

Topological Trajectory Prediction with Homotopy Classes

Jennifer Wakulicz, Ki Myung Brian Lee, Teresa Vidal-Calleja, Robert Fitch

Abstract—Trajectory prediction in a cluttered environment is key to many important robotics tasks such as autonomous navigation. However, there are an infinite number of possible trajectories to consider. To simplify the space of trajectories under consideration, we utilise *homotopy classes* to partition the space into countably many mathematically equivalent classes. All members within a class demonstrate identical high-level motion with respect to the environment, i.e., travelling above or below an obstacle. This allows high-level prediction of a trajectory in terms of a sparse label identifying its homotopy class. We therefore present a light-weight learning framework based on variable-order Markov processes to learn and predict homotopy classes and thus high-level agent motion. By informing a Gaussian mixture model (GMM) with our homotopy class predictions, we see great improvements in low-level trajectory prediction compared to a naive GMM on a real dataset.

I. INTRODUCTION

Understanding the dynamics of mobile agents is important for robots to operate autonomously in environments with pedestrians, vehicles, animals, or other moving entities. Autonomous driving, service robot navigation through crowds, and similar applications all rely on trajectory prediction to help avoid dynamic obstacles. Predicting trajectories is also an important form of *intention inference*, where the intention to move in a certain direction or to a given location is inferred [1–4]. Examples involving pedestrian trajectories include goal-directed inference [1, 2], modelling social interactions [3], and industrial applications with cobots [5, 6].

We are interested in leveraging environmental context to predict the abstract motion of an agent. Obstacles in the environment, for example, place constraints on the trajectories we wish to predict. The environment, therefore, is a rich source of information for trajectory prediction that can be seen to split the set of all possible trajectories into equivalence classes. These equivalence classes summarise the agent’s high-level motion in a sparse identifier, sometimes referred to as an *h-signature*, that specifies properties such as whether the trajectory passes above or below a certain obstacle, and how many times it does so. Having this identifier is powerful, as it gives a general idea of the motion in a sparse representation that can be used as the basis for developing efficient computational frameworks. For example, *h-signatures* have been used in various path planning contexts [7–10].

In trajectory prediction, topological concepts like *h-signatures* allow the extraction of salient, predictive features of trajectories that arise from obstacles or other features of the environment. Existing trajectory prediction methods, however, typically are based on a geometric representation (e.g., [4, 11]) and can be computationally prohibitive for long-term prediction. These methods operate on full trajectories and consider all differences between them, including minor differences. The number of small and insignif-




Fig. 1. Illustration of the problem considered in this paper. The task is to predict high-level motions then use them to inform lower level trajectory predictions.

icant variations, even between topologically equivalent trajectories, can easily become computationally overwhelming. These geometric algorithms thus trade-off prediction horizon with computational efficiency [2, 4, 11, 12].

In this paper, we propose a new topological approach to trajectory prediction by leveraging the salient features of paths summarised as by *h-signatures*. *h-signatures* are both compact in size and provide global information about a trajectory as they only encode how a trajectory moves past an obstacle as opposed to an overall trajectory. The use of *h-signatures* in trajectory prediction can therefore circumvent the aforementioned trade-off that other methods face.

We thus present our method for topology-informed trajectory prediction, which introduces the notion of *partial h-signatures*, defined for incomplete trajectories. The concept of partial *h-signatures* decomposes trajectory prediction into high-level and low-level prediction problems. As initial solutions to these problems, we present a high-level prediction algorithm using a variable-order Markov Process (VOMP) [13] and a low-level prediction algorithm based on a hierarchical Gaussian mixture model (GMM). We demonstrate the behaviour of our method with a simple synthetic dataset, and then extensively evaluate its performance compared to a baseline GMM method without topological knowledge using the ATC shopping mall dataset [14]. Results show that our method performs as well or better in terms of geometric error, and considerably better in identifying the correct high-level motion. This indicates that the *h-signature* is a meaningful high-level representation for predicting agent trajectories, thus useful for applications such as target search [15].

II. RELATED WORK

Trajectory prediction is a well-studied problem in robotics. Model-based approaches use an estimated or learned dynamic model of an agent’s motion to predict an agent’s state forward in time [15, 16]. Model-free learning based methods can be used where dynamics are cumbersome to model. For example, by treating trajectory prediction as a sequence generation problem, recurring neural networks (RNNs) and long-short term memory networks (LSTMs)

have been implemented in various contexts with large success [12, 17–19]. Other approaches have introduced mixture models such as GMMs and the kernel trajectory map [11] to capture the multi-modal nature of an agent’s possible future paths, showing greater prediction accuracy over single-mode approaches [20–22].

Beyond removing the burden of acquiring a dynamic model, the power of model-free learning methods lie in their ability to incorporate contextual information into the model. For example, social interactions between agents have been captured in various architectures to improve pedestrian trajectory prediction in crowded environments [3, 21, 23, 24]. Where crowds are extremely dense and interactions are complex, bulk properties of crowds have been learned instead [25, 26]. Even the intent of an agent may be learned, and has been proposed as a powerful contextual cue for long-term trajectory prediction problems [1, 3, 4, 27].

Contextual information common and relevant to all trajectory prediction scenarios is the environment within which the agent is travelling. Obstacles, unreachable regions and road structures – the *topological features* of an environment – all dictate how an agent can move through an environment. It seems then that topological information should be a fruitful contextual cue for trajectory prediction. Indeed, topology-aware learning methods have been proposed in the past [28, 29]. While these previous works use topological information to inform local motion models used in prediction, our paper leverages this contextual cue to develop a model-free prediction framework.

III. BACKGROUND AND PROBLEM FORMULATION

A. Homotopy Theory in Robotics

Two paths τ_1, τ_2 in a topological space \mathcal{D} with common start and ending points are *homotopic* if there exists a continuous transformation or deformation from one to the other [30]. Sets of paths homotopic to one another are named *homotopy classes*. Non-trivial homotopy classes arise as a result of obstacles in the space, as deforming some paths into others would require moving through an obstacle and breaking the continuity requirement. This notion is depicted in Fig. 2. Paths τ_1 and τ_2 are homotopic as they can be continuously deformed into each other, indicated by the dotted paths between them. These paths are however not homotopic to τ_3 , as any deformation into τ_3 would require moving through the obstacle \mathcal{O} .

In the context of robotics, obstacles in the environment split the space of an agent’s possible trajectories between points from a single homotopy class into a countable number of homotopy classes. The obstacles therefore dictate the number of unique ways an agent may travel through the space from one point to another. For example, in Fig. 2 (a) the obstacle \mathcal{O} splits the space of trajectories into those that move ‘above’ the obstacle, those that move ‘below’, and those that wind around the obstacle any number of times before moving to the end point. Such abstraction of the high-level motions available to an agent is a powerful tool for navigation, prediction and tracking tasks often encountered in robotics [7, 8, 10, 28, 29].




Fig. 2. Illustrations of homotopy classes and h -signature calculation.

B. h -signatures as Homotopy Invariants

To identify which homotopy class a trajectory τ belongs to, one must compute a *homotopy invariant* – a unique identifier $h(\tau)$ of a trajectory’s homotopy class such that $h(\tau_1) = h(\tau_2)$ if and only if τ_1 and τ_2 are homotopic. A commonly used homotopy invariant is detailed in [30] and coined the *h -signature* of a trajectory in [9]. To compute the h -signature of a path, non-intersecting rays are drawn from within each obstacle to the boundary of the environment. Then a ‘word’ is constructed by following the path and appending letter ‘ n ’ if the path crosses the ray corresponding to the n -th obstacle from left to right, and ‘ $-n$ ’ if τ crosses from right to left. The construction of $h(\gamma)$ is given in Fig. 2 (b) to illustrate. The final word may then be *reduced* by cancelling all consecutive appearances of n and $-n$. For example, $h(\gamma)$ in Fig. 2 (b) reduces from $(1, 2, -2)$ to (1) .

C. Problem Formulation

Consider an agent traversing through a planar environment \mathcal{D} containing n obstacles $\mathcal{O} = \{\mathcal{O}_1, \dots, \mathcal{O}_n\}$. The agent’s trajectory is denoted $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$. We assume that the agent begins at a start location on the boundary $\delta\mathcal{D}$ of the environment and travels to an end location on $\delta\mathcal{D}$. We are given a dataset of K historical, fully observed trajectories $\mathcal{X} = \{X_{\text{obs}}^1, \dots, X_{\text{obs}}^K\}$, with full knowledge of obstacles. Further, in the online setting, we have a partial noisy measurement of trajectory $\mathbf{Y}_{\text{obs}} = \{\mathbf{y}_1, \dots, \mathbf{y}_{T_{\text{obs}}}\}$ up to time T_{obs} , with a known sensor model $P(\mathbf{y}_t | \mathbf{x}_t)$.

We are then interested in predicting the agent’s high-level motion through the environment. An agent’s ‘high-level motion’ can be abstracted from its trajectory in many ways. For example, one might cluster trajectories with a GMM and treat each cluster as a unique high-level motion. We propose that the best abstraction is given by a trajectory’s homotopy class and thus that prediction of a trajectory’s h -signature is the most appropriate solution. Formally, the problem is:

Problem 1. (High-level prediction) Given the partial, noisy measurements of the trajectory \mathbf{Y}_{obs} , predict the h -signature \hat{h} of the robot’s future full trajectory.

While the solution to Problem 1 can be used as a tool in many robotics problems, here we are interested in how the h -signature may be used to produce a topology-informed low-level prediction for a trajectory:

Problem 2. (Low-level prediction) Given the predicted h -signature \hat{h} associated with an agent’s partially observed measurements \mathbf{Y}_{obs} , predict the full trajectory \mathbf{X} .

IV. TOPOLOGY-INFORMED TRAJECTORY PREDICTION

A. Overview

Problems 1 and 2 are challenging to solve directly without modification, as illustrated in the factor graphs shown in Fig. 3. In a pedantic Bayesian formulation (Fig. 3 (a)), one would first predict the underlying trajectory for the entire duration, and subsequently predict the corresponding h -signature. This is because the measurements are conditionally independent of the h -signature given the trajectory, and the only given relationship between the trajectory and the h -signature is the computation process outlined in Sec. III-B. In other words, in this view, low-level trajectory prediction precedes high-level prediction, limiting its effectiveness.

Instead, we propose to circumvent the low-level prediction by introducing the notion of partial h -signature (green lines, Fig. 3 (b)). Unlike the usual h -signature, partial h -signatures can be obtained from an incomplete trajectory, as we detail in Sec. IV-B. Given the partial h -signature, we predict the full h -signature. This is achieved with a VOMP model (red line, Fig. 3 (b)) trained on a dataset (Sec. IV-C). The full h -signature can then be used to predict the low-level trajectory, with a model learnt using the historical dataset. As we readily have access to trajectories and their associated h -signature, this can be as simple as learning a mixture of experts for low-level trajectory prediction within each homotopy class. To this end, we demonstrate the use of a hierarchical GMM in Sec. IV-E.

B. Partial h -signatures


As described in Sec. III-A, homotopy classes exist only for paths between two fixed start and end points. In the context of this paper trajectories are assumed to start and end on different boundary points of the environment. Similar to [10], we ensure our description of homotopy classes is valid by applying a quotient map, mapping all boundary points to a single quotient point while preserving the topology of the space. Then, the set of all homotopy classes here is over paths between the quotient point and itself.

With this in mind, the notion of a *partial h -signature* is ill-defined in a topological sense. Much like a full h -signature, it is calculated by constructing a ‘word’ according to a path’s ray crossings. However, it is extracted from an incomplete trajectory, i.e. one that has not yet returned to the quotient point. It is thus crucial to note that a partial h -signature is not an identifier of a homotopy class but rather a *predictor*.


To predict full h -signatures from partial ones, a notion of compatibility between the two is needed. For a given partial h -signature p , compatible full h -signatures are those whose prefix is p . In other words, the set of all h -signatures compatible with p is $\mathcal{H}(p) = \{h \mid \exists p', h = pp'\}$.

C. High-level Prediction using Topological Variable Order Markov Processes

As h -signatures in an environment with n obstacles are simply ‘words’ constructed from an alphabet $\mathcal{A} = \{1, \dots, n, -1, \dots, -n\}$, prediction of a full h -signature given partial h -signature can be viewed as a sequence completion



(a) Pedantic model.



(b) Proposed model.

Fig. 3. Probabilistic graphical models depicting the problem. Directed arrows from A to B imply availability of a model of B given A . Dashed arrows are ignored during inference.

problem. The sparse nature of the h -signature as a representation of high-level motion allows for relatively simple techniques to be used for sequence generation. We propose a VOMP [13] for this purpose. Like LSTMs, VOMPs can learn dependencies in data that are of varying length. However, VOMPs are capable of producing probabilistic predictions where LSTMs are not. This is crucial for our approach, as a probability distribution over full h -signatures is required to produce probabilistic low-level predictions (blue lines, Fig. 3 (b)).

We represent a VOMP with a probabilistic suffix automaton (PSA) as in [13]. Here, a PSA state is a partial h -signature p of up to some maximum length $L > 0$ constructed from alphabet \mathcal{A} as $p = a_1 \dots a_l$, $0 \leq l \leq L$. Transitions between two states p and p' are allowed only if there exists some $a \in \mathcal{A}$ such that p' is a *suffix* of ap . Allowed transitions have associated with them a probability that the transition will occur. Transition probabilities learned offline can be used to produce probability distributions over future states online. An example PSA is drawn in Fig. 4 (b).

We follow [13] closely to train a VOMP over h -signatures offline, and adapt their online prediction process to better suit the prediction of h -signatures. Our algorithms for online prediction and offline learning are detailed below.

1) *Online Prediction*: The trained VOMP outputs probabilities of arbitrarily long h -signatures. To find the probability of any h -signature $P(h)$ one simply takes a corresponding walk through the PSA, multiplying transition probabilities from state to state. However, we are interested in finding the conditional distribution $P(h \mid p)$ over possible full h -signatures h given a partial h -signature p extracted from a partially observed trajectory. Letting $\mathcal{H}(p)$ be the set of all full h -signatures compatible with p , the conditional probability $P(h \mid p)$ of any $h \in \mathcal{H}(p)$ can be calculated in typical fashion,

$$P(h \mid p) = \frac{P(h)}{\sum_{h \in \mathcal{H}(p)} P(h)}. \quad (1)$$

Algorithm 1 PSA offline learning

Inputs: $\epsilon, L, \mathcal{A}, h$ -signature data

Output: Trained PSA

- 1: initialise tree T
 - 2: $\mathcal{P} \leftarrow \{a \mid a \in \mathcal{A}, P(a) \geq \epsilon\}$
 - 3: **while** \mathcal{P} not empty **do**:
 - 4: $p \leftarrow \mathcal{P}.\text{pop}()$
 - 5: **if** $\mathcal{E}(p, \text{suffix}(p)) \geq \epsilon$ **then**:
 - 6: add path to p to T
 - 7: **if** $|p| \leq L$ **then**:
 - 8: $\mathcal{P} \leftarrow \mathcal{P} \cup \{ap \mid a \in \mathcal{A}, P(ap) \geq \epsilon\}$
 - 9: **for** all leaves r in T **do**:
 - 10: **if** longest prefix(r) not in T **then**:
 - 11: add path to r to T
 - 12: PSA \leftarrow leaves of T
-

However, the set $\mathcal{H}(p)$ of compatible h -signatures is of infinite size. Just as agents may walk paths of varying lengths through an environment, the lengths of compatible full h -signatures varies. For example, if an agent is observed passing from left to right above the first obstacle in an environment their partial h -signature is (1). Later, the agent may or may not pass above any other obstacle. Then, their full h -signature may be (1) or (1, ...).

To handle these nuances, we assume that the longest possible h -signature an agent will take through the environment is the maximum length h -signature present in training data. That is, for any $h \in \mathcal{H}(p)$, $P(h) = 0$ if $|h| > \max(|p|)$. Then, we weigh h -signature probabilities in Eqn. 1 by the probability of observing an h -signature of that length in data. Thus, the VOMP is queried in order to make the adjusted calculation

$$P(h \mid p) = \frac{P(h) \cdot P(|h|)}{\sum_{h \in \mathcal{H}(p)} P(h) \cdot P(|h|)}. \quad (2)$$

2) *Offline Learning*: To facilitate online prediction of full h -signatures given partial h -signatures, a prediction suffix tree (PST) (Fig. 4 (a)) is built over the alphabet \mathcal{A} admitted by all possible h -signatures given the environment. The PST is then converted to a PSA (Fig 4 (b)) as per [13]. A high-level overview of this process is provided in Alg. 1.

To construct the PST, it is first initialised with root node corresponding to the ‘empty’ h -signature labelled (). Paths to suffixes are then successively added to the tree if the suffix has sufficiently strong predictive power. Specifically, a child node labelled with partial h -signature ap is added to parent node p if some measure of statistical difference \mathcal{E} between $P(\cdot \mid p)$ and $P(\cdot \mid ap)$ is above a user-defined threshold ϵ . The metric used is the KL divergence scaled by the probability of observing ap ,

$$\mathcal{E}(ap, p) = P(ap) D_{KL}(P(\cdot \mid ap) \parallel P(\cdot \mid p)). \quad (3)$$

This scaling factor $P(ap)$ serves to avoid the addition of suffixes that have very low probability of occurring, yet give large KL divergence. Suffixes up to length L are tested for their predictive power and added.




Fig. 4. Equivalent methods of representing a VOMP. Values along edges between states are transition probabilities.

To calculate $P(a \mid p)$ and $P(p)$ from training data, Laplace’s rule of succession is used. Denoting $\mathcal{F}(p)$ as the frequency with which partial h -signature p appears in data, $\mathcal{F}(p^C)$ is then used to denote the frequency of the complement event; the frequency of observing any other partial h -signature of length $|p|$. Then,

$$P(p) \approx \frac{\mathcal{F}(p) + 1}{\mathcal{F}(p^C) + |\mathcal{A}|}. \quad (4)$$

Similarly, denoting the frequency with which letter a follows p in observations by $\mathcal{F}(a|p)$ and the frequency of observing any other letter after p by $\mathcal{F}(a^C \mid p)$,

$$P(a \mid p) \approx \frac{\mathcal{F}(a|p) + 1}{\mathcal{F}(a^C \mid p) + |\mathcal{A}|}. \quad (5)$$

After construction, taking the leaves of the PST gives the states of the corresponding PSA. However, this simple action may not always admit a valid transition between all states. In this case, leaves must be added to the PST to ensure a complete PSA. Nodes are added to the PST until, for every leaf in the PST, the longest prefix of the leaf exists in the PST. When this condition is true, the leaves of the PST are guaranteed to give a complete PSA [13]. Transition probabilities between states of the PSA correspond directly to the transition probabilities of the PST found via Eqn. 5 and are used to calculate probabilities online.

D. Illustration of h -signature Prediction

Here, an illustration of VOMP’s high-level motion prediction process is given. A simplistic toy environment and dataset was created, shown in Fig. 5 (a). Trajectory data was created by running Dijkstra’s algorithm on a graph over the environment to find shortest-distance paths from randomly selected points on the left boundary to random points on all other boundaries. Thus, in this simple dataset all trajectories move from left to right, and the set of possible homotopy classes are denoted by h -signatures $\{(), (1), (1, 2)\}$.

Prediction of h -signatures over time is shown in the following figures Fig. 5 (b) - (d). At time $t = 8$, the partial h -signature is (), and the VOMP predicts that all h -signature trajectories are possible in the future, with probabilities $P(()) = 0.28$, $P((1)) = 0.37$, $P((1, 2)) = 0.35$. At $t = 17$ the () class is predicted with probability 0 now the observed trajectory has passed above the first obstacle, and $P((1)) = 0.52$, $P((1, 2)) = 0.48$. In the final time step the




Fig. 5. VOMP h -signature prediction for a test trajectory over time. Black trajectory is the observed test trajectory. Coloured regions correspond to straight-line path completions that lie within predicted homotopy classes. Colour coding of homotopy classes is as in the training set. Alpha of the regions reflects the VOMP output probability for that class.

VOMP correctly assesses that the trajectory will have full h -signature $(1, 2)$ with probability 1 as the trajectory has passed the second obstacle. Coloured regions correspond to regions in which one may expect the agent to be in the future, with probabilities indicated by opacity. These regions demonstrate the predictive power of the VOMP output.

E. Low-level Prediction using hierarchical Gaussian Mixture Models

We present a hierarchical GMM as an example implementation of the low-level prediction algorithm for retrieving a probability distribution $P(\mathbf{X} | h)$ over the trajectory \mathbf{X} given the final h -signature h . To this end, we simply cluster the trajectories from the training dataset into their homotopy classes (i.e. having the same h -signature h), and fit a GMM for each class,

$$P(\mathbf{X} | h) = \sum_c w^{(c,h)} \mathcal{N}(\mathbf{X} | \mathbf{M}^{(c,h)}, \Sigma^{(c,h)}). \quad (6)$$

Here, for a given h -signature h , $w^{(c,h)}$, $\mathbf{M}^{(c,h)}$ and $\Sigma^{(c,h)}$ are the weight, mean and covariance respectively of a component $c \in [1, N_C]$. $\mathcal{N}(\mathbf{X} | \mathbf{M}, \Sigma)$ is the multivariate normal distribution over \mathbf{X} with mean vector \mathbf{M} and covariance Σ . In doing so, the low-level patterns in the trajectories can be captured by computing the full covariance matrix Σ , thereby capturing correlations between positions at particular times.

From this distribution, one can derive a fully probabilistic prediction of the trajectory given position measurements \mathbf{Y}_{obs} and partial h -signature p . To simplify computation, we treat the partial h -signature p as an additional independent measurement to \mathbf{Y}_{obs} . Then, the partial h -signature is conditionally independent of \mathbf{X} given the full h -signature h , and we can easily condition on the partial h -signature as a weighted sum $P(\mathbf{X} | p) = \sum_h P(\mathbf{X} | h)P(h | p)$. This effectively only scales the weights of the GMM by $P(h | p)$ given by the VOMP.

Subsequently, we can further condition on the actual measurements \mathbf{Y}_{obs} . Following standard methods,

$$P(\mathbf{X} | p, \mathbf{Y}_{\text{obs}}) = \sum_{c,h} \hat{w}_{\text{obs}}^{(c,h)} \mathcal{N}(\mathbf{X} | \hat{\mathbf{M}}_{\text{obs}}^{(c,h)}, \hat{\Sigma}_{\text{obs}}^{(c,h)}). \quad (7)$$

Here, the conditional mean and covariance $\hat{\mathbf{M}}_{\text{obs}}^{(c,h)}$, $\hat{\Sigma}_{\text{obs}}^{(c,h)}$ are calculated in the same manner as standard conditional Gaussian distribution [31, Sec. 8.1.3] given measurements. The conditional weights are calculated as

$$\hat{w}_{\text{obs}}^{(c,h)} \propto w_t^{(c,h)} P(h | p) \mathcal{N}(\mathbf{Y}_{\text{obs}} | \mathbf{M}^{(c,h)}, \Sigma^{(c,h)} + \sigma_{\mathbf{Y}}^2 I), \quad (8)$$

followed by normalisation. The last term is the marginal likelihood of observing \mathbf{Y}_t within each mixture component.

V. EXPERIMENTAL RESULTS

A. Experimental Setup

We experimentally illustrate the benefits of high-level prediction by comparison against a standard GMM without topological knowledge. We use the ATC shopping mall dataset [14] shown in Fig. 6 (a). We retrieved a portion of the dataset containing 17558 trajectories. We selected 9230 trajectories among the dataset trajectories contained within the area shown in Fig. 6 (a) that satisfy the border crossing assumption. Among these trajectories, we randomly selected 5538 trajectories to create a training dataset, and 1000 trajectories for testing. Trajectories are interpolated over a discrete number of timesteps and used to evaluate the prediction performance of the two pipelines.

Three metrics are used for comparison. Most immediately, we consider the standard *average displacement error* (ADE), which measures the deterministic error in the maximum likelihood prediction. This is defined as

$$\text{ADE} = \frac{1}{T} \sum_t \|\mathbf{x}_t - \hat{\mathbf{M}}_{t,\text{obs}}^{(c^*,h^*)}\|, \quad (9)$$

where (c^*, h^*) denotes the index of the component with the highest weight.

Since the framework is fully probabilistic, we also need to account for the uncertainty estimates produced. We also consider the *average Mahalanobis distance* (AMD), which accounts for the weights and covariances in the GMM:

$$\text{AMD} = \sum_{c,h} \frac{\hat{w}_{\text{obs}}^{(c,h)}}{T} \sum_t \Theta_t^{(c,h)\text{T}} (\hat{\Sigma}_{t,\text{obs}}^{(c,h)})^{-1} \Theta_t^{(c,h)}, \quad (10)$$

where $\Theta_t^{(c,h)} = (\mathbf{x}_t - \hat{\mathbf{M}}_{t,\text{obs}}^{(c,h)})$.

Lastly, we would like to characterise the performance of the high-level prediction made. To do so, we compare the *KL divergence* (KLD) between the posterior GMM weights given the full trajectory and a partial trajectory, defined as

$$\text{KLD} = \sum_{c,h} \hat{w}_T^{(c,h)} (\log \hat{w}_T^{(c,h)} - \log \hat{w}_{\text{obs}}^{(c,h)}). \quad (11)$$

For fairness, we use the same number of mixtures for the naive GMM as in the topology-informed GMM. All GMMs used are trained using the implementation in the SCIPY library, using the default arguments.

B. Results

Figure 7 shows the results of comparisons between naive GMM (red) and our topology-informed approach (green) in terms of ADE (top), AMD (middle), and KLD (bottom). In terms of ADE, our approach outperforms the naive approach after some time, namely showing a 34.8% improvement in the median around halfway through the trajectory ($t = 56$), and up to 69.4% over time. This increase in performance is because the topology-informed approach gains more information as the target crosses obstacles, and this topological information improves the predictive power of the VOMP.

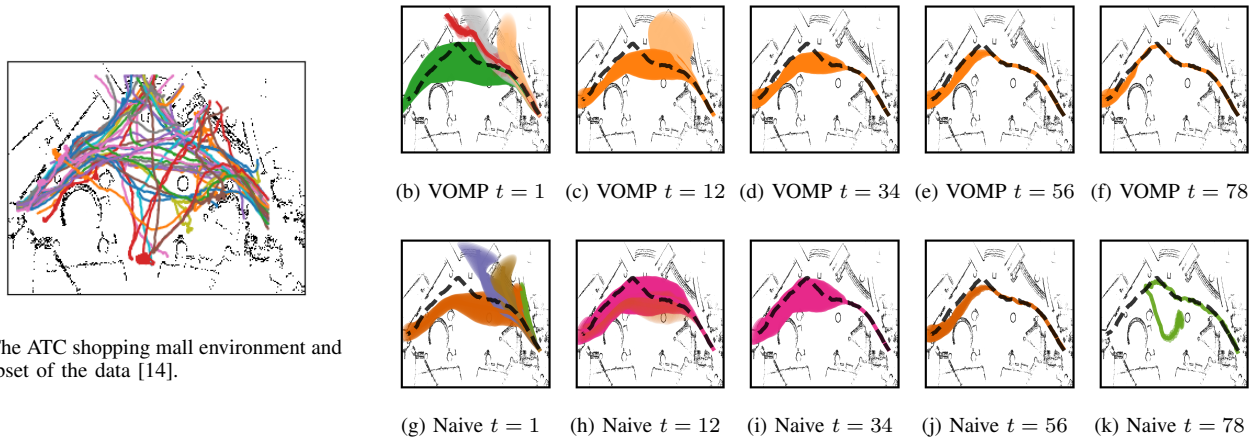


Fig. 6. Snapshots of predictions output by a GMM with h -signature context given by the VOMP (b) - (f) and without (g) - (k). Black dashed trajectory is the ground truth test trajectory. Shaded regions show the variance of around the mean trajectory in solid colour. Transparency is proportional to weight. Colours in (b) - (f) indicate h -signatures, while colours in (g) - (k) indicate mixture components.

Since the naive GMM was given the same number of components as the topology-informed one calculated based on the number of homotopy classes, it is unsurprising that the naive GMM performs relatively well. In practice, a major challenge in deploying GMMs is selecting the right number of mixture components. In fact, the naive GMM is initialised with a number directly informed by the number of homotopy classes present in the data, providing a small topological cue to the baseline method.

Nonetheless, the AMD metric shows that our approach performs significantly better than or equal to the naive approach. Specifically, we see a 72.3% improvement over the naive halfway through the trajectory, with a maximum of 80.3% improvement over time. This implies that the uncertainty predictions from the topology-informed approach are more *consistent* with the actual error compared to the naive approach. This is because our approach pre-clusters trajectories that are ‘similar’, improving the fitness of the GMM model within each cluster.

Further, our approach consistently outperforms in terms of KLD by a great margin, with 100% improvement achieved by halfway through the trajectory. This means that early trajectory GMM weights of our approach better represented the final GMM weights. Viewing GMM components as an alternative description of high-level motion to h -signatures, this implies that h -signatures are indeed better descriptors.

Figure 6 shows the behaviour of our approach and the baseline GMM. Early on, we see both methods give highest weight to components that reasonably predict the general motion of the trajectory. However, the naive approach begins to produce poor predictions in an area of the environment where many training paths diverge. This ‘crossroads’ is visible in the sample data plotted in Fig. 6 (a). Additionally, we see the naive GMM re-weights the highest probability component often, in-line with its performance in the KLD metric. In comparison, our approach continues to correctly predict the same h -signature across time, even when reaching the aforementioned crossroads. This demonstrates the power of topological information in imbuing low-level predictions

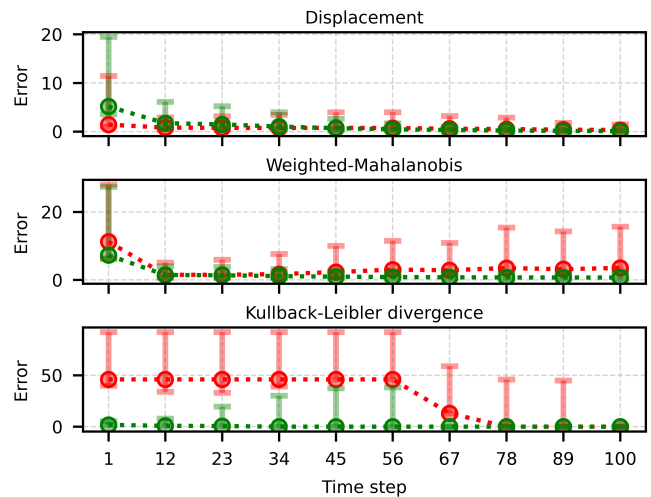


Fig. 7. Experimental results. Red: Naive GMM, Green: Our approach. Lower is better. Markers show the median. Error bars show 25% and 75% quantiles.

with robustness to deviations in trajectories.

VI. CONCLUSION

We presented a framework for trajectory prediction using homotopy classes. The core of the framework is the notion of partial h -signature, which allows prediction of the full trajectory at both abstract and geometric levels. Under this framework, we presented VOMP and hierarchical GMMs as minimal implementations of the high-level and low-level prediction components, the combination of which was shown to outperform baselines without topological knowledge.

Our framework opens many exciting opportunities for future work. The high-level and low-level components are modular, and can be replaced with many other algorithms of varying sophistication and efficiency. Further, we expect the saliency and sparseness of h -signatures will be instrumental when directly used in planning problems such as target search in cluttered environments [15]. We believe similar topological features can be defined in other application domains such as marine buoys [32, 33].

REFERENCES

- [1] G. Best and R. Fitch, "Bayesian intention inference for trajectory prediction with an unknown goal destination," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Hamburg, Germany: IEEE, Sept. 2015, pp. 5817–5823.
- [2] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine, "Precog: Prediction conditioned on goals in visual multi-agent settings," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, October 2019.
- [3] S. Eiffert, K. Li, M. Shan, S. Worrall, S. Sukkarieh, and E. Nebot, "Probabilistic Crowd GAN: Multimodal Pedestrian Trajectory Prediction using a Graph Vehicle-Pedestrian Attention Network," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5026–5033, Oct. 2020, arXiv:2006.12906 [cs].
- [4] K. Katuwandeniya, S. H. Kiss, L. Shi, and J. Valls Miro, "Multi-modal scene-compliant user intention estimation in navigation," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 1001–1006.
- [5] K. H. Dinh, O. Oguz, G. Huber, V. Gabler, and D. Wollherr, "An approach to integrate human motion prediction into local obstacle avoidance in close human-robot collaboration," in *2015 IEEE International Workshop on Advanced Robotics and its Social Impacts*, 2015, pp. 1–6.
- [6] S. El Zaatari, M. Marei, W. Li, and Z. Usman, "Cobot programming for collaborative industrial tasks: An overview," *Robotics and Autonomous Systems*, vol. 116, pp. 162–180, 2019.
- [7] S. Bhattacharya, V. Kumar, and M. Likhachev, "Search-based path planning with homotopy class constraints," in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. Atlanta, Georgia: AAAI Press, July 2010, pp. 1230–1237.
- [8] S. Bhattacharya, R. Ghrist, and V. Kumar, "Persistent Homology for Path Planning in Uncertain Environments," *IEEE Transactions on Robotics*, vol. 31, no. 3, pp. 578–590, June 2015.
- [9] S. Bhattacharya, S. Kim, H. Heidarsson, G. S. Sukhatme, and V. Kumar, "A topological approach to using cables to separate and manipulate sets of objects," *The International Journal of Robotics Research*, vol. 34, no. 6, pp. 799–815, May 2015.
- [10] S. McCammon and G. A. Hollinger, "Topological path planning for autonomous information gathering," *Autonomous Robots*, 2021, publisher: Springer US.
- [11] W. Zhi, L. Ott, and F. Ramos, "Kernel Trajectory Maps for Multi-Modal Probabilistic Motion Prediction," in *Proceedings of the Conference on Robot Learning*, L. P. Kaelbling, D. Kravic, and K. Sugiura, Eds., vol. 100. PMLR, Nov. 2020, pp. 1405–1414.
- [12] S. Eiffert, H. Kong, N. Pirmarzdashti, and S. Sukkarieh, "Path Planning in Dynamic Environments using Generative RNNs and Monte Carlo Tree Search," in *2020 IEEE International Conference on Robotics and Automation*. Paris, France: IEEE, May 2020, pp. 10 263–10 269.
- [13] D. Ron, Y. Singer, and N. Tishby, "The Power of Amnesia," in *Advances in Neural Information Processing Systems*, J. Cowan, G. Tesauero, and J. Alsppector, Eds., vol. 6. Morgan-Kaufmann, 1993.
- [14] D. Brscic, T. Kanda, T. Ikeda, and T. Miyashita, "Person position and body direction tracking in large public spaces using 3D range sensors," *Trans. on Human-Mach. Sys.*, vol. 43, pp. 522 – 534, 2013.
- [15] J. Wakulicz, H. Kong, and S. Sukkarieh, "Active Information Acquisition under Arbitrary Unknown Disturbances," in *2021 IEEE International Conference on Robotics and Automation*. Xi'an, China: IEEE, May 2021, pp. 8429–8435.
- [16] W. Zhi, R. Senanayake, L. Ott, and F. Ramos, "Spatiotemporal Learning of Directional Uncertainty in Urban Environments With Kernel Recurrent Mixture Density Networks," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4306–4313, Oct. 2019.
- [17] F. Althe and A. de La Fortelle, "An LSTM network for highway trajectory prediction," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems*. Yokohama: IEEE, Oct. 2017, pp. 353–359.
- [18] Z. Shi, M. Xu, Q. Pan, B. Yan, and H. Zhang, "LSTM-based Flight Trajectory Prediction," in *2018 International Joint Conference on Neural Networks*. Rio de Janeiro: IEEE, July 2018, pp. 1–8.
- [19] J. Bi, H. Yuan, K. Xu, H. Ma, and M. Zhou, "Large-scale Network Traffic Prediction With LSTM and Temporal Convolutional Networks," in *2022 International Conference on Robotics and Automation*. Philadelphia, PA, USA: IEEE, May 2022, pp. 3865–3870.
- [20] J. Wiest, M. Hoffken, U. Kresel, and K. Dietmayer, "Probabilistic trajectory prediction with Gaussian mixture models," in *2012 IEEE Intelligent Vehicles Symposium*. Alcal de Henares, Madrid, Spain: IEEE, June 2012, pp. 141–146.
- [21] B. Ivanovic and M. Pavone, "The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, October 2019.
- [22] A. Zyner, S. Worrall, and E. Nebot, "Naturalistic Driver Intention and Path Prediction Using Recurrent Neural Networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 4, pp. 1584–1594, Apr. 2020.
- [23] A. Vemula, K. Muelling, and J. Oh, "Social Attention: Modeling Attention in Human Crowds," in *2018 IEEE International Conference on Robotics and Automation*. Brisbane, QLD: IEEE, May 2018, pp. 4601–4607.
- [24] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human Trajectory Prediction in Crowded Spaces," in *2016 IEEE Conference on Computer Vision and Pattern Recognition*. Las Vegas, NV, USA: IEEE, June 2016, pp. 961–971.
- [25] S. H. Kiss, K. Katuwandeniya, A. Alempijevic, and T. Vidal-Calleja, "Probabilistic Dynamic Crowd Prediction for Social Navigation," in *2021 IEEE International Conference on Robotics and Automation*. Xi'an, China: IEEE, May 2021, pp. 9269–9275.
- [26] —, "Constrained Gaussian Processes With Integrated Kernels for Long-Horizon Prediction of Dense Pedestrian Crowd Flows," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7343–7350, July 2022.
- [27] Y. Yao, E. Atkins, M. Johnson-Roberson, R. Vasudevan, and X. Du, "BiTraP: Bi-Directional Pedestrian Trajectory Prediction With Multi-Modal Goal Estimation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1463–1470, Apr. 2021.
- [28] F. T. Pokorny, M. Hawasly, and S. Ramamoorthy, "Topological trajectory classification with filtrations of simplicial complexes and persistent homology," *International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 204–223, 2016, ISBN: 0278364915586.
- [29] J. Frederico Carvalho, M. Vejdemo-Johansson, F. T. Pokorny, and D. Kravic, "Long-term Prediction of Motion Trajectories Using Path Homology Clusters," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Macau, China: IEEE, Nov. 2019, pp. 765–772.
- [30] A. Hatcher, *Algebraic Topology*. Cambridge University Press, 2002.
- [31] K. B. Petersen and M. S. Pedersen, *The Matrix Cookbook*. Technical University of Denmark, 2007, vol. 16, no. 4.
- [32] D. Roemmich, M. H. Alford, H. Claustre, K. Johnson, B. King, J. Moum, P. Oke, W. B. Owens, S. Pouliquen, S. Purkey, M. Scanderberg, T. Suga, S. Wijffels, N. Zilberman, D. Bakker, M. Baringer, M. Belbeoch, H. C. Bittig, E. Boss, P. Calil, F. Carse, T. Carval, F. Chai, D. O. Conchubhair, F. d'Ortenzio, G. Dall'Olmo, D. Desbruyeres, K. Fennel, I. Fer, R. Ferrari, G. Forget, H. Freeland, T. Fujiki, M. Gehlen, B. Greenan, R. Hallberg, T. Hibiya, S. Hosoda, S. Jayne, M. Jochum, G. C. Johnson, K. Kang, N. Kolodziejczyk, A. Körtzinger, P.-Y. L. Traon, Y.-D. Lenn, G. Maze, K. A. Mork, T. Morris, T. Nagai, J. Nash, A. N. Garabato, A. Olsen, R. R. Pattabhi, S. Prakash, S. Riser, C. Schmechtig, C. Schmid, E. Shroyer, A. Sterl, P. Sutton, L. Talley, T. Tanhua, V. Thierry, S. Thomalla, J. Toole, A. Troisi, T. W. Trull, J. Turton, P. J. Velez-Belchi, W. Walczowski, H. Wang, R. Wanninkhof, A. F. Waterhouse, S. Waterman, A. Watson, C. Wilson, A. P. S. Wong, J. Xu, and I. Yasuda, "On the future of Argo: A global, full-depth, multi-disciplinary array," *Frontiers in Marine Science*, vol. 6, p. 439, 2019.
- [33] E. Aro, "The utility of an autonomous multi-robot system of underwater floats," *IFAC Proc. Vol.*, vol. 45, no. 28, pp. 55–59, 2012.