

# Microsoft proposes a new way to train AI: eliminating lengthy prompts while maintaining performance.

Microsoft's OPCD method helps AI memorize knowledge without lengthy prompts, reducing costs and increasing efficiency.

During the development of AI applications, businesses often have to use **very long system prompts** to fine-tune the model's behavior. These prompts contain everything from internal knowledge and operating rules to problem-specific instructions.

However, when deployed on a large scale, this approach begins to reveal its limitations. Having to 'cram' so much information into each AI call significantly increases response times and also drives up the processing cost per query.

To address this problem, researchers at Microsoft proposed a new training method called **On-Policy Context Distillation (OPCD)**. The core idea is simple: instead of repeatedly providing the same amount of information in the prompt, train the model to **'remember' that knowledge internally**.

OPCD leverages the model's own responses during training to 'distill' knowledge. This allows the model to learn how to respond appropriately to specific applications without losing its general capabilities.

## Why do long prompts become a burden?

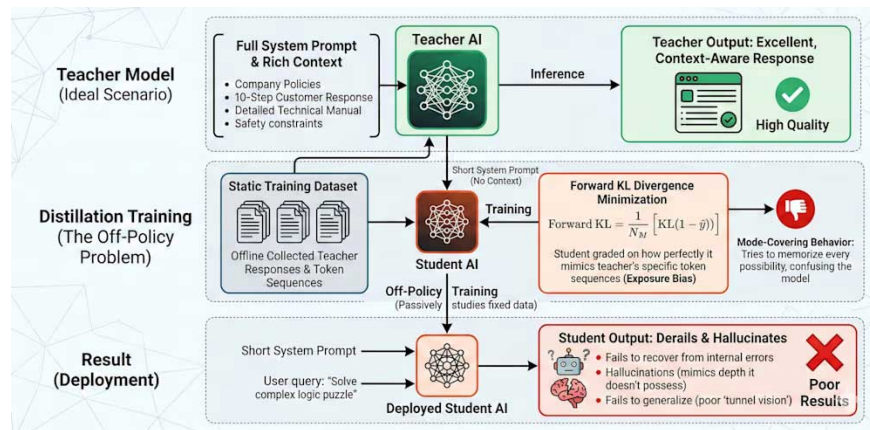
In-context learning mechanisms allow AI behavior to be adjusted in real time without retraining. However, the drawback is that all this knowledge is only temporary.

This means that each time the system is used, it has to reload all the data, such as company policies, technical documents, or customer history. Repeating this process not only slows down the system but also easily causes information interference.

Distillation solutions operate on a 'teacher-student' model. The 'teacher' model is provided with a full, lengthy prompt and can generate accurate, detailed answers. Meanwhile, the 'student' model only receives simple questions without the full context.

The 'student's' task is to observe how the 'teacher' responds and learn from it. Over time, it will 'compress' all the logic and knowledge from the long prompt into its own parameters. When put into practical use, the 'student'

model can operate much faster without having to repeat the initial massive amount of data.



## The problem with the old method

Despite its effectiveness, the traditional distillation method still has limitations.

Firstly, there is the phenomenon of 'exposure bias'. During the learning process, the model only sees the correct answers from the 'teacher,' but in real-world operation, it must make its own decisions. Because it is not used to handling errors, the model easily 'goes astray.'

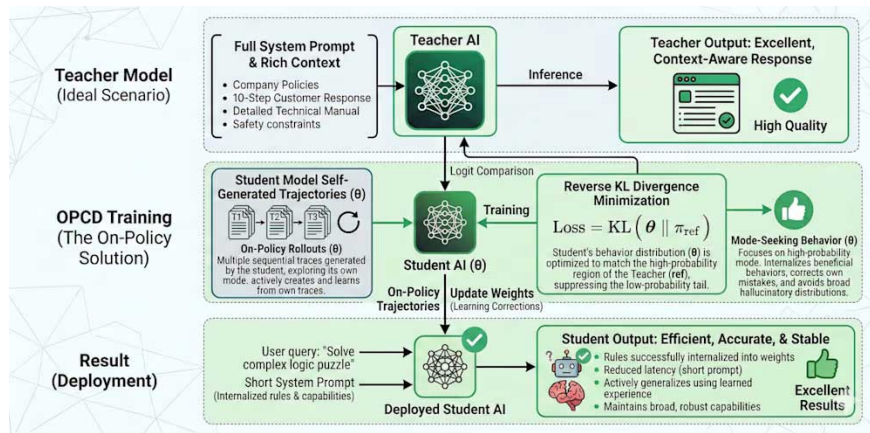
Secondly, the assessment method based on similarity to the 'mentor' causes the model to attempt to encompass too many possibilities. This makes the inferences scattered and inaccurate.

As a result, in practice, AI can produce 'illusions'—providing false information but appearing very confident, or struggling to handle new situations.

OPCDs are designed to overcome these weaknesses.

Instead of learning from a fixed dataset, the 'student' model will generate its own answers during the training process. At this point, the 'teacher' acts as a direct instructor, evaluating each step the 'student' takes.

At each token generation step, the system compares the 'student's' reasoning with how the 'teacher'—who has full context—would do it. As a result, the model learns not only the final outcome but also the thought process.



## Practical benefits for businesses

This new approach offers very clear benefits when implemented in practice:

The model no longer relies on long prompts, resulting in faster response times and significantly reduced processing costs. Simultaneously, its reasoning capabilities are more stable because it has been "trained" in conditions closer to real-world usage.

More importantly, businesses can build customized AI models to meet their specific needs without sacrificing performance or accuracy.

OPCD shows a new direction in AI development: **it's not about adding more data to the prompt, but about teaching the model to understand and remember information from within**.

If widely adopted, this approach could help make enterprise AI systems faster, cheaper, and more reliable in the future.

You finished reading the article "**Microsoft proposes a new way to train AI: eliminating lengthy prompts while maintaining performance.**" edited by the [TipsMake](#) team. We hope this article has provided you with many useful tech tips and tricks. You can search for similar articles on tips and guides. Thank you for reading and for following us regularly.