No Loss, No Gain: Gated Refinement and Adaptive Compression for Prompt Optimization

Wenhang Shi, Yiren Chen, Shuqing Bian, Xinyi Zhang,

Kai Tang, Pengfei Hu, Zhe Zhao, Wei Lu, Xiaoyong Du

Renmin University of China

NeurIPS 2025

November 03, 2025

1. Background: Automatic Prompt Optimization (APO)

Optimize the prompt using Optimizer LLMs to improve the performance of Base LLM on the target task.

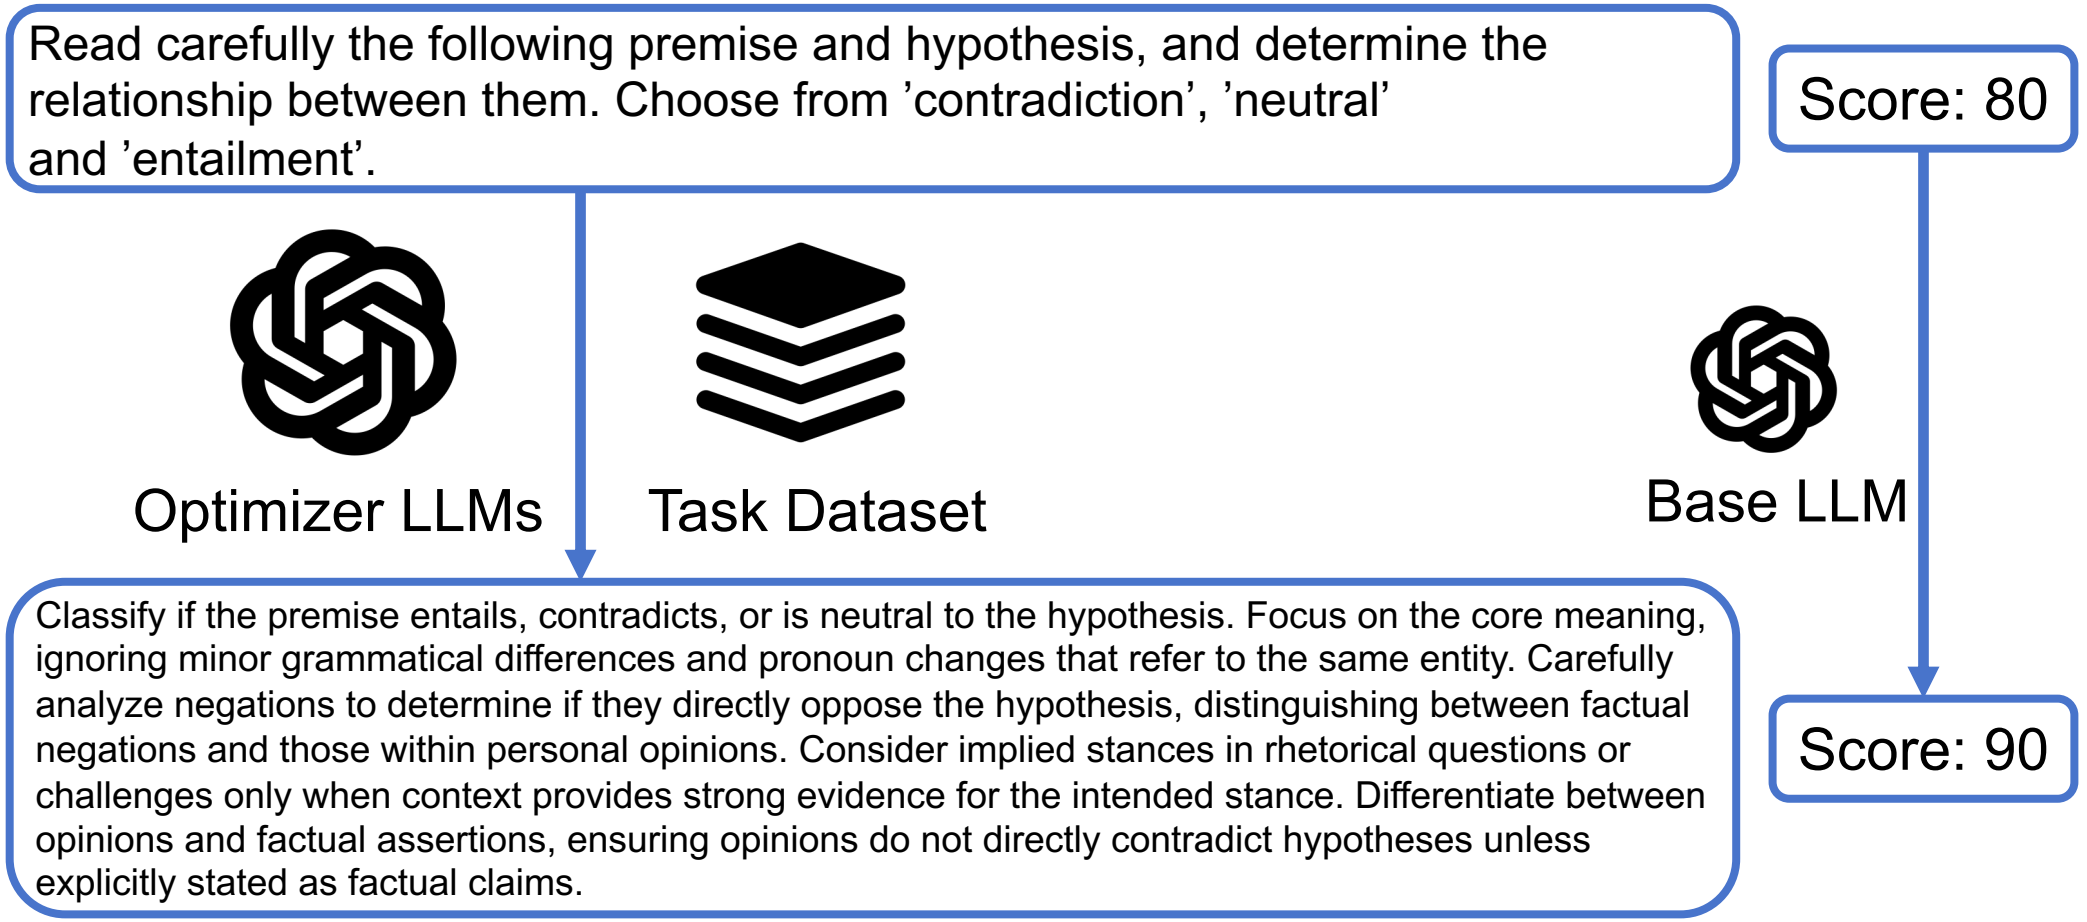

2. Previous Methods

Core Strategy: Employ LLMs as optimizers to iteratively expand and select prompt candidates based on validation set performance.

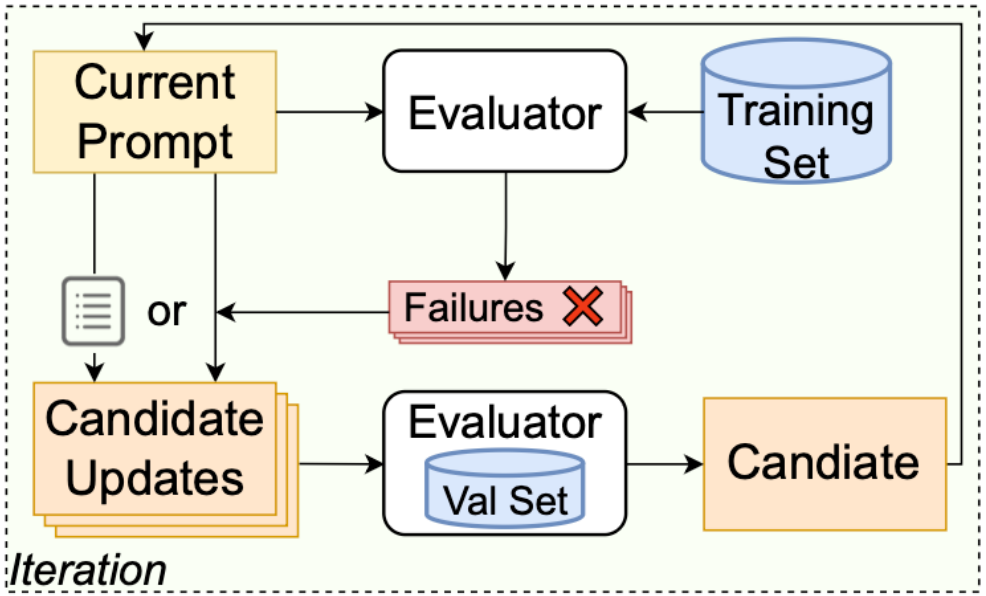

(a) Expansion and Selection

- **Search-Based**: Generates prompts using heuristics such as text edits or paraphrasing

Limitation: Lack clear optimization guidance -> random, semantically similar modifications

- **Reflection-Based**: Generate prompts based on analyses of failed training samples

Limitation: Overly aggressive and biased updates -> prompt overcorrection, semantic drift

2. Previous Methods

Optimization Problems

1. **Inefficiency**: Generate a large number of prompts at each step to secure improvement

2. Prone to Local Optima: Performance plateaus after only a few update steps

- **Search-Based**: hard to achieve progress in discrete prompt space via minimal changes
- **Reflection-Based**: incorporate instance-specific information, which lacks generalization

3. Our Method: Gated Refinement and Information Distillation (GRACE)

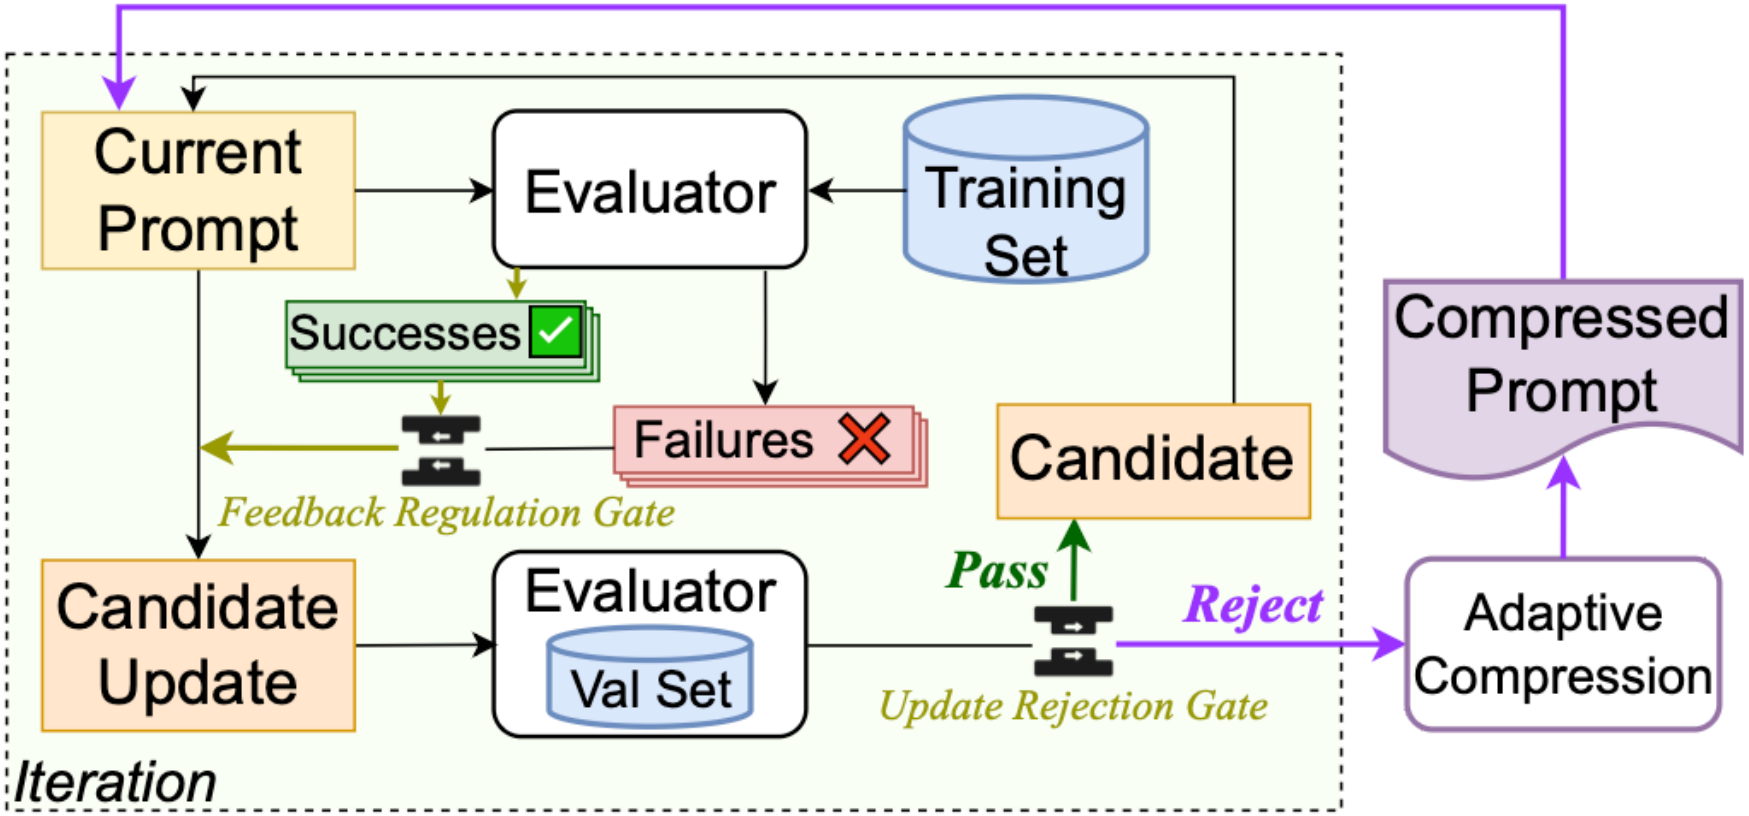

1) Generate an effective update via a feedback regulation gate
2) Selectively adopt updates via an update rejection gate
3) Activate adaptive compression when optimization stagnates (rejecting K updates)

3.1 Gated Refinement: Two-Stage Information Filtering for Stable Updates

1) Generate an effective update via a feedback regulation gate

- Using successful samples as a regularization gate
- Leverage known effective patterns to control and balance the content and magnitude of update
- Refine strong or biased signals from failed samples

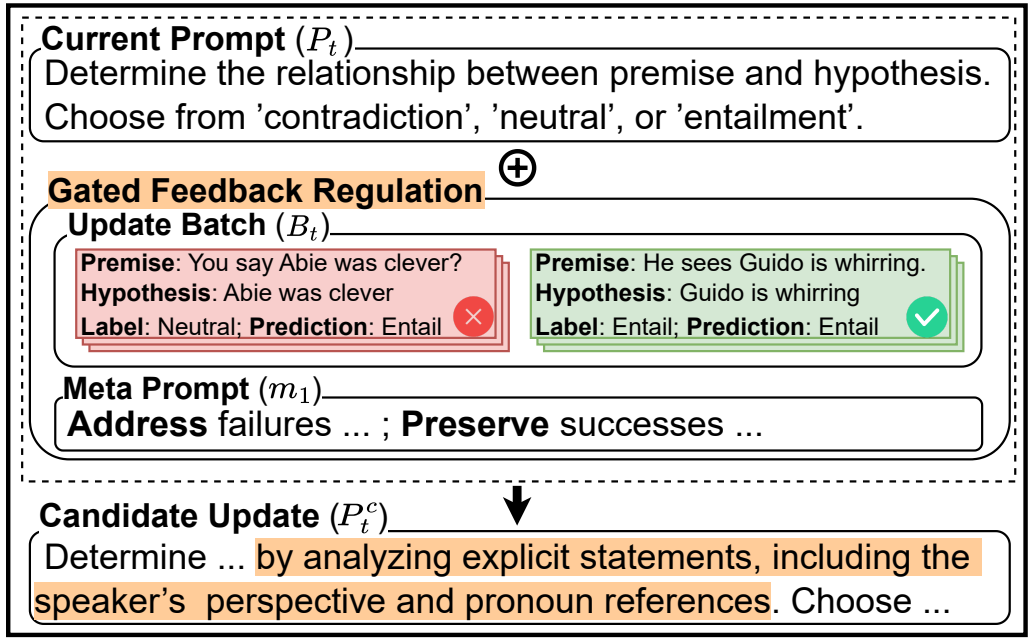

3.1 Gated Refinement: Two-Stage Information Filtering for Stable Updates

2) Selectively adopt updates via an update rejection gate

- Reject updates failing to improve performance

$$\mathcal{P}_{t+1} = \underset{\mathcal{P} \in \{\mathcal{P}_t, \mathcal{P}_t^c\}}{argmax} f_{\mathcal{B}}(\mathcal{P}, D_{val})$$

- Discard unnecessary or detrimental updates
- Ensure only meaningful and beneficial information is incorporated into the prompt updates

3.1 Gated Refinement: Two-Stage Information Filtering for Stable Updates

The two-stage gating mechanism introduces controlled information loss to enable more targeted and stable prompt updates, thereby enhancing optimization efficiency

However, during optimization process

- Enriched information shifts from general guidance to case-specific details
- Over-specification traps the optimization in local optima

3.2 Adaptive Compression: Information Distillation for Escaping Local Optima

3) Activate adaptive compression when optimization stagnates (rejecting K updates)

- Simplify by merging or removing redundant elements
- Abstract away concrete, instance-specific instructions into more broadly applicable guidance.

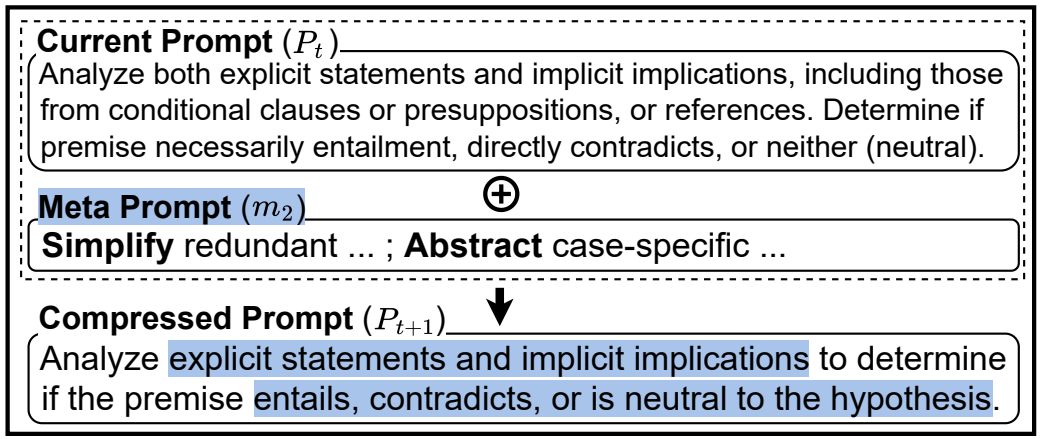

- Help escape from local optima
- Provide a better optimization starting point
- Open up new directions for gated refinement

4. Experiments

4.1 Comparisons with baselines on three types of tasks

- **Better Final Performance**: 4.7%, 4.4% and 2.7% improvements over SoTA methods

	BBH	Domain-Specific Tasks				General NLP Tasks			
	Avg.	NCBI	Biosses	MedQA	Avg.	Subj	TREC	CB	Avg.
Task (ZS)	77.45	60.83	72.50	84.75	72.69	64.20	66.20	89.29	73.23
Task (FS)	72.73	64.90	65.00	79.25	69.72	85.00	69.80	92.86	82.55
CoT (ZS)	77.74	60.02	72.50	85.75	72.76	59.10	64.80	94.64	72.85
CoT (FS)	79.62	64.69	67.50	84.50	72.23	84.00	73.40	94.64	84.01
EvoPrompt	81.15	70.96	70.00	84.75	75.24	92.30	85.40	89.29	89.00
OPRO	85.51	69.47	72.50	85.50	75.82	94.60	86.40	89.29	90.10
APO	88.14	**73.83**	67.50	85.50	75.61	94.80	90.60	96.43	93.94
PromptAgent	89.42	71.81	75.00	86.00	77.60	91.50	90.20	94.64	92.11
GRACE	**94.13**	**73.83**	**85.00**	**86.50**	**82.00**	**95.70**	**94.20**	**100**	**96.63**

4. Experiments

4.2 Optimization Convergence Curve

Higher Efficiency

- Faster Convergence: 80 vs. 300 prompts
- More Sustained Updates: escaping local optima

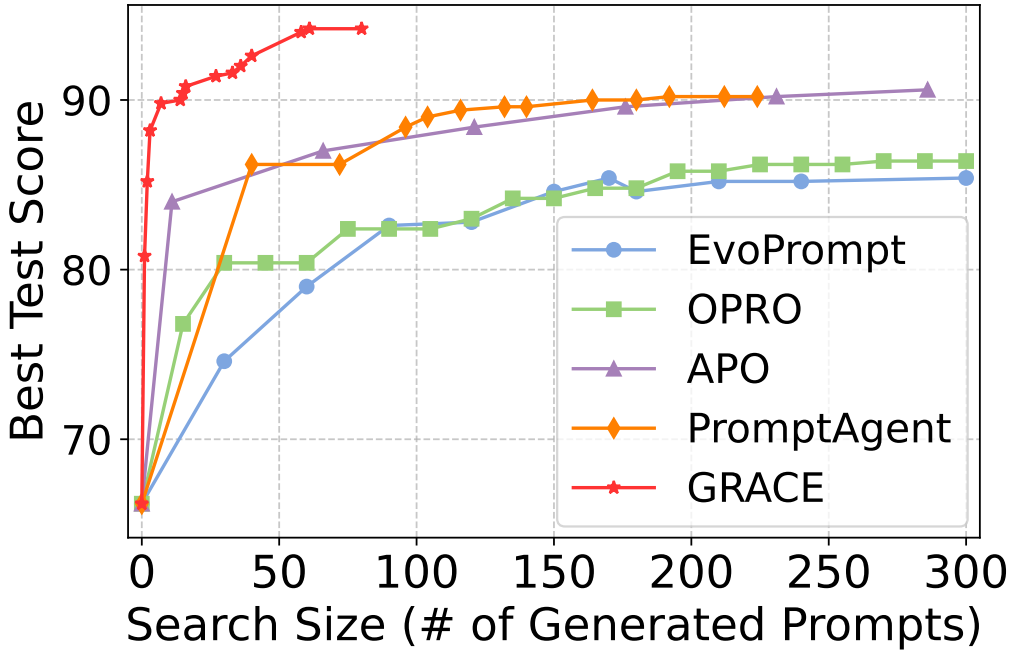

4. Experiments

4.3 Optimization Trace

Green: beneficial -> update
Red: harmful -> rejected
Blue: compressed

State	Prompt	Score
Step 0 Initial	Read carefully the following premise and hypothesis, and determine the relationship between them. Choose from 'contradiction', 'neutral', or 'entailment'.	89.3
Step 1 Parent 0	Read ... Determine their relationship by analyzing explicit statements, including the speaker's perspective and pronoun references. Choose from ...	91.1
Step 2 Parent 1	Read ... Determine their relationship by analyzing whether the premise directly supports (entailment), contradicts, or neither (neutral). Pay attention to: Whether statements are presented as facts, hypotheticals, or opinions; Whether questions or possibilities in the premise justify the hypothesis.	92.9
Step 3 Parent 2	Read ... Analyze both explicit statements and implicit implications, including those from conditional clauses, presuppositions, or references. Determine if the premise necessarily supports (entailment), directly contradicts (contradiction), or neither (neutral).	**94.6**
Step 4 Parent 3	Read ... Analyze both explicit statements and implicit implications, including those from conditional clauses (noting their pragmatic implications), presuppositions, and references (resolving coreference and speaker identity). Distinguish between factual assertions and subjective opinions. Determine if ...	92.9
Step 5 Parent 3	Read ... Analyze both ... Distinguish between assertions of belief/opinion and objective facts. For conditional statements, evaluate whether the premise provides evidence beyond hypothetical scenarios. When resolving references or coreferences, rely solely on explicit information. Determine if ...	91.1
Step 9 Parent 3	Read ... Analyze both explicit statements and implicit implications to determine if the premise entails, contradicts, or is neutral toward the hypothesis.	94.6
Step 13 Parent 12	Read ... Analyze both explicit statements and implicit implications, including beliefs, hypothetical scenarios, and conditional statements. Determine if ... by evaluating factual support, direct opposition, or lack of relevant information.	**96.4**

5. Conclusion

- We propose GRACE, an efficient prompt optimization framework, which strategically introduces information loss to achieve stable and sustained prompt improvements and effectively escape local optima.
- GRACE demonstrates substantial gains in both performance and efficiency across 11 tasks spanning three practical and distinct domains.