

# Webtalk

## A Natural Language Interface to the Web via Reinforcement Learning

We implemented a novel system for interacting with web pages through natural language. The system uses policy gradient reinforcement learning to find a policy mapping sentences to actions to be performed. We formally evaluated its ability to interact with a simple web form and informally demonstrated its utility on more complicated real-world pages. The system demonstrates the viability of using reinforcement learning to make the web accessible via natural language.



The input to the training algorithm was a website and a corpus of English sentences gathered from Amazon Mechanical Turk paired with interactions (e.g. “click on the continue button”). The algorithm gradually improves a policy, which maps sentences to actions, by gradient descent. This policy trained on a handmade website performed surprisingly well when applied to real-world sites like Hipmunk flight search or the New York Times’ homepage, and we believe that the technique could perform even better with a more representative training corpus and website.

Talking to a webpage is convenient, but would be invaluable for those who are less able to use mice and keyboards. We believe our project has demonstrated the potential of reinforcement learning for this task.