernovae observations⁴⁰ and early Universe measurements of the Hubble constant that combine cosmic microwave background data with the Λ cold dark matter model (where Λ is the cosmological constant) of cosmology⁴¹ is due to systematic errors or discrepancies that may usher in new discoveries^{42,43}.

Real-time detection of GWs and neutrinos

Real-time detection of GWs is essential to trigger time-sensitive searches to find their electromagnetic and/or astroparticle counterparts. Because of the computationally intensive nature and poor scalability of GW detection algorithms^{10–14}, the imminent increase in the number of GW observations owing to the continuous increase in sensitivity of ground-based GW detectors, and the fact that available computing resources for these analyses is likely to remain the same, new solutions are needed to meet the increased demand for low-latency analyses. To harness the recent developments in data science that have revolutionized big-data analyses in other science domains⁴⁴, DL prototypes have been developed for real-time detection and parameter estimation of non-spinning black holes on quasi-circular orbits that describe a 2D signal manifold (parametrized by the masses of the binary components (m_1, m_2)), both in the context of simulated LIGO noise⁴⁵ and raw advanced LIGO noise⁴⁶. These studies, which depended critically on the use of the Blue Waters supercomputer for data generation, curation and for the training and testing of DL models, have demonstrated that DL enables the detection of moderately eccentric, non-spinning BBH mergers and spin-precessing BBH mergers on quasi-circular orbits⁴⁵, higher-order waveform modes of non-spinning BBHs on eccentric orbits⁴⁷ and GW de-noising of true BBH mergers reported by the advanced LIGO and Virgo detectors^{48–51}. These studies have sparked the interest of the GW astrophysics community, leading to several developments in machine learning and DL for GW data analysis and source modelling^{52–55}.

Combining DL and HPC enables the design and training of neural network models using terabyte-size datasets of modelled waveforms, achieving state-of-the-art accuracy in actual detection scenarios. This approach reduces the training stage from weeks to minutes, enabling detailed uncertainty quantification studies to assess the robustness of DL models. The first generation of neural network models designed at the interface of DL and HPC was introduced in REF⁵⁶, in which a template bank of over ten million modelled waveforms was used to train and validate a DL model within 10 hours. Furthermore, once fully trained, this DL algorithm may be used for real-time GW parameter estimation using a single GPU. This algorithm has been applied to estimate the astrophysical parameters of the catalogue of BBH mergers reported in REF¹, finding that DL results are consistent with those produced by Bayesian analyses. Future

developments in this field concern the use of Bayesian neural networks to enable parameter estimation analyses endowed with full posterior distributions⁵⁷.

There is a pressing need to expand the scope of existing DL algorithms to cover the entire signal manifold that is available to existing GW detectors, and to enable real-time detection and characterization of BNS and NSBH mergers. These neural network models will have to be trained with significantly longer waveform signals. Since existing DL algorithms^{46,56} are capable of processing GW signals within milliseconds, one expects that even if DL algorithms for BNS observations are ~1,000 times slower than those used for BBH detection, one may still be able to extract BNS and NSBH signals from GW data in low latency.

In addition to compact binary mergers, core-collapse supernovae^{58–61} that occur within our Galaxy may be observed in electromagnetic and gravitational waves and neutrinos^{62,63}. For core-collapse supernovae produced in the Large Magellanic Cloud, a neutrino burst may be observed, though GWs may be too weak to be observed by second-generation GW detectors. For more distant core-collapse supernovae, only electromagnetic signatures may be observed⁶⁴.

As in the case of electromagnetic and gravitational observatories, DL has been explored for neutrino detection. DL is particularly suited for neutrino searches since neural networks are powerful feature extractors, and the core data challenge in neutrino detection concerns the correct categorization of high-level structures such as clusters, jets, tracks, showers and rings to tell apart signals from background noise. Recent studies provide evidence that convolutional neural networks may be used as event classifiers, achieving promising results separating signals from background noise⁶⁵. Furthermore, graph neural networks provide an ideal framework to also take into account the topology of the detector to enhance neutrino detection. These studies showcase cutting-edge applications of computer vision to address contemporary challenges in MMA and high-energy physics⁶⁶.

Numerical relativity simulations

MMA sources encompass the GW emission from BNS and NSBH mergers, the γ -ray emission from the compact object formed after merger, and the optical and infrared after-glow from the radioactive decay of r-process elements in ejected material. In the same way that numerical relativity simulations of BBH mergers have been critical for the detection and characterization of these sources with GW observations, numerical relativity simulations of BNS and NSBH mergers are critical to obtain insights into the physical processes that may lead to the production of electromagnetic and astroparticle counterparts, and to better interpret MMA observations⁶⁷. These modelling efforts do not currently benefit from DL, but recent studies have suggested the possibility to improve the efficiency and robustness of simulations, enabling the inclusion of detailed microphysics^{68–72}, and a significance increase in the speed with which partial differential equations are solved^{73,74}.

Modelling the GW emission of MMA sources requires numerical relativity⁷⁵ and the inclusion of

magnetic fields to describe the formation of the jet and ejected material⁷⁶. Modelling the r-process elements requires neutrino transport and nuclear reaction networks⁷⁷. Whereas calculations with general relativistic magneto-hydrodynamics are now routine^{78,79}, the inclusion of more sophisticated neutrino and microphysics is in its infancy. Using state-of-the-art numerical relativity software stacks, one can visualize the late inspiral phase of a BNS merger (FIG. 1a) in which the stars exchange matter due to tidal interactions. Upon merger, a hyper-massive neutron star is formed (FIG. 1b) and matter is ejected at high velocities.

Accurate modelling of MMA sources relies on solving a realistic nuclear equation of state with the density, temperature and composition dependence, to capture not only the inspiral and merger effects but also, more importantly, to study the remnant with its subsequent black-hole-disk evolution^{80,81}. In the late merger phase, the inclusion of neutrino radiation transport and realistic magnetic fields are crucial to reliably simulate the short γ -ray burst and kilonova scenarios⁸², whereas electromagnetic radiative transport is needed to provide realistic estimates of the disk geometry, mass accretion rate, and electromagnetic luminosity and spectra⁸³.

Astrophysics modelling involves a wide range of length and timescales. Often, the relevant scales can be modelled using mesh refinement⁸⁴. A promising application of machine learning is to model subgrid physics that cannot currently be modelled from first principles at computational costs⁸⁵ acceptable for parameter study simulations of turbulent motion during supernovae or neutron star collision⁸⁶. Subgrid models tuned by a small number of high-resolution simulations avoid these costs⁸⁷.

Cyber-infrastructure requirements

Designing DL algorithms for real-time MMA analyses includes data curation, model training and inference, and result analysis. Training is the most computationally intensive step: it requires a significant amount of inter-process/thread communication and sequential calculations. Mature models can be trained using many central processing units (CPUs)/GPUs (horizontal scaling) with large batch sizes. New models, however, require a rounded balance of horizontal and vertical scaling, for example, accelerated CPU/GPU processing. In general, inference is a parallel process: each input can be processed completely independently of any other input. Depending on the vertical scale of computation, the bottlenecks are often data retrieval, input/output latency and bandwidth. Once fully trained, DL algorithms can process large datasets in real time^{45,46,56}.

To enable real-time MMA analyses, the cyber-infrastructure must support interactive access to large collections of disparate datasets, and the ability to apply DL algorithms to subsets of the data. These requirements suggest a low-latency interconnect, HPC-like system, with adequately provisioned data and computing resources.

At the time of writing, the Extreme Science and Engineering Discovery Environment (XSEDE) provides 2.8 PFlop s⁻¹ of single-precision GPU computing power through the Bridges and Comet clusters, both of which are nearing retirement. New resources are expected

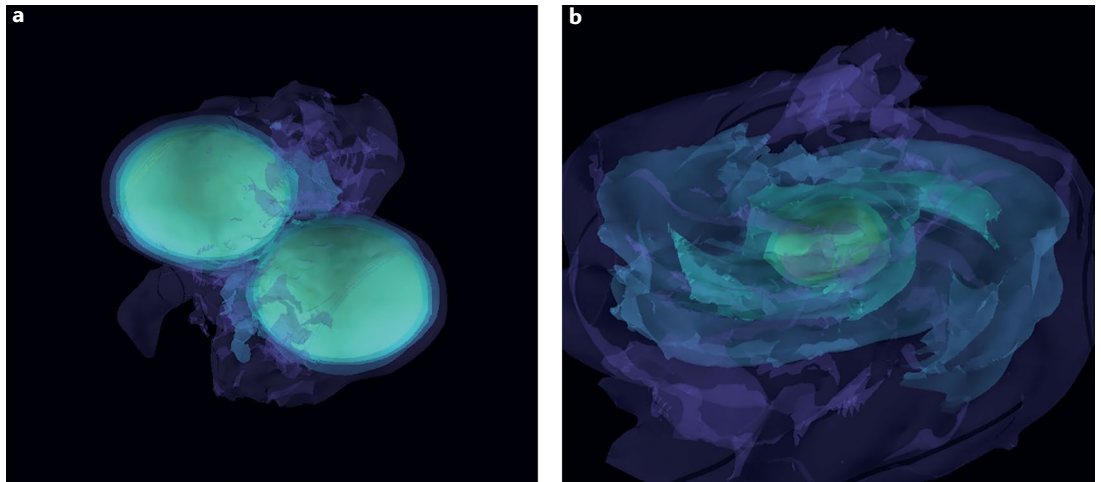


Fig. 1 | Visualization of the results of a numerical relativity simulation of two neutron stars before and after a merger. a | Two neutron stars moving at a fraction of the speed of light shortly before they merge. Tidal forces start to tear apart the neutron stars and fling material away from them. Neutron star mergers may power short γ -ray bursts, as observationally confirmed with the binary neutron star merger GW170817 (REF. 3). **b** | The neutron star collision forms a hypermassive neutron star, which is surrounded by neutron-rich ejecta. The radioactive decay of the neutron-rich nuclei in the ejecta may produce a short-lived optical/near-infrared weak supernova-like signal, known as kilonovae²⁰. The observed features in GW170817 are broadly consistent with a kilonova model, indicating that r -process elements with a mass of $0.05 M_{\odot}$ were synthesized in that event^{21,105}. This numerical relativity simulation was performed using the Einstein Toolkit¹⁰⁶, with the addition of the WhiskyTHC¹⁰⁷ hydrodynamics code, and the CTGamma space-time evolution code. The physics in this simulation includes the full general relativity space-time evolution, a finite-temperature nuclear equation of state and a neutrino evolution model. The simulation was performed and visualized on the Extreme Science and Engineering Discovery Environment (XSEDE) Stampede2 supercomputer.

to become available on the 2020 time frame, but it is unknown what the configuration of these resources will be, and whether they will effectively support the MMA computing needs. The National Science Foundation (NSF)-funded Frontera supercomputer will provide only 8 PFlop s^{-1} of 32-bit precision computational power through a small GPU section of the system, which will be readily consumed by basic MMA machine learning applications. Other NSF-funded investments include the DL project at the National Centre for Supercomputing Applications, which consists of 16 IBM Power9 servers with four NVIDIA V100 GPUs, and that supports TensorFlow and other DL frameworks in a container-based environment. The Department of Energy (DOE) supercomputing platforms Summit and Titan at Oak Ridge National Laboratory are equipped with over 27,000 NVIDIA Volta V100 GPUs and over 18,000 K20X GPUs, respectively. Piz Daint, Europe's fastest supercomputer, is equipped with 5,320 hybrid compute nodes, each equipped with two Intel Xeon E5-2690 v3 processors and a single NVIDIA Tesla P100. The system also has 1,431 GPU-less XC40 nodes, each of which has two Intel Xeon E5-2695 v4 processors. According to EuroHPC, a joint initiative between the European Union and European countries to spearhead innovation for the European HPC ecosystem, the computation capabilities in the European Union do not match the current demand for computing and data needs from the academic and industry sectors in Europe⁸⁸. Furthermore, the fastest European supercomputers depend on non-European technology. To address these problems and develop a strategic plan for innovation and competitiveness, EuroHPC aims to buy

and install two pre-exascale machines by 2020, and three or four additional petascale machines. It is expected that by 2023, Europe will have two exascale machines and one post-exascale machine that may support the first hybrid HPC-quantum computing infrastructure in Europe.

Across the spectrum of computational sciences, the supply of large-scale computing resources does not meet the demand^{18,19,88}. Funding-agency research grants for computational projects support people, but often do not guarantee computer resources. Deployment and operation of sufficient agency resources to cover the majority of the predicted demand over the next few years might best be served by organizing around individual resources scaled around $50\text{--}100 \text{ PFlops s}^{-1}$ double-precision, $10\text{--}100$ petabytes storage, and external network data intake capacity sufficient for the analysis of the expected observational data. The estimated demand will probably require multiple instances of resources/centres at this scale. These should provide common HPC, machine learning/DL and data analytics software environments.

The landscape of existing and planned NSF-funded HPC cyber-infrastructure indicates a pressing need to expand these resources to meet the MMA needs^{13,89}. Given the much broader applicability of machine learning/DL approaches, the agencies should invest in the resources and services necessary to enable the success of MMA and other community efforts reliant on these approaches. These resources include appropriate storage and data-sharing environments that can support multidisciplinary collaborations, along with an adequate networking infrastructure to support both bulk data transport and rapid notifications. We also advocate

for the exploration of newer technologies, such as tensor processing units and quantum computers, and their applicability to these challenges.

Future needs of MMA

MMA heavily relies on sharing real-time alerts among various instruments. These alerts would be prioritized by every instrument, given their observing capabilities and science goals. A common tool to help identify potential host galaxies of GW sources is highly desirable⁹⁰, as well as coordination between groups to determine which sources may be optimally followed up with wide-field, all-sky survey instruments, or by targeting opportunity observations with narrow field-of-view telescopes in which the number of observable sources is significantly limited⁹¹. DL classifiers could automate these time-sensitive tasks in existing services such as the NASA/Infrared Processing and Analysis Center Extragalactic Database (NED) Gravitational Wave Follow-up (GWF) Service⁹² and the Astrophysical Multi-Messenger Observatory Network (AMON)⁹³.

The engagement of the next generation of telescopes (LSST, Thirty Meter Telescope, Extremely Large Telescope and so on) in MMA should be defined, as should observing policies for existing large telescopes: the observing strategy, mechanism and software infrastructure should be flexible enough to allow a rapid reaction to alerts regarding important potential targets. Pre-defining the amount of time and resources each next-generation telescope and survey will spend on MMA follow-ups, involving the transient-sky and static-sky science communities, would avoid potential conflicts and waste of resources^{7,94–96}. Archives for large telescopes should ideally support real-time ingestion of freshly acquired data and, as much as possible, make MMA follow-up data public immediately after receipt.

Ideally, a common, public database would be built to centralize all information available for every recorded event. This will minimize the duplication of common tasks (such as data collection, galaxy cross-matching, ranking and so on) among follow-up teams. The centralized database could be coupled with classification algorithms and catalogues cross-matching. In this sense, the infrastructure of the LSST data management system and the alert brokers under development are good examples^{97,98}. The automation of telescopes is an important part of maximizing the number of observed MMA events (for example, the Astronomical Event Observatory Network programme will provide programmable access to a number of telescopes⁹⁹).

Finally, it is crucial that funding agencies explore the best way to manage competitive proposals pursuing the same scientific goals. Given the scientific value of some MMA events, several teams are likely to simultaneously request the same observations from the same facilities, creating unneeded duplication and the potential for conflict. Although competition can be valuable on a scientific level, on the infrastructure level it often leads to inefficiency. Given the increasing numbers of MMA events available owing to improving instrument sensitivities, this inefficiency could become a significant drain on the limited observing resources.

Whereas different approaches may be used to manage the technical issues relevant to these problems, the scientific community may still need to reach consensus on the common requirements and science priorities (for instance, within the Astro2020: Decadal Survey on Astronomy and Astrophysics¹⁰⁰) that would allow the funding agencies to plan a common infrastructure and provide guidance to both the review and decision process regarding competing usage proposals.

Community building

By the very nature of the field, the science goals of MMA cannot be realized by any one individual or small team. Community efforts are essential to building the necessary software to manage the needs of large-scale collaborations¹⁰¹. Successful examples in astronomy (for example, the Astropy community project, the yt community project and the LSST–Dark Energy Science Collaboration) can act as models for building an MMA software community.

Other data-intensive scientific domains, such as high-energy physics (HEP), share similar challenges in building and sustaining a cohesive community, workforce training, development and advancement, access and delivery of large amounts of data, low-latency data analysis and using machine learning techniques. In HEP, 18 workshops in two years engaged key national and international partners from HEP, computer science, industry and data science to generate over eight community position papers, including a software institute strategic plan¹⁰² and a community white paper¹⁰³ as a roadmap for HEP software and computing research and development over the next decade. The MMA community should consider what can be learned from this intensive community effort and use the transferable solutions.

A similar effort was already set in motion with the NSF-funded project Community planning for Scalable Cyber-infrastructure to support Multi-Messenger Astrophysics (SCiMMA)¹⁰⁴, which seeks to identify the key questions and cyber-infrastructure projects required by the MMA community to take full advantage of current and imminent next-generation facilities. SCiMMA held a number of workshops to bring community members together and is now writing white papers on various subtopics, including systems development; data management, communication and collaboration; analysis, inference and machine learning; modelling and theory; education, training and workforce development; and management. SCiMMA will also develop a strategic plan for a scalable cyber-infrastructure institute for MMA laying out its proposed mission, identifying the highest-priority areas for cyber-infrastructure research and development for the US-based MMA community, and presenting a strategy for managing and evolving a set of services for the broader community.

Conclusions

For the challenges discussed in the previous section, we make a number of recommendations, summarized in BOX 1. We identified two pillars that require immediate attention to fully realize the MMA science programme: first, increase the speed and depth of signal-processing

Box 1 | Summary of recommendations to realize the science goals of multi-messenger astrophysics

Numerical relativity

- Improve coupling and cross-talk between numerical relativity and modelling tools and build a pipeline for predicting observables
- Accelerate community efforts to develop and release open-source versions of modelling codes, especially microphysics and transport packages
- Explore publicly funded opportunities to develop numerical relativity software critical for MMA interpretations such as through the National Science Foundation (NSF) or the Department of Energy (DOE) office of science

Gravitational wave astrophysics and neutrino physics

- Accelerate the convergence of deep learning (DL) and high-performance computing (HPC) to reduce training stage from weeks to minutes of DL models that cover the complete signal manifold of black hole mergers
- Design Bayesian neural network models for gravitational wave (GW) parameter estimation covering the full signal manifold of black hole mergers
- Design neural network models for real-time detection of compact binaries involving neutron stars and core-collapse supernovae
- Design and deploy production-scale neural network classifiers to categorize particle interactions in neutrino detectors
- Accelerate the development of graph neural networks that are tailored for the specific topology of existing detectors

Electromagnetic surveys

- Continue the development and adoption of software, including the robotization of telescopes, to facilitate the rapid, automated, coordinated response to GW triggers from multiple telescopes observing across the electromagnetic (EM) spectrum and leverage machine learning for decision-making and optimization
- Continue the development of algorithms and software for the classification and identification of EM counterparts of GW events, including the integration of DL in event selection and classification from images and hybrid data sources (direct integration of images and catalogue data)
- Continue the development of simulations and emulations in astrophysics modelling to fully explore the parameter space of the EM counterparts of GW events
- Create and facilitate automated access to comprehensive catalogues of archival and new observations to be used in event selection and classification

- Create communication channels between concurrently observing facilities to share plans for observing targets

Cyber-infrastructure for DL

- Develop public training and test datasets, and a mock system infrastructure to test trained models
- Funding agencies (including NSF and DOE in the United States) should agree on a split of responsibility for the common cyber-infrastructure needed to support all projects and researchers, rather than individual projects
- Funding agencies should ensure that the majority of their own computational research funding also covers the proposed computing resources (through their own systems and inter-agency agreements)
- Maintaining a current and consistent DL and workflow software stack are essential computational resources for theoretical research and practical applications at scale, and provide critical tools for community training and education
- Development of open-source software products and scalable and efficient DL algorithms that maximize use of HPC resources are needed
- Integrating scalable machine learning with developments in computational methods and HPC platforms has the potential to benefit many key areas of interest to funding agencies

Community building

- Develop trust between projects, individual researchers and these two groups; then formalize that trust into agreements about data collection, data access, credit and authorship, and so on
- Build on the example of successfully governed, open-source community projects to develop community-maintained software systems
- Organize workshops to train users and encourage participation
- Develop mechanisms to provide incentives (opportunities to represent scientific communities at conferences with invited talks and keynotes), rewards (including fellowships, and support to attend and participate in prestigious summer and winter schools) and recognition (having leadership roles in scientific communities) for success
- Develop rules and guidelines for data sharing and publication to assure ethical behaviour and avoid waste of EM observational resources

DL; deep learning; DOE, Department of Energy; EM, electromagnetic; GW, gravitational wave; HPC, high-performance computing; NSF, National Science Foundation.

algorithms for real-time detection of GWs and their astroparticle and electromagnetic counterparts; and second, identify the cyber-infrastructure resources to simulate and search for MMA events in the ever increasing and disparate datasets. Furthermore, DL methods to accelerate the solution of partial differential equations could be used to greatly increase the speed, accuracy and robustness of numerical relativity simulations of MMA sources^{73,74}.

There is concern regarding the availability of future cyber-infrastructure facilities for MMA data analytics. Furthermore, whereas DOE HPC platforms are particularly suited for DL analytics, they are highly over-subscribed. Therefore, we recommend reaching out to diverse advanced cyber-infrastructure centres in the United States (Summit, Sierra, Frontera, Bridges-2, Theta, Lassen and so on) and Europe (Piz Daint, SuperMUC-NG and so on) to obtain computational resources for DL at scale. This is critical because the

convergence of DL with HPC is essential to train deeper and more accurate neural network models with terabyte-size training datasets. This approach is essential to ensure that DL algorithms characterize the sources' parameter space that is available to observatories. Once these models are fully trained, they can be used for real-time MMA searches using minimal computational resources.

We also recognize the need for policy making regarding data acquisition and data sharing at astronomical observatories and for facilitating the interaction among the different MMA subfields. The MMA community should put in place mechanisms to reward its members, allow them to develop and provide career paths for its members beyond academia. We suggest collaborating with industry partners who, in addition to co-funding data science workshops and bootcamps for MMA researchers, could also recruit members of this community.

Published online 3 October 2019

1. Abbott, B. P. 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 (2019).
2. Arnett, W. D., Bahcall, J. N., Kirshner, R. P. & Woosley, S. E. Supernova 1987A. *Annu. Rev. Astron. Astrophys.* **27**, 629–700 (1989).
3. Abbott, B. P. et al. GW170817: observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.* **119**, 161101 (2017).
4. Abbott, B. P. et al. Estimating the contribution of dynamical ejecta in the kilonova associated with GW170817. *Astrophys. J. Lett.* **850**, L39 (2017).
5. IceCube Collaboration Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert. *Science* **361**, 147–151 (2018).
6. Large Synoptic Survey Telescope LSST system and survey key numbers. *LSST* <https://www.lsst.org/scientists/keynumbers> (2018).
7. Abell, P. A. et al. LSST Science Book, version 2.0. Preprint at <https://arxiv.org/abs/0912.0201> (2009).
8. Robertson, B. E. et al. Galaxy formation and evolution science in the era of the Large Synoptic Survey Telescope. *Nat. Rev. Phys.* **1**, 450–462 (2019).
9. Owen, B. J. & Sathyaprakash, B. S. Matched filtering of gravitational waves from inspiraling compact binaries: computational cost and template placement. *Phys. Rev. D* **60**, 022002 (1999).
10. Harry, I., Privitera, S., Bohé, A. & Buonanno, A. Searching for gravitational waves from compact binaries with precessing spins. *Phys. Rev. D* **94**, 024012 (2016).
11. Huerta, E. A. et al. Complete waveform model for compact binaries on eccentric orbits. *Phys. Rev. D* **95**, 024038 (2017).
12. Huerta, E. A. et al. BOSS-LDG: a novel computational framework that brings together blue waters, open science grid, shifter and the LIGO data grid to accelerate gravitational wave discovery. In *2017 IEEE 13th International Conference on e-Science 335–344* (IEEE, 2017).
13. Huerta, E. A., Haas, R., Jha, S., Neubauer, M. & Katz, D. S. Supporting high-performance and high-throughput computing for experimental science. *Comput. Softw. Big Sci.* **3**, 5 (2019).
14. Weitzel, D. et al. Data access for LIGO on the OSG. In *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact 24*, 1–6 (PEARC, 2017).
15. Abbott, B. P. et al. Observing gravitational-wave transient GW150914 with minimal assumptions. *Phys. Rev. D* **93**, 122004 (2016).
16. Jones, P. W., Osipov, A. & Rokhlin, V. Randomized approximate nearest neighbors algorithm. *Proc. Natl Acad. Sci. USA* **108**, 15679–15686 (2011).
17. Liang, S., Liu, Y., Wang, C. & Jian, L. Design and evaluation of a parallel k-nearest neighbor algorithm on CUDA-enabled GPU. In *2010 IEEE 2nd Symposium on Web Society 53–60* (IEEE, 2010).
18. Andre, J. C. et al. Big data and extreme-scale computing: pathways to convergence toward a shaping strategy for a future software and data ecosystem for scientific inquiry. *Int. J. High Perform. Comput. Appl.* **32**, 435–479 (2018).
19. Engineering National Academies of Sciences and Medicine *Future Directions for NSF Advanced Computing Infrastructure Support. U.S. Science and Engineering in 2017–2020* (The National Academies Press, 2016).
20. Metzger, B. D. & Berger, E. What is the most promising electromagnetic counterpart of a neutron star binary merger? *Astrophys. J.* **746**, 48 (2012).
21. Siegel, D. M. & Metzger, B. D. Three-dimensional grmhd simulations of neutrino-cooled accretion disks from neutron star mergers. *Astrophys. J.* **858**, 52 (2018).
22. Abbott, B. P. et al. Prospects for observing and localizing gravitational-wave transients with advanced LIGO and advanced Virgo. *Living Rev. Relativ.* **21**, 3 (2018).
23. Drout M. R. et al. Light curves of the neutron star merger GW170817/SSS17a: implications for r-process nucleosynthesis. *Science* **358**, 1570–1574 (2017).
24. Mooley, K. P. et al. A mildly relativistic wide-angle outflow in the neutron-star merger event GW170817. *Nature* **554**, 207–210 (2018).
25. Andreoni, I. et al. Mary, a pipeline to aid discovery of optical transients. *Publ. Astron. Soc. Aust.* **34**, e037 (2017).
26. Sedaghat, N. & Mahabal, A. Effective image differencing with convolutional neural networks for real-time transient hunting. *Mon. Not. R. Astron. Soc.* **476**, 5365–5376 (2018).
27. Jones, D. O. et al. Measuring dark energy properties with photometrically classified Pan-STARRS supernovae. II. Cosmological parameters. *Astrophys. J.* **857**, 51 (2018).
28. Kessler, R. et al. Results from the supernova photometric classification challenge. *Publ. Astron. Soc. Pac.* **122**, 1415–4131 (2010).
29. Scolnic, D. et al. How many kilonovae can be found in past, present, and future survey data sets? *Astrophys. J. Lett.* **852**, L3 (2018).
30. Setzer, C. N. et al. Serendipitous discoveries of kilonovae in the LSST main survey: maximising detections of sub-threshold gravitational wave events. *Mon. Not. R. Astron. Soc.* **485**, 4260–4273 (2019).
31. Schutz, B. F. Determining the Hubble constant from gravitational wave observations. *Nature* **323**, 310–311 (1986).
32. Soares-Santos, M. et al. First measurement of the Hubble constant from a dark standard siren using the Dark Energy Survey galaxies and the LIGO/Virgo binary–black-hole merger GW170814. *Astrophys. J.* **876**, L7 (2019).
33. Abbott, B. P. et al. A gravitational-wave standard siren measurement of the Hubble constant. *Nature* **551**, 85–88 (2017).
34. Cowperthwaite, P. S. et al. The electromagnetic counterpart of the binary neutron star merger LIGO/Virgo GW170817. II. UV, optical, and near-infrared light curves and comparison to kilonova models. *Astrophys. J. Lett.* **848**, L17 (2017).
35. Fishbach, M. et al. A standard siren measurement of the Hubble constant from GW170817 without the electromagnetic counterpart. *Astrophys. J. Lett.* **871**, L13 (2019).
36. D. Sánchez, H., Huertas-Company, M., Bernardi, M., Tuccillo, D. & Fischer, J. L. Improving galaxy morphologies for SDSS with deep learning. *Mon. Not. R. Astron. Soc.* **476**, 3661–3676 (2018).
37. Khan, A. et al. Deep learning at scale for the construction of galaxy catalogs in the Dark Energy Survey. *Phys. Lett. B* **795**, 248–258 (2019).
38. Eisenstein, D. J. et al. SDSS-III: massive spectroscopic surveys of the distant Universe, the Milky Way, and extra-solar planetary systems. *Astron. J.* **142**, 72 (2011).
39. Dark Energy Survey Collaboration et al. The Dark Energy Survey: more than dark energy — an overview. *Mon. Not. R. Astron. Soc.* **460**, 1270–1299 (2016).
40. Riess, A. G., Casertano, S., Yuan, W., Macri, L. M. & Scolnic, D. Large Magellanic Cloud cepheid standards provide a 1% foundation for the determination of the Hubble constant and stronger evidence for physics beyond Λ CDM. *Astrophys. J.* **876**, 85 (2019).
41. Aghanim, N. et al. Planck 2018 results. VI. Cosmological parameters. Preprint at <https://arxiv.org/abs/1807.06209> (2018).
42. Freedman, W. L. Cosmology at a crossroads. *Nat. Astron.* **1**, 0121 (2017).
43. Poulin, V., Smith, T. L., Karwal, T. & Kamionkowski, M. Early dark energy can resolve the Hubble tension. *Phys. Rev. Lett.* **122**, 221301 (2019).
44. Lecun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
45. George, D. & Huerta, E. A. Deep neural networks to enable real-time multimessenger astrophysics. *Phys. Rev. D* **97**, 044039 (2018).
46. George, D. & Huerta, E. A. Deep Learning for real-time gravitational wave detection and parameter estimation: results with advanced LIGO data. *Phys. Lett. B* **778**, 64–70 (2018).
47. Rebei, A. et al. Fusing numerical relativity and deep learning to detect higher-order multipole waveforms from eccentric binary black hole mergers. *Phys. Rev. D* **100**, 044025 (2019).
48. George, D., Shen, H. & Huerta, E. A. Classification and unsupervised clustering of ligo data with deep transfer learning. *Phys. Rev. D* **97**, 101501 (2018).
49. Shen, H., George, D., Huerta, E. A. & Zhao, Z. Denoising gravitational waves with enhanced deep recurrent denoising auto-encoders. In *2019 IEEE International Conference on Acoustics, Speech and Signal Processing 3257–3241* (IEEE, 2019).
50. Shen, H., George, D., Huerta, E. A. & Zhao, Z. Denoising gravitational waves using deep learning with recurrent denoising autoencoders. Preprint at <https://arxiv.org/abs/1711.09919> (2017).
51. Wei, W. & Huerta, E. A. Gravitational wave denoising of binary black hole mergers with deep learning. Preprint at <https://arxiv.org/abs/1901.00869> (2019).
52. Chua, A. J. K., Galley, C. R. & Vallisneri, M. ROMAN: Reduced-order modeling with artificial neurons. *Phys. Rev. Lett.* **122**, 211101 (2019).
53. Dreissigacker, C., Sharma, R., Messenger, C. & Prix, R. Deep-learning continuous gravitational waves. *Phys. Rev. D* **100**, 044009 (2019).
54. Gabbard, H., Williams, M., Hayes, F. & Messenger, C. Matching matched filtering with deep networks for gravitational-wave astronomy. *Phys. Rev. Lett.* **120**, 141103 (2018).
55. Nakano, H. et al. Comparison of various methods to extract ringdown frequency from gravitational wave data. *Phys. Rev. D* **99**, 124032 (2019).
56. Shen, H., Huerta, E. A. & Zhao, Z. Deep learning at scale for gravitational wave parameter estimation of binary black hole mergers. Preprint at <https://arxiv.org/abs/1903.01998> (2019).
57. Springenberg, J. T., Klein, A., Falkner, S. & Hutter, F. Bayesian optimization with robust Bayesian neural networks. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)* (eds Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I. & Garnett, R.) 4134–4142 (Curran Associates, 2016).
58. Burrows, A., Hayes, J. & Fryxell, B. A. On the nature of core-collapse supernova explosions. *Astrophys. J.* **450**, 830 (1995).
59. Burrows, A., Radice, D. & Vartanyan, D. Three-dimensional supernova explosion simulations of 9-, 10-, 11-, 12-, and 13- M_{\odot} stars. *Mon. Not. R. Astron. Soc.* **485**, 3153–3168 (2019).
60. Mösta, P. et al. r-process nucleosynthesis from three-dimensional magnetorotational core-collapse supernovae. *Astrophys. J.* **864**, 171 (2018).
61. Radice, D., Morozova, V., Burrows, A., Vartanyan, D. & Nagakura, H. Characterizing the gravitational wave signal from core-collapse supernovae. *Astrophys. J. Lett.* **876**, L9 (2019).
62. Janka, H.-T., Melson, T. & Summa, A. Physics of core-collapse supernovae in three dimensions: a sneak preview. *Annu. Rev. Nucl. Part. Sci.* **66**, 341–375 (2016).
63. Woosley, S. & Janka, T. The physics of core-collapse supernovae. *Nat. Phys.* **1**, 147–154 (2005).
64. Gossan, S. E. et al. Observing gravitational waves from core-collapse supernovae in the advanced detector era. *Phys. Rev. D* **93**, 042002 (2016).
65. Aurisano, A. et al. A convolutional neural network neutrino event classifier. *J. Instrum.* **11**, P09001 (2016).
66. Choma, N. et al. Graph neural networks for icecube signal classification. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA) 386–391* (IEEE, 2018).
67. Hinderer, T. et al. Discerning the binary neutron star or neutron star–black hole nature of GW170817 with gravitational wave and electromagnetic measurements. Preprint at <https://arxiv.org/abs/1808.03836> (2018).
68. Kim, B. et al. Deep fluids: a generative network for parameterized fluid simulations. *Comput. Graph Forum* **38**, 59-70 (2019).
69. Ling, J., Kurzwski, A. & Templeton, J. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166 (2016).
70. Maulik, R., San, O., Rasheed, A. & Vedula, P. Subgrid modelling for two-dimensional turbulence using neural networks. *J. Fluid Mech.* **858**, 122–144 (2019).
71. Viganò, D. & Palenzuela, C. Fitting of extended sub-grid scale models in compressible turbulent MHD. Preprint at <https://arxiv.org/abs/1904.04099> (2019).
72. Xie, C., Wang, J., Li, K. & Ma, C. Artificial neural network approach to large-eddy simulation of compressible isotropic turbulence. *Phys. Rev. E* **99**, 053113 (2019).
73. Berg, J. & Nyström, K. A unified deep artificial neural network approach to partial differential equations in complex geometries. *Neurocomputing* **317**, 28–41 (2018).
74. Weinan, E., Han, J. & Jentzen, A. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat.* **5**, 349–380 (2017).
75. Duez, M. D. & Zlochower, Y. Numerical relativity of compact binaries in the 21st century. *Rep. Prog. Phys.* **82**, 016902 (2019).

76. Baiotti, L. & Rezzolla, L. Binary neutron star mergers: a review of Einstein's richest laboratory. *Rep. Prog. Phys.* **80**, 096901 (2017).
77. Lippuner, J. & Roberts, L. F. Skynet: a modular nuclear reaction network library. *Astrophys. J. Suppl. Ser.* **233**, 18 (2017).
78. Paschalidis, V., Ruiz, M. & Shapiro, S. L. Relativistic simulations of black hole–neutron star coalescence: the jet emerges. *Astrophys. J.* **806**, L14 (2015).
79. Ruiz, M., Lang, R. N., Paschalidis, V. & Shapiro, S. L. Binary neutron star mergers: a jet engine for short gamma-ray bursts. *Astrophys. J.* **824**, L6 (2016).
80. Fernández, R. et al. Long-term GRMHD simulations of neutron star merger accretion disks: implications for electromagnetic counterparts. *Mon. Not. R. Astron. Soc.* **482**, 3373 (2019).
81. Nouri, F. H. et al. Evolution of the magnetized, neutrino-cooled accretion disk in the aftermath of a black hole–neutron star binary merger. *Phys. Rev. D* **97**, 083014 (2018).
82. Radice, D. et al. Binary neutron star mergers: mass ejection, electromagnetic counterparts, and nucleosynthesis. *Astrophys. J.* **869**, 130 (2018).
83. Kasen, D., Badnell, N. R. & Barnes, J. Opacities and spectra of the r-process ejecta from neutron star mergers. *Astrophys. J.* **774**, 25 (2013).
84. Berger, M. J. & Colella, P. Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.* **82**, 64–84 (1989).
85. Chen, T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. Neural ordinary differential equations. *Advances in Neural Information Processing Systems* **31**, 6571–6583 (2018).
86. Radice, D. et al. Neutrino-driven convection in core-collapse supernovae: high-resolution simulations. *Astrophys. J.* **820**, 76 (2016).
87. Giacomazzo, B., Zrake, J., Duffell, P., MacFadyen, A. I. & Perna, R. Producing magnetar magnetic fields in the merger of binary neutron stars. *Astrophys. J.* **809**, 39 (2015).
88. EuroHPC Leading the way in the European supercomputing. *EuroHPC* <https://eurohpc-ju.europa.eu/#inline-nav-1> (2018).
89. Huerta, E. A. et al. BOSS-LDG: a novel computational framework that brings together blue waters, open science grid, shifter and the LIGO data grid to accelerate gravitational wave discovery. In *2017 IEEE 13th International Conference on e-Science* 335–344 (IEEE, 2017).
90. Arcavi, I. et al. Optical follow-up of gravitational-wave events with las cumbres observatory. *Astrophys. J.* **848**, L33 (2017).
91. Coughlin, M. W. et al. Optimizing searches for electromagnetic counterparts of gravitational wave triggers. *Mon. Not. R. Astron. Soc.* **478**, 692–702 (2018).
92. California Institute of Technology NED gravitational wave follow-up service. *NED* <https://ned.ipac.caltech.edu/gwf/overview> (2019).
93. Pennsylvania State University Astrophysical multimessenger observatory network. *AMON* <https://www.amon.psu.edu/amon-system/> (2019).
94. Cowperthwaite, P. S. et al. *Astro 2020 Science White Paper: Joint Gravitational Wave and Electromagnetic Astronomy with LIGO and LSST in the 2020's*. Preprint at <https://arxiv.org/abs/1904.02718> (2019).
95. Marshall, P. et al. *Science-Driven Optimization of the LSST Observing Strategy*. Preprint at <https://arxiv.org/abs/1708.04058>, 10.5281/zenodo.842713 (2017).
96. A. Kinney et al. The W. M. Keck Observatory Scientific Strategic Plan. *Keck Observers' Newsletter* <https://www2.keck.hawaii.edu/inst/newsletters/Vol20/scientificstrategicplan.pdf> (2016).
97. Narayan, G. et al. Machine learning-based brokers for real-time classification of the LSST alert stream. *Astrophys. J. Suppl.* **236**, 9 (2018).
98. Smith, K. W. et al. Lasair: the transient alert broker for LSST:UK. *Res. Notes AAS* **3**, 26 (2019).
99. AEON Team Astronomical Event Observatory Network *NOAO* <http://ast.noao.edu/data/aeon> (2018).
100. The National Academies of Sciences, Engineering, and Medicine. The decadal survey on astronomy and astrophysics (astro2020). *The National Academies of Sciences, Engineering, and Medicine* <http://sites.nationalacademies.org/DEPS/astro2020/index.htm> (2019).
101. Katz, D. S. et al. Community organizations: changing the culture in which research software is developed and sustained. *Comp. Sci. Eng.* **21**, 8–24 (2019).
102. Elmer, P., Neubauer, M. & Sokoloff, M. D. Strategic plan for a Scientific Software Innovation Institute (S2I2) for high energy physics. Preprint at <https://arxiv.org/abs/1712.06592> (2017).
103. Albrecht, J. et al. A roadmap for HEP software and computing R&D for the 2020s. *Comput. Softw. Big Sci.* **3**, 7 (2019).
104. Allen, G. et al. Multi-messenger astrophysics: harnessing the data revolution. Preprint at <https://arxiv.org/abs/1807.04780> (2018).
105. Hotokezaka, K., Beniamini, P. & Piran, T. Neutron star mergers as sites of r-process nucleosynthesis and short gamma-ray bursts. *Int. J. Mod. Phys. D* **27**, 1842005 (2018).
106. Löffler, F. et al. The Einstein toolkit: a community computational infrastructure for relativistic astrophysics. *Class. Quantum Gravity* **29**, 115001 (2012).
107. Radice, D. & Rezzolla, L. THC: a new high-order finite-difference high-resolution shock-capturing code for special-relativistic hydrodynamics. *Astron. Astrophys.* **547**, A26 (2012).

Acknowledgements

The authors gratefully acknowledge support from NVIDIA, Argonne Leadership Computing Facility, Oak Ridge Leadership Computing Facility, and the National Science Foundation through grant NSF-1848815. Artwork in this manuscript was supported in part by the National Science Foundation through grants ACI-1238993, NSF-1550514 and TG-PHY160053.

Author contributions

E.A.H. led and coordinated the writing of this Expert Recommendation. All authors contributed to developing the ideas, and writing and reviewing this manuscript. S.R. produced the artwork in figure 1.

Competing interests

The authors declare no competing interests.

Peer review information

Nature Reviews Physics thanks Brant Robertson, Viviana Acquaviva and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Publisher's note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.