



# CURIOSITY

TIME IS MONEY: NO MORE THUMB TWIDDLING

JORDI BIEGER, ICELANDIC INSTITUTE FOR INTELLIGENT MACHINES

# THUMB TWIDDLING


Q: What does AERA do if after 20 hours of training you give it the opportunity to reflect on what it learned?

A:



# CURIOSITY

The drive to learn and explore the unknown.



# INSUFFICIENT KNOWLEDGE AND RESOURCES

- Any intelligent system must carefully manage the limited amount of knowledge and resources it has.
- These include:
  - Time
  - Input
  - Energy

# SCHMIDHUBER'S FORMAL THEORY OF ARTIFICIAL CURIOSITY

## Intelligent Agent

Predictor /  
Compressor

Reinforcement  
Learner

- Use the predictor's improvements as rewards for the reinforcement learner.
- Seek out good learning situations and utilize "input" scarce resource.
- Explains art, music, humor, science, fun and active learning.
- How to measure learning progress?

# MEASURING PROGRESS


- Oudeyer and Kaplan point out that simply measuring the change of prediction over time is not robust.
- Learning progress should be considered on a per-situation basis.
- Their Intelligent Adaptive Curiosity system learns situation classifiers on the fly and only considers prediction improvements within each situation.

# CURIOSITY IN AERA

- Curiosity improves prediction.
- All AERA models are predictors. This allows for fine-grained progress measurements.
- Better prediction improves AERA overall, and also results in better management of time and energy (computation).
- Curiosity-fueled falsification or vindication of models can cause deletion or compression, which frees up time and energy.

# STEUNEBRINK ET AL.'S WORK-PLAY-DREAM FRAMEWORK

- Exploration-Exploitation Tradeoff
- Work:
  - Exploitation
- Play:
  - Exploration
  - Supervised
- Dream:
  - Batch processing
  - Buffer flushing
  - Conservative exploration
  - Planning





# DISCUSSION

- Why does AERA need this type of curiosity and downtime, while it seems that e.g. NARS doesn't?
- Is there a curiosity mechanism in NARS (implicit or explicit)?
- What about OpenCog? Does it twiddle its thumbs, does it have curiosity, does it require downtime?
- Ideas for the future?