Teaching Functions with Gaussian Process Regression

Maya Malaviya (mmalaviy@stevens.edu)
Department of Computer Science, Stevens Institute of Technology
601 Hudson St, Hoboken, NJ 07030 USA

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

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.

Keywords: pedagogy; gaussian processes

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, & Griffiths, 2014). Teachers also rely on this situational understanding when choosing advice, tailoring their choices to assist a learner's beliefs and goals (Rafferty, Brunskill, Griffiths, & Shafto, 2011; Rafferty, LaMar, & 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, & Gweon, 2020; Sumers, Ho, Hawkins, Narasimhan, & Griffiths, 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, Griffiths, Williams, & Kalish, 2015; Schulz, Tenenbaum, Duvenaud, Speekenbrink, & Gershman, 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 that estimates the underlying function given a set of points and an artificial teacher agent that generates a useful set of points that a learner will observe.

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 $D \in \{(x_1, y_1), \ldots, (x_n, y_n)\}$.

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

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 teaching the 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 periodic kernel function for k, a standard option for capturing functions that repeat themselves, with parameters p (periodicity) and ℓ (within-period smoothness).

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

Simulated Teaching Strategies

A teacher model is specified by how it generates and assesses candidate points to give to the learner. Each of the following teacher models knew the underlying f^* and a set of points which the learner already observed D.

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

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

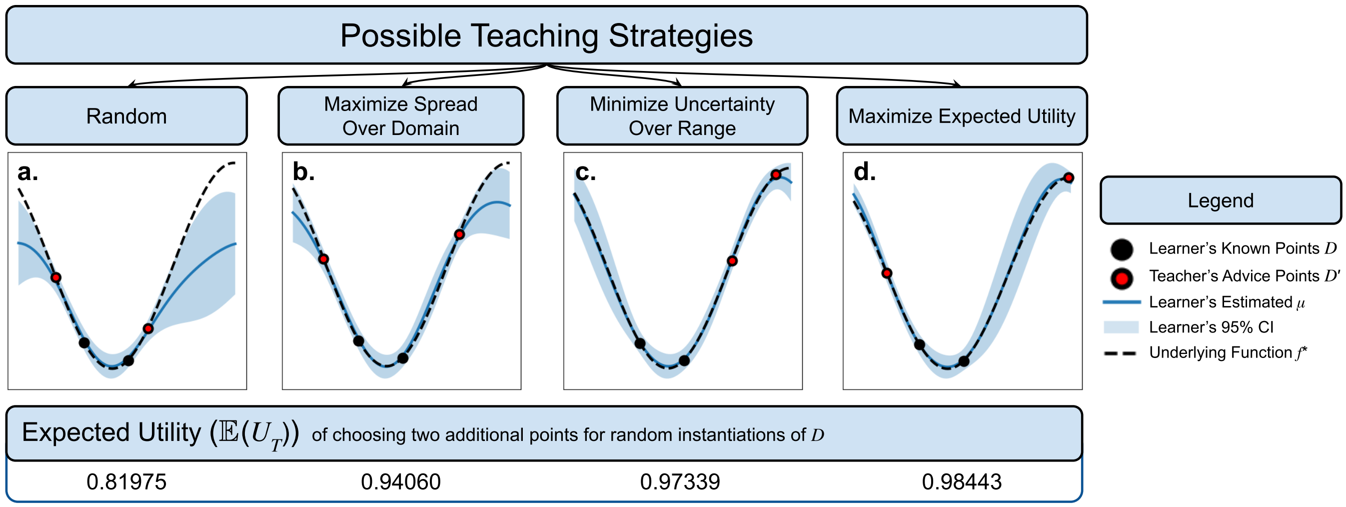

Figure 1: We selected one example of a sinusoidal target function (dotted line) and learners' already-observed points (black dots) to capture the results of three teacher strategies. Teacher choices for subsequent points are shown as red dots. The learner's resulting $\mathcal{GP}$ is represented by a mean function μ (blue line) and the 95% confidence interval (blue shading) given the kernel function k. The approximated expected utilities $\mathbb{E}(U_T)$ quantify the probability that each teaching policy chooses points that induce the right function f^* given a random instantiation of D. Note that because we specified 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.

$(x_{\text{midpt}}, y_{\text{midpt}})$ 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, when very few points can be taught, it could fail to capture important parts of the function where there are rapid, patterned, or unintuitive changes.

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 $x_{\text{uncert}} \in \mathbf{x}^*$ with the highest standard deviation and adds its coordinate $(x_{\text{uncert}}, y_{\text{uncert}})$ to D'. This is repeated n times, with each iteration updating $D \leftarrow D'$. Inferring areas where learners are most uncertain and selecting points to clarify is a reasonable pedagogical approach. But, there is a possibility that the selected points may neglect one section of the function (e.g., Figure 1c where μ is accurate but still has high uncertainty on most of the left side of the graph).

Maximizing Utility of D' The teacher performs approximate inference on the learner's representation of the function distribution. After simulating the learner's $\mathcal{GP}$, the teacher samples multiple candidate points and determines the $x_{\text{util}} \in \mathbf{x}^*$ with the highest utility and adds its coordinate $(x_{\text{util}}, y_{\text{util}})$ to D'. This is repeated n times, with each iteration updating $D \leftarrow D'$. Because the utility function is built in to this teacher, we expect that, given enough candidate samples to select from, it should be the optimal strategy, and we saw that its utility was significantly higher than random ($p = 0.02$) and domain spread ($p = 0.01$) strategies. Qualitatively, this was the only strategy which caused μ to realize the full amplitude of the function on the right side of the domain (Figure 1d).

Discussion

Gaussian processes can model how people generate informative teaching examples that support learning with little data, even in continuous spaces. We modeled strategies for teachers to generate points for learners, two of which used information about a learner's $\mathcal{GP}$ to make pedagogical choices. In future work, we plan on simulating teachers that, rather than assuming the learner's kernel, will instead *infer* the learner's kernel. This could give rise to new teaching strategies where teachers are sensitive to the learner's distribution over kernels. For instance, we would expect that teachers select points that correct the learner's $\mathcal{GP}$ away from a mistaken prior. 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. We will consider additional tractable alternative teaching heuristics, like prioritizing local minima and maxima coordinates. Perhaps a single heuristic is sufficient to capture human choices, but people could flexibly employ many heuristics that approximate inference of a learner's $\mathcal{GP}$ and inference of optimal points to teach. Further work could go beyond visualized functions, to understand how we teach functions that describe the behavior of complex causal variables. This task only considers teaching through giving example points, but mapping language inputs to some underlying strategies could broaden the scope of our model, since describing abstract underlying rules can sometimes be more clear than demonstrative examples (Sumers, Ho, Hawkins, & Griffiths, 2023). Ultimately, we hope that this paradigm can provide insight into computationally tractable methods for teaching and reasoning about complex, continuous domains.

References

Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2020, February). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*, *4*(22), 144–152. doi: 10.1038/s41562-019-0748-6

Gweon, H. (2021, October). Inferential social learning: cognitive foundations of human social learning and teaching. *Trends in Cognitive Sciences*, *25*(10), 896–910. doi: 10.1016/j.tics.2021.07.008

Lucas, C. G., Griffiths, T. L., Williams, J. J., & Kalish, M. L. (2015, October). A rational model of function learning. *Psychonomic Bulletin & Review*, *22*(5), 1193–1215. doi: 10.3758/s13423-015-0808-5

Rafferty, A. N., Brunskill, E., Griffiths, T. L., & Shafto, P. (2011). Faster teaching by pomdp planning. In G. Biswas, S. Bull, J. Kay, & A. Mitrovic (Eds.), *Artificial intelligence in education* (p. 280–287). Berlin, Heidelberg: Springer. doi: 10.1007/978-3-642-21869-9$_3$7

Rafferty, A. N., LaMar, M. M., & Griffiths, T. L. (2015). Inferring learners' knowledge from their actions. *Cognitive Science*, *39*(3), 584–618. doi: 10.1111/cogs.12157

Schulz, E., Tenenbaum, J. B., Duvenaud, D., Speekenbrink, M., & Gershman, S. J. (2017, December). Compositional inductive biases in function learning. *Cognitive Psychology*, *99*, 44–79. doi: 10.1016/j.cogpsych.2017.11.002

Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014, June). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, *71*, 55–89. doi: 10.1016/j.cogpsych.2013.12.004

Sumers, T. R., Ho, M. K., Hawkins, R. D., & Griffiths, T. L. (2023, March). Show or tell? exploring when (and why) teaching with language outperforms demonstration. *Cognition*, *232*, 105326. doi: 10.1016/j.cognition.2022.105326

Sumers, T. R., Ho, M. K., Hawkins, R. D., Narasimhan, K., & Griffiths, T. L. (2021, May). Learning rewards from linguistic feedback. *Proceedings of the AAAI Conference on Artificial Intelligence*, *35*(7), 6002–6010. doi: 10.1609/aaai.v35i7.16749