PoliFormer: Scaling On-Policy RL with Transformers Results in Masterful Navigators

Kuo-Hao Zeng, Zichen "Charles" Zhang, Kiana Ehsani, Rose Hendrix, Jordi Salvador,
Alvaro Herrasti, Ross Girshick, Aniruddha Kembhavi, Luca Weihs

poliformer.allen.ai

Scaling On-Policy RL Training with Transformers

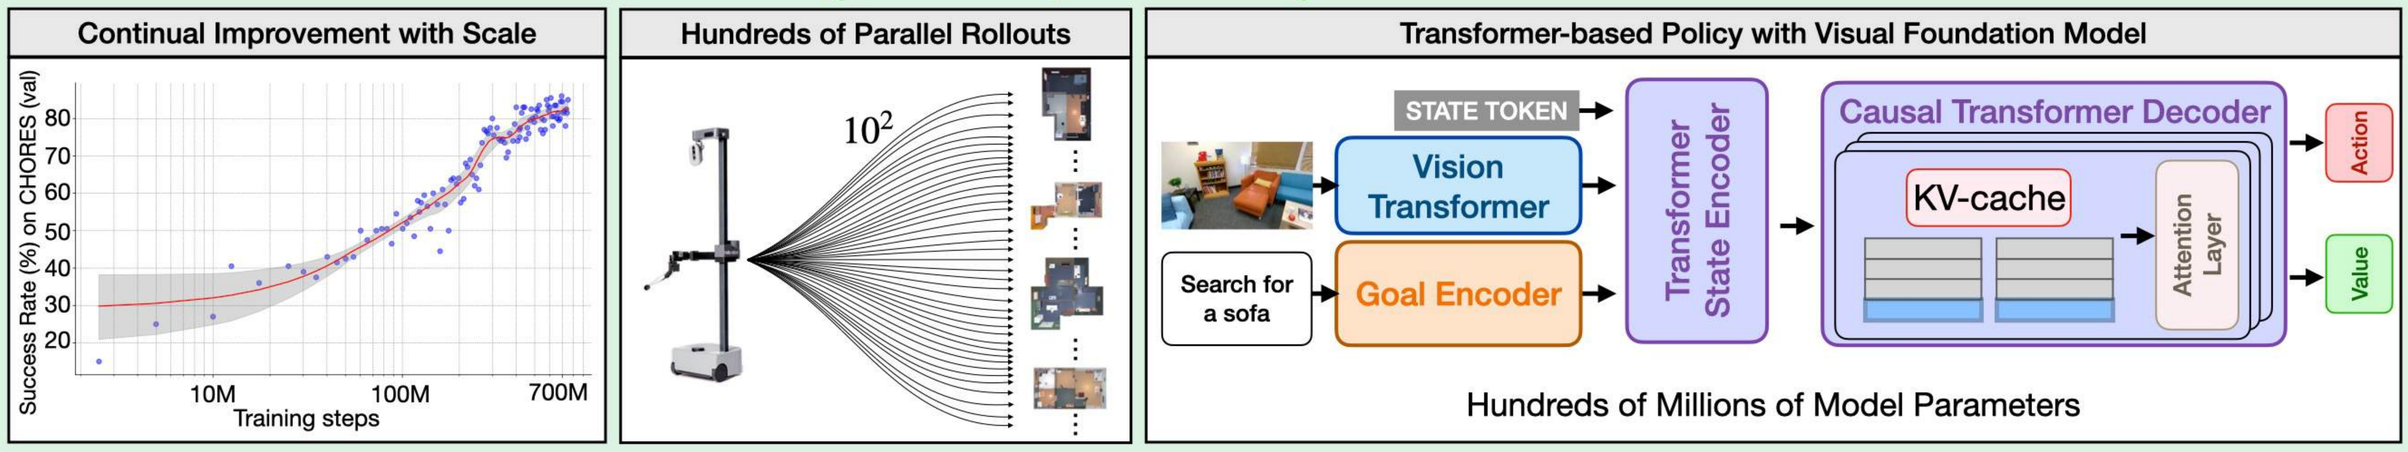

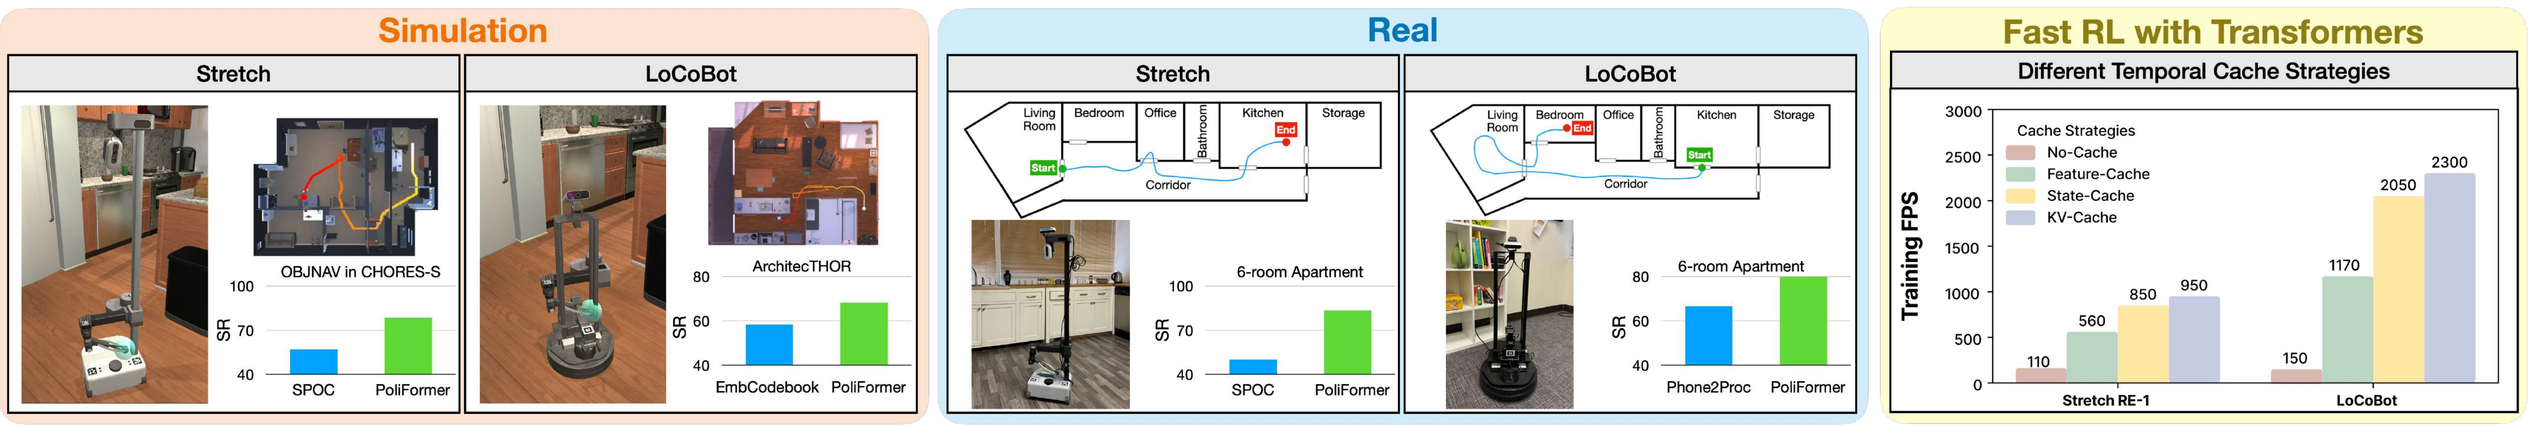

Zero-Shot Downstream Tasks

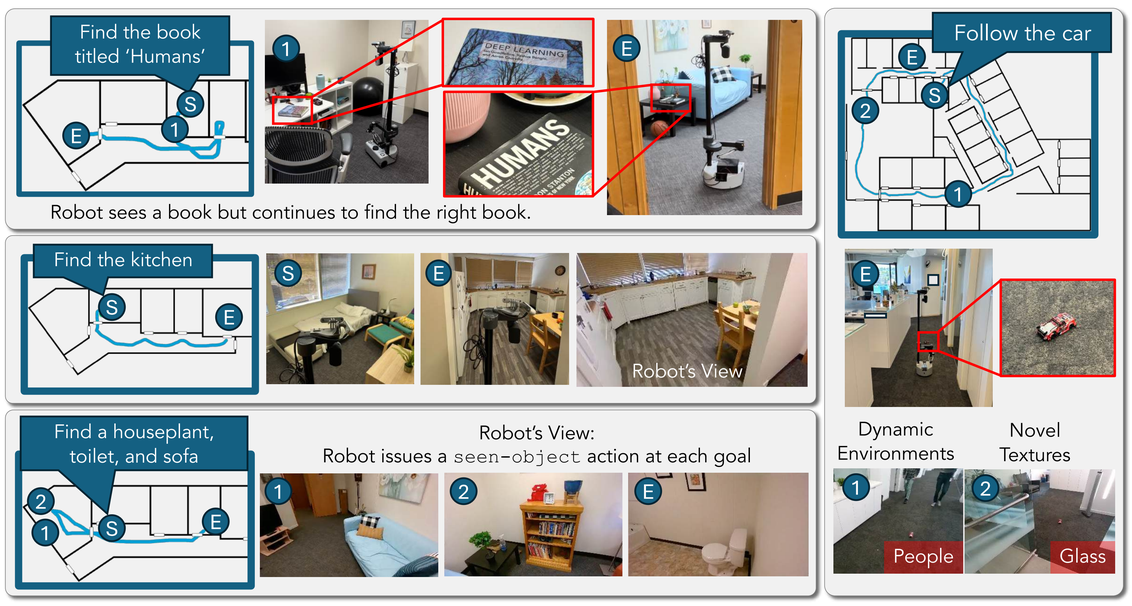

Scale Model and Data

Scale in Architecture

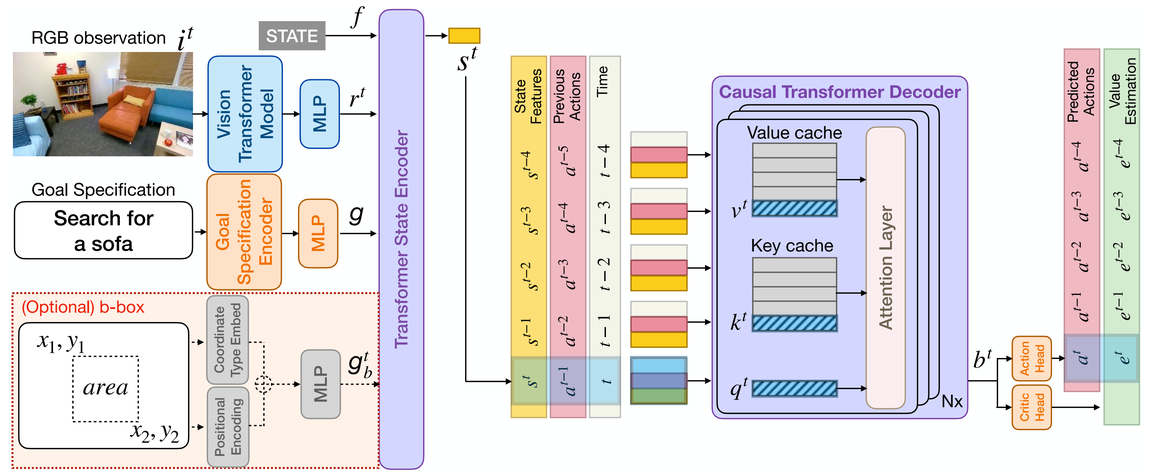

Vision Transformer Model and Transformer State Encoder

Causal Transformer Decoder with KV-Cache

Quantitative Results

Simulation

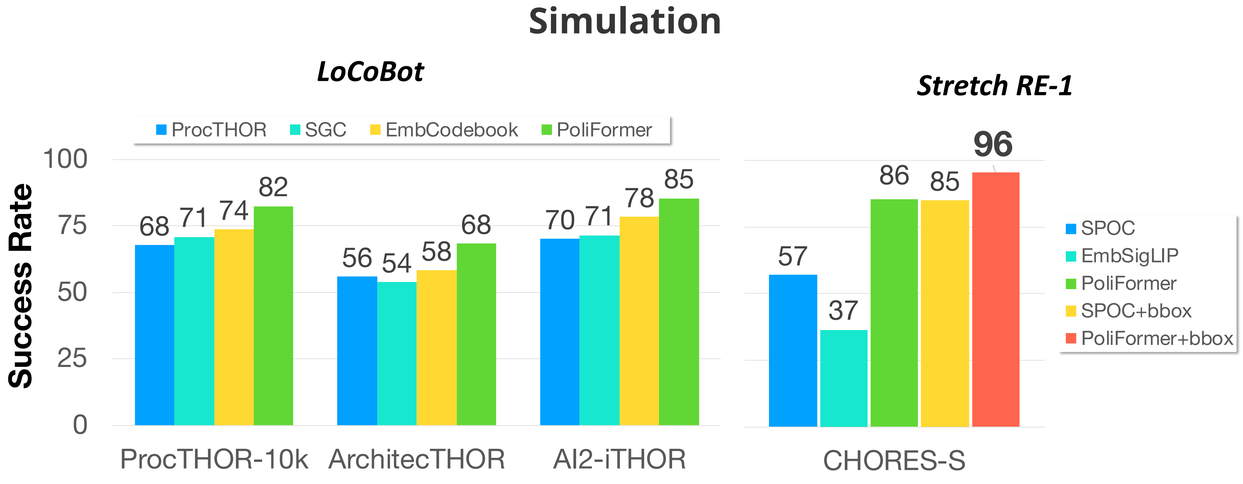

Real-World (zero-shot)

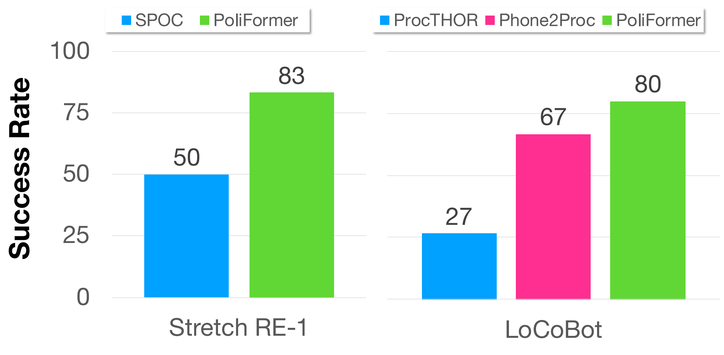

SoTA results in **simulation and the real-world**, across **multiple embodiments**.

Scale in Diverse Environment Interactions

Hundreds of parallel rollouts and large batch sizes lead to high training throughput and allow us to train for **700M** training interactions.

150k ProcTHOR houses with **~40k** Objaverse assets.