Latent learning progress (LLP) guides hierarchical goal selection in humans

Gaia Molinaro | Anne Collins | Pierre-Yves Oudeyer | Cédric Colas

Humans are **autotelic agents** – but how do we select which goals to pursue?

- New paradigm where **goal selection** is the dependent variable
  - Deterministic feedback
  - Various difficulty levels
  - Hierarchical relationships among goals
- Performance and **learning progress** are important signals for goal selection
- Learning progress ~ derivative of performance

Ten et al., 2021,Colas et al., 2022, Karayanni & Nelken, 2022, Poli et al., 2022, Molinaro & Collins, 2023, Chu et al., 2024

LLP is updated before changes in performance occur

No changes in performance,
but **latent learning progress** occurs

Standard
learning progress
only kicks in here

Goal difficulty and hierarchy impact goal selection

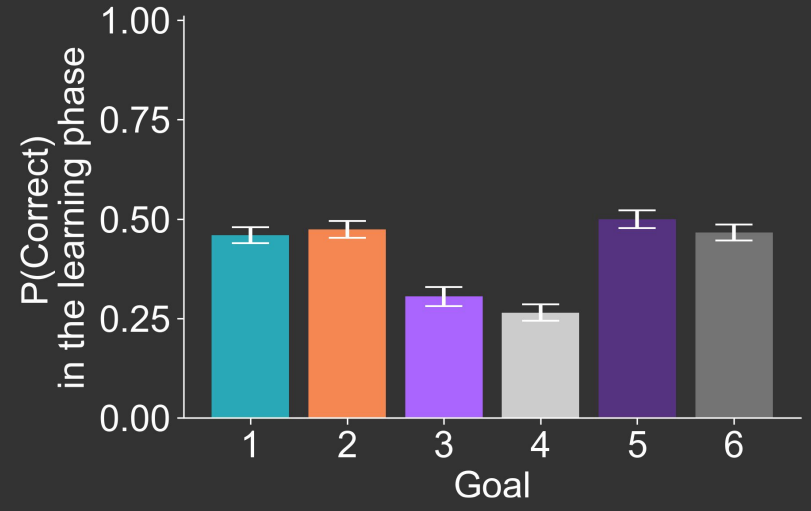

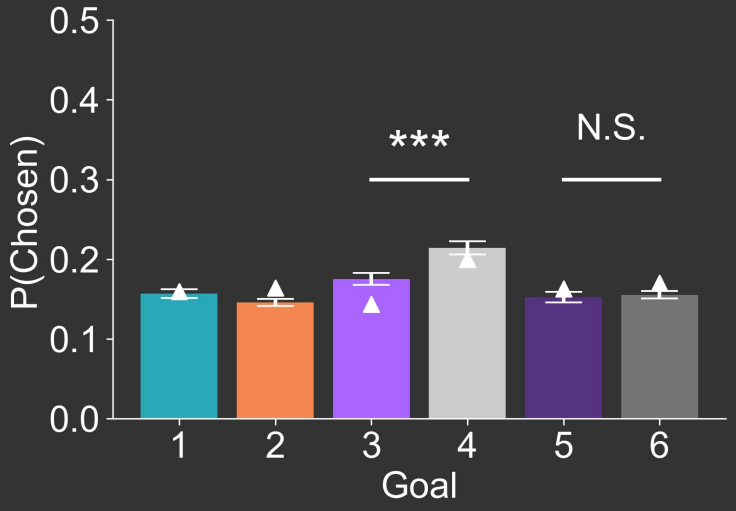

Harder goals G3 and G4
(4 vs 2 ingredients)

Hierarchical structures can be
leveraged in G3 and G5

Latent learning progress guides goal selection

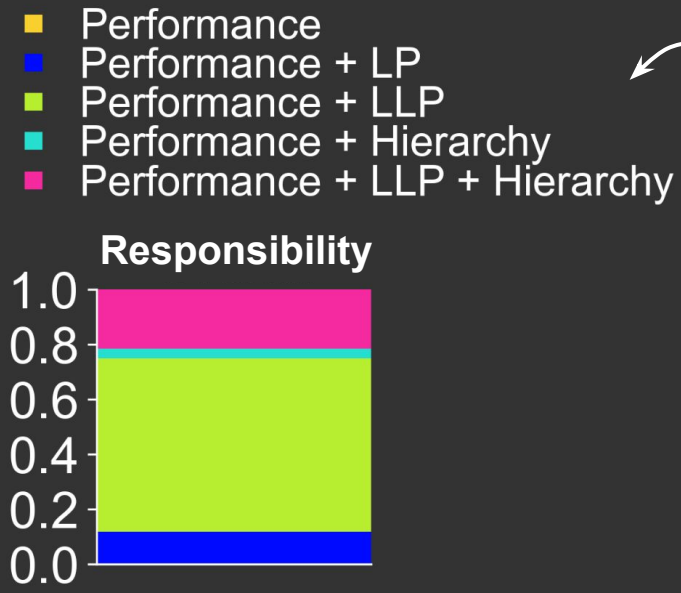

Comparing various signals for goal selection through computational modeling

- At least in certain settings, **LLP** better explains human goal selection than LP
- Hierarchy also likely plays a role
  - Directly on goal selection
  - Indirectly, through learning
- More work is needed to capture the richness in individual strategies
- LLP may be a useful signal for autotelic machines

Piray et al., 2019