
Optimizing Likelihood-free Inference using Self-supervised Neural Symmetry Embeddings

Deep Chatterjee[†]

Massachusetts Institute of Technology
Cambridge, MA, USA

Philip C. Harris^{*}

Institute for Artificial Intelligence and Fundamental Interactions (IAIFI)
Massachusetts Institute of Technology
Cambridge, MA, USA

Maanas Goel

University of Pennsylvania
Philadelphia, PA, USA

Malina Desai

Massachusetts Institute of Technology
Cambridge, MA, USA

Michael W. Coughlin

University of Minnesota
Minneapolis, MN, USA

Erik Katsavounidis

Massachusetts Institute of Technology
Cambridge, MA, USA

Abstract

Likelihood-free inference is quickly emerging as a powerful tool to perform fast/effective parameter estimation. We demonstrate a technique of optimizing likelihood-free inference to make it even faster by marginalizing symmetries in a physical problem. In this approach, physical symmetries, for example, time-translation are learned using joint-embedding via self-supervised learning with symmetry data augmentations. Subsequently, parameter inference is performed using a normalizing flow where the embedding network is used to summarize the data before conditioning the parameters. We present this approach on two simple physical problems and we show faster convergence in a smaller number of parameters compared to a normalizing flow that does not use a pre-trained symmetry-informed representation.

1 Introduction

Parameter estimation is fundamental to experimental science. Traditional approaches involve stochastic sampling, like markov-chain monte carlo (MCMC) or nested sampling, to obtain posterior probability densities and confidence intervals on parameters. However, with the increase in number and flavors of analyses, complexity of experiments, and the demand for near real-time inference, stochastic sampling is turning out to be a challenge both in terms of time and required compute. Likelihood-free, or simulation-based inference (LFI/SBI) has emerged as a potential breakthrough in this regard (see Refs. [1, 2] for a review). Simulations to generate data, \mathbf{d} , corresponding to physical parameters, Θ , exists for many experiments. The idea of LFI is to learn $p(\Theta, \mathbf{d})$, $p(\Theta|\mathbf{d})$, or $p(\mathbf{d}|\Theta)$

[†]deep1018@mit.edu; ^{*}pcharris@mit.edu




Figure 1: **Left:** Example signal and data of a damped harmonic oscillator (SHO) with representative values of $\Theta = \{\omega_0, \beta\}$. The data is shown by the solid dots to represent sampling a time-series. **Right:** An example of a sine-gaussian (SG) pulse with parameters, $\Theta = \{f_0, \tau\}$. In either case, different times of arrival, shown in faded curves, do not affect the inference of parameters.

from the pairs $\{\Theta_i, \mathbf{d}_i\}$ as desired (see Sec. 2.1 of [3] for a nice description). In this work, we focus on the construction of the posterior, $p(\Theta|\mathbf{d})$.

A critical aspect of the application of LFI arises from the difficulty in performing high dimensional fits. Often times, posterior extraction requires a large amount of training samples, and large neural network to ensure a fully accurate description of the model. As an example, current state of the art networks for gravitational-wave (GW) parameter estimation require over 100 million parameters to ensure an accurate description of the GW parameters [4]. One potential way to mitigate the size of LFI models is to project out parameters that are known to be invariant to the desired intrinsic parameters. For example, in the case of pulse-like time-series, the time of arrival may not be as important as the physical system/process that produces the pulse. However, a conventional application of LFI, unaware of this symmetry, considers data corresponding to time translated signals as *unique* and requires that the arrival time be determined, or marginalized over, with the other desired physical parameters. This necessitates a larger model and expensive compute to train a model that learns to marginalize over time (or some other variables, λ), despite the fact that the desired parameter inference is time-(or λ -) invariant. If the interest is limited to the physical, intrinsic parameters, this is superfluous.

In this paper, we utilize self-supervision to project away the invariant parameters so as to reduce the LFI to the extraction of only the intrinsic parameters. Here, we marginalize over this symmetry by learning a *similarity* embedding through self-supervised learning (SSL). In the instance of time translation invariance we rely on SSL to build a reduced representation where instances of data with different arrival times are represented as similar values in an embedded space. We train a representation network by minimizing the VICReg similarity loss [5]. A similar technique was used recently to summarize cosmological data [6]. SSL has also been applied to learn symmetries in partial differential equations (PDEs) [7]. There are other techniques, not involving SSL to summarize data for LFI on timeseries e.g. YuleNet [8]. Once we have constructed the SSL space, the subsequent LFI is conditioned on the embedded space. We demonstrate this scheme considering two simple physical systems – a damped harmonic oscillator (subsequently called SHO) that is started at different times, and a sine-gaussian pulse (subsequently called SG) that has an arbitrary time of arrival in a window. We show that we are able to obtain posteriors which are accurate and require a significantly smaller number of training parameters compared to the case when the data is not represented, or similarity pre-training is not performed. We point out that while *signals* with identical physical parameters, translated in time is an exact symmetry, the corresponding *data* is not due to the presence of noise. This motivates “similar” representation as opposed to “same.”

2 Marginalizing Time-translation in Physical problems

Consider a SHO described by physical parameters, $\Theta = \{\omega_0, \beta\}$, the natural frequency and damping coefficient respectively. We would like to infer on the parameters given a time-series, \mathbf{d} , with SHO

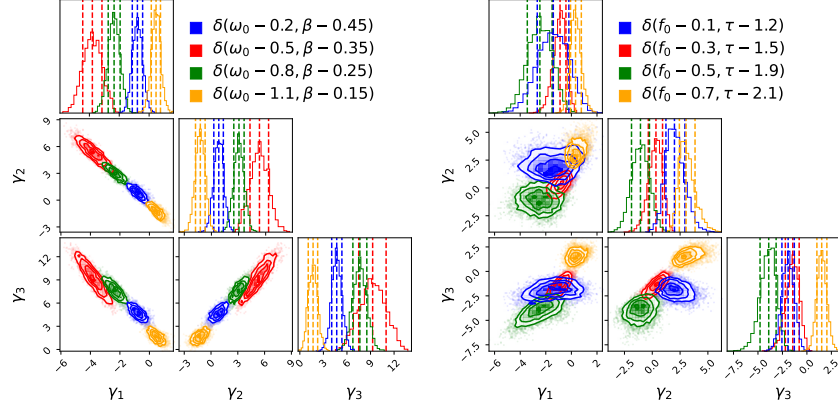


Figure 2: SHO (left) and SG (right) data representations. The signals have intrinsic parameters denoted by the delta function, but have augmented time shifts sampled uniformly from a window. We observe that the SSL construction leads to separation based on intrinsic parameters where the variations in the time of arrival is reduced to a substantially smaller region than the intrinsic parameters, approaching the time-project space.

amplitudes as shown in the left panel of Fig. 1; white noise is additionally added. If the oscillator is started at a different time, represented by the shifted solid lines in the figure, the inference on ω_0 and β remains the same. The same applies in the context of inferring properties of a SG pulse, represented in the right panel of Fig. 1, where the parameters of interest are the central frequency, f_0 , and width, τ . In the standard scheme of LFI, the parameters are sampled from prior distribution, $\Theta \sim p(\Theta)$, following which the data, \mathbf{d} , is generated, and the model is provided $\{\Theta_i, \mathbf{d}_i\}$ pairs. This, however, leaves learning nuisance parameters, like time shifts, and marginalizing over them, up to the model, which could be expensive in terms of trainable parameters and compute.

2.1 Representing time-shifted data

We project the time shifted data instances into a 3D embedded space via a neural network, $\Gamma \equiv h \circ f$, such that the data at different times, with the same parameters Θ , is mapped to similar regions in this embedded space. We follow the scheme of self-supervised learning as follows:

- We consider two batches – in the first case the data, \mathbf{d} , has a fixed reference time of arrival, in the second, the data, \mathbf{d}' , is augmented by shifting start time by a suitable time-shift prior.
- In our implementation, f is a 1D convolutional resnet followed by a fully connected contraction layer to represent each batch in 3 dimensions i.e. $\gamma \in \mathbb{R}^3$. The representation is then passed through a fully-connected expander network, h , to output $x \in \mathbb{R}^{12}$ i.e.,

$$\gamma = f(\mathbf{d}); \gamma' = f(\mathbf{d}'); x = h(\gamma); x' = h(\gamma'). \quad (1)$$

Ideally, the dimensionality of the representation is determined based on a systematic hyper-parameter tuning. However, in this case, our choices dimensions was based on performance from trials on a handful of cases.

- We use the VICReg loss between the expanded outputs to minimize regularized variance and covariance of the embedded batches individually, and the invariance between the batches using mean-squared error. The net loss is a weighted sum of the three terms,

$$\mathcal{L}_{\text{VICReg}}(x, x') = \lambda_1 \text{MSE}(x, x') + \lambda_2 [\text{Var}(x + \epsilon) + \text{Var}(x' + \epsilon)] + \lambda_3 [\text{Cov}(x) + \text{Cov}(x')], \quad (2)$$

where $\lambda_{1,2,3}$ are the relative weighting of the three terms that can be adjusted during training. For this work we found that increasing the reconstruction and variance weights compared to the covariance term in the initial epochs followed by setting the weights equal helped with the convergence.

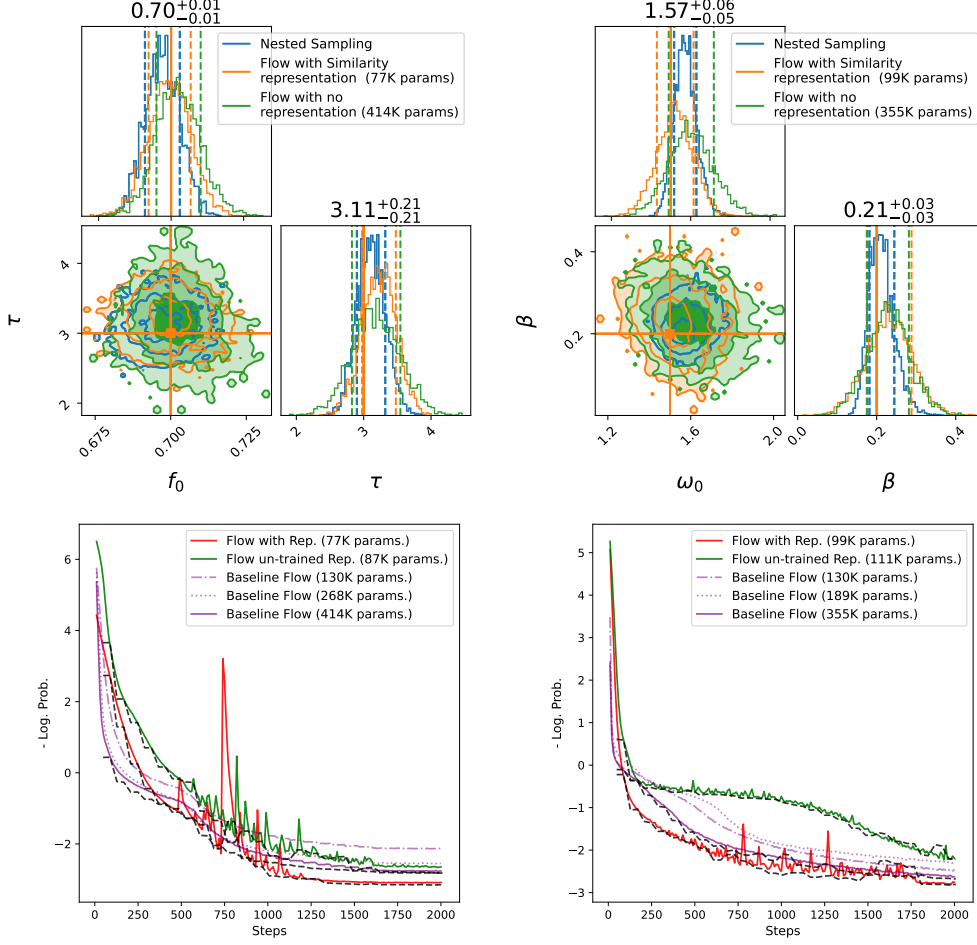


Figure 3: **Top:** We show consistency between posteriors from a NF with a pre-trained Γ , a baseline NF without any data summary, and results from nested sampling. The \pm uncertainties in the individual panels are quoted based on the nested sampling result. **Bottom:** Loss curves (black dashed lines for validation) for the different cases. The NF with a pre-trained Γ converges in a smaller number of trainable parameters, and sometimes more rapidly, compared to other cases.

- Once trained, data from same physical parameters are represented as clusters embedded in this space irrespective of the time of arrival. This is shown in Fig. 2, where different physical parameters are represented as roughly separated clusters irrespective of their time of arrival. We freeze part of the network f at this point.^a

3 LFI with Normalizing Flows

Next, we perform the conventional LFI by training a normalizing flow (NF) [9]. We use a base 2D standard normal distribution with affine autoregressive transforms i.e. a masked autoregressive flow (MAF) [10].^b We condition the normalizing flow with the representation network, f . Hence, the pairs provided during training are $\{\Theta_i, f(\mathbf{d}_i)\} \equiv \{\Theta_i, \gamma_i\}$. The network is trained by maximizing the log-likelihood. In Fig. 3 we show an example of posterior samples from testing data for both SHO and SG models. We compare the same with posteriors samples obtained by training a baseline NF which does not have any such representation network, and also with results from nested sampling.^c We compare the widths of the posterior obtained using our trained model and stochastic sampling and

^aHere, we freeze the conv. layers of f , and left the fully connected layers free to train in the next step.

^bWe use the implementation from `nflows` library with minor changes.

^cWe use the `bilby` inference library with `dynesty` sampler.




Figure 4: **Left:** Results from 1000 tests data instances for which posteriors were drawn. It plots fraction of times the true parameters of the SHO model lie in credible interval. In each case 3000 posterior samples are drawn from the trained model. The gray bands denote the 1-, 2-, 3- σ confidence intervals in decreasing order of opacity. **Right:** Same plot as the left for the SG model and testing set.

show that they are consistent. We also compute the Cramer-Rao bound on the parameter uncertainty by expanding the likelihood to quadratic order in the argument of the exponential. We show that the results obtained from the trained flow are consistent with CRB for the SG and SHO models (see Appendix A.1). In Fig. 4 we show the result of 1000 test data instances for which posteriors were drawn from the trained model. The plot highlights the consistency between the confidence intervals and the fraction of times the true parameters lie inside the confidence intervals. The diagonal trend in the figure, along with the consistency of individual posteriors mentioned in Appendix A.1 show consistency across the entire testing dataset. The gray bands on both panels denote the 1-, 2-, 3- σ confidence intervals.

Comparing number of trainable parameters – the baseline NF model, without any data summary, has $\mathcal{O}(10^6)$ while the same for the NF with pre-trained Γ is significantly lower i.e. $\mathcal{O}(10^5)$. For a batch size of 10^3 , a forward pass involves $\sim 1.1 \times 10^8$ multiply-accumulate operations (MACs) for the NF with the representation, while $\sim 4.1 \times 10^8$ MACs for the baseline NF. This may translate to longer training times, for example, timing 100 epochs a NVIDIA GeForce GTX 1080 Ti GPU with SGD optimizer takes ~ 20 minutes for the NF with representation while ~ 80 minutes for the baseline NF. We found 100 epochs of the similarity pre-training, which takes ~ 10 minutes, was sufficient to reduce the individual loss terms in Eq. (2). We found that freezing all trainable parameters of representation network parameters at this stage led to slightly larger posterior widths compared to nested sampling. However, letting the network further train with the flow, at least leaving the fully connected layers of the representation network unfrozen, led to similar performance as nested sampling. Therefore, the prescription presented here generally implies smaller networks and more efficient memory and compute requirements.

4 Application in Physics and Astronomy

While in this work, we have considered projecting out the time of arrival as a specific example of a symmetry that does not impact the measurement of other parameters, it should be noted that the technique presented here could be used in general to marginalize over parameters that are not of interest. Many problems physics and astronomy exhibit a set of *intrinsic* vs. *extrinsic* parameters. For example, in GW data analysis, the intrinsic parameters include the masses and spins of merging binary black holes, while extrinsic parameters involve distance, inclination, sky-location etc. Likewise, a kilonova lightcurve is intrinsically characterized by the ejecta mass and lanthanide fraction etc., and extrinsically by redshift, extinction, time of arrival etc. If the interest is only limited to a subset of parameters, or to only intrinsic properties, a representation using self-supervised learning can be used for efficient LFI.

This work is supported by NSF HDR Institute Grant PHYS-2117997, “Accelerated AI Algorithms for Data-Driven Discovery.”

A Appendix

A.1 Analytic approximation of the posterior

The signal models are,

$$y_i \equiv y(t_i) = \begin{cases} \exp(-\beta\omega_0 t_i) \cos(\omega_0 t_i \sqrt{1-\beta^2}) & \text{for SHO} \\ \exp\left(-\frac{t_i^2}{\tau^2}\right) \sin(2\pi f_0 t_i), & \text{for SG.} \end{cases} \quad (3)$$

Assuming white noise drawn from $\mathcal{N}(0, \sigma)$, (with σ known) the likelihood is,

$$\mathcal{L} \propto \exp\left[-\frac{1}{2} \sum_i \left(\frac{y_i - \hat{y}_i}{\sigma}\right)^2\right], \quad (4)$$

where \hat{y} is the truth i.e. signal corresponding to true parameters $\{\hat{\omega}_0, \hat{\beta}\}$ ($\{\hat{f}_0, \hat{\tau}\}$) for SHO (SG). Expanding the argument of the exponential to quadratic order in $\{\omega_0, \beta\}$ (or $\{f_0, \tau\}$) around the true value, the likelihood is approximated by gaussian with widths,

$$\Delta\omega_0 \approx \frac{\sigma}{\sqrt{\sum_i^N t_i^2 e^{-2\hat{\beta}\hat{\omega}_0 t_i} \left[\sqrt{1-\hat{\beta}^2} \sin(\hat{\omega}_0 t_i \sqrt{1-\hat{\beta}^2}) + \hat{\beta} \cos(\hat{\omega}_0 t_i \sqrt{1-\hat{\beta}^2}) \right]^2}}, \quad (5a)$$

$$\Delta\beta \approx \frac{\sigma}{\sqrt{\sum_i^N \hat{\omega}_0^2 t_i^2 e^{-2\hat{\beta}\hat{\omega}_0 t_i} \left[\frac{\hat{\beta} \sin(\hat{\omega}_0 t_i \sqrt{1-\hat{\beta}^2})}{\sqrt{1-\hat{\beta}^2}} - \cos(\hat{\omega}_0 t_i \sqrt{1-\hat{\beta}^2}) \right]^2}}, \quad (5b)$$

$$\Delta\tau \approx \frac{\sigma \hat{\tau}^3}{\sqrt{\sum_i^N 4 t_i^4 \exp(-2t_i^2/\hat{\tau}^2) \sin^2(2\pi \hat{f}_0 t_i)}}, \quad (5c)$$

$$\Delta f \approx \frac{\sigma}{\sqrt{\sum_i^N 4\pi^2 t_i^2 \exp(-2t_i^2/\hat{\tau}^2) \cos^2(2\pi \hat{f}_0 t_i)}}. \quad (5d)$$

By assuming flat priors, the posterior is proportional to the likelihood, having widths approximately given by Eqs. (5). For the examples in Fig. 3, the value of $\sigma = 0.4$ in either case. The true SG parameters are $\{\hat{f}_0 = 0.7, \hat{\tau} = 0.3\}$; for SHO they are $\{\hat{\omega}_0 = 1.5, \hat{\beta} = 0.2\}$. From Eqs. (5), we get $\Delta f_0 \approx 0.006$ $\Delta\tau \approx 0.2$ for SG, and $\Delta\omega_0 \approx 0.04$ $\Delta\beta \approx 0.03$ for SHO models, satisfying Cramer-Rao bound for the one-sigma widths shown in Fig. 3.

A.2 Code Availability

The codes used in the analysis, including trained weights to reproduce figures in the draft, are provided publicly and tagged. Online rendering of notebooks are available here: <https://github.com/ML4GW/summer-projects-2023/tree/neurips-2023/symmetry-informed-flows>.

References

- [1] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, may 2020.
- [2] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference, 2021.
- [3] Justin Alsing, Tom Charnock, Stephen Feeney, and Benjamin Wandelt. Fast likelihood-free cosmology with neural density estimators and active learning. *Monthly Notices of the Royal Astronomical Society*, jul 2019.
- [4] Maximilian Dax, Stephen R. Green, Jonathan Gair, Jakob H. Macke, Alessandra Buonanno, and Bernhard Schölkopf. Real-time gravitational wave science with neural posterior estimation. *Physical Review Letters*, 127(24), dec 2021.

- [5] Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning, 2022.
- [6] Aizhan Akhmetzhanova, Siddharth Mishra-Sharma, and Cora Dvorkin. Data compression and inference in cosmology with self-supervised machine learning, 2023.
- [7] Grégoire Mialon, Quentin Garrido, Hannah Lawrence, Danyal Rehman, Yann LeCun, and Bobak T. Kiani. Self-supervised learning with lie symmetries for partial differential equations, 2023.
- [8] Pedro L. C. Rodrigues and Alexandre Gramfort. Learning summary features of time series for likelihood free inference, 2020.
- [9] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows, 2016.
- [10] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation, 2018.