LINNA: Likelihood Inference Neural Network Accelerator

Chun-Hao To [*1] **Eduardo Rozo** [*2]

Abstract

Bayesian posterior inference of modern multi-probe cosmological analyses incurs massive computational costs. These computational costs have severe environmental impacts and the long wall-clock time slows scientific productivity. To address these difficulties, we introduce LINNA: the Likelihood Inference Neural Network Accelerator. Relative to the baseline of modern survey cosmological analyses, LINNA reduces the computational cost associated with posterior inference by a factor of 8–50. To accomplish these reductions, LINNA automatically builds training data sets, creates neural network emulators, and produces a Markov chain that samples the posterior. We explicitly verify that LINNA accurately reproduces the first-year DES cosmological constraints derived from a variety of different data vectors with our default code settings, without needing to retune the algorithm every time. Further, we find that LINNA is sufficient for enabling accurate and efficient sampling for LSST Y10 multi-probe analyses.

1. Introduction

Modern cosmological analyses typically employ a Bayesian framework, where posteriors are sampled using Markov Chain Monte Carlo (MCMC) techniques. Despite significant progress in the development of MCMC sampling algorithms (e.g. Jia & Seljak, 2022; Karamanis et al., 2021), the computational requirements associated with performing these calculations can be significant. This is especially true for multi-probe analyses (e.g. DES Collaboration et al., 2021; van Uitert et al., 2018; Joudaki et al., 2018), which typically require the introduction of many tens of nuisance parameters to account for systematic effects presented in the data. As a specific example, the joint analysis of clusters, weak lensing, and galaxy clustering of To et al. 2021a required 26 nuisance parameters. To achieve a well-sampled posterior, the chains in that analysis required at least five million likelihood evaluations, totaling $\sim 50k$ CPU hours. That is, a single run had a wall-clock time of *three weeks* using a state-of-the-art computing cluster with 100 cores.

The primary difficulty in this type of analysis is that the theory model used to evaluate the likelihood is computationally expensive. One way to bypass this limitation is to rely on fast surrogates of said model (e.g. Auld et al., 2008; Agarwal et al., 2014; Pellejero-Ibañez et al., 2020; Manrique-Yus & Sellentin, 2020; Spurio Mancini et al., 2022; DeRose et al., 2021). However, the surrogates are typically built in pre-defined parameter spaces that are narrower than those used in most cosmological analyses. That is, the use of surrogates that are constructed *a priori* either limits the choice of priors that can be employed, or requires extrapolations of the surrogate outside the parameter regions over which the surrogate was calibrated.

To address these shortcomings, we present a new posterior inference tool: the Likelihood Inference Neural Network Accelerator (LINNA). Given the standard inputs of MCMC samplers, LINNA can accurately generate MCMC samples of posteriors for different combinations of cosmological probes without the need to tune for each combination. Our contributions are as follows:

- We design LINNA to enable fast and accurate posterior inference with arbitrary priors through the use of emulators of the theory model. The emulators of the theory model are based on a neural network that is trained automatically using automatically generated training data.
- We demonstrate that LINNA succeeds in accurately recovering the parameter posteriors derived from three different combinations of two point correlations of galaxies, shear and clusters, measured in the Dark Energy Survey (Drlica-Wagner et al., 2018) year 1 analyses (DES Y1).
- We show that LINNA accurately reproduces the forecasted constraints for Rubin Observa- tory's Legacy Survey of Space and Time (LSST) Year 10 data set.

[*]Equal contribution [1]Center for Cosmology and AstroParticle Physics (CCAPP), Ohio State University, Columbus, OH 43210, USA [2]Department of Physics, University of Arizona, Tucson, AZ 85721, USA. Correspondence to: Chun-Hao To <to.87@osu.edu>, Eduardo Rozo <erozo@arizona.edu>.

ICML 2022 Workshop on Machine Learning for Astrophysics, Baltimore, Maryland, USA, 2022.

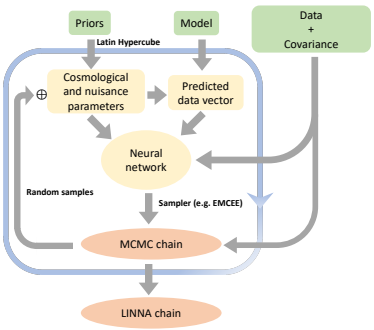

Figure 1. Outline of LINNA structure. Green blocks indicate inputs provided by users and orange blocks represent outputs of LINNA. Blocks within the blue rectangular are run iteratively. The symbol $\bigoplus$ represents the procedure of expanding the training data by randomly sampling the MCMC chain from the previous iteration.

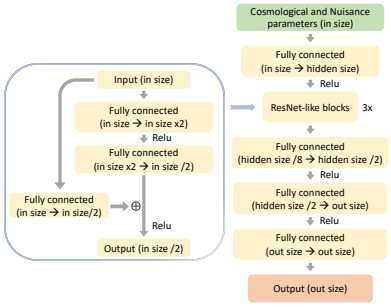

Figure 2. The architecture of the neural network in LINNA. Green blocks represent input of the network and orange blocks represent the output. Numbers in paratheses show the size of input and output vectors of each block. Hidden size is set to the minimum value of 32 times output vector size and 1000. Relu represents the Rectified Linear Unit, a widely used nonlinear function. The symbol $\bigoplus$ denotes element-wise addition.

2. Methods

LINNA design. LINNA uses neural networks to create emulators of computationally expensive theory models. This emulator can be used in combination with standard MCMC sampling algorithms to quickly recover accurate approximations of the parameter posteriors. Because the posteriors are sampled using standard techniques, we may readily use well-tested convergence tests of the resulting chains.

While the concept of LINNA is simple, there is one major challenge. In high-dimensional parameter spaces, the volume of reasonably defined priors can greatly exceed the volume of the posteriors. Consequently, uniform sampling of the prior is extremely inefficient: it is likely that only a few training points fall within one- to three- sigma region of the posteriors, where high accuracy of the neural network emulators is most desirable. We solve this problem using an iterative approach. In particular, a crude neural network is trained on a training set that uniformly samples the priors. This neural network is then fed into a posterior sampling tool to produce a MCMC chain, from which additional training data is generated. In this way, we construct a training set that covers the entire prior volume and is also capable of accurately reconstructing the parameter posteriors with comparatively little training data.

In the first few iterations, the neural network emulators are rough at best, so their high-confidence regions can be offset from those of the true posterior. To overcome this problem, for the first few iterations of the algorithm, we enlarge the parameter posteriors using a temperature parameter T such that

$$\chi^2_T = \chi^2_{\rm NN}/T^2, \tag{1}$$

where $\chi^2_{\rm NN}$ is the χ^2 value of the data calculated using neural network emulators. The resulting broadening allows for the posteriors of the neural network emulator to be offset from the true posterior while still completely covering the same. In this way, when we generate new training data from the surrogate posterior we are able to improve the performance of the network over the true parameter posterior, even if the latter was offset from that of the approximate model. Initially, the temperature T should be set to a large value to ensure that the training data spans a wide enough parameter space. In the subsequent iterations, T can be gradually decreased due to the increasing accuracy of the neural network model over the region of interest. Empirically, we find that setting $T = 4$ and 2 for the first two iterations and $T = 1$ for all subsequent iterations leads to good posteriors.

We summarize the workflow for LINNA in figure 1. First, we use Latin Hypercube sampling of the prior volume space to generate an initial set of training points. At each training point, we use the provided theory model to evaluate the expectation value of the observable. The resulting training points and model expectations are used to train a neural network to arrive at our first emulator. This emulator in combination with the provided likelihood function, which is scaled by the appropriate temperature parameter, is fed to a standard sampling algorithm to arrive at MCMC samples of the posterior. The resulting chain is then randomly sampled to expand the number of points in the parameter space at which the neural network is trained. The procedure is iterated, decreasing the temperature from iteration to iteration as described above.

We run LINNA using the two publicly available posterior sampling code EMCEE (Foreman-Mackey et al., 2013) and ZEUS (Karamanis et al., 2021). The convergence criteria are set based on the integrated autocorrelation time τ (Goodman & Weare, 2010).

Emulators of theory models. The emulator of theory mod-

els in LINNA is a neural network, whose architecture is shown in figure 2.

The loss function is judiciously chosen based on the task at hand, namely recovering accurate posteriors from the data. To that end, the loss function is motivated by two observations. First, the accuracy requirement of the emulator is set by the noise of the data, i.e., noisy data do not require accurate models. Consequently, instead of using a uniformly weighted loss function, we weight the differences between the training data and the emulator by the inverse of the covariance matrix of the data. Second, for the purposes of sampling posteriors, the emulator needs only to achieve high accuracy over high-likelihood regions. For this reason, we upweight the importance of those points in the parameter space that are a good fit to the data by dividing their contribution to the loss function by the χ^2 of the model. Our loss function is therefore defined as

$$Loss = \left\langle \frac{(NN - M)^T C^{-1} (NN - M)}{(M - d)^T C^{-1} (M - d)} \right\rangle , \qquad (2)$$

where NN denotes the output of the neural network, M represents the (actual) model prediction, C is the covariance matrix, d is the data vector, and $\langle \dots \rangle$ refers to the mean over training points.

3. Results and Discussion

We now test the performance of LINNA in a real-world scenario. Specifically, we compare the posterior estimated using LINNA to that estimated using EMCEE to sample the posterior function implemented in COSMOLIKE (Krause & Eifler, 2017) (hereafter, the "*brute force method*"). Consequently, for the purpose of this comparison, the results presented here are obtained using EMCEE as the LINNA sampler and COSMOLIKE as the LINNA theory model. In practice, further improvements can be achieved using faster posterior sampling algorithms such as ZEUS. We perform all analyses using 128 CPU cores and 1 NVIDIA GeForce RTX 2080 Ti GPU. We consider three survey cosmology data vectors measured from first-year observations of DES. Ranked by their computational cost, the three data sets are:

1. 6×2pt+N (To et al., 2021a): a joint analysis of all data vectors in 3×2pt and 4×2pt+N.
2. 4×2pt+N (To et al., 2021a;b): a joint analysis of galaxy clustering, cluster–galaxy cross correlations, cluster clustering, cluster lensing, and cluster abundances.
3. 3×2pt (Abbott et al., 2018): a joint analysis of galaxy clustering, galaxy–galaxy lensing, and cosmic shear.

In these analyses, we sample 32 (28, 26) cosmological and nuisance parameters for 6×2pt+N (4×2pt+N, 3×2pt) with the same priors presented in (To et al., 2021a).

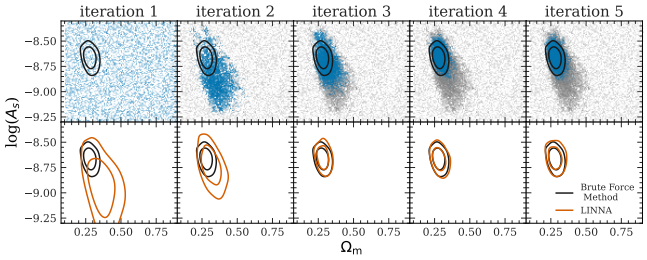

Figure 3. An illustration of the LINNA algorithm in the A_s–Ω_m parameter subspace, when applied to the 6×2pt+N data set. Top row: distributions of training data. Black contours show the 68% and 95% contours of the targeted posteriors estimated using the *brute force method*. Bottom rows: comparison of LINNA posteriors (orange) to the target posterior (black) at each iteration. We note that while we only show Ω_m and A_s in this plot, LINNA shows strong consistency with the *brute force method* after three iterations in all other cosmological and nuisance parameters as well.

6×2pt+N analysis. Figure 3 shows the 68% and 95% constraints on A_s and Ω_m using LINNA at each iteration. The blue points in the top row correspond to the $10,000$ newly added training data points at each iteration, while the grey points represent the cumulative training data from previous iterations. In the first iteration, the training set is generated by uniformly sampling the prior volume using a Latin Hypercube. For the purposes of Latin Hypercube sampling, the boundaries of the parameters with Gaussian priors are set to $\pm n\sigma$ around their mean values, where n is a user-specified parameter. Throughout the analysis, we adopt $n = 3$ because the Gaussian priors are only applied to nuisance parameters, whose priors are conservatively chosen. We see that LINNA converges in just 3 iterations, though we advocate running an additional high-precision iteration to verify the convergence. For this reason, we adopt four iterations as our standard LINNA setting.

Figure 4a compares the final constraints obtained with LINNA (orange contours) to those derived using the *brute force method* (black contours). The comparison is restricted to the most important parameters in the analyses: Ω_m, A_s, two nonlinear intrinsic alignment parameters (A_{IA}, η_{IA}), and four richness–mass relation parameters ($\ln\lambda_0, A_{\ln\lambda}, B_{\ln\lambda}, \sigma_{\ln\lambda}$). We find excellent agreement between the two posteriors, with a shift in the mean parameter constraints of less than 0.2σ. The one-sigma errors on individual parameters all agree to better than 12%. Critically, LINNA ran $\sim$ 10 hours, compared to $\sim$ 3 weeks for the *brute force method*. As noted earlier, the run time can be further reduced by decreasing the number of training data and by replacing EMCEE with ZEUS in LINNA. Specifically, LINNA can reach similar performances on all constrained cosmological and nuisance parameters with the run time

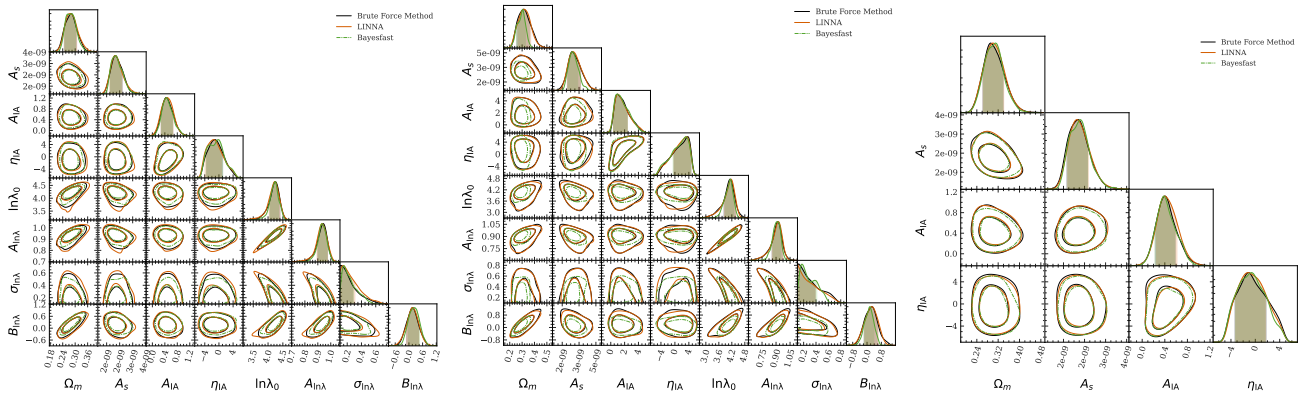

(a) 6×2pt+N data vector. (b) 4×2pt+N data vector. (c) 3×2pt data vector.

Figure 4. Comparison of LINNA, the *brute force method*, and BAYESFAST. Plots show constraints on important parameters when fitting to data after marginalizing over additional nuisance parameters. Black contours are obtained with the *brute force method* (EMCEE+COSMOLIKE), orange contours are obtained with LINNA, and green contours are obtained using BAYESFAST. Contours show 68% and 95% constraints.

down to 5.5 hours.

We also compare our results to those obtained using BAYESFAST (Jia & Seljak, 2022) [1], another posterior inference acceleration code currently in development (green contours in figure 4a). Using the same MCMC convergence criteria and hardware as LINNA, we find that BAYESFAST converges in 5.8 hours, a run time comparable to that of LINNA using ZEUS with 2000 training points. However, figure 4a demonstrates that the accuracy in the recovered posteriors of BAYESFAST is slightly worse than LINNA.

4×2pt+N and 3×2pt analyses. Figure 4b compares the posteriors derived for the 4×2pt+N data vector using LINNA (orange contours), the *brute force method* (black contours) and BAYESFAST (green contours). As before, we restrict the figure to the most important parameters for the 4×2pt+N data vector. All samplers are run with exactly the same setting as for the 6×2pt+N analysis. The agreement between LINNA and the *brute force method* is remarkable and demonstrates that LINNA's hyperparameters do not need to be retuned when running in subsets of the original data used to tune LINNA. This is not a trivial statement: BAYESFAST posteriors, which were only modestly biased for 6×2pt+N data set, are now much more strongly biased than before. Figure 4c compares the posteriors derived for the 3×2pt data vector from LINNA (orange contours), the *brute force method* (black contours), and BAYESFAST (green contours). Here, LINNA and BAYESFAST are both in excellent agreement with the *brute force method*.

As was the case for 6×2pt+N, using LINNA for these analyses leads to significant reductions in run time relative to the *brute force method*. With the baseline setting, LINNA takes 11.5 (5.4) hours to sample the 4×2pt+N (3×2pt) posteriors, while the *brute force method* takes ~ 4 (1) days respectively. Again, LINNA can be further sped up while maintaining the same accuracy on all constrained parameters by decreasing the number of training data and by using faster samplers. With the same fast setting as used for 6×2pt+N, the run time of LINNA is 4.2 and 3.5 hours for 4×2pt+N and 3×2pt respectively.

6×2pt+N LSST Y10 simulated analysis. We construct the data vector using COSMOLIKE, using parameters described in (The LSST Dark Energy Science Collaboration et al., 2018). We perform the same comparison as being done for the DES Y1 analysis. We find that the level of agreement is consistent with what we found in analyzing DES Y1 data. These results indicate that LINNA is expected to be applicable on the analysis of data with LSST Y10 precision.

4. Conclusion

We present a new posterior sampling tool LINNA that combines a neural network and standard MCMC posterior samplers to accelerate the posterior inferences of survey cosmology analyses. With the specially designed iterative training process and the judiciously chosen loss function, we find that this combination can reduce the computational time of posterior inferences for various different survey cosmology analyses by a factor 8–50. We test the performance of LINNA using different combinations of DES Y1 data and LSST Y10 simulated data and find strong consistency with the traditional brute force method. We conclude that LINNA is widely applicable for various cosmological analyses in current and future surveys.

[1] https://github.com/HerculesJack/bayesfast

References

Abbott, T. M. C., Abdalla, F. B., Alarcon, A., Aleksić, J., Allam, S., Allen, S., Amara, A., Annis, J., Asorey, J., Avila, S., Bacon, D., Balbinot, E., Banerji, M., Banik, N., Barkhouse, W., Baumer, M., Baxter, E., Bechtol, K., Becker, M. R., Benoit-Lévy, A., Benson, B. A., Bernstein, G. M., Bertin, E., Blazek, J., Bridle, S. L., Brooks, D., Brout, D., Buckley-Geer, E., Burke, D. L., Busha, M. T., Campos, A., Capozzi, D., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Castander, F. J., Cawthon, R., Chang, C., Chen, N., Childress, M., Choi, A., Conselice, C., Crittenden, R., Crocce, M., Cunha, C. E., D'Andrea, C. B., da Costa, L. N., Das, R., Davis, T. M., Davis, C., De Vicente, J., DePoy, D. L., DeRose, J., Desai, S., Diehl, H. T., Dietrich, J. P., Dodelson, S., Doel, P., Drlica-Wagner, A., Eifler, T. F., Elliott, A. E., Elsner, F., Elvin-Poole, J., Estrada, J., Evrard, A. E., Fang, Y., Fernandez, E., Ferté, A., Finley, D. A., Flaugher, B., Fosalba, P., Friedrich, O., Frieman, J., García-Bellido, J., Garcia-Fernandez, M., Gatti, M., Gaztanaga, E., Gerdes, D. W., Giannantonio, T., Gill, M. S. S., Glazebrook, K., Goldstein, D. A., Gruen, D., Gruendl, R. A., Gschwend, J., Gutierrez, G., Hamilton, S., Hartley, W. G., Hinton, S. R., Honscheid, K., Hoyle, B., Huterer, D., Jain, B., James, D. J., Jarvis, M., Jeltema, T., Johnson, M. D., Johnson, M. W. G., Kacprzak, T., Kent, S., Kim, A. G., King, A., Kirk, D., Kokron, N., Kovacs, A., Krause, E., Krawiec, C., Kremin, A., Kuehn, K., Kuhlmann, S., Kuropatkin, N., Lacasa, F., Lahav, O., Li, T. S., Liddle, A. R., Lidman, C., Lima, M., Lin, H., MacCrann, N., Maia, M. A. G., Makler, M., Manera, M., March, M., Marshall, J. L., Martini, P., McMahon, R. G., Melchior, P., Menanteau, F., Miquel, R., Miranda, V., Mudd, D., Muir, J., Möller, A., Neilsen, E., Nichol, R. C., Nord, B., Nugent, P., Ogando, R. L. C., Palmese, A., Peacock, J., Peiris, H. V., Peoples, J., Percival, W. J., Petravick, D., Plazas, A. A., Porredon, A., Prat, J., Pujol, A., Rau, M. M., Refregier, A., Ricker, P. M., Roe, N., Rollins, R. P., Romer, A. K., Roodman, A., Rosenfeld, R., Ross, A. J., Rozo, E., Rykoff, E. S., Sako, M., Salvador, A. I., Samuroff, S., Sánchez, C., Sanchez, E., Santiago, B., Scarpine, V., Schindler, R., Scolnic, D., Secco, L. F., Serrano, S., Sevilla-Noarbe, I., Sheldon, E., Smith, R. C., Smith, M., Smith, J., Soares-Santos, M., Sobreira, F., Suchyta, E., Tarle, G., Thomas, D., Troxel, M. A., Tucker, D. L., Tucker, B. E., Uddin, S. A., Varga, T. N., Vielzeuf, P., Vikram, V., Vivas, A. K., Walker, A. R., Wang, M., Wechsler, R. H., Weller, J., Wester, W., Wolf, R. C., Yanny, B., Yuan, F., Zenteno, A., Zhang, B., Zhang, Y., Zuntz, J., and Dark Energy Survey Collaboration. Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. Phys. Rev. D, 98(4):043526, August 2018. doi: 10.1103/PhysRevD.98.043526.

Agarwal, S., Abdalla, F. B., Feldman, H. A., Lahav, O., and Thomas, S. A. PkANN - II. A non-linear matter power spectrum interpolator developed using artificial neural networks. MNRAS, 439(2):2102–2121, April 2014. doi: 10.1093/mnras/stu090.

Auld, T., Bridges, M., and Hobson, M. P. COSMONET: fast cosmological parameter estimation in non-flat models using neural networks. MNRAS, 387(4):1575–1582, July 2008. doi: 10.1111/j.1365-2966.2008.13279.x.

DeRose, J., Chen, S.-F., White, M., and Kokron, N. Neural Network Acceleration of Large-scale Structure Theory Calculations. *arXiv e-prints*, art. arXiv:2112.05889, December 2021.

DES Collaboration, Abbott, T. M. C., Aguena, M., Alarcon, A., Allam, S., Alves, O., Amon, A., Andrade-Oliveira, F., Annis, J., Avila, S., Bacon, D., Baxter, E., Bechtol, K., Becker, M. R., Bernstein, G. M., Bhargava, S., Birrer, S., Blazek, J., Brandao-Souza, A., Bridle, S. L., Brooks, D., Buckley-Geer, E., Burke, D. L., Camacho, H., Campos, A., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Castander, F. J., Cawthon, R., Chang, C., Chen, A., Chen, R., Choi, A., Conselice, C., Cordero, J., Costanzi, M., Crocce, M., da Costa, L. N., da Silva Pereira, M. E., Davis, C., Davis, T. M., De Vicente, J., DeRose, J., Desai, S., Di Valentino, E., Diehl, H. T., Dietrich, J. P., Dodelson, S., Doel, P., Doux, C., Drlica-Wagner, A., Eckert, K., Eifler, T. F., Elsner, F., Elvin-Poole, J., Everett, S., Evrard, A. E., Fang, X., Farahi, A., Fernandez, E., Ferrero, I., Ferté, A., Fosalba, P., Friedrich, O., Frieman, J., García-Bellido, J., Gatti, M., Gaztanaga, E., Gerdes, D. W., Giannantonio, T., Giannini, G., Gruen, D., Gruendl, R. A., Gschwend, J., Gutierrez, G., Harrison, I., Hartley, W. G., Herner, K., Hinton, S. R., Hollowood, D. L., Honscheid, K., Hoyle, B., Huff, E. M., Huterer, D., Jain, B., James, D. J., Jarvis, M., Jeffrey, N., Jeltema, T., Kovacs, A., Krause, E., Kron, R., Kuehn, K., Kuropatkin, N., Lahav, O., Leget, P. F., Lemos, P., Liddle, A. R., Lidman, C., Lima, M., Lin, H., MacCrann, N., Maia, M. A. G., Marshall, J. L., Martini, P., McCullough, J., Melchior, P., Mena-Fernández, J., Menanteau, F., Miquel, R., Mohr, J. J., Morgan, R., Muir, J., Myles, J., Nadathur, S., Navarro-Alsina, A., Nichol, R. C., Ogando, R. L. C., Omori, Y., Palmese, A., Pandey, S., Park, Y., Paz-Chinchón, F., Petravick, D., Pieres, A., Plazas Malagón, A. A., Porredon, A., Prat, J., Raveri, M., Rodriguez-Monroy, M., Rollins, R. P., Romer, A. K., Roodman, A., Rosenfeld, R., Ross, A. J., Rykoff, E. S., Samuroff, S., Sánchez, C., Sanchez, E., Sanchez, J., Sanchez Cid, D., Scarpine, V., Schubnell, M., Scolnic, D., Secco, L. F., Serrano, S., Sevilla-Noarbe, I., Sheldon, E., Shin, T., Smith, M., Soares-Santos, M., Suchyta, E., Swanson,

M. E. C., Tabbutt, M., Tarle, G., Thomas, D., To, C., Troja, A., Troxel, M. A., Tucker, D. L., Tutusaus, I., Varga, T. N., Walker, A. R., Weaverdyck, N., Weller, J., Yanny, B., Yin, B., Zhang, Y., and Zuntz, J. Dark Energy Survey Year 3 Results: Cosmological Constraints from Galaxy Clustering and Weak Lensing. *arXiv e-prints*, art. arXiv:2105.13549, May 2021.

Drlica-Wagner, A., Sevilla-Noarbe, I., Rykoff, E. S., Gruendl, R. A., Yanny, B., Tucker, D. L., Hoyle, B., Carnero Rosell, A., Bernstein, G. M., Bechtol, K., Becker, M. R., Benoit-Lévy, A., Bertin, E., Carrasco Kind, M., Davis, C., de Vicente, J., Diehl, H. T., Gruen, D., Hartley, W. G., Leistedt, B., Li, T. S., Marshall, J. L., Neilsen, E., Rau, M. M., Sheldon, E., Smith, J., Troxel, M. A., Wyatt, S., Zhang, Y., Abbott, T. M. C., Abdalla, F. B., Allam, S., Banerji, M., Brooks, D., Buckley-Geer, E., Burke, D. L., Capozzi, D., Carretero, J., Cunha, C. E., D'Andrea, C. B., da Costa, L. N., DePoy, D. L., Desai, S., Dietrich, J. P., Doel, P., Evrard, A. E., Fausti Neto, A., Flaugher, B., Fosalba, P., Frieman, J., García-Bellido, J., Gerdes, D. W., Giannantonio, T., Gschwend, J., Gutierrez, G., Honscheid, K., James, D. J., Jeltema, T., Kuehn, K., Kuhlmann, S., Kuropatkin, N., Lahav, O., Lima, M., Lin, H., Maia, M. A. G., Martini, P., McMahon, R. G., Melchior, P., Menanteau, F., Miquel, R., Nichol, R. C., Ogando, R. L. C., Plazas, A. A., Romer, A. K., Roodman, A., Sanchez, E., Scarpine, V., Schindler, R., Schubnell, M., Smith, M., Smith, R. C., Soares-Santos, M., Sobreira, F., Suchyta, E., Tarle, G., Vikram, V., Walker, A. R., Wechsler, R. H., Zuntz, J., and DES Collaboration. Dark Energy Survey Year 1 Results: The Photometric Data Set for Cosmology. ApJS, 235(2):33, April 2018. doi: 10.3847/1538-4365/aab4f5.

Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J. emcee: The MCMC Hammer. PASP, 125(925): 306, March 2013. doi: 10.1086/670067.

Goodman, J. and Weare, J. Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, 5(1):65–80, January 2010. doi: 10.2140/camcos.2010.5.65.

Jia, H. and Seljak, U. BayesFast: A Fast and Scalable Method for Cosmological Bayesian Inference. *in prep.*, 2022.

Joudaki, S., Blake, C., Johnson, A., Amon, A., Asgari, M., Choi, A., Erben, T., Glazebrook, K., Harnois-Déraps, J., Heymans, C., Hildebrand t, H., Hoekstra, H., Klaes, D., Kuijken, K., Lidman, C., Mead, A., Miller, L., Parkinson, D., Poole, G. B., Schneider, P., Viola, M., and Wolf, C. KiDS-450 + 2dFLenS: Cosmological parameter constraints from weak gravitational lensing tomography and overlapping redshift-space galaxy clustering. MNRAS, 474(4):4894–4924, March 2018. doi: 10.1093/mnras/stx2820.

Karamanis, M., Beutler, F., and Peacock, J. A. zeus: A python implementation of ensemble slice sampling for efficient bayesian parameter inference. *arXiv preprint arXiv:2105.03468*, 2021.

Krause, E. and Eifler, T. cosmolike - cosmological likelihood analyses for photometric galaxy surveys. MNRAS, 470(2):2100–2112, September 2017. doi: 10.1093/mnras/stx1261.

Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.

Manrique-Yus, A. and Sellentin, E. Euclid-era cosmology for everyone: neural net assisted MCMC sampling for the joint 3×2 likelihood. MNRAS, 491(2):2655–2663, January 2020. doi: 10.1093/mnras/stz3059.

Pellejero-Ibañez, M., Angulo, R. E., Aricó, G., Zennaro, M., Contreras, S., and Stücker, J. Cosmological parameter estimation via iterative emulation of likelihoods. MNRAS, 499(4):5257–5268, December 2020. doi: 10.1093/mnras/staa3075.

Spurio Mancini, A., Piras, D., Alsing, J., Joachimi, B., and Hobson, M. P. COSMOPOWER: Emulating cosmological power spectra for accelerated Bayesian inference from next-generation surveys. MNRAS, January 2022. doi: 10.1093/mnras/stac064.

The LSST Dark Energy Science Collaboration, Mandelbaum, R., Eifler, T., Hložek, R., Collett, T., Gawiser, E., Scolnic, D., Alonso, D., Awan, H., Biswas, R., Blazek, J., Burchat, P., Chisari, N. E., Dell'Antonio, I., Digel, S., Frieman, J., Goldstein, D. A., Hook, I., Ivezić, Ž., Kahn, S. M., Kamath, S., Kirkby, D., Kitching, T., Krause, E., Leget, P.-F., Marshall, P. J., Meyers, J., Miyatake, H., Newman, J. A., Nichol, R., Rykoff, E., Sanchez, F. J., Slosar, A., Sullivan, M., and Troxel, M. A. The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document. *arXiv e-prints*, art. arXiv:1809.01669, September 2018.

To, C., Krause, E., Rozo, E., Wu, H., Gruen, D., Wechsler, R. H., Eifler, T. F., Rykoff, E. S., Costanzi, M., Becker, M. R., Bernstein, G. M., Blazek, J., Bocquet, S., Bridle, S. L., Cawthon, R., Choi, A., Crocce, M., Davis, C., DeRose, J., Drlica-Wagner, A., Elvin-Poole, J., Fang, X., Farahi, A., Friedrich, O., Gatti, M., Gaztanaga, E., Giannantonio, T., Hartley, W. G., Hoyle, B., Jarvis, M., MacCrann, N., McClintock, T., Miranda, V., Pereira, M. E. S., Park, Y., Porredon, A., Prat, J., Rau, M. M.,

Ross, A. J., Samuroff, S., Sánchez, C., Sevilla-Noarbe, I., Sheldon, E., Troxel, M. A., Varga, T. N., Vielzeuf, P., Zhang, Y., Zuntz, J., Abbott, T. M. C., Aguena, M., Amon, A., Annis, J., Avila, S., Bertin, E., Bhargava, S., Brooks, D., Burke, D. L., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Chang, C., Conselice, C., da Costa, L. N., Davis, T. M., Desai, S., Diehl, H. T., Dietrich, J. P., Everett, S., Evrard, A. E., Ferrero, I., Flaugher, B., Fosalba, P., Frieman, J., García-Bellido, J., Gruendl, R. A., Gutierrez, G., Hinton, S. R., Hollowood, D. L., Honscheid, K., Huterer, D., James, D. J., Jeltema, T., Kron, R., Kuehn, K., Kuropatkin, N., Lima, M., Maia, M. A. G., Marshall, J. L., Menanteau, F., Miquel, R., Morgan, R., Muir, J., Myles, J., Palmese, A., Paz-Chinchón, F., Plazas, A. A., Romer, A. K., Roodman, A., Sanchez, E., Santiago, B., Scarpine, V., Serrano, S., Smith, M., Suchyta, E., Swanson, M. E. C., Tarle, G., Thomas, D., Tucker, D. L., Weller, J., Wester, W., Wilkinson, R. D., and DES Collaboration. Dark Energy Survey Year 1 Results: Cosmological Constraints from Cluster Abundances, Weak Lensing, and Galaxy Correlations. Phys. Rev. Lett., 126(14):141301, April 2021a. doi: 10.1103/PhysRevLett.126.141301.

To, C.-H., Krause, E., Rozo, E., Wu, H.-Y., Gruen, D., DeRose, J., Rykoff, E., Wechsler, R. H., Becker, M., Costanzi, M., Eifler, T., da Silva Pereira, M. E., Kokron, N., and DES Collaboration. Combination of cluster number counts and two-point correlations: validation on mock Dark Energy Survey. MNRAS, 502(3):4093–4111, April 2021b. doi: 10.1093/mnras/stab239.

van Uitert, E., Joachimi, B., Joudaki, S., Amon, A., Heymans, C., Köhlinger, F., Asgari, M., Blake, C., Choi, A., Erben, T., Farrow, D. J., Harnois-Déraps, J., Hildebrandt, H., Hoekstra, H., Kitching, T. D., Klaes, D., Kuijken, K., Merten, J., Miller, L., Nakajima, R., Schneider, P., Valentijn, E., and Viola, M. KiDS+GAMA: cosmology constraints from a joint analysis of cosmic shear, galaxy-galaxy lensing, and angular clustering. MNRAS, 476(4): 4662–4689, June 2018. doi: 10.1093/mnras/sty551.