Decision-Focused Sequential Experimental Design: A Directional Uncertainty-Guided Approach

Mo Liu, UNC-Chapel Hill, Co-authors: Beichen Wan, Paul Grigas, Zuo-Jun Max Shen

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

What is sequential experimental design?

Experimental design: $x_1, x_2, \ldots, x_m$. Also called: feature, covariates, group
Outcome vector for one design: $c \in \mathbb{R}^d$. Also called: Label, Response, cost

- The outcome c is random
- The conditional distribution of $c|x$ is unknown

Goal: Predict $\mathbb{E}[c|x]$ for every possible design x

Prediction models

- Hypothesis class of prediction models: $\mathcal{H} = \{h_1, h_2, \ldots, h_n\}$
- Goal: Identify the best prediction model $h^* \in \mathcal{H}$
  such that $h^*(x) \approx \mathbb{E}[c|x]$ for every possible design x

Difficulty: Designing an experiment and observe its outcome is costly and time-consuming

- Examples: medical trials, LLM labeling

Solution:

- Reduce the number of required experiments using sequential experimental design

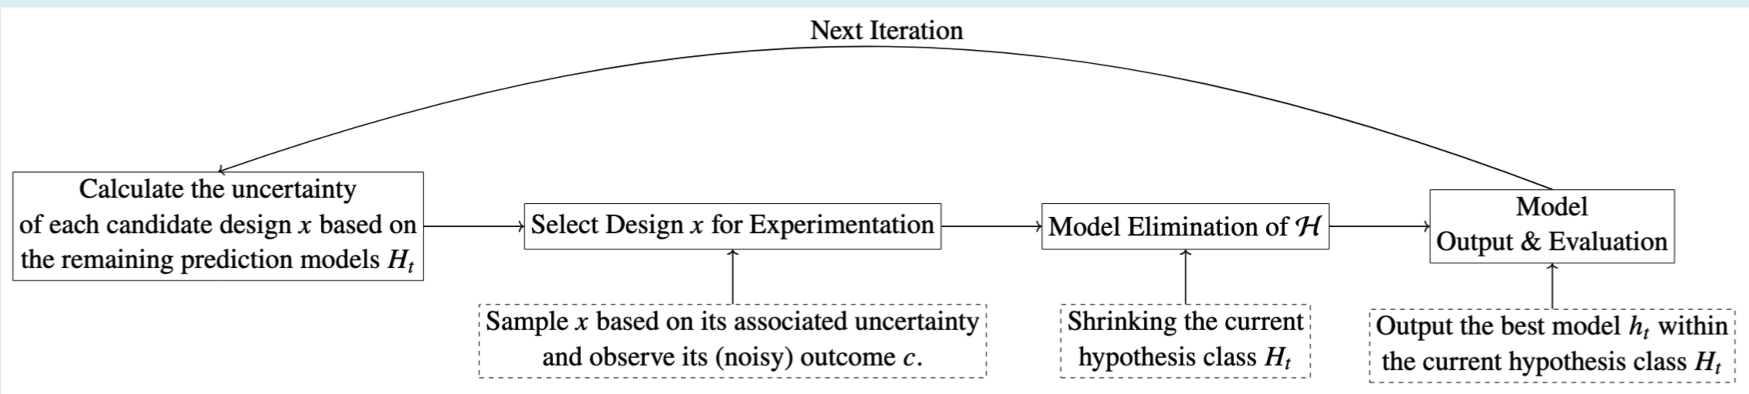

What is decision-focused learning?

The outcome vector c is used as coefficients for a downstream optimization problem:

$$\min_{w \in S} \mathbb{E}[c^T w|x] = \min_{w \in S} \mathbb{E}[c^T|x] w$$

- w is a decision vector within a feasible region S, which can be a polyhedron or an integer set

Examples of downstream decisions:

- Choosing a medical treatment for a patient given their medical features
- Planning a vehicle route on a specific rainy day
- Assigning jobs to different agents based on task features

Predict–Then–Optimize Framework

- Given a group x, and a prediction $h(x)$ for $\mathbb{E}[c|x]$
- To obtain the optimal decision w, we use the plug-in approach:

$$\min_{w \in S} h(x)^T w$$

Let $w^*(c)$ denote the optimal decision given the true cost vector c

- Smart predict-then-optimize (SPO) loss

$$\ell_{\text{SPO}}(h(x), c) = c^T w^*\big(h(x)\big) - c^T w^*(c)$$

The SPO loss measures the decision error induced by using the prediction

- Best prediction model is defined by

$$h^* \in \arg\min_{h \in \mathcal{H}} R_{\text{SPO}}(h), \text{ where } R_{\text{SPO}}(h) = \mathbb{E}[\ell_{\text{SPO}}(h(x), c)]$$

Why are traditional uncertainty-quantification methods not suitable?

Given a hypothesis class H, how to quantify the uncertainty of each design x?

- The uncertainty measure must align with the downstream **SPO loss**.

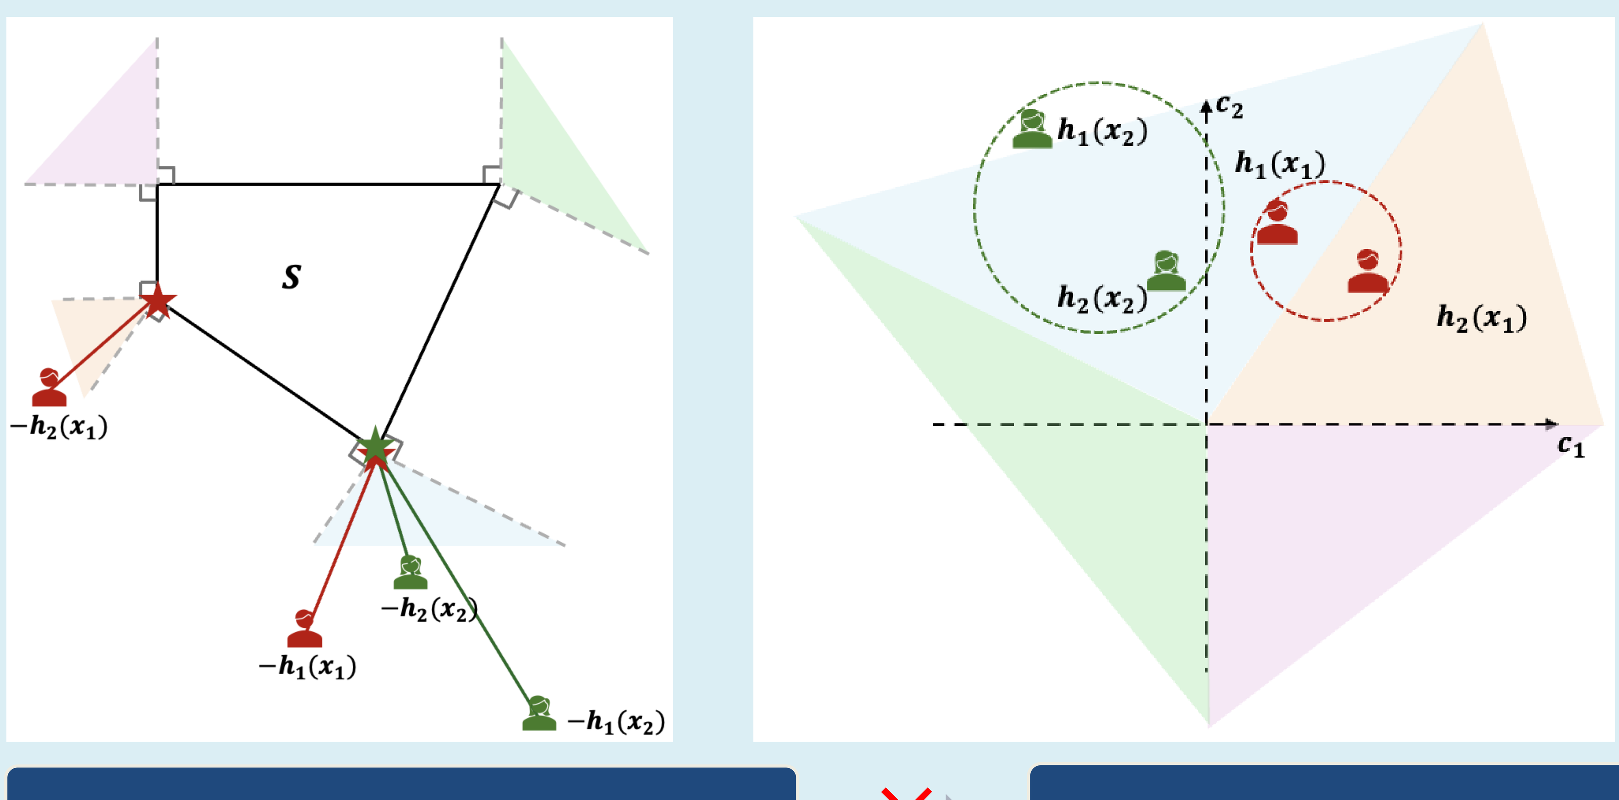

A smaller uncertainty based on L2 norm ✗ → A smaller uncertainty for SPO loss

A new uncertainty criterion: Directional Uncertainty

Given a set of candidate prediction models H and design x:

- Traditional uncertainty: $\max_{h_1, h_2 \in H} \|h_1(x) - h_2(x)\|$
- **Directional uncertainty**:

$$\max_{h_1, h_2 \in H} \left\| \frac{h_1(x)}{\|h_1(x)\|} - \frac{h_2(x)}{\|h_2(x)\|} \right\|$$

Example: Suppose $c \in \mathbb{R}^2$. There are four candidate prediction models and three possible designs:

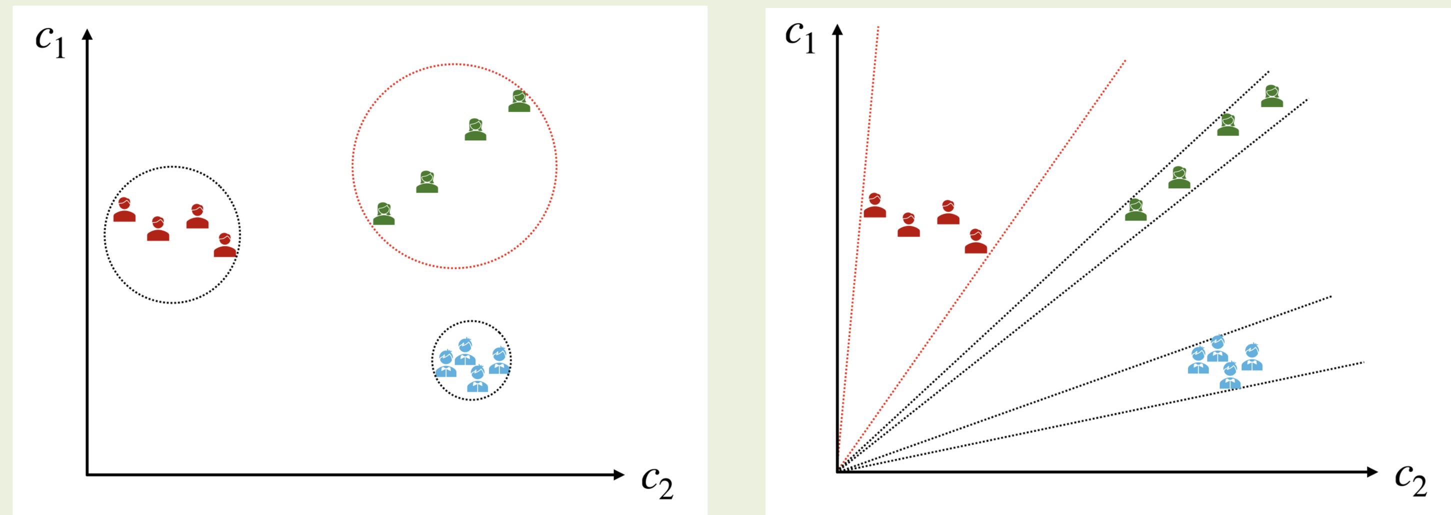

- Traditional criterion: Green design has the largest uncertainty
- Directional uncertainty: Red design has the largest uncertainty

Why does directional uncertainty better align with SPO loss?
Fact: SPO loss is scale invariant:

$$\ell_{\text{SPO}}(\hat{c}, c) = \ell_{\text{SPO}}(\alpha \hat{c}, c) \text{ for all } \hat{c} \text{ and } \alpha > 0.$$

Theoretical Guarantees

Directional margin condition:
There exists a constant $\eta > 0$, such that for all $h \in \mathcal{H}$, and $x \in \mathcal{X}$, it holds that $\frac{\nu_S(h(x))}{\|h(x)\|} \geq \eta$.

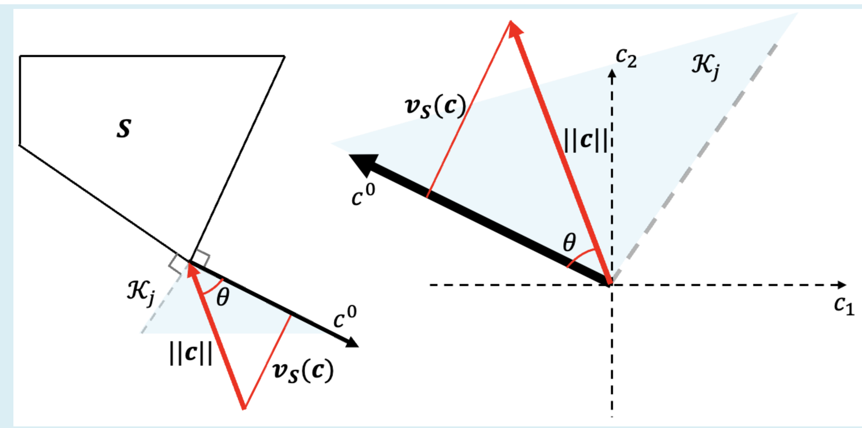

✓ The Directional margin condition holds for **any** hypothesis class after some transformations.

Lipschitz Lemma: Under the directional margin condition, there exists a constant $L > 0$, such that $\ell_{\text{SPO}}(\cdot, c)$ is L-$Lipschitz$.

Convergence Guarantee: Consider the Importance-Weighted Sequential Design based on Directional Uncertainty (IWSD-DU). Let $\delta \in (0,1)$ be a given parameter. For all $T \geq 1$, with probability at least $1 - \delta$, the SPO risk satisfies that

$$R_{\text{SPO}}(h_T) - R_{\text{SPO}}(h^*) \leq 4L \sqrt{\frac{\log(2T|\mathcal{H}|/\delta)}{T}}$$

Comparison with decision-blind design:
Under certain conditions, the IWSD-DU algorithm has an earlier stopping time than the design blind designs.

Numerical results on real-world data

Task Assignment Problem for LLMs:

Different LLMs have different (random) performance ($-c_{i,j}$) on different tasks. Given the category of the text and the capacity of each LLM, what is the best task assignment?

$$\begin{aligned} \min_{w \in \mathbb{R}^{m \times n}} : \ & \mathbb{E}\Big[\sum_{i=1}^{m} \sum_{j=1}^{n} c_{i,j} w_{i,j} \Big| x\Big] \\ \text{subject to } & \sum_{i=1}^{m} w_{i,j} \leq A_j, \quad \forall j = 1, \ldots, n, \\ & \sum_{j=1}^{n} w_{i,j} = B_i, \quad \forall i = 1, \ldots, m, \\ & w_{i,j} \in \{0,1\}, \quad \forall i = 1, \ldots, m, \ j = 1, \ldots, n. \end{aligned}$$

- ❑ Text summarization
- ❑ Text extension
- ❑ Text translation
- ❑ Bullet point generation
- ❑ Sentiment analysis

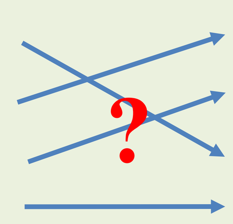

- ChatGPT5
- Gemma3
- Qwen1.5
- Deepseek-llm

Experimental design: Select one text from one category
Outcome of one experiment: Ask a human annotator to evaluate how well each LLM performs on each task
Goal: Use a small number of human-labeled experiments, to find a good prediction model to predict the performance $c_{i,j}$.

- This prediction model is plugged-in for future task assignment problem for each category.

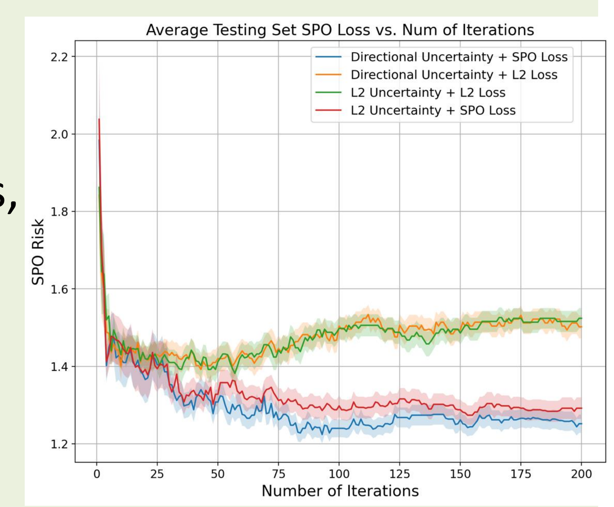

Numerical results: The IWSD-DU algorithm has a **lower** (testing set) SPO risk, given the same size of training set, compared with all benchmark design methods.