Informed Meta-Learning

Learning the inductive bias from human expert knowledge

Katarzyna Kobalczyk   Mihaela van der Schaar

Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, United Kingdom

Introduction

Informed machine-learning: from knowledge to inductive biases. Machine learning practitioners map expert knowledge to prior probabilities over functions by manual model design. Regions of the function space with a non-zero prior probability define the hypothesis space and the relative probabilities of prior solutions dictate the inductive biases of the model. Knowledge about different learning tasks is associated with distinct inductive biases.

Pros: Models with well-designed task-specific biases require less training data and are more robust.

Cons: Handcrafted biases require significant engineering efforts and need to be customized for every new learning problem.

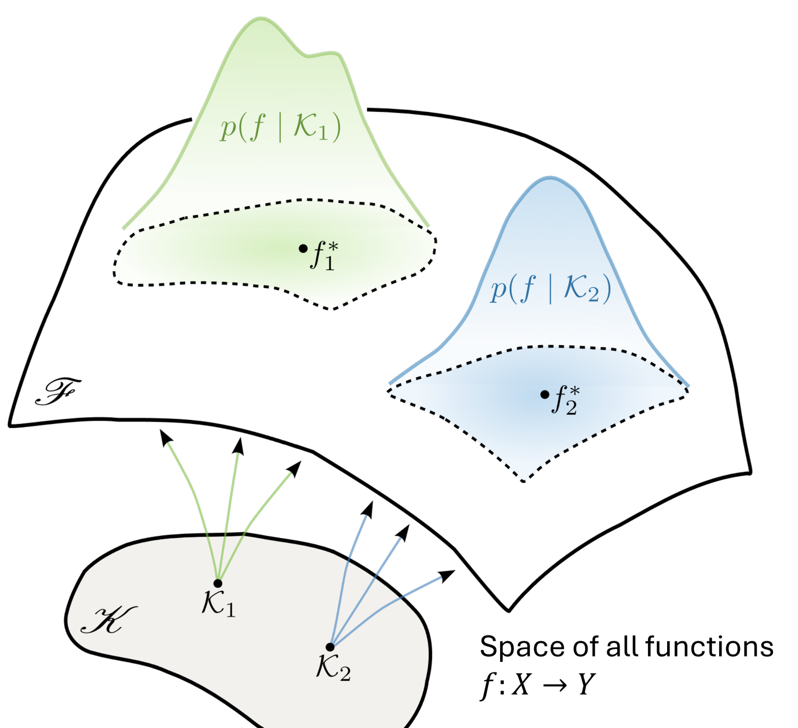

Meta-learning as inductive bias learning. Instead of manual inductive bias specification, the paradigm of meta-learning proposes that inductive biases are meta-learned by training over a distribution of related tasks.

Pros: Meta-learning enables flexible learning of complex inductive biases otherwise difficult to design by hand.

Cons: Success of meta-learners depends on the similarity of tasks at meta-train and test time. Heterogenous task distributions fail to provide strong inductive biases.

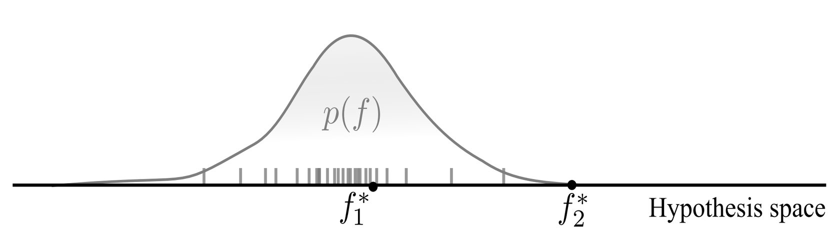

Concentrated meta distribution, provides strong inductive biases, so long as the task of interest belongs to the same environment as the tasks used during meta-training.

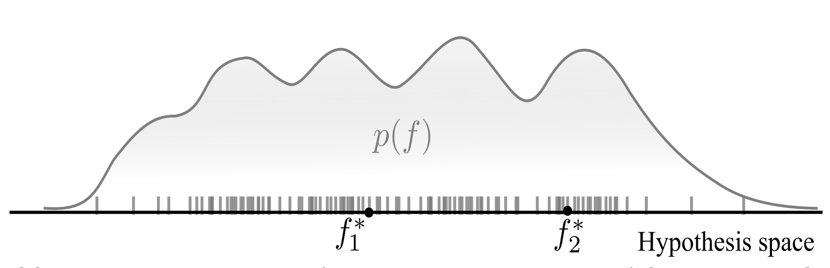

Heterogeneous environments span a wide range of functions, and thus support a wider range of tasks, at the cost of weaker inductive biases

Informed meta-learning: towards automatic inductive bias specification. We want to retain the flexibility of learning fine-grained and less formally stringent inductive biases and at the same time be able to guide the learner to the space of solutions agreeing with the prior knowledge of domain experts.

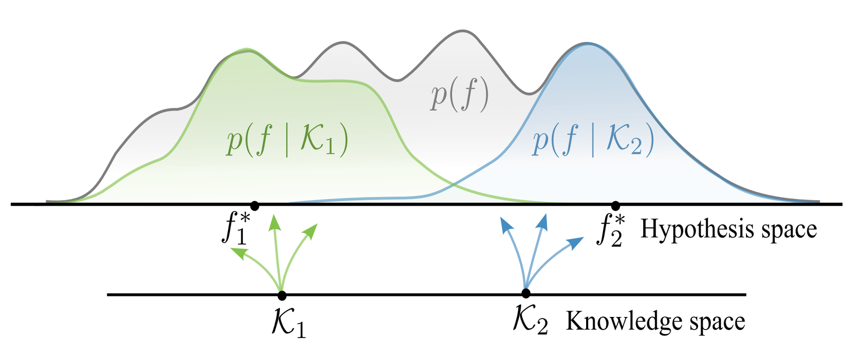

Informed meta-learning allows to condition the meta-distribution on the expert knowledge, concentrating the meta-distribution around the task-relevant regions.

In informed meta-learning the mapping from knowledge to inductive biases, $\mathcal{K} \to p(f \mid \mathcal{K})$, is meta-learned by training over multiple learning tasks and their representations of knowledge.

Method

Informed Neural Processes. We instantiate an informed meta-learner by extending the family of Neural Processes—probabilistic, fully amortized meta-learners.

Model. (I)NPs model the distribution over functions through a fixed dimensional latent variable z sampled from a variational distribution q. INPs extend NPs by additionally conditioning q on the information contained in expert knowledge. We model the predictive posterior distribution as:

$$p(y \mid x, \mathcal{D}_C, \mathcal{K}) := \int p(y \mid x, z) q(z \mid \mathcal{D}_C, \mathcal{K}) dz.$$

Training. INPs are trained in an episodic fashion over a distribution of learning task $\mathcal{T}_j = \{\mathcal{D}_C, \mathcal{D}_T\}$ and their associated knowledge representations $\mathcal{K}_j$. Denoting by r_C and r_T the context and target data representations and by k the knowledge embedding vector of a single task, parameters of the model are learned by maximising the expectation of ELBO over all training tasks,

$$\log p(y_T \mid x_T, r_C, k) \geq \mathbb{E}_{q(z \mid r_T, k)} \left[\log p(y_T \mid x_T, z)\right] - D_{\mathrm{KL}}\left(q(z \mid r_T, k) \mid\mid q(z \mid r_C, k)\right)$$

Results

Data efficiency, robustness and uncertainty.

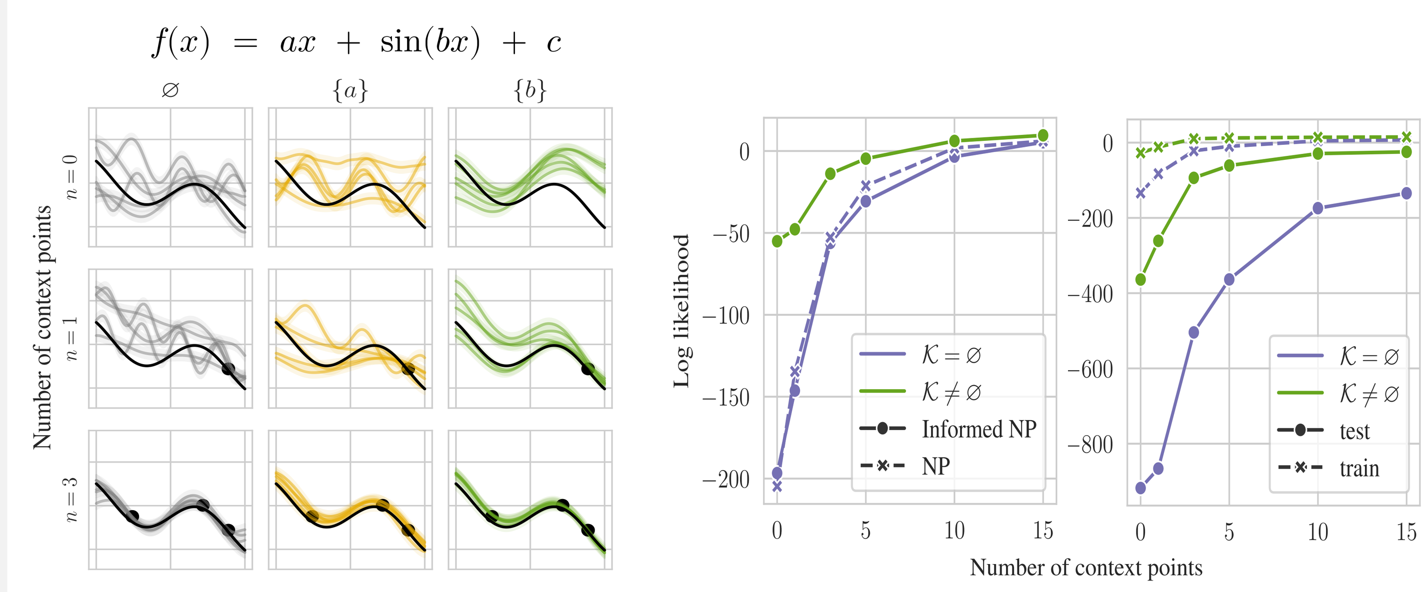

On a series of controlled synthetic experiments, we:

- Demonstrate how successful integration of oracle knowledge about the data generating process **improves data efficiency** and **mitigate the adverse effects of train-test distribution shift.**
- We derive methods to quantitatively measure the **impact of expert knowledge on the reduction in epistemic uncertainty.**

From natural language to functional priors. In practical scenarios, predictive functions are difficult to model with closed-form mathematical expressions. The benefit of meta-learned prior is their **functional flexibility**. We show how INPs learn to map descriptions in **natural language** to functional priors.

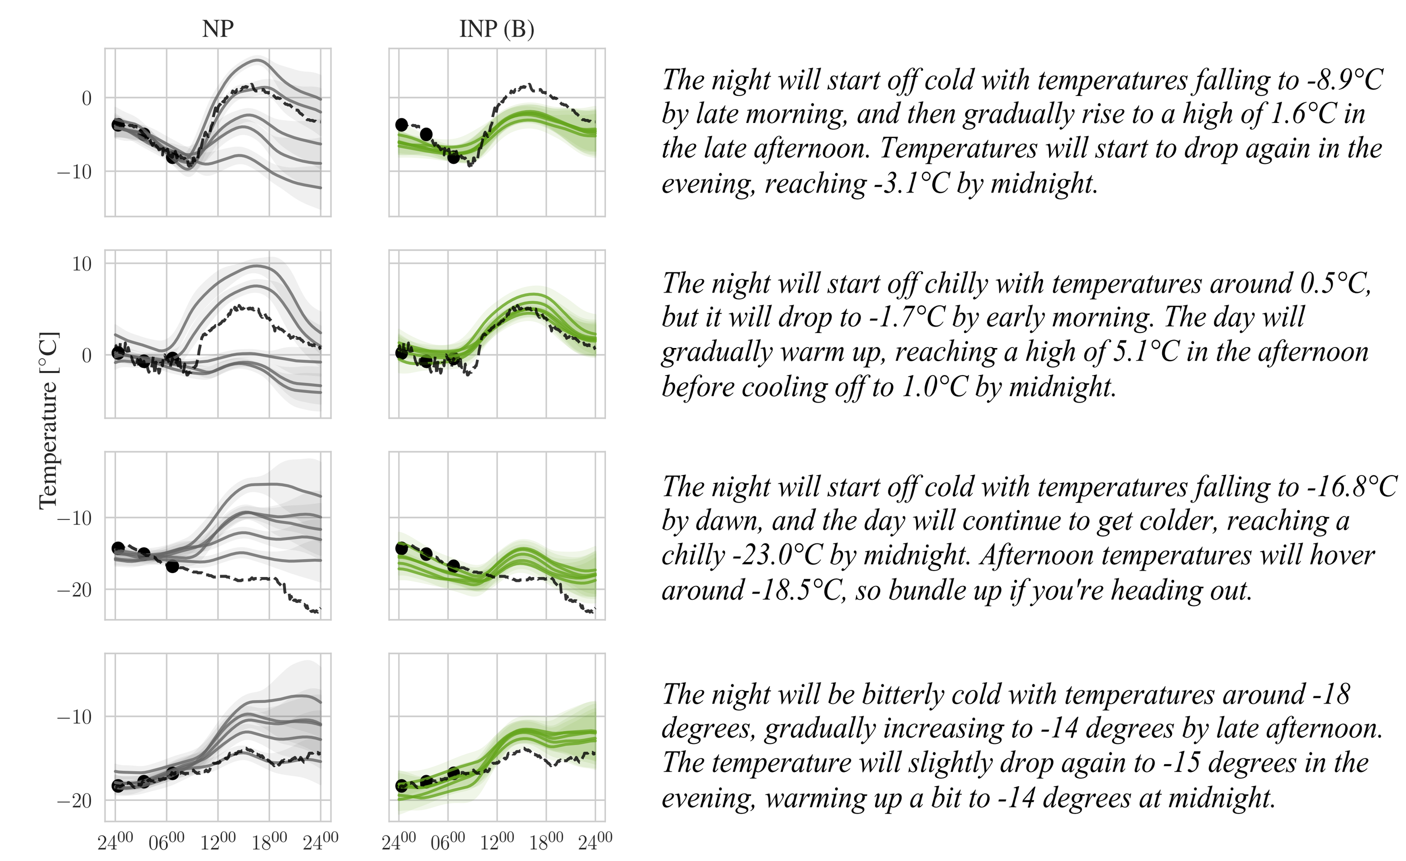

Beyond 1D regression. We extend INPs to the task of few-shot classification with expert knowledge in the form class-level features or descriptions in natural language. Successful alignment of knowledge and data representations facilitates robust generalization to new, previously unseen classes, zero-shot classification and improved few-shot classification accuracy.

Sample Images	Sample image captions	GPT-generated class description
	1. *A large bird with a white belly, black and white wings with a long beak.* 2. *This bird is white and grey in color with a curved beak, and black eye rings.* 3. *A large bird with a white belly and face, black back and wings, and peach bill.* 4. *Bird has gray body feathers, white breast feather, and long beak* 5. *A medium sised bird with black wings, and a bill that curves downwards*	*This bird breed is a medium to large size, characterised by its grey body feathers, contrasting white belly and face, black back and wings, distinctive black eye rings, and a long, downward-curving peach bill.*
	1. *This big bird has a sharp beak and has black covering its body.* 2. *An all black bird with a distinct thick, rounded bill.* 3. *This entirely black bird has long and wide rectrices relative to the size of its body.* 4. *A black bird with a long tail and large beak.* 5. *This black bird has sparse plumage and a thick brown beak.*	*This bird breed is large and entirely black with sparse plumage, characterised by its thick brown beak, long tail, and wide rectrices relative to its body size.*

Read the full paper: