

Error-guided likelihood-free MCMC

Volodimir Begy^{1,2} and Erich Schikuta¹

¹ University of Vienna, Faculty of Computer Science, Research Group Workflow Systems and Technology, Währinger Straße 29, 1090, Vienna, Austria

² CERN, 1211 Geneva 23, Switzerland
volodimir.begy@cern.ch

Abstract. This work presents a novel posterior inference method for models with intractable evidence and likelihood functions. Error-guided likelihood-free MCMC, or EG-LF-MCMC in short, has been developed for scientific applications, where a researcher is interested in obtaining approximate posterior densities over model parameters, while avoiding the need for expensive training of component estimators on full observational data or the tedious design of expressive summary statistics, as in related approaches. Our technique is based on two phases. In the first phase, we draw samples from the prior, simulate respective observations and record their errors ϵ in relation to the true observation. We train a classifier to distinguish between corresponding and non-corresponding (ϵ, θ) -tuples. In the second stage the said classifier is conditioned on the smallest recorded ϵ value from the training set and employed for the calculation of transition probabilities in a Markov Chain Monte Carlo sampling procedure. By conditioning the MCMC on specific ϵ values, our method may also be used in an amortized fashion to infer posterior densities for observations, which are located a given distance away from the observed data. We evaluate the proposed method on benchmark problems with semantically and structurally different data and compare its performance against the state of the art approximate Bayesian computation (ABC).

Keywords: Likelihood-free inference · Markov Chain Monte Carlo · artificial neural networks.

1 Motivation

In natural sciences the studied phenomena are typically formalized by mathematical models. These models may come in various flavors, ranging from systems of differential equations to complex stochastic computer simulations. With the advancement of computing technologies it has become increasingly common to specify mechanistic models using complex computer programs. This gives rise to the negative side effect that the likelihood functions and the model evidence are typically intractable, meaning, they have no closed-form expressions due to the model complexity. Researchers often wish to evaluate theories or refine the developed statistical models. An applicable tool for this is posterior inference, whereby

a practitioner seeks a conditional probability density function over model parameters given observed data. In modern scientific applications, posterior inference is a non-trivial task. As an example, consider a computer science researcher, who has experimentally gathered some metrics describing the performance of a data grid, which is a large-scale distributed data-intensive computing system. This production system is abstracted by a simulator with numerous, potentially latent parameters. The researcher would like to relate the observed data \mathbf{x} to concrete values of the simulator parameters $\boldsymbol{\theta}$, by obtaining a posterior $p(\boldsymbol{\theta} | \mathbf{x})$. Due to the intractability of the likelihood function $p(\mathbf{x} | \boldsymbol{\theta})$ and model evidence $p(\mathbf{x})$, this task cannot be performed analytically by the means of the Bayes’ theorem:

$$p(\boldsymbol{\theta} | \mathbf{x}) = \frac{p(\mathbf{x} | \boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{x})}. \quad (1)$$

This issue has motivated the field of likelihood-free or simulation-based inference [4], in which researchers develop new inference methods, which do not rely on the direct evaluation of likelihood functions and model evidence. Typically, such novel inference techniques are developed using tools from soft computing and statistics.

We propose a method for the construction of the error-conditioned posterior $p(\boldsymbol{\theta} | \epsilon)$ instead of the traditional posterior $p(\boldsymbol{\theta} | \mathbf{x})$. Error-guided likelihood-free MCMC can be used for 2 different scientific purposes by employing a single trained classifier in an amortized fashion. Firstly, it allows to obtain the posterior $p(\boldsymbol{\theta} | \epsilon_{min})$, where ϵ_{min} is the smallest recorded error from the training set in relation to the true observation \mathbf{x}_o . The parameters $\boldsymbol{\theta}$ drawn from $p(\boldsymbol{\theta} | \epsilon_{min})$ aim at reproducing the observed data \mathbf{x}_o . Secondly, the practitioner may be interested in deriving posterior densities over model parameters, which allow to generate observations a given distance away from the observed data. This is achievable using our approach by conditioning the MCMC procedure on any ϵ value, which is covered by the range of the training set.

2 Related Work

Arguably the most popular method for likelihood-free posterior sampling is approximate Bayesian computation (ABC) [2, 16]. Despite being conceptually simple, it is relatively effective. In rejection ABC a certain $\boldsymbol{\theta}$ is drawn from the prior $p(\boldsymbol{\theta})$ and a respective observation $\mathbf{x} \sim p(\mathbf{x} | \boldsymbol{\theta})$ gets simulated. Then, a selected distance function $d(\mathbf{x}, \mathbf{x}_o)$ calculates the distance between the observed and simulated data. The observations may also be compressed into lower dimensions by summary statistics. The construction of expressive summary statistics is a non-trivial task [1, 7, 18]. In case the distance is below a selected threshold ϵ , the generating parameter $\boldsymbol{\theta}$ is accepted. This procedure is repeated, until the collected sample is large enough to approximate the posterior density. Typically, ABC is relatively inefficient in high-dimensional $\boldsymbol{\theta}$ -spaces, since most simulated

observations get rejected. We note that in the ABC literature ϵ stands for the acceptance threshold, while in our formalism ϵ denotes concrete distances between the candidate and true observations.

In the recent years, plenty of more advanced likelihood-free inference methods have emerged. Our work is inspired by Likelihood-free MCMC with Amortized Approximate Ratio Estimators, or AALR-MCMC [12]. AALR-MCMC trains an artificial neural network classifier to learn a mapping between the θ and \mathbf{x} spaces by distinguish corresponding (\mathbf{x}, θ) -tuples from non-corresponding ones. In case of high dimensionality of the \mathbf{x} space, the observed data can be compressed by a summary statistic. It is demonstrated that an optimally trained classifier outputs $\frac{p(\mathbf{x}, \theta)}{p(\mathbf{x}, \theta) + p(\mathbf{x})p(\theta)}$, allowing us to calculate the likelihood-to-evidence ratio for a given input pair (\mathbf{x}, θ) . In order to obtain a posterior, an MCMC sampling scheme is used. While constructing a Markov Chain, the probability of the transition from the current state θ_t to a proposed state θ' is calculated using the estimated likelihood ratio $\hat{r}(\mathbf{x} | \theta_t, \theta')$, which is extracted from the converged classifier. A major strength of this technique is its ability to perform amortized inference for different observations \mathbf{x} using a single trained classifier.

A family of approaches constructs the approximate prior using sequential procedures [6]. In SNPE-A [19] the authors refine a proposal prior $\tilde{p}(\mathbf{x})$ by iteratively re-training a parameterized conditional density estimator $q_\phi(\theta | \mathbf{x})$, which is realized using single-component Mixture Density Networks [3]. Once $\tilde{p}(\mathbf{x})$ converges, the final approximate posterior is derived after training an MDN with multiple components. SNPE-B [15] extends upon this work by training the MDN using an importance-weighted loss function. This allows $q_\phi(\theta | \mathbf{x})$ to output posterior densities without additional analytical transformations. However, the authors in [9] argue, that such importance-weighting gives rise to a high variance of the density estimator parameters. Furthermore, the loss function includes a term, which calculates the Kullback-Leibler divergence between the distributions over MDN parameters from the previous and the current optimization rounds. Thus, the optimization procedure primarily focuses on refining highly uncertain parameters. The authors improve the simulation efficiency by discarding proposed θ , which are deemed as nonsensical by a binary classifier. SNPE-C/APT [9] optimizes a conditional estimator, which directly targets the posterior, by minimizing the estimated negative log density of a proposal posterior over samples drawn from a proposal prior. SNL [20] opts for learning an estimator of the likelihood function $q_\phi(\mathbf{x} | \theta)$ and obtains the approximate posterior at round r by $\hat{p}_r \propto q_\phi(\mathbf{x}_o | \theta)p(\theta)$. Hermans et al. [12] demonstrate that AALR-MCMC can also be executed as a sequential procedure by using the approximated posterior from the round r as the prior in the round $r + 1$. Such implementation makes sense when the capacity of the classifier is not sufficient for it to converge in a single round.

Durkan et al. [5] have shown that AALR-MCMC and SNPE-C are instances of contrastive learning and presented a generalized algorithm for Sequential Contrastive Likelihood-free Inference.

3 The Method

EG-LF-MCMC consists of two phases. Firstly, we draw θ samples from the prior $p(\theta)$, generate observations $\mathbf{x} \sim p(\mathbf{x} | \theta)$ from the implicit model and record their errors ϵ in relation to the true observation \mathbf{x}_o by applying a distance function $d(\mathbf{x}, \mathbf{x}_o)$. We train a classifier $\mathbf{s}(\epsilon, \theta)$ to distinguish between corresponding and non-corresponding (ϵ, θ) -tuples. In the second stage the output of the classifier is transformed into estimated ratios $\hat{r}(\epsilon | \theta', \theta_t)$ in order to evaluate the transition probabilities for subsequent state samples θ' and θ_t in an MCMC sampling scheme. The method may be run iteratively, by setting the posterior $p(\theta | \epsilon)$ from the round r as the prior at the round $r + 1$.

Algorithm 1 Training of $\mathbf{s}(\epsilon, \theta)$

Inputs: Generative model $p(\mathbf{x} | \theta)$
 Prior $p(\theta)$
 True observation \mathbf{x}_o
 Distance function $d(\mathbf{x}, \mathbf{x}_o)$

Outputs: Classifier $\mathbf{s}_\phi(\epsilon, \theta)$

Hyperparameters: Batch-size M

- 1: **while** $\mathbf{s}_\phi(\epsilon, \theta)$ **not converged do**
- 2: **Sample** $\theta \leftarrow \{\theta_m \sim p(\theta)\}_{m=1}^M$
- 3: **Sample** $\theta' \leftarrow \{\theta'_m \sim p(\theta)\}_{m=1}^M$
- 4: **Simulate** $\mathbf{x} \leftarrow \{\mathbf{x}_m \sim p(\mathbf{x} | \theta_m)\}_{m=1}^M$
- 5: **Simulate** $\mathbf{x}' \leftarrow \{\mathbf{x}'_m \sim p(\mathbf{x} | \theta'_m)\}_{m=1}^M$
- 6: $\epsilon \leftarrow \{d(\mathbf{x}_m, \mathbf{x}_o)\}_{m=1}^M$
- 7: $\epsilon' \leftarrow \{d(\mathbf{x}'_m, \mathbf{x}_o)\}_{m=1}^M$
- 8: $\mathcal{L}_a \leftarrow \text{BCE}(\mathbf{s}_\phi(\epsilon, \theta), \mathbb{1}) + \text{BCE}(\mathbf{s}_\phi(\epsilon', \theta), \mathbb{0})$
- 9: $\mathcal{L}_b \leftarrow \text{BCE}(\mathbf{s}_\phi(\epsilon', \theta'), \mathbb{1}) + \text{BCE}(\mathbf{s}_\phi(\epsilon, \theta'), \mathbb{0})$
- 10: $\mathcal{L} \leftarrow \mathcal{L}_a + \mathcal{L}_b$
- 11: $\phi \leftarrow \text{optimizer}(\phi, \nabla_\phi \mathcal{L})$
- 12: **end while**
- 13: **return** $\mathbf{s}_\phi(\epsilon, \theta)$

Our method builds upon AALR-MCMC and modifies it in the following ways: (1) the classifier is trained on errors ϵ in relation to \mathbf{x}_o instead of original or compressed observational data \mathbf{x} ; (2) the MCMC procedure is conditioned on ϵ instead of \mathbf{x} . When inferring a posterior $p(\theta | \epsilon_{min})$, which aims at reproducing \mathbf{x}_o , we pass the smallest simulated ϵ value to the MCMC. Alternatively, if one wishes to infer a posterior for observations, which are a given distance away from \mathbf{x}_o , it is possible to specify any ϵ value, which is covered by the range of the training set.

The function $d(\mathbf{x}, \mathbf{x}_o) = \epsilon$ can be flexibly specified taking into account the semantics of the studied problem. We have found the L_1 distance to work well on a wide range of problems:

$$d(\mathbf{x}, \mathbf{x}_o) = \sum_{i=1}^N |\mathbf{x} - \mathbf{x}_o|_i \quad (2)$$

Training the classifier on ϵ brings several advantages. In problems with high-dimensional observations, the size of the training set becomes drastically reduced, resulting into memory and computational efficiency. It becomes easier to specify a sufficient classifier architecture. These advantages can also be obtained by compressing the data using a summary statistic. However, we argue that it is easier to specify a semantically viable distance function than an expressive summary statistic for non-trivial high-dimensional data. Lastly, the novel amortization over ϵ may come handy in specific scientific use-cases.

The training procedure for the classifier $\mathbf{s}(\epsilon, \boldsymbol{\theta})$ is adapted from AALR-MCMC and is presented in the Algorithm 1, where BCE stands for the binary cross entropy loss function. In [12] the authors have shown that the classifier trained by AALR-MCMC to optimality outputs $\frac{p(\mathbf{x}, \boldsymbol{\theta})}{p(\mathbf{x}, \boldsymbol{\theta}) + p(\mathbf{x})p(\boldsymbol{\theta})}$ for a given $(\mathbf{x}, \boldsymbol{\theta})$ -pair and that this output can be transformed into the likelihood-to-evidence ratio:

$$\frac{p(\mathbf{x} | \boldsymbol{\theta})}{p(\mathbf{x})} = r(\mathbf{x} | \boldsymbol{\theta}) \quad (3)$$

Since our procedure replaces \mathbf{x} by ϵ , a classifier optimally trained by EG-LF-MCMC outputs $\frac{p(\epsilon, \boldsymbol{\theta})}{p(\epsilon, \boldsymbol{\theta}) + p(\epsilon)p(\boldsymbol{\theta})}$, which is transformed into the ratio:

$$\frac{p(\epsilon | \boldsymbol{\theta})}{p(\epsilon)} = r(\epsilon | \boldsymbol{\theta}) \quad (4)$$

In order to perform the MCMC sampling, we adapt the likelihood-free Metropolis-Hastings sampler [10, 17] from AALR-MCMC for the use of our classifier $\mathbf{s}(\epsilon, \boldsymbol{\theta})$, as summarized by the Algorithm 2. We note that in the common case of a uniform prior the terms $\log p(\boldsymbol{\theta}')$ and $\log p(\boldsymbol{\theta}_t)$ from the 5th line can be removed. The term $\frac{q(\boldsymbol{\theta}_t | \boldsymbol{\theta}')}{q(\boldsymbol{\theta}' | \boldsymbol{\theta}_t)}$ from the 6th line can be removed in case of a symmetric transition distribution. These modifications will speed up the computation.

Algorithm 2 Error-guided likelihood-free Metropolis-Hastings

Inputs: Initial parameter θ_0
Prior $p(\theta)$
Transition distribution $q(\theta)$
Classifier $\mathbf{s}(\epsilon, \theta)$
Error ϵ

Outputs: Markov chain $\theta_{0:n}$

Hyperparameters: Markov chain size n

- 1: $t \leftarrow 0$
- 2: $\theta_t \leftarrow \theta_0$
- 3: **for** $t < n$ **do**
- 4: $\theta' \sim q(\theta | \theta_t)$
- 5: $\lambda \leftarrow (\log \hat{r}(\epsilon | \theta') + \log p(\theta')) - (\log \hat{r}(\epsilon | \theta_t) + \log p(\theta_t))$
- 6: $\rho \leftarrow \min(\exp(\lambda) \frac{q(\theta_t | \theta')}{q(\theta' | \theta_t)}, 1)$
- 7: $\theta_{t+1} \leftarrow \begin{cases} \theta' & \text{with probability } \rho \\ \theta_t & \text{with probability } 1 - \rho \end{cases}$
- 8: $t \leftarrow t + 1$
- 9: **end for**
- 10: **return** $\theta_{0:n}$

4 Experiments

We compare EG-LF-MCMC to rejection ABC, since both approaches are similar in that they operate on calculating a distance between a candidate observation \mathbf{x} and the true observation \mathbf{x}_o . In particular, we evaluate the ability of $p(\theta | \epsilon_{min})$ from EG-LF-MCMC and $p(\theta | \mathbf{x})$ from ABC to recover the true generating parameter θ^* , which gives rise to the true observation \mathbf{x}_o .

		Circle	Linear
EG-LF-MCMC	n simulations	12,033,424	12,652,661
	n training epochs	75	36
ABC	n simulations	134,582,382	298,734,182
	ϵ	100	200

Table 1. Further implementation details

On each task, both methods employ the L_1 distance to calculate the discrepancy between the candidate and the true observations. The EG-LF-MCMC classifiers are realized using a multilayer perceptron with 4 hidden layers, each having 128 hidden units and SELU [14] non-linearities. The nets are trained using the Adam optimizer [13] with a learning rate of 0.0001 and a batch size of

64. We project the training data onto the $[0, 1]$ interval, as the experiments have indicated that this stabilizes the training. Furthermore, after such projection, the MCMC procedure can always be conditioned on $\epsilon_{min} = 0$, without the extra effort of scanning the training set for the actual value of the smallest recorded ϵ . We use the PyTorch framework [21] for the implementation and the training of the neural nets. Both methods draw 1 million samples in order to approximate the posterior. EG-LF-MCMC additionally performs 10,000 burn-in steps. Other variable implementation metrics for both methods are reported in the Table 1. Since both approaches share a single prior, we pick the value for the acceptance threshold in ABC based on the analysis of the distributions of ϵ from the training set of EG-LF-MCMC, which are depicted in the Fig. 1.

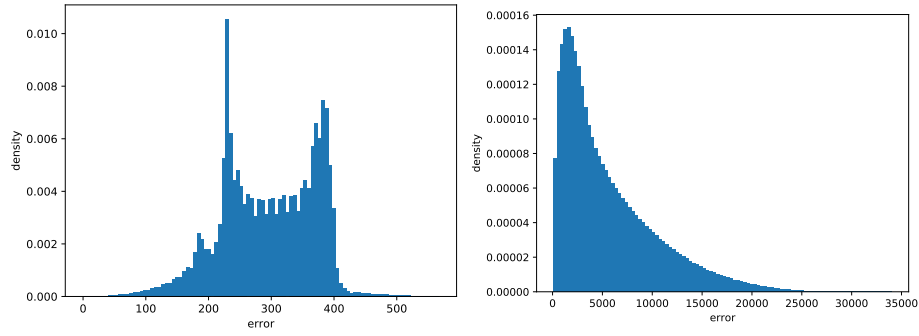


Fig. 1. Distribution of ϵ for the circle (left) and the linear (right) benchmark problems

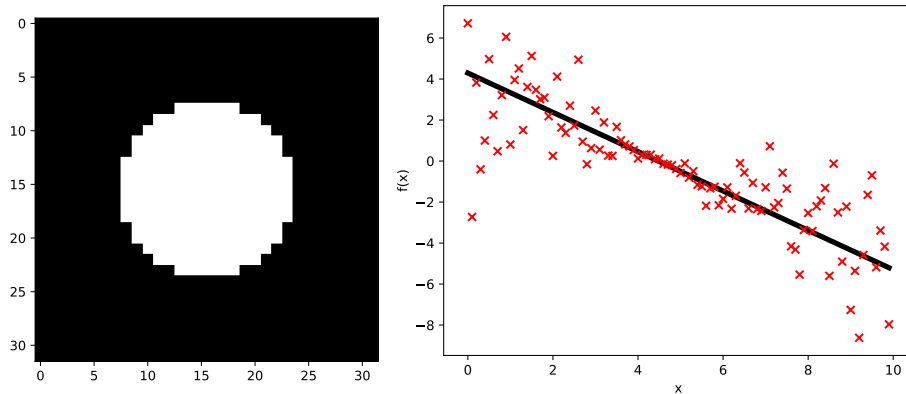


Fig. 2. Example \mathbf{x}_o for the circle (left) and the linear (right) benchmark problems

Circle image simulator The first benchmark problem is based on a deterministic simulator of 2-dimensional circle images, while both x and y coordinates are constrained by the interval $[-1, 1]$ [11]. The images are black and white, having a size of 32×32 pixels. The simulator is configured by 3 parameters $\theta = (x, y, r)$, where x and y are the coordinates of the circle’s center, and r is its radius. The true observation \mathbf{x}_o is generated from $\theta^* = (x = 0, y = 0, r = 0.5)$ and is depicted on the left side of the Fig. 2. We assume a uniform prior with the range of $[-1, 1]$ for x and y and the range of $[0, 1]$ for r . Certain simulated observations from the training set of EG-LF-MCMC are able to achieve $\epsilon = 0$ and exactly reproduce \mathbf{x}_o . The obtained posteriors are presented in the Fig. 3.

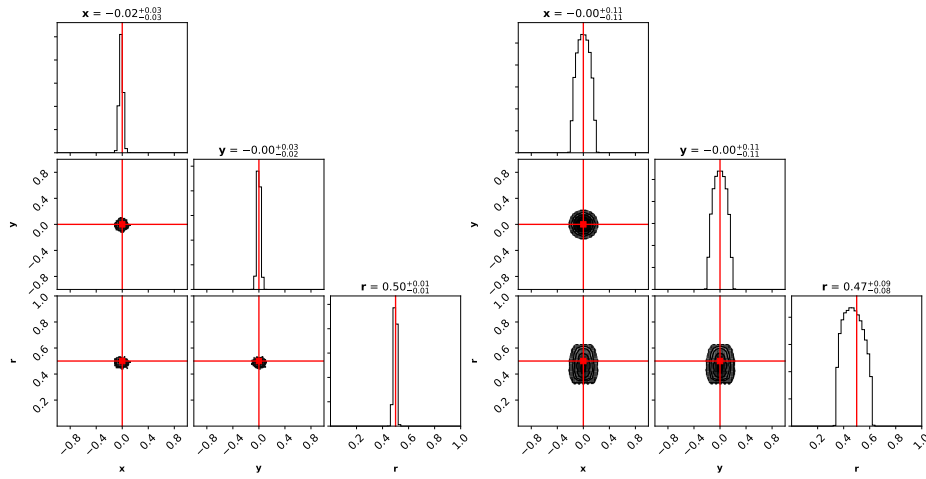


Fig. 3. Posteriors for the circle problem. The true generating parameter θ^* is depicted in red. *Left:* $p(\theta | \epsilon)$ from EG-LF-MCMC. *Right:* $p(\theta | \mathbf{x})$ from ABC

Stochastic linear model Secondly, we compare the two approaches on the benchmark problem of a stochastic linear model, adapted from [8, 12]. The model is specified as:

$$y = mx + b + \epsilon_1 |f(mx + b)| + 0.05\epsilon_2 \quad (5)$$

where ϵ_1 and ϵ_2 are drawn from the standard normal $\mathcal{N}(0, 1)$. In this model $mx + b$ calculates a deterministic value for y based on a standard linear function. The expression $\epsilon_1 |f(mx + b)|$ adds a stochastic error to it based on the magnitude of $mx + b$. The term $0.05\epsilon_2$ adds a further error. We generate a 100-dimensional observation \mathbf{x}_o by applying $\theta^* = (m = -0.9594, b = 4.294, f = 0.534)$ and measuring the values $(y_1, y_2, \dots, y_{100})$ at 100 equidistant x points. An example \mathbf{x}_o for this model is depicted on the right side of the Fig. 2. We assume a uniform

prior with the range of $[-5, 5]$ for m and the range of $[0, 10]$ for b and f . The smallest ϵ recorded by the training set of EG-LF-MCMC is 94.7. The obtained posteriors are presented in the Fig. 4.

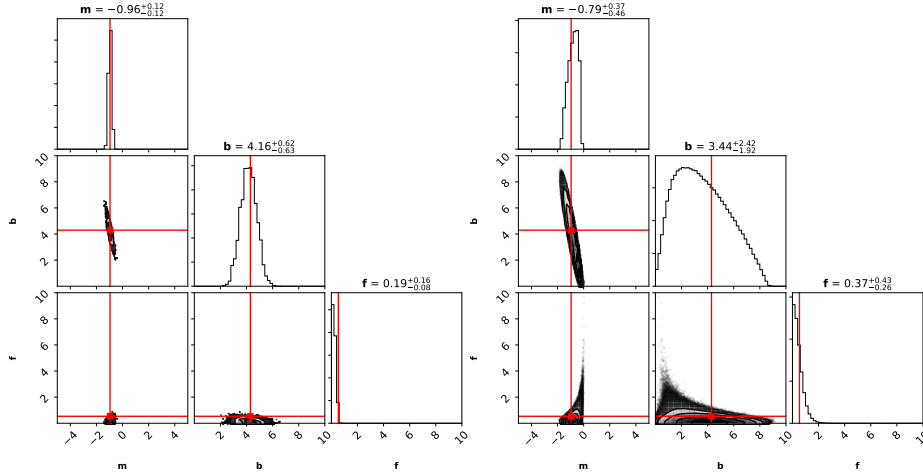


Fig. 4. Posteriors for the linear problem. The true generating parameter θ^* is depicted in red. *Left:* $p(\theta | \epsilon)$ from EG-LF-MCMC. *Right:* $p(\theta | \mathbf{x})$ from ABC

Discussion Despite requiring drastically less calls to the implicit generative model, as can be seen from Table 1, the posteriors obtained by EG-LF-MCMC are significantly more certain about the true generating parameter θ^* on both benchmark problems. Concerning the linear problem, both approaches recover f^* , but tend to underestimate it. We explain this behavior as follows. The term f controls the magnitude of the stochastic error in the observations. Since ϵ_1 is drawn from a standard normal, the error $\epsilon_1 |f(mx+b)|$ can shift an observation’s value y_x , where $1 \leq x \leq 100$, in both positive and negative directions. Thus, even if the true value f^* is inferred, across the true observation \mathbf{x}_o and a candidate observation \mathbf{x} the values y_x may be stochastically shifted in opposite directions, leading to a greater error. Since both EG-LF-MCMC and ABC try to infer the posteriors by minimizing the error, it is safer to assume a smaller value for f . This way \mathbf{x} will stay closer to \mathbf{x}_o on average, across different simulation runs.

4.1 Amortization over ϵ

We demonstrate the ability of EG-LF-MCMC to amortize the inference over ϵ , which is a completely new feature compared to AALR-MCMC. For this purpose, we turn again to the circle problem. We use the same trained classifier (amortization) as in the previously presented inference for θ^* in order to infer

approximate posterior densities for observations, which are a given distance ϵ away from \mathbf{x}_o . We denote such observations by \mathbf{x}_ϵ . Concretely, we would like \mathbf{x}_ϵ to differ from \mathbf{x}_o by 100 pixels. For this purpose, we condition the MCMC on $\epsilon = 100$. The sampled multivariate chain is depicted in Fig. 5.




Fig. 5. Circle problem: approximate posterior density $p(\boldsymbol{\theta} | \epsilon = 100)$

The approximated posterior densities make sense semantically: in order for the images to differ by 100 pixels, \mathbf{x}_ϵ needs to be generated with the coordinates and the radius, which revolve around, but tend to diverge from $\boldsymbol{\theta}^*$. To further verify the accuracy of our approach, we draw 10,000 $\boldsymbol{\theta}$ samples from the resulting Markov chain, simulate respective observations $\mathbf{x}_\epsilon \sim p(\mathbf{x} | \boldsymbol{\theta})$ and each time record the error ϵ in relation to the true observation \mathbf{x}_o . Finally, we calculate the average and the standard deviation from the sample of recorded ϵ values, which turn out to be 101.46 and 3.52 respectively, confirming the accuracy of EG-LF-MCMC.

5 Conclusions and Future Work

We have presented a novel method for likelihood-free inference, named error-guided likelihood-free MCMC. It is demonstrated that EG-LF-MCMC outperforms ABC in the ability to infer the true generating model parameter $\boldsymbol{\theta}^*$ on problems with semantically and structurally different high-dimensional observational data. The circle simulator generates deterministic images, while the linear model produces stochastic numerical samples. Furthermore, we have presented the ability of EG-LF-MCMC to amortize the inference over variable ϵ values. To the best of our knowledge, this feature is novel in the likelihood-free inference

literature, and it may enable novel opportunities for various scientific applications.

In the next phases of our research we plan to analytically study the connection between $p(\boldsymbol{\theta}|\epsilon)$ and the respective $p(\boldsymbol{\theta}|\mathbf{x})$. We hypothesize that they are equivalent when $\epsilon = 0$ and in other cases $p(\boldsymbol{\theta}|\epsilon)$ approximates $p(\boldsymbol{\theta}|\mathbf{x})$.

Acknowledgment

The computational results presented have been achieved in part using the Vienna Scientific Cluster (VSC).

References

1. Aeschbacher, S., Beaumont, M.A., Futschik, A.: A novel approach for choosing summary statistics in approximate bayesian computation. *Genetics* **192**(3), 1027–1047 (2012). <https://doi.org/10.1534/genetics.112.143164>, <https://www.genetics.org/content/192/3/1027>
2. Beaumont, M.A., Zhang, W., Balding, D.J.: Approximate bayesian computation in population genetics. *Genetics* **162**(4), 2025–2035 (2002), <https://www.genetics.org/content/162/4/2025>
3. Bishop, C.M.: Mixture density networks (1994), <http://publications.aston.ac.uk/id/eprint/373/>
4. Cranmer, K., Brehmer, J., Louppe, G.: The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences* (2020). <https://doi.org/10.1073/pnas.1912789117>, <https://www.pnas.org/content/early/2020/05/28/1912789117>
5. Durkan, C., Murray, I., Papamakarios, G.: On contrastive learning for likelihood-free inference. *International Conference on Machine Learning* (2020)
6. Durkan, C., Papamakarios, G., Murray, I.: Sequential Neural Methods for Likelihood-free Inference. *arXiv e-prints arXiv:1811.08723* (Nov 2018)
7. Fearnhead, P., Prangle, D.: Constructing summary statistics for approximate bayesian computation: semi-automatic approximate bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **74**(3), 419–474 (2012). <https://doi.org/10.1111/j.1467-9868.2011.01010.x>, <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9868.2011.01010.x>
8. Foreman-Mackey, D.: emcee. <https://emcee.readthedocs.io/en/v2.2.1/user/line/> (2012–2013)
9. Greenberg, D., Nonnenmacher, M., Macke, J.: Automatic posterior transformation for likelihood-free inference. *Proceedings of Machine Learning Research*, vol. 97, pp. 2404–2414. PMLR, Long Beach, California, USA (09–15 Jun 2019), <http://proceedings.mlr.press/v97/greenberg19a.html>
10. Hastings, W.K.: Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**(1), 97–109 (04 1970). <https://doi.org/10.1093/biomet/57.1.97>, <https://doi.org/10.1093/biomet/57.1.97>
11. Hermans, J., Begy, V.: Hypothesis. <https://github.com/montefiore-ai/hypothesis> (2019)
12. Hermans, J., Begy, V., Louppe, G.: Likelihood-free mcmc with amortized approximate ratio estimators. *International Conference on Machine Learning* (2020)

13. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv e-prints arXiv:1412.6980 (Dec 2014)
14. Klambauer, G., Unterthiner, T., Mayr, A., Hochreiter, S.: Self-normalizing neural networks. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 971–980. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/6698-self-normalizing-neural-networks.pdf>
15. Lueckmann, J.M., Goncalves, P.J., Bassetto, G., Öcal, K., Nonnenmacher, M., Macke, J.H.: Flexible statistical inference for mechanistic models of neural dynamics. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 1289–1299. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/6728-flexible-statistical-inference-for-mechanistic-models-of-neural-dynamics.pdf>
16. Marin, J.M., Pudlo, P., Robert, C.P., Ryder, R.J.: Approximate bayesian computational methods. *Statistics and Computing* **22**(6), 1167–1180 (2012)
17. Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of state calculations by fast computing machines. *The journal of chemical physics* **21**(6), 1087–1092 (1953)
18. Nunes, M.A., Balding, D.J.: On optimal selection of summary statistics for approximate bayesian computation. *Statistical Applications in Genetics and Molecular Biology* **9**(1) (06 Sep 2010). <https://doi.org/https://doi.org/10.2202/1544-6115.1576>, <https://www.degruyter.com/view/journals/sagmb/9/1/article-sagmb.2010.9.1.1576.xml.xml>
19. Papamakarios, G., Murray, I.: Fast ε -free inference of simulation models with bayesian conditional density estimation. In: *Advances in Neural Information Processing Systems*. pp. 1028–1036 (2016)
20. Papamakarios, G., Sterratt, D., Murray, I.: Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. *Proceedings of Machine Learning Research*, vol. 89, pp. 837–848. PMLR (16–18 Apr 2019), <http://proceedings.mlr.press/v89/papamakarios19a.html>
21. 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., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 32, pp. 8024–8035. Curran Associates, Inc. (2019), <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>