

Pareto Prompt Optimization

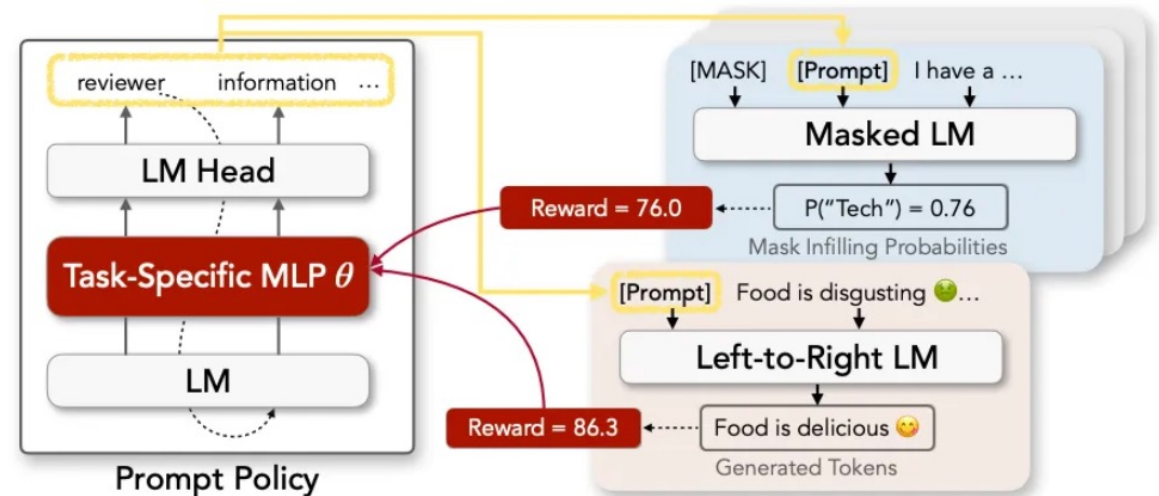
Guang Zhao, Byung-Jun Yoon, Gilchan Park, Shantenu Jha, Shinjae Yoo, Xiaoning Qian



@BrookhavenLab

What is Prompt Optimization

- A **prompt** is the input text to the LLM.
- Prompt optimization aims to automatically find effective prompts that produce desirable results from an LLM.




Workflow for RLPrompt

Source: Deng, Mingkai, et al. "Rlprompt: Optimizing discrete text prompts with reinforcement learning."

Multi objective Prompt Optimization


- Example: text summarization
 - Accuracy vs conciseness
- No single prompt excels in all objectives simultaneously
- We aim for a set of prompts from *Pareto Front*



HyperVolume Metric illustration in Objective space

Multi objective Prompt Optimization


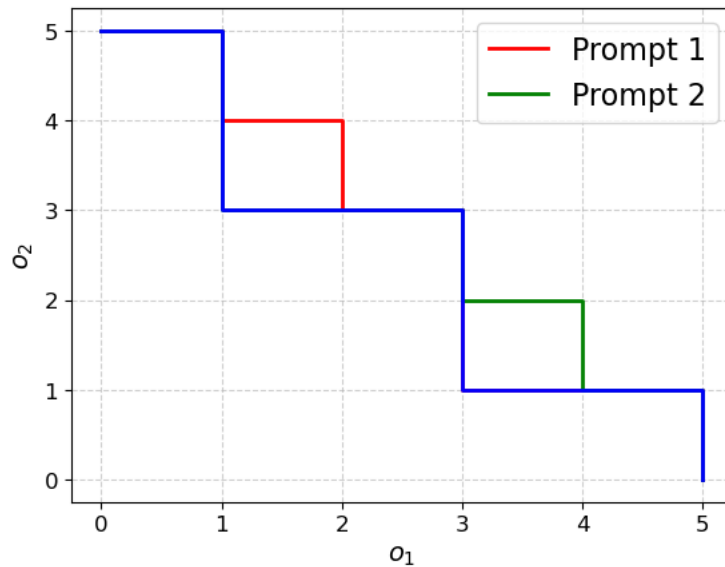
- RL or EA usually relies on single objective criterion.
- Example: HyperVolume.



HyperVolume Metric illustration in Objective space

Decision-making criteria


- Requires accurate objective values
- Sensitive to non-linear transformations.



Effect of Non-Linear Transformation on Hypervolume Metric

Challenges of Multi-Objective Optimization in Text Generation

- Accurate objective values is often unavailable.
- Objectives may be measured by different metrics which have non-linear relationships.



Non-Linear relationship between different metrics.

ParetoPrompt

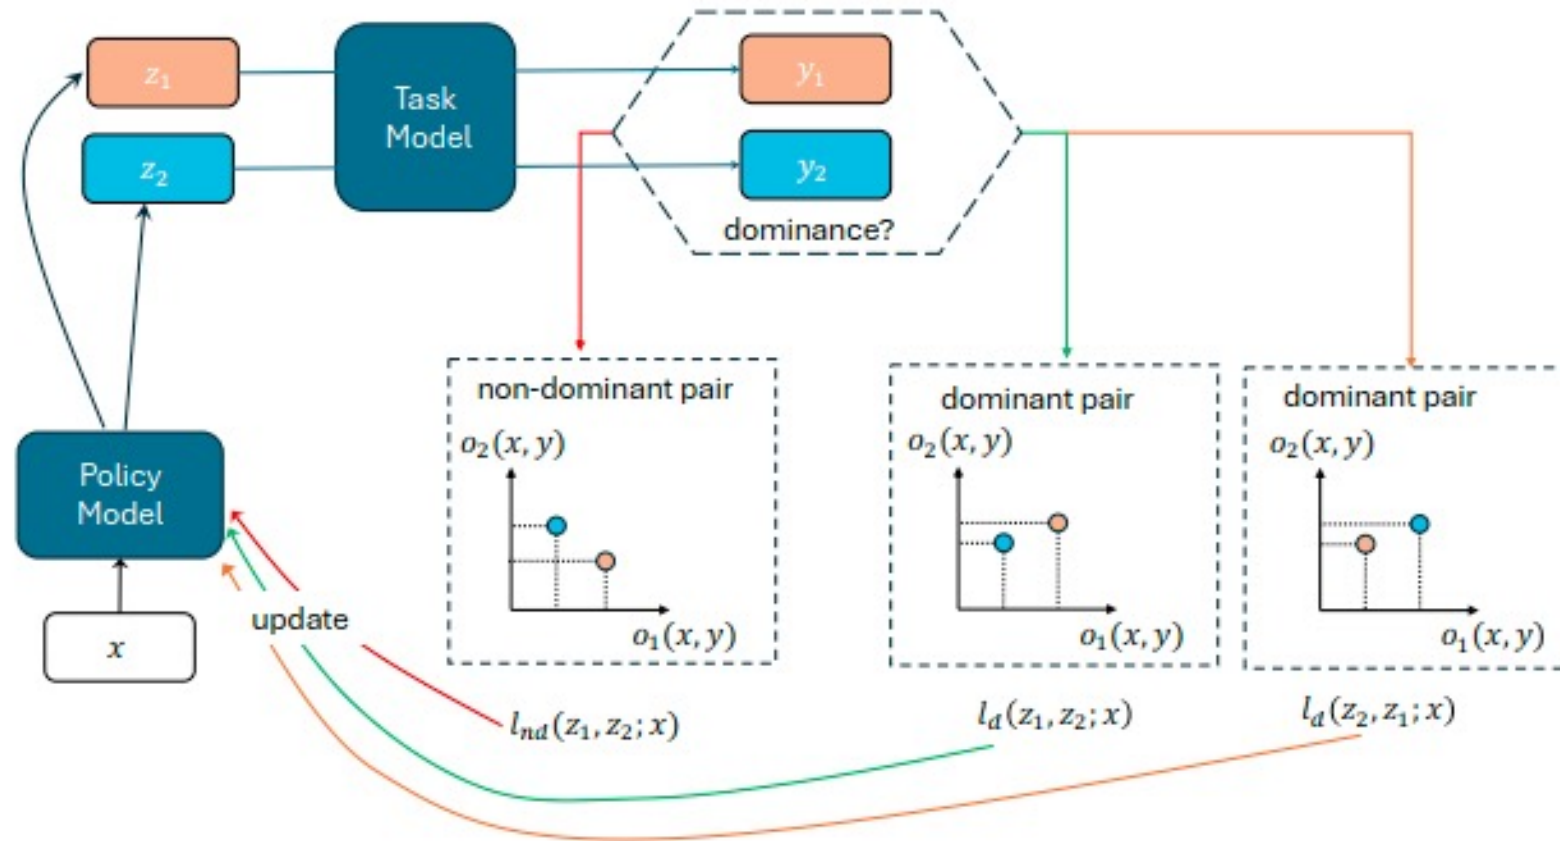
We propose ParetoPrompt, a **multi-objective** prompt optimization method driven by RL.

Our algorithm optimizes prompts based solely on the relative ranking of corresponding objectives, without requiring **precise objective values**.

- **Advantages:**

- Avoids assumptions about how to combine objectives.
- Remains *robust* to changes in how each objective is scored (non-linear transformations, etc.).

ParetoPrompt Algorithm



ParetoPrompt Workflow

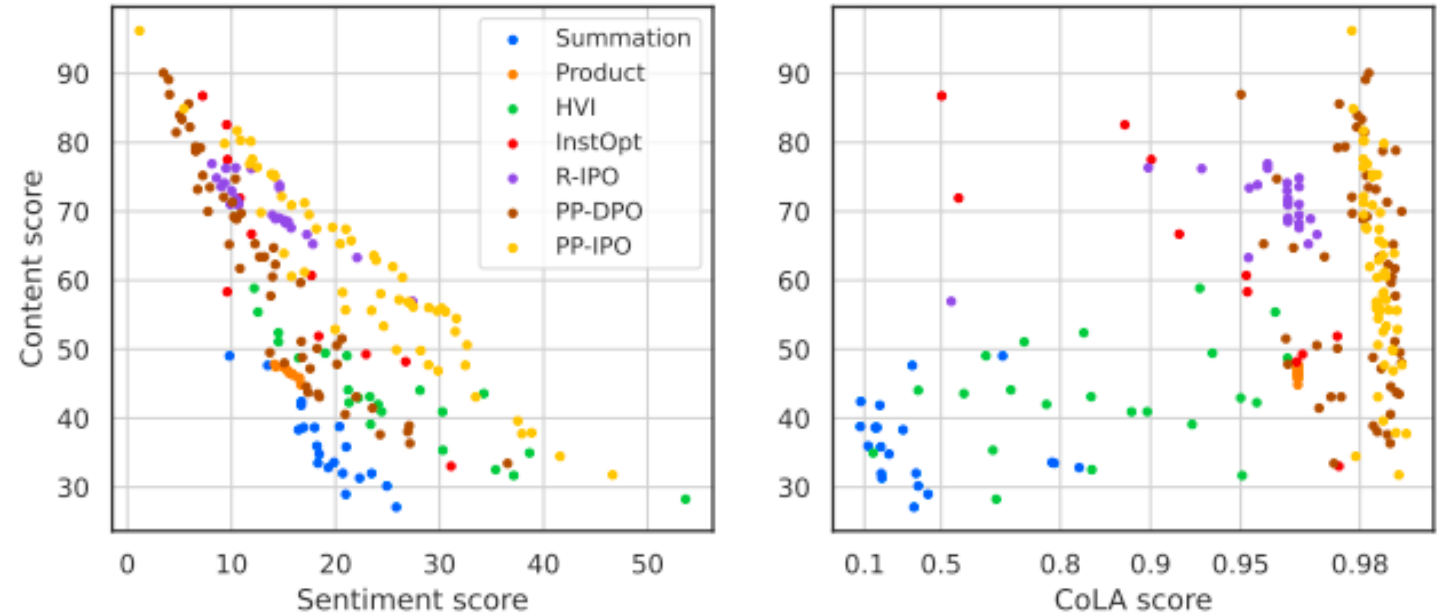
Few-shot Text Classification with Misaligned Metrics

Dataset	Metric	Summation	Product	HVI	InstOpt	R-IPO	PP-DPO	PP-IPO
MR	C-HV	76.5(8.4)	63.3(5.2)	77.2(6.9)	70.6(9.1)	75.5(8.5)	80.9(4.3)	83.0(4.6)
	P-HV	54.4(11.1)	46.5(11.2)	55.3(12.0)	67.2(0.1)	61.0(5.1)	66.8(11.5)	68.3(7.1)
	Diff-HV	22.1	16.8	21.9	3.4	14.5	14.1	14.7
SST-5	C-HV	32.0(2.1)	27.4(1.3)	36.2(0.8)	31(3.7)	29.5(7.0)	34.1(2.5)	35.3(0.5)
	P-HV	23.6(3.1)	20.4(1.2)	24.4(2.5)	32.6(1.9)	29.9(1.3)	34.3(0.4)	34.0(0.5)
	Diff-HV	8.4	7.0	11.8	-1.6	-0.4	-0.2	1.3
Yahoo	C-HV	39.8(1.9)	27.5(1.2)	42.2(2.6)	40.7(3.9)	24.9(1.3)	47.1(1.7)	48.2(0.5)
	P-HV	21.2(4.8)	17.0(7.4)	28.3 5.4	31.5(4.2)	20.6(4.3)	31.3(5.2)	35.5(2.7)
	Diff-HV	18.6	10.5	13.9	9.2	4.3	15.8	12.7
Yelp-5	C-HV	35.5(2.8)	27.0(6.7)	40.4(2.6)	24.1(1.4)	40.0(3.7)	41.3(3.1)	40.6(2.1)
	P-HV	21.9(2.1)	19.7(4.8)	20.5(3.1)	35.1(2.5)	37.1(4.4)	39.7(2.5)	39.8(2.3)
	Diff-HV	13.6	7.3	19.9	-11.0	2.9	1.6	0.8

THREE-OBJECTIVE Text Style Transfer task

Three objectives:

1. Style
2. Content similarity
3. Prompt fluency



Comparative Performance Across Three-Objective Prompt Optimization Methods

ParetoPrompt

We prove that even without the specific objective value in the prompt optimization:

- ParetoPrompt achieves comparable performance with the algorithms relying on specific objective values.
- ParetoPrompt robust performance even when the training metric differs from the evaluation metrics used during testing.

Thank you

Guang Zhao
Email: gzhao@bnl.gov