










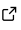
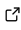
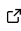
Caustics: A Python Package for Accelerated Strong Gravitational Lensing Simulations

Connor Stone ^{1,2,3*}, Alexandre Adam ^{1,2,3*}, Adam Coogan ^{1,2,3,a*}, M. J. Yantovski-Barth ^{1,2,3}, Andreas Filipp ^{1,2,3}, Landung Setiawan ⁴, Cordero Core ⁴, Ronan Legin ^{1,2,3}, Charles Wilson ^{1,2,3}, Gabriel Missael Barco ^{1,2,3}, Yashar Hezaveh ^{1,2,3,5,6,7}, and Laurence Perreault-Levasseur ^{1,2,3,5,6,7}

1 Ciela Institute - Montréal Institute for Astrophysical Data Analysis and Machine Learning, Montréal, Québec, Canada 2 Department of Physics, Université de Montréal, Montréal, Québec, Canada 3 Mila - Québec Artificial Intelligence Institute, Montréal, Québec, Canada 4 eScience Institute Scientific Software Engineering Center, 1410 NE Campus Pkwy, Seattle, WA 98195, USA 5 Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, 10010, New York, NY, USA 6 Perimeter Institute for Theoretical Physics, Waterloo, Canada 7 Trottier Space Institute, McGill University, Montréal, Canada a Work done while at UdeM, Ciela, and Mila  Corresponding author * These authors contributed equally.

DOI: [10.21105/joss.07081](https://doi.org/10.21105/joss.07081)

Software

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

Editor: [Ivelina Momcheva](#)  

Reviewers:

- [@andigu](#)
- [@Jammy2211](#)

Submitted: 19 June 2024

Published: 22 November 2024

License

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

Summary

Gravitational lensing is the deflection of light rays due to the gravity of intervening masses. This phenomenon is observed at a variety of configurations, involving any non-uniform mass such as planets, stars, galaxies, clusters of galaxies, and even the large-scale structure of the Universe. Strong lensing occurs when the distortions are significant and multiple images of the background source are observed. The lens and lensed object(s) must be aligned closely on the sky. As the discovery of lens systems has grown to the low thousands, these systems have become pivotal for precision measurements in astrophysics, notably for phenomena including dark matter (e.g. [Hezaveh et al., 2016](#); [Vegetti & Vogelsberger, 2014](#)), supernovae (e.g. [Rodney et al., 2021](#)), quasars (e.g. [Peng et al., 2006](#)), the first stars (e.g. [Welch et al., 2022](#)), and the Universe's expansion rate (e.g. [K. C. Wong et al., 2020](#)). With future surveys expected to discover hundreds of thousands of lensing systems, the modelling and simulation of such systems must be done at orders of magnitude larger scale than ever before. Here we present *caustics*, a Python package designed to facilitate machine learning and Bayesian methods to handle the extensive computational demands of modelling such a vast number of lensing systems.

Statement of need

The next generation of astronomical surveys, such as the Legacy Survey of Space and Time, the Roman Core Community Surveys, and the Euclid wide survey, are expected to uncover hundreds of thousands of gravitational lenses ([Collett, 2015](#)), dramatically increasing the scientific potential of gravitational lensing studies. Currently, analyzing a single lensing system can take several days or weeks, which will soon be infeasible to scale. Thus, advancements such as GPU acceleration and/or automatic differentiation are needed to reduce the analysis timescales. Machine learning will be critical to achieve the necessary speed to process these lenses. It will also be needed to meet the complexity of strong lens modelling. Literature on machine learning applications in strong gravitational lensing underscores this need ([Brehmer et al., 2019](#); [Chianese et al., 2020](#); [Coogan et al., 2020](#); [Karchev, Coogan, et al., 2022](#); [Karchev,](#)

Anau Montel, et al., 2022; Mishra-Sharma & Yang, 2022). caustics is built with the future of lensing in mind, using PyTorch (Paszke et al., 2019) to accelerate the low-level computation and enable deep learning algorithms, which rely on automatic differentiation.

Several other simulation packages for strong gravitational lensing are already publicly available. The well-established lenstronomy package has been in use since 2018 (Birrer et al., 2021); GLAMER is a C++-based code for modelling complex and large dynamic range fields (Metcalf & Petkova, 2014); PyAutoLens is also widely used (Nightingale et al., 2021); GIGA-Lens is a specialized JAX-based (Bradbury et al., 2018) gravitational lensing package (Gu et al., 2022); and Herculeans is a more general JAX-based lensing simulator package (Galan et al., 2022); among others (GRAVLENS Keeton, 2011; LENSTOOL Kneib et al., 2011; SLITRONOMY Galan et al., 2021; paltax Wagner-Carena et al., 2024). There are also several in-house codes developed for specialized analysis which are then not publicly released (e.g. Suyu & Halkola, 2010). The development of caustics has been primarily focused on three aspects: processing speed, user experience, and flexibility. Processing speed is comparable to the widely used lenstronomy when on CPU, and can be over 1000 times faster on GPU depending on configuration as seen in Figure 1. The user experience is streamlined by providing three interfaces to the code: configuration file, object-oriented, and functional. Flexibility is achieved by a determined focus on minimalism in the core functionality of caustics and encouraging user extension.




Figure 1: Runtime comparisons for a simple lensing setup. We compare the amount of time taken (y-axis) to generate a certain number of lensing realizations (x-axis) where a Sérsic model is lensed by an SIE mass distribution. For CPU calculations we use an Intel Gold 6148 Skylake and for the GPU we use a NVIDIA V100. All tests were done at 64-bit precision. On the left, the lensing system is sampled at 128-pixel resolution only at pixel midpoints. On the right, a more realistic simulation includes upsampled pixels and PSF convolution. From the two tests we see varying performance enhancements from compiled, unbatched, batched, multi-threaded, and GPU processing setups.

caustics fills a timely need for a differentiable lensing simulator. Several other fields have already benefited from such simulators, for example gravitational wave analysis (Coogan et al., 2022; Edwards et al., 2024; K. W. K. Wong et al., 2023), astronomical image photometry (Stone et al., 2023), point spread function modelling (Desdoigts et al., 2023); time series analysis (Million et al., 2024), and even generic optimization for scientific problems (Nikolic, 2018). With caustics it will now be possible to analyze over 100,000 lenses in a timely manner (Hezaveh et al., 2017; Perreault Levasseur et al., 2017).

Scope

caustics is a gravitational lensing simulator. The purpose of the project is to streamline the simulation of strong gravitational lensing effects on the light of a background source. The primary focus is on all transformations between the source plane(s) and the image plane

through the lensing plane(s). There is minimal effort on modelling the observational elements of the atmosphere or telescope optics. A variety of parametric lensing profiles are included, such as: Singular Isothermal Ellipsoid (SIE), Elliptical Power Law (EPL), Pseudo-Jaffe, Navarro-Frenk-White (NFW), and External Shear. Additionally, it offers non-parametric representations such as a gridded convergence or a potential field and pixelized sources.

Once a lensing system has been defined, caustics can then perform various computational operations on the system such as raytracing through the lensing system, forwards and backwards. Users can compute the lensing potential, convergence, deflection field, time delay field, shear, and magnification. All of these operations can readily be performed in a multi-plane setting to account for interlopers or multiple sources.

With these building blocks in place, one can construct fast and accurate simulators used to produce training sets for machine learning models or for inference on real-world systems. Neural networks have become a widespread tool for amortized inference of gravitational lensing parameters (Hezaveh et al., 2017) or in the detection of gravitational lenses (Huang et al., 2021; Petrillo et al., 2017), but they require large and accurate training sets that can be created quickly with caustics. The simulators are differentiable, enabling algorithms such as recurrent inference machines (Adam et al., 2023) and diffusion models (Adam et al., 2022; Remy et al., 2023).

The scope of caustics ends at lensing simulation, thus it does not include functionality to optimize or sample the resulting functions. Users are encouraged to use already existing optimization and sampling codes such as `scipy.optimize` (Virtanen et al., 2020), `emcee` (Foreman-Mackey et al., 2013), `dynesty` (Speagle, 2020), `Pyro` (Bingham et al., 2019), and `torch.optim` (Paszke et al., 2019). Interfacing with these codes is easy and demonstrations are included in the documentation.

Further, caustics does not implement simulators for all possible lensing problems (AGN microlensing, multi-source lensing, supernova time-delay cosmography, etc.). Instead it is formatted much like PyTorch where one constructs a class (Module) and builds a forward model function by calling individual (often functional) components defined within caustics. In this way it can always be adapted to the specific needs of a lensing problem.

Acknowledgements

This research was enabled by a generous donation by Eric and Wendy Schmidt with the recommendation of the Schmidt Futures Foundation. We acknowledge the generous software support of the University of Washington Scientific Software Engineering Center (SSEC) at the eScience Institute, via matching through the Schmidt Futures Virtual Institute for Scientific Software (VISS). CS acknowledges the support of a NSERC Postdoctoral Fellowship and a CITA National Fellowship. This research was enabled in part by support provided by Calcul Québec and the Digital Research Alliance of Canada. The work of A.A. and R.L. was partially funded by NSERC CGS D scholarships. Y.H. and L.P. acknowledge support from the National Sciences and Engineering Council of Canada grants RGPIN-2020-05073 and 05102, the Fonds de recherche du Québec grants 2022-NC-301305 and 300397, and the Canada Research Chairs Program. Thanks to Simon Birrer for communications regarding benchmarking of `lenstronomy`. Thanks to James Nightingale and Andi Gu for thoughtful reviews of the codebase, and thanks to Ivelina Momcheva for a thorough editorial review.

References

Adam, A., Coogan, A., Malkin, N., Legin, R., Perreault-Levasseur, L., Hezaveh, Y., & Bengio, Y. (2022). Posterior samples of source galaxies in strong gravitational lenses

- with score-based priors. *Machine Learning and the Physical Sciences Workshop*, E1. <https://doi.org/10.48550/arXiv.2211.03812>
- Adam, A., Perreault-Levasseur, L., Hezaveh, Y., & Welling, M. (2023). Pixelated reconstruction of foreground density and background surface brightness in gravitational lensing systems using recurrent inference machines. *The Astrophysical Journal*, 951(1), 6. <https://doi.org/10.3847/1538-4357/acf84>
- Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P., Horsfall, P., & Goodman, N. D. (2019). Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, 20(1), 973–978.
- Birrer, S., Shajib, A. J., Gilman, D., Galan, A., Aalbers, J., Million, M., Morgan, R., Pagano, G., Park, J. W., Teodori, L., Tessore, N., Ueland, M., Vyvere, L. V. de, Wagner-Carena, S., Wempe, E., Yang, L., Ding, X., Schmidt, T., Sluse, D., ... Amara, A. (2021). Lenstronomy II: A gravitational lensing software ecosystem. *Journal of Open Source Software*, 6(62), 3283. <https://doi.org/10.21105/joss.03283>
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leery, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., & Zhang, Q. (2018). *JAX: Composable transformations of Python+NumPy programs* (Version 0.3.13). <http://github.com/google/jax>
- Brehmer, J., Mishra-Sharma, S., Hermans, J., Louppe, G., & Cranmer, K. (2019). Mining for dark matter substructure: Inferring subhalo population properties from strong lenses with machine learning. *The Astrophysical Journal*, 886(1), 49. <https://doi.org/10.3847/1538-4357/ab4c41>
- Chianese, M., Coogan, A., Hofma, P., Otten, S., & Weniger, C. (2020). Differentiable strong lensing: uniting gravity and neural nets through differentiable probabilistic programming. *Monthly Notices of the Royal Astronomical Society*, 496(1), 381–393. <https://doi.org/10.1093/mnras/staa1477>
- Collett, T. E. (2015). The population of galaxy–galaxy strong lenses in forthcoming optical imaging surveys. *The Astrophysical Journal*, 811(1), 20. <https://doi.org/10.1088/0004-637X/811/1/20>
- Coogan, A., Edwards, T. D. P., Chia, H. S., George, R. N., Freese, K., Messick, C., Setzer, C. N., Weniger, C., & Zimmerman, A. (2022). Efficient gravitational wave template bank generation with differentiable waveforms. *Physical Review D*, 106(12), 122001. <https://doi.org/10.1103/PhysRevD.106.122001>
- Coogan, A., Karchev, K., & Weniger, C. (2020). Targeted likelihood-free inference of dark matter substructure in strongly-lensed galaxies. *arXiv e-Prints*, arXiv:2010.07032. <https://doi.org/10.48550/arXiv.2010.07032>
- Desdoigts, L., Pope, B. J. S., Dennis, J., & Tuthill, P. G. (2023). Differentiable optics with ∂ Lux: I—deep calibration of flat field and phase retrieval with automatic differentiation. *Journal of Astronomical Telescopes, Instruments, and Systems*, 9(2), 028007. <https://doi.org/10.1117/1.JATIS.9.2.028007>
- Edwards, T. D. P., Wong, K. W. K., Lam, K. K. H., Coogan, A., Foreman-Mackey, D., Isi, M., & Zimmerman, A. (2024). Differentiable and hardware-accelerated waveforms for gravitational wave data analysis. *Physical Review D*, 110(6), 064028. <https://doi.org/10.1103/PhysRevD.110.064028>
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. (2013). emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 125(925), 306. <https://doi.org/10.1086/670067>
- Galan, A., Peel, A., Joseph, R., Courbin, F., & Starck, J.-L. (2021). SLITRONOMY: Towards

- a fully wavelet-based strong lensing inversion technique. *Astronomy & Astrophysics*, 647, A176. <https://doi.org/10.1051/0004-6361/202039363>
- Galan, A., Vernardos, G., Peel, A., Courbin, F., & Starck, J.-L. (2022). Using wavelets to capture deviations from smoothness in galaxy-scale strong lenses. *Astronomy & Astrophysics*, 668, A155. <https://doi.org/10.1051/0004-6361/202244464>
- Gu, A., Huang, X., Sheu, W., Aldering, G., Bolton, A. S., Boone, K., Dey, A., Filipp, A., Jullo, E., Perlmutter, S., Rubin, D., Schlafly, E. F., Schlegel, D. J., Shu, Y., & Suyu, S. H. (2022). GIGA-Lens: Fast Bayesian inference for strong gravitational lens modeling. *The Astrophysical Journal*, 935(1), 49. <https://doi.org/10.3847/1538-4357/ac6de4>
- Hezaveh, Y., Dalal, N., Holder, G., Kisner, T., Kuhlen, M., & Perreault Levasseur, L. (2016). Measuring the power spectrum of dark matter substructure using strong gravitational lensing. *Journal of Cosmology and Astroparticle Physics*, 2016(11), 048. <https://doi.org/10.1088/1475-7516/2016/11/048>
- Hezaveh, Y., Perreault Levasseur, L., & Marshall, P. J. (2017). Fast automated analysis of strong gravitational lenses with convolutional neural networks. *Nature*, 548(7669), 555–557. <https://doi.org/10.1038/nature23463>
- Huang, X., Storfer, C., Gu, A., Ravi, V., Pilon, A., Sheu, W., Venguswamy, R., Banka, S., Dey, A., Landriau, M., Lang, D., Meisner, A., Moustakas, J., Myers, A. D., Sajith, R., Schlafly, E. F., & Schlegel, D. J. (2021). Discovering new strong gravitational lenses in the DESI Legacy Imaging Surveys. *The Astrophysical Journal*, 909(1), 27. <https://doi.org/10.3847/1538-4357/abd62b>
- Karchev, K., Anau Montel, N., Coogan, A., & Weniger, C. (2022). Strong-lensing source reconstruction with denoising diffusion restoration models. *arXiv e-Prints*, arXiv:2211.04365. <https://doi.org/10.48550/arXiv.2211.04365>
- Karchev, K., Coogan, A., & Weniger, C. (2022). Strong-lensing source reconstruction with variationally optimized Gaussian processes. *Monthly Notices of the Royal Astronomical Society*, 512(1), 661–685. <https://doi.org/10.1093/mnras/stac311>
- Keeton, C. R. (2011). *GRAVLENS: Computational methods for gravitational lensing*. Astrophysics Source Code Library, record ascl:1102.003.
- Kneib, J.-P., Bonnet, H., Golse, G., Sand, D., Jullo, E., & Marshall, P. (2011). *LENSTOOL: A gravitational lensing software for modeling mass distribution of galaxies and clusters (strong and weak regime)*. Astrophysics Source Code Library, record ascl:1102.004.
- Metcalfe, R. B., & Petkova, M. (2014). GLAMER - I. A code for gravitational lensing simulations with adaptive mesh refinement. *Monthly Notices of the Royal Astronomical Society*, 445(2), 1942–1953. <https://doi.org/10.1093/mnras/stu1859>
- Million, M., Michalewicz, K., Dux, F., Courbin, F., & Marshall, P. J. (2024). Image deconvolution and point-spread function reconstruction with STARRED: A wavelet-based two-channel method optimized for light-curve extraction. *The Astronomical Journal*, 168(2), 55. <https://doi.org/10.3847/1538-3881/ad4da7>
- Mishra-Sharma, S., & Yang, G. (2022). Strong lensing source reconstruction using continuous neural fields. *Machine Learning for Astrophysics*, 34. <https://doi.org/10.48550/arXiv.2206.14820>
- Nightingale, James. W., Hayes, R. G., Kelly, A., Amvrosiadis, A., Etherington, A., He, Q., Li, N., Cao, X., Frawley, J., Cole, S., Enia, A., Frenk, C. S., Harvey, D. R., Li, R., Massey, R. J., Negrello, M., & Robertson, A. (2021). 'PyAutoLens': Open-source strong gravitational lensing. *Journal of Open Source Software*, 6(58), 2825. <https://doi.org/10.21105/joss.02825>
- Nikolic, B. (2018). Acceleration of Non-Linear Minimisation with PyTorch. *arXiv e-Prints*,

- arXiv:1805.07439. <https://doi.org/10.48550/arXiv.1805.07439>
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., ... Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems 32* (pp. 8024–8035). Curran Associates, Inc. <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- Peng, C. Y., Impey, C. D., Rix, H.-W., Kochanek, C. S., Keeton, C. R., Falco, E. E., Lehár, J., & McLeod, B. A. (2006). Probing the coevolution of supermassive black holes and galaxies using gravitationally lensed quasar hosts. *The Astrophysical Journal*, *649*(2), 616–634. <https://doi.org/10.1086/506266>
- Perreault Levasseur, L., Hezaveh, Y. D., & Wechsler, R. H. (2017). Uncertainties in parameters estimated with neural networks: Application to strong gravitational lensing. *The Astrophysical Journal Letters*, *850*(1), L7. <https://doi.org/10.3847/2041-8213/aa9704>
- Petrillo, C. E., Tortora, C., Chatterjee, S., Vernardos, G., Koopmans, L. V. E., Verdoes Kleijn, G., Napolitano, N. R., Covone, G., Schneider, P., Grado, A., & McFarland, J. (2017). Finding strong gravitational lenses in the Kilo Degree Survey with Convolutional Neural Networks. *Monthly Notices of the Royal Astronomical Society*, *472*(1), 1129–1150. <https://doi.org/10.1093/mnras/stx2052>
- Remy, B., Lanusse, F., Jeffrey, N., Liu, J., Starck, J.-L., Osato, K., & Schrabback, T. (2023). Probabilistic mass-mapping with neural score estimation. *Astronomy & Astrophysics*, *672*, A51. <https://doi.org/10.1051/0004-6361/202243054>
- Rodney, S. A., Brammer, G. B., Pierel, J. D. R., Richard, J., Toft, S., O'Connor, K. F., Akhshik, M., & Whitaker, K. E. (2021). A gravitationally lensed supernova with an observable two-decade time delay. *Nature Astronomy*, *5*, 1118–1125. <https://doi.org/10.1038/s41550-021-01450-9>
- Speagle, J. S. (2020). DYNESTY: a dynamic nested sampling package for estimating Bayesian posteriors and evidences. *Monthly Notices of the Royal Astronomical Society*, *493*(3), 3132–3158. <https://doi.org/10.1093/mnras/staa278>
- Stone, C. J., Courteau, S., Cuillandre, J.-C., Hezaveh, Y., Perreault-Levasseur, L., & Arora, N. (2023). ASTROPHOT: fitting everything everywhere all at once in astronomical images. *Monthly Notices of the Royal Astronomical Society*, *525*(4), 6377–6393. <https://doi.org/10.1093/mnras/stad2477>
- Suyu, S. H., & Halkola, A. (2010). The halos of satellite galaxies: the companion of the massive elliptical lens SL2S J08544-0121. *Astronomy & Astrophysics*, *524*, A94. <https://doi.org/10.1051/0004-6361/201015481>
- Vegetti, S., & Vogelsberger, M. (2014). On the density profile of dark matter substructure in gravitational lens galaxies. *Monthly Notices of the Royal Astronomical Society*, *442*(4), 3598–3603. <https://doi.org/10.1093/mnras/stu1284>
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*, 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Wagner-Carena, S., Lee, J., Pennington, J., Aalbers, J., Birrer, S., & Wechsler, R. H. (2024). A strong gravitational lens is worth a thousand dark matter halos: Inference on small-scale structure using sequential methods. *The Astrophysical Journal*, *975*(2), 297. <https://doi.org/10.3847/1538-4357/ad6e70>

- Welch, B., Coe, D., Zackrisson, E., de Mink, S. E., Ravindranath, S., Anderson, J., Brammer, G., Bradley, L., Yoon, J., Kelly, P., Diego, J. M., Windhorst, R., Zitrin, A., Dimauro, P., Jiménez-Teja, Y., Abdurro'uf, Nonino, M., Acebron, A., Andrade-Santos, F., ... Vikaeus, A. (2022). JWST imaging of Earendel, the extremely magnified star at redshift $z = 6.2$. *The Astrophysical Journal Letters*, 940(1), L1. <https://doi.org/10.3847/2041-8213/ac9d39>
- Wong, K. C., Suyu, S. H., Chen, G. C.-F., Rusu, C. E., Million, M., Sluse, D., Bonvin, V., Fassnacht, C. D., Taubenberger, S., Auger, M. W., Birrer, S., Chan, J. H. H., Courbin, F., Hilbert, S., Tihhonova, O., Treu, T., Agnello, A., Ding, X., Jee, I., ... Meylan, G. (2020). H0LiCOW - XIII. A 2.4 per cent measurement of H_0 from lensed quasars: 5.3σ tension between early- and late-Universe probes. *Monthly Notices of the Royal Astronomical Society*, 498(1), 1420–1439. <https://doi.org/10.1093/mnras/stz3094>
- Wong, K. W. K., Isi, M., & Edwards, T. D. P. (2023). Fast gravitational-wave parameter estimation without compromises. *The Astrophysical Journal*, 958(2), 129. <https://doi.org/10.3847/1538-4357/acf5cd>