Adaptive Labeling for Efficient Out-of-distribution Model Evaluation

NeurIPS 2024

Daksh Mittal Yuanzhe Ma Shalmali Joshi Hongseok Namkoong

Columbia University

Rigorous Empirical Evaluation Forms the Basis of Engineering Progress

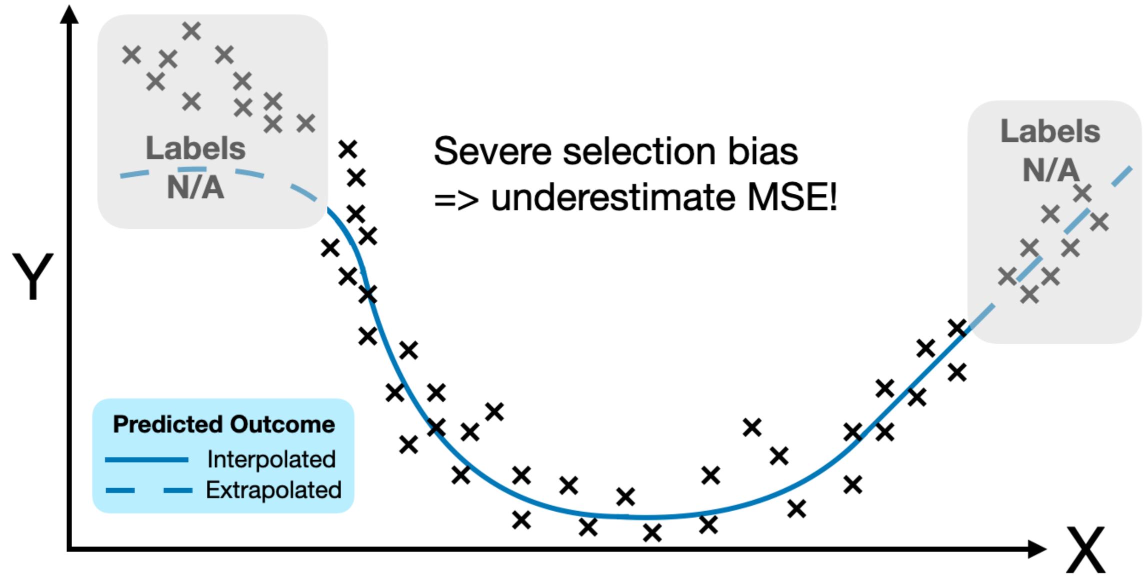

- Initial available data $\mathcal{D}^0 = (\mathcal{X}^0, \mathcal{Y}^0)$ – suffers selection bias
- AI Model $\psi(\cdot)$ trained and evaluated using dataset $\mathcal{D}^0$
- Distribution seen during deployment P_X
- Naïve evaluation using $\mathcal{D}^0$ fails to capture performance across P_X
- Ground truth labels are costly

Problem : Efficiently evaluate model $\psi(\cdot)$ on P_X while acquiring minimal number of labels

MDP for Adaptive Labeling

- $Y = f^*(X) + Noise$, f^*– Unknown
- **Bayesian framework :** Prior μ over f
- Sequentially acquire data in batches (and update the beliefs over f)

States – Posterior beliefs over f that is $\mu_t = \mu(\cdot\,|\mathcal{D}^{0:t})$

Actions – Batch $\mathcal{X}^t$ selected to be labeled in time period t

Reward/Cost – Minimize Variance of MSE of model at end of horizon T

MDP

MDP for Adaptive Labeling – three critical components

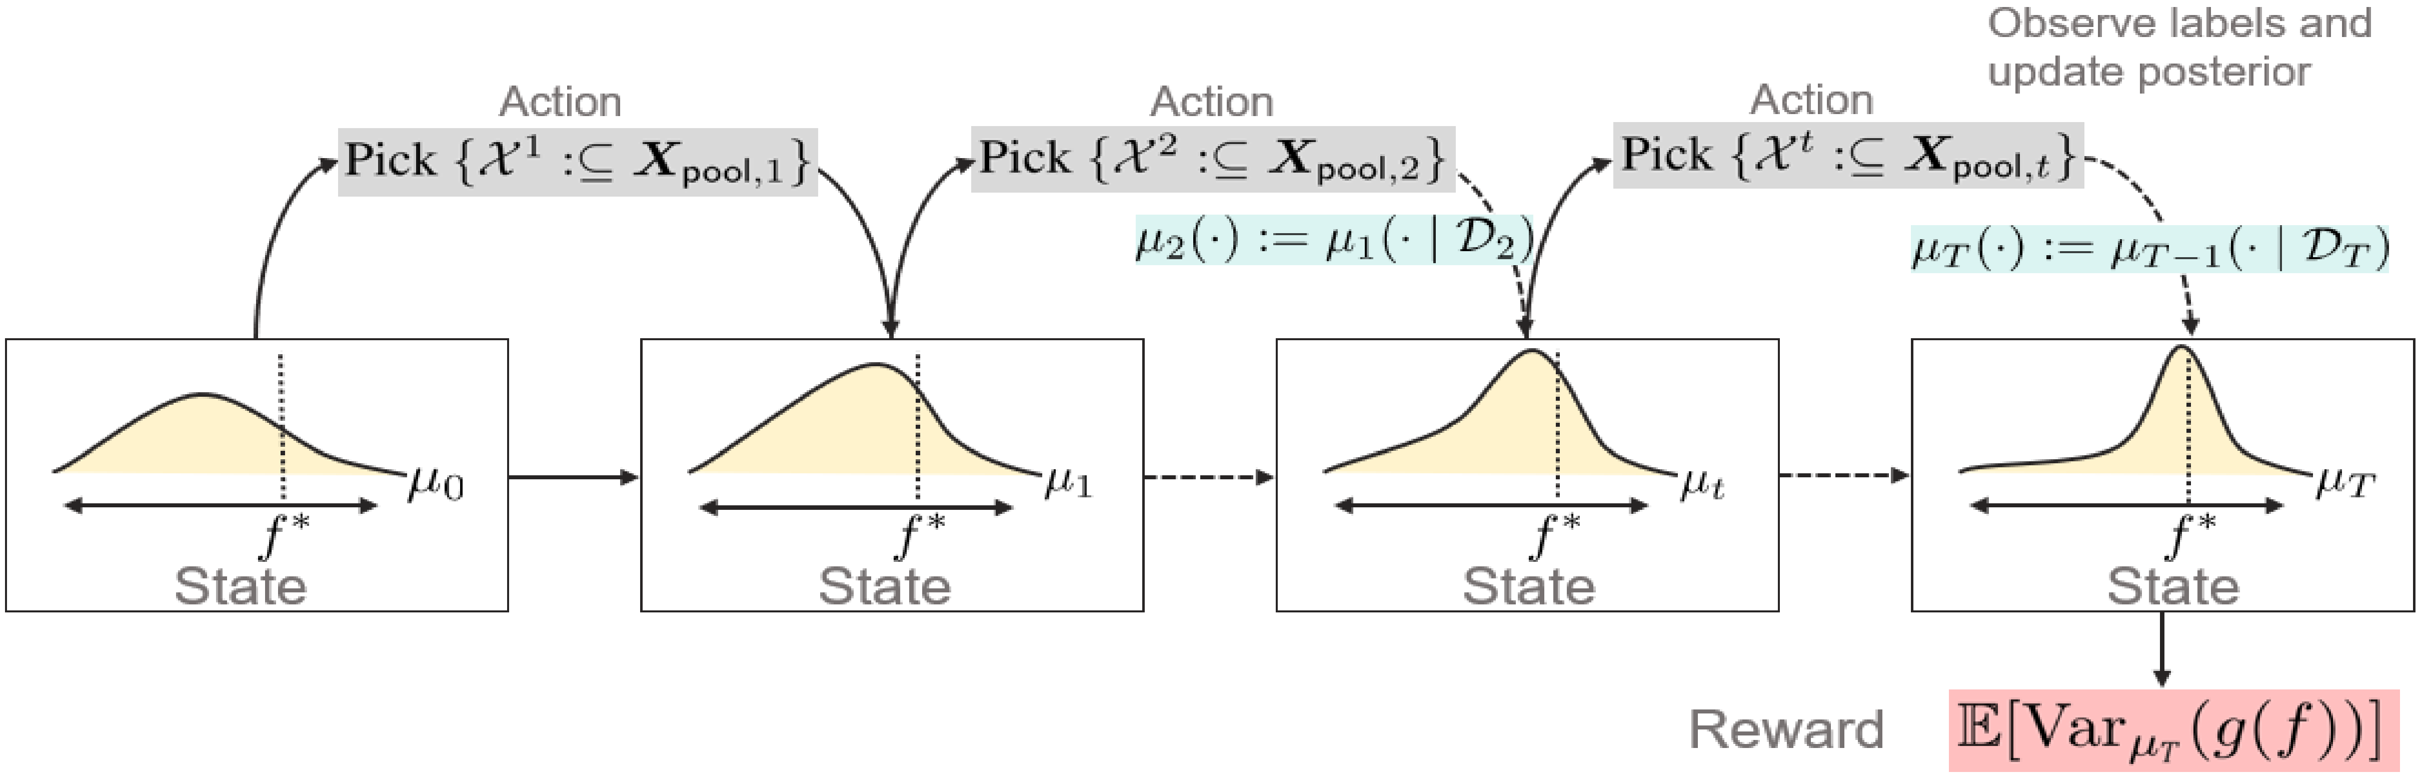

- Posteriors through **Uncertainty Quantification** – GPs, Deep Learning based methods (Ensemble)
- Actions sampled using policy π_θ – parametrized through **K-subset sampling**
- **Reward evaluation** – $E[Var_{f\sim\mu_T}(g(f))]$ where $g(f) = MSE\ of\ model\ \psi(.)\ under\ f\ and\ P_X$

Solving the MDP

- Employ **one-step lookaheads**
- Policy gradients through PATHWISE gradients rather than high-variance REINFORCE
- Smoothed pipeline to enable PATHWISE gradients

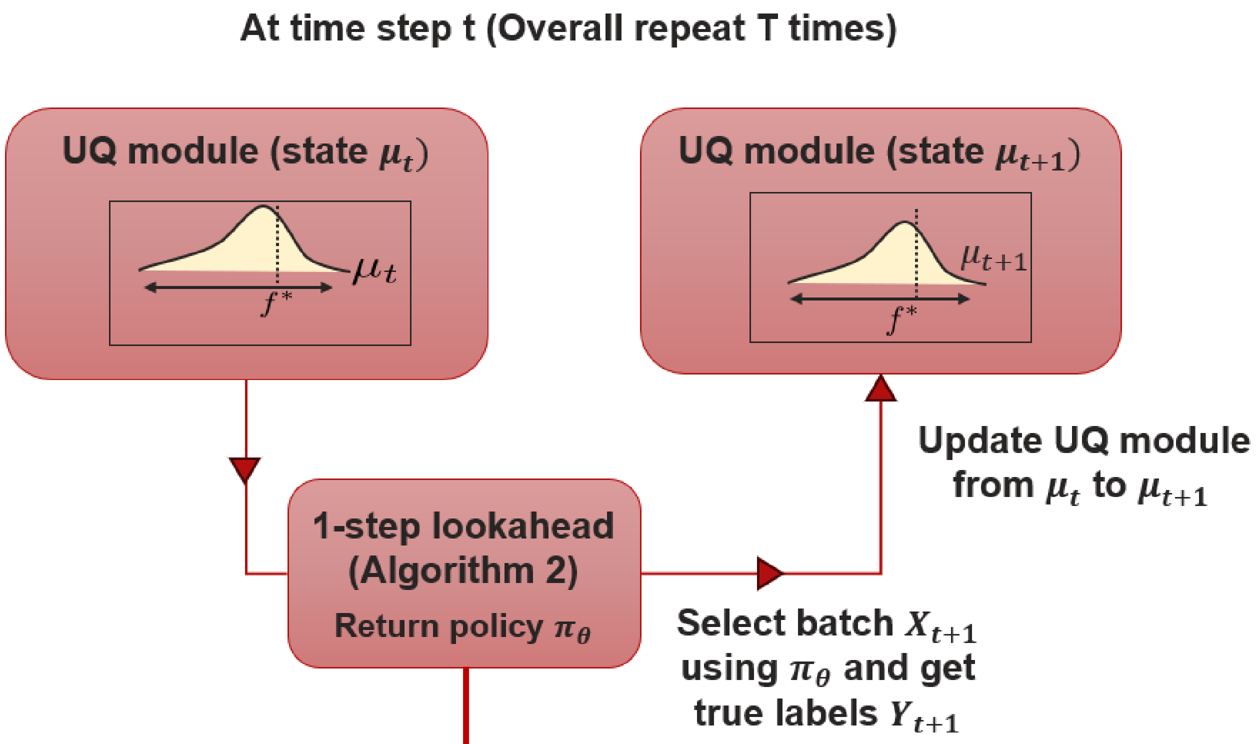

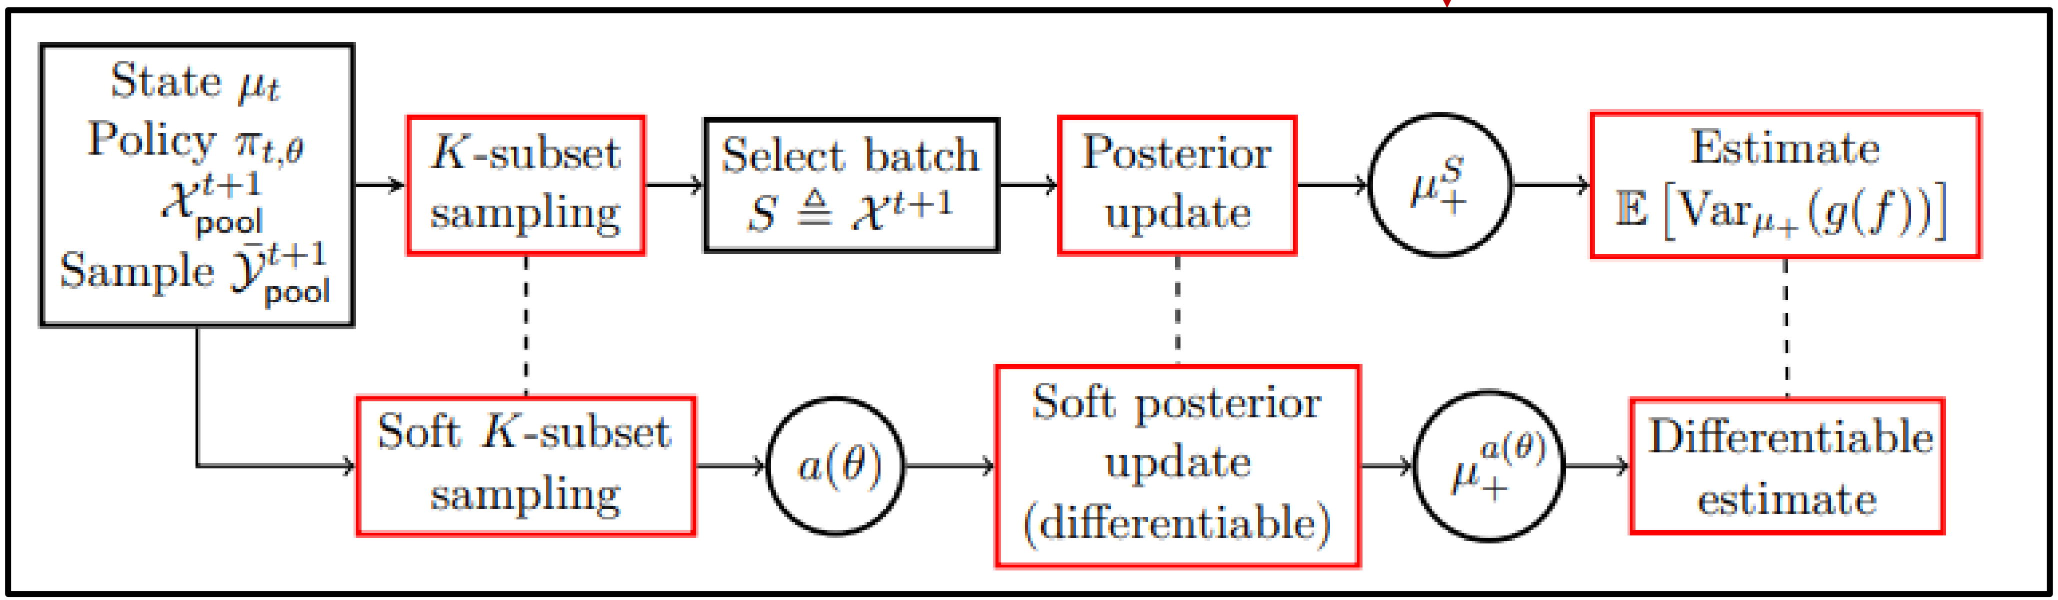

Our algorithm outperforms other baselines

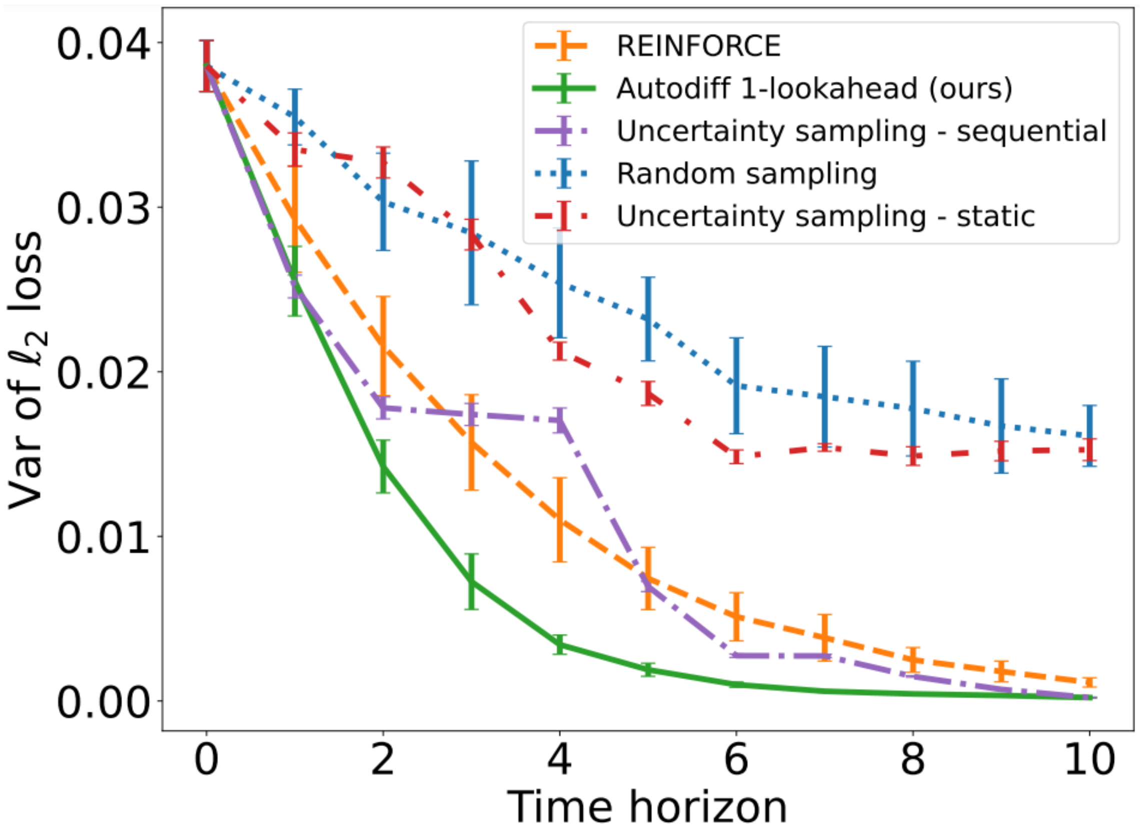

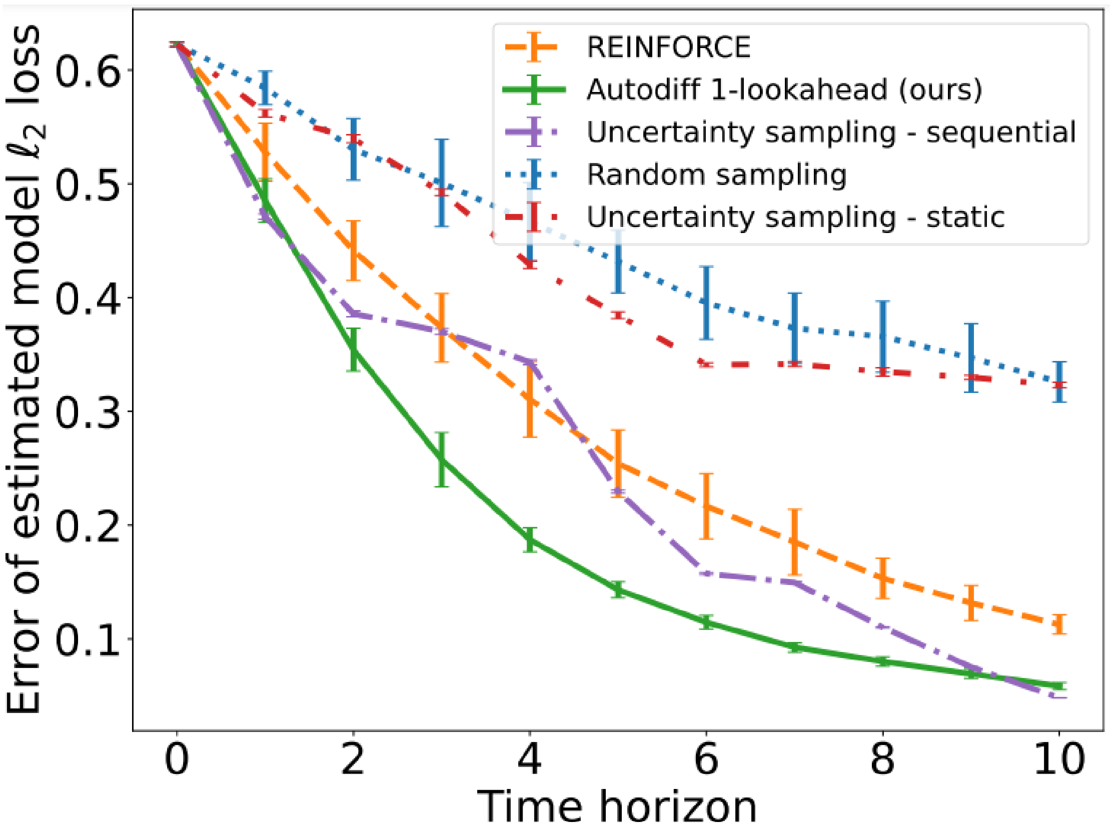

- Similar results for real datasets and deep learning based uncertainty quantification methodologies