

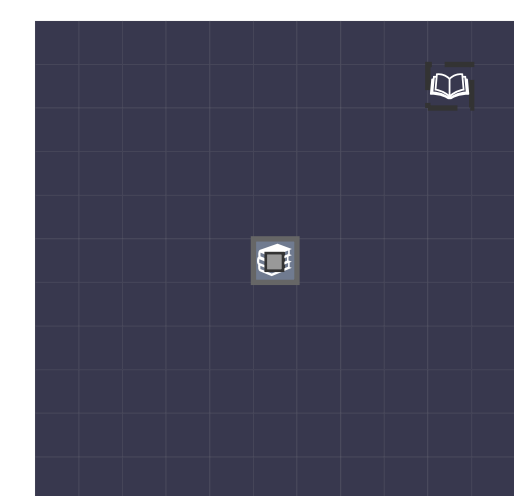
Informed by curiosity studies across psychology, philosophy, and marketing, Ady, Shariff, Günther, and Pilarski (2022) described five properties expected to benefit curious machine learners. Prototyped here are three of the five properties.



The learner in this grid can discover what consistently induces curiosity, analogous to turning on Netflix or visiting a bookstore. With the three properties below implemented *together*, the learner satisfies its curiosity and then learns to return to the 'bookstore' without getting stuck reading the same book.



Visit counts with all three properties intact



Time step 0



Time step 3



Time step 7



Time step 16

Directedness toward inostensible referents


When curiosity is induced, the agent selects a sequence of actions likely to satisfy its curiosity using its knowledge of what it doesn't know.

Cessation when satisfied

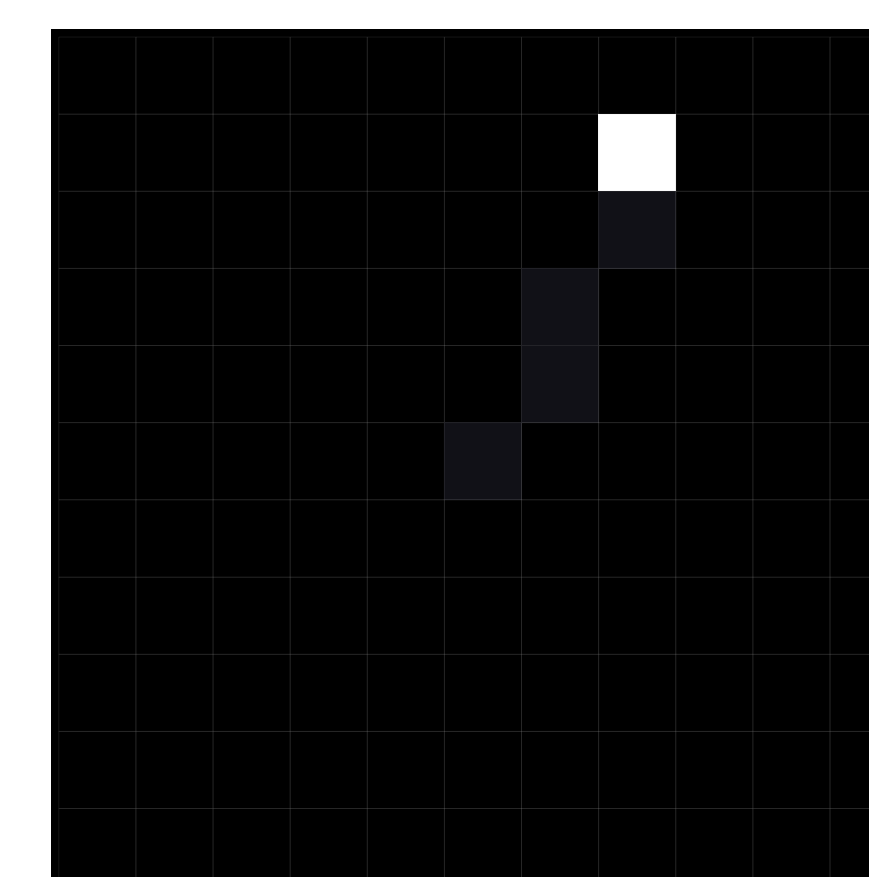
Once a learner has found a situation that satisfies their curiosity, there's no need to observe the same situation again, so curiosity's effect ends.

Voluntary exposure


Learners develop enduring preferences for consistently curiosity-inducing contexts, choosing to partake in activities likely to induce curiosity.



Without directedness towards inostensible referents



Without cessation when satisfied



Without voluntary exposure

Specific curiosity is a holistic pursuit.

Ablation studies suggest:

When any one property is missing, curious behaviour is noticeably impaired.

Nadia Ady [✉ nadia.ady@aalto.fi](mailto:nadia.ady@aalto.fi)



UNIVERSITY OF ALBERTA



HELSINKI INSTITUTE FOR INFORMATION TECHNOLOGY



cat image by
calmreak on freepik