
Bayesian multi-band fitting of alerts for kilonovae detection

Biswajit Biswas

AstroParticule et Cosmologie
Université Paris Cité, CNRS
F-75013, Paris, France
biswas@apc.in2p3.fr

Junpeng Lao

Google Switzerland
Gustav-Gull-Platz 1
8004 Zurich, Switzerland

Eric Aubourg

AstroParticule et Cosmologie
Université Paris Cité, CNRS
F-75013, Paris, France

Alexandre Boucaud

AstroParticule et Cosmologie
Université Paris Cité, CNRS
F-75013, Paris, France

Axel Guinot

AstroParticule et Cosmologie
Université Paris Cité, CNRS
F-75013, Paris, France

Emille E. O. Ishida

LPC, Université Clermont Auvergne
CNRS/IN2P3
F-63000, Clermont-Ferrand, France

Cécile Roucelle

AstroParticule et Cosmologie
Université Paris Cité, CNRS
F-75013, Paris, France

Abstract

In the era of multi-messenger astronomy, early classification of photometric alerts from wide-field and high-cadence surveys is a necessity to trigger spectroscopic follow-ups. These classifications are expected to play a key role in identifying potential candidates that might have a corresponding gravitational wave (GW) signature. Machine learning classifiers using features from parametric fitting of light curves are widely deployed by broker software to analyze millions of alerts, but most of these algorithms require as many points in the filter as the number of parameters to produce the fit, which increases the chances of missing a short transient. Moreover, the classifiers are not able to account for the uncertainty in the fits when producing the final score. In this context, we present a novel classification strategy that incorporates data-driven priors for extracting a joint posterior distribution of fit parameters and hence obtaining a distribution of classification scores. We train and test a classifier to identify kilonovae events which originate from binary neutron star mergers or neutron star black hole mergers, among simulations for the Zwicky Transient Facility observations with 19 other non-kilonovae-type events. We demonstrate that our method can estimate the uncertainty of misclassification, and the mean of the distribution of classification scores as point estimate obtains an AUC score of 0.96 on simulated data. We further show that using this method we can process the entire alert stream in real-time and bring down the sample of probable events to a scale where they can be analyzed by domain experts.

1 Introduction

The detection of the binary neutron star (BNS) merger event GW170817 [Abbott et al., 2017] marked the dawn of multi-messenger astronomy as the electromagnetic (EM) counterpart of the event was detected following the initial GW detection by LIGO and VIRGO collaborations. This EM counterpart, so-called kilonova (KN), originates from BNS and possibly from neutron star black hole

mergers and is a source of heavy element production via the r-process[Villar et al., 2017]. Moreover, detecting such events along with the GW counterpart gives us a way to evaluate the value of Hubble’s Constant (H_0) [Schutz, 1986] and addresses the tension between the Cosmic Microwave Background and Type Ia supernovae (SNIa) measurements of H_0 [Freedman, 2021].

To a certain extent, GW detectors can localize events in the sky. However, a large sky fraction needs to be scanned to look for EM counterparts within the localization error boxes obtained from GW signal. Although spectroscopic measurement is necessary to obtain an accurate spectrum of these events, wide-field surveys observe a much bigger patch of the sky in broadband filters so they can be used to select targets for deployment of the expensive spectroscopic follow-up. Hence considerable efforts are being put into identifying potential events using the Zwicky Transient Facility [ZTF, Bellm et al., 2018, Andreoni et al., 2021], Dark Energy Camera[Flaugher et al., 2015, Garcia et al., 2020] and also preparing for the arrival of data from the Rubin Observatory Legacy Survey of Space and Time [LSST, Ivezić et al., 2019].

Whenever a 5σ signal is detected above the background, ZTF creates alerts containing a 30-day history of the event. Since around one million such alerts are generated each night by ZTF, intermediary broker software is required to process the incoming alert stream and identify events of interest. Typical methods of classification rely on parameters obtained from performing parametric fits of the light curves, independently on each filter. A major challenge in identifying KNe events in the alert stream is the limited availability of photometric detection points in each broadband filter. As we move from ZTF to LSST, we will have more photometric bands, but fewer points per band. Hence we need to prepare for the complexity of classification with scarce data.

In the context of the FINK broker [Möller et al., 2020] that is being developed for LSST and is currently processing alerts from the ZTF, we present a new approach to combine information from multiple bands and jointly fit all the filters of an alert by performing a Bayesian fit using Monte Carlo Markov Chain (MCMC) with a data-driven prior. By performing a simultaneous fit, we remove the constraint of a minimum number of points and also provide a distribution of the scores of the classification. Although our approach can be extended to any parametric fitting, in this work we use the parametric fitting scheme presented in Biswas, B. et al. [2023].

2 Dataset

To train and test our classifier we use simulations from Muthukrishna et al. [2019] generated with SNANA simulator [Kessler et al., 2009] and models from the Photometric LSST Astronomical Classification Challenge [PLAsTiCC, Kessler et al., 2019]. This set consists of 42568 events among which 38000 are non-KNe representing 19 other types of events, while the remaining 4568 KNe simulations comprise 2 KN models (see Muthukrishna et al. [2019] for details). Moreover, to represent a more recent ZTF cadence, another 1000 KNe events from Stachie et al. [2020] were added to this dataset.

To mimic alert data, we consider 30 days of data around the observation with the highest signal-to-noise ratio, starting 10 days before and running up to 20 days after. This assumption is not exactly true for alerts because we will have information only up to the day of the alert. However, it serves as a necessary first test to evaluate our method before testing on a more complex setting. This combined dataset is then split into a training set of 3280 KNe and 19000 non-KNe and a test set of 2288 KNe and 19000 non-KNe. For more details, readers can refer to Biswas, B. et al. [2023].

3 Methodology

To demonstrate the performance of our approach we use the method proposed by Biswas, B. et al. [2023] where each filter of the light curve is represented as a linear combination of three basis vectors generated by performing principal component analysis on noiseless light curve realizations. Figure 1a shows the set of templates that are used as basis vectors to fit the light curves from day -50 to 50 , where day 0 is approximately the maximum observed flux. The optimized coefficients of this linear combination constitute a set of features that are fed to a random forest [RF, Ho, 1995] classifier to classify alerts as KNe by the authors. A limitation of such approaches is that the parametric fit is performed independently on each filter. It requires at least as many points in each filter as the

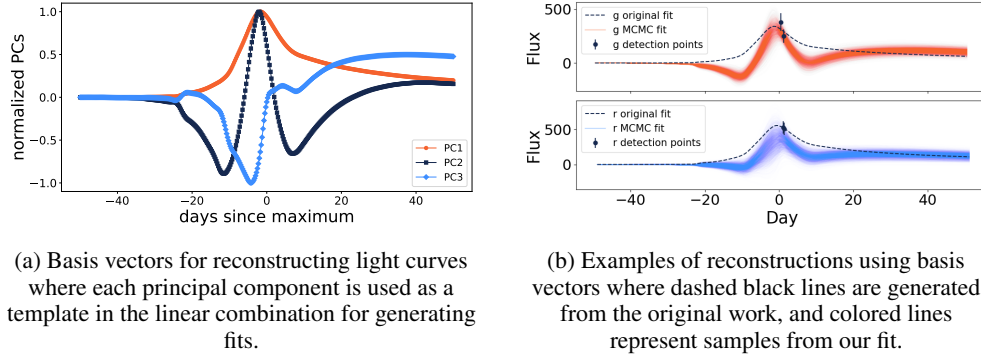


Figure 1: Parametric fitting of light curves.

number of free parameters. Moreover, the uncertainty of the fit is not taken into account during the classification process since the fit is deterministic.

In this work, we first develop a data-driven prior for the joint distribution of parameters of all the filters. We combine the prior with Gaussian log-likelihood for the observed points to generate MCMC chains to obtain a posterior distribution for the parameters. Figure 1b shows an example of the distribution of fits we obtained from our algorithm instead of the deterministic fit proposed by [Biswas, B. et al., 2023]. The authors generate fits even when there are fewer points than the number of free parameters by using strong regularization, but they still require at least 2 points in each band. As we move towards LSST, such a constraint on all the 6 LSST bands will result in missing faster transients like KNe. Hence, motivating a joint multi-band Bayesian fitting to eliminate this constraint.

3.1 Generating features

To build the data-driven prior, we need to identify a sample of events for which we are confident that we have good fits. Ideally, the goal would be to construct this prior from real data but due to the lack of availability to real KNe events, we rely on the features generated by Biswas, B. et al. [2023] on the realistic training set simulations. These features were generated only on events with multiple points in each filter thus ensuring reliable fits. As KNe events are short and more difficult to fit, we extract a joint multi-band covariance matrix from the features of only the KNe events to act as a prior for our Bayesian fit.

Although we have all the necessary components, before running the MCMC chains we must normalize the light curve data in the same manner as Biswas, B. et al. [2023]. We thus divide the data of each band by the maximum observed flux in that band. Once the parameters of the fit are obtained, we include this information, once again, in the classifier by passing the normalization factor as a separate feature. For each event, after a warmup phase of 500 steps, we generate 4 MCMC chains with 500 samples in each chain which provides us with a posterior distribution of 2000 samples per event.

3.2 Classification

One approach for classification could be to use summary statistics such as the mean and standard deviation of the posterior distribution to obtain a discrete score for the event. But in a high-dimension space, if the distribution is non-Gaussian, these summary statistics will not encompass all the information. Instead, we chose to obtain a distribution of classification scores by classifying each posterior sample individually.

Along with the fit parameters, we also pass as features to the classifier the maximum value of flux in each band and the log-likelihood of the fit. We compared an RF using `scikit-learn` [Pedregosa et al., 2011] and XGBoost classifier [Chen and Guestrin, 2016] with 30 trees and a max depth of 42 as no significant performance gain was observed with a more complicated model. For training, we used a balanced set of 200k samples drawn from the complete pool of samples and found that the performances were comparable (Figure 2). Hence, we hereafter consider only the results from the RF classifier. Figure 3 shows the distribution of scores for some events with mean scores around 0.5 and some around 1. As we do not observe a bimodal distribution for the scores, we use the




Figure 2: Receiver operating characteristics curve.




Figure 3: Examples of distribution of scores for 5 events with mean score close to 1 and 5 events near the classification threshold of 0.5.

mean and standard deviation of the distribution for each alert as summary statistics to draw our final conclusions.

4 Results


Upon using the classifier on the test sample, we obtain the ROC curve shown in figure 2 using the mean as a point estimate of the classification score. The exact choice of threshold would depend upon the application of such a classifier. For example, as there are around one million alerts from ZTF every night, we might want to restrict ourselves to a smaller false positive (FP) rate so that we have a reliable sample that can be followed up by experts. On the other hand, if we have a GW localization in a particular region of the sky, our goal would be to not miss the event, even if it is at the cost of a high FP rate.

Since the classification threshold should ideally be dynamic, we present results with the threshold set to 0.5. Under this setting, we obtain a precision of 0.45 and recall of 0.85. Even though the value of precision might seem fairly low, we can bring down the number of potential events to a scale where they can be analyzed by domain experts while maintaining a high value for the recall, i.e., without missing many actual KNe events. In Figure 4 we further observe that the FP events have relatively lower values of scores and a higher value of uncertainty associated with the classification hence validating the uncertainty quantification. A direct comparison with the results of Biswas, B. et al. [2023] where the authors require at least 2 points in each band, we obtain a drop in performance in terms of both precision and recall. However, such a comparison can be misleading as our sample contains events with less than 2 data points per band as well since there is no constraint on the minimum number of points.

5 Conclusion and Outlook

As we prepare for the arrival of LSST, we need robust classification algorithms that can operate on a small number of points for early classification of short-term transients. Our Bayesian fitting approach for KN classification takes a step in this direction by jointly fitting all filters, thereby combining information from all photometric bands to generate a parametric fit. Our algorithm is not only able to incorporate information from a data-driven prior but also estimate the uncertainty of the fits.

As the next step, we must retrain the classifier on more alert-like curves with no data after the day of the alert. Although a slight drop in performance is expected because there will be incomplete data about the light curve, we expect the results to generalize fairly well. Once trained, it is straightforward to integrate the method proposed here within the FINK broker to be tested on real ZTF alerts before being extended for the 6 bands of LSST. Additionally, the computational cost needs to be taken into account when pushing the module into production. We were able to generate features for around 1000 events per minute using a single CPU. The algorithm being written in the JAX[Bradbury et al., 2018] framework, can easily be parallelized for faster computing efficiency. Although our algorithm is capable of processing the entire ZTF alert stream, processing the LSST stream, which will be at least one order of magnitude larger, would still be a challenge. However, we can easily introduce



(a) Distribution of scores and sigma for True positives.

(b) Distribution of scores and sigma for False positives.

Figure 4: Distribution of mean(μ) and sigma(σ) for each event in test set.

appropriate selections to discard a major fraction of events at an earlier stage to significantly reduce the computational cost. For LSST, in particular, this module will be used in combination with a series of selection cuts which would result in a feasible mechanism to analyze the alert stream and filter out interesting events in real-time so that spectroscopic follow-up can be triggered on these events.

Acknowledgments and Disclosure of Funding

We deeply thank Daniel Muthukrishna and Cosmin Stachie for providing a set of simulated KNe data used to train the model. This project has received funding from the European Unions Horizon 2020 research and innovation program under the Marie Sklodowska-Curie grant agreement No 945304 Cofund AI4theSciences hosted by PSL University. This work was developed within the FINK community and made use of the FINK community broker resources. FINK is supported by LSST-France and CNRS/IN2P3.

References

- B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, and et al. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017. doi: 10.1103/PhysRevLett.119.161101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101>.
- I. Andreoni, M. W. Coughlin, E. C. Kool, M. M. Kasliwal, H. Kumar, V. Bhalerao, A. S. Carracedo, A. Y. Q. Ho, P. T. H. Pang, D. Saraogi, K. Sharma, V. Shenoy, E. Burns, T. Ahumada, S. Anand, L. P. Singer, D. A. Perley, K. De, U. C. Fremling, E. C. Bellm, M. Bulla, A. Crellin-Quick, T. Dietrich, A. Drake, D. A. Duev, A. Goobar, M. J. Graham, D. L. Kaplan, S. R. Kulkarni, R. R. Laher, A. A. Mahabal, D. L. Shupe, J. Sollerman, R. Walters, and Y. Yao. Fast-transient Searches in Real Time with ZTFReST: Identification of Three Optically Discovered Gamma-Ray Burst Afterglows and New Constraints on the Kilonova Rate. *ApJ*, 918(2):63, Sept. 2021. doi: 10.3847/1538-4357/ac0bc7.
- E. C. Bellm, S. R. Kulkarni, M. J. Graham, R. Dekany, R. M. Smith, R. Riddle, F. J. Masci, G. Helou, T. A. Prince, S. M. Adams, C. Barbarino, T. Barlow, J. Bauer, R. Beck, J. Belicki, R. Biswas, N. Blagorodnova, D. Bodewits, B. Bolin, V. Brinnel, T. Brooke, B. Bue, M. Bulla, R. Burruss, S. B. Cenko, C.-K. Chang, A. Connolly, M. Coughlin, J. Cromer, V. Cunningham, K. De, A. Delacroix, V. Desai, D. A. Duev, G. Eadie, T. L. Farnham, M. Feeney, U. Feindt, D. Flynn, A. Franckowiak, S. Frederick, C. Fremling, A. Gal-Yam, S. Gezari, M. Giomi, D. A. Goldstein, V. Z. Golkhou, A. Goobar, S. Groom, E. Hacquins, D. Hale, J. Henning, A. Y. Q. Ho, D. Hover, J. Howell, T. Hung, D. Huppenkothen, D. Imel, W.-H. Ip, eljko Ivezi, E. Jackson, L. Jones, M. Juric, M. M. Kasliwal, S. Kaspi, S. Kaye, M. S. P. Kelley, M. Kowalski, E. Kramer,

- T. Kupfer, W. Landry, R. R. Laher, C.-D. Lee, H. W. Lin, Z.-Y. Lin, R. Lunnan, M. Giomi, A. Mahabal, P. Mao, A. A. Miller, S. Monkewitz, P. Murphy, C.-C. Ngeow, J. Nordin, P. Nugent, E. Ofek, M. T. Patterson, B. Penprase, M. Porter, L. Rauch, U. Rebbapragada, D. Reiley, M. Rigault, H. Rodriguez, J. van Roestel, B. Rusholme, J. van Santen, S. Schulze, D. L. Shupe, L. P. Singer, M. T. Soumagnac, R. Stein, J. Surace, J. Sollerman, P. Szkody, F. Taddia, S. Terek, A. V. Sistine, S. van Velzen, W. T. Vestrand, R. Walters, C. Ward, Q.-Z. Ye, P.-C. Yu, L. Yan, and J. Zolkower. The zwicky transient facility: System overview, performance, and first results. *Publications of the Astronomical Society of the Pacific*, 131(995):018002, dec 2018. doi: 10.1088/1538-3873/aaecbe. URL <https://dx.doi.org/10.1088/1538-3873/aaecbe>.
- Biswas, B., Ishida, E. E. O., Peloton, J., Möller, A., Pruzhinskaya, M. V., de Souza, R. S., and Muthukrishna, D. Enabling the discovery of fast transients - a kilonova science module for the fink broker. *A&A*, 677:A77, 2023. doi: 10.1051/0004-6361/202245340. URL <https://doi.org/10.1051/0004-6361/202245340>.
- J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4232-2. doi: 10.1145/2939672.2939785. URL <http://doi.acm.org/10.1145/2939672.2939785>.
- B. Flaugher, H. T. Diehl, K. Honscheid, T. M. C. Abbott, O. Alvarez, R. Angstadt, J. T. Annis, M. Antonik, O. Ballester, L. Beaufore, G. M. Bernstein, R. A. Bernstein, B. Bigelow, M. Bonati, D. Boprie, D. Brooks, E. J. Buckley-Geer, J. Campa, L. Cardiel-Sas, F. J. Castander, J. Castilla, H. Cease, J. M. Cela-Ruiz, S. Chappa, E. Chi, C. Cooper, L. N. da Costa, E. Dede, G. Derylo, D. L. DePoy, J. de Vicente, P. Doel, A. Drlica-Wagner, J. Eiting, A. E. Elliott, J. Emes, J. Estrada, A. F. Neto, D. A. Finley, R. Flores, J. Frieman, D. Gerdes, M. D. Gladders, B. Gregory, G. R. Gutierrez, J. Hao, S. E. Holland, S. Holm, D. Huffman, C. Jackson, D. J. James, M. Jonas, A. Karcher, I. Karliner, S. Kent, R. Kessler, M. Kozlovsky, R. G. Kron, D. Kubik, K. Kuehn, S. Kuhlmann, K. Kuk, O. Lahav, A. Lathrop, J. Lee, M. E. Levi, P. Lewis, T. S. Li, I. Mandrichenko, J. L. Marshall, G. Martinez, K. W. Merritt, R. Miquel, F. Muñoz, E. H. Neilsen, R. C. Nichol, B. Nord, R. Ogando, J. Olsen, N. Palaio, K. Patton, J. Peoples, A. A. Plazas, J. Rauch, K. Reil, J.-P. Rheault, N. A. Roe, H. Rogers, A. Roodman, E. Sanchez, V. Scarpine, R. H. Schindler, R. Schmidt, R. Schmitt, M. Schubnell, K. Schultz, P. Schurter, L. Scott, S. Serrano, T. M. Shaw, R. C. Smith, M. Soares-Santos, A. Stefanik, W. Stuermer, E. Suchyta, A. Sypniewski, G. Tarle, J. Thaler, R. Tighe, C. Tran, D. Tucker, A. R. Walker, G. Wang, M. Watson, C. Weaverdyck, W. Wester, R. Woods, B. Yanny, and T. D. Collaboration). The dark energy camera. *The Astronomical Journal*, 150(5):150, oct 2015. doi: 10.1088/0004-6256/150/5/150. URL <https://dx.doi.org/10.1088/0004-6256/150/5/150>.
- W. L. Freedman. Measurements of the Hubble Constant: Tensions in Perspective. *ApJ*, 919(1):16, Sept. 2021. doi: 10.3847/1538-4357/ac0e95.
- A. Garcia, R. Morgan, K. Herner, A. Palmese, M. Soares-Santos, J. Annis, D. Brout, A. K. Vivas, A. Drlica-Wagner, L. Santana-Silva, D. L. Tucker, S. Allam, M. Wiesner, J. García-Bellido, M. S. S. Gill, M. Sako, R. Kessler, T. M. Davis, D. Scolnic, J. Casares, H. Chen, C. Conselice, J. Cooke, Z. Doctor, R. J. Foley, J. Horvath, D. A. Howell, C. D. Kilpatrick, C. Lidman, F. Olivares E., F. Paz-Chinchón, J. Pineda-G., J. Quirola-Vásquez, A. Rest, N. Sherman, T. M. C. Abbott, M. Aguena, S. Avila, E. Bertin, S. Bhargava, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, M. Costanzi, L. N. da Costa, S. Desai, H. T. Diehl, J. P. Dietrich, P. Doel, S. Everett, B. Flaugher, P. Fosalba, D. Friedel, J. Frieman, E. Gaztanaga, D. W. Gerdes, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, S. R. Hinton, D. L. Hollowood, K. Honscheid, D. J. James, K. Kuehn, N. Kuropatkin, O. Lahav, M. Lima, M. A. G. Maia, M. March, J. L. Marshall, F. Menanteau, R. Miquel, R. L. C. Ogando, A. A. Plazas, A. K. Romer, A. Roodman, E. Sanchez, V. Scarpine, M. Schubnell, S. Serrano, I. Sevilla-Noarbe, M. Smith, E. Suchyta, M. E. C. Swanson, G. Tarle, D. Thomas, T. N. Varga, A. R. Walker, J. Weller, and DES Collaboration. A DESGW Search for the Electromagnetic Counterpart to the LIGO/Virgo Gravitational-wave Binary Neutron Star Merger Candidate S190510g. *ApJ*, 903(1):75, Nov. 2020. doi: 10.3847/1538-4357/abb823.

- T. K. Ho. Random decision forests. In *Proceedings of 3rd international conference on document analysis and recognition*, volume 1, pages 278–282. IEEE, 1995.
- Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew, and et al. LSST: From Science Drivers to Reference Design and Anticipated Data Products. *ApJ*, 873:111, Mar. 2019. doi: 10.3847/1538-4357/ab042c.
- R. Kessler, J. P. Bernstein, D. Cinabro, B. Dilday, J. A. Frieman, S. Jha, S. Kuhlmann, G. Miknaitis, M. Sako, M. Taylor, and J. Vanderplas. Snana: A public software package for supernova analysis. *Publications of the Astronomical Society of the Pacific*, 121(883):1028, aug 2009. doi: 10.1086/605984. URL <https://dx.doi.org/10.1086/605984>.
- R. Kessler, G. Narayan, A. Avelino, E. Bachelet, R. Biswas, P. J. Brown, D. F. Chernoff, A. J. Connolly, M. Dai, S. Daniel, R. Di Stefano, M. R. Drout, L. Galbany, S. González-Gaitán, M. L. Graham, R. Hložek, E. E. O. Ishida, J. Guillochon, S. W. Jha, D. O. Jones, K. S. Mandel, D. Muthukrishna, A. O’Grady, C. M. Peters, J. R. Pierel, K. A. Ponder, A. Prša, S. Rodney, V. A. Villar, LSST Dark Energy Science Collaboration, and Transient and Variable Stars Science Collaboration. Models and Simulations for the Photometric LSST Astronomical Time Series Classification Challenge (PLAsTiCC). *PASP*, 131(1003):094501, Sept. 2019. doi: 10.1088/1538-3873/ab26f1.
- D. Muthukrishna, G. Narayan, K. S. Mandel, R. Biswas, and R. Hloek. Rapid: Early classification of explosive transients using deep learning. *Publications of the Astronomical Society of the Pacific*, 131(1005):118002, sep 2019. doi: 10.1088/1538-3873/ab1609. URL <https://dx.doi.org/10.1088/1538-3873/ab1609>.
- A. Möller, J. Peloton, E. E. O. Ishida, C. Arnault, E. Bachelet, T. Blaineau, D. Boutigny, A. Chauhan, E. Gangler, F. Hernandez, J. Hrivnac, M. Leoni, N. Leroy, M. Moniez, S. Pateyron, A. Ramparison, D. Turpin, R. Ansari, T. AllamăJr, A. Bajat, B. Biswas, A. Boucaud, J. Bregeon, J.-E. Campagne, J. Cohen-Tanugi, A. Coleiro, D. Dornic, D. Fouchez, O. Godet, P. Gris, S. Karpov, A. NebotăGomez-Moran, J. Neveu, S. Plaszczyński, V. Savchenko, and N. Webb. fink, a new generation of broker for the LSST community. *Monthly Notices of the Royal Astronomical Society*, 501(3):3272–3288, 11 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa3602. URL <https://doi.org/10.1093/mnras/staa3602>.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- B. F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323(6086):310–311, Sept. 1986. doi: 10.1038/323310a0.
- C. Stachie, M. W. Coughlin, N. Christensen, and D. Muthukrishna. Using machine learning for transient classification in searches for gravitational-wave counterparts. *Monthly Notices of the Royal Astronomical Society*, 497(2):1320–1331, 06 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1776. URL <https://doi.org/10.1093/mnras/staa1776>.
- V. A. Villar, J. Guillochon, E. Berger, B. D. Metzger, P. S. Cowperthwaite, M. Nicholl, K. D. Alexander, P. K. Blanchard, R. Chornock, T. Eftekhari, W. Fong, R. Margutti, and P. K. G. Williams. The Combined Ultraviolet, Optical, and Near-infrared Light Curves of the Kilonova Associated with the Binary Neutron Star Merger GW170817: Unified Data Set, Analytic Models, and Physical Implications. *ApJ*, 851(1):L21, Dec. 2017. doi: 10.3847/2041-8213/aa9c84.