

Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <http://orca.cf.ac.uk/127213/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Huerta, E. A., Allen, Gabrielle, Andreoni, Igor, Antelis, Javier M., Bachelet, Etienne, Berriman, G. Bruce, Bianco, Federica B., Biswas, Rahul, Carrasco Kind, Matias, Chard, Kyle, Cho, Minsik, Cowperthwaite, Philip S., Etienne, Zachariah B., Fishbach, Maya, Forster, Francisco, George, Daniel, Gibbs, Tom, Graham, Matthew, Gropp, William, Gruendl, Robert, Gupta, Anushri, Haas, Roland, Habib, Sarah, Jennings, Elise, Johnson, Margaret W. G., Katsavounidis, Erik, Katz, Daniel S., Khan, Asad, Kindratenko, Volodymyr, Kramer, William T. C., Liu, Xin, Mahabal, Ashish, Marka, Zsuzsa, McHenry, Kenton, Miller, J. M., Moreno, Claudia, Neubauer, M. S., Oberlin, Steve, Olivas, Alexander R., Petravick, Donald, Rebei, Adam, Rosofsky, Shawn, Ruiz, Milton, Saxton, Aaron, Schutz, Bernard F., Schwing, Alex, Seidel, Ed, Shapiro, Stuart L., Shen, Hongyu, Shen, Yue, Singer, Leo P., Sipocz, Brigitta M., Sun, Lunan, Towns, John, Tsokaros, Antonios, Wei, Wei, Wells, Jack, Williams, Timothy J., Xiong, Jinjun and Zhao, Zhizhen 2019. Enabling real-time multi-messenger astrophysics discoveries with deep learning. *Nature Reviews Physics* 1 (10) , pp. 600-608. 10.1038/s42254-019-0097-4 file

Publishers page: <http://dx.doi.org/10.1038/s42254-019-0097-4> <<http://dx.doi.org/10.1038/s42254-019-0097-4>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



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

E. A. Huerta^{1,*}, 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 Förster¹¹, Daniel George¹², Tom Gibbs¹³, Matthew Graham², William Gropp¹, 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. Olivas 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. Sipócz²², Lunan Sun¹, John Towns¹, Antonios Tsokaros¹, Wei Wei¹, Jack Wells²³, Timothy J. Williams¹⁴, Jinjun Xiong⁸, and Zhizhen Zhao¹

* e-mail: elihu@illinois.edu

¹University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA

²California Institute of Technology, 1200 E California Blvd, Pasadena, California 91125, USA

³Tecnologico de Monterrey, School of Engineering and Sciences, Zapopan, Jalisco 45138, Mexico

⁴Las Cumbres Observatory, 6740 Cortona Drive, Suite 102, Goleta, California 93117, USA

⁵University of Delaware, Newark, Delaware, 19716, USA

⁶The Oskar Klein Centre for Cosmoparticle Physics, Stockholm University, AlbaNova, Stockholm SE-106 91, Sweden

⁷The University of Chicago, Chicago, Illinois 60605, USA

⁸IBM T.J. Watson Research Center, Yorktown Heights, New York 10598, USA

⁹Observatories of the Carnegie Institute for Science, 813 Santa Barbara Street, Pasadena, California 91101-1232, USA

¹⁰West Virginia University, Morgantown, West Virginia 26506, USA

¹¹Center for Mathematical Modeling, Beaucheff 851, 7th floor, Santiago, Chile

¹²Google X, Mountain View, California 94043, USA

¹³NVIDIA, 2788 San Tomas Expressway Santa Clara, California, 95050

¹⁴Argonne National Laboratory, Leadership Computing Facility, Lemont, Illinois 60439, USA

¹⁵Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139, USA

¹⁶Columbia University in the City of New York, New York 10027, USA

¹⁷Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

¹⁸Universidad de Guadalajara, Centro Universitario de Ciencias Exactas e Ingenieria, Guadalajara, Jalisco 44430, Mexico

¹⁹University of Maryland, College Park, Maryland 80742, USA

²⁰Cardiff University, Cardiff CF24 3AA, United Kingdom

²¹NASA Goddard Space Flight Center, Code 661, 8800 Greenbelt Rd., Greenbelt, MD 20771, USA

²²University of Washington, 3910 15th Avenue NE, Seattle, Washington 98195, USA

²³Oak Ridge National Laboratory, Oak Ridge, TN 37831

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 two 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 gravitational wave 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), have enabled groundbreaking discoveries, such as the observation of supernovae 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 kilometer-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 $110 - 3840 \text{ Gpc}^{-3} \text{ yr}^{-1}$ for BNS, and $9.7 - 101 \text{ Gpc}^{-3} \text{ yr}^{-1}$ for BBH, at the 90% confidence level. Upper limits at the 90% confidence level for the neutron star-black hole (NSBH) merger rate is $610 \text{ Gpc}^{-3} \text{ yr}^{-1}$ (REF.¹). Furthermore, next-generation electromagnetic surveys will significantly increase the survey area, field of view, alert production of new and unexpected astronomical events, volumes of image data and catalog 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 hr period. Its alert latency will be 60 seconds, producing up to 10 million alerts per night. With a field of view of 9.6 deg^2 , it will survey a total area of $18,000 \text{ deg}^2$, and will complete about 1000 visits per night. By the end of the survey, the final raw image data will be 60 PB, requiring 0.4 exabytes 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 data sets of modeled 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 analyses¹²⁻¹⁴. 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 labor- 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 neighbor methods¹⁶, whose computational-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 centers, 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 with 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 (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 imaging pre-selected galaxies identified as plausible hosts. In both cases, these campaigns use automated discovery pipelines that perform image subtraction and artifact 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 catalog of potential host galaxies, which are accompanied with 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 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 GW observations are combined with photometric redshift catalogs. 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 catalogs 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 catalogs 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 which 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 astro-particle counterparts. Because of the computationally-intensive nature and poor-scalability of GW detection algorithms¹⁰⁻¹⁴, the imminent increase in the number of GW observations thanks 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 detectors⁴⁸⁻⁵¹. 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 modeling⁵²⁻⁵⁵.

Combining DL and HPC enables the design and training of neural network models using TB-size data sets of modeled 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 REF.⁵⁶, in which a template bank of over

ten million modeled 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 catalog of BBH mergers reported in REF.¹, finding that deep learning 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 $\sim 1,000$ 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 a 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 an 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 get insights into the physical processes that may lead to the production of electromagnetic and astro-particle counterparts, and to better interpret MMA observations⁶⁷. These modeling 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}.

Modeling 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. Modeling 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 hypermassive neutron star is formed (FIG. 1b) and matter is ejected at high velocities.

Accurate modeling 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, 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 modeling involves a wide range of length and time scales. Often, the relevant scales can be modeled using mesh refinement⁸⁴. A promising application of machine learning is to model sub-grid physics that cannot currently be modeled from first principles at computational costs⁸⁵ acceptable for parameter study simulations of turbulent motion during supernovae or neutron star collision⁸⁶. Sub-grid 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 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 PFlops^{-1} of single precision GPU computing power through the Bridges and Comet clusters, both of which are nearing retirement. New resources are expected 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 PFlops^{-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 Center for Supercomputing Applications, which consists of 16 IBM Power9 servers with four NVIDIA V100 GPUs, and which supports TensorFlow and other DL frameworks in a container-based environment. The Department of Energy 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 3-4 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 50-100 PFlops/s double precision, 10-100 petabytes storage, and external network data intake capacity sufficient for the analysis of the expected observational data. The estimated demand will likely require multiple instances of resources/centers 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, 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 desirable⁹³, as well as coordination between groups to determine which sources may be optimally followed-up with wide-field, all-sky survey instruments, or through target of opportunity observations with narrow field of view telescopes where the number of observable sources is significantly limited⁹⁴. DL classifiers could automate these time sensitive tasks in existing services such as the NASA/IPAC Extragalactic Database (NED) Gravitational Wave Follow-up (GWF) Service⁹⁵ and the Astrophysical Multi-Messenger Observatory Network (AMON)⁹⁶.

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 and static sky science communities, would avoid potential conflicts and waste of resources^{7,97-99}. Archives for large telescopes should ideally support real-time ingestion of freshly-acquired data and, and as much as possible, make MMA follow-up data public immediately after receipt.

Ideally, a common, public database would be build to centralize all information available for every recorded event. This will




Figure 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 $0.05M_{\odot}$ were synthesized in that event^{21,90}. This numerical relativity simulation was performed using the `Einstein Toolkit`⁹¹, with the addition of the `WhiskyTHC`⁹² hydrodynamics code, and the `CTGamma` space-time evolution code. The physics in this simulation includes the fully 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.

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 catalogs cross-matching. In this sense, the infrastructure of the LSST data management system and the alert brokers under development are good examples^{100,101}. The automation of telescopes is an important part of maximizing the number of observed MMA events (for example the Astronomical Event Observatory Network program will provide programmable access to a number of telescopes¹⁰²).

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 thanks 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 Astrophysics¹⁰³) 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 collaborations¹⁰⁴. 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 Plan¹⁰⁵ and a Community White paper¹⁰⁶ 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)¹⁰⁷, 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 sub-topics, including: systems development; data management, communication, and collaboration; analysis, inference, and machine learning; modeling 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 cyberinfrastructure 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 Table 1. We identified two pillars that require immediate attention to fully realize the MMA science program: first, increase the speed and depth of signal processing algorithms for real-time detection of GWs and their astro-particle and electromagnetic counterparts; and second, identify the cyber-infrastructure resources to simulate and search for MMA events in the ever increasing and disparate data sets. 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 oversubscribed. Therefore, we recommend reaching out to diverse advanced cyber-infrastructure centers in the US (Summit, Sierra, Frontera, Bridges-2, Theta, Lassen, etc.) and Europe (Piz Daint, SuperMUC-NG, etc.) 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 TB-size training data sets. 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 sub-fields. 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.

Area	Recommendation
Numerical relativity	<p>Improve coupling and cross-talk between numerical relativity and modeling tools and build a pipeline for predicting observables</p> <p>Accelerate community efforts to develop and release open source versions of modeling codes, especially microphysics and transport packages</p> <p>Explore publicly funded opportunities to develop numerical relativity software critical for MMA interpretations such as through the NSF or the DOE office of science</p>
Gravitational wave astrophysics and neutrino physics	<p>Accelerate the convergence of DL and HPC to reduce training stage from weeks to minutes of DL models that cover the complete signal manifold of black hole mergers</p> <p>Design Bayesian neural network models for GW parameter estimation covering the full signal manifold of black hole mergers</p> <p>Design neural network models for real-time detection of compact binaries involving neutron stars and core-collapse supernovae</p> <p>Design and deploy production-scale neural network classifiers to categorize particle interactions in neutrino detectors</p> <p>Accelerate the development of graph neural networks that are tailored for the specific topology of existing detectors</p>
Electromagnetic surveys	<p>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 EM spectrum and leverage machine learning for decision making and optimization</p> <p>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 catalog data)</p> <p>Continue the development of simulations and emulations in astrophysics modeling to fully explore the parameter space of the EM counterparts of GW events</p> <p>Create and facilitate automated access to comprehensive catalogs of archival and new observations to be used in event selection and classification</p> <p>Create communication channels between concurrently observing facilities to share plans for observing targets</p>
Cyber-infrastructure for DL	<p>Develop public training and test data sets, and a mock system infrastructure to test trained models</p> <p>Funding agencies (including NSF and DOE in the US) should agree on a split of responsibility for the common cyber-infrastructure needed to support all projects and researchers, rather than individual projects</p> <p>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)</p> <p>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</p> <p>Development of open-source software products and scalable and efficient DL algorithms which maximize use of HPC resources are needed</p> <p>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.</p>
Community building	<p>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.</p> <p>Build on the example of successfully-governed, open source community projects to develop community maintained software systems</p> <p>Organize workshops to train users and encourage participation</p> <p>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</p> <p>Develop rules and guidelines for data sharing and publication to assure ethical behavior and avoid waste of EM observational resources</p>

Table 1. Summary of recommendations to realize the science goals of Multi-Messenger Astrophysics. Notes: NSF, National Science Foundation; DOE, Department of Energy; HPC, High Performance Computing; GW, gravitational waves; EM, electromagnetic; DL, deep learning

References

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.* **X9**, 031040 (2019). [1811.12907](#).
2. Arnett, W. D., Bahcall, J. N., Kirshner, R. P. & Woosley, S. E. Supernova 1987a. *Annual Review of Astronomy and Astrophysics* **27**, 629–700 (1989).
3. Abbott, B. P. *et al.* GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physical Review Letters* **119**, 161101 (2017). [1710.05832](#).
4. Abbott, B. P. *et al.* Estimating the Contribution of Dynamical Ejecta in the Kilonova Associated with GW170817. *Astrophys. J.* **850**, L39 (2017). [1710.05836](#).
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. <https://www.lsst.org/scientists/keynumbers>.
7. Collaboration, L. S., Abell, P. A. *et al.* LSST Science Book, Version 2.0 (2009). [0912.0201](#).
8. Robertson, B. E. *et al.* Galaxy formation and evolution science in the era of the Large Synoptic Survey Telescope. *Nature 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). [gr-qc/9808076](#).
10. Huerta, E. A. *et al.* Complete waveform model for compact binaries on eccentric orbits. *Phys. Rev. D* **95**, 024038 (2017). [1609.05933](#).
11. Harry, I., Privitera, S., Bohé, A. & Buonanno, A. Searching for gravitational waves from compact binaries with precessing spins. *Phys. Rev. D* **94**, 024012 (2016). [1603.02444](#).
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 *Proceedings, 13th International Conference on e-Science: Auckland, New Zealand, October 24-27, 2017*, 335–344 (2017). [1709.08767](#).
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). [1810.03056](#).
14. Weitzel, D. *et al.* Data Access for LIGO on the OSG. *arXiv e-prints* arXiv:1705.06202 (2017). [1705.06202](#).
15. Abbott, B. P. *et al.* Observing gravitational-wave transient GW150914 with minimal assumptions. *Phys. Rev. D* **93**, 122004 (2016). [1602.03843](#).
16. Jones, P. W., Osipov, A. & Rokhlin, V. Randomized approximate nearest neighbors algorithm. *Proceedings of the National Academy of Sciences* **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 (2010).
18. Asch, M. *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. *The International Journal of High Performance Computing Applications* **32**, 435–479 (2018).
19. National Academies of Sciences, Engineering, and Medicine. *Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020* (The National Academies Press, Washington, DC, 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). [1108.6056](#).
21. Siegel, D. M. & Metzger, B. D. Three-dimensional grmhd simulations of neutrino-cooled accretion disks from neutron star mergers. *The Astrophysical Journal* **858**, 52 (2018).
22. Abbott, B. P. *et al.* Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO and Advanced Virgo. *Living Reviews in Relativity* **19** (2016). [1304.0670](#).
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). [1710.05443](#).

24. Mooley, K. P. *et al.* A mildly relativistic wide-angle outflow in the neutron-star merger event GW170817. *Nature* **554**, 207–210 (2018). [1711.11573](#).
25. Andreoni, I. *et al.* Mary, a Pipeline to Aid Discovery of Optical Transients. *Publications of the Astronomical Society of Australia* **34**, e037 (2017). [1708.04629](#).
26. Sedaghat, N. & Mahabal, A. Effective image differencing with convolutional neural networks for real-time transient hunting. *MNRAS* **476**, 5365–5376 (2018). [1710.01422](#).
27. Kessler, R. *et al.* Results from the Supernova Photometric Classification Challenge. *Publications of the Astronomical Society of the Pacific* **122**, 1415 (2010). [1008.1024](#).
28. Jones, D. O. *et al.* Measuring Dark Energy Properties with Photometrically Classified Pan-STARRS Supernovae. II. Cosmological Parameters. *Astrophys. J.* **857**, 51 (2018). [1710.00846](#).
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). [1710.05845](#).
30. Setzer, C. N. *et al.* Serendipitous Discoveries of Kilonovae in the LSST Main Survey: Maximising Detections of Sub-Threshold Gravitational Wave Events. *arXiv e-prints* arXiv:1812.10492 (2018). [1812.10492](#).
31. Schutz, B. F. Determining the Hubble constant from gravitational wave observations. *Nature* **323**, 310 (1986).
32. The DES Collaboration, the LIGO Scientific Collaboration & the Virgo Collaboration. 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. *arXiv e-prints* arXiv:1901.01540 (2019). [1901.01540](#).
33. 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. *The Astrophysical Journal Letters* **848**, L17 (2017).
34. Abbott, B. P. *et al.* A gravitational-wave standard siren measurement of the Hubble constant. *Nature* **551**, 85–88 (2017).
35. Fishbach, M. *et al.* A standard siren measurement of the Hubble constant from GW170817 without the electromagnetic counterpart. *ArXiv e-prints* arXiv:1807.05667 (2018). [1807.05667](#).
36. Khan, A. *et al.* Deep learning at scale for the construction of galaxy catalogs in the Dark Energy Survey. *Physics Letters B* **795**, 248–258 (2019). [1812.02183](#).
37. Domínguez Sánchez, H., Huertas-Company, M., Bernardi, M., Tuccillo, D. & Fischer, J. L. Improving galaxy morphologies for SDSS with Deep Learning. *MNRAS* **476**, 3661–3676 (2018).
38. Eisenstein, D. J. *et al.* SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way, and Extra-Solar Planetary Systems. *The Astronomical Journal* **142**, 72 (2011). [1101.1529](#).
39. Dark Energy Survey Collaboration *et al.* The Dark Energy Survey: more than dark energy - an overview. *MNRAS* **460**, 1270–1299 (2016). [1601.00329](#).
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). [1903.07603](#).
41. Aghanim, N. *et al.* Planck 2018 results. VI. Cosmological parameters (2018). [1807.06209](#).
42. Poulin, V., Smith, T. L., Karwal, T. & Kamionkowski, M. Early dark energy can resolve the hubble tension. *Phys. Rev. Lett.* **122**, 221301 (2019).
43. Freedman, W. L. Cosmology at a crossroads. *Nature Astronomy* **1**, 0121 (2017).
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). [1701.00008](#).
46. George, D. & Huerta, E. A. Deep Learning for real-time gravitational wave detection and parameter estimation: Results with Advanced LIGO data. *Physics Letters B* **778**, 64–70 (2018). [1711.03121](#).
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). [1807.09787](#).
48. Wei, W. & Huerta, E. A. Gravitational Wave Denoising of Binary Black Hole Mergers with Deep Learning. *Physics Letters B* **800**, 135081 (2020). [1901.00869](#).

49. Shen, H., George, D., Huerta, E. A. & Zhao, Z. Denoising gravitational waves with enhanced deep recurrent denoising auto-encoders. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3237–3241 (2019).
50. Shen, H., George, D., Huerta, E. A. & Zhao, Z. Denoising Gravitational Waves using Deep Learning with Recurrent Denoising Autoencoders. *ArXiv e-prints* (2017). [1711.09919](#).
51. George, D., Shen, H. & Huerta, E. A. Classification and unsupervised clustering of LIGO data with Deep Transfer Learning. *Phys. Rev. D* **97**, 101501 (2018). [1706.07446](#).
52. Chua, A. J. K., Galley, C. R. & Vallisneri, M. Reduced-order modeling with artificial neurons for gravitational-wave inference. *Phys. Rev. Lett.* **122**, 211101 (2019). [1811.05491](#).
53. Gabbard, H., Williams, M., Hayes, F. & Messenger, C. Matching Matched Filtering with Deep Networks for Gravitational-Wave Astronomy. *Physical Review Letters* **120**, 141103 (2018). [1712.06041](#).
54. Nakano, H. *et al.* Comparison of various methods to extract ringdown frequency from gravitational wave data. *Phys. Rev. D* **99**, 124032 (2019). [1811.06443](#).
55. Dreissigacker, C., Sharma, R., Messenger, C., Zhao, R. & Prix, R. Deep-Learning Continuous Gravitational Waves. *Phys. Rev. D* **100**, 044009 (2019). [1904.13291](#).
56. Shen, H., Huerta, E. A. & Zhao, Z. Deep Learning at Scale for Gravitational Wave Parameter Estimation of Binary Black Hole Mergers. *arXiv e-prints* arXiv:1903.01998 (2019). [1903.01998](#).
57. Springenberg, J. T., Klein, A., Falkner, S. & Hutter, F. Bayesian optimization with robust bayesian neural networks. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I. & Garnett, R. (eds.) *Advances in Neural Information Processing Systems 29*, 4134–4142 (Curran Associates, Inc., 2016).
58. Burrows, A., Hayes, J. & Fryxell, B. A. On the nature of core collapse supernova explosions. *Astrophys. J.* **450**, 830 (1995). [astro-ph/9506061](#).
59. Burrows, A., Radice, D. & Vartanyan, D. Three-dimensional supernova explosion simulations of 9-, 10-, 11-, 12-, and 13- M_{\odot} stars. *Mon. Not. Roy. Astron. Soc.* **485**, 3153–3168 (2019). [1902.00547](#).
60. Radice, D., Morozova, V., Burrows, A., Vartanyan, D. & Nagakura, H. Characterizing the Gravitational Wave Signal from Core-Collapse Supernovae. *Astrophys. J.* **876**, L9 (2019). [1812.07703](#).
61. Mösta, P. *et al.* R-process Nucleosynthesis from Three-Dimensional Magnetorotational Core-Collapse Supernovae. *Astrophys. J.* **864**, 171 (2018). [1712.09370](#).
62. Janka, H. T., Melson, T. & Summa, A. Physics of Core-Collapse Supernovae in Three Dimensions: a Sneak Preview. *Ann. Rev. Nucl. Part. Sci.* **66**, 341–375 (2016). [1602.05576](#).
63. Woosley, S. & Janka, T. The physics of core-collapse supernovae. *Nature Phys.* **1**, 147 (2005). [astro-ph/0601261](#).
64. Gossan, S. E. *et al.* Observing Gravitational Waves from Core-Collapse Supernovae in the Advanced Detector Era. *Phys. Rev. D* **93**, 042002 (2016). [1511.02836](#).
65. Aurisano, A. *et al.* A Convolutional Neural Network Neutrino Event Classifier. *JINST* **11**, P09001 (2016). [1604.01444](#).
66. Choma, N. *et al.* Graph neural networks for icecube signal classification. In Wani, M., Sayed-Mouchaweh, M., Lughofer, E., Gama, J. & Kantardzic, M. (eds.) *Proceedings - 17th IEEE International Conference on Machine Learning and Applications, ICMLA 2018*, Proceedings - 17th IEEE International Conference on Machine Learning and Applications, ICMLA 2018, 386–391 (Institute of Electrical and Electronics Engineers Inc., 2019).
67. Hinderer, T. *et al.* Distinguishing the nature of comparable-mass neutron star binary systems with multimessenger observations: GW170817 case study. *Phys. Rev. D* **100**, 06321 (2019). [1808.03836](#).
68. Kim, B. *et al.* Deep fluids: A generative network for parameterized fluid simulations. *Computer Graphics Forum* **38**, 59–70 (2019).
69. Ling, J., Kurzwaski, A. & Templeton, J. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *Journal of Fluid Mechanics* **807**, 155–166 (2016).
70. Maulik, R., San, O., Rasheed, A. & Vedula, P. Subgrid modelling for two-dimensional turbulence using neural networks. *Journal of Fluid Mechanics* **858**, 122–144 (2019).
71. Viganò, D., Aguilera-Miret, R. & Palenzuela, C. Extension of the subgrid-scale gradient model for compressible magnetohydrodynamics turbulent instabilities. *Physics of Fluids* **31**, 105102 (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. *Communications in Mathematics and Statistics* **5**, 349–380 (2017).
75. Duez, M. D. & Zlochower, Y. Numerical relativity of compact binaries in the 21st century. *Reports on Progress in Physics* **82**, 016902 (2019). [1808.06011](https://arxiv.org/abs/1808.06011).
76. Baiotti, L. & Rezzolla, L. Binary neutron star mergers: a review of einstein’s richest laboratory. *Reports on Progress in Physics* **80**, 096901 (2017).
77. Lippuner, J. & Roberts, L. F. Skynet: A modular nuclear reaction network library. *The Astrophysical Journal Supplement Series* **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). [1410.7392](https://arxiv.org/abs/1410.7392).
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). [1604.02455](https://arxiv.org/abs/1604.02455).
80. Fernández, R. *et al.* Long-term GRMHD Simulations of Neutron Star Merger Accretion Disks: Implications for Electromagnetic Counterparts. *ArXiv e-prints* (2018). [1808.00461](https://arxiv.org/abs/1808.00461).
81. Hossein Nouri, F. *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). [1809.11161](https://arxiv.org/abs/1809.11161).
83. Kasen, D., Badnell, N. R. & Barnes, J. Opacities and Spectra of the r-process Ejecta from Neutron Star Mergers. *ApJ* **774**, 25 (2013). [1303.5788](https://arxiv.org/abs/1303.5788).
84. Berger, M. & Colella, P. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics* **82**, 64 – 84 (1989).
85. Chen, T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. Neural ordinary differential equations. *arXiv preprint arXiv:1806.07366* (2018).
86. Radice, D. *et al.* Neutrino-Driven Convection in Core-Collapse Supernovae: High-Resolution Simulations. *Astrophys. J.* **820**, 76 (2016). [1510.05022](https://arxiv.org/abs/1510.05022).
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). [1410.0013](https://arxiv.org/abs/1410.0013).
88. EuroHPC. EuroHPC: Leading the way in the European Supercomputing. <https://eurohpc-ju.europa.eu/index.html>. 2018.
89. Huerta, E. A. *et al.* Boss-Ildg: 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 (e-Science)*, 335–344 (2017).
90. Hotokezaka, K., Beniampini, P. & Piran, T. Neutron star mergers as sites of r-process nucleosynthesis and short gamma-ray bursts. *International Journal of Modern Physics D* **27**, 1842005 (2018). [1801.01141](https://arxiv.org/abs/1801.01141).
91. Löffler, F. *et al.* The Einstein Toolkit: a community computational infrastructure for relativistic astrophysics. *Classical and Quantum Gravity* **29**, 115001 (2012). [1111.3344](https://arxiv.org/abs/1111.3344).
92. 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). [1206.6502](https://arxiv.org/abs/1206.6502).
93. Arcavi, I. *et al.* Optical Follow-up of Gravitational-wave Events with Las Cumbres Observatory. *Astrophys. J.* **848**, L33 (2017). [1710.05842](https://arxiv.org/abs/1710.05842).
94. Coughlin, M. W. *et al.* Optimizing searches for electromagnetic counterparts of gravitational wave triggers. *Mon. Not. Roy. Astron. Soc.* **478**, 692–702 (2018). [1803.02255](https://arxiv.org/abs/1803.02255).

95. NED. California Institute of Technology NED gravitational wave follow-up service. <https://ned.ipac.caltech.edu/gwf/overview>. 2019.
96. AMON. Pennsylvania State University Astrophysical multimessenger observatory network. <https://www.amon.psu.edu/amon-system/>. 2019.
97. Marshall, P. *et al.* Science-Driven Optimization of the LSST Observing Strategy (2017). [1708.04058](https://arxiv.org/abs/1708.04058).
98. W. M. Keck Observatory Scientific Strategic Plan. The W. M. Keck Observatory Scientific Strategic Plan. https://www.ucolick.org/home/about/governance/Keck_Scientific_Strategic_Plan_Final.pdf (2016). Online from July 2016.
99. Cowperthwaite, P. S. *et al.* Astro 2020 Science White Paper: Joint Gravitational Wave and Electromagnetic Astronomy with LIGO and LSST in the 2020's (2019). [1904.02718](https://arxiv.org/abs/1904.02718).
100. Narayan, G. *et al.* Machine Learning-based Brokers for Real-time Classification of the LSST Alert Stream. *Astrophys. J. Suppl.* **236**, 9 (2018). [1801.07323](https://arxiv.org/abs/1801.07323).
101. Smith, K. W. *et al.* Lasair: The Transient Alert Broker for LSST:UK. *Research Notes of the American Astronomical Society* **3**, 26 (2019).
102. AEON Team. Astronomical event observatory network (2018). <http://ast.noao.edu/data/aeon>.
103. The National Academies of Sciences, Engineering, and Medicine. The National Academies of Sciences, Engineering, and Medicine. The decadal survey on astronomy and astrophysics (astro2020). <https://sites.nationalacademies.org/DEPS/astro2020/index.htm>. 2019.
104. Katz, D. S. *et al.* Community organizations: Changing the culture in which research software is developed and sustained. *Comp. in Sci. & Eng.* **21**, 8–24 (2019).
105. Elmer, P., Neubauer, M. & Sokoloff, M. D. Strategic Plan for a Scientific Software Innovation Institute (S2I2) for High Energy Physics. *arXiv e-prints* arXiv:1712.06592 (2017). [1712.06592](https://arxiv.org/abs/1712.06592).
106. Albrecht, J. *et al.* A Roadmap for HEP Software and Computing R&D for the 2020s. *arXiv e-prints* arXiv:1712.06982 (2017). [1712.06982](https://arxiv.org/abs/1712.06982).
107. Allen, G. *et al.* Multi-Messenger Astrophysics: Harnessing the Data Revolution. *arXiv e-prints* arXiv:1807.04780 (2018). [1807.04780](https://arxiv.org/abs/1807.04780).

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.

Reviewer 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.