On the emergence of collective intelligence

Anil Yaman

Computational Intelligence Group
Vrije Universiteit Amsterdam
https://www.anilyaman.com/
a.yaman@vu.nl

Heterodox Methods for Interpretable and Efficient Artificial Intelligence
13/06/2022

Collective emergent intelligence

The group level (intelligent) behavior that emerges through the interactions of collection of individuals

- Biological neural networks, immune system, animals societies, …

What are the local rules that allow the emergence of intelligent group level behavior?

- Discovering and understanding (interpretability)
- Mimicking the CI for solving computational tasks

How to communicate, specialize, divide the labor and cooperate?

Social learning strategies

How to learn efficiently as a population?

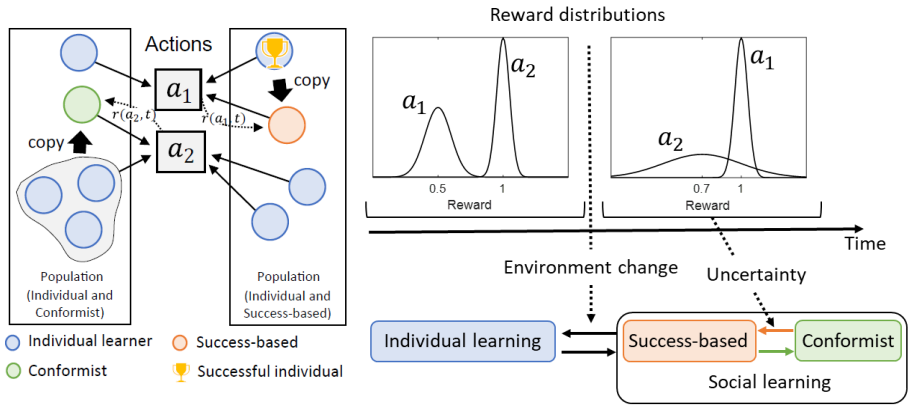

(Yaman A. et al., 2022, PLOS Computational Biology)

Biologically inspired learning

How to learn without a global loss function?

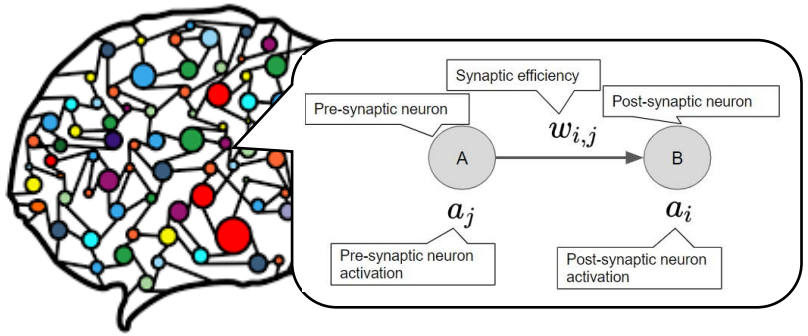

(Yaman A. et al., 2020, Evolutionary Computation)
(Yaman A. et al., 2019, GECCO)
(Yaman A. et al., 2018, GECCO)

Phoenix project (FET-Open H2020)

How to build swarm of sensor networks to explore environments?

(Yaman A. et al., 2017, EvoStar)

Social norms

How can cooperation emerge in self-interested lifetime-learning agents?

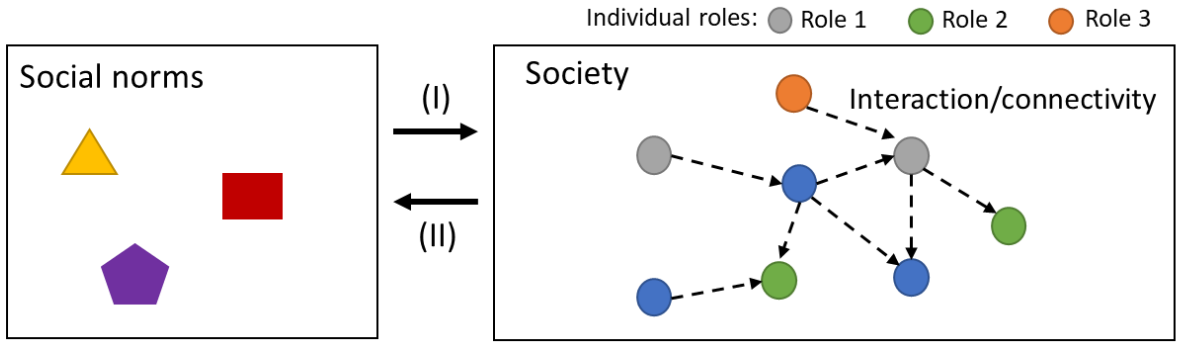

(Yaman A. et al., 2022, in submission)

Distributed evolutionary learning

How to optimize networks using a distributed approach?

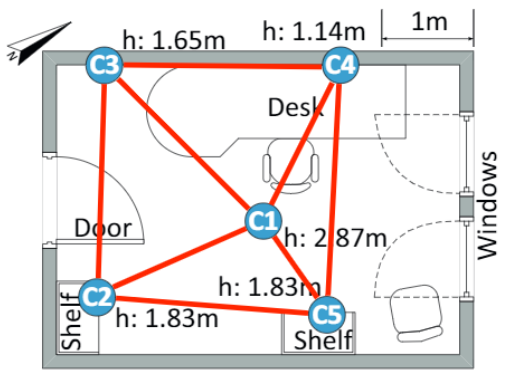

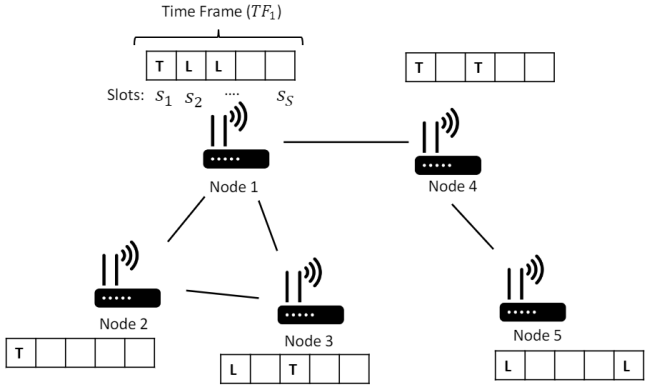

(Yaman A. et al., 2022, Under review)
(Yaman A., Iacca G., 2020, Applied Soft Computing)

How to learn efficiently as a population?

Individual learning: can improve the behavior but involves a learning cost

Social learning (copying the behaviors of others): cannot innovate but provides sample-efficient learning by reducing the learning cost

However, when to copy? (social information can be less accurate)

who to copy from? (identifying individuals with reliable knowledge)

Information is a public good, what about free-riding?

- **Success-based:** copy the most successful individual
- **Conformist:** copy the majority

Yaman A, Bredeche N, Çaylak O, Leibo JZ, Lee SW. Meta-control of social learning strategies. 2022, PLOS Computational Biology

Evolutionary dynamics of social learning

(a) Binary decision-making task

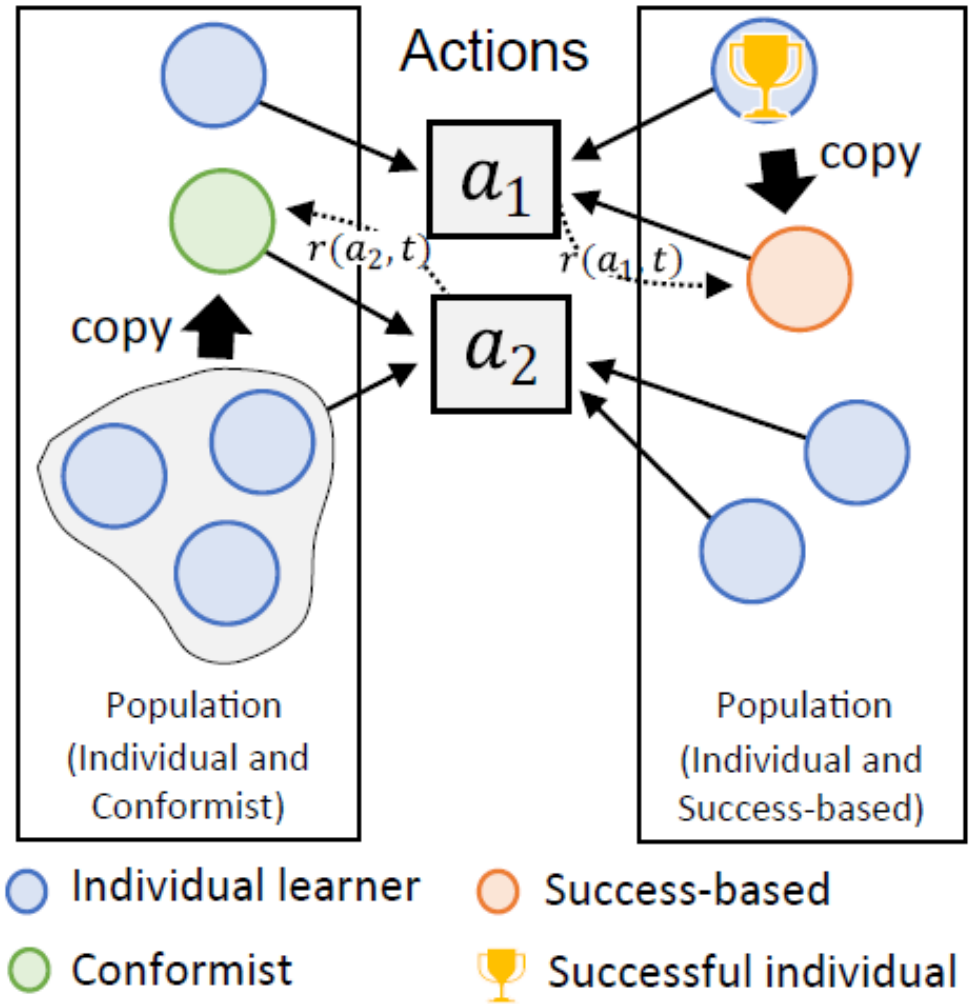

(b) Individual vs. conformist social learning

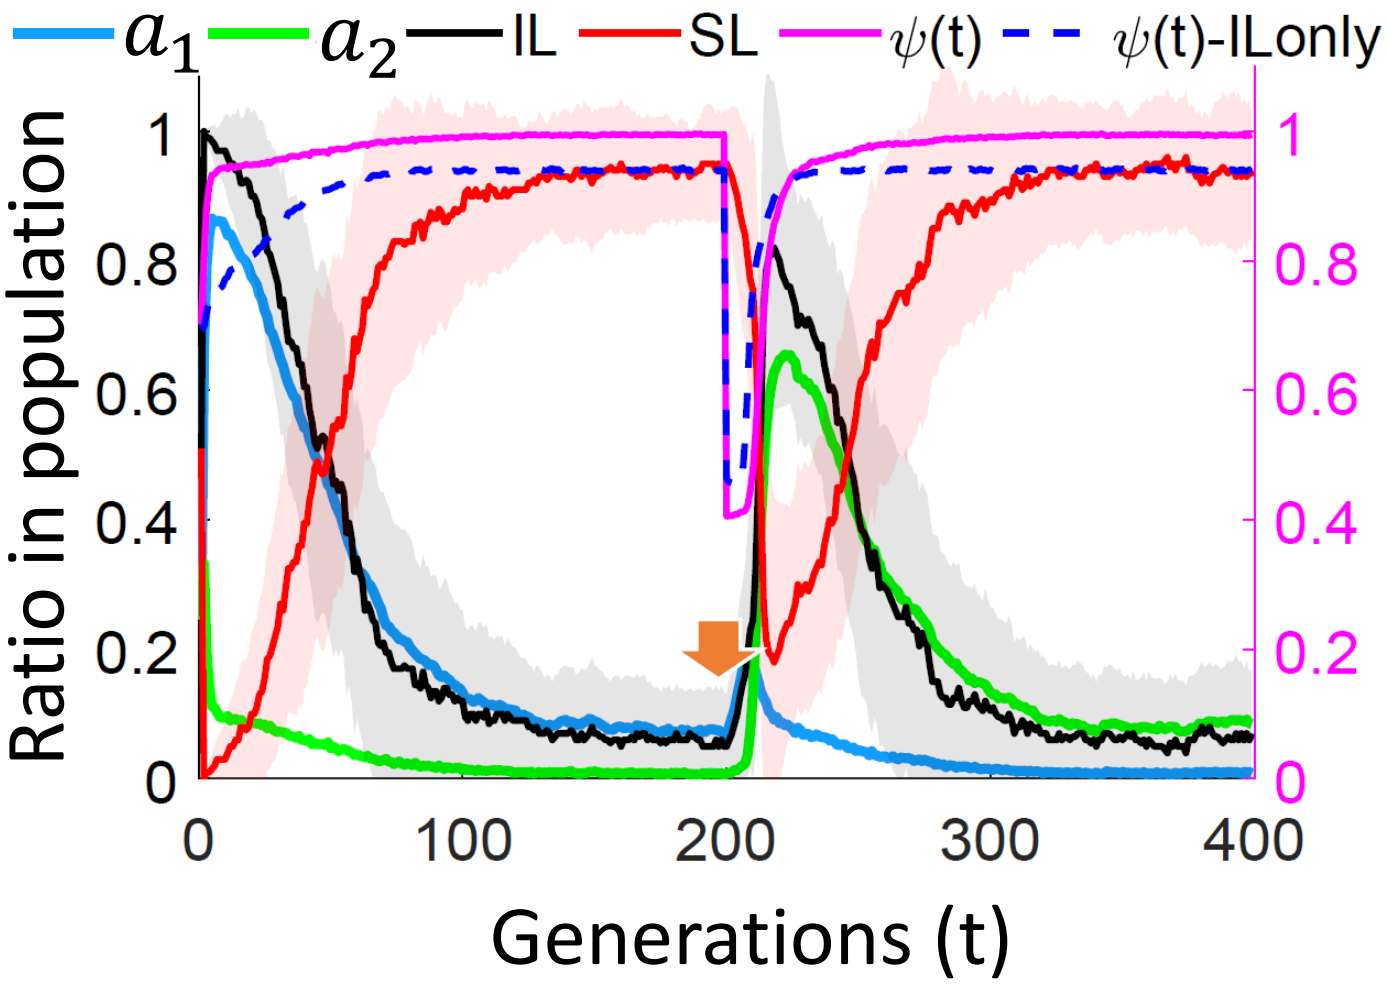

What are the trade-offs of SL strategies?

(a) Binary decision-making task

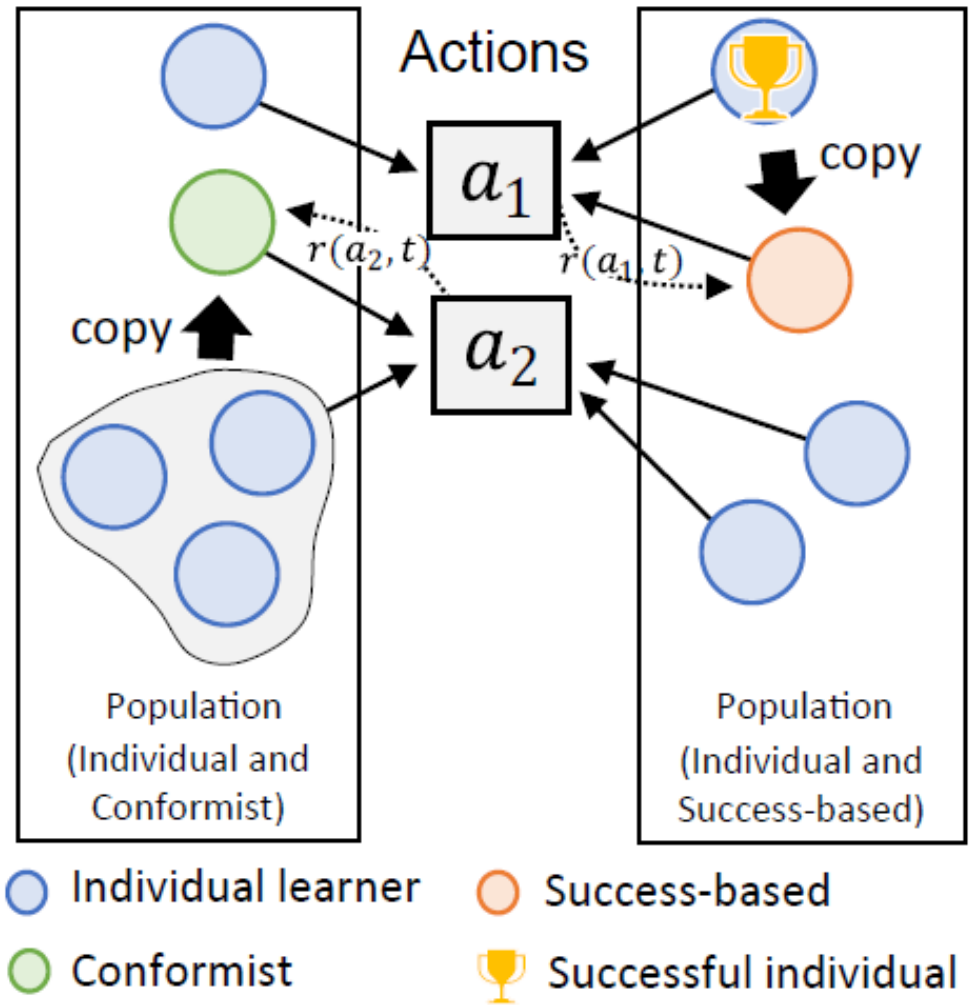

(b) SLSs on environments with low and high uncertainty

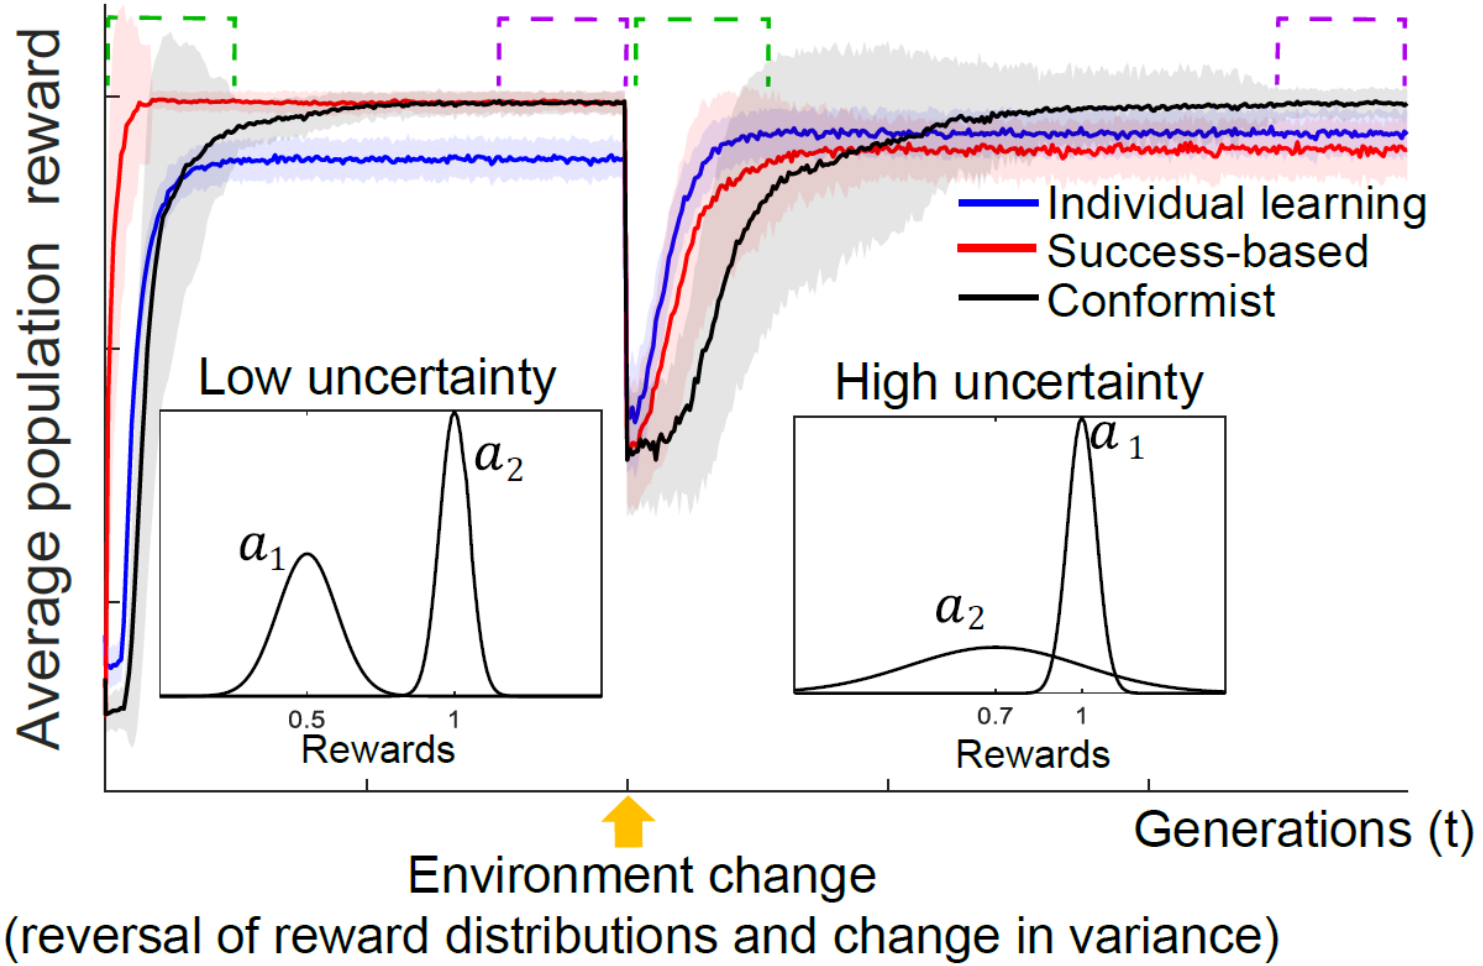

Optimum distribution prediction uncertainty (ODPU)

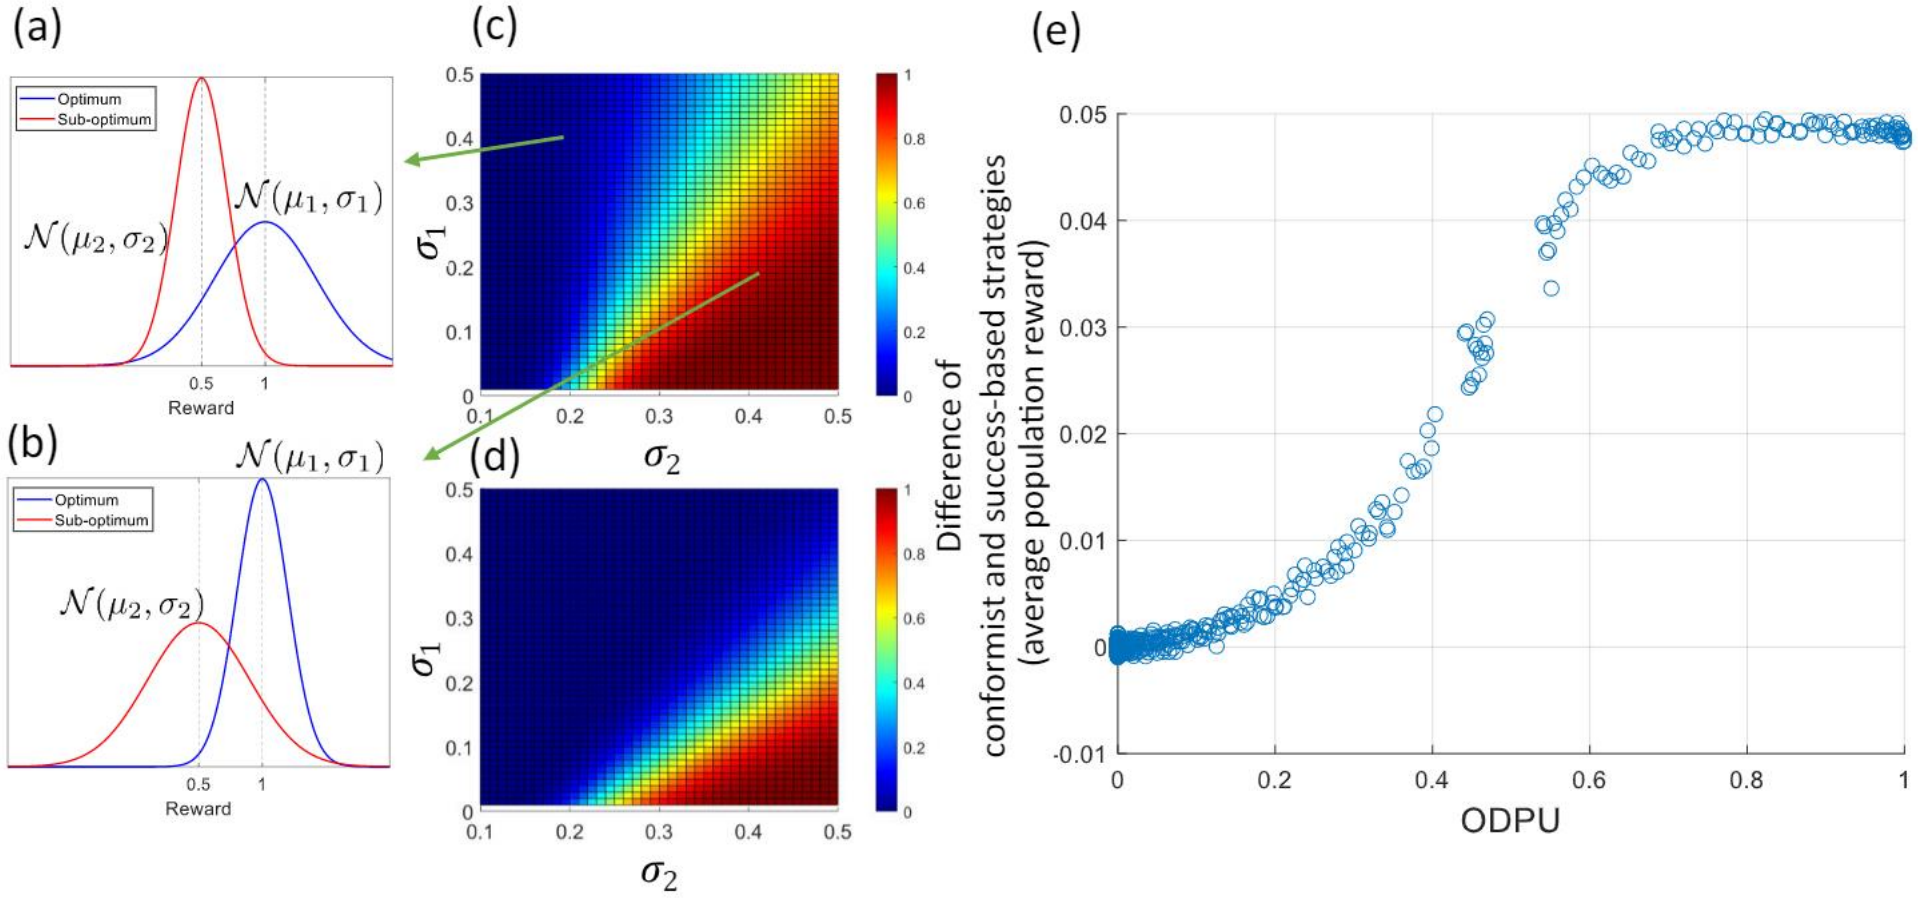

(a) and (b): σ_1= 0.4, σ_2 = 0.2 and σ_1= 0.2, σ_2 = 0.4
(c) and (d) show the ODPU, formalized as the probability of sampling the highest reward value from the sub-optimum distribution
In (c) the ratios of optimum and sub-optimum: 0.05 and 0.95
In (d) the ratios of optimum and sub-optimum: are 0.5 and 0.5
(e) the relation between the ODPU and the difference in average rewards

Meta-control hypothesis

- Different learning strategies have different sensitivities to uncertainty (e.g. Pavlovian, model-based and model-free [1])
- Meta-social learning provides effective and sample-efficient learning in social context
- Can have a high impact in real-world applications of multi-agent systems
- Can help form hypothesis for arbitration mechanisms in social learning context

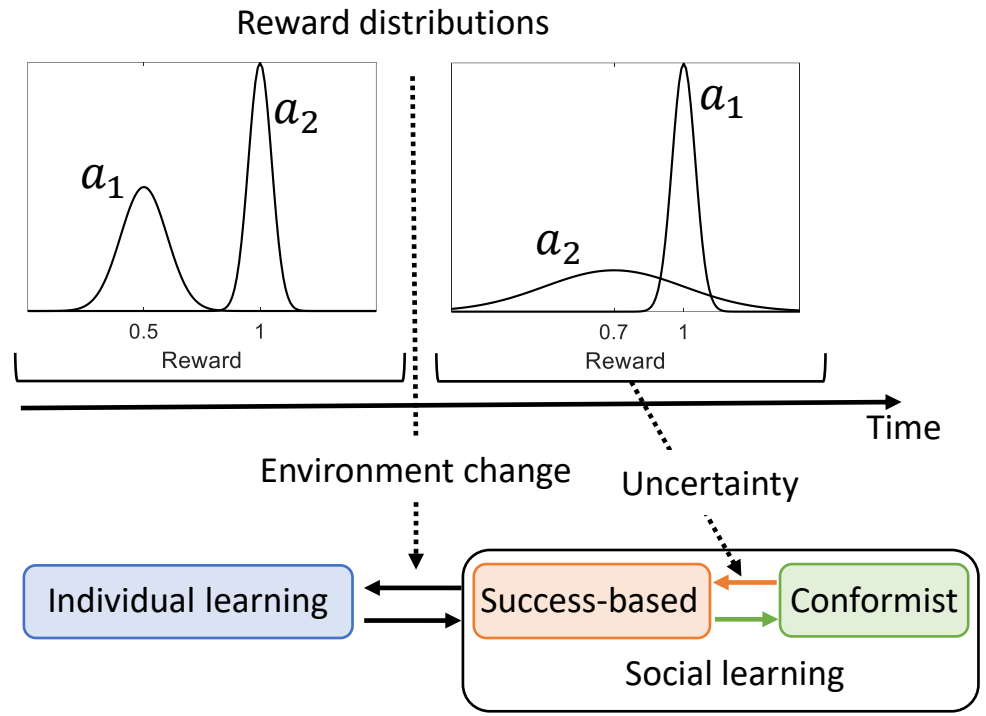

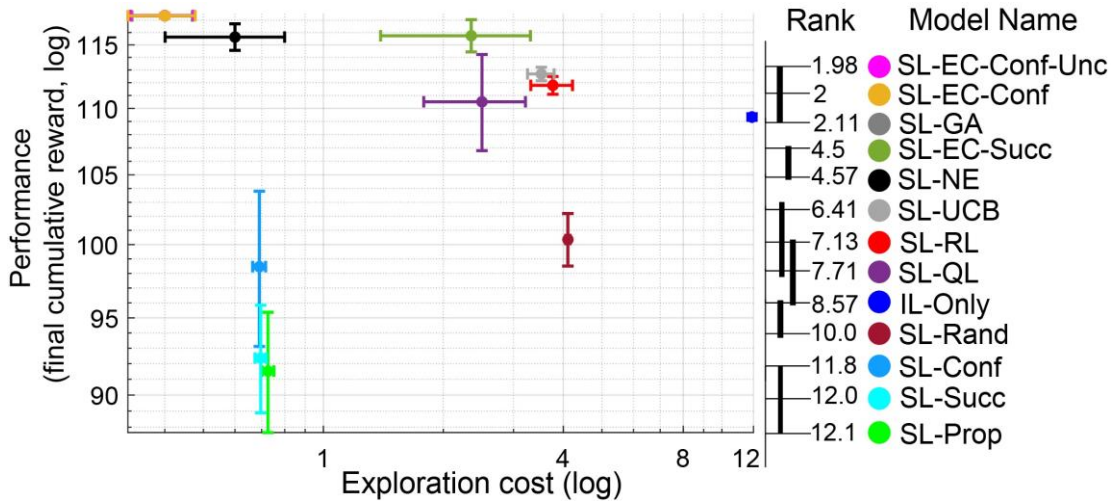

[1] Sang Wan Lee, Shinsuke Shimojo, and John P O'Doherty. Neural computations underlying arbitration between model-based and model-free learning. Neuron, 81(3):687–699, 2014

Learning division of labor in self-interested lifetime learning agents

(a) Spatial role distribution on grid environment

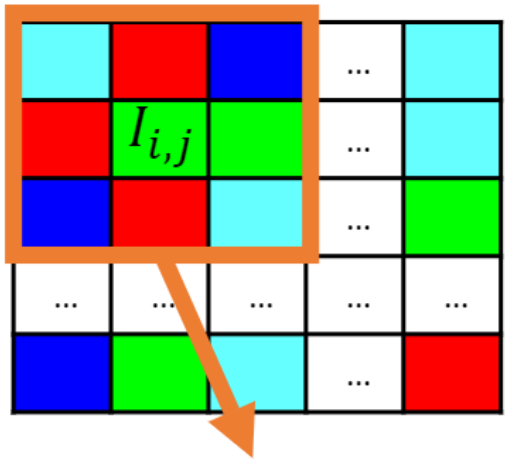

Neighborhood (N) of focal agent: $I_{i,j}$

(b) Settlement maintenance

Roles (ρ)	$r_t^{(i,j)}$
Cleaner (C)	0
Forager (F)	3
Hunter (H)	6, if 5 or more hunters, 0, otherwise.
Soldier (S)	0

Role functions

- C: mitigates waste accumulation
- S: protects against adversarial attacks

(c) Common pasture

Roles (ρ)	$r_t^{(i,j)}$
Worker (W)	0
Herder (C - considerate)	5
Herder (G - greedy)	10

Role functions

- W: increases speed of resource recovery

Learning division of labor in self-interested lifetime learning agents

Roles (ρ)		$r_t^{(i,j)}$
	Cleaner (C)	0
	Forager (F)	3
	Hunter (H)	6, if 5 or more hunters, 0, otherwise.
	Soldier (S)	0

Role functions

- C: mitigates waste accumulation
- S: protects against adversarial attacks

without social sanctions (selfish)

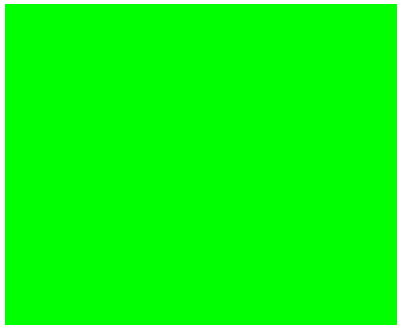

with social sanctions

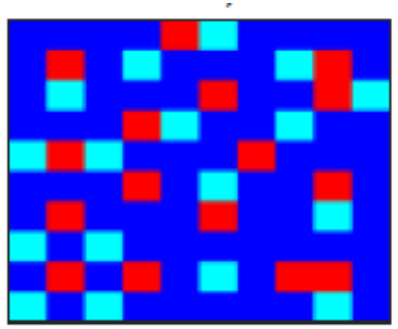

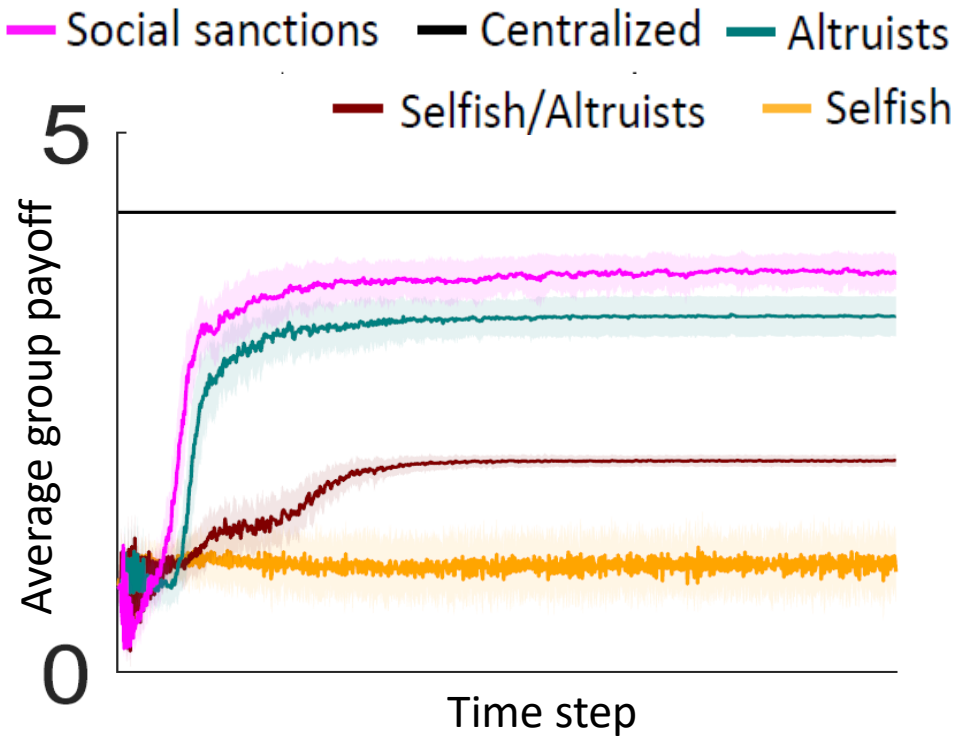

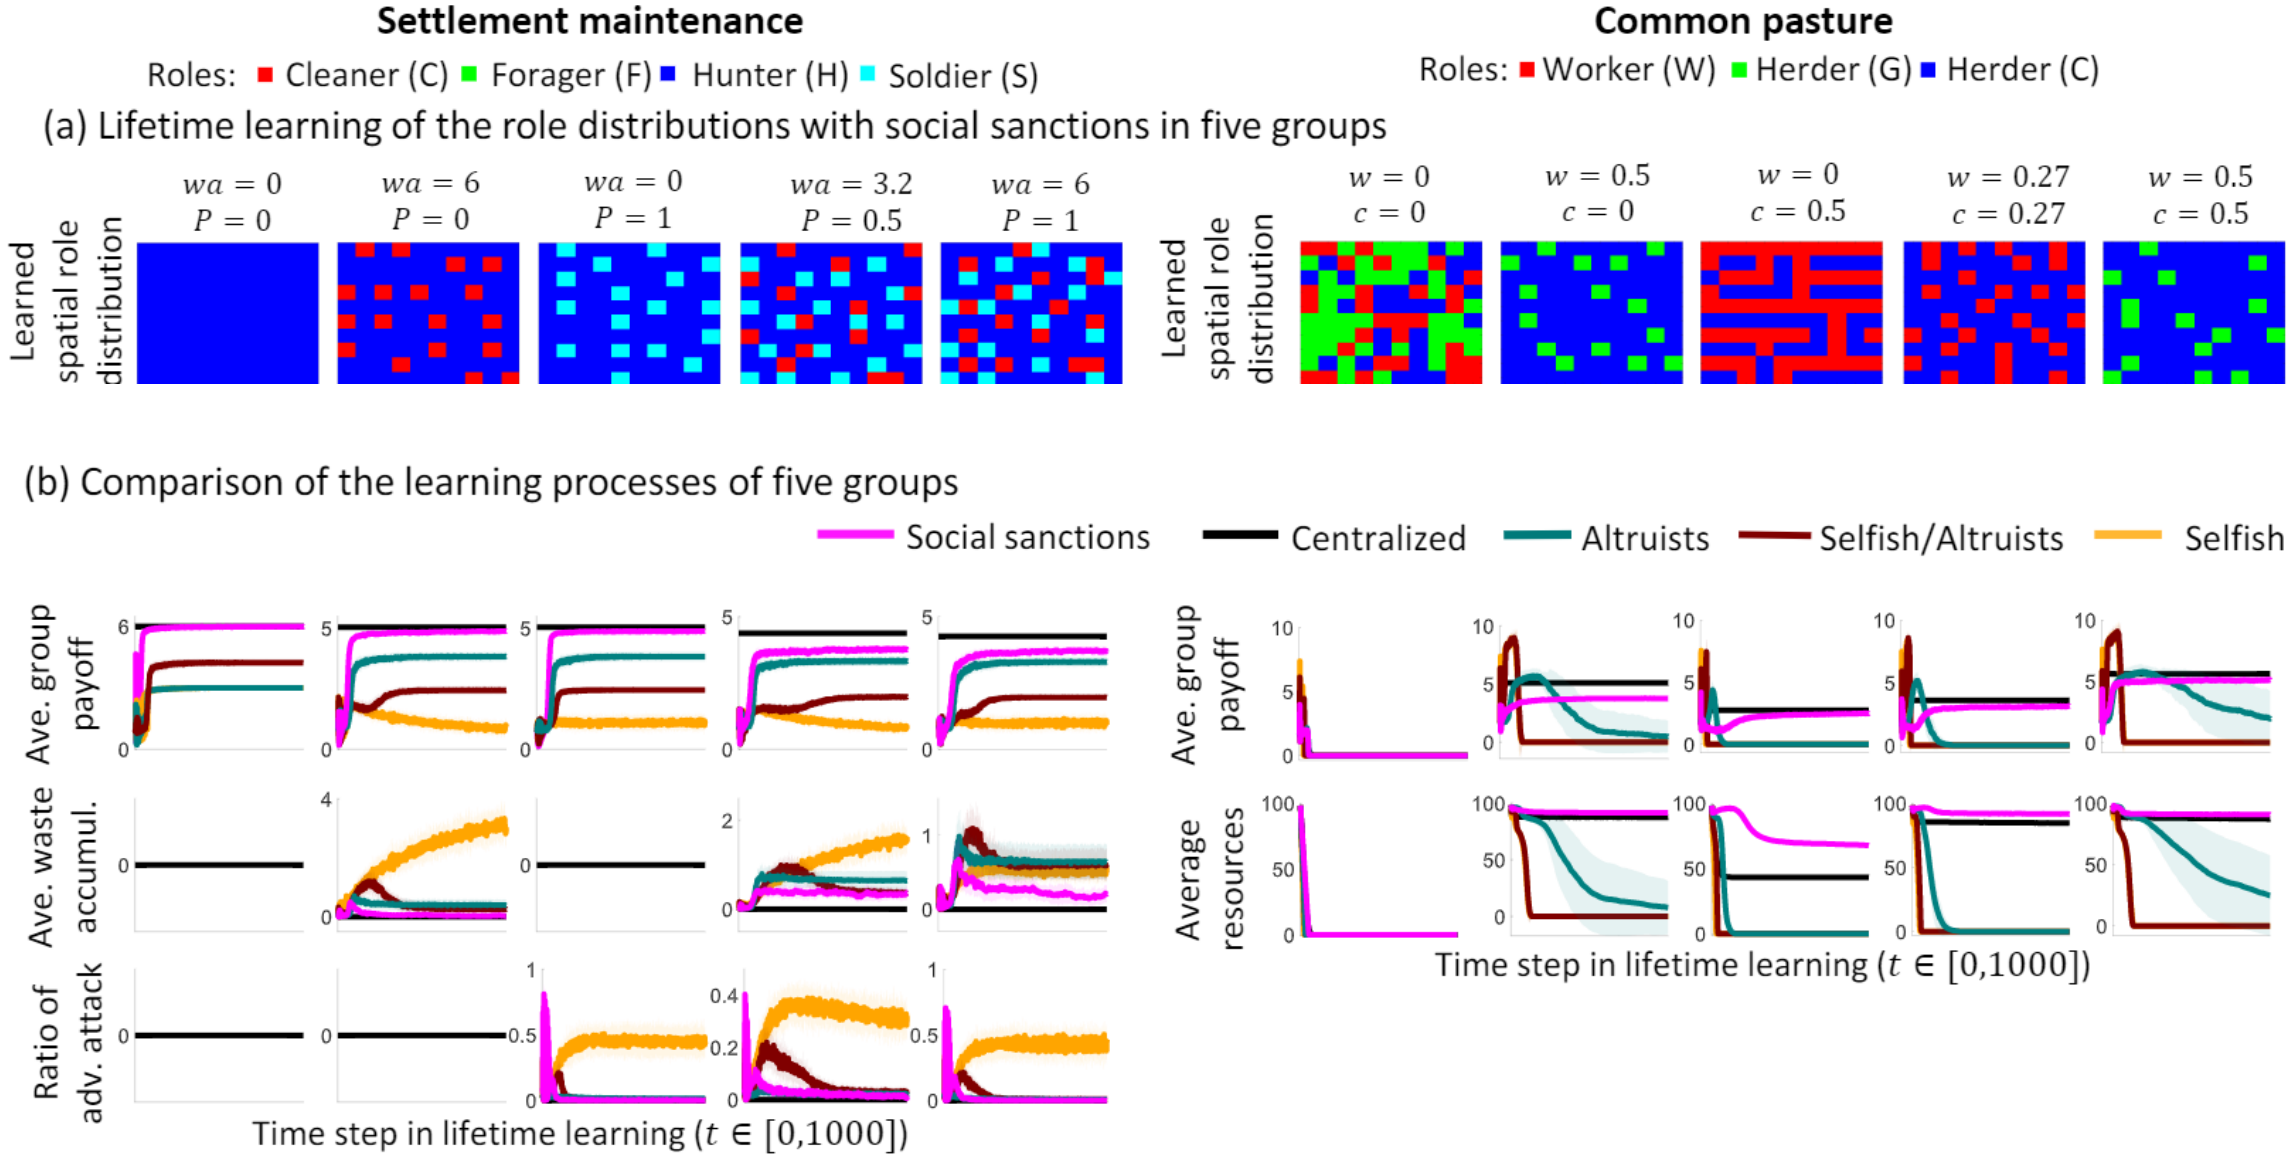
Settlement maintenance
Roles: Cleaner (C) Forager (F) Hunter (H) Soldier (S)
Common pasture
Roles: Worker (W) Herder (G) Herder (C)
(a) Lifetime learning of the role distributions with social sanctions in five groups
Learned spatial role distribution
$wa = 0$ $P = 0$
$wa = 6$ $P = 0$
$wa = 0$ $P = 1$
$wa = 3.2$ $P = 0.5$
$wa = 6$ $P = 1$
Learned spatial role distribution
$w = 0$ $c = 0$
$w = 0.5$ $c = 0$
$w = 0$ $c = 0.5$
$w = 0.27$ $c = 0.27$
$w = 0.5$ $c = 0.5$
(b) Comparison of the learning processes of five groups
Social sanctions
Centralized
Altruists
Selfish/Altruists
Selfish
Ave. group payoff
Ave. waste accumul.
Ratio of adv. attack
Time step in lifetime learning ($t \in [0,1000]$)
Ave. group payoff
Average resources
Time step in lifetime learning ($t \in [0,1000]$)

Future directions

- Social learning, cultural evolution and social norms
- Emergence of language/communication
- Human-AI interaction/cooperation

Thank you!