Adaptive Transductive Inference via Sequential Experimental Design with Contextual Retention

(NeurIPS 2024 BDU Workshop)

Tareq Si Salem

Huawei Technologies, Paris Research Center, France

Problem Statement

Modern machine learning (ML) struggles with limited data and dynamic real-world conditions.

- Data acquisition for many ML tasks is often constrained by resource limitations, including sample size, temporal factors, and computational power.
- The deployment of ML models is challenged by the phenomenon of concept drift, where the underlying relationships between input features and target variables evolve over time.

This research aims to address the following question:

How can we effectively optimize data acquisition, data freshness, and model selection methodologies in dynamic environments characterized by concept drifts?

Contributions

This work makes several key contributions to the study of active learning in dynamic environments with concept drift.

- **Novel Framework for Dynamic Active Learning.** We introduce a novel framework for active learning in dynamic environments, incorporating concept drift and data freshness. This framework addresses the challenge of limited-capacity context retention by integrating data collection, freshness decisions, and model retraining strategies.
- **Theoretical Analysis and Decoupled Approach.** We provide a rigorous theoretical analysis, identifying the inherent trade-off between variance and bias. To address this, we propose a decoupled approach: a variance reduction policy for experimental design and a bias reduction policy for data freshness parameter selection.
- **Unified Framework with Dynamic Regret Guarantees.** We present a unified framework for instantiating concrete policies, leveraging online mirror descent (OMD). We establish dynamic regret guarantees, extending beyond static settings, and adapt OMD for both full-information and bandit settings.

Main Assumptions

Assumption 1. *(Compact Experiments and Query Sets) Experiments $\mathbf{x} \in \mathcal{X}$ and $\mathbf{x}' \in \mathcal{Z}$ are uniformly bounded under the ℓ_2 norm by $D_{\mathcal{X}}$ and $D_{\mathcal{Z}}$, respectively. Formally, $\|\boldsymbol{x}\|_2 \leq D_{\mathcal{X}}, \|\boldsymbol{x}'\|_2 \leq D_{\mathcal{Z}}$ for all $\boldsymbol{x} \in \mathcal{X}, \boldsymbol{x}' \in \mathcal{Z}$.*

Assumption 2. *(Compact Parameter Set) We assume that the true model parameters $\beta_t^\star$ for $t \in [T]$, are uniformly bounded. Specifically, there exists a positive constant $B^\star$ such that $\|\beta_t^\star\|_2 \leq B^\star$ for all $t \in [T]$.*

Assumption 3. *(Invertible Design Matrices) The matrix $\sum_{\boldsymbol{x}\in\mathcal{X}} \boldsymbol{x}\boldsymbol{x}^\intercal$ is non-singular, meaning that there exists a positive constant $\omega \in \mathbb{R}_{>0}$ such that the following inequality holds: $\sum_{\boldsymbol{x}\in\mathcal{X}} \boldsymbol{x}\boldsymbol{x}^\intercal \succeq |\mathcal{X}|\,\omega \boldsymbol{I} \succ 0$.*

Assumptions 1–2 guarantee the compactness of the experimental design space, the query space, and the model space. This compactness assumption is frequently employed in the analysis of learning problems, facilitating the establishment of various theoretical properties. Furthermore, Assumption 3 ensures the invertibility of the covariance matrix $\boldsymbol{V}(\boldsymbol{\pi})$, enabling the well-definedness of the inference model.

Proposed Approach

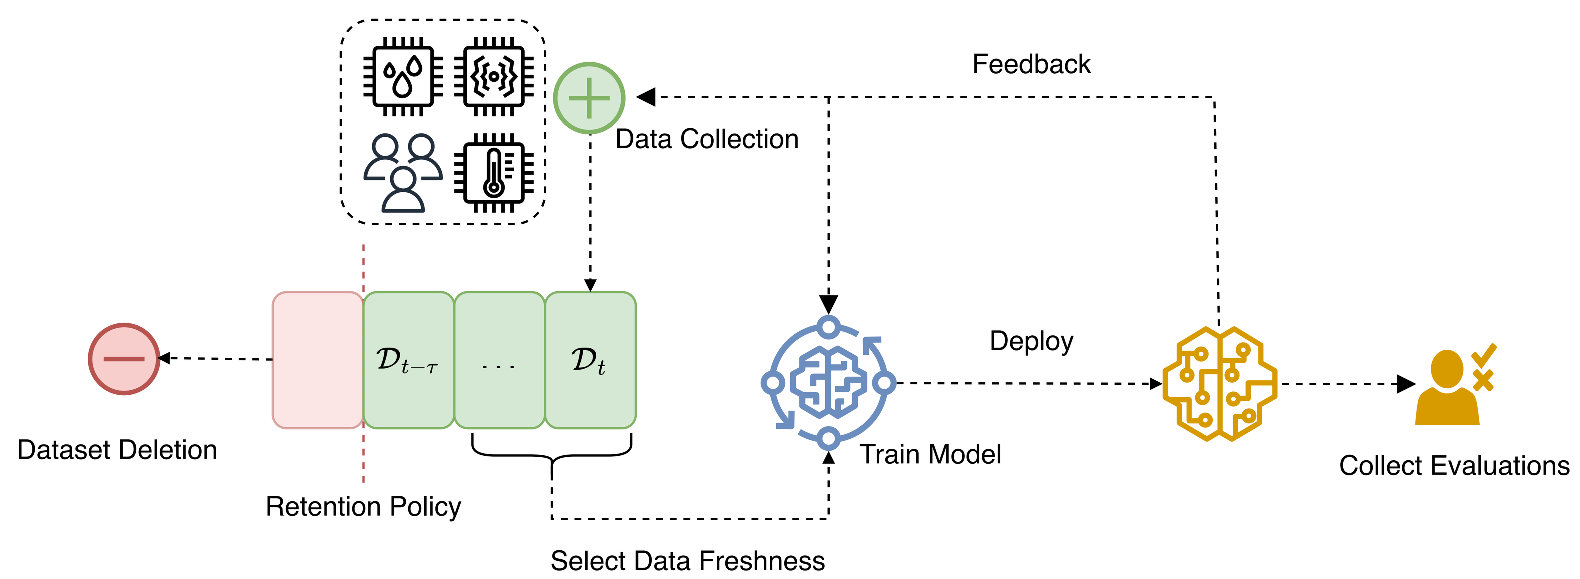

Data Collection. The policy has access to a pool of experiments $\mathcal{X} \subset \mathbb{R}^d$ to collect labels from a variety of experimental sources, such as sensors, surveys, and databases. A data retention policy purges datasets exceeding an age threshold $\tau \in \mathbb{N}$. The set of all feasible experimental designs is

$$\triangle_{\mathcal{X}} \triangleq \left\{\boldsymbol{\pi} \in [0,1]^{\mathcal{X}} : \|\boldsymbol{\pi}\|_1 = 1\right\}.$$

The collected labels is some noisy perturbation of a true underlying linear model:

$$y_t = \beta_t^\star \cdot \boldsymbol{x} + \xi \sim N(\beta_t^\star \cdot \boldsymbol{x}, \sigma^2)$$

Model Retraining. The LSE model is trained only on the data selected at timeslot t, i.e., the datasets $\mathcal{D}_{t-\tau_t}, \mathcal{D}_{t-\tau_t+1}, \ldots, \mathcal{D}_t$. The estimator is given by

$$\hat{\beta}_t = \frac{1}{M}\boldsymbol{V}^{-1}(\boldsymbol{\pi}_{t-\tau_t:t})\boldsymbol{X}^\intercal \boldsymbol{y}\left(\mathcal{D}_{t-\tau_t:t}\right),$$

where $\boldsymbol{X} = (\boldsymbol{x}^\intercal)_{\boldsymbol{x}\in\mathcal{X}}$, $\boldsymbol{V}(\boldsymbol{\pi}) = \boldsymbol{X}^\intercal \mathrm{diag}(\boldsymbol{\pi})\boldsymbol{X}$, and $\boldsymbol{y}(\mathcal{D}) = (\sum_{(\boldsymbol{x},y_i)\in\mathcal{D}_{\boldsymbol{x}}} y_i)_{\boldsymbol{x}\in\mathcal{X}}$.

Model Deployment. During deployment, the trained model $\hat{\beta}_t$ is used to predict labels for experiments $\boldsymbol{x}_t \in \mathcal{Z} \subset \mathbb{R}^d$. User feedback in the form of prediction errors is collected to refine the model.

Performance Metric. We compare the performance of the sequence of designs and data-freshness parameters w.r.t. the best design in hindsight and data-freshness window size after seeing all the queries in terms of the EPE:

$$\mathfrak{R}_T(\boldsymbol{\mathcal{P}}) \triangleq \mathbb{E}\left[\sum_{t=\tau+1}^{T} f_t(\boldsymbol{\pi}_{t-\tau_t}, \ldots, \boldsymbol{\pi}_t) - \min_{\{\boldsymbol{\pi}_t^\star\}_{t=1}^T \in \triangle_{\mathcal{X}}^T, \{\tau_t^\star\}_{t=1}^T \in \mathcal{T}^T} \sum_{t=\tau+1}^{T} f_t(\boldsymbol{\pi}^\star_{t-\tau_t^\star}, \ldots, \boldsymbol{\pi}_t^\star)\right],$$

Taming the Regret

We propose a policy $\boldsymbol{\mathcal{P}}^{\mathrm{Entropic-VBR}}$ that decides data-freshness τ_t and experimental designs $\boldsymbol{\pi}_t$.

Under proper assumption, we establish the following regret guarantees:

$$\mathfrak{R}_T(\boldsymbol{\mathcal{P}}^{\mathrm{Entropic-VBR}}) = \mathcal{O}\left(\sqrt{\log(1/\sigma)P_T^{\star,\mathrm{v}}T} + \sqrt{\log(T)P_T^{\star,\mathrm{b}}T} + P_T^{\star,\mathrm{v}}\right),$$

under path lengths $P_T^{\star,v}$ and $P_T^{\star,b}$.

This results suggests that, as the path length grows sublinearly, the average regret of the policy approaches zero.

Numerical Demonstration

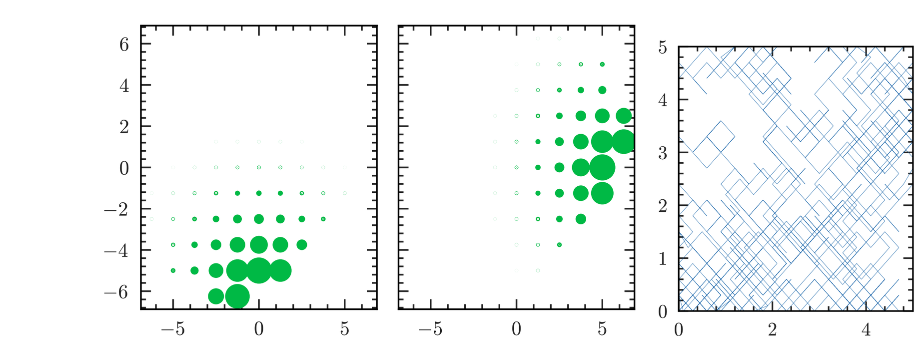

Query distributions, and true model drift (10^3 initial iterations)

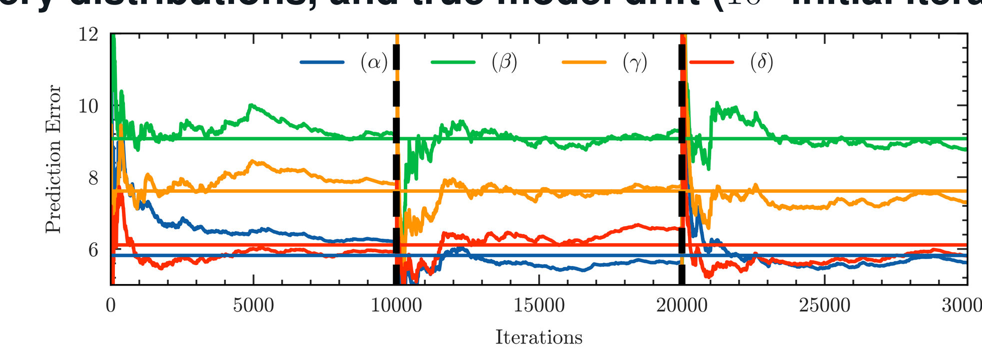

Time-averaged prediction error for three intervals: $0 \leq t \leq \frac{1}{3}T$, $\frac{1}{3}T < t \leq \frac{2}{3}T$, and $\frac{2}{3}T < t \leq T$

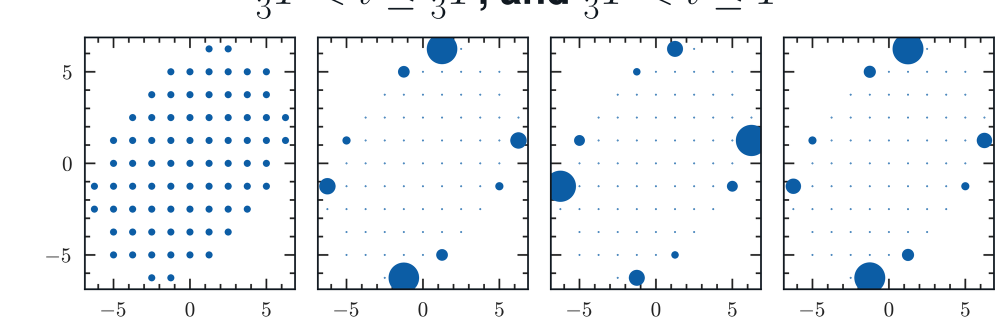

Initial and selected designs at $t = \frac{1}{3}T$, $t = \frac{2}{3}T$, and $t = T$

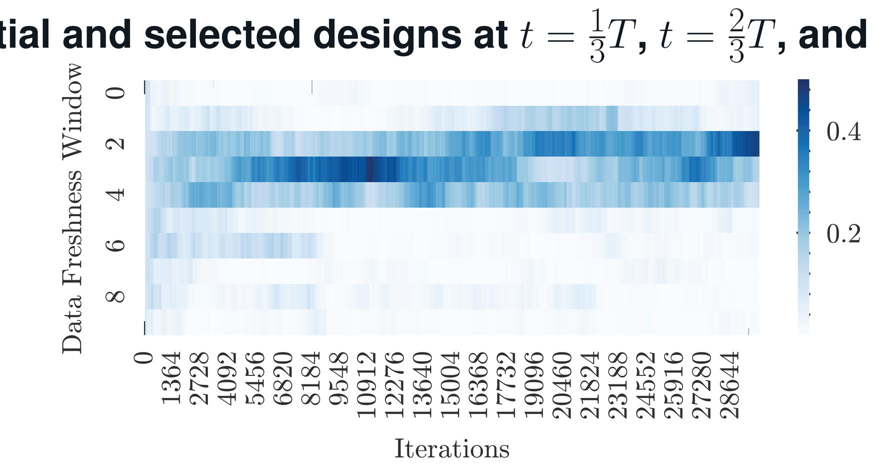

Evolution of the policy-learned data-freshness distribution over time, illustrating the non-trivial nature of the optimal window size.

The performance evaluation shows that a uniform experimental design with the maximum freshness window is the least effective baseline (β). Optimizing the data-freshness window improves performance (γ), and further gains are achieved by optimizing the experimental design (δ). Our proposed policy (α) outperforms the baseline, demonstrating its ability to identify optimal sequences of designs and freshness windows.

Conclusion

This work has presented a comprehensive investigation of active learning in dynamic environments characterized by concept drift.

Potential avenues for future research:

- **Beyond Non-linear Relationships.** To improve the flexibility and accuracy of the model, future work could explore non-linear models, such as those based on kernel methods.
- **General Noise Models.** Expanding the framework to accommodate more general noise models would increase its applicability to a wider range of real-world scenarios.