BAYESIAN UNCERTAINTY QUANTIFICATION MEETS TOPOLOGY

Anonymous authors
Paper under double-blind review

ABSTRACT

Computational topology recently started to emerge as an overarching paradigm for characterising the 'shape' of high-dimensional data, leading to powerful algorithms in (un)supervised representation learning. While capable of capturing prominent features at multiple scales, topological methods cannot readily quantify the *uncertainty* of their respective descriptors. We develop a novel approach that bridges this gap, making it possible to employ topology-based loss functions to perform parameter estimation with Bayesian uncertainty quantification. Our method affords easy integration into topological machine learning algorithms. We demonstrate its efficacy for parameter estimation in different simulation settings.

1 INTRODUCTION

Topological machine learning methods enable describing the 'shape' of data at multiple scales while remaining impervious to many different types of noise. This led to strong hybrid models that combine geometry and topology in different domains, including computer vision (Hu et al., 2019; Waibel et al., 2022), graph learning (Horn et al., 2022; Yan et al., 2022; Zhao et al., 2020), time series analysis (Zeng et al., 2021), and unsupervised representation learning (Carrière et al., 2020; Moor et al., 2020). However, despite the advantageous properties of such integrations, topological methods cannot readily perform uncertainty quantification. Such scenarios are common occurrences in real-world data analysis tasks, and are frequently handled with methods from Bayesian statistics. For instance, assuming a known prior distribution $p(\theta)$ and real-world data y, one could employ Bayes' formula to obtain the *posterior* distribution $p(\theta|y) = \frac{p(y \mid \theta)}{p(y)} p(\theta)$. The posterior distribution allows for a comprehensive description of parameter uncertainties. Regrettably, the application of this formulation to complex generative processes presents substantial challenges. In these instances, the likelihood function $p(y \mid \theta)$, and consequently, the marginal likelihood $p(y) = \int p(y \mid \theta) p(\theta) \mathrm{d}\theta$, frequently become *intractable*. The enormous complexity in such systems often requires methods that capture structural properties of the underlying process that are invariant to isometries (such as rotations and translations), making computational topology a well-suited tool. Consequently, this necessitates an alternative framework for uncertainty quantification that satisfies two key requirements:

(i) It must effectively utilise topological information in the observed data y.
(ii) It should offer a theoretically sound notion of posterior beliefs.

Both criteria are critical to overcoming the limitations of topological machine learning methods in complex scenarios and enhancing their applicability to real-world data analysis tasks where uncertainty quantification is important. To effectively articulate our posterior beliefs, it is essential to establish the pertinent constraints. Primarily, in the absence of a likelihood function, we posit that access to information on the parameter θ is *solely available* through samples from generative simulation models, which we denote by $x \sim p(x \mid \theta)$. Additionally, we assume that we have access to a *loss function* of the form $(x, y) \mapsto \ell(x, y)$, which measures to what extent a generated sample x is 'close' to the observed data y. In the general setting as outlined above, we endeavour to demonstrate that the distribution $\pi(\theta \mid y) = \frac{\int \exp(-\ell(y,x)) p(x \mid \theta) p(\theta) \mathrm{d}x}{\iint \exp(-\ell(y,x)) p(x \mid \theta) p(\theta) \mathrm{d}x \mathrm{d}\theta}$ offers an alternative for encapsulating the inherent uncertainty in such generative models. As the **central contributions** of this paper, we thus (i) leverage Bayesian inference and topology via efficient geometrical–topological loss terms that are capable of handling different modalities (point clouds and images), (ii) demonstrate how sampling algorithms can leverage topologically informed loss functions to perform Bayesian

inference, and (iii) demonstrate the utility of our approach to uncertainty quantification on several complex generative models. As a main limiting factor we would like to point out that although computational topology seems to be the appropriate tool for certain complex models, the resulting performance for parameter estimation highly depends on the setting. A careful assessment of the given scenario is recommended beforehand. Moreover, we do not claim that our methodology outperforms benchmark Bayesian approaches under the assumption of a known likelihood function.

2 Background: Computational Topology

Before we introduce our approach for Bayesian uncertainty quantification, we first provide a brief overview of relevant methods from computational topology. Such methods recently emerged, making use of geometrical–topological properties of data to describe their overall 'shape.' The concept of *persistent homology* is of particular interest in the context of machine learning since it leads to efficient descriptors of structured and unstructured data while at the same time satisfying invariance and robustness guarantees (Hensel et al., 2021). In the context of analysing complex generative models, invariance to isometries, for instance, is of particular interest. Persistent homology obtains multi-scale geometrical–topological shape descriptors of data by means of combinatorial data structures. Given a metric such as the Euclidean distance, *persistent homology* assigns a family of shape descriptors, the *persistence diagrams*, to point clouds (Edelsbrunner & Harer, 2010); the process can be extended to other modalities. The calculation of persistence diagrams turns out to be stable with respect to the Hausdorff distance between point clouds, making this method highly robust in practice (Cohen-Steiner et al., 2007). See Fig. 1 for an overview of this process.

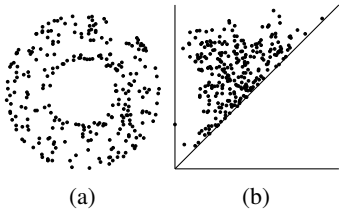

Figure 1: Given a point cloud (a), *persistent homology* lets us obtain shape descriptors known as *persistence diagrams* (b).

Persistent homology originated from *simplicial homology*, a combinatorial method to obtain simple shape descriptors from a simplicial complex K, i.e. from a high-dimensional generalisation of a graph. Although simplicial homology is a powerful topological tool that captures the shape of the underlying object of interest, a single simplicial complex often does not provide sufficient insight into complex data. The main idea is thus to obtain a nested sequence of simplicial complexes from the given data that encodes topological information about the data *on multiple scales*. The notion of 'scale' and how to define such a sequence in practice depends on the given problem; we will describe two constructions that are relevant in the context of this paper. Given a point cloud X, the *Vietoris-Rips complex* at stage ϵ associated to X is the abstract simplicial complex in which a k-simplex is defined for every subset of X, consisting of $k+1$ points in X that have diameter at most 2ϵ. By varying ϵ one obtains a sequence of nested simplicial complexes associated to X. The second construction is given by *cubical complexes*. Since we only use cubical complexes for greyscale images, we restrict ourselves to this particular setting, and refer to the literature for a more comprehensive introduction (Rieck et al., 2020b; Wagner et al., 2012). Any greyscale image A can be interpreted as a graph, where the nodes are given by the non-zero pixels, such that there is an edge between two nodes if and only if the two nodes correspond to neighbouring pixels. For a threshold ϵ, we obtain a new image from A by setting all pixels in A to zero, which admit a pixel value less than ϵ. Varying ϵ leads to a nested sequence of the graphs corresponding to the respective images, and therefore we obtain a nested sequence corresponding to A. Subsequently, we explain how to derive *multi-scale topological features* from such a nested sequence, also known as a *filtration*:[1] Let $\emptyset = \mathrm{K}_0 \subseteq \mathrm{K}_1 \subseteq \cdots \subseteq \mathrm{K}_{m-1} \subseteq \mathrm{K}_m = \mathrm{K}$ be a nested sequence of simplicial or cubical complexes. The main idea of persistent homology involves tracking topological features, measured by means of *homology groups*, alongside this filtration. The family of boundary operators $\partial(\cdot)$, together with the inclusion homomorphism, induces a homomorphism between corresponding *homology groups* of the filtration, i.e. $\iota_d^{i,j} \colon H_d(\mathrm{K}_i) \to H_d(\mathrm{K}_j)$. This homomorphism yields a sequence of homology groups for every dimension d. Given indices $i \leq j$, the dth *persistent homology group* is then defined as $H_d^{i,j} := \ker \partial_d(\mathrm{K}_i) / \left(\operatorname{im} \partial_{d+1}(\mathrm{K}_j) \cap \ker \partial_d(\mathrm{K}_i)\right)$. It can be seen as the homology group that contains all homology classes created in K_i that are still present in K_j. Typically, the filtration

[1] For a more in-depth discussion of computational topology, in particular in the context of machine learning, we refer the reader to Hensel et al. (2021).

of K has associated values $a_0 \leq a_1 \leq \cdots \leq a_{m-1} \leq a_m$ (such as edge weights of a graph or pairwise distances of a point cloud). This permits us to summarise multi-scale topological activity in a *persistence diagram*: for each dimension d and each pair $i \leq j$, we store the pair $(a_i, a_j) \in \mathbb{R}^2$ with multiplicity $\mu_{i,j}^{(d)} := \left(\beta_d^{i,j-1} - \beta_d^{i,j}\right) - \left(\beta_d^{i-1,j-1} - \beta_d^{i-1,j}\right)$ in a set (in practice, $\mu_{i,j}^{(d)} = 0$ for many pairs). The resulting set of points is called the dth *persistence diagram* $\mathcal{D}_d$. Given a point $(x, y) \in \mathcal{D}_d$, the quantity $\operatorname{pers}(x, y) := |y - x|$ is referred to as its *persistence*. High persistence is commonly considered to correspond to *features*, while low persistence is seen to indicate *noise* Edelsbrunner & Harer (2010), but recent work shows that in many data sets, low persistence values rather indicate 'low reliability,' which may still play an important role for downstream analyses (Bendich et al., 2016; Rieck et al., 2020b).

3 Methods

Numerous theoretical justifications have been proposed for employing probability distributions as a measure of uncertainty, including seminal arguments such as the 'Dutch Book' theorem (De Finetti et al., 1937). While these justifications provide compelling reasons for probabilistic approaches to uncertainty quantification, within the scope of this paper, we will assert that probability distributions represent a cogent, effective means for quantifying uncertainty. Our focus is not to debate the philosophical underpinnings of this stance, but rather to leverage its practical utility.

3.1 Bayesian Inference

Within the Bayesian framework, unknown quantities, such as parameters denoted by θ, are furnished with a probability distribution $p(\theta)$, referred to as the *prior distribution*. Subsequent inference about the parameters θ given observed data y is performed via the Bayesian update procedure. This process involves adjusting the prior distribution $p(\theta)$ through the multiplication by the ratio of the likelihood to the evidence, i.e. $p(\theta|y) = \frac{p(y|\theta)}{p(y)} p(\theta)$. However, applying this procedure in our particular context presents notable challenges. The primary issue stems from the fact that the likelihood function $p(y \,|\, \theta)$ is typically unknown and poses considerable difficulties for estimation, particularly when the data exhibits intricate geometrical–topological properties. To circumvent this problem, recent research has shown that the use of the likelihood function of the data can be avoided, producing so-called generalised Bayesian posteriors (Bissiri et al., 2016). One approach to construct such a generalisation to the Bayesian posterior can be found through the lens of viewing Bayesian inference as the solution to an optimisation problem (Csiszár, 1975; Donsker & Varadhan, 1975; Zellner, 1988). Specifically, it can be seen as finding a distribution, q, that balances prior information as measured by the Kullback–Leibler (KL) divergence between the prior and the posterior candidate q (which we want to *minimise*) and the expected log-likelihood of new observations y (which we want to *maximise*):

$$\operatorname{argmin}_{q \in \mathcal{P}(\Theta)} \left\{ \mathbb{E}_q[-\log p(y|\theta)] + \mathrm{KL}(q(\theta), p(\theta)) \right\}. \tag{1}$$

The solution to this problem is the traditional Bayesian posterior $q^*(\theta) = p(\theta|y)$. Recently, generalisations of this approach have been suggested, demonstrating that coherent belief updates can also occur if instead of the negative log-likelihood loss, $-\log p(y \,|\, \theta)$ any loss $l(y, \theta)$ is taken (Bissiri et al., 2016; Knoblauch et al., 2022). The solution to the concomitant optimisation problem

$$\operatorname{argmin}_{q \in \mathcal{P}(\Theta)} \left\{ \mathbb{E}_q[l(y, \theta)] + \mathrm{KL}(q(\theta), p(\theta)) \right\} \tag{2}$$

is given by

$$q^*(\theta) = \frac{\exp(-l(y, \theta)) p(\theta)}{\int \exp(-l(y, \theta)) p(\theta) \mathrm{d}\theta} =: \pi_{\mathrm{GBI}}(\theta \,|\, y). \tag{3}$$

However, the problem is that even in this general formulation, the loss formally scores observed data y with respect to a parameter θ. A common approach is to use expectations over samples from a generative model, i.e. simulate $x \sim f(x \,|\, \theta)$ and approximate

$$l(y, \theta) := E_{p(x \,|\, \theta)} \left[\ell(y, x) \right] \approx \frac{1}{N} \sum_{i=1}^{N} \ell(y, x_i), \tag{4}$$

with $x_i \sim p(x \mid \theta)$. Replacing the loss $l(y, \theta) = E_{p(x \mid \theta)}[\ell(y, x)]$ with a sample average, however, changes[2] the posterior from Eq. (3). We will demonstrate in the following that we can directly construct a valid posterior from the loss function $\ell(y, x)$.

3.2 Comparison-Based Posteriors

Uncertainty estimates predicated on distance functions have a long-standing history in Bayesian statistics as a means of *approximating* posterior distributions, with notable use in the realm of approximate Bayesian computation (ABC, (Tavaré et al., 1997; Pritchard et al., 1999; Beaumont et al., 2002)). In the ensuing discussion, we will illustrate that a modification of the aforementioned approach, specifically, a generalised posterior from approximate Bayesian computation such as Schmon et al. (2021), can furnish us with a theoretically sound form of posterior update. The proposed approach involves using a *loss function* $(x, y) \mapsto \ell(y, x)$ to construct a joint posterior involving simulated data $x \sim p(x \mid \theta)$ and real data y in the form

$$\pi(\theta, x \mid y) \propto \exp\left(-\ell(y, x)\right) p(x \mid \theta) p(\theta). \tag{5}$$

This approach, however, encounters several obstacles. Firstly, it is not clear that Eq. (5) defines a reasonable notion of uncertainty. In addition, since x is just *simulated* data, not *observed* data, we need to average over possible values, leading to

$$\pi(\theta \mid y) = \frac{\int_{\mathsf{X}} \exp\left(-\ell(y, x)\right) p(x \mid \theta) p(\theta) \mathrm{d}x}{\iint_{\mathsf{X} \times \Theta} \exp\left(-\ell(y, x)\right) p(x \mid \theta) p(\theta) \mathrm{d}x \mathrm{d}\theta}. \tag{6}$$

It turns out that this is is a valid (generalised) posterior under the loss associated with the topological properties of our data, namely:

Proposition 3.1. *The comparison-based posterior*

$$\pi(\theta, x \mid y) = \frac{\exp\left(-\ell(y, x)\right) p(x \mid \theta) p(\theta)}{\int_{\mathsf{X}} \int_{\Theta} \exp\left(-\ell(y, x)\right) p(x \mid \theta) p(\theta) \mathrm{d}\theta \mathrm{d}x} \tag{7}$$

is the solution $q^*(\theta, x) = \pi(\theta, x \mid y)$ *to the optimisation problem*

$$q^* = \operatorname{argmin}_{q \in \mathcal{P}(\Theta \times \mathsf{X})} \left\{ \mathbb{E}_q[\ell(y, x)] + \mathrm{KL}\left(q(\theta, x), p(x|\theta)\pi(\theta)\right) \right\}. \tag{8}$$

The calculations follow usual arguments and can be found in the appendix. In Eq. (8), the first term in the objective function measures the expected loss under the posterior, where the loss function $\ell(y, x)$ quantifies the discrepancy between the observed data y and the simulated data x. This is precisely where we can use improved inductive biases or expert knowledge in employing loss functions such as the topological losses introduced in Section 3.4. The second term encourages the posterior to stay close to the prior, as in the traditional Bayesian setting, however, with a significant difference: the prior distribution $p(\theta)$ over the parameters θ together with the simulation model $p(x \mid \theta)$ naturally implies prior beliefs over what the data should look like. Thus, it is sensible to consider the KL-divergence between the joint beliefs $p(x \mid \theta)p(\theta)$ and $q(x, \theta)$.

Our approach can be seen as a generalisation of the traditional Bayesian posterior in the sense that it allows for a more flexible modelling of the data generation process. Instead of assuming that the observed data is generated exactly according to the model $p(y \mid \theta)$, we allow for the possibility that the data is generated from a distribution that produces data *close* to y in terms of the loss function ℓ. This makes our approach potentially more robust to model misspecification and more suitable for complex data with intricate topological properties.

3.3 Inference for Comparison-Based Posteriors

Evaluating posteriors such as those given by Eq. (6) analytically is typically *infeasible*. An alternative approach may involve direct targeting of a variational distribution as shown in Eq. (8). However, this strategy encounters two main difficulties: firstly, the distribution $p(x \mid \theta)$ is often intractable, rendering the computation of the KL divergence term non-trivial. Secondly, in order to compute

[2]This occurs even in the case where the estimator is unbiased since this property is lost under the non-linear (exponential) transformation by virtue of Jensen's inequality.

Algorithm 1 Importance sampling estimation of the mean

Require: Observed data y, test function h
1: **for** $i = 1 : n$ **do**
2: Sample $\theta_i \sim p(\theta)$
3: Sample $x_i \sim p(x \mid \theta_i)$
4: **end for**
5: **return**

$$\hat{h}(\theta) = \frac{\sum_{i=1}^{n} h(\theta_i) e^{-w\ell(y,x_i)}}{\sum_{i=1}^{n} e^{-w\ell(y,x_i)}}$$

the gradients, we must propagate them through the simulation models. This poses a challenge as $p(x \mid \theta)$ often represents a 'black-box' model, not necessarily implemented within a framework that supports automatic differentiation. Consequently, this typically precludes the practical application of gradient-based optimisation techniques. However, it turns out that Monte Carlo algorithms, such as self-normalised importance sampling or so-called pseudo-marginal Markov chain Monte Carlo (Andrieu & Roberts, 2009) are able to compute expectations with respect to posteriors of the form in Eq. (6).

Algorithm 1 shows a self-normalised *importance sampling* procedure. Almost-sure convergence can be shown under mild regularity conditions (e.g. Owen, 2013, Theorem 9.2) for a test function of interest h, i.e.

$$\frac{\sum_{i=1}^{n} h(\theta_i) e^{-\ell(y,x_i)}}{\sum_{i=1}^{n} e^{-\ell(y,x_i)}} \xrightarrow{n\to\infty} \frac{\iint h(\theta) e^{-\ell(y,x)} p(x \mid \theta) p(\theta) \mathrm{d}x \mathrm{d}\theta}{\iint e^{-\ell(y,x)} p(x \mid \theta) p(\theta) \mathrm{d}x \mathrm{d}\theta} = E_{\pi(\theta \mid y)}[h(\boldsymbol{\theta})]. \tag{9}$$

An alternative is to use *Markov chain Monte Carlo* (MCMC), where we produce a Markov chain converging to the distribution (6); see Algorithm 2 in the appendix. Such algorithms indeed converge to the desired target distribution (Andrieu & Roberts, 2009, Theorem 1) under mild assumptions.

3.4 Losses Based on Geometry & Topology

Having introduced multiple inference algorithms and explained how to obtain topological features, we now discuss how to derive topology-based loss functions. We start with discussing metrics in computational topology. It turns out that persistence diagrams can be endowed with a metric by using optimal transport. Given two diagrams $\mathcal{D}$ and $\mathcal{D}'$ containing features of the same dimensionality, their pth *Wasserstein distance* is defined as

$$\mathrm{W}_p(\mathcal{D}, \mathcal{D}') := \left(\inf_{\eta \colon \mathcal{D} \to \mathcal{D}'} \sum_{x \in \mathcal{D}} \|x - \eta(x)\|_\infty^p \right)^{\frac{1}{p}}, \tag{10}$$

where $\eta(\cdot)$ denotes a bijection. Since $\mathcal{D}$ and $\mathcal{D}'$ generally have different cardinalities, we consider them to contain an infinite number of points of the form (τ, τ), i.e. tuples of zero persistence; this is akin to requiring each diagram to contain the projections of points to the diagonal, originating from the *other* diagram. A suitable $\eta(\cdot)$ can thus always be found. Solving Eq. (10) is practically feasible using modern optimal transport algorithms Flamary et al. (2021). While this is the most 'principled' approach—in the sense that such a loss formulation forms a proper metric in the mathematical sense—alternative formulations to Eq. (10) exist, and our framework is fundamentally compatible with all of them. Along these lines, other topological descriptors, such as *Betti curves* Rieck et al. (2020a) or *persistence images* Adams et al. (2017), might be used instead of persistence diagrams. These descriptors are often easier to compute *and* afford the use of fast L_1 or L_2 distances as proxies for Eq. (10). For instance, the L_2 distance between persistence images is known to share some advantageous properties with the Wasserstein distance Adams et al. (2017) but it only handles a 'discretised' version of the data, so some information is invariably lost (or, to briefly consider an optimistic point of view, the persistence image is *smoother* than a persistence diagram). As another alternative to the previously-described descriptors, we could also employ a *kernel*, i.e. a similarity measure between persistence diagrams. These similarity measures are not metrics in the mathematical sense, lacking the requirement of the identity of indiscernibles, but are easier to compute (unlike the persistence images, they do not require vectorisations) and require fewer parameters Kwitt et al. (2015); Reininghaus et al. (2015).

A loss function based on the Wasserstein distance. To leverage the power of computational topology and Bayesian inference, we propose combining the Bayesian framework for updating posterior probability estimates with a *topological loss function*. Given observed data y and simulated data x sampled from $p(x \mid \theta)$, let $\mathcal{D}_x$ and $\mathcal{D}_y$ be the persistence diagrams corresponding to x and y, respectively. We then define a *topological loss function* $\ell_T \colon \mathsf{Y} \times \mathsf{X} \to \mathbb{R}$, by setting

$$\ell_T(y, x) := \mathrm{W}_p(\mathcal{D}_y, \mathcal{D}_x). \tag{11}$$

This formulation applies to different data modalities. If the data can be represented as a point cloud, we calculate persistence diagrams from the Vietoris–Rips complex. For (greyscale) image data, we obtain persistence diagrams via *cubical complexes*, enabling topological feature calculations for more complicated modalities such as MRI data Rieck et al. (2020b). As we outlined above, other choices for the loss function would be possible, and we may now continue with our Bayesian updating procedure. Note that unless otherwise mentioned, we use $p = 2$ to compute the topological loss.

Properties. The Wasserstein distance between persistence diagrams is *stable* in the sense that the geometric distance constitutes an upper bound of the topological distance. Given two point clouds Y, X of the *same cardinality*, we have $\mathrm{W}_p(\mathcal{D}_y, \mathcal{D}_x) \leq C_p \, \mathrm{W}_p(\mathsf{Y}, \mathsf{X})$, where C_p refers to a constant that only depends on p and $\mathrm{W}_p(\mathsf{Y}, \mathsf{X})$ is defined similarly to Eq. (10), but calculated on the *point clouds themselves* Skraba & Turner (2022). Since $\mathrm{W}_p(\cdot, \cdot) \leq \mathrm{W}_{p'}(\cdot, \cdot)$ for $p' \leq p$, low values of p are desirable in terms of stronger stability; we find that $p = 2$ provides a suitable compromise solution. Moreover, our loss function ℓ_T is invariant under isometries (Edelsbrunner & Harer, 2010) and stable under subsampling (Moor et al., 2020).

A loss function based on the Hausdorff distance. As a baseline and computationally simpler comparison partner, we also define a geometry-based loss function based on the *Hausdorff distance* between point clouds X and Y. Given a metric space (M, d) and two non-empty subsets $\mathsf{X}, \mathsf{Y} \subseteq M$, we define our loss as

$$\ell_G(\mathsf{X}, \mathsf{Y}) := \inf\{\epsilon \geq 0 \,|\, \mathsf{X} \subset \mathsf{Y}_\epsilon, \mathsf{Y} \subset \mathsf{Y}_\epsilon\}, \tag{12}$$

where $\mathsf{X}_\epsilon = \cup_{x \in \mathsf{X}} \{m \in M;\ d(m, x) \leq \epsilon\}$ denotes the ϵ-thickening of X in M. While this loss does not satisfy invariance properties, it is more efficient to compute in practice. However, as we will see in the experimental section, its utility in complex data generation scenarios is limited.

3.5 Related Work

A range of Bayesian inference methods have been developed to address the challenges involved with uncertainty quantification when dealing with complex models whose likelihoods are analytically intractable. While each of these approaches offers a distinct solution, they also present their own unique sets of challenges and limitations.

Approximate Bayesian computation (ABC) is a well-established approach (Tavaré et al., 1997; Pritchard et al., 1999; Beaumont et al., 2002). ABC employs a pre-defined threshold to accept parameters that yield a small enough distance between the simulated and observed data. More recently, Wasserstein ABC (Bernton et al., 2019) has capitalised on optimal transport theory to define a distance function between simulated and observed data. This form of rejection-based ABC comes with several potential problems such as the choice of the rejection threshold: a small threshold will lead to high numerical costs whereas a threshold too large can result in incorrect coverage of credible regions (Frazier et al., 2018). To make matters worse, rejection ABC algorithms are known to exhibit non-standard asymptotic behaviour in case of model misspecification (Frazier et al., 2019).

Several authors have suggested constructing a generalised posterior along the lines of Eq. (3) using distance functions, such as *maximum mean discrepancy* (MMD, Park et al. (2015); Chérief-Abdellatif & Alquier (2020)). However, these methods, in contrast to our approach, use a theoretical loss, for instance, $l(y, \theta) = E_{p(x|\theta)}[\ell(y, x)]$, which is then *approximated* using empirical averages, thus potentially introducing a source of bias or instability. By contrast, our approach, is more related to (particle) filtering (e.g. Doucet et al., 2009), where inference is carried out over the latent state simultaneously. Miller & Dunson (2018) introduced a discrepancy-based Bayesian procedure they termed 'coarsening,' where the observed data is considered as a coarsened version of latent data, thereby enhancing robustness. The work that aligns most closely with our approach to uncertainty

quantification is that of Schmon et al. (2021), which introduces (5) as a generalised version of ABC, although not in the context of topological data analysis.

Finally, topological data analysis (TDA) has already started to see some use in the context of analysing generative models. Some works are concerned with fitting the parameters of previously-defined prior distributions of persistence diagrams (Maroulas et al., 2022; Oballe, 2020), thus demonstrating the expressivity of TDA for *describing* complex systems. We are interested in parameter estimates of black-box models by exploiting this expressivity. Along these lines, Topaz et al. (2015) use persistent homology for a qualitative analysis of the topological characteristics of swarming models (such as Vicsek model). In comparison, our approach is intended to leverage topology and Bayesian inference for *quantitative* studies, enabling parameter estimation. More recently, Thorne et al. (2022), used topological statistics in combination with *rejection sampling*. This method constitutes the first use of TDA in the setting of ABC; it involves numerous parameter choices (intrinsic to the definition of such statistics) that our method does not necessitate. Moreover, the rejection sampling approach is itself parameter-driven. Finding a suitable rejection threshold requires additional domain knowledge. Our approach stands out by integrating a loss function to construct a posterior distribution directly, without resorting to approximation, thereby offering a theoretically-grounded form of posterior belief.

4 Experiments

As a consistency check, we use our proposed methods for inferring the radius parameters of two objects, namely the 2-sphere and the 2-torus. Table 3 in the appendix depicts the results of the inference process. For the remainder of this section, we empirically validate the utility of the proposed method in the form of three main experiments that compare topological losses and respective geometrical losses, by using the results from Section 3. Moreover, we provide a comparison to standard rejection sampling procedures with respect to summary statistics, see Appendix A.

4.1 Vicsek Model: Swarm Behaviour

The Vicsek model is an agent-based model that is particularly useful to study collective motion like swarming and flocking. The underlying idea is that each agent aligns at any time step with its nearest neighbours, before the alignment is perturbed by an additional random term. The resulting angle of the alignment determines the direction of movement for the subsequent time step, where the movement takes place with constant speed which is independent of time and individual. Following Vicsek et al. (1995), the formula for the alignment angle update of an individual i is given by $\theta(t+1) = \langle\theta(t)\rangle_r + \Delta\theta$, where $\langle\theta(t)\rangle_r$ denotes the average of the angles of individuals with distance to agent i at most r, and $\Delta\theta$ is uniformly sampled from the interval $[-\eta/2, \eta/2]$ for $0 \leq \eta \leq 2\pi$. The new position of the agent is then determined by moving with constant speed towards the direction of the updated angle. Although the formula for each agent is simple, collective behaviour of the system over time is highly chaotic due to the complex dependencies between agents.

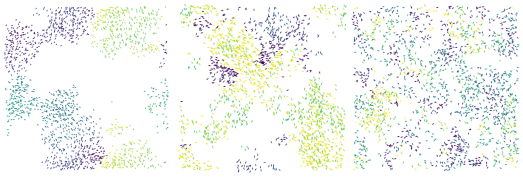

Figure 2: Left to right: simulations of the Vicsek model for $\eta = 0.15, 0.3$ and 0.6, respectively. For lower values of η swarming behaviour emerges, whereas for higher values of η swarms tend to merge and the model becomes more ergodic.

To demonstrate the utility of topological losses for parameter inference in a setting where the respective likelihood function is *not analytically tractable*, we infer the noise parameter η after a certain number of iterations (time steps) of the model. We distribute 2000 agents across a square of edge length $L = 25$, where opposite edges of the square are identified with the orientation being preserved; the 'world' thus constitutes the surface of 2-torus. For the inference of η, we use 250 simulations for each sampling method, aiming to infer η after $5, 10,$ and 50 iterations of the model, respectively, repeating the inference 5 times. Since the Vicsek model gives rise to point clouds, we calculate Vietoris–Rips complexes and compare the corresponding persistence diagrams. As a baseline, we use the Hausdorff distance between the point clouds; Fig. 3 depicts the results. In our experiments, the results for the topological loss with the MCMC sampling procedure outperforms in all of the settings, and leading to highly accurate estimates even after 50 time steps. By contrast, we note that the estimators that were calculated with the Hausdorff loss lead

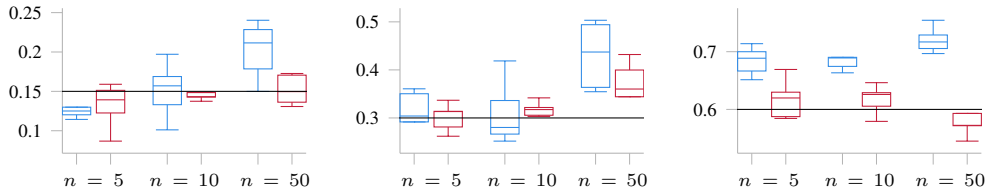

(a) Importance sampling

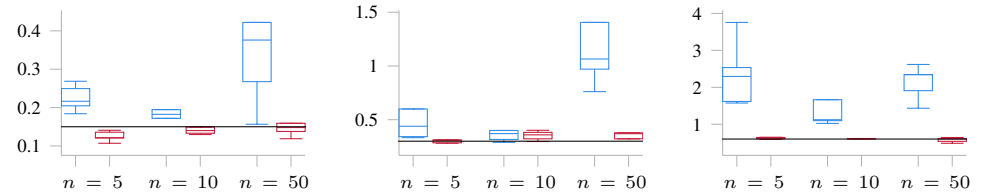

(b) MCMC sampling

Figure 3: Summary of parameter estimates for the noise parameter η of the Vicsek model, using 5 repetitions of the inference procedure. Each subplot depicts estimates for $\eta \in \{0.15, 0.30, 0.60\}$, respectively. The true parameter is shown as a solid black line. Estimates based on the geometrical loss ℓ_G are less accurate than the topological loss ℓ_T, in particular for larger number of iterations. Please refer to Table 4 for the raw values.

to particularly poor results in the regimes that are closer to being ergodic, i.e. for larger values of η. Finally, we remark that our approach with the proposed topological loss outperforms standard rejection sampling in combination with certain summary statistics, see Appendix A.

4.2 Fluid Dynamics: Lattice Boltzmann Simulation

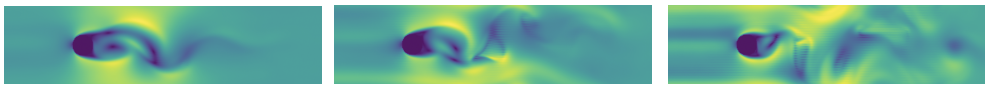

(a) low turbulence (b) medium turbulence (c) high turbulence

Figure 4: Left to right: simulations of the fluid model after 3000 iterations, for $\eta = 0.2, 0.3$ and 0.4, respectively. For lower values of η, the fluid flow is less turbulent than for higher values of η.

The lattice Boltzmann methods (LBM) are a collection of algorithms that are used for fluid simulation. Instead of the conservation of macroscopic properties of the dynamical system, LBM models the fluid over a discrete lattice, by performing local propagation and collision processes, for all particles simultaneously. For a thorough introduction, see Krüger et al. (2017). After initialising the state of each particle equally, a portion of normally distributed noise is added to each state, independently. This portion of randomness is controlled by a global parameter $\eta \geq 0$, where higher values of η correspond to a higher portion of noise in the initialisation. The crucial observation is that the degree of turbulence in the system can be modelled by varying η, as is shown in Fig. 4. Our objective is therefore to infer the parameter η, which corresponds to a given observation after a certain amount of iterations in the model, where we fix this number of iterations to be $t = 3000$ timesteps. As the topological loss, we use the Wasserstein

Table 1: Estimating the randomness parameter of the lattice Boltzmann model after 3000 simulation iterations, using 5 repetitions of the inference procedure. The topological loss together with importance sampling outperforms all of the other configurations.

		Randomness parameter η		
Loss	*Sampling method*	0.20	0.30	0.40
ℓ_G	Importance sampling	0.16 ± 0.02	0.22 ± 0.04	0.35 ± 0.13
	MCMC	0.15 ± 0.10	0.91 ± 1.64	0.38 ± 0.25
ℓ_T	**Importance sampling**	$\mathbf{0.21 \pm 0.01}$	$\mathbf{0.27 \pm 0.00}$	$\mathbf{0.38 \pm 0.01}$
	MCMC	2.28 ± 1.40	1.71 ± 0.81	2.35 ± 1.10

distance between the persistence diagrams of the *cubical complexes* corresponding to the observed and the simulated state, respectively, while the Hausdorff distance between the states serves as a baseline comparison. Table 1 shows the results; we observe that the topological loss together with importance sampling outperforms the geometrical loss in both configurations, by far. Again, rejection sampling in combination with summary statistics performs less reliably, see Appendix A.

4.3 Percolation Model: Multi-Scale Structures

In statistical physics, percolation refers to the behaviour of a network when links are added to it. Percolation can be viewed as a process of geometric phase transition since at a critical state of the network, disconnected components merge into large connected clusters by the addition of a small fraction of links. In our setting, the percolation model is probabilistic in the sense that links are added randomly, with some probability p, and the goal is to infer p. We realise such a network as a square 2D greyscale image of pixel size n^2, for some positive integer n.[3] We assign a value v to each pixel in the image, where $v = 0$ with probability $1 - p$, and v is sampled uniformly from the set $\{1, \dots, v_{\max}\}$ with probability p. Here, $v_{\max}$ is the maximum realisable greyscale value, which we set to 50 for our experiments. For the inference of p, we used 250 simulations for both the importance sampling and the MCMC sampling. Since we are dealing with greyscale images, we obtain persistence diagrams via *cubical complexes*. We compare our proposed topological loss ℓ_T to several other losses that are commonly used in imaging processing, including MSE, RMSE, the universal image quality index (UQI), the relative average spectral error (RASE), the spatial correlation coefficient (SCC), and the pixel-based visual information fidelity (VIFP). Since we find that the results of SCC and VIFP outperform other scores, and since SCC and VIFP lead to comparable results, we only show the results for SCC (see Zhou et al. (1998) for a construction of SCC) in comparison to the ones obtained by the topological loss. Table 2 shows the results. We again observe that our topology-based loss ℓ_T, together with MCMC sampling, outperforms the other methods by far. However, in this particular experiment rejection sampling together with a mean statistic performs even better, we refer to Appendix A for a discussion.

Table 2: Estimating the parameter of the Percolation model, using 5 repetitions of the inference procedure. The topological loss together with either importance sampling or MCMC sampling outperforms all other configurations.

		Randomness parameter p		
Loss	*Sampling method*	0.15	0.30	0.60
ℓ_G	Importance sampling	0.18 ± 0.00	0.30 ± 0.47	0.59 ± 0.00
	MCMC	3.18 ± 2.55	1.30 ± 0.47	2.18 ± 1.47
ℓ_T	**Importance sampling**	$\mathbf{0.15 \pm 0.00}$	$\mathbf{0.29 \pm 0.00}$	$\mathbf{0.59 \pm 0.01}$
	MCMC	$\mathbf{0.15 \pm 0.02}$	$\mathbf{0.29 \pm 0.00}$	$\mathbf{0.59 \pm 0.00}$

5 Discussion

We presented a novel approach for uncertainty estimation in topological data analysis. Our method, which builds on the framework of generalised Bayesian inference, introduces a topology-based loss function into the construction of the posterior distribution. This approach offers a theoretically grounded form of posterior belief that improves upon some of the challenges inherent in existing distance-based methods, which rely on geometrical features. We have empirically shown that our method outperforms such geometrical approaches for parameter estimation in the setting of complex systems, which suggests future applications in other complex scenarios such as the life sciences.

Our work contributes to the ongoing evolution of Bayesian inference methods for complex models. It presents a new avenue of exploration that merges the advantages of both topological data analysis and Bayesian statistics. We believe that the principles and techniques introduced in this paper have broad applicability and offer a compelling new approach for uncertainty quantification in a wide range of data analysis tasks. Future research will focus on further refining this method, focusing on efficient estimates in sparse regimes. We also plan on assessing the performance arising from other topology-based formulations, which are geared towards specific modalities such as meshes (Turner et al., 2014) or time series (Zeng et al., 2021).

[3]Note that alternatively, a greyscale image can be interpreted as an undirected graph, by defining the set of vertices to be the set of pixels, and adding edges between neighbouring pixels if both of their values are non-zero.

REFERENCES

Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18(8): 1–35, 2017.

Christophe Andrieu and Gareth O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697 – 725, 2009. doi: 10.1214/07-AOS574. URL https://doi.org/10.1214/07-AOS574.

Mark A Beaumont, Wenyang Zhang, and David J Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.

Paul Bendich, J. S. Marron, Ezra Miller, Alex Pieloch, and Sean Skwerer. Persistent homology analysis of brain artery trees. *Annals of Applied Statistics*, 10(1):198–218, 2016.

Espen Bernton, Pierre Jacob, Mathieu Gerber, and Christian Robert. Approximate bayesian computation with the wasserstein distance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(2):235–269, 2019.

Pier Giovanni Bissiri, Chris C Holmes, and Stephen G Walker. A general framework for updating belief distributions. *Journal of the royal statistical society. series b, statistical methodology*, 78(5): 1103, 2016.

Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Theo Lacombe, Martin Royer, and Yuhei Umeda. PersLay: A neural network layer for persistence diagrams and new graph topological signatures. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2786–2796. PMLR, 2020.

Badr-Eddine Chérief-Abdellatif and Pierre Alquier. Mmd-bayes: Robust bayesian estimation via maximum mean discrepancy. In *Symposium on Advances in Approximate Bayesian Inference*, pp. 1–21. PMLR, 2020.

David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, January 2007. doi: 10.1007/s00454-006-1276-5.

Imre Csiszár. I-divergence geometry of probability distributions and minimization problems. *The annals of probability*, pp. 146–158, 1975.

Bruno De Finetti, Henry E Kyburg, and Howard E Smokler. Foresight: Its logical laws, its subjective sources. *Breakthroughs in statistics*, 1:134–174, 1937.

Monroe D Donsker and SR Srinivasa Varadhan. Asymptotic evaluation of certain markov process expectations for large time, i. *Communications on Pure and Applied Mathematics*, 28(1):1–47, 1975.

Arnaud Doucet, Adam M Johansen, et al. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering*, 12(656-704):3, 2009.

Herbert Edelsbrunner and John Harer. *Computational topology: An introduction*. American Mathematical Society, Providence, RI, USA, 2010.

Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. POT: Python optimal transport. *Journal of Machine Learning Research*, 22(78): 1–8, 2021.

David T Frazier, Gael M Martin, Christian P Robert, and Judith Rousseau. Asymptotic properties of approximate bayesian computation. *Biometrika*, 105(3):593–607, 2018.

David T Frazier, Christian Robert, and Judith Rousseau. Model misspecification in abc: Consequences and diagnostics. *Journal of the Royal Statistical Society: Series B*, 2019.

Felix Hensel, Michael Moor, and Bastian Rieck. A survey of topological machine learning methods. *Frontiers in Artificial Intelligence*, 4, 2021. ISSN 2624-8212. doi: 10.3389/frai.2021.681108.

Max Horn, Edward De Brouwer, Michael Moor, Yves Moreau, Bastian Rieck, and Karsten Borgwardt. Topological graph neural networks. In *International Conference on Learning Representations (ICLR)*, 2022.

Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Jeremias Knoblauch, Jack Jewson, and Theodoros Damoulas. An optimization-centric view on bayes' rule: Reviewing and generalizing variational inference. *Journal of Machine Learning Research*, 23 (132):1–109, 2022.

Timm Krüger, Halim Kusumaatmaja, Alexandr Kuzmin, Orest Shardt, Goncalo Silva, and Erlend Magnus Viggen. The lattice boltzmann method. *Springer International Publishing*, 10(978-3): 4–15, 2017.

Roland Kwitt, Stefan Huber, Marc Niethammer, Weili Lin, and Ulrich Bauer. Statistical topological data analysis - A kernel perspective. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 28*, pp. 3070–3078. Curran Associates, Inc., 2015.

Vasileios Maroulas, Cassie Putman Micucci, and Farzana Nasrin. Bayesian topological learning for classifying the structure of biological networks. *Bayesian Analysis*, 17(3):711–736, 2022. doi: 10.1214/21-BA1270.

Jeffrey W Miller and David B Dunson. Robust bayesian inference via coarsening. *Journal of the American Statistical Association*, 2018.

Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning (ICML)*, number 119 in Proceedings of Machine Learning Research, pp. 7045–7054. PMLR, 2020.

Christopher A. Oballe. *Bayesian Topological Machine Learning*. PhD thesis, University of Tennessee, Knoxville, 2020.

Art B. Owen. *Monte Carlo theory, methods and examples*. 2013. URL https://artowen.su.domains/mc/.

Mijung Park, Wittawat Jitkrittum, and Dino Sejdinovic. K2-abc: approximate bayesian computation with kernel embeddings. *arXiv preprint arXiv:1502.02558*, 2015.

Jonathan K Pritchard, Mark T Seielstad, Anna Perez-Lezaun, and Marcus W Feldman. Population growth of human y chromosomes: a study of y chromosome microsatellites. *Molecular biology and evolution*, 16(12):1791–1798, 1999.

Jan Reininghaus, Stefan Huber, Ulrich Bauer, and Roland Kwitt. A stable multi-scale kernel for topological machine learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4741–4748, Red Hook, NY, USA, June 2015. Curran Associates, Inc. doi: 10.1109/CVPR.2015.7299106.

Bastian Rieck, Filip Sadlo, and Heike Leitte. Topological machine learning with persistence indicator functions. In Hamish Carr, Issei Fujishiro, Filip Sadlo, and Shigeo Takahashi (eds.), *Topological Methods in Data Analysis and Visualization V*, pp. 87–101. Springer, Cham, Switzerland, 2020a. ISBN 978-3-030-43036-8. doi: 10.1007/978-3-030-43036-8_6.

Bastian Rieck, Tristan Yates, Christian Bock, Karsten Borgwardt, Guy Wolf, Nick Turk-Browne, and Smita Krishnaswamy. Uncovering the topology of time-varying fMRI data using cubical persistence. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6900–6912. Curran Associates, Inc., 2020b.

Sebastian M Schmon, Patrick W Cannon, and Jeremias Knoblauch. Generalized posteriors in approximate bayesian computation. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2021. URL https://openreview.net/forum?id=tKrg5DAyeWq.

Primoz Skraba and Katharine Turner. Wasserstein stability for persistence diagrams, 2022.

Simon Tavaré, David J Balding, Robert C Griffiths, and Peter Donnelly. Inferring coalescence times from dna sequence data. *Genetics*, 145(2):505–518, 1997.

Thomas Thorne, Paul D W Kirk, and Heather A Harrington. Topological approximate bayesian computation for parameter inference of an angiogenesis model. *Bioinformatics*, 38(9):2529–2535, 2022. doi: 10.1093/bioinformatics/btac118.

Chad M Topaz, Lori Ziegelmeier, and Tom Halverson. Topological data analysis of biological aggregation models. *PloS one*, 10(5):e0126383, 2015.

Katharine Turner, Sayan Mukherjee, and Doug M. Boyer. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4):310–344, 2014. doi: 10.1093/imaiai/iau011.

Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical review letters*, 75(6):1226, 1995.

Hubert Wagner, Chao Chen, and Erald Vuçini. Efficient computation of persistent homology for cubical data. In Ronald Peikert, Helwig Hauser, Hamish Carr, and Raphael Fuchs (eds.), *Topological Methods in Data Analysis and Visualization II: Theory, Algorithms, and Applications*, pp. 91–106. Springer, Heidelberg, Germany, 2012. doi: 10.1007/978-3-642-23175-9_7.

Dominik J. E. Waibel, Scott Atwell, Matthias Meier, Carsten Marr, and Bastian Rieck. Capturing shape information with multi-scale topological loss terms for 3d reconstruction. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li (eds.), *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pp. 150–159, Cham, Switzerland, 2022. Springer. doi: 10.1007/978-3-031-16440-8_15.

Zuoyu Yan, Tengfei Ma, Liangcai Gao, Zhi Tang, Yusu Wang, and Chao Chen. Neural approximation of graph topological features. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 33357–33370. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/d7ce06e9293c3d8e6cb3f80b4157f875-Paper-Conference.pdf.

Arnold Zellner. Optimal information processing and bayes's theorem. *The American Statistician*, 42 (4):278–280, 1988.

Sebastian Zeng, Florian Graf, Christoph Hofer, and Roland Kwitt. Topological attention for time series forecasting. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 24871–24882. Curran Associates, Inc., 2021.

Qi Zhao, Ze Ye, Chao Chen, and Yusu Wang. Persistence enhanced graph neural network. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2896–2906. PMLR, 2020.

Jie Zhou, Daniel L Civco, and John A Silander. A wavelet transform method to merge landsat tm and spot panchromatic data. *International journal of remote sensing*, 19(4):743–757, 1998.

A ADDITIONAL RESULTS

In the interest of reproducibility and to improve the understanding of our method, we also provide more detailed insights into our experiments.

Synthetic data parameter estimates. As a consistency check and illustrative example, we use our proposed methods for inferring the radius parameters of two objects, namely the 2-sphere and the 2-torus. The 2-sphere is the surface of a closed 3-dimensional ball of radius r, while the 2-torus is constructed using an *inner radius* r and an *outer radius* R. The point clouds were sampled uniformly from the respective surface of known parameter.

Table 3: Estimating the radius parameters of a 2D sphere (r) and a 2D torus (r, R). For the more complex torus parameter estimation, we observe that our topological loss ℓ_T provides more reliable estimates than the geometrical loss ℓ_G. The results show the mean of five runs for each parameter choice.

		Sphere radius r			Torus radii (r, R)					
Loss	*Sampling method*	1.00	5.00	10.00	$(1, 2)$		$(3, 5)$		$(5, 10)$	
ℓ_G	Importance sampling	0.97	4.94	9.91	0.94	1.96	2.99	4.95	4.98	9.98
	MCMC	1.00	4.99	10.00	0.99	2.04	2.77	4.96	5.01	10.00
ℓ_T	Importance sampling	1.01	5.02	9.98	1.02	1.96	3.01	5.05	4.95	10.04
	MCMC	0.99	4.99	10.00	1.04	1.89	2.99	4.97	5.01	10.04

Table 3 depicts the results of the inference process. We observe that both types of losses are capable of inferring the right parameters of these simple geometric objects, with our topology-based loss ℓ_T providing slightly more reliable estimates in the case of a torus. Table 3 shows the raw values of learning radius parameters for synthetic data sets (2D spheres and 2D tori). This experiment primarily shows that for such simple shapes, geometry-based and topology-based losses perform equivalently. For more complex data sets, however, we find that the improved robustness of our topology-based loss term helps in inferring the ground truth parameters.

Swarm behaviour of the Vicsek Model. Fig. 5 shows simulations of the Vicsek model for different noise parameters η. The higher η, the more 'chaotic' the behaviour of the resulting complex system.

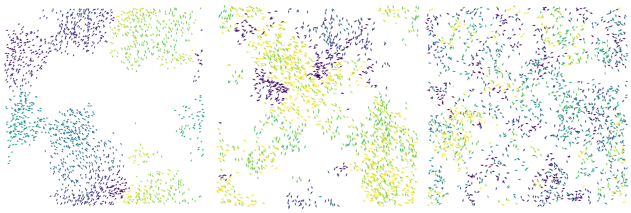

Figure 5: Left to right: simulations of the Vicsek model for $\eta = 0.15, 0.3$ and 0.6, respectively. For lower values of η swarming behaviour emerges, whereas for higher values of η swarms tend to merge and the model becomes more ergodic.

Vicsek model parameter estimates. Accompanying Fig. 3, we show the raw parameter estimates of η, the noise parameter of the model in Table 4. Moreover, we show the results of a *rejection sampling* procedure with summary statistics (mean and standard deviation, respectively), in Fig. 6. In the latter setting we used the L_2 distance between the respective summary statistics as the loss that determines the results of our method. We observe that there is no statistic that clearly discriminates the other one, and that the performance of the results highly depends on the parameter and timestep which is considered. We find that the overall accuracy and reliability of the respective experiments by using topological losses is significantly higher; see Fig. 3 for a comparison.

Table 4: Noise parameter estimation after $n = \{5, 10, 50\}$ iterations of the Vicsek model. Estimates of η based on a topological loss are always closer to the ground truth value than estimates obtained via a geometrical loss.

			True parameter η		
	Loss	*Sampling method*	0.15	0.30	0.60
$n = 5$	Hausdorff distance	Importance sampling	0.17	0.35	0.67
		MCMC	0.13	0.36	1.15
	Topological distance	Importance sampling	0.14	0.29	0.59
		MCMC	0.12	0.33	0.63
$n = 10$	Hausdorff distance	Importance sampling	0.11	0.39	0.70
		MCMC	0.18	0.41	1.95
	Topological distance	Importance sampling	0.18	0.27	0.65
		MCMC	0.13	0.34	0.62
$n = 50$	Hausdorff distance	Importance sampling	0.12	0.29	0.72
		MCMC	0.28	1.40	2.28
	Topological distance	Importance sampling	0.16	0.31	0.61
		MCMC estimate	0.21	0.24	0.59

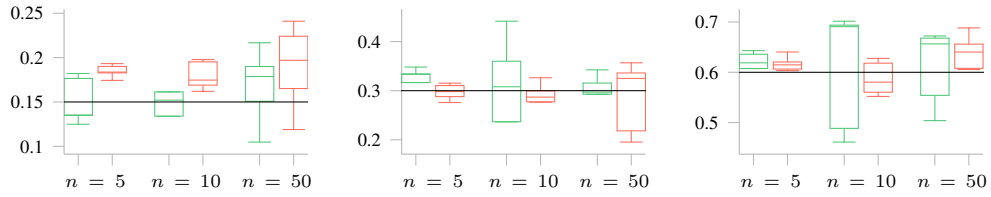

(a) Rejection sampling

Figure 6: Summary of parameter estimates for the noise parameter η of the Vicsek model with respect to summary statistics and a standard rejection sampling procedure, using 5 repetitions of the inference procedure. Each subplot depicts estimates for $\eta \in \{0.15, 0.30, 0.60\}$, respectively. The true parameter is shown as a solid black line. Estimates based on the mean statistic L_2 loss and standard deviation statistic L_2 loss both are significantly less accurate than the respective topological loss, see Fig. 3b.

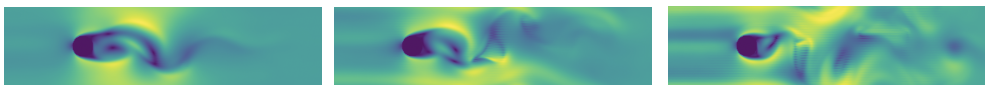

(a) low turbulence (b) medium turbulence (c) high turbulence

Figure 7: Left to right: simulations of the fluid model after 3000 iterations, for $\eta = 0.2, 0.3$ and 0.4, respectively. For lower values of η the fluid flow is less turbulent than for higher values of η.

Lattice Boltzmann model for fluid dynamics. The degree of randomness in the local propagation and collision steps in the lattice Boltzmann method (LBM) model controls the turbulent behaviour of the global system, as is illustrated in Fig. 7. This is the parameter η that we estimate in our experiments. The results are shown in Fig. 8 and Fig. 9, respectively. Using importance sampling, the topological loss outperforms the geometrical loss in terms of accuracy, as can be seen from the median and interquartile range of the boxplots in Fig. 8. Moreover, the topological loss in combination with importance sampling also outperforms the geometrical loss with MCMC sampling, as evidenced from Table 1. The topological loss in combination with MCMC sampling, however, is far off in terms of accuracy. The latter indicates that the convergence rates of MCMC sampling and importance sampling may be very different, and depend on the given setting. Finally, we note that rejection

sampling with both the mean and standard deviation summary statistic does not perform reliably, and its accuracy highly depends on the true parameter in the underlying observation.

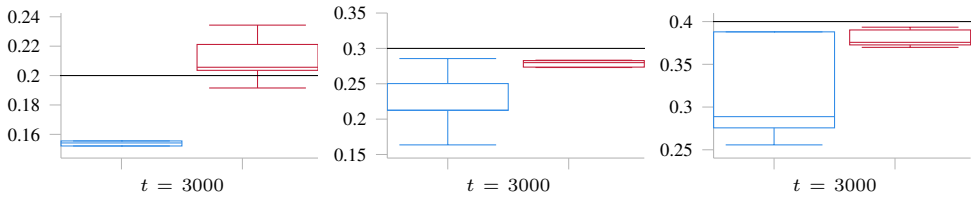

(a) Importance sampling

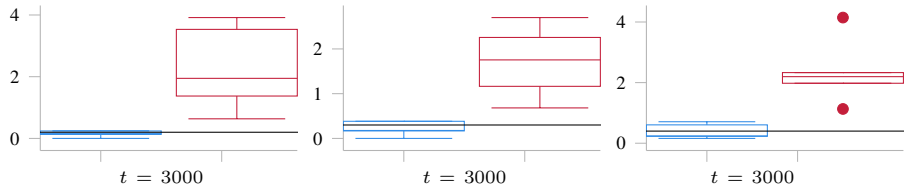

(b) MCMC sampling

Figure 8: Summary of parameter estimates for the parameter η of the fluid model, using 5 repetitions of the inference procedure. Each subplot depicts estimates for $\eta \in \{0.20, 0.30, 0.40\}$, respectively. The true parameter is shown as a solid black line. Estimates based on the geometrical loss ℓ_G are less accurate than the topological loss ℓ_T.

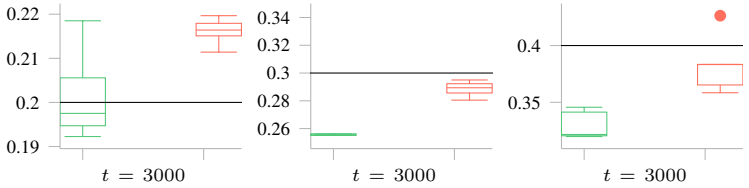

Rejection sampling

Figure 9: Summary of parameter estimates for the parameter η of the fluid dynamics model with respect to summary statistics and a standard rejection sampling procedure, using 5 repetitions of the inference procedure. Each subplot depicts estimates for $\eta \in \{0.20, 0.30, 0.40\}$, respectively. The true parameter is shown as a solid black line. Estimates based on the mean statistic L_2 loss and on the standard deviation statistic L_2 loss tend to be less reliable than estimates based on the topological loss, see Fig. 8.

Multi-scale structures in percolation models. Fig. 10 shows an illustration of our proposed percolation model. The higher the value of p, the more likely it is for a pixel to be non-zero. Consequently, there are more non-zero pixels for higher values of p. As discussed in the main paper, the topological distance outperforms the SCC distance significantly, with both importance sampling and MCMC sampling. However, the rejection sampling procedure together with the mean statistic performs even more accurate in this experiment, as can be seen by comparing Fig. 11 and Fig. 12. This is not surprising: once the pixel size converges to infinity, the mean of the pixel values will converge to the expected pixel value, where the latter only depends on p (since the maximum greyscale value is fixed in our experiments, and therefore does not have any discriminative power for inference). Therefore, at least in a large-pixel regime we can infer the true parameter of the observation from the mean of its pixel values. Note that this is due to the additional structure in this experiment, and does not hold for complex systems that cannot be described in such a simple way, as has been seen in the previous experiments. Finally, rejection sampling with the standard deviation statistic leads to highly erroneous results, as can also be seen from Fig. 12. For complex systems it is therefore very difficult to determine the 'right' summary statistic which contains a sufficient amount of information to infer the underlying parameter. This choice is not necessary when using topological losses, which makes the latter an appropriate generic choice, in many applications.

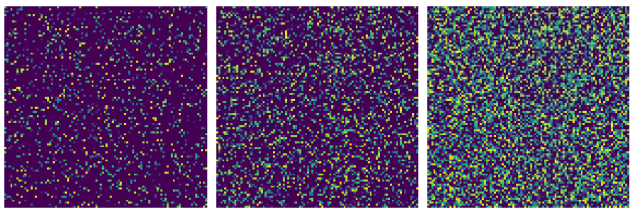

Figure 10: Left to right: samples of the proposed percolation model for $p = 0.15, 0.3$ and 0.6, respectively. All samples admit a fixed maximum greyscale value of 50. Although the geometric distance between two samples of fixed p can be large due to the uniform sampling of pixels, the overall topological structure (which is captured by persistent homology) for a given p is more stable with respect to sampling.

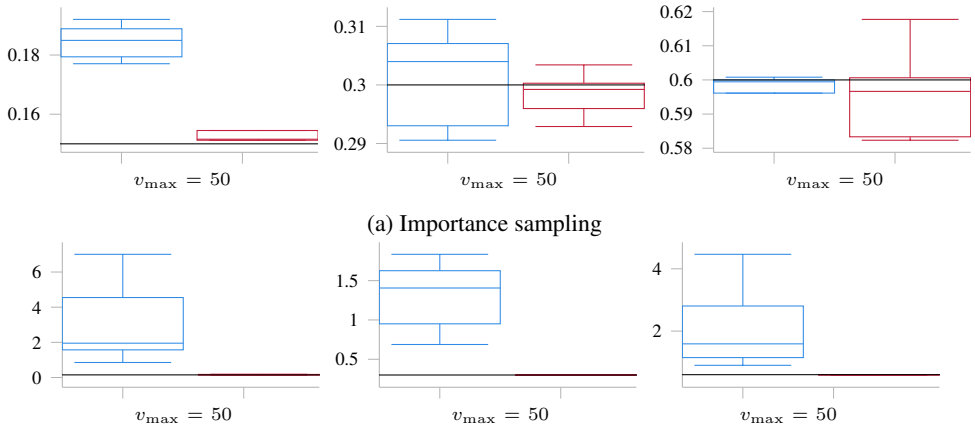

(a) Importance sampling

(b) MCMC sampling

Figure 11: Summary of parameter estimates for the parameter p of the Percolation model, using 5 repetitions of the inference procedure. Each subplot depicts estimates for $p \in \{0.15, 0.30, 0.60\}$, respectively. The true parameter is shown as a solid black line. Estimates based on the geometrical loss ℓ_G are less accurate than the topological loss ℓ_T, in particular for larger number of iterations.

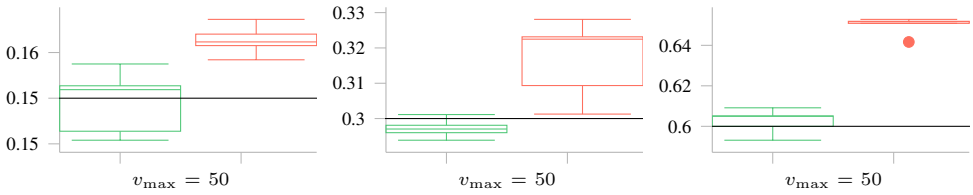

(a) Rejection sampling

Figure 12: Summary of parameter estimates for the parameter p of the Percolation model with respect to summary statistics and a standard rejection sampling procedure, using 5 repetitions of the inference procedure. Each subplot depicts estimates for $p \in \{0.15, 0.30, 0.60\}$, respectively. The true parameter is shown as a solid black line. Estimates based on the mean statistic L_2 loss are more accurate than the standard deviation statistic L_2 loss.

B FURTHER DETAILS AND PROOF

Algorithm 2 MCMC approximation of the target

Require: Observed data y, start value θ_0, proposal distribution $q(\vartheta \,|\, \theta)$

Sample $X_0 \sim p(x \,|\, \theta_0)$
for $i = 1 : K$ **do**
  Sample $U \sim \mathsf{Unif}[0, 1]$
  Sample $\theta' \sim q(\cdot \,|\, \theta_{i-1})$
  Sample $X' \sim p(x \,|\, \theta')$
  Compute

$$a(\theta_{i-1}, \theta') = \frac{e^{-\ell(y,X')} p(\theta')}{e^{-\ell(y,X_{i-1})} p(\theta_{i-1})} \frac{q(\theta_{i-1} \,|\, \theta')}{q(\theta' \,|\, \theta_{i-1})}$$

  if $U \leq a(\theta_{i-1}, \theta')$ **then**
    Set $\theta_i = \theta', X_i = X'$
  else
    Set $\theta_i = \theta_{i-1}, X_i = X_{i-1}$
  end if
end for
return

$$\boldsymbol{\theta} = (\theta_{1:K}).$$

Proof about the comparison-based posterior

Proposition B.1. *The comparison-based posterior*

$$\pi(\theta, x \,|\, y) = \frac{\exp\left(-\ell(y, x)\right) p(x \,|\, \theta) p(\theta)}{\int_{\mathsf{X}} \int_{\Theta} \exp\left(-\ell(y, x)\right) p(x \,|\, \theta) p(\theta) \mathrm{d}\theta \mathrm{d}x}$$

is the solution $q^*(\theta, x) = \pi(\theta, x \,|\, y)$ *to the optimisation problem*

$$q^* = \mathrm{argmin}_{q \in \mathcal{P}(\Theta \times \mathsf{X})} \left\{ \mathbb{E}_q[\ell(y, x)] + \mathrm{KL}\left(q(\theta, x), p(x|\theta)\pi(\theta)\right) \right\}.$$

Proof.

$$
\begin{aligned}
q_B &= \operatorname{argmin}_q \left\{ \mathbb{E}_q \left[\ell(y,x) \right] + KL\left(q(x,\theta), p(x|\theta)\pi(\theta) \right) \right\} \\
&= \operatorname{argmin}_q \left\{ \iint \ell(y,x) q(x,\theta) \mathrm{d}x\mathrm{d}\theta + \iint \log \frac{q(x,\theta)}{p(x|\theta)\pi(\theta)} q(x,\theta) \mathrm{d}x\mathrm{d}\theta \right\} \\
&= \operatorname{argmin}_q \left\{ \iint \log\left(\exp \ell(y,x) \right) q(x,\theta) \mathrm{d}x\mathrm{d}\theta + \iint \log \frac{q(x,\theta)}{p(x|\theta)\pi(\theta)} q(x,\theta) \mathrm{d}x\mathrm{d}\theta \right\} \\
&= \operatorname{argmin}_q \left\{ \iint \log \left(\frac{q(x,\theta)}{p(x|\theta)\pi(\theta)\exp\left(-\ell(y,x)\right)} \right) q(x,\theta) \mathrm{d}x\mathrm{d}\theta \right\} \\
&= \operatorname{argmin}_q \left\{ \iint \log \left(\frac{q(x,\theta)}{p(x|\theta)\pi(\theta)\exp\left(-\ell(y,x)\right) Z^{-1}} \right) q(x,\theta) \mathrm{d}x\mathrm{d}\theta \right\} - \log Z \\
&= \operatorname{argmin}_q KL\left\{ q(x,\theta), p(x|\theta)\pi(\theta)\exp\left(-\ell(y,x)\right) Z^{-1} \right\}.
\end{aligned}
$$

Hence, the minimiser q is such that

$$
q(x,\theta) = \frac{p(x \mid \theta)\pi(\theta)\exp\left(-\ell(y,x)\right)}{Z}, \quad q(\theta) = \int q(x,\theta)\mathrm{d}x
$$

and

$$
Z := \iint \exp\left(-\ell(y,x)\right) p(x \mid \theta)\pi(\theta)\mathrm{d}x\mathrm{d}\theta.
$$

□