Teaching Functions with Gaussian Process Regression

Maya Malaviya, Mark K. Ho
Department of Computer Science
Stevens Institute of Technology
601 Hudson St, Hoboken, NJ 07030 USA
mmalaviy@stevens.edu, mho4@stevens.edu

Abstract

Humans are remarkably adaptive instructors who adjust advice based on their estimations about a learner's prior knowledge and current goals. Many topics that people teach, like goal-directed behaviors, causal systems, categorization, and time-series patterns, have an underlying commonality: they map inputs to outputs through an unknown function. This project builds upon a Gaussian process (GP) regression model that describes learner behavior as they search the hypothesis space of possible underlying functions to find the one that best fits their current data. We extend this work by implementing a teacher model that reasons about a learner's GP regression in order to provide specific information that will help them form an accurate estimation of the function.

Introduction

Learners recognize when they are in a pedagogical scenario and are able to infer more information as a result, even when given very little data (Shafto, Goodman, and Griffiths 2014). Teachers also rely on this situational understanding when choosing advice, tailoring their choices to assist a learner's beliefs and goals (Rafferty et al. 2011; Rafferty, LaMar, and Griffiths 2015), even from a young age (see natural pedagogy, Gweon 2021). However, the majority of previous task paradigms examine pedagogy in small and discrete domains, like categorization or feature-learning (Bridgers, Jara-Ettinger, and Gweon 2020; Sumers et al. 2021). Other work, which models teacher choices in more complex tasks like algebra teaching (Rafferty et al. 2011), only allow teachers to choose from a small pool of discrete actions. This project posits an extension to a Gaussian process regression model that allows examination of pedagogical reasoning in continuous and open-ended task domains. Function learning research characterizes how people narrow down the theoretically infinite space of function hypotheses to more interpretable representations. Function learning has been modeled as a Gaussian process regression, and compositional biases describe human patterns of function hypothesis generation and learning (Lucas et al. 2015; Schulz et al. 2017). But, this body of research has not examined the role of pedagogy in guiding learner hypothesis generation.

Model

We took inspiration from a visual function completion task in which human participants observed an image with a few dots placed along a domain (Schulz et al. 2017). Participants drew a line that represented the function which they believed had produced those dots. We modeled an artificial learner agent who generates an estimate of the underlying function given a set of points, like the aforementioned human participants, with a Gaussian process. We also modeled an artificial teacher agent that generates a useful set of points for a learner to observe, so they may better learn the function.

Function Learning & Teaching

Function learning can be formalized as a Bayesian inference problem, in which a learner updates a belief distribution about a continuous function f conditioned on data points $D \in \{(x_1, y_1), \dots, (x_n, y_n)\}$.

$$P_L(f|D) \propto P(D|f)P(f)$$

In this example, we assume teachers are trying to teach a target function f^* and are tasked with giving a set of n example points $D' = \{(x'_1, y'_1), ..., (x'_n, y'_n)\}$ that help the learner learn the function for some *target inputs* $\mathbf{x}^* \in \mathbb{R}^m$. The teacher's utility function is based on how similar the learner's inferred function is to the true function at the target inputs. Given a function f, a target function f^*, and target inputs $\mathbf{x}^*$, the teacher's function-wise utility is based on the mean-squared error (MSE):

$$U_T(f; f^*, \mathbf{x}^*) = \exp\left(-MSE(f^*(\mathbf{x}^*), f(\mathbf{x}^*))\right)$$

Then, the expected utility for providing the learner with teaching points D', given target function f^* and target inputs $\mathbf{x}^*$ is:

$$U_T(D'; f^*, \mathbf{x}^*) = \mathbb{E}_{P_L(f|D')}[U_T(f; f^*, \mathbf{x}^*)]$$

Gaussian Processes

A Gaussian process (GP) defines a distribution over functions, parameterized by a mean function μ, which specifies the expected output function, and kernel function k which specifies the covariance of outputs. We model learners performing Bayesian updates on a $\mathcal{GP}$ as they gather more data in a process called Gaussian process regression. Let $f : \mathcal{X} \to \mathbb{R}$ be a function drawn from a $\mathcal{GP}$. We chose the radial basis kernel function for k, a standard option for a kernel function that captures a decay in covariance as the distance between inputs increases.

$$f \sim \mathcal{GP}(\mu, k) \qquad k(\mathbf{x_i}, \mathbf{x_j}) = \exp\left(-\frac{\|\mathbf{x_i} - \mathbf{x_j}\|^2}{2\sigma^2}\right)$$

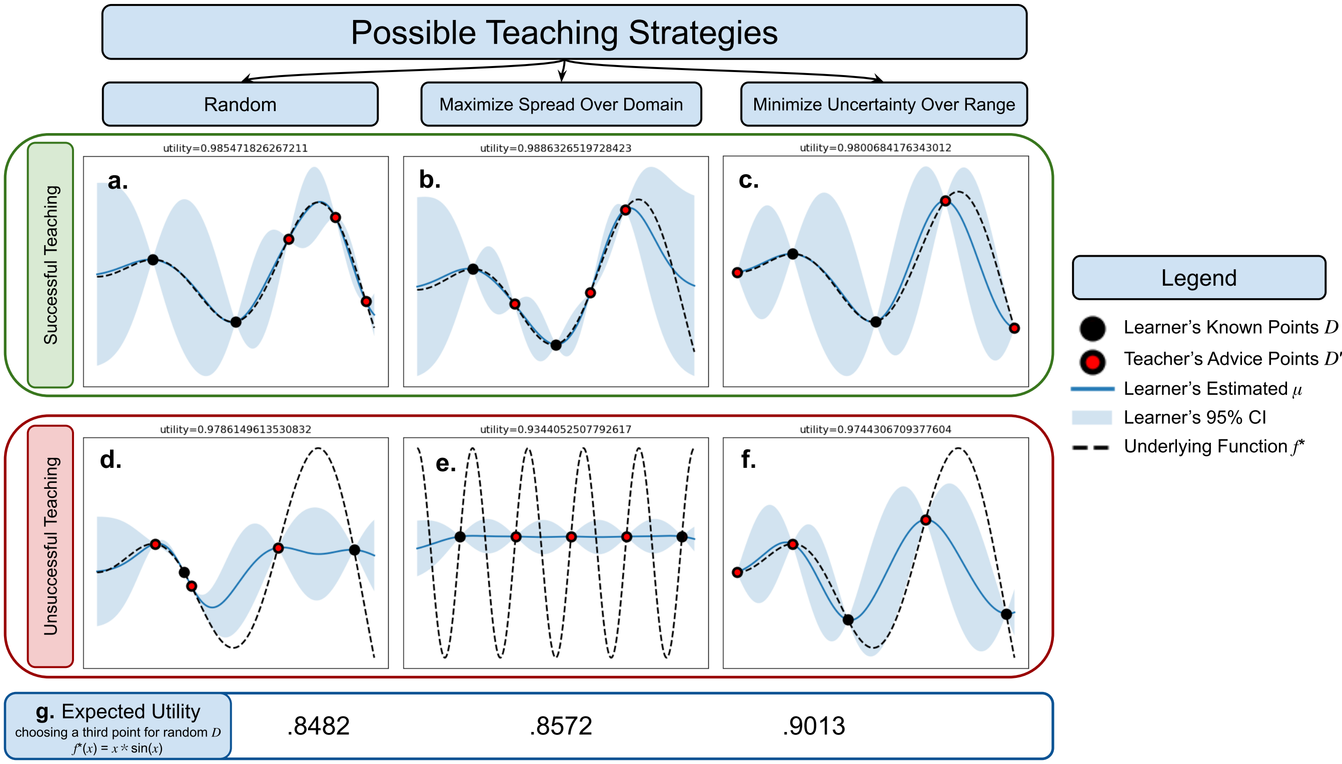

Figure 1: We selected examples of target functions (dotted line) and learners' already-observed points (black dots) to capture the successes and limits of three teacher strategies. Teacher choices for subsequent points are shown as red dots. The learner's resulting $\mathcal{GP}$ is represented by a mean function μ (blue line) and the 95% confidence interval (blue shading) given the kernel function k. The approximated expected utilities $\mathbb{E}(U_T)$ in **g.** quantify the probability that each teaching policy chooses points that induce the right function $f^*(x) = x * \sin x$ given a random known D. Note that by specifying just one f^* and kernel k for a demonstrative example, $\mathbb{E}(U_T)$ does not necessarily generalize to other functions or reflect human choices.

Simulated Teaching Strategies

A teacher model is specified by how it generates and assesses candidate points to give to the learner. Each of the considered teacher models knew the underlying f^* and a set of points which the learner already observed D. Then, they generated candidate sets of n points and selected the D' with the highest utility U_T.

Random Sampling The teacher selects n coordinates uniformly at random to build D'. Random point selection can lead to successful teaching (Figure 1a), but not reliably, since it could select a redundant point that falls too close to something the learner already knows (Figure 1d).

Maximizing Spread Over Domain The teacher calculates the length of the intervals between known $x_i \in D$ along the specified domain $\mathbf{x}^*$. They select the interval $[x_a, x_b]$ with max $|x_b - x_a|$, calculate the $x_{\text{midpt}} = a + \frac{b-a}{2}$ and add $(x_{\text{midpt}}, y_{\text{midpt}})$ to D'. This is repeated n times, with each iteration updating $D \leftarrow D'$. Selecting points that maximize spread across the domain can prevent redundant points, but could fail to capture important parts of the function where there are rapid or patterned changes (e.g., Figure 1e where periodicity lines up with the spread over target domain so only one y value is seen).

Minimizing Uncertainty Over Range The teacher performs approximate inference on the learner's representation of the function distribution. After simulating the learner's $\mathcal{GP}$, the teacher determines the $x_{\text{u}} \in \mathbf{x}^*$ with the highest standard deviation and add its coordinate $(x_{\text{u}}, y_{\text{u}})$ to D'. This is repeated n times, with each iteration updating $D \leftarrow D'$. Inferring where learners are most uncertain and selecting those points is a more reliable pedagogical approach, but could possibly be misleading (e.g., Figure 1f, where local minima and maxima are not yet known).

Discussion

Gaussian processes can help model how people generate informative teaching examples that support learning with little data. In future work, we will collect human data for the function teaching task and perform model fitting and comparison to examine which strategies best capture participants' pedagogical choices. Perhaps a single heuristic is sufficient to capture human choices, but people could flexibly employ many heuristics that approximate inference over a learner's GP regression. We will also consider further teaching heuristics, like prioritizing local minima and maxima coordinates, and more model-based teacher strategies which consider the learner's compositional priors via the kernel function.

References

Bridgers, S.; Jara-Ettinger, J.; and Gweon, H. 2020. Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*, 4(22): 144–152.

Gweon, H. 2021. Inferential social learning: cognitive foundations of human social learning and teaching. *Trends in Cognitive Sciences*, 25(10): 896–910.

Lucas, C. G.; Griffiths, T. L.; Williams, J. J.; and Kalish, M. L. 2015. A rational model of function learning. *Psychonomic Bulletin & Review*, 22(5): 1193–1215.

Rafferty, A. N.; Brunskill, E.; Griffiths, T. L.; and Shafto, P. 2011. Faster Teaching by POMDP Planning. In Biswas, G.; Bull, S.; Kay, J.; and Mitrovic, A., eds., *Artificial Intelligence in Education*, Lecture Notes in Computer Science, 280–287. Berlin, Heidelberg: Springer. ISBN 978-3-642-21869-9.

Rafferty, A. N.; LaMar, M. M.; and Griffiths, T. L. 2015. Inferring Learners' Knowledge From Their Actions. *Cognitive Science*, 39(3): 584–618.

Schulz, E.; Tenenbaum, J. B.; Duvenaud, D.; Speekenbrink, M.; and Gershman, S. J. 2017. Compositional inductive biases in function learning. *Cognitive Psychology*, 99: 44–79.

Shafto, P.; Goodman, N. D.; and Griffiths, T. L. 2014. A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71: 55–89.

Sumers, T. R.; Ho, M. K.; Hawkins, R. D.; Narasimhan, K.; and Griffiths, T. L. 2021. Learning Rewards From Linguistic Feedback. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7): 6002–6010.