


Each following page shows the success rate of the model at different time-steps during an in-context rollout within a novel scene (shown in both third person and birds eye view). Path taken by agent is highlighted in blue.


It can be seen that as a model collects more experience it starts to take the shortest paths to (the same) goals sampled at different points in time.

It can also be seen that when it goes significantly beyond it's training context length (4k), it performance starts to degrade again.

Mug




Success
Num steps=3874




Success
Num steps=4530


obj_names: mug
previous_episode_success: 1.00



Success
Num steps=5066




Success
Num steps=6168





Gelatin Box

Gelatin Box




Failure
Num steps=387




Failure
Num steps=955


Agent fails to find the gelatin box in the scene



Success
Num steps=1332



Success
Num steps=5453




Success
Num steps=7660


Agent finds the gelatin box and navigates to it successfully



Agent not looking
at the object



Failure
Num steps=8760



Failure
Num steps=9327