Causal-Informed Hybrid Online Adaptive Optimization for Ad Load Personalization in Large-Scale Social Networks

Aakash Mishra, Qi Xu, Zhigang Hua,
Keyu Nie, Vishwanath Sangale,
Vishal Vaingankar, Jizhe Zhang, Ren Mao

Introduction & Motivation

- **The Challenge:** Personalizing ad load in large-scale social networks requires balancing user experience and conversions under operational constraints.
- **The Trade-off:** Showing too few ads under utilizes conversions, while too many degrade engagement and retention.
- **Complexity:** This is a high-dimensional, constrained online optimization problem where decisions must adapt rapidly to dynamic user behavior.

Problem Statement

Current methods face distinct limitations:

- **Traditional Primal-Dual:** Enforces constraints reliably but adapts slowly in dynamic environments. Inherently exploitative.
- **Bayesian Optimization (BO):** Enables exploration under uncertainty but suffers from slow convergence in high-dimensional spaces.

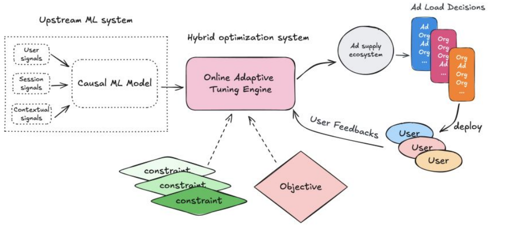

Figure 1: Causal-informed hybrid online optimization system for ad load personalization

Proposed Framework CTR-CBO

Cohort-Based Trust Region Contextual Bayesian Optimization We propose a hybrid framework combining Primal-Dual methods with BO, enhanced by upstream Causal ML.

- **Hybrid Approach:** Combines spatial partitioning (MoRBO) and temporal adaptation (PDCBO).
- **Causal Integration:** Leverages upstream causal ML to estimate counterfactual treatment effects, informing Gaussian Process Regression (GPR) surrogates.
- **Architecture:**
  - **Primal Step:** Maximizes hypervolume improvement with a trust-region approach.
  - **Dual Step:** Ensures time-average constraint satisfaction weighted across cohorts.

Methods & Algorithm

Local GP Modeling Partition policy space into trust regions delineated by causal user-cohort sensitivity. Fit local Gaussian Process Regressors for ads score and constraint GPs per cohort.

Kernel Formulation To model the exponential relationship at the session level, we define the kernel as:

$$k_{sigmoid}(x_i, x_j) = \sigma_f^2 \cdot \frac{1}{1 + e^{-(a \cdot x_i^T x_j + b)}}$$

Primal Update (Multi-Objective) Compute hypervolume improvement (HVI) and update parameters (theta).

$$\theta_{k,t} = arg\,max_{\theta \in Tr_k}\{HVI_k(\theta, z_t) + \eta \lambda_t^T c_{k,t}(\theta, z_t)\}$$

Dual Update (Constraint Satisfaction)
Update dual variables λ for time-average constraint satisfaction

$$\lambda_{t+1} = [\lambda_t + \sum_{k=1}^{K} w_k c_{k,t}(\theta_{k,t}, z_t) + \epsilon e]_+$$

Experimental Results

Validated on a billion-user social network and synthetic datasets.

- **Social Media A/B Test:** CTR-CBO required significantly fewer iterations to converge compared to CBO (~2 iterations vs. ~9).
- **Synthetic Dataset:** CTR-CBO outperformed CBO in achieving convergence to policy thresholds (1% ads score increase for 1.5% impression increase).

Proxy Model Accuracy The GP surrogates showed strong predictive performance against actual A/B test results

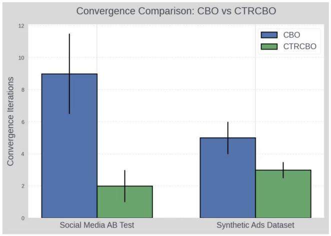

Figure 2: Comparison of convergence iterations between CBO and CTRCBO.

Conclusion

Performance: CTR-CBO outperforms naive CBO for online policy tuning based on causal model cohorts.
Impact: Demonstrates faster convergence, robust constraint satisfaction, and improved personalization metrics.
Scalability: The framework effectively handles large-scale ads supply systems by combining the stability of primal-dual methods with the adaptive exploration of BO.

References

[1] N. Buchbinder et al., "Online primal-dual algorithms for maximizing ad allocations," ESA, 2007.

[2] S. Daulton et al., "Multi-objective bayesian optimization over high-dimensional search spaces," UAI, 2022.

[3] A. Goli et al., "Personalizing ad load to optimize subscription and ad revenues," *Marketing Science*, 2022.

[4] H. Sagtani et al., "Ad-load balancing via off-policy learning in a content marketplace," *arXiv*, 2023.

[5] W. Shi et al., "Ads supply personalization via doubly robust learning," 2024.

[6] Q. Xu et al., "Large-scale sponsored search ad allocation with online adaptation," *Nature Scientific Reports*, 2024.

[7] W. Xu et al., "Primal-dual contextual bayesian optimization for control system online optimization," CDC, 2023.