# Prompting Decision Transformer for Few-Shot Policy Generalization

Mengdi Xu[1], Yikang Shen[2], Shun Zhang[3], Yuchen Lu[2],
Ding Zhao[1], Josh Tenenbaum[4], Chuang Gan[3,5]

CMU[1], Mila[2], MIT-IBM Watson AI Lab[3], MIT[4], Umass Amherst[5]

# Few-Shot Policy Generalization

How to achieve **few-shot policy generalization** via only **offline data**?

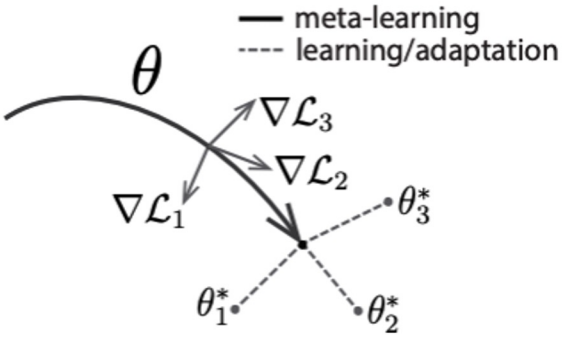


**Algorithm**:
MAML [Finn, 2017]

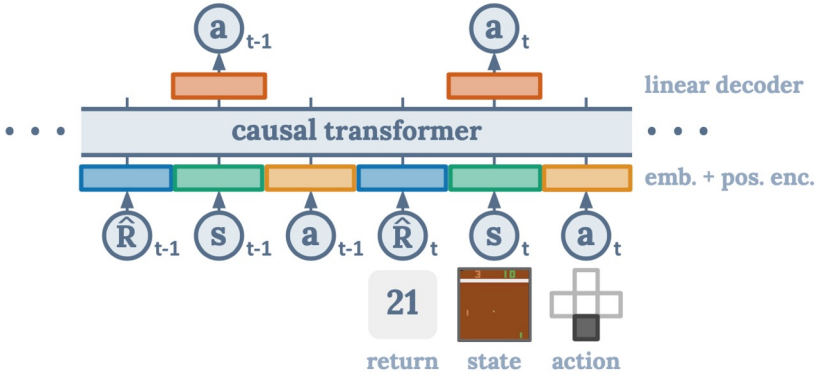


**Model Architecture**:
Decision Transformer [Chen, 2021]

***Will a better model architecture help achieve few-shot generalization?***

# Prompt Decision Transformer (Prompt-DT)

- **A policy conditioned on short trajectory prompts (2-15 steps) composed of few-shot demonstrations.**

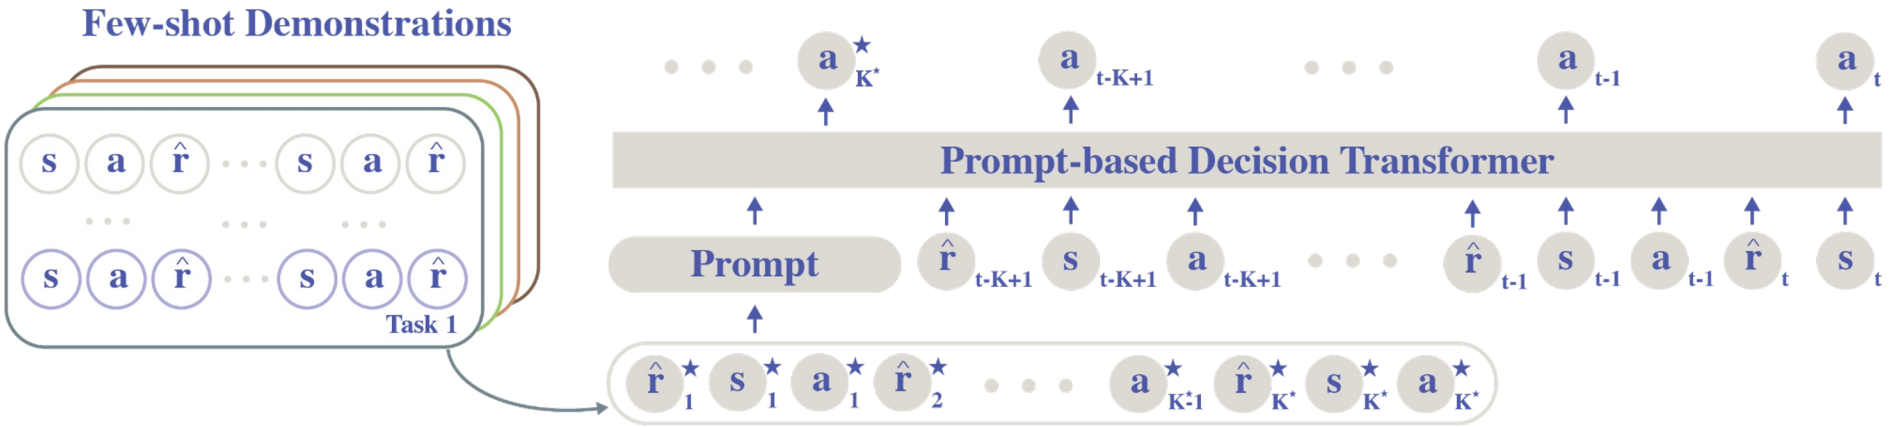


Prompt-DT could successfully **solve the unseen task** with a short trajectory prompt that only contains **2-15 timesteps**, demonstrating **strong data efficiency**.

# Prompt Decision Transformer

- Evaluation on MuJoCo fine-grained control benchmarks.


Prompt-DT


Prompt-DT

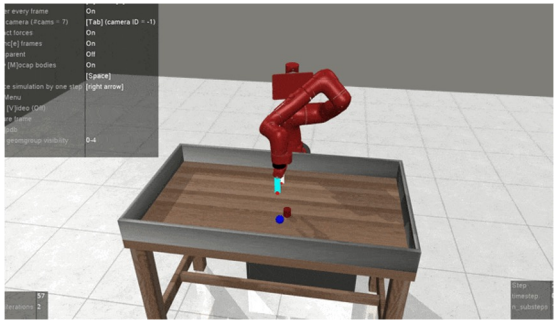

Prompt-DT


MT-ORL(no prompt)


MT-ORL(no prompt)

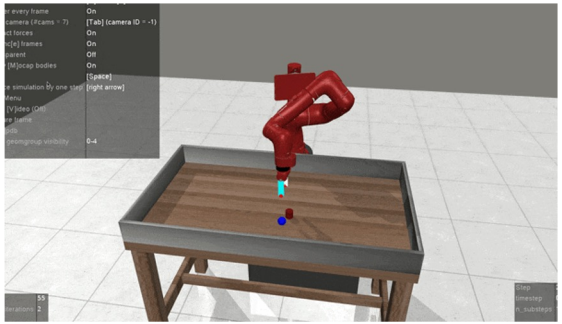

MT-ORL(no prompt)

# Prompt Decision Transformer

## Thank you for listening!

- See our paper for more about the training and testing algorithms, experiments, and ablations on the effect of prompt quality and quantity.

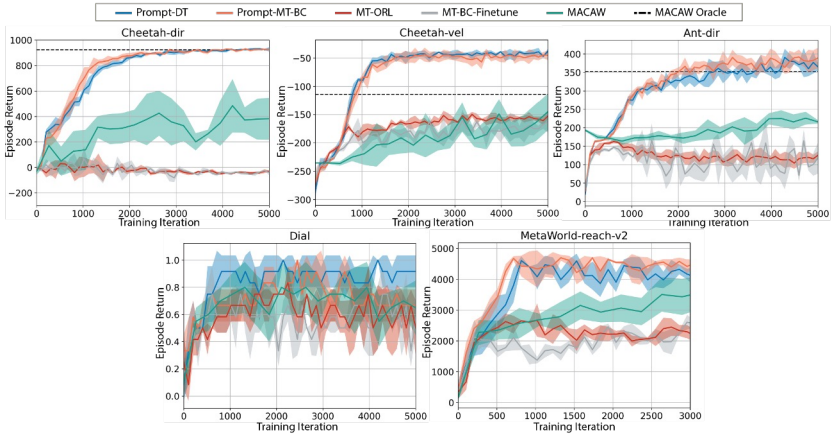


Project website:
mxu34.github.io/PromptDT/

Scan for visualization!