DeepMind

Model-Value Inconsistency

Angelos Filos*, Eszter Vértes*, **Zita Marinho***, Gregory Farquhar, Diana Borsa, Abram Friesen, Feryal Behbahani, Tom Schaul, André Barreto, Simon Osindero

ICML 2022

Key Observation

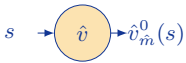

Key Observation

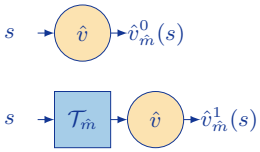

Key Observation

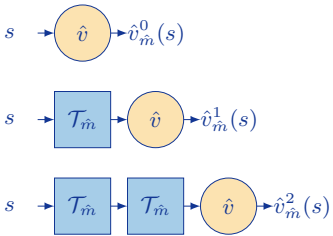

Key Observation

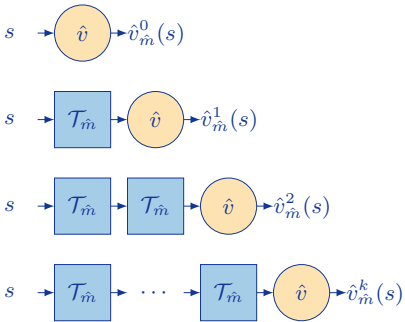

Figure 1: Ensemble value predictions from a single model $\hat{m}$ and value function $\hat{v}$.

Key Observation

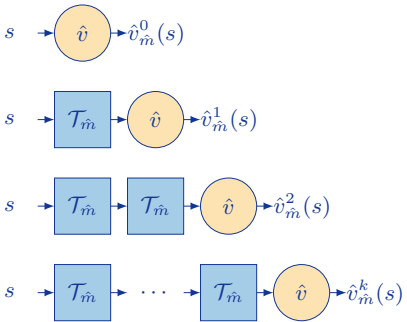

Figure 1: Ensemble value predictions from a single model $\hat{m}$ and value function $\hat{v}$.

Implicit Value Ensemble (IVE)

$$\{\hat{v}^i_{\hat{m}}\}_{i=0}^{k} \triangleq \underbrace{\{\hat{v}, \mathcal{T}^{\pi}_{\hat{m}}\hat{v}, \ldots, (\mathcal{T}^{\pi}_{\hat{m}})^k\hat{v}\}}_{k+1 \text{ value predictions}}$$

Key Observation

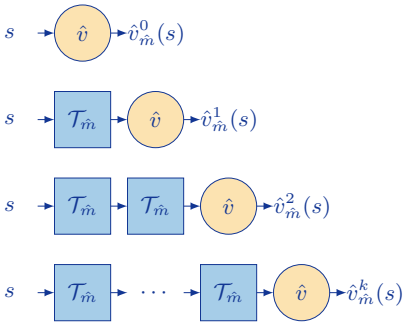

Figure 1: Ensemble value predictions from a single model $\hat{m}$ and value function $\hat{v}$.

Implicit Value Ensemble (IVE)

$$\{\hat{v}^i_{\hat{m}}\}_{i=0}^k \triangleq \underbrace{\{\hat{v}, \mathcal{T}^\pi_{\hat{m}}\hat{v}, \ldots, (\mathcal{T}^\pi_{\hat{m}})^k\hat{v}\}}_{k+1 \text{ value predictions}}$$

Model-Value Inconsistency

Implicit Value Ensemble disagreement, e.g.,

$$\sigma\text{-IVE}(k) \triangleq \texttt{std}\{\hat{v}^i_{\hat{m}}\}_{i=0}^k$$

tl;dr

Research Goal

Determine how effective model-value inconsistency is as a signal for epistemic uncertainty.

tl;dr

Research Goal

Determine how effective model-value inconsistency is as a signal for epistemic uncertainty.

Key Takeaway

"Your model-based agent is secretly a performing ensemble of value functions"

Value Uncertainty Quantification

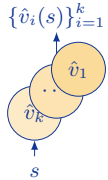

(a) Value Ensemble (EVE)

Value Uncertainty Quantification

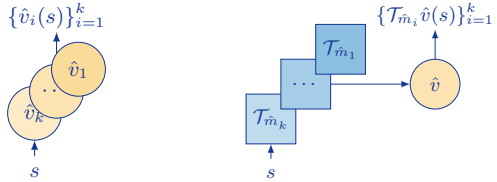

(a) Value Ensemble (EVE) **(b)** Model Ensemble (EMVE)

Value Uncertainty Quantification

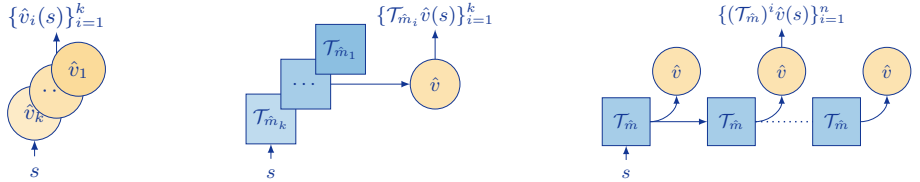

(a) Value Ensemble (EVE) **(b)** Model Ensemble (EMVE) **(c)** Implicit Value Ensemble (IVE)

Figure 2: Value uncertainty quantification in scalable epistemic uncertainty-aware RL agents.

Value Uncertainty Quantification

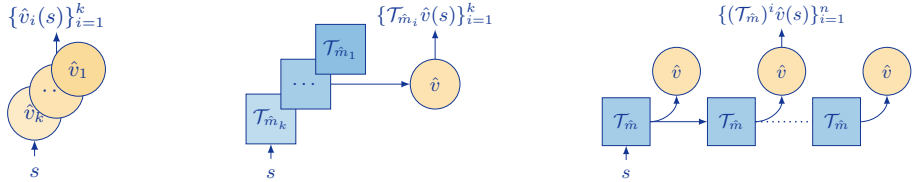

(a) Value Ensemble (EVE) **(b)** Model Ensemble (EMVE) **(c)** Implicit Value Ensemble (IVE)

Figure 2: Value uncertainty quantification in scalable epistemic uncertainty-aware RL agents.

Key Difference

IVE relies on *single point estimates* of a learned model $\hat{m}$ and a value function $\hat{v}$.

IVE as proxy for Epistemic Uncertainty

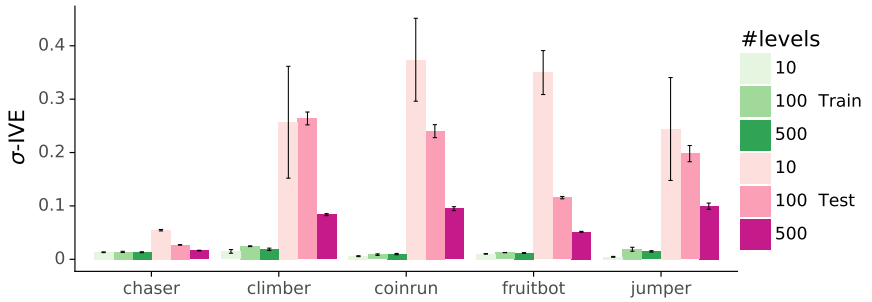

Figure 3: σ-IVE(5) computed using the model of the Muesli agent while evaluating on both training (green) and unseen test levels (pink). Bars, error-bars show mean and ste across 3 seeds.

Conclusion

Investigate the qualitative and quantitative properties of IVE:

- Disagreement among IVE members is an effective signal for epistemic uncertainty
- it is a good signal to guide exploration
- IVEs can robustify value-based planning, and help deal with distribution shifts.

Thank you!

This is joint work with these wonderful collaborators.


Angelos Filos


Eszter Vértes


Zita Marinho


Greg Farquhar


Diana Borsa


Abram Friesen


Feryal Behbahani


Tom Schaul


André Barreto


Simon Osindero

Check our paper for more details or reach out to
(afilos, evertes, zmarinho) @deepmind.com