



The QA-GUI Environment Equivalence: Turning Static Health QA into Scalable GUI Agent RL Data

scalingX

RsiingSI · Independent Researcher
20261017

The bottleneck for GUI agents long ago stopped being the model. It is environment data.

Recent work on environment synthesis is plentiful — AutoWebWorld, CUA-GYM, InfiniteWeb — but all of it does essentially the same thing: build from scratch, or reconstruct from websites. The web domain is crowded, while other modalities and other domains are only just beginning.

We want to ask a different question: **can a QA corpus serve directly as the seed for GUI environments?**

Every QA pair is in fact a flattened GUI environment — the question is the instruction, the context is the interface state, the answer is the reward. Today's QA benchmarks only flatten it into text in, text out. Unfold it, and it

becomes an executable environment that can be operated across devices.

If this road works, the scale of environment data is a different order entirely. LifeAgentBench has 20k+ entries, MedQA 10k+, PubMedQA 270k. Health is only the first validation scenario.

But the core difficulties sit exactly here: which QA are worth unfolding? How are cross-device states generated automatically? How is the verifier synthesized? And most importantly — can the unfolded environment actually train an agent, or is it only a good-looking shell?

The experiments are not done yet. But I think this direction deserves a serious run.

Paper and code will be released when ready. DM if you are interested.