Integrating Inference and Experimental Design for Contextual Behavioral Model Learning

Gongtao Zhou Haoran Yu

zhougt3@qq.com, yhrhawk@gmail.com

School of Computer Science & Technology, Beijing Institute of Technology

Background

Portfolio Choice:

- Trading platforms utilize robo-advisors to interact with investors and elicit hidden information from their choices across portfolios.

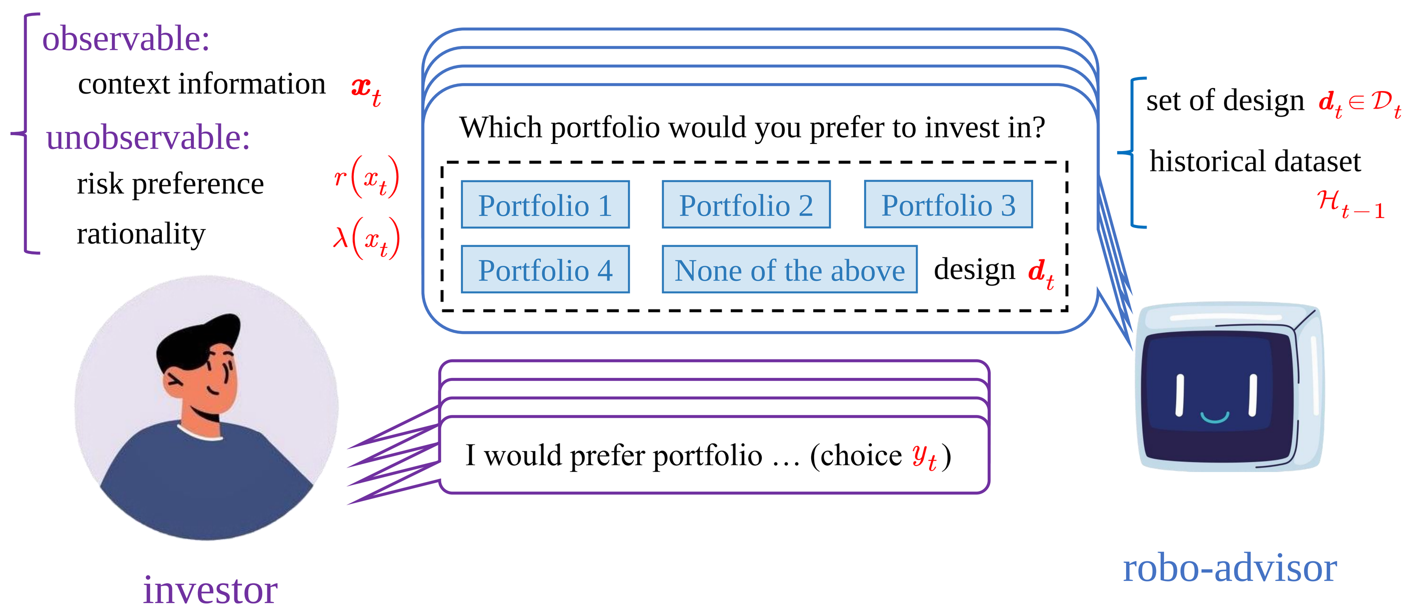

- Inferring the relationship between an investor's observable and hidden information depends on the platform's experimental design.

Our Goal: *Design experiment to learn contextual behavioral model.*

- Collect the most informative data for learning network weights.

Problem Formulation

- The platform strategy is based on:
  - t: period of platform interaction $t = 1, 2, .., T$
  - $\boldsymbol{x}_t$: investor contextual information
  - $\boldsymbol{d} = \{(m_1, v_1), ..., (m_K, v_K)\}$: designed investment portfolios
    - $*$ m, v: return mean and return variance
    - $*$ K: the number of portfolios in design
  - $\mathcal{H}_{t-1} \triangleq \{(\mathbf{x}_1, \mathbf{d}_1, y_1), ..., (\mathbf{x}_{t-1}, \mathbf{d}_{t-1}, y_{t-1})\}$: historical dataset
  - $\boldsymbol{\theta}$: contextual behavioral model weights

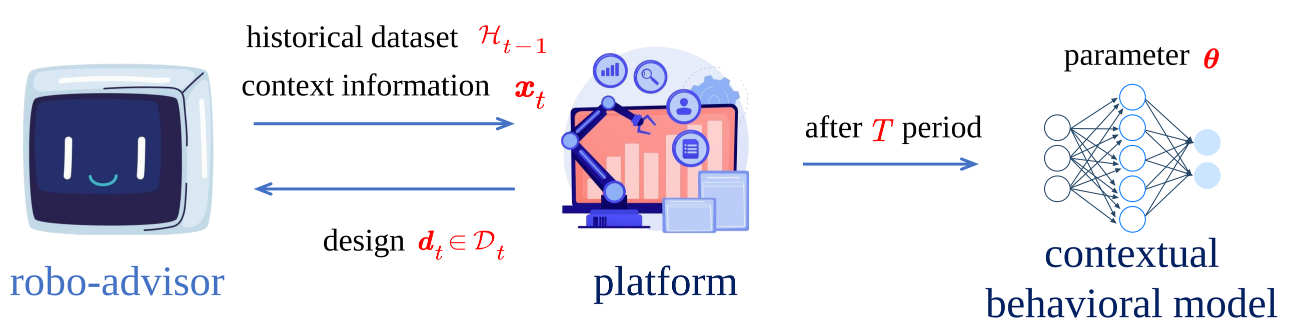

Platform's Experimental Design Problem

Given $\mathcal{H}_{t-1}$ and $\mathbf{x}_t$, how to optimize the design $\mathbf{d}_t \in \mathcal{D}_t$ for training an accurate θ-parameterized neural network?

Platform's Experimental Design Solution

- Expected information gain (EIG): quantify design quality

$$\text{EIG}(\boldsymbol{d}) \triangleq \mathbb{E}_{p(y|\boldsymbol{x},\boldsymbol{d})}[H[p(\boldsymbol{\theta})] - H[p(\boldsymbol{\theta}|\boldsymbol{x}, y, \boldsymbol{d})]]$$

  - $p(\theta)$ is prior; $p(\boldsymbol{\theta}|\boldsymbol{x}, y, \boldsymbol{d})$ is posterior.
  - The uncertainty of $\boldsymbol{\theta}$ decreases as new data is collected.

- Bayesian neural networks (BNN): calculate EIG

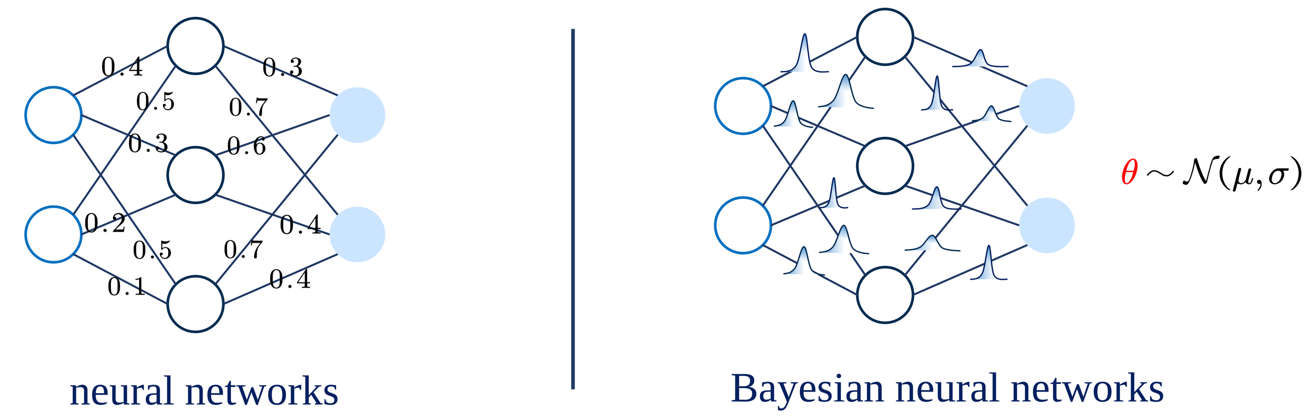

  - EIG calculations necessitate non-fixed model weights.
  - We use variational inference $q(\boldsymbol{\theta}|\boldsymbol{\phi})$ to approximate $p(\boldsymbol{\theta}|\boldsymbol{x}, y, \boldsymbol{d})$.

Our Solution:

- Inference: learn $\boldsymbol{\phi}$ by training Bayesian neural network
- Design: choose the design to maximize the estimated EIG

I-ID-LP

- **I**ntegrated **I**nference-and-**D**esign with **L**earnable **P**riors (I-ID-LP)
  - Learnable prior: a random initialized prior distribution $p(\boldsymbol{\theta})$ will disrupt the inference step during the early periods
  - Integrated Inference-and-Design: separating inference and design may lead to suboptimal performance, as the optimization of $\boldsymbol{\phi}$ fails to maximize information gain

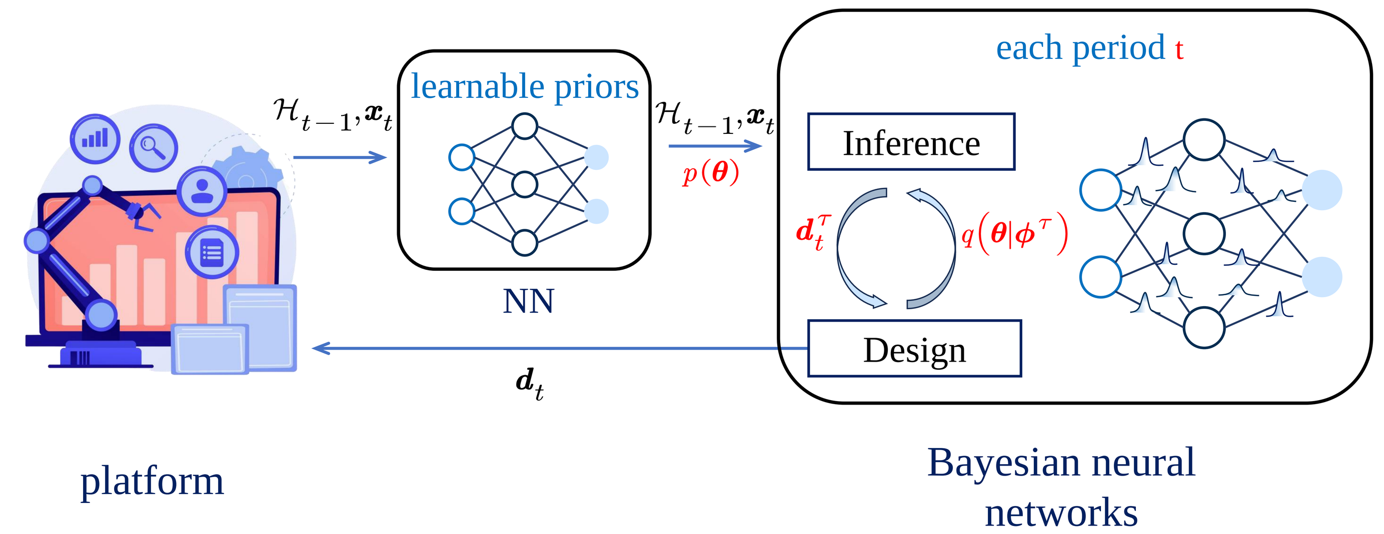

Iterative Optimization

Loss function for iterative optimization:

- Optimize ϕ^τ while keeping $\boldsymbol{d}_t = \boldsymbol{d}_t^{\tau-1}$:

$$\phi^\tau = \arg\min_{\boldsymbol{\phi}} \text{KL}[q(\boldsymbol{\theta}|\boldsymbol{\phi})||p(\boldsymbol{\theta}|\mathcal{H}_{t-1})] - \mathbb{E}_{q(\boldsymbol{\theta}|\boldsymbol{\phi})}[\log \text{KL}[p(y_t|\boldsymbol{x}_t, \boldsymbol{d}_t^{\tau-1}, \boldsymbol{\theta})||p(y_t|\boldsymbol{x}_t, \boldsymbol{d}_t^{\tau-1})]]$$

- Optimize $\boldsymbol{d}_t^\tau$ while keeping $\boldsymbol{\phi} = \boldsymbol{\phi}^\tau$:

$$\boldsymbol{d}_t^\tau = \arg\max_{\boldsymbol{d}_t} \int q(\boldsymbol{\theta}|\boldsymbol{\phi}^\tau) \log \text{KL}[p(y_t|\boldsymbol{x}_t, \boldsymbol{d}_t, \boldsymbol{\theta})||p(y_t|\boldsymbol{x}_t, \boldsymbol{d}_t)]d\boldsymbol{\theta}$$

Experiments & Results

Comparison Among ID Methods:

- NLL Loss: the loss of predicting investor choice y in $\mathcal{H}_{test}$.
- MSEs of Predicting $r(\boldsymbol{x})$ and $\lambda(\boldsymbol{x})$: the mean squared errors between the outputs of NN and the actual $r(\boldsymbol{x})$, $\lambda(\boldsymbol{x})$ of investors in $\mathcal{H}_{test}$.

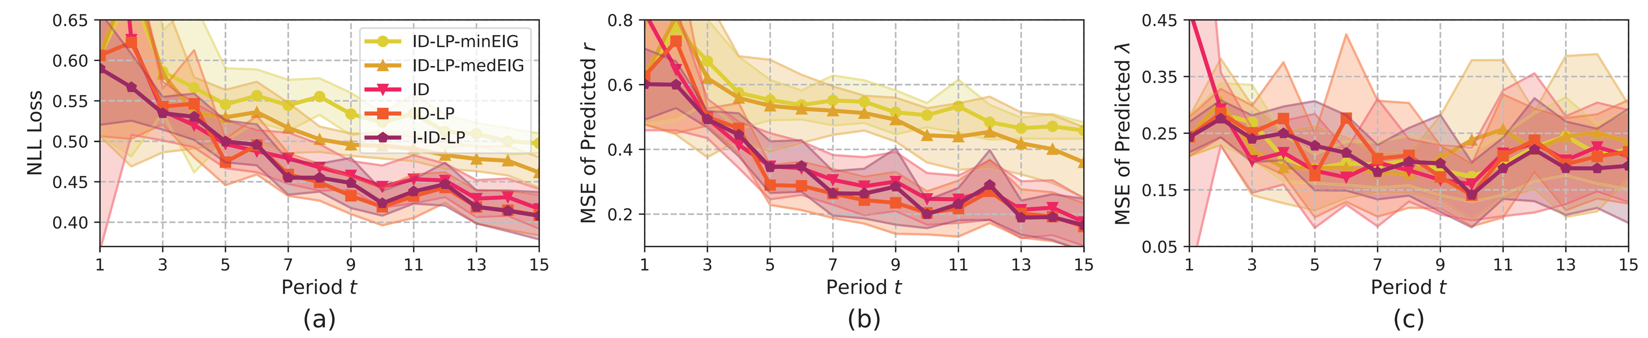

Comparison of ID Methods Under Setting A

Comparison with Other Methods

Table: Negative Log-Likelihood Loss on Testing Dataset

Method	Setting A	Setting B	Setting C
ID	0.416 ± 0.025	0.319 ± 0.017	0.397 ± 0.010
ID-LP	0.409 ± 0.025	0.306 ± 0.018	0.390 ± 0.007
I-ID-LP	**0.408** ± **0.029**	**0.297** ± **0.010**	**0.387** ± **0.007**
ID-LP-minEIG	0.498 ± 0.012	0.427 ± 0.026	0.565 ± 0.113
ID-LP-medEIG	0.461 ± 0.019	0.331 ± 0.010	0.435 ± 0.047
PreEntropy	0.557 ± 0.111	0.366 ± 0.028	0.634 ± 0.216
MinMaxPro	0.611 ± 0.122	0.366 ± 0.030	0.705 ± 0.196
LaplaceInfer	0.462 ± 0.026	0.327 ± 0.012	0.411 ± 0.013
Random	0.459 ± 0.022	0.329 ± 0.012	0.425 ± 0.035
MaxMean	0.468 ± 0.003	0.404 ± 0.045	0.510 ± 0.012
MaxVar	0.560 ± 0.008	0.340 ± 0.008	0.624 ± 0.010
MaxMean+Var	0.468 ± 0.003	0.336 ± 0.004	0.501 ± 0.003
MaxMean−Var	0.512 ± 0.013	0.803 ± 0.424	0.562 ± 0.063

More results on other experiments can be found in our paper.