Transductive Active Learning

with Application to Safe Bayesian Optimization

Jonas Hübotter, Bhavya Sukhija, Lenart Treven, Yarden As, Andreas Krause

Paper Code

ETH zürich

Background

- Active learning is a powerful paradigm for data selection that commonly aims to learn f globally on $\mathcal{X}$
- In many real-world problems,
  - *i* the domain is so large that learning f globally is hopeless; or
  - *ii* agents have limited information / access to $\mathcal{X}$

How can we find agents that solve tasks efficiently by learning in a *directed* manner and *extrapolating* beyond their limited information?

Transductive Active Learning

"only learn what is needed"

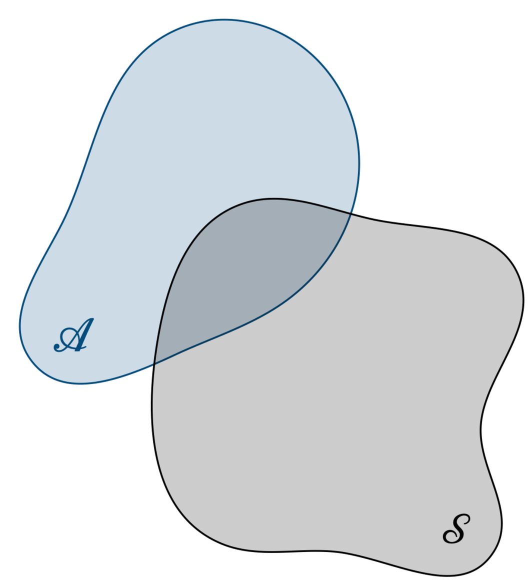

- Sample space $\mathcal{S} \subseteq \mathcal{X}$
- Target space $\mathcal{A} \subseteq \mathcal{X}$
- Unknown function f over $\mathcal{X}$

Goal: Learn f within $\mathcal{A}$ by sampling from $\mathcal{S}$

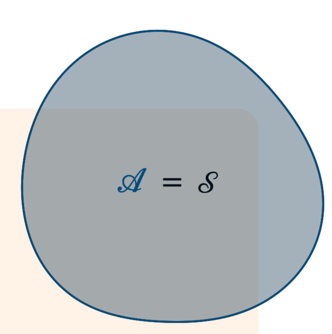

(Inductive) Active Learning
"learn as much as you can"

⇝ TAL generalizes AL to goal-orientation ($\mathcal{A}$) and extrapolation ($\mathcal{S}$)

- We model f by a Gaussian process
- The marginal variance $\sigma_n^2(\boldsymbol{x})$ of the GP after n samples is a proxy for the approximation error at $\boldsymbol{x}$

New goal: Reduce uncertainty $\sigma_n^2(\boldsymbol{x})$ at $\boldsymbol{x} \in \mathcal{A}$

An Algorithmic Framework for TAL

Proposal: Select samples to minimize posterior "uncertainty" within $\mathcal{A}$

⇝ Quantifying "uncertainty" by entropy yields **ITL**:

$$\boldsymbol{x}_n = \operatorname*{arg\,min}_{\boldsymbol{x} \in \mathcal{S}} \mathrm{H}[\boldsymbol{f}(\mathcal{A}) \mid D_{n-1}, (\boldsymbol{x}, f(\boldsymbol{x}) + \varepsilon)] = \operatorname*{arg\,max}_{\boldsymbol{x} \in \mathcal{S}} \mathrm{I}(\boldsymbol{f}(\mathcal{A}); (\boldsymbol{x}, f(\boldsymbol{x}) + \varepsilon) \mid D_{n-1})$$

⇝ Quantifying "uncertainty" by variance yields **VTL**:

$$\boldsymbol{x}_n = \operatorname*{arg\,min}_{\boldsymbol{x} \in \mathcal{S}} \operatorname{tr} \mathrm{Var}[\boldsymbol{f}(\mathcal{A}) \mid D_{n-1}, (\boldsymbol{x}, f(\boldsymbol{x}) + \varepsilon)]$$

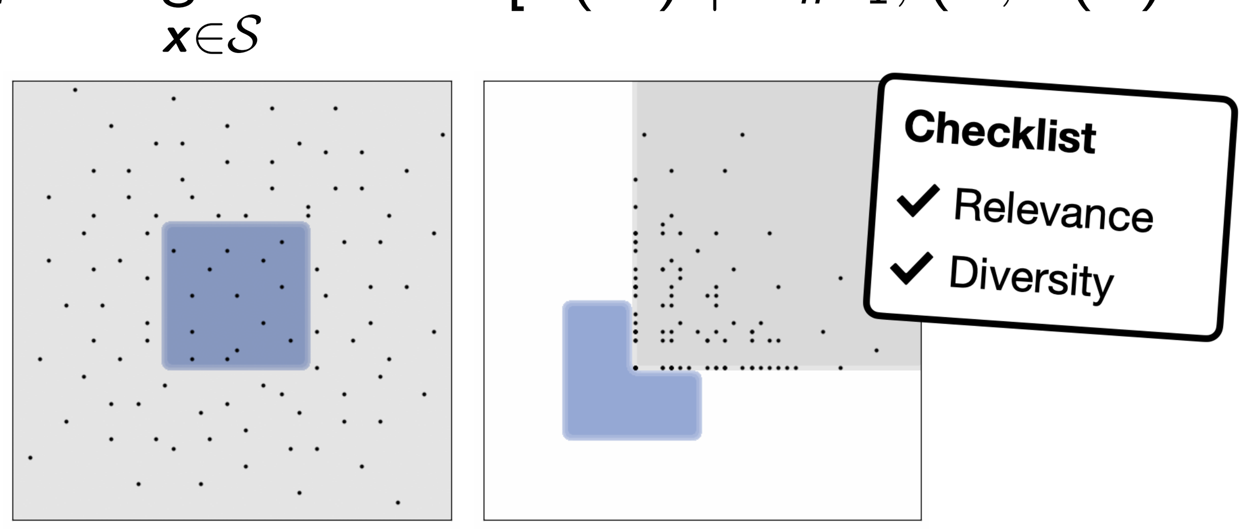

RBF kernel

Theory: Convergence Guarantees for ITL & VTL

How much can be learned about $\mathcal{A}$ from $\mathcal{S}$?

Generalization bound. For every $\boldsymbol{x} \in \mathcal{A}$:

$$\sigma_n^2(\boldsymbol{x}) \leq \underbrace{\mathrm{Var}[f(\boldsymbol{x}) \mid \boldsymbol{f}(\mathcal{S})]}_{\text{irreducible}} + \underbrace{C \log(n)/\sqrt{n}}_{\text{reducible}}$$

Approximation error bound. If $f \in \mathcal{H}_k(\mathcal{X})$ then for every $\boldsymbol{x} \in \mathcal{A}$ with probability $1 - \delta$:

$$|f(\boldsymbol{x}) - \mu_n(\boldsymbol{x})|^2 \leq \beta_n^2(\delta) \left[\text{irreducible} + C \log(n)/\sqrt{n}\right]$$

where $\mu_n(\boldsymbol{x})$ is the prediction and $\beta_n(\delta)$ the CI width

Example: Safe Bayesian Optimization

Task: Under constraint $c^\star$ inducing *true* safe set $\mathcal{S}^\star = \{\boldsymbol{x} \mid c^\star(\boldsymbol{x}) \geq 0\}$, find $\arg\max_{\boldsymbol{x} \in \mathcal{S}^\star} f^\star(\boldsymbol{x})$.

Calibrated model:

- $l_n^f(\boldsymbol{x}) \leq f^\star(\boldsymbol{x}) \leq u_n^f(\boldsymbol{x})$
- $l_n^c(\boldsymbol{x}) \leq c^\star(\boldsymbol{x}) \leq u_n^c(\boldsymbol{x})$

Estimated safe sets:

- $\mathcal{S}_n = \{\boldsymbol{x} \mid l_n^c(\boldsymbol{x}) \geq 0\}$
- $\hat{\mathcal{S}}_n = \{\boldsymbol{x} \mid u_n^c(\boldsymbol{x}) \geq 0\}$

⇝ $\mathcal{S}_n \subseteq \mathcal{S}^\star \subseteq \hat{\mathcal{S}}_n$

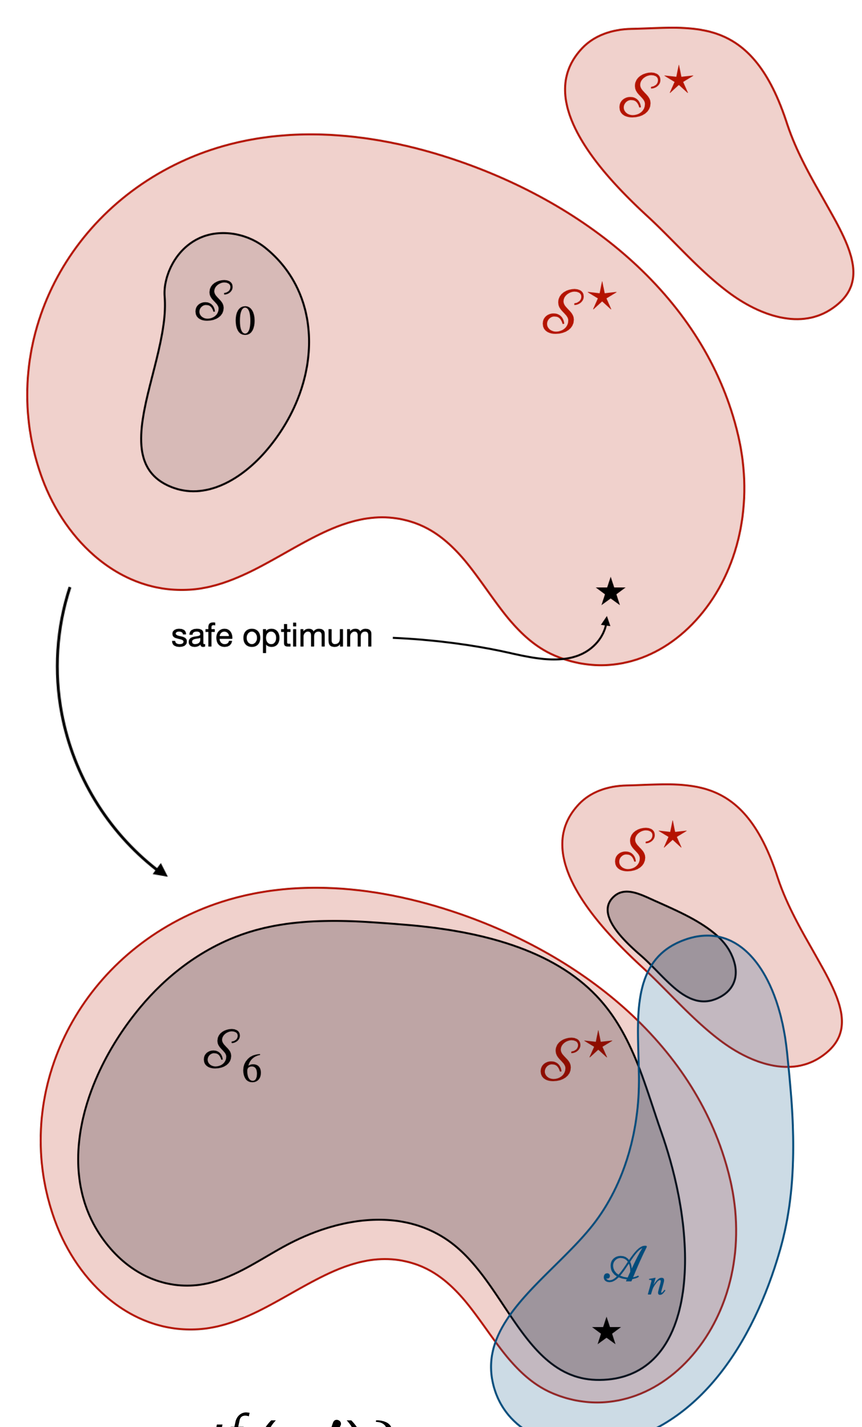

Learn *potential maximizers*
$\mathcal{A}_n = \{\boldsymbol{x} \in \hat{\mathcal{S}}_n \mid u_n^f(\boldsymbol{x}) \geq \max_{\boldsymbol{x}' \in \mathcal{S}_n} l_n^f(\boldsymbol{x}')\}$

Applying TAL: If $f^\star, c^\star$ are sufficiently regular, ITL & VTL find the safe reachable optimum.

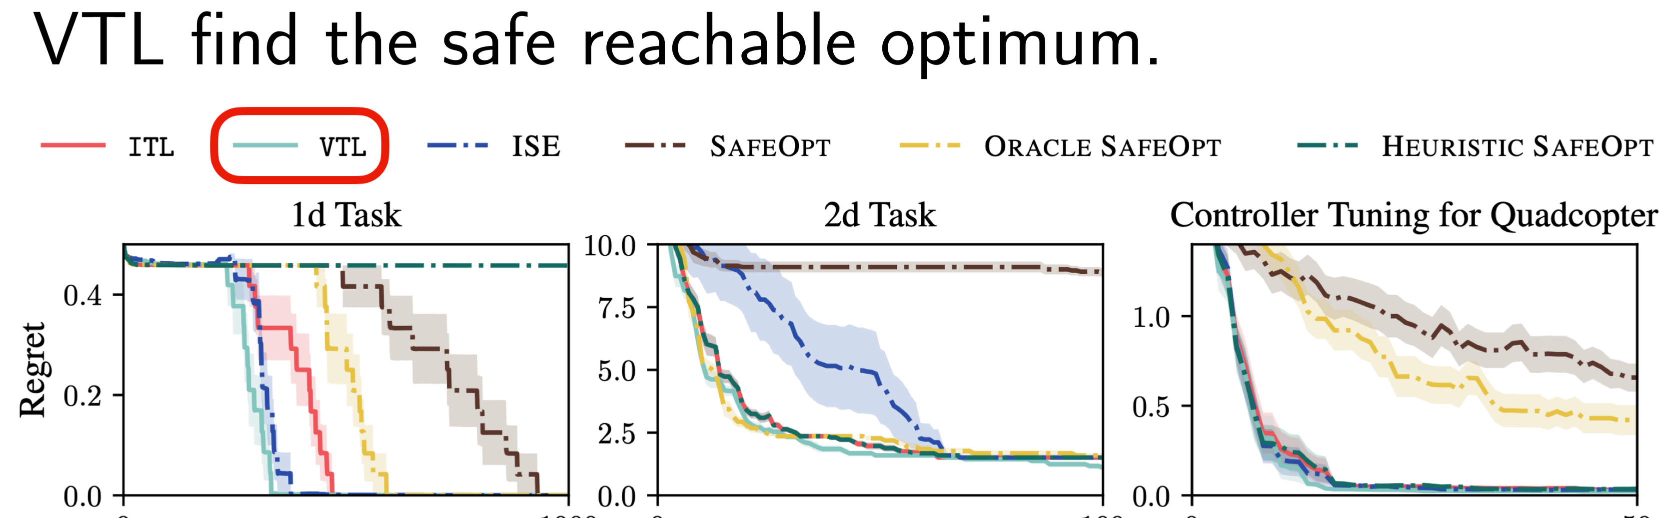

⇝ Framing Safe BO as TAL allows retrieving only the information that is needed to find the safe optimum

Key Takeaways

- Transductive Active Learning is a powerful paradigm for learning under resource constraints such as limited interaction time & limited access
- TAL is widely applicable beyond Safe BO, e.g., in active fine-tuning of large NNs: