Meta-Learning Hypothesis Spaces for Sequential Decision-making

Parnian Kassraie, Jonas Rothfuss, Andreas Krause



ETH zürich
Learning &
Adaptive Systems

Motivation: Sequential Decision-making

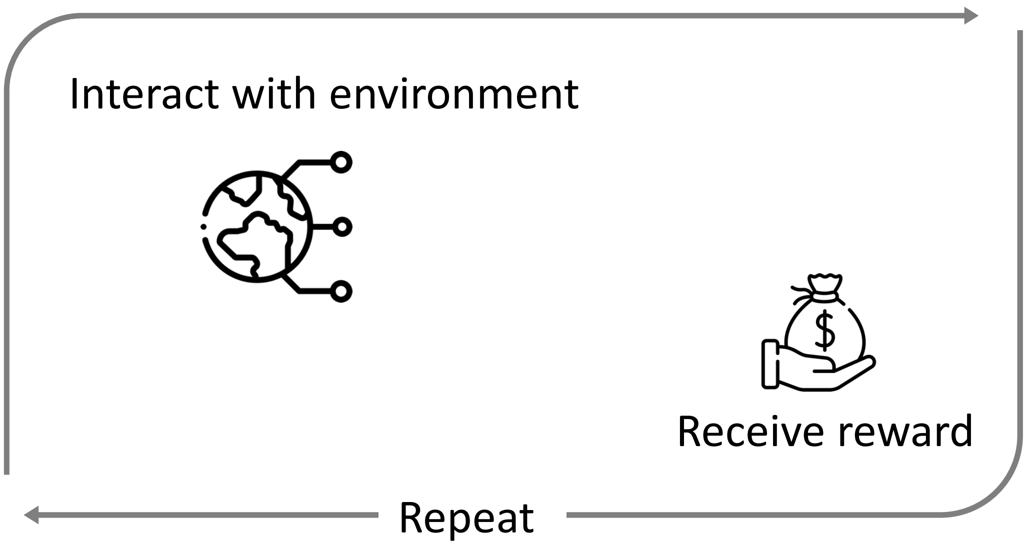

Sequential Decision Problems:

- Bandits / Bayesian Optimization
- Active Learning
- Model-based RL

Motivation: Sequential Decision-making

Interact with environment

Receive reward

Repeat

Sequential Decision Problems:

- Bandits / Bayesian Optimization
- Active Learning
- Model-based RL

explore

exploit

Confidence sets are great for guiding explorations!

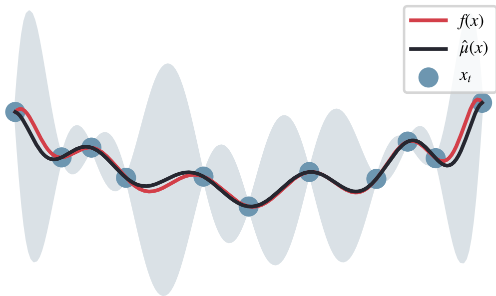

width ⟷ current uncertainty

center ⟷ current knowledge

Hypothesis Spaces and Confidence Sets

Invalid

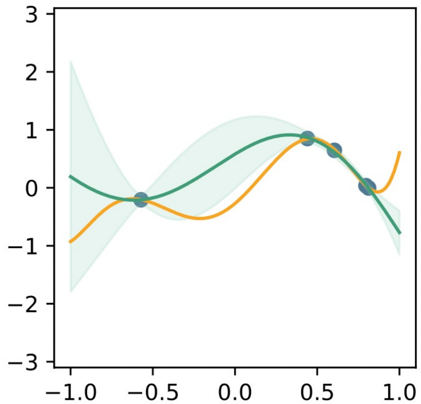

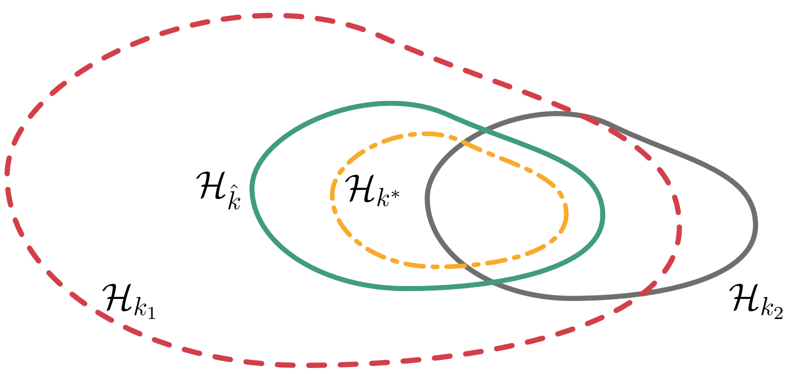

$\mathcal{C}_{t-1}(k_2; \boldsymbol{x})$ Invalid

$\mathcal{C}_{t-1}(k^*; \boldsymbol{x})$ True sets (Valid)

Hypothesis Spaces and Confidence Sets

Invalid   Valid but too wide

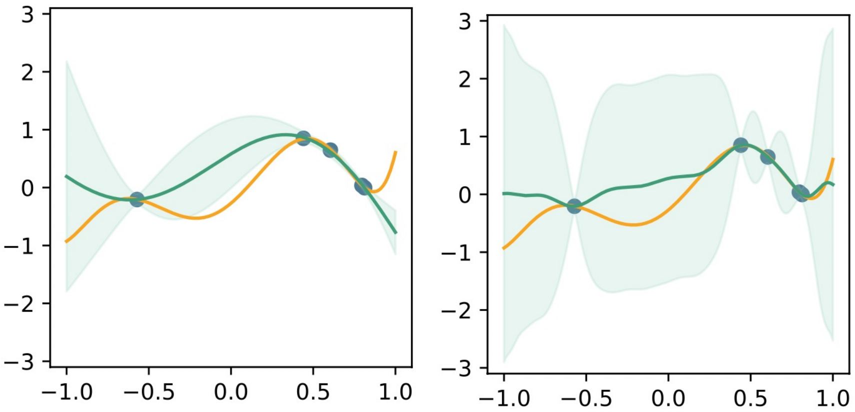

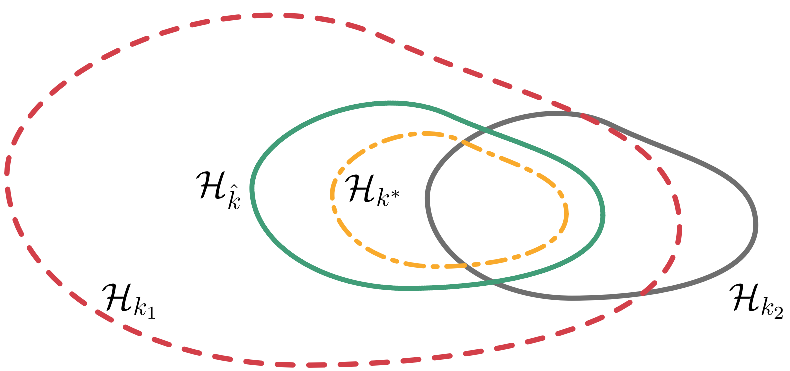

$\mathcal{C}_{t-1}(k_2; \boldsymbol{x})$ Invalid

$\mathcal{C}_{t-1}(k_1; \boldsymbol{x})$ Valid but too wide

$\mathcal{C}_{t-1}(k^*; \boldsymbol{x})$ True sets (Valid)

Hypothesis Spaces and Confidence Sets

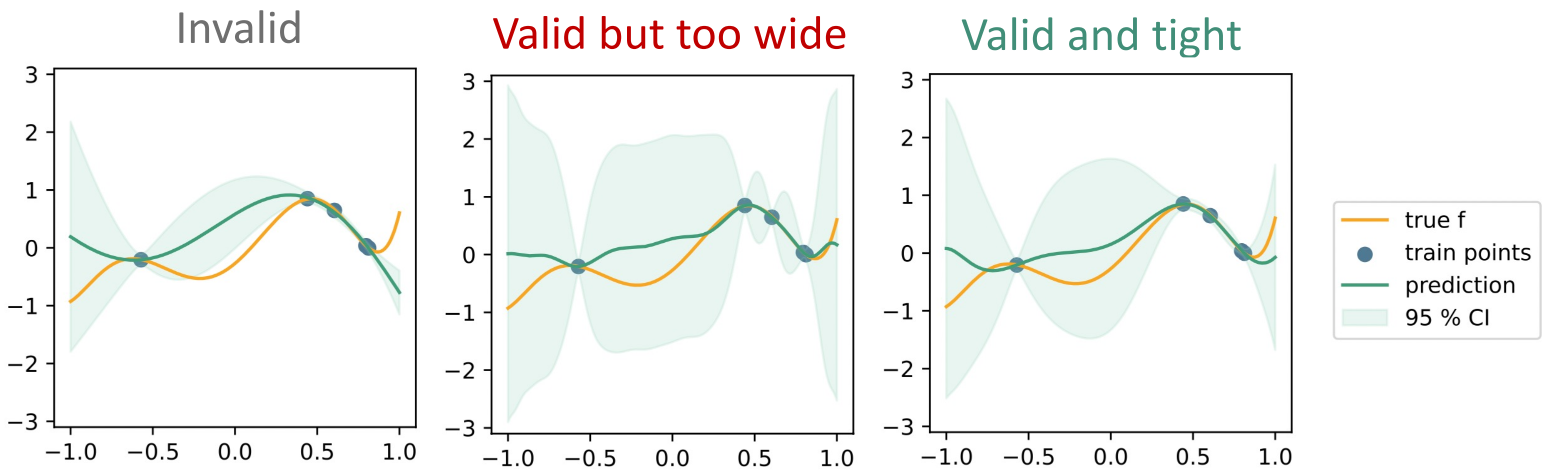

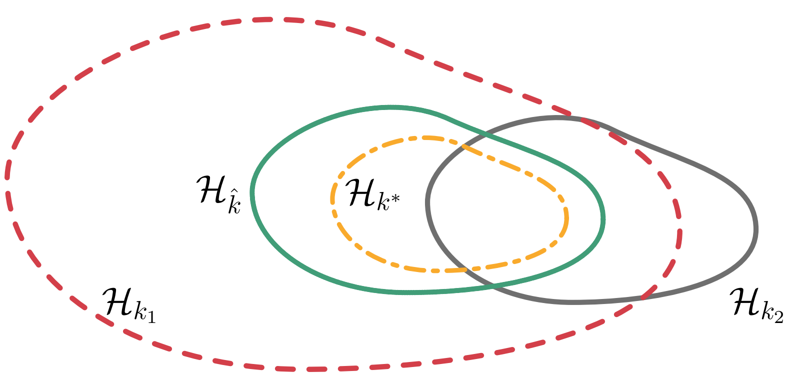

$\mathcal{C}_{t-1}(k_2; \boldsymbol{x})$ Invalid

$\mathcal{C}_{t-1}(k_1; \boldsymbol{x})$ Valid but too wide

$\mathcal{C}_{t-1}(\hat{k}; \boldsymbol{x})$ Valid and tight

$\mathcal{C}_{t-1}(k^*; \boldsymbol{x})$ True sets (Valid)

Hypothesis Spaces and Confidence Sets

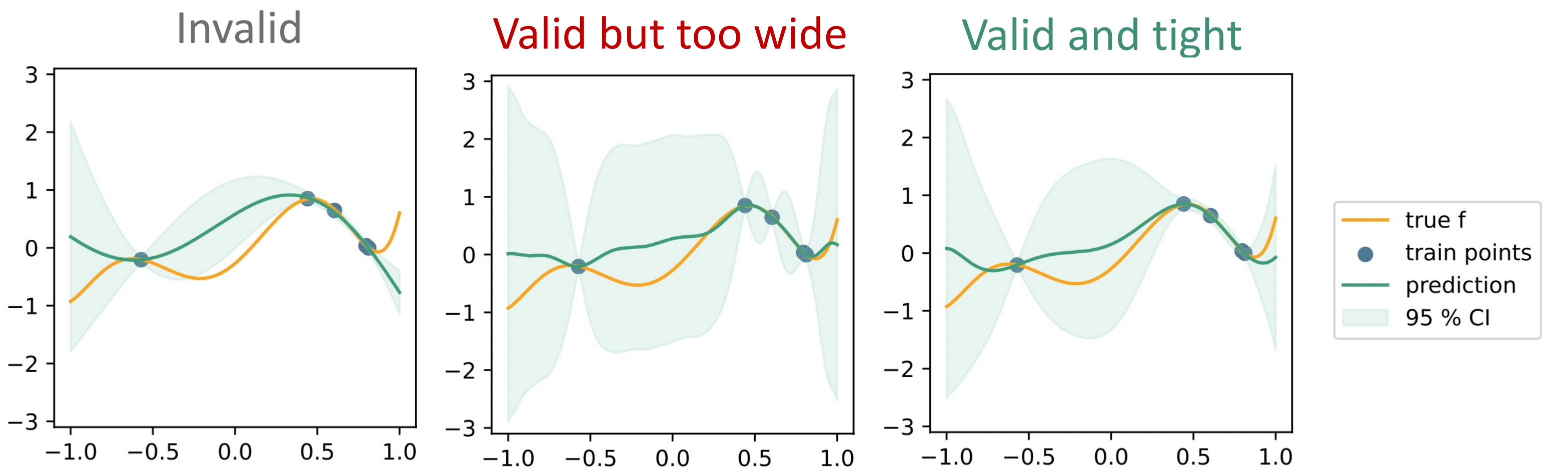

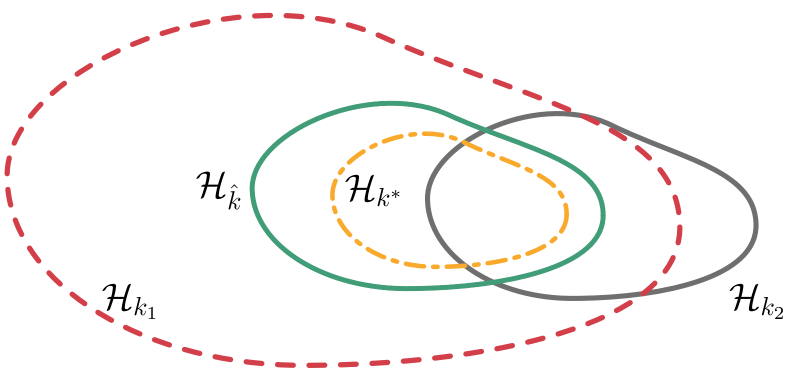

$\mathcal{C}_{t-1}(k_2; \boldsymbol{x})$ Invalid

$\mathcal{C}_{t-1}(k_1; \boldsymbol{x})$ Valid but too wide

$\mathcal{C}_{t-1}(\hat{k}; \boldsymbol{x})$ Valid and tight

$\mathcal{C}_{t-1}(k^*; \boldsymbol{x})$ True sets (Valid)

How can we find a good $\mathcal{H}_{\hat{k}}$?

Our approach: Meta-learning $\mathcal{H}_{k^*}$

Data from similar tasks

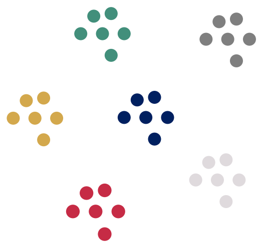

Our approach: Meta-learning $\mathcal{H}_{k^*}$

Data from similar tasks

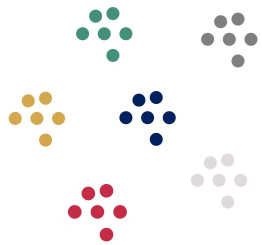

Pool of candidate kernels / features

$k_1(\boldsymbol{x}, \boldsymbol{x}')$ $k_3(\boldsymbol{x}, \boldsymbol{x}')$ $k_5(\boldsymbol{x}, \boldsymbol{x}')$

$k_2(\boldsymbol{x}, \boldsymbol{x}')$ $k_4(\boldsymbol{x}, \boldsymbol{x}')$

(True) kernel has sparse additive structure

$$k^*(\boldsymbol{x}, \boldsymbol{x}') = \sum_{j=1}^{p} \eta_j^* k_j(\boldsymbol{x}, \boldsymbol{x}')$$

[this holds for all Mercer kernels]

Our approach: Meta-learning $\mathcal{H}_{k^*}$

Data from similar tasks

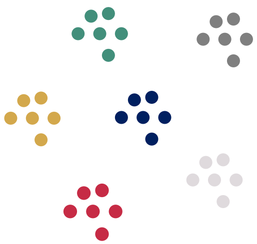

Pool of candidate kernels / features

$k_1(\boldsymbol{x}, \boldsymbol{x}')$ $k_3(\boldsymbol{x}, \boldsymbol{x}')$ $k_5(\boldsymbol{x}, \boldsymbol{x}')$ $k_2(\boldsymbol{x}, \boldsymbol{x}')$ $k_4(\boldsymbol{x}, \boldsymbol{x}')$

(True) kernel has sparse additive structure

$$k^*(\boldsymbol{x}, \boldsymbol{x}') = \sum_{j=1}^{p} \eta_j^* k_j(\boldsymbol{x}, \boldsymbol{x}')$$

[this holds for all Mercer kernels]

Meta-KeL

Shrinks hypothesis space by eliminating kernels / features that are w.h.p. not active in the meta-training tasks

$\hat{k}(\boldsymbol{x}, \boldsymbol{x}')$

Can be reduced to a group lasso problem!

Properties of the meta-learned kernel

Theorem (Informal)

Under mild regularity assumptions on the meta-data, with probability greater than $1-\delta$,

- *The meta-learned hypothesis space contains all relevant hypotheses, i.e.,* $\mathcal{H}_{k^*} \subseteq \mathcal{H}_{\hat{k}}$
- *The confidence sets of* $\hat{k}$ *are calibrated, i.e.,* $\forall f \in \mathcal{H}_{k^*}$:

$$\mathbb{P}\left(\forall \boldsymbol{x} \in \mathcal{X}, \forall t \geq 1: \; f(\boldsymbol{x}) \in \mathcal{C}_{t-1}(\hat{k}; \boldsymbol{x})\right) \geq 1-\delta.$$

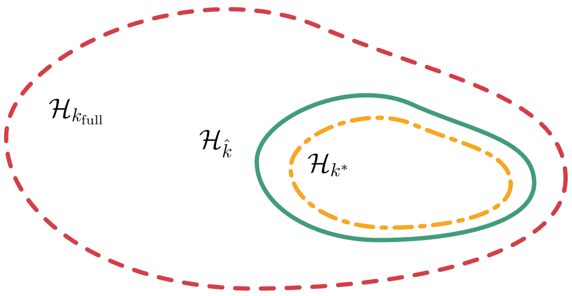

Properties of the meta-learned kernel

Theorem (Informal)

Under mild regularity assumptions on the meta-data, with probability greater than $1 - \delta$,

- *The meta-learned hypothesis space contains all relevant hypotheses, i.e., $\mathcal{H}_{k^*} \subseteq \mathcal{H}_{\hat{k}}$*
- *The confidence sets of $\hat{k}$ are calibrated, i.e., $\forall f \in \mathcal{H}_{k^*}$:*

$$\mathbb{P}\left(\forall \boldsymbol{x} \in \mathcal{X}, \forall t \geq 1 : f(\boldsymbol{x}) \in \mathcal{C}_{t-1}(\hat{k}; \boldsymbol{x})\right) \geq 1 - \delta.$$

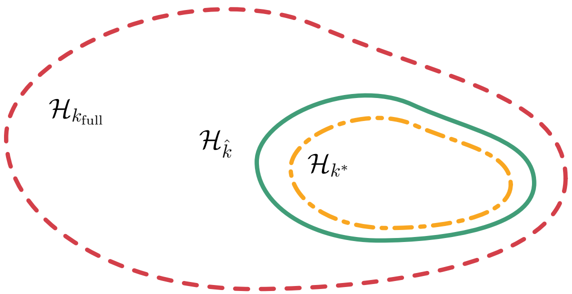

+ The meta-learned confidence bounds approach the oracle bounds as the amount of meta-training data grows

Application: Bayesian Optimization

f is the objective function of a BO problem.

Regret $R_T = \sum_{t=1}^{T} [f(\boldsymbol{x}^*) - f(\boldsymbol{x}_t)]$ Goal $R_T/T \to 0 \text{ as } T \to \infty$

Application: Bayesian Optimization

f is the objective function of a BO problem.

Regret $R_T = \sum_{t=1}^{T} [f(\boldsymbol{x}^*) - f(\boldsymbol{x}_t)]$ Goal $R_T/T \to 0 \text{ as } T \to \infty$

GP-UCB Policy: [GP-UCB, Srinivas et al.]

Corollary

Provided that there is enough meta-data,

- *The learner achieves sublinear regret, w.h.p.*
- *This guarantee is tight compared to the one for the Oracle learner, and approaches it at a $\mathcal{O}(1/\sqrt{mn})$ rate.*

Poster session:
Thu 21 Jul 6 p.m. EDT — 8:30 p.m. EDT

Meta-Learning Hypothesis Spaces for Sequential Decision-making

Parnian Kassraie, Jonas Rothfuss, Andreas Krause



ETH zürich
Learning & Adaptive Systems