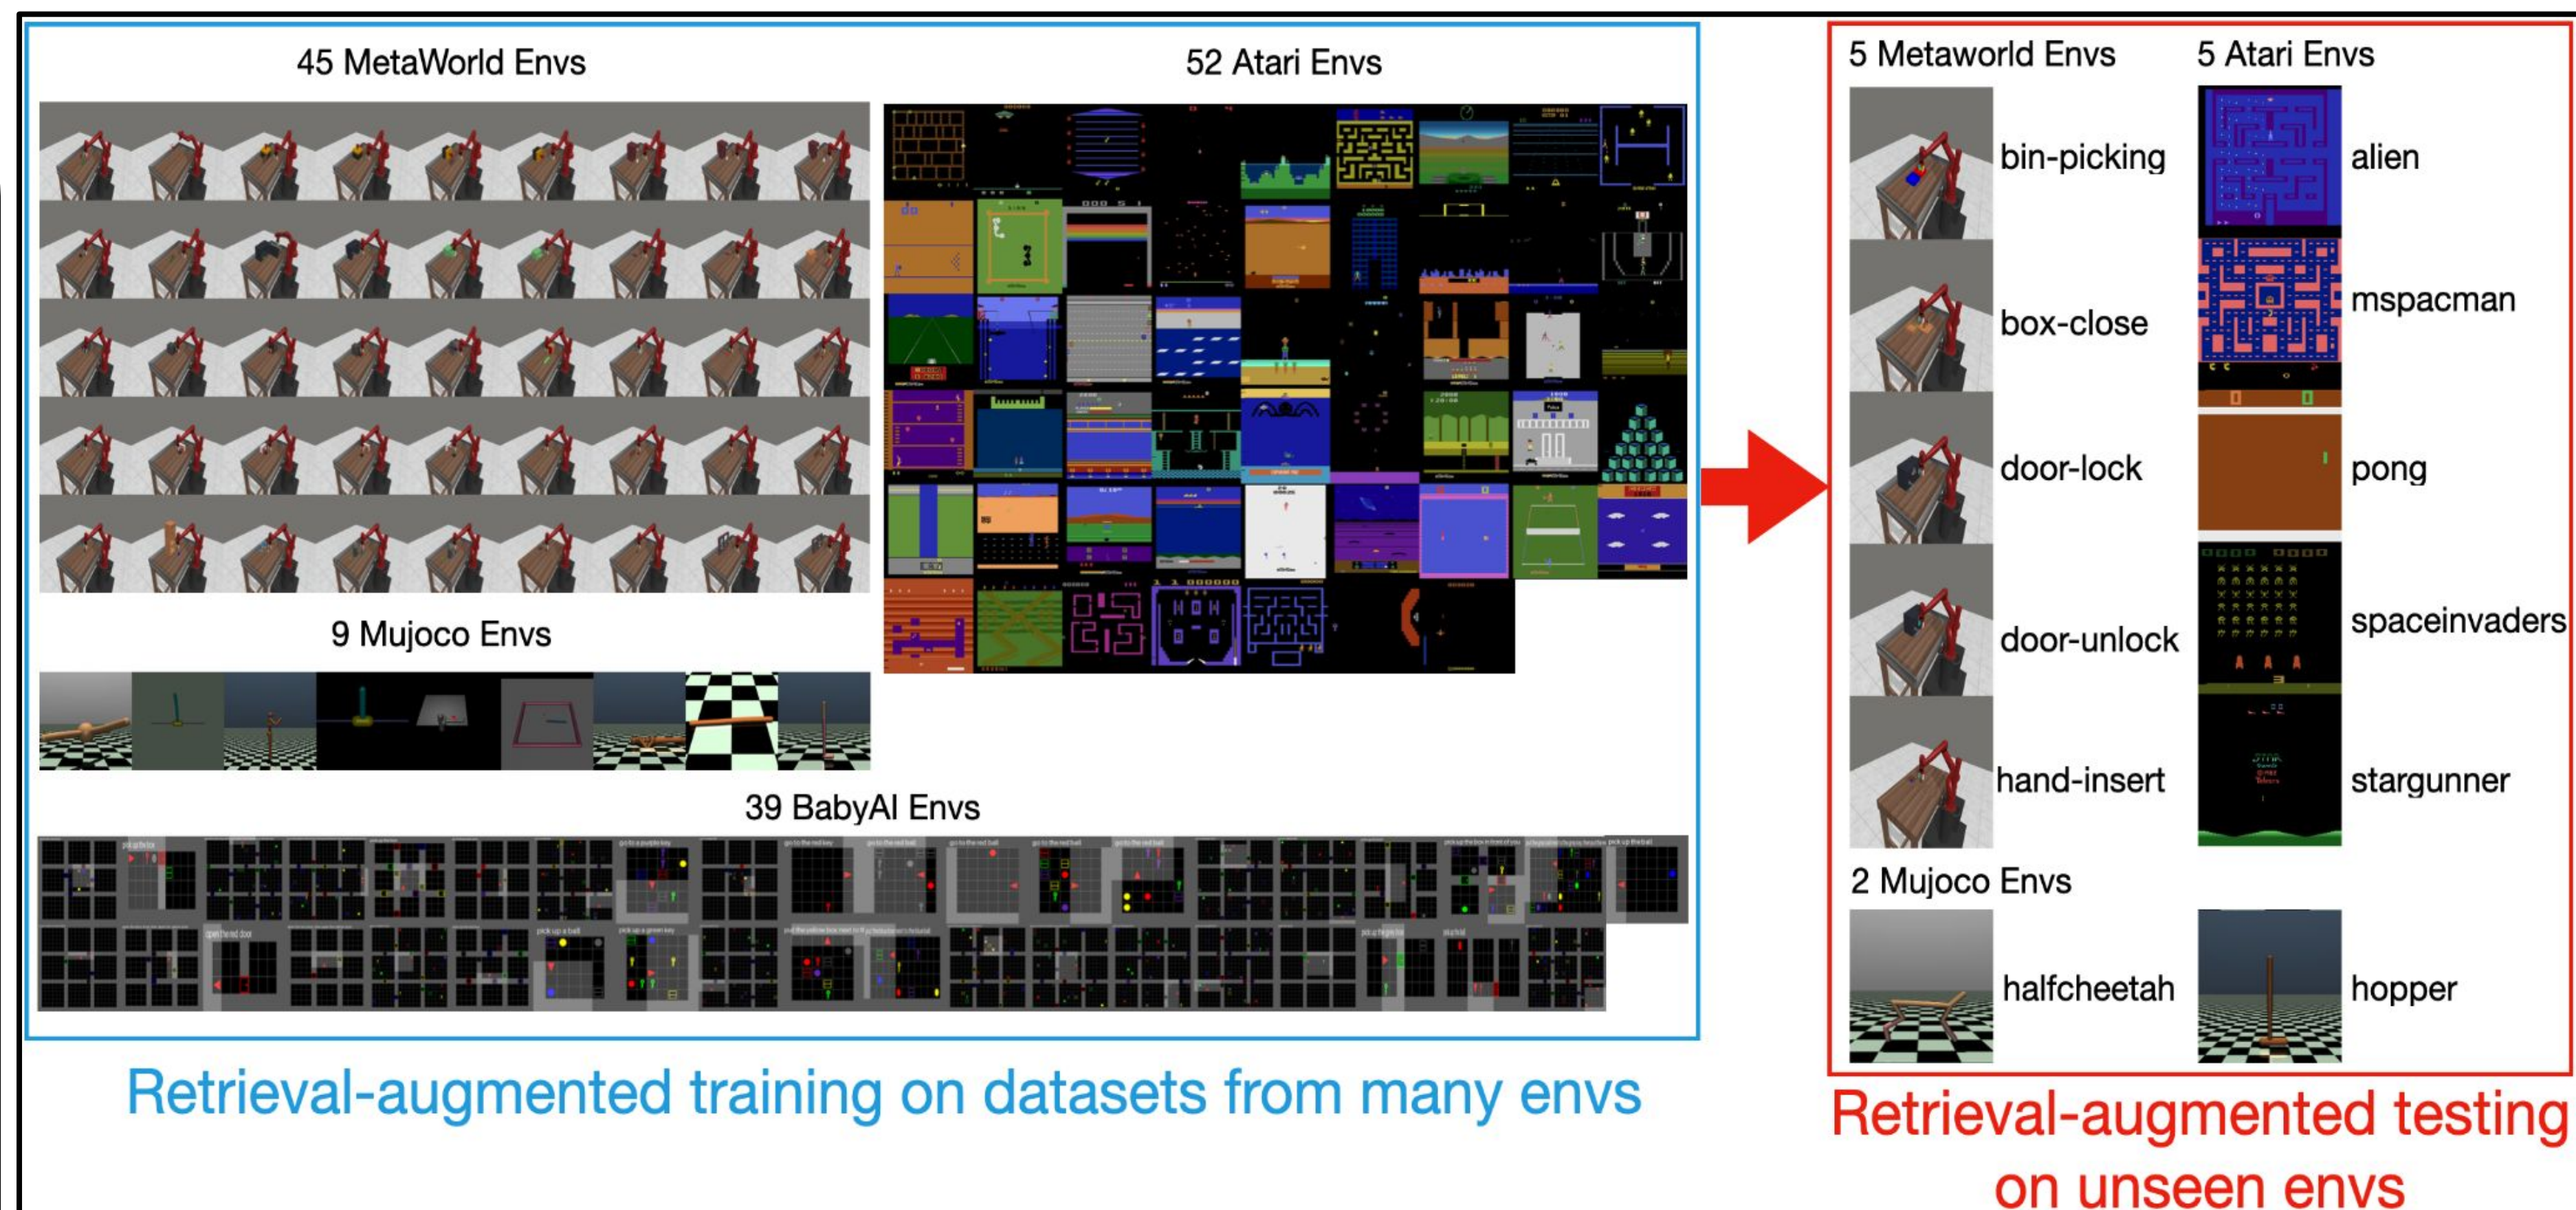


REGENT: A Retrieval-Augmented Generalist Agent That Can Act In-Context in New Environments

Kaustubh Sridhar, Souradeep Dutta, Dinesh Jayaraman, Insup Lee

Problem Setting

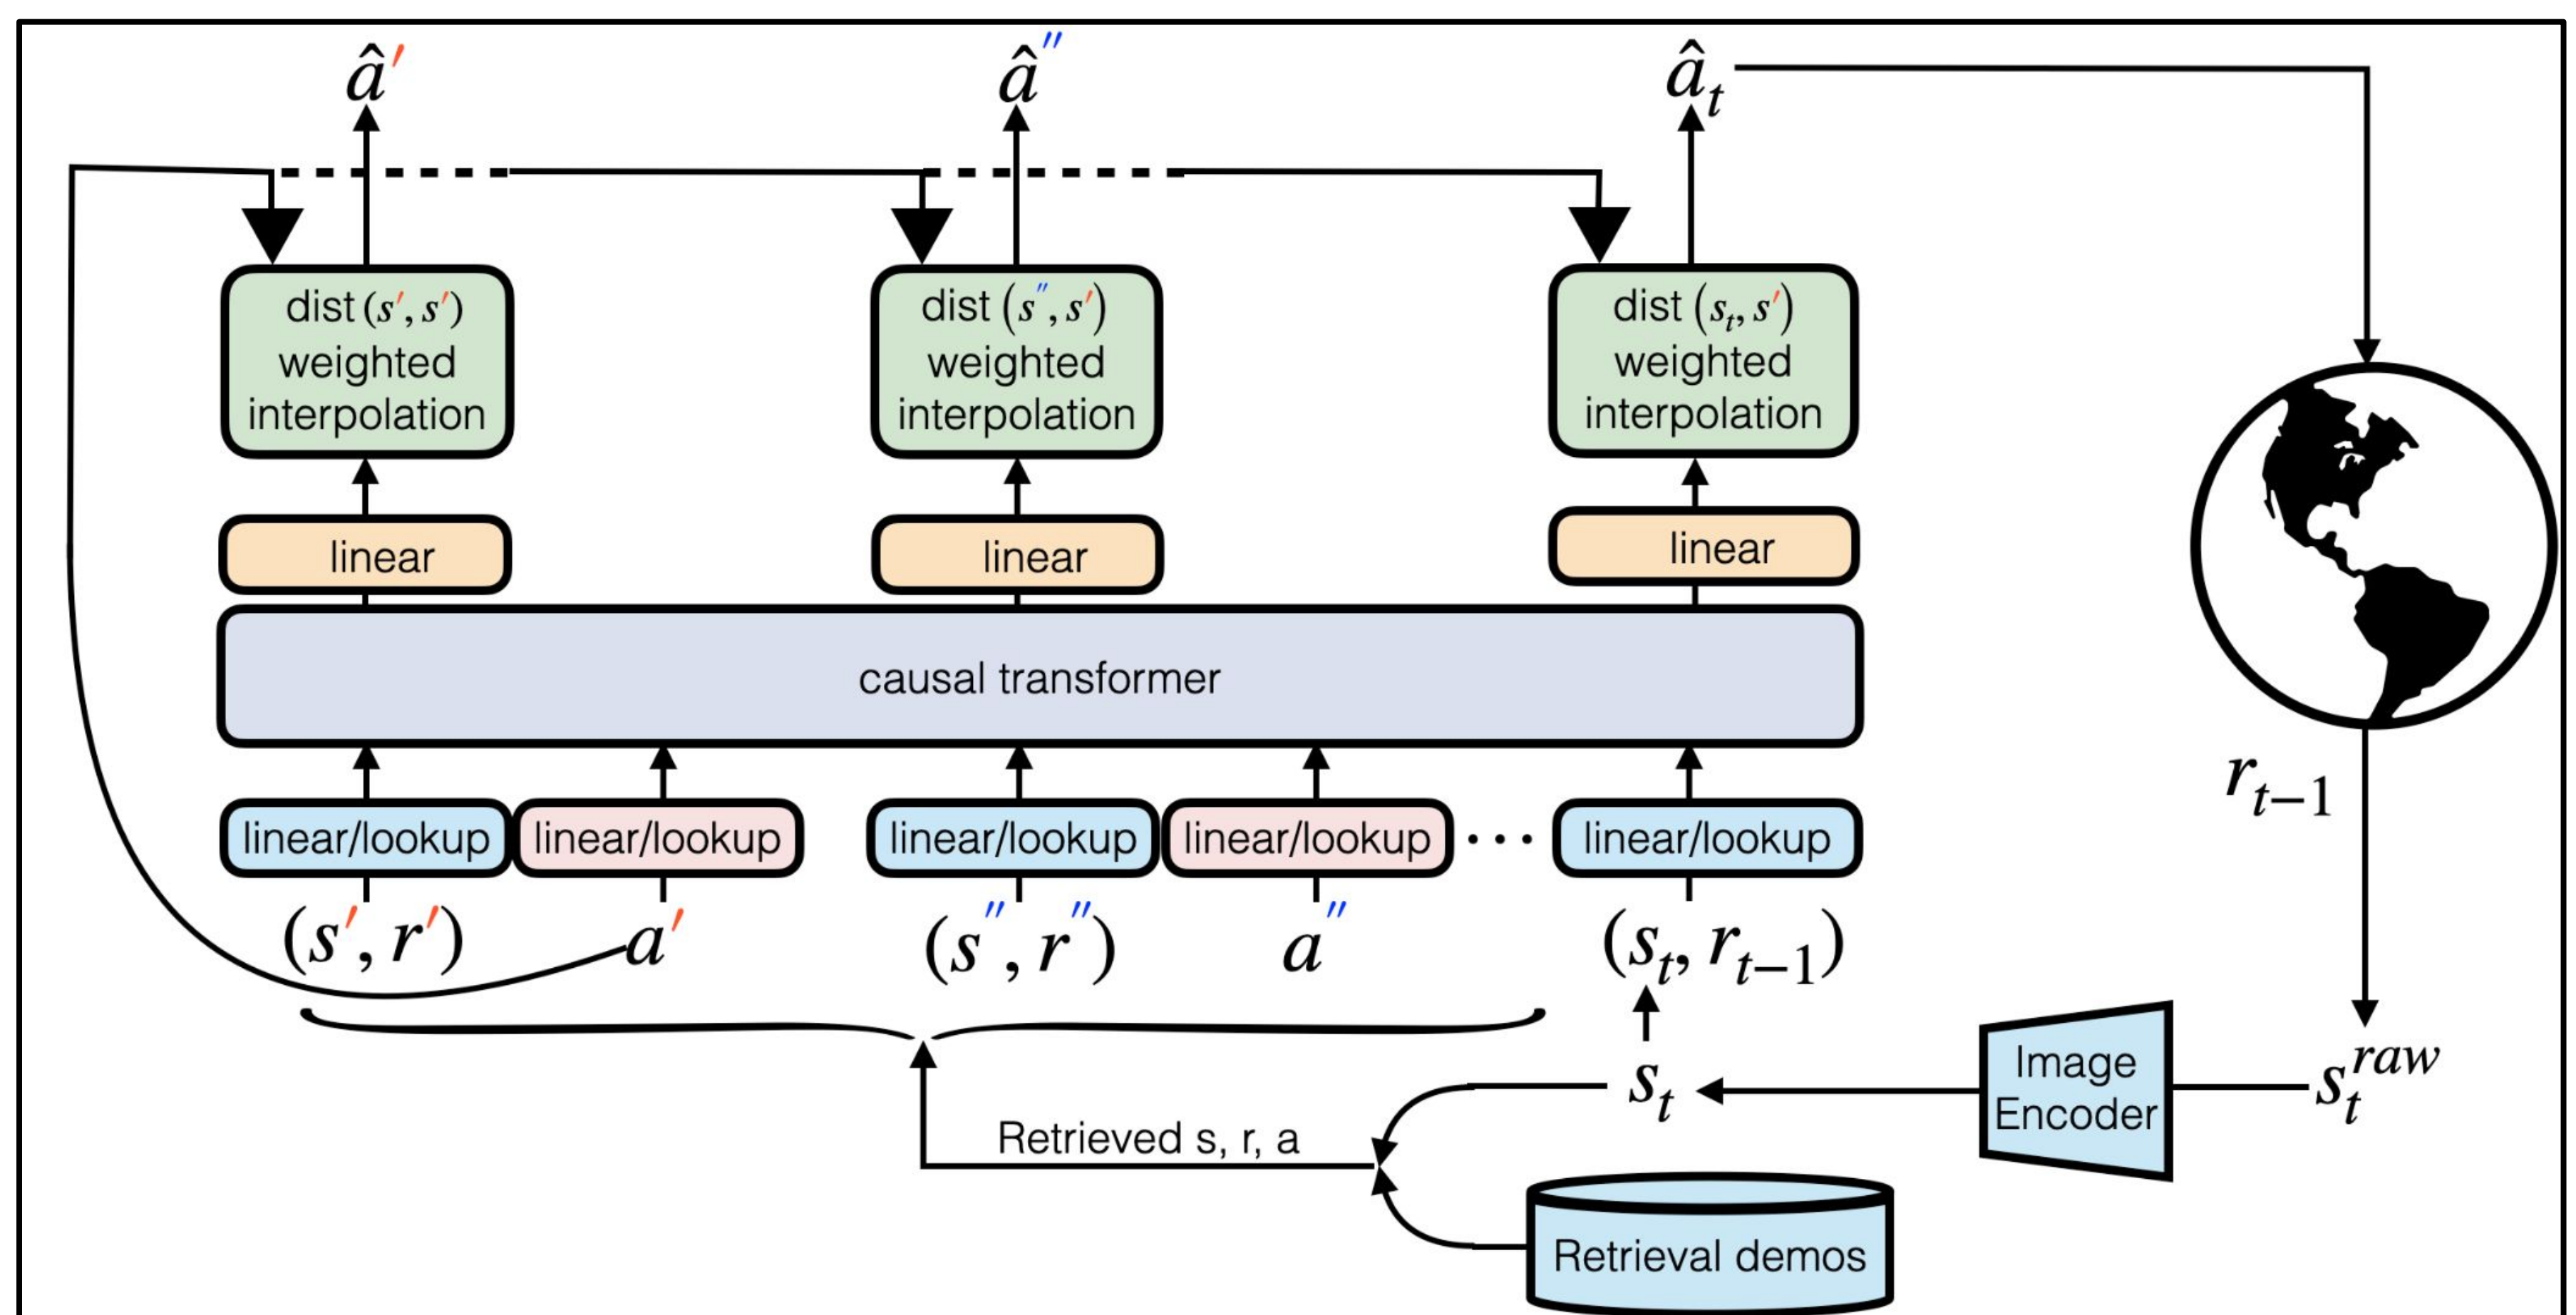


Is scaling current agent architectures the most effective way to build (in-context learning capable) generalist agents?

Key Idea: Retrieval

Retrieval offers a powerful bias for adaptation. A simple retrieval-based nearest neighbor agent (R&P) competes with SOTA generalist agents! From this starting point, we construct a retrieval-augmented agent (REGENT).

Architecture of REGENT & Retrieve-&-Play



Key Ingredients

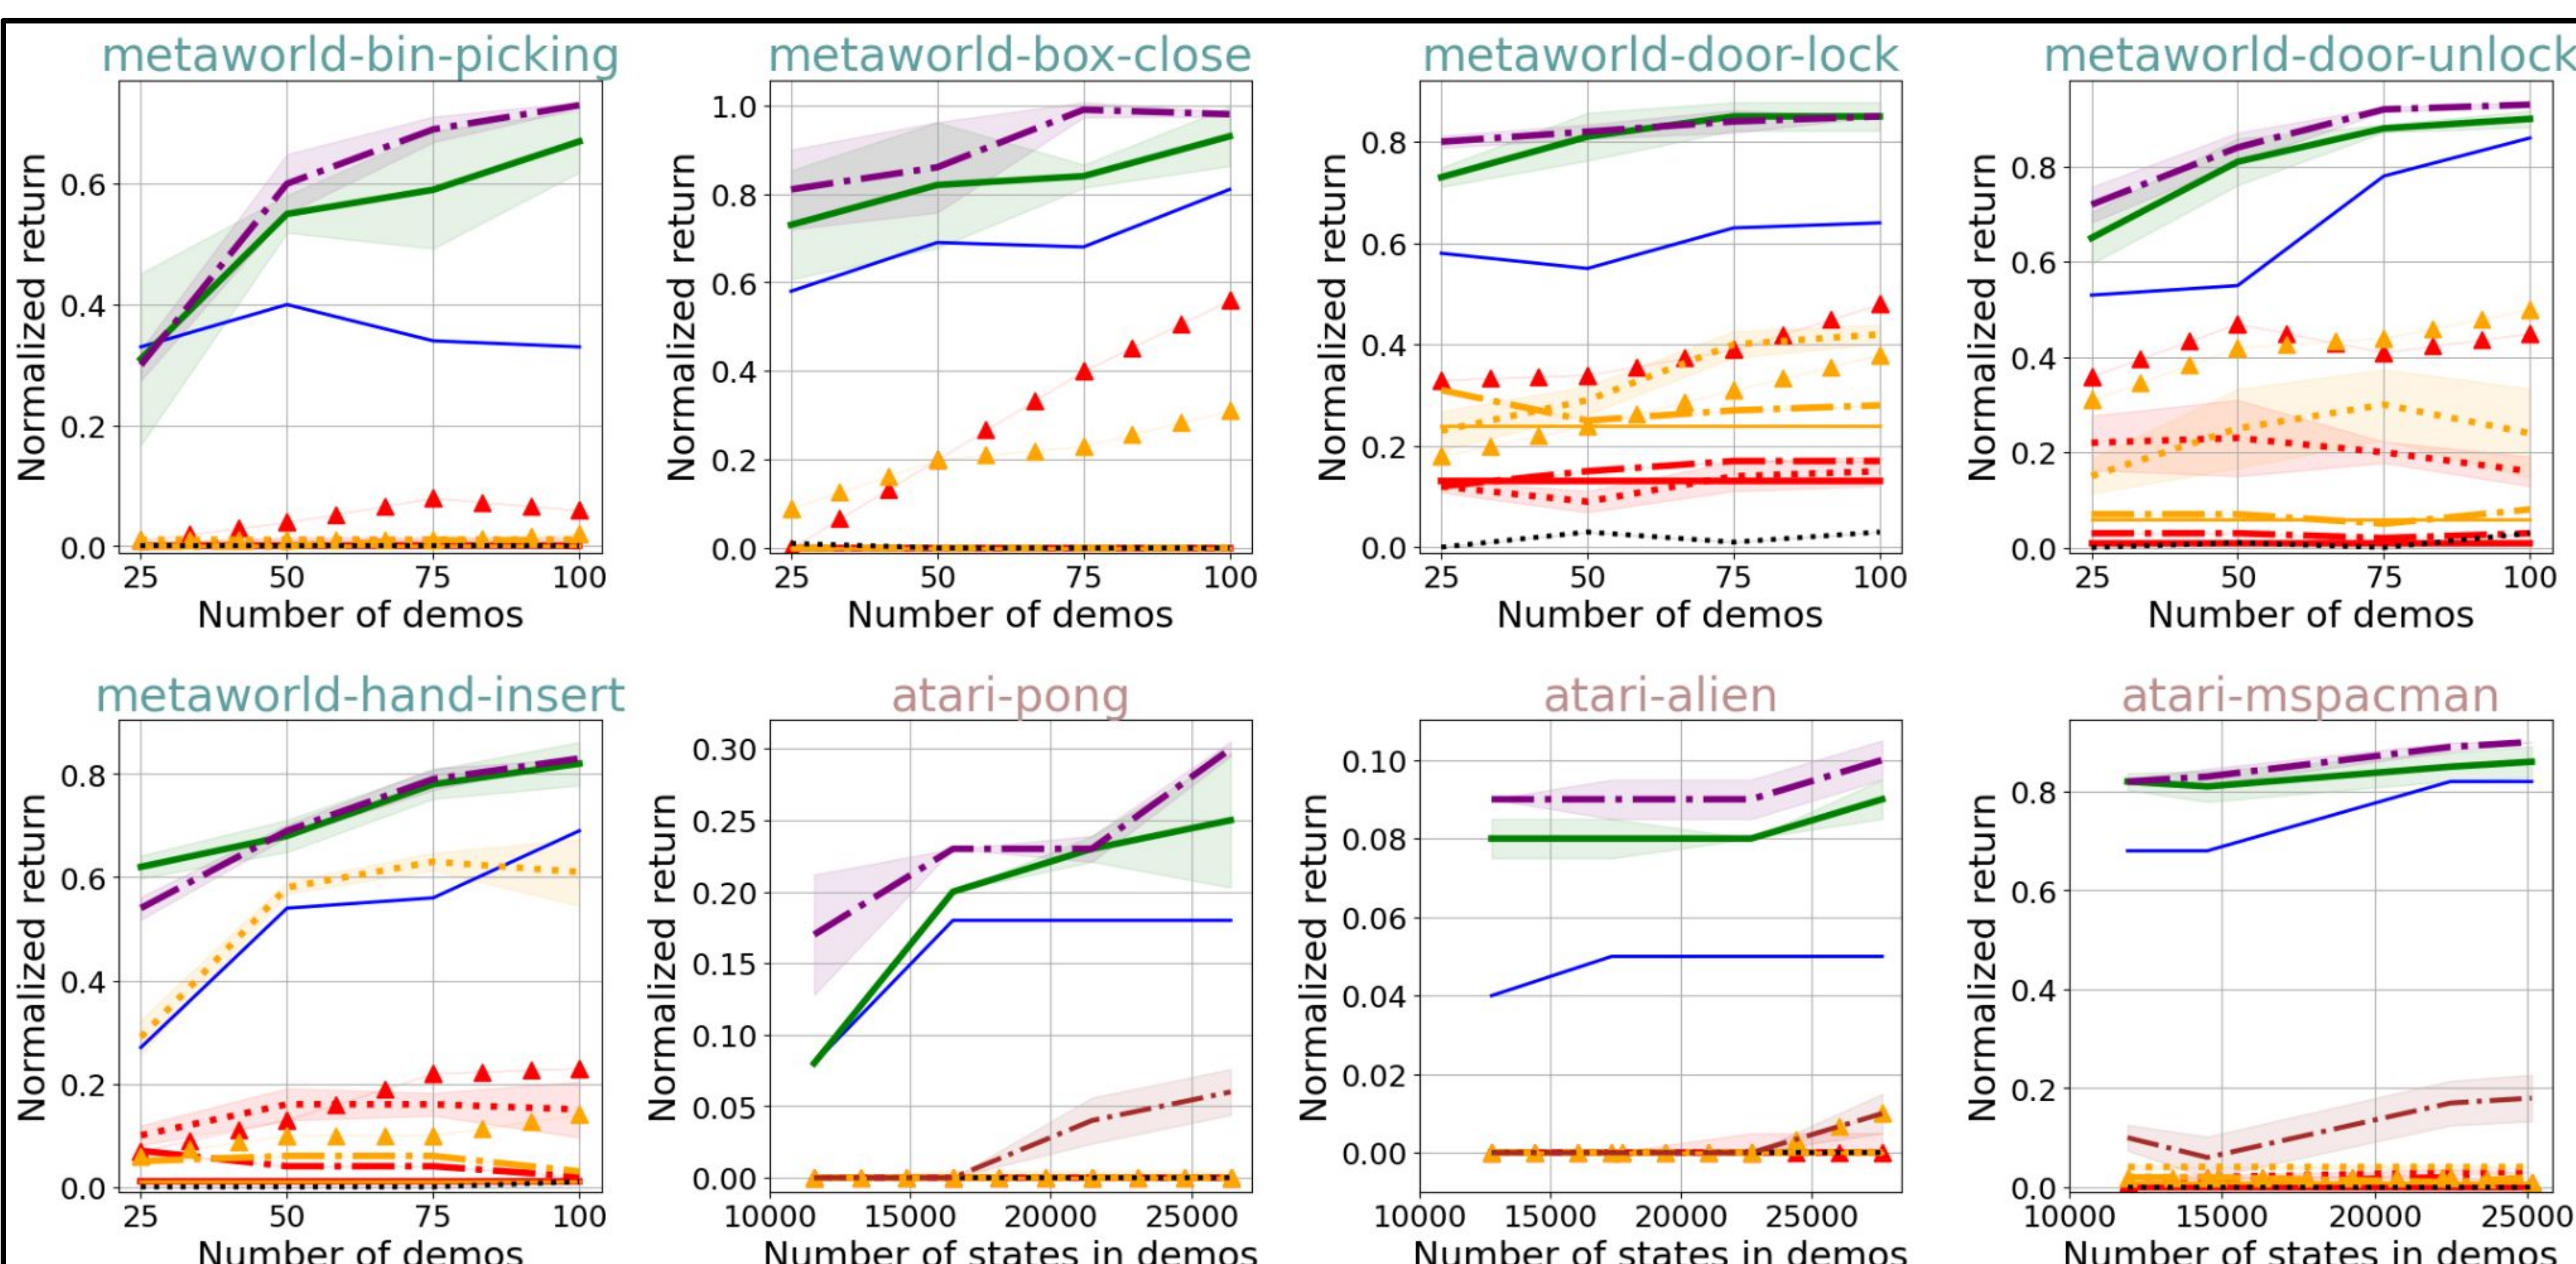
- interpolation between R&P and the output of the transformer
- retrieval with l_2 distance with image embeddings or proprioceptive obs

$$\pi_{\text{REGENT}}^\theta(s_t, r_{t-1}, c_t) = e^{-\lambda d(s_t, s')} \pi_{\text{R\&P}}(s_t, c_t) + (1 - e^{-\lambda d(s_t, s')}) \sigma(\pi_\theta(s_t, r_{t-1}, c_t))$$

$$\sigma(x) = \begin{cases} \text{Softmax}(x), & \text{if action space is discrete} \\ L \times \text{MixedReLU}(x), & \text{if action space is continuous} \end{cases}$$

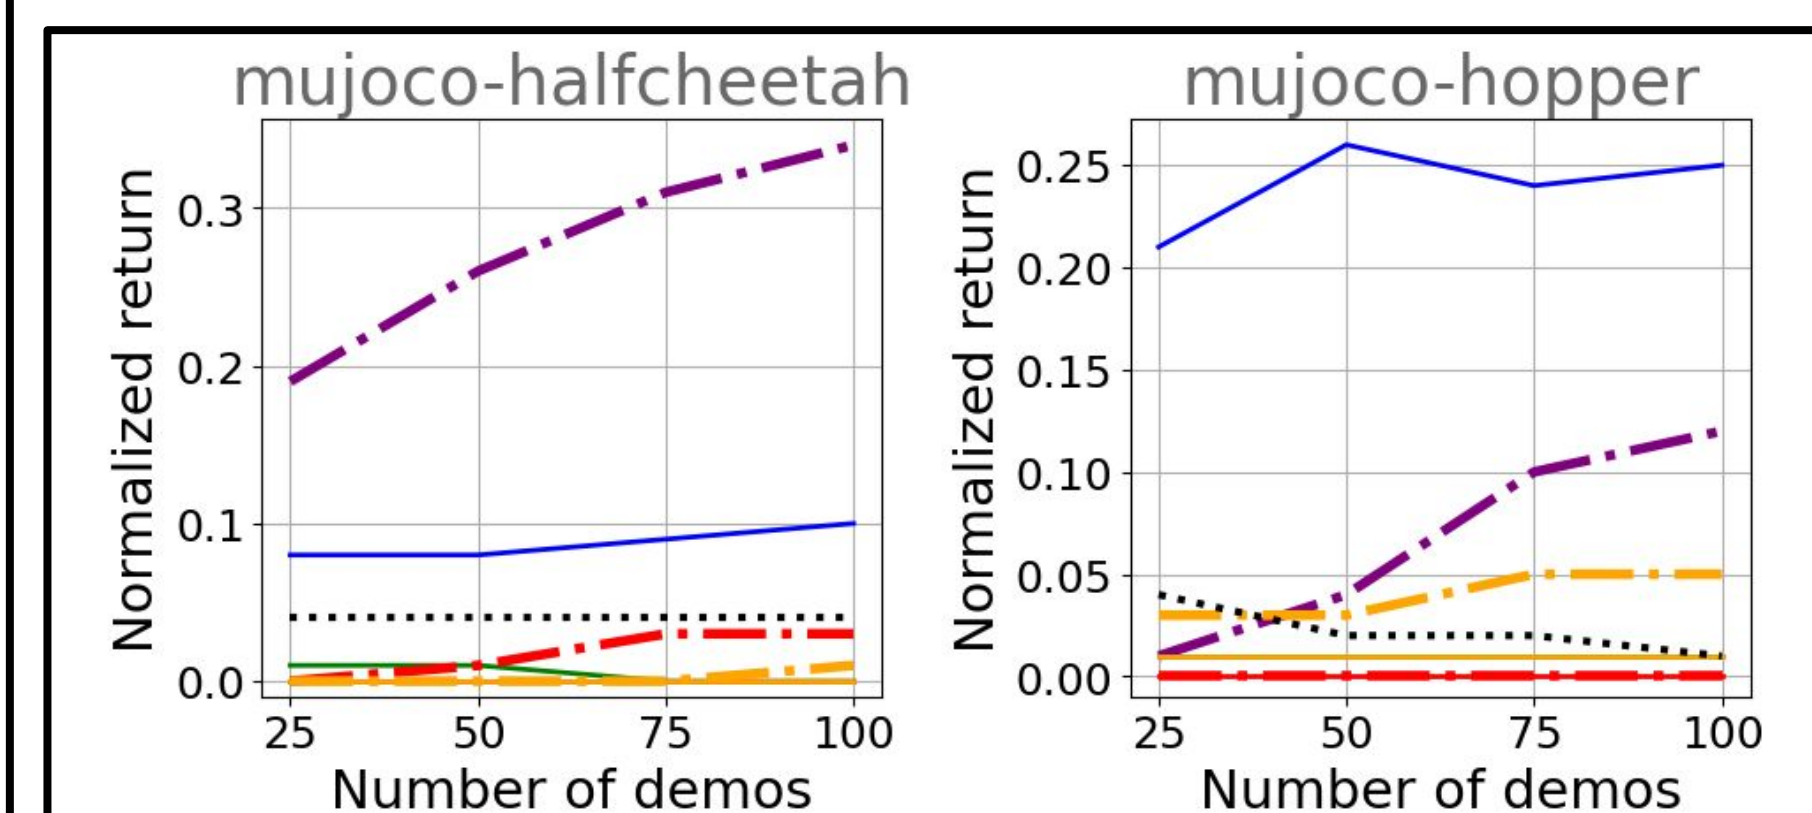
- minimize loss on all actions
- adequate retrieved context, correct context order, expert demos
- 100k transitions in each of the 145 pretrain envs

Experimental Evaluation in Unseen Environments



Limitations/ Future Work

Generalization to new embodiments!



Much more at

bit.ly/regent-research

