
Learning from lions: inferring the utility of agents from their trajectories

Adam D. Cobb

Department of Engineering Science
University of Oxford
acobb@robots.ox.ac.uk

Andrew Markham

Department of Computer Science
University of Oxford
andrew.markham@cs.ox.ac.uk

Stephen J. Roberts

Department of Engineering Science
University of Oxford
sjrob@robots.ox.ac.uk

Abstract

We build a model using Gaussian processes to infer a spatio-temporal vector field from observed agent trajectories. Significant landmarks or influence points in agent surroundings are jointly derived through vector calculus operations that indicate presence of sources and sinks. We evaluate these influence points by using the Kullback–Leibler divergence between the posterior and prior Laplacian of the inferred spatio-temporal vector field. Through locating significant features that influence trajectories, our model aims to give greater insight into underlying causal utility functions that determine agent decision-making. A key feature of our model is that it infers a joint Gaussian process over the observed trajectories, the time-varying vector field of utility and canonical vector calculus operators. We apply our model to both synthetic data and lion GPS data collected at the Buby Valley Conservancy in southern Zimbabwe.

1 Introduction

Inferring the possible cause of an agent’s behaviour from observing its actions is an important area of research across multiple domains. Generic solutions to this problem typically exploit Inverse Reinforcement Learning (IRL) [1] and include learning preference value functions from choices [2], as well as utility, value or policy surfaces from observed actions in a space. In this paper we consider observed trajectories, influenced by a time- and space-varying utility function, and look to infer the spatio-temporal utility function from observed tracks. Our work is motivated by a desire to further understand animal interaction with an environment, fuelled by the existence of long-term GPS data. Instead of making deductions about likely tracks, given our knowledge of features in the environment, we look at inverting the task by inferring these significant features in the landscape from tracks. Our model constructs Gaussian processes (GPs) for jointly inferring a spatio-temporal vector field and canonical vector calculus operations, allowing us to estimate a time-varying map of sources and sinks of potential influence.

Representing the utility function using a GP model is applied to IRL applications in both [3, 4], which improve on previous IRL techniques that constrain the reward to be a linear combination of features [5, 6, 7]. Instead of treating the problem of inferring characteristics of the utility function through finding an optimal policy that matches observed agent demonstrations, we take advantage of using a GP model to infer the distributions over trajectories and their joint derivatives. These distributions

provide an estimate of significant attractors and repellers in an agent’s surrounding which influence the observed agent’s dynamics.

The use of GPs to model vector fields can be seen in [8] to directly model magnetic fields. They use divergence-free and curl-free kernels to enforce linear constraints on their models. In [9], a method for incorporating any linear constraint is presented. To the best of our knowledge, apart from the above work on modeling vector fields directly with Gaussian processes, the derivative properties of GPs are not currently exploited to infer vector calculus operations from lower order observations.

Previous techniques for studying how animals interact with the environment rely on using GPS data to build density maps to construct probability distributions of their likely locations [10, 11]. Although more recent approaches have incorporated time [12] into these models, current methods do not focus on inferring the driving force behind animal actions. This work looks to address this gap in the literature, providing not only a tool to aid in animal behaviour studies, but also as an addition to the growing body of work in inverse reinforcement learning.

The rest of the paper is organised as follows: Section 2 introduces the GP and its associated derivative manipulations for vector calculus. The model is discussed in Section 3, which is followed by a description of the synthetic data experiments in Section 4. The results of applying the model to the lion GPS data are presented in Section 5. Finally, we discuss the implications of our results and comment on future work in Section 6.

2 Preliminaries

As a requirement for our model, we introduce the Gaussian process, defined by its mean function $\mu(\mathbf{x})$ and covariance function $\mathbf{K}(\mathbf{x}, \mathbf{x}')$ [13]. The mean and covariance functions are parameterized by their hyperparameters and encode prior information into the model, such as smoothness, periodicity and any known underlying behaviour. In our work, $\mu(\mathbf{x})$ is set to zero as our data is preprocessed to have a zero mean. We define a function, distributed via a GP, as follows:

$$\mathbf{f}(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), \mathbf{K}(\mathbf{x}, \mathbf{x}')). \quad (1)$$

Manipulating Gaussian identities using Bayes’ rule gives formulae for the GP’s posterior mean,

$$\mathbb{E}[f(\mathbf{x}_*)] = \mathbf{k}_{\mathbf{x}_* \mathbf{x}}^\top (\mathbf{K}_{\mathbf{xx}} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, \quad (2)$$

and posterior covariance,

$$\mathbb{V}[f(\mathbf{x}_*)] = k_{\mathbf{x}_* \mathbf{x}_*} - \mathbf{k}_{\mathbf{x}_* \mathbf{x}}^\top (\mathbf{K}_{\mathbf{xx}} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_{\mathbf{x}_* \mathbf{x}}, \quad (3)$$

where \mathbf{x}_* is a test point under question and σ^2 is the noise variance hyperparameter.

Any affine transformation of Gaussian distributed variables remain jointly Gaussian distributed. As differentiation is an affine operation, applying this property to any collection of random variables distributed by a GP, gives jointly Gaussian distributed derivatives, $\mathbf{f}'(\mathbf{x})$. For a test point \mathbf{x}_* and corresponding output $f(\mathbf{x}_*)$, the derivatives associated with the GP in Equation (1) are distributed with posterior mean,

$$\mathbb{E}\left[\frac{\partial^n f(\mathbf{x}_*)}{\partial \mathbf{x}_*^n}\right] = \frac{\partial^n \mathbf{k}_{\mathbf{x}_* \mathbf{x}}^\top}{\partial \mathbf{x}_*^n} (\mathbf{K}_{\mathbf{xx}} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, \quad (4)$$

and posterior covariance,

$$\mathbb{V}\left[\frac{\partial^n f(\mathbf{x}_*)}{\partial \mathbf{x}_*^n}\right] = \frac{\partial^{2n} k_{\mathbf{x}_* \mathbf{x}'_*}}{\partial \mathbf{x}_*^n \partial \mathbf{x}'_*^n} - \frac{\partial^n \mathbf{k}_{\mathbf{x}_* \mathbf{x}}^\top}{\partial \mathbf{x}_*^n} (\mathbf{K}_{\mathbf{xx}} + \sigma^2 \mathbf{I})^{-1} \frac{\partial^n \mathbf{k}_{\mathbf{x}_* \mathbf{x}}}{\partial \mathbf{x}_*^n}. \quad (5)$$

We define Equation (4) as the predictive mean of the n^{th} derivative with respect to any test points \mathbf{x}_* and Equation (5) as its corresponding variance.

The choice of covariance selected throughout this paper is the squared exponential kernel,

$$k_{\text{SE}}(\mathbf{x}_*, \mathbf{x}_i) = l^2 \exp\left(-\frac{1}{2}(\mathbf{x}_* - \mathbf{x}_i)^\top \mathbf{\Lambda}^{-1}(\mathbf{x}_* - \mathbf{x}_i)\right), \quad (6)$$

with \mathbf{x}_* and \mathbf{x}_i corresponding to a test point and training point respectively. The hyperparameter $\mathbf{\Lambda}$, is a diagonal matrix of input scale lengths, where each element determines the relevance of its

corresponding dimension and the output scale, denoted by l^2 , controls the magnitude of the output [14]. Although the choice of kernel is somewhat arbitrary, our choice is motivated by the desire to obtain smooth measures over arbitrary derivative functions and the ease by which the kernel can be differentiated and used in combination with Equations (4) and (5). The following formulae define the squared exponential kernel for the first and second order derivatives [15]:

$$\frac{\partial k_{\text{SE}}(\mathbf{x}_*, \mathbf{x}_i)}{\partial \mathbf{x}_*} = -\mathbf{\Lambda}^{-1}(\mathbf{x}_* - \mathbf{x}_i) k_{\text{SE}}(\mathbf{x}_*, \mathbf{x}_i) \quad (7)$$

$$\frac{\partial^2 k_{\text{SE}}(\mathbf{x}_*, \mathbf{x}_i)}{\partial \mathbf{x}_*^2} = \mathbf{\Lambda}^{-1} ((\mathbf{x}_* - \mathbf{x}_i)(\mathbf{x}_* - \mathbf{x}_i)^\top \mathbf{\Lambda}^{-1} - \mathbf{I}) k_{\text{SE}}(\mathbf{x}_*, \mathbf{x}_i). \quad (8)$$

We note at this point that estimating derivatives using a joint GP model over the function and derivatives [16, 17] offers a benign noise escalation in comparison to numerical differentiation.

2.1 Vector calculus with GPs

We define each dimension of the vector field, $\mathbf{V}_t = [V_x \ V_y]_t^\top$, at time t to be modelled by a multi-input, multi-output GP with a three-dimensional input tuple consisting of $\mathcal{X} = (x, y, t)$, being the spatial and time components of our trajectory. This GP is constructed by introducing a separable kernel [18], such that

$$\begin{bmatrix} V_x \\ V_y \end{bmatrix} \sim \mathcal{GP} \left(\begin{bmatrix} \mu_x \\ \mu_y \end{bmatrix}, \begin{bmatrix} k_x(\mathcal{X}, \mathcal{X}') & 0 \\ 0 & k_y(\mathcal{X}, \mathcal{X}') \end{bmatrix} \right)$$

contains an independently learned kernel for each output dimension. Applying Equation (4) to this GP model, by jointly inferring the derivatives in the x and y directions, gives the time dependent posterior for each random variable in the following tuple:

$$\left(V_x, V_y, \frac{\partial V_x}{\partial x}, \frac{\partial V_y}{\partial y}, \frac{\partial V_x}{\partial y}, \frac{\partial V_y}{\partial x} \right)_t.$$

We combine these predictive derivatives using Equations (9) and (10) to infer probability distributions over the divergence and curl of the vector field \mathbf{V} :

$$\nabla \cdot \mathbf{V} = \frac{\partial V_x}{\partial x} + \frac{\partial V_y}{\partial y}, \quad (9)$$

$$\nabla \times \mathbf{V} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \frac{\partial}{\partial x} & \frac{\partial}{\partial y} & \frac{\partial}{\partial z} \\ V_x & V_y & 0 \end{vmatrix} = \mathbf{k} \left(\frac{\partial V_y}{\partial x} - \frac{\partial V_x}{\partial y} \right). \quad (10)$$

By simple application of the appropriate operators, we may readily define the time-varying spatial Laplacian:

$$\nabla \cdot (\nabla \psi) \equiv \nabla^2 \psi, \quad (11)$$

where ψ is a scalar potential function. This (time-varying) Laplacian is of key importance as it defines *sources* and *sinks* in the spatial domain.

3 Model

Our model builds upon the theory introduced in section 2. The objective is to design a model that can indicate influential features in an agent's landscape from observing the trajectories.

A trajectory ζ_a for agent a is defined as a collection of timestamped locations $\mathbf{x} \in \mathcal{R}^2$ or tuples (\mathbf{x}, t) . The elements of the vector \mathbf{x} are referred to as x and y for the purposes of this model. We also make the assumption that each agent acts according to a utility or potential function $\phi(\mathbf{x}, t)$, which is dependent on time and space. Whilst interacting with the environment, each agent aims to maximise this utility function at all times.

3.1 Fitting to the agent trajectories

The first stage of the model uses Equations (2) and (3) to fit a GP to each trajectory ζ_a . Using a single GP with a separable kernel, Equation (12) defines our GP prior, which is a joint distribution over the path and its derivatives with respect to time, where we denote \ddot{f}_x and \ddot{f}_y as the second-order time derivatives in the x and y directions:

$$p(\vec{\mathbf{f}}, \ddot{f}_x, \ddot{f}_y) = \mathcal{GP} \left(\mathbf{0}, \begin{bmatrix} k_x(t, t') & 0 \\ 0 & k_y(t, t') \end{bmatrix} \right). \quad (12)$$

For each of the x and y components of $\vec{\mathbf{f}}$, a set of hyperparameters are learned for the separable kernel. The input space of this GP model is time t and the output space consists of the x and y components of $\vec{\mathbf{f}}$. If an agent trajectory, ζ_a , consists of N data points, we can use the posterior, $p(\vec{\mathbf{f}}, \ddot{f}_x, \ddot{f}_y \mid \zeta_a)$, to infer expectations for the derivatives at each of the n data points.

Second-order derivatives are inferred at this stage, as we make the tacit assumption that an agent acts in accordance with a second-order dynamical system, i.e. the agent obeys Newtonian dynamics. This assumption means that an agent’s underlying utility induces a “force” of influence on the agent, thus generating an acceleration (we here take the ‘influence-mass’ of the agent as unity). More formally, this induces an acceleration equal to the derivative of the agent utility:

$$\ddot{f}_x = \frac{\partial \phi}{\partial x}, \quad \ddot{f}_y = \frac{\partial \phi}{\partial y}.$$

When dealing with a multi-agent system of M homogeneous agents, a trajectory model can be calculated for each agent to form the set of joint distributions,

$$\left\{ p(\vec{\mathbf{f}}, \ddot{f}_x, \ddot{f}_y \mid \zeta_a) \right\}_{a=0}^M.$$

From this GP model we are able to jointly predict the velocity and acceleration at any point on an agent’s trajectory for M agents. The next layer of our model combines these to construct a probability distribution over the extended agent environment.

3.2 Inferring the vector field and Laplacian

In order to infer the gradient of the utility function, $\nabla \phi(\mathbf{x}, t)$, the set of inferred second derivatives for all M agents is propagated through a second GP, which also has a separable kernel model, as below:

$$p\left(\vec{\mathbf{V}}(\mathbf{x}, t), \frac{\partial V_x}{\partial x}, \frac{\partial V_y}{\partial y}, \frac{\partial V_x}{\partial y}, \frac{\partial V_y}{\partial x} \mid \{\ddot{f}_x, \ddot{f}_y, \zeta_a\}_{a=0}^M\right) \\ = \mathcal{GP} \left(\begin{bmatrix} \mu_x^{\text{post}} \\ \mu_y^{\text{post}} \end{bmatrix}, \begin{bmatrix} k_x^{\text{post}}(\zeta_a, \zeta'_a) & 0 \\ 0 & k_y^{\text{post}}(\zeta_a, \zeta'_a) \end{bmatrix} \right). \quad (13)$$

The vector $\vec{\mathbf{V}}(\mathbf{x}, t) = [V_x \ V_y]^\top$ consists of two random variables that model the acceleration in the two axes and the superscript label ‘post’ refers to the calculated posterior mean and covariance. Equation (13) combines the M multiple agent paths into one model and enables predictions to be made at different points in space that are not constrained to a single agent trajectory as in Equation (12). The input-output pairs for this GP model are the x, y and t values in each ζ_a that correspond to the \ddot{f}_x and \ddot{f}_y values.

The Newtonian assumption made in Section 3.1 is formally included as

$$\nabla \phi(\mathbf{x}, t) \propto \vec{\mathbf{V}}(\mathbf{x}, t).$$

The partial derivatives from Equation (13) can thence be combined to calculate a distribution of the divergence of $\vec{\mathbf{V}}$. It follows that this divergence is proportional to the Laplacian under the same assumption,

$$\nabla^2 \phi(\mathbf{x}, t) \propto \nabla \cdot \vec{\mathbf{V}}(\mathbf{x}, t).$$

In particular, our interest lies in the estimation of the Laplacian of the utility function, as it indicates *sources* and *sinks* of the potential function in the environment. In this context, we regard *sinks* as agent *attractors* and *sources* as agent *repellers*.

3.3 Metric for locating significant attractors and repellers: Kullback–Leibler divergence

Our metric of change from prior field to posterior field is measured via the Kullback-Leibler (KL) divergence [19]. The motivation for selecting the KL divergence comes from its ability measure a distance between two distributions by taking into account both the mean and variance. This provides a natural indication of the informativeness of spatial locations, at given times, and in the context of our application, offers a measure of *trajectory-influencing locations*.

Given the model at time t , each point in space has an associated potential field distribution, defined via the GP posterior as a univariate normal distribution. The KL divergence can be readily calculated as the difference between two univariate normal distributions, namely the prior and posterior [20], as below:

$$D_{\text{KL}}(p_{\text{prior}} \parallel p_{\text{posterior}}) = \frac{1}{2} \left(\frac{\sigma_{\text{pr}}^2}{\sigma_{\text{po}}^2} + \frac{(\mu_{\text{po}} - \mu_{\text{pr}})^2}{\sigma_{\text{po}}^2} - 1 + \ln \left(\frac{\sigma_{\text{po}}}{\sigma_{\text{pr}}} \right) \right), \quad (14)$$

where

$$p_{\text{prior}} = \mathcal{N}(\mu_{\text{pr}}, \sigma_{\text{pr}}^2), \quad p_{\text{posterior}} = \mathcal{N}(\mu_{\text{po}}, \sigma_{\text{po}}^2).$$

We refer back to Equation (9) and (11) in order to calculate the following prior at location \mathcal{X} in space-time:

$$\left(\frac{\partial V_x}{\partial x} + \frac{\partial V_y}{\partial y} \right) \sim \mathcal{N} \left(0, \frac{\partial^2 k_x(\mathcal{X}, \mathcal{X}')}{\partial x^2} + \frac{\partial^2 k_y(\mathcal{X}, \mathcal{X}')}{\partial y^2} \right),$$

where, $\frac{\partial^2 k_x(\mathcal{X}, \mathcal{X}')}{\partial x^2} + \frac{\partial^2 k_y(\mathcal{X}, \mathcal{X}')}{\partial y^2} = \frac{h_x^2}{\lambda_x^2} + \frac{h_y^2}{\lambda_y^2}$. The hyperparameters h_x and λ_x are the output and input scale lengths of the x -part of the separable kernel in the GP model, with h_y and λ_y corresponding to the y -part.

As our interest lies in determining attractors and repellers in the field, a further addition to the KL divergence in Equation (14) is to multiply it by the sign of the posterior mean of the Laplacian. This multiplication carries over the prediction of negative sinks and positive sources, whilst measuring the divergence from the zero prior mean. We refer to this trivial extension as the *signed KL divergence*.

3.4 Hyperparameter optimisation

Optimisation of the hyperparameters is achieved through maximisation of the GPs marginal log-likelihood using scaled conjugate gradient approaches. We initialise hyperparameters for spatial models with domain knowledge, where available. In our example, lions are dynamic objects with maximal accelerations, which places natural priors on the relationships between spatial movement and time variation.

4 Application to synthetic data

We apply our model to synthetic data where the true utility function and its derivatives are known, so as to test the performance of the approach. The experiment consists of a multi-agent system of four identical agents, whose dynamics we observe. Our goal is to infer, from trajectories alone, the underlying potential value function.

4.1 Utility function

The utility function for the test example is defined as a two-component Gaussian mixture model, where each Gaussian has a time-varying covariance that changes harmonically according to:

$$\phi(\mathbf{x}, t) = \mathcal{N}(\boldsymbol{\mu}_1, (\sin t + a)\mathbf{I}) + \mathcal{N}(\boldsymbol{\mu}_2, (\cos t + b)\mathbf{I}), \quad (15)$$

where a and b are constants and $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2$ are the distribution means. We also denote \mathbf{I} as the identity matrix. The first and second order derivatives are calculated using the known derivatives of a multivariate normal distribution $p(\mathbf{x})$ [21]:

$$\frac{\partial p(\mathbf{x})}{\partial \mathbf{x}} = -p(\mathbf{x})\boldsymbol{\Sigma}^{-1}(\mathbf{x} - \mathbf{m}),$$

$$\frac{\partial^2 p(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} = p(\mathbf{x}) \left(\Sigma^{-1} (\mathbf{x} - \mathbf{m})(\mathbf{x} - \mathbf{m})^\top \Sigma^{-1} - \Sigma^{-1} \right).$$

These derivatives are then used in conjunction with the vector calculus operations in Equations (9) and (11) to calculate the true Laplacian of the utility function.

4.2 Agent model

Agents are modelled according to a second order dynamical system, whereby, at each time step t , the acceleration $\ddot{\mathbf{x}}$, velocity $\dot{\mathbf{x}}$ and position \mathbf{x} are given by:

$$\begin{aligned} \ddot{\mathbf{x}}_{t+1} &= \nabla \phi(\mathbf{x}_t, t) + \mathbf{n}, \\ \dot{\mathbf{x}}_{t+1} &= \dot{\mathbf{x}}_t + \eta \ddot{\mathbf{x}}_{t+1}, \\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \eta \dot{\mathbf{x}}_{t+1}. \end{aligned}$$

The random variable, \mathbf{n} , corresponds to additive noise distributed according to a two-dimensional multivariate Gaussian and the variable η is the update increment.

4.3 Experimental results

Four agents were initialised at locations $[2, 1]$, $[2, -1]$, $[-2, 1]$ and $[-2, -1]$ with a velocity of zero. The parameters of the utility function in Equation (15) were $a = b = 2.1$ and $\boldsymbol{\mu}_1 = [-1.5, 0]$, $\boldsymbol{\mu}_2 = [1.5, 0]$. The model stepped through 200 time-steps and the inferred Laplacian of the utility function was calculated from the four agent trajectories in $\{\zeta_a\}_{a=1}^4$. The subset of frames displayed in Figure 1 show how the attractors of the true utility functions are tracked across time by both the posterior inferred Laplacian and the signed KL divergence. The uncertainty in our model is displayed through the signed KL divergence. The black and white markers denote the location of the dominant attractor, at that time step, given by the ground truth in the first row. For all three rows, regions of attraction are indicated by blue and it can be seen that the model is able to track the true attractors across time.

In Figure 2, the left-hand plot displays the mean KL divergence taken over time, showing the emergence of the average location and size of the attractors. The right-hand plot shows each agent's path and how their trajectories are drawn to the sink locations. This figure demonstrates that our model has correctly inferred the sink locations without prior knowledge of their location.

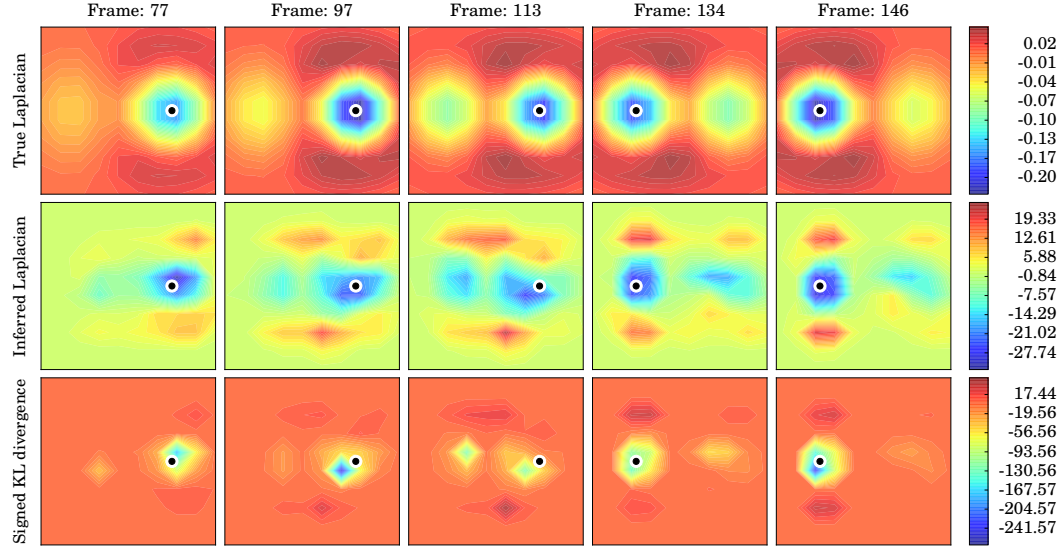


Figure 1: Top row: Laplacian of true utility function. Middle row: inferred Laplacian of utility function. Bottom row: signed KL divergence. The location of the global minimum of the true Laplacian is indicated via the black and white marker. The predicted locations of the sinks, given by both the signed KL divergence and the posterior Laplacian, align well with their true locations.

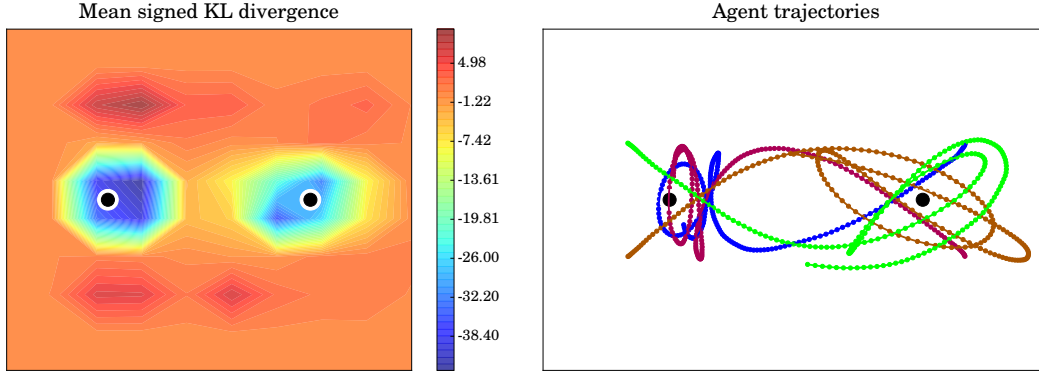


Figure 2: The mean signed KL divergence across time is shown in the left plot and the corresponding agent trajectories are shown in the right plot. Each agent path is denoted via a separate color. The true locations of the attractors are indicated by the black markers. Taking the average over the 200 frames of data allows us to see the mean inferred position of attractors in the agents’ field, which coincide with the known positions.

5 Lion trajectories

In this section, we examine how our model can be used to determine drivers in the landscape that impact an animal’s behaviour. Previous techniques for studying how animals interact with the environment rely on using GPS data to build density maps to construct probability distributions of their likely locations [10, 11]. Although more recent approaches have incorporated time [12] into these models, current methods do not focus on inferring the driving force behind animal actions and instead simply show where an animal is likely to be found. Therefore we apply our model to a subset of lion GPS data to show a possible way of inferring the location of influential features in a lion’s landscape. This perspective differs from traditional ecological techniques, which rely on knowledge of the environment to test a hypothesis. In particular, our focus is to build from work by Valeix et al. (2010) [22] that show waterholes are key attractive features on lions’ trajectories, using knowledge of waterhole locations. Inverting this perspective, applying our model to this data set aims to infer locations of waterholes that are important to a particular lion, without supplying this as a prior.

The model is applied to GPS data collected from lions living at the Buby Valley Conservancy (BVC) in southern Zimbabwe. Figures 3 and 4 display the results of applying the model to 200 data points covering 11 days and 16 hours of a particular lion’s GPS readings. The x and y locations have been normalised in keeping with the relative geographic information system (GIS) data. In both figures the locations of waterholes are included to show their correspondence with predicted regions of attraction. Figure 3 shows the model inferring attractor locations at different time steps, where the units are in hours since the start of the data set. These regions can be seen to overlap the known waterhole locations at certain time steps. In Figure 4 the left-hand plot superimposes the GIS data of waterholes on the mean signed KL divergence produced from the lion trajectory, where the mean is taken across time. The corresponding lion trajectory GPS readings are shown in the right-hand plot. The mean signed KL divergence indicates that there are two waterholes located in areas near local minima, leading to the reasonable assumption that they may behave as influential attractive features. In this way, we can infer decision-making strategies behind observed trajectories. Richer models could investigate further effects such as location of prey and other lions in the ecosystem.

6 Conclusion

Through applying layers of Gaussian Processes combined with vector calculus, the model provides a useful way of inferring significant features in an agent’s surroundings that influence its behaviour. We demonstrate the results of applying this model to a synthetic data set, where the true attractors are known. Our model successfully finds the attractors that determine the agents’ trajectories. Furthermore, applying the model to a subset of lion GPS data allows us to determine a spatio-temporal influence measure which highlights the two waterholes that the lion visits most frequently.




Figure 3: Top row: inferred Laplacian of utility function. Bottom row: signed KL divergence. Each column represents the model at a given hour since the start of the subset of lion data. The relative location of the waterholes are denoted by the black and white markers. Blue regions indicate attractive regions. The signed KL divergence infers attractors occurring at different times that coincide with the waterholes, without prior knowledge of the waterhole locations.

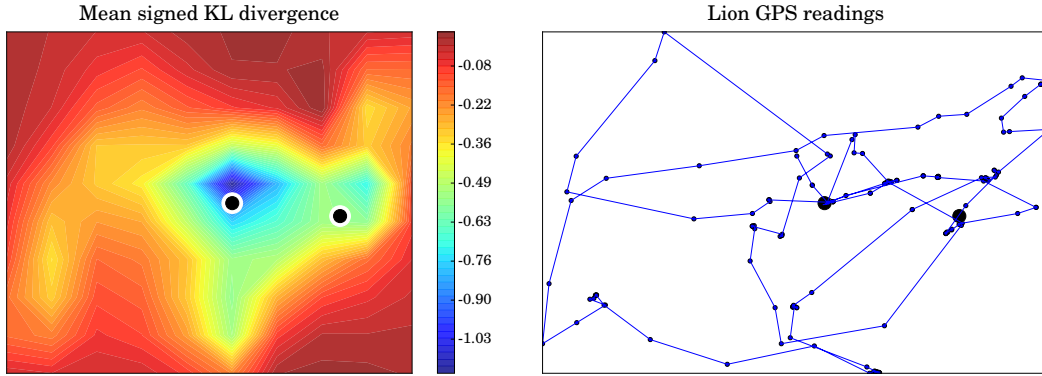


Figure 4: The left plot displays the mean of the signed KL divergence across time to give an average estimate of the lion attractor locations. The blue areas indicate large negative values corresponding to the sinks. The right-hand plot contains the lion trajectory data. The black markers give the relative locations of waterholes, provided by the GIS data. This data set covers 200 GPS readings that were collected over 11 days and 16 hours. The spacial scale for this figure is $15 \times 15 \text{ km}^2$.

Unlike previous work in this domain, such information was not provided to the model. Furthermore, unlike density estimation based methods, our model makes predictions that depend on dynamic decisions made by the lion.

Further work will extend this model to propagate uncertainty from the GPS readings through to the final Laplacian predictions, adapting techniques used in [23, 24]. We aim to extend our model to considerably larger data sets through utilising sparse GP models as needed [25].

Acknowledgements

We would like to thank the shareholders of Buby Valley Conservancy for allowing us to conduct field research on the conservancy. We would also like to thank the management staff of Buby Valley Conservancy for assistance in the field.

References

- [1] Stuart Russell. Learning agents for uncertain environments. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 101–103. ACM, 1998.
- [2] Wei Chu and Zoubin Ghahramani. Preference Learning with Gaussian Processes. In *Proceedings of the 22nd international conference on Machine learning*, pages 137–144. ACM, 2005.

- [3] Sergey Levine, Zoran Popovic, and Vladlen Koltun. Nonlinear Inverse Reinforcement Learning with Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 19–27, 2011.
- [4] Qifeng Qiao and Peter A Beling. Inverse Reinforcement Learning with gaussian Process. In *American Control Conference (ACC), 2011*, pages 113–118. IEEE, 2011.
- [5] Pieter Abbeel and Andrew Y Ng. Apprenticeship Learning via Inverse Reinforcement Learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM, 2004.
- [6] Deepak Ramachandran and Eyal Amir. Bayesian Inverse Reinforcement Learning. *Urbana*, 51(61801):1–4, 2007.
- [7] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum Entropy Inverse Reinforcement Learning. In *AAAI*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.
- [8] Niklas Wahlström, Manon Kok, Thomas B Schön, and Fredrik Gustafsson. Modeling Magnetic Fields Using Gaussian Processes. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 3522–3526. IEEE, 2013.
- [9] Carl Jidling, Niklas Wahlström, Adrian Wills, and Thomas B Schön. Linearly constrained Gaussian processes. *arXiv preprint arXiv:1703.00787*, 2017.
- [10] Jon S Horne, Edward O Garton, Stephen M Krone, and Jesse S Lewis. Analyzing animal movements using Brownian bridges. *Ecology*, 88(9):2354–2363, 2007.
- [11] Peter N Laver and Marcella J Kelly. A Critical Review of Home Range Studies. *Journal of Wildlife Management*, 72(1):290–298, 2008.
- [12] Andrew J Lyons, Wendy C Turner, and Wayne M Getz. Home range plus: a space-time characterization of movement over real landscapes. *Movement Ecology*, 1(1):2, 2013.
- [13] Carl Edward Rasmussen. *Gaussian Processes for Machine Learning*. 2006.
- [14] Stephen Roberts, M Osborne, M Ebden, Steven Reece, N Gibson, and S Aigrain. Gaussian processes for time-series modelling. *Phil. Trans. R. Soc. A*, 371(1984):20110550, 2013.
- [15] Andrew McHutchon. Differentiating Gaussian Processes. <http://mlg.eng.cam.ac.uk/mchutchon/DifferentiatingGPs.pdf>, 2013.
- [16] PR Brook, A Karastergiou, S Johnston, M Kerr, RM Shannon, and SJ Roberts. Emission-rotation correlation in pulsars: new discoveries with optimal techniques. *Monthly Notices of the Royal Astronomical Society*, 456(2):1374–1393, 2016.
- [17] Tracy Holsclaw, Bruno Sansó, Herbert KH Lee, Katrin Heitmann, Salman Habib, David Higdon, and Ujjaini Alam. Gaussian Process Modeling of Derivative Curves. *Technometrics*, 55(1):57–67, 2013.
- [18] Mauricio A Alvarez, Lorenzo Rosasco, Neil D Lawrence, et al. Kernels for Vector-Valued Functions: a Review. *Foundations and Trends® in Machine Learning*, 4(3):195–266, 2012.
- [19] Solomon Kullback and Richard A Leibler. On Information and Sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- [20] John Duchi. *Derivations for Linear Algebra and Optimization*. Berkeley, California, 2007.
- [21] Kaare Brandt Petersen, Michael Syskind Pedersen, et al. The Matrix Cookbook. *Technical University of Denmark*, 7:15, 2008.
- [22] Marion Valeix, Andrew J Loveridge, Zeke Davidson, Hillary Madzikanda, Hervé Fritz, and David W Macdonald. How key habitat features influence large terrestrial carnivore movements: waterholes and African lions in a semi-arid savanna of north-western Zimbabwe. *Landscape Ecology*, 25(3):337–351, 2010.
- [23] Andreas Damianou and Neil Lawrence. Deep Gaussian Processes. In *Artificial Intelligence and Statistics*, pages 207–215, 2013.
- [24] Agathe Girard, Carl Edward Rasmussen, Joaquin Quinero Candela, and Roderick Murray-Smith. Gaussian Process Priors With Uncertain Inputs Application to Multiple-Step Ahead Time Series Forecasting. *Advances in neural information processing systems*, pages 545–552, 2003.
- [25] Michalis K Titsias. Variational Learning of Inducing Variables in Sparse Gaussian Processes. In *AISTATS*, volume 5, pages 567–574, 2009.