

30 APR 2026

Rapid iteration in agentic loops

Taking output as input

While working with AI-agents to build things I often iterate through versions of the thing I am building. You try, experiment, fail, improve, and again. In the manual days before AI, I checked results by hand. But what if we learn agents to 'see' themselves?

Imagine you add a button somewhere. After writing the code you test drive the new program to see how it looks. Does it align well, does it fit on small screens, does it work? You can easily understand that the more complex software becomes, the more time consuming this process becomes.

Let agents see

To speed things up I gave the AI agent the ability to rapidly see the work it has created, on different screen sizes and to automatically test for alignment, overflow, legibility (a11y) and consistency. I am leveraging headless browsers and local computer vision (opencv) models to do this. I trained the setup by analysing beer bottles, having the agent build a 3D model from a photo. This exercise really improves your understanding of the image
→ data → AI-context → action flow

Faster 'loop'

By automating the checking of iterative versions I reduced 80% of the friction during development. In addition to visual checks I also feed the agent log output and automated "smoke tests". As a fun experiment I once operated the agentic loop while walking (Dutch: 'loop') on a treadmill: the faster the 'loop' the fitter I become, ha!

Conclusion

Agentic workflows likely are at its worse now: things will only improve as computers and models will get faster. But you should not wait for this 'mythical day of frictionlessness': tackle your bottle necks now! Energy not spent on repetitive tasks can be used elsewhere, where humans can make a real difference!