Title: Autonomous learning agents and what they can teach us about quantum mechanics, biology and more

Date: Sep 19, 2017 03:30 PM

URL: http://pirsa.org/17090018

Abstract: <p>What can machine learning teach us about quantum mechanics? I will begin with a brief overview of attempts to bring together the two fields, and the insights this may yield. I will then focus on one particular framework, Projective Simulation, which describes physics-based agents that are capable of learning by themselves. These agents can serve as toy models for studying a wide variety of phenomena, as I will illustrate with examples from quantum experiments and biology.</p>

Outline

Quantum machine learning

- Quantum-enhanced machine learning
- Machine learning applied to quantum systems

Autonomous learning agents

- General features
- The projective simulation framework
- Applications:
  - quantum experiments
  - a toy model of agency
  - interaction and collective behaviour
  - learning structured models

Quantum Machine Learning

What is artificial intelligence?

- knowledge representation
- automated reasoning
- computer vision, perception
- robotics
- natural language processing
- machine learning

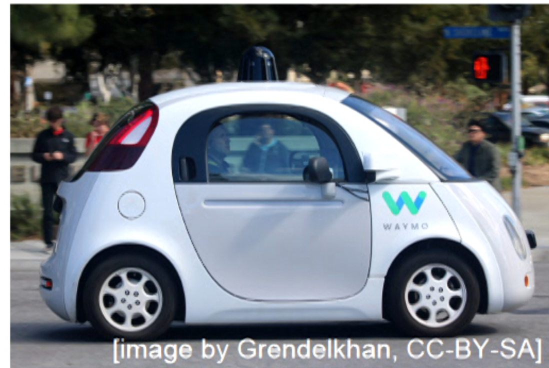

[image by Grendelkhan, CC-BY-SA]

[image by Atomic Taco, CC-BY-SA]

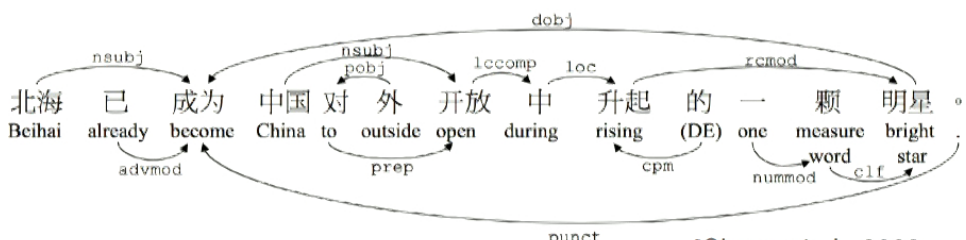

[Chang et al, 2009. Discriminative Reordering with Chinese Grammatical Relations Features]

Quantum-enhanced machine learning

Annealing and sampling

- computational advantage in optimization
- content-addressable memory

Algorithms for universal quantum computers

- amplitude-amplification [Aïmeur, Brassard, Gambs, 2013]
- amplitude-encoding: quantum linear algebra
  - quantum linear system solver [Harrow, Hassidim, Lloyd, 2009]
  - density matrix exponentiation [Lloyd, Mohseni, Rebentrost, 2014]

Reviews: Biamonte et al, Nature 549, 195 (2017);
Dunjko and Briegel, arXiv:1709.02779;
Schuld, Sinayskiy, Petruccione, Contemp Phys 56, 172 (2015)

Machine learning applied to quantum systems

- estimation and metrology: optimization
- condensed matter and many-body physics
  - detect and learn phases
  - learn efficient representations of states
- control and design

Reviews: Biamonte et al, Nature 549, 195 (2017);
Dunjko and Briegel, arXiv:1709.02779;
Schuld, Sinayskiy, Petruccione, Contemp Phys 56, 172 (2015)

Autonomous Learning Agents

Autonomous learning agents as toy models

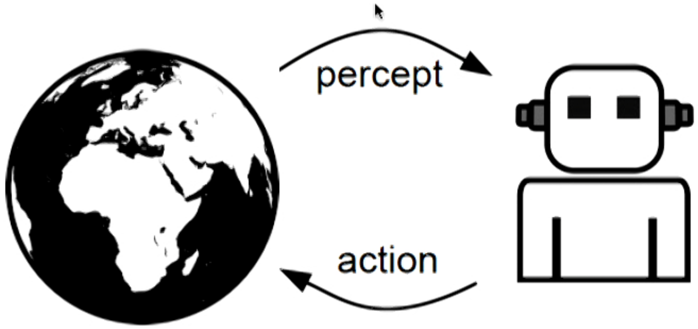

Autonomous learning agents as toy models

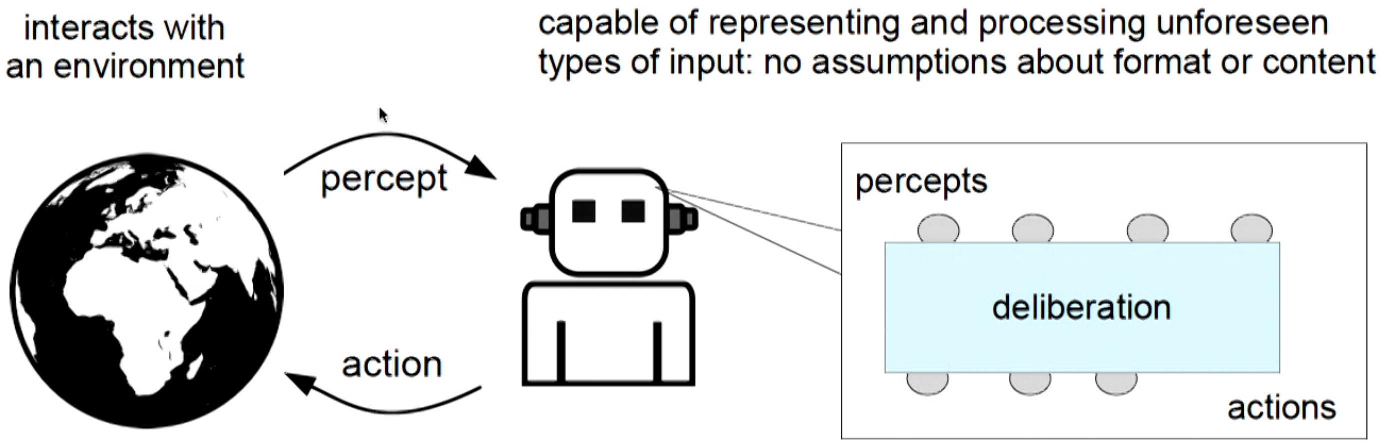

Autonomous learning agents as toy models

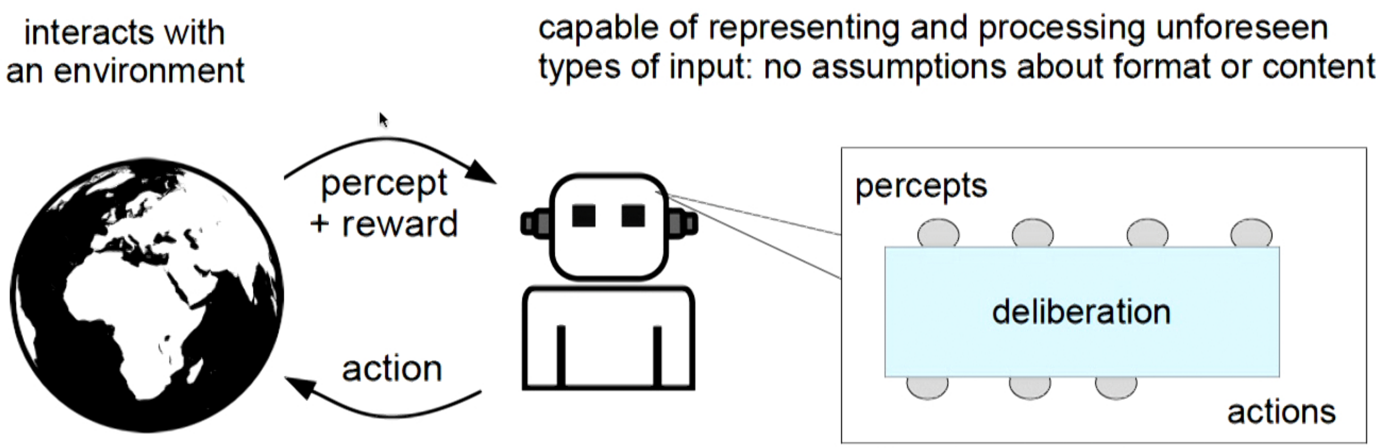

Agent is capable of making decisions that are not pre-programmed:

- non-trivial internal structure and dynamics
- able to modify deliberation structure, based on feedback from the environment

Autonomous learning agents as toy models

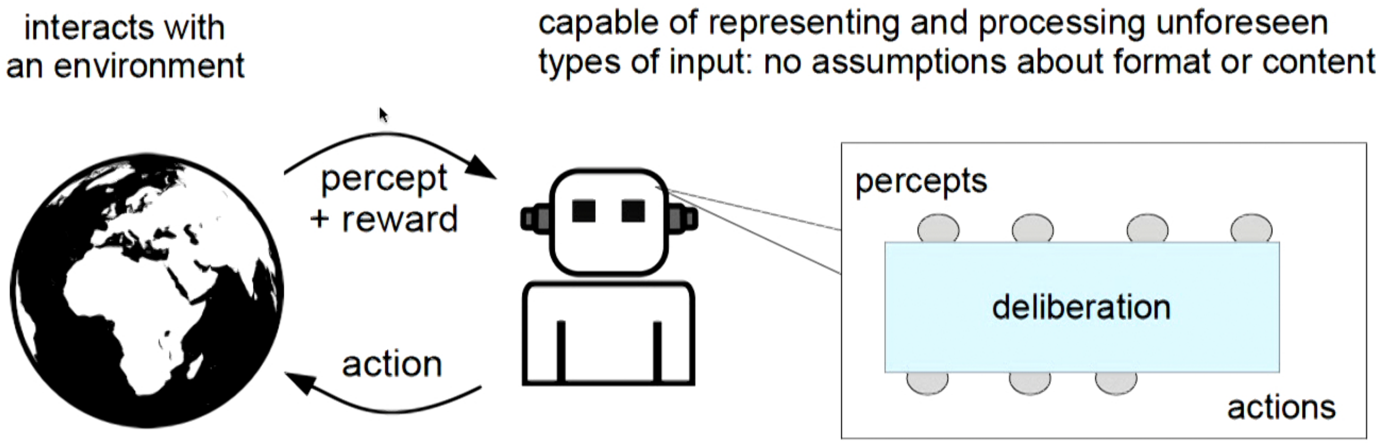

Agent is capable of making decisions that are not pre-programmed:
- non-trivial internal structure and dynamics
- able to modify deliberation structure, based on feedback from the environment

desired features of deliberation process:
- toy model: internal mechanisms simple enough that we can understand how decisions are reached and why the agent behaves as it does: few components, simple rules
- capable of structuring acquired knowledge

Projective simulation: a framework for autonomous learning agents

episodic and compositional memory:
clips of previous experiences

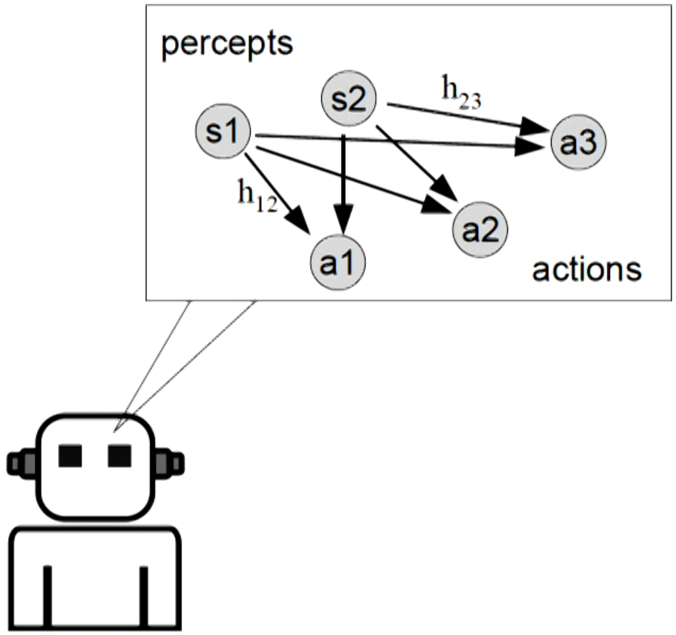

[Briegel and de las Cuevas, Sci Rep 2, 400 (2012)]

Projective simulation: a framework for autonomous learning agents

episodic and compositional memory:
clips of previous experiences

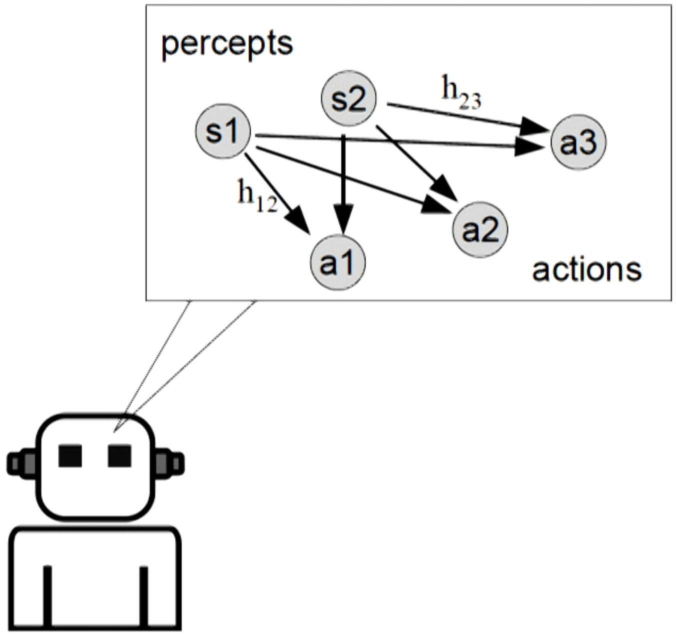

h-matrix:
record of percept-action connections

$$h = \begin{pmatrix} 1.0 & 2.5 & 1.0 \\ 1.2 & 1.4 & 1.0 \end{pmatrix} \begin{array}{l} \leftarrow \text{percept 1} \\ \leftarrow \text{percept 2} \end{array}$$

[Briegel and de las Cuevas, Sci Rep 2, 400 (2012)]

Projective simulation: a framework for autonomous learning agents

episodic and compositional memory:
clips of previous experiences

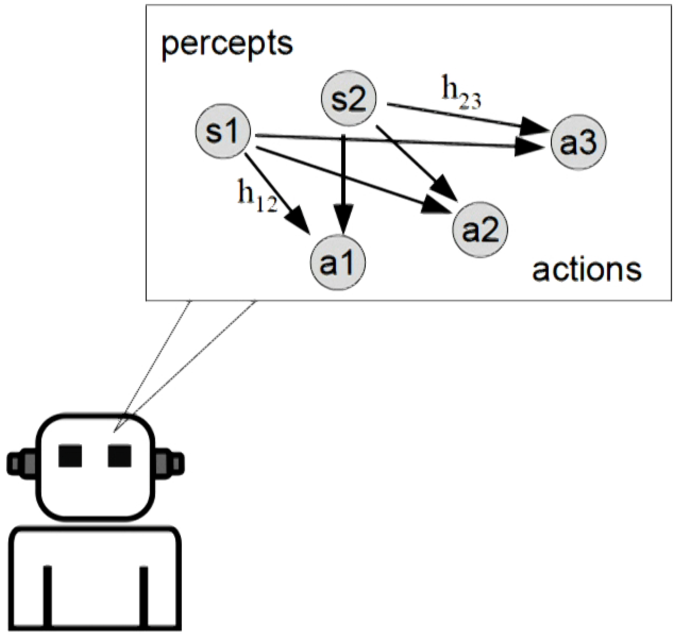

h-matrix:
record of percept-action connections

$$h = \begin{array}{cl} \begin{pmatrix} 1.0 & 2.5 & 1.0 \\ 1.2 & 1.4 & 1.0 \end{pmatrix} & \begin{array}{l} \leftarrow \text{percept 1} \\ \leftarrow \text{percept 2} \end{array} \end{array}$$

Deliberation:
given a percept, choose an action

$$Prob(a_i | s_1) = f(1.0 \quad 2.5 \quad 1.0)$$

[Briegel and de las Cuevas, Sci Rep 2, 400 (2012)]

Projective simulation: a framework for autonomous learning agents

episodic and compositional memory:
clips of previous experiences

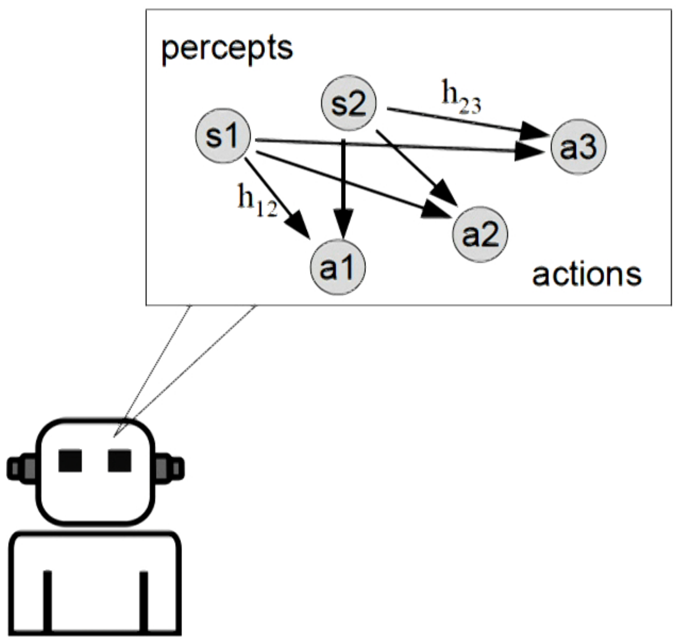

h-matrix:
record of percept-action connections

$$h = \begin{pmatrix} 1.0 & 2.5 & 1.0 \\ 1.2 & 1.4 & 1.0 \end{pmatrix} \begin{array}{l} \leftarrow \text{percept 1} \\ \leftarrow \text{percept 2} \end{array}$$

Deliberation:
given a percept, choose an action

$$Prob(a_i|s_1) = f(1.0 \quad 2.5 \quad 1.0)$$

Reinforcement:
given percept, action and reward R,

$$h^{t+1}(a_i, s_j) = h^t(a_i, s_j) + R_{ij}^t$$

[Briegel and de las Cuevas, Sci Rep 2, 400 (2012)]

Applications of PS agents to reinforcement-learning tasks

The mountain car problem

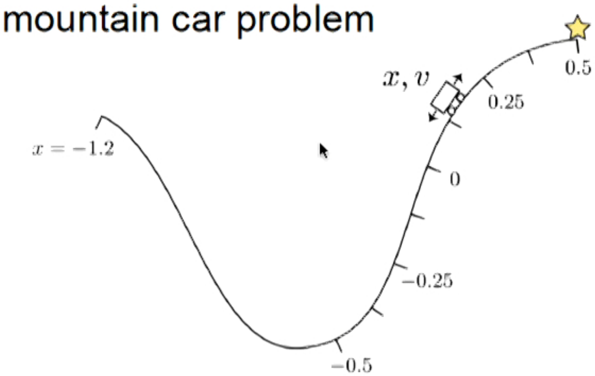

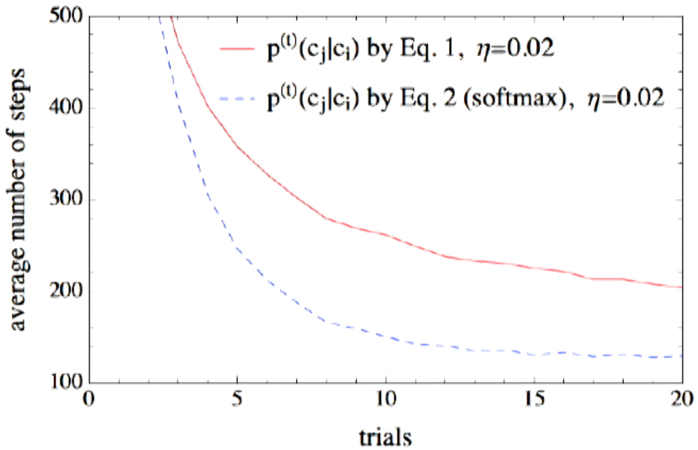

[Melnikov et al, arXiv:1405.5459]

Applications of PS agents to reinforcement-learning tasks

The mountain car problem

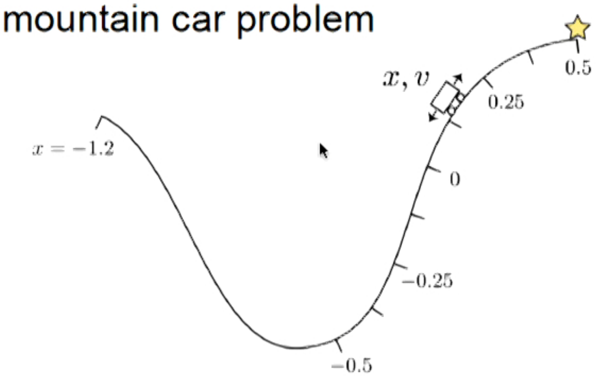

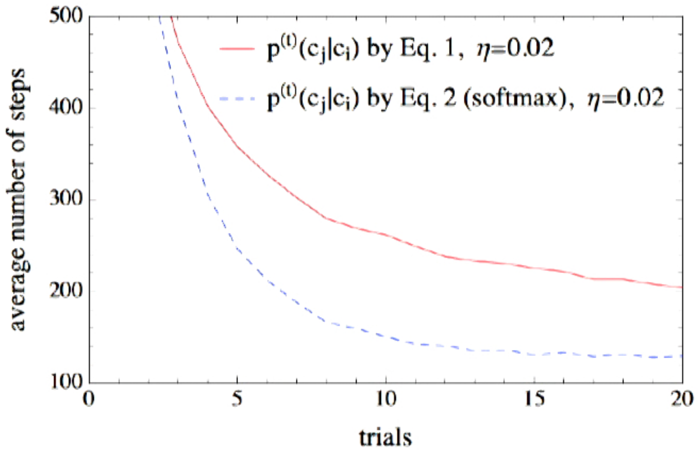

[Melnikov et al, arXiv:1405.5459]

Autonomous learning in robotics

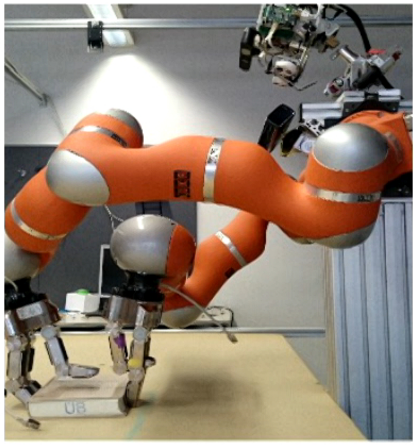

[Hangl et al, IROS 2016, arXiv:1603.00794]

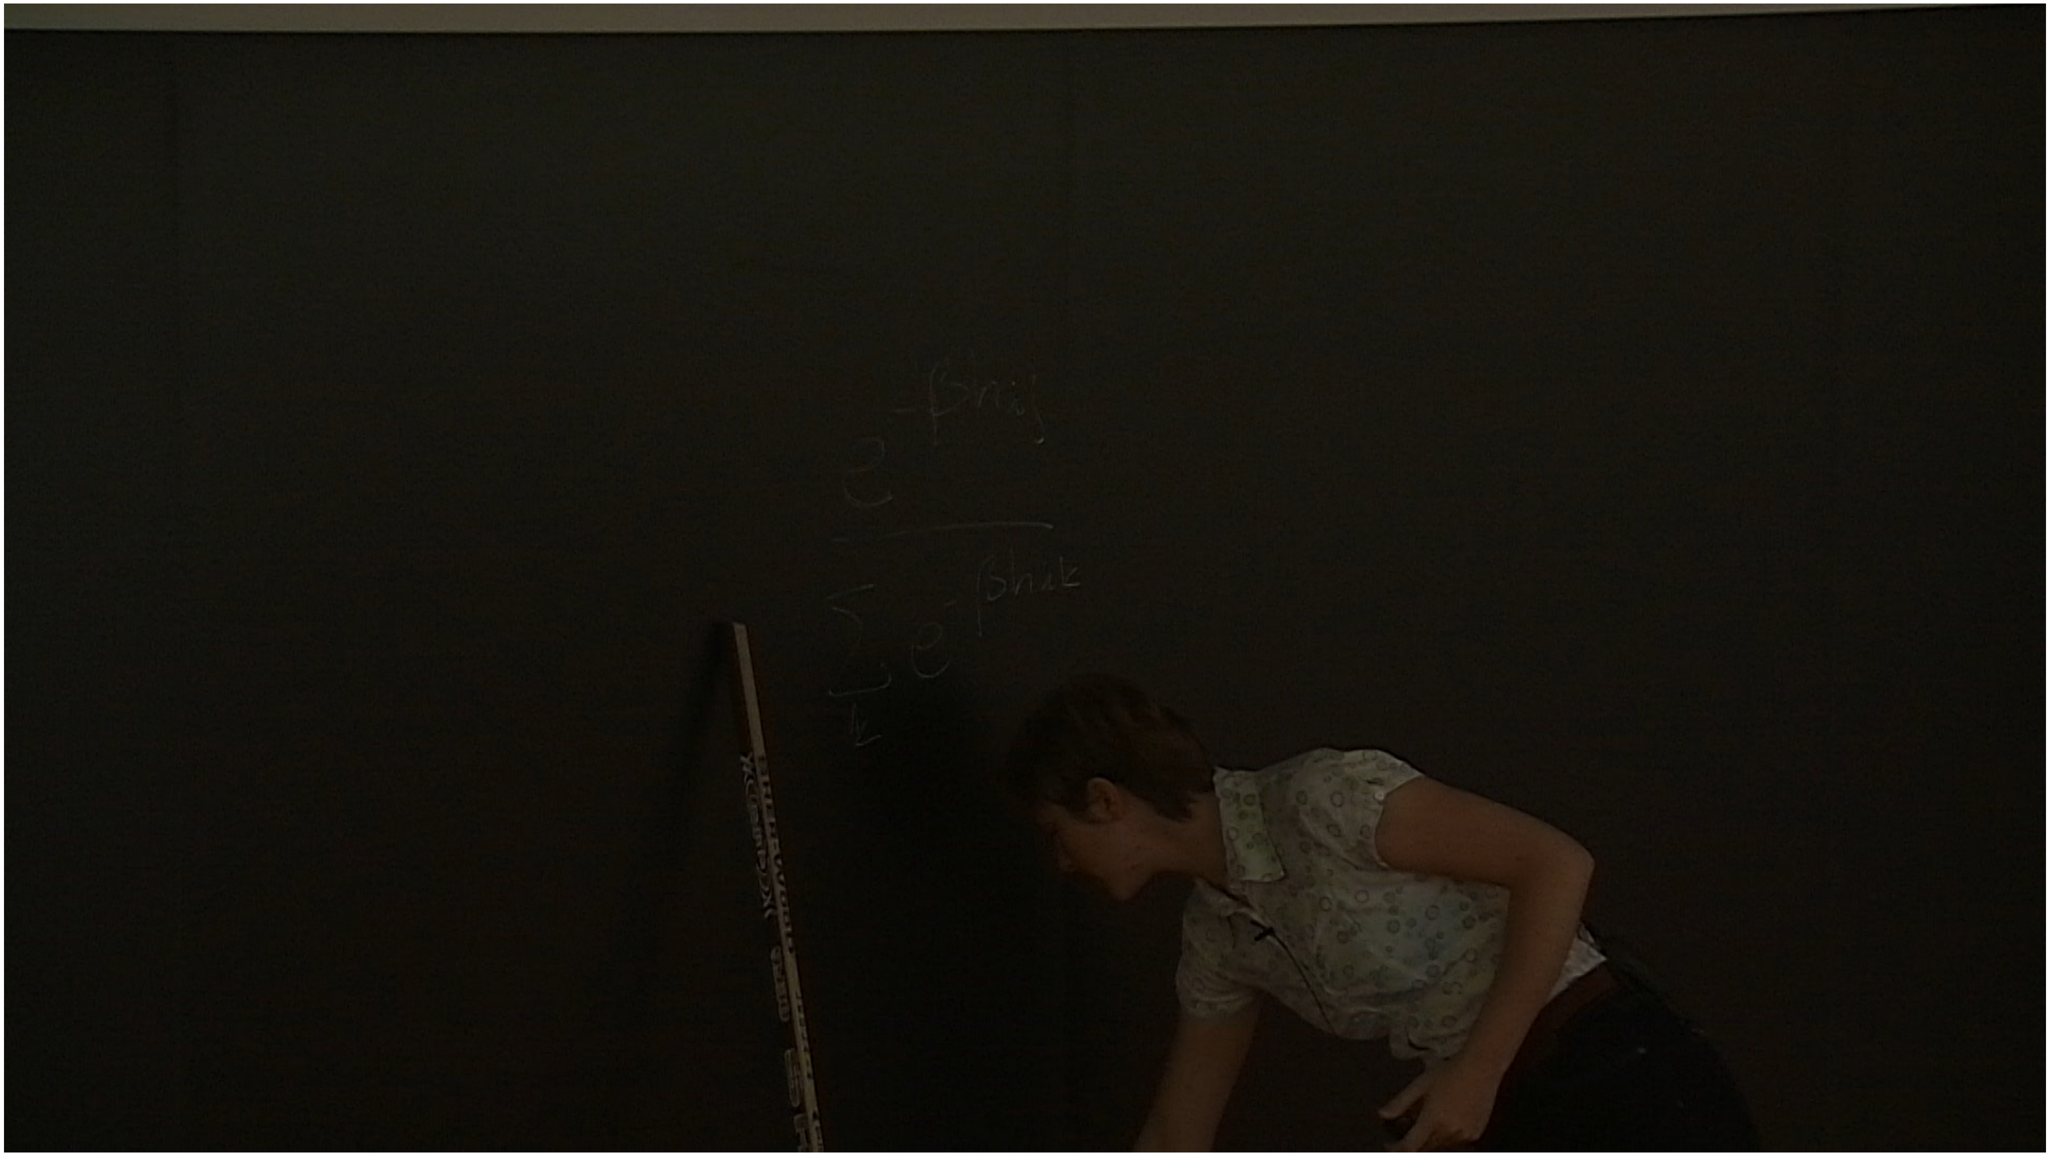
$\frac{e^{-\beta h_{ij}}}{\sum_k e^{-\beta h_{ik}}}$

PS agents for quantum experiments

Measuring a spin in a stray magnetic field

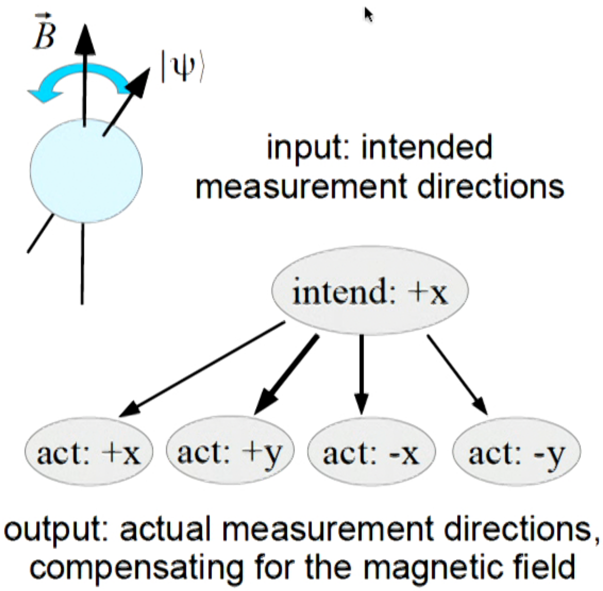

[Tiersch et al, Sci Rep 5, 12874 (2015)]

PS agents for quantum experiments

Measuring a spin in a stray magnetic field

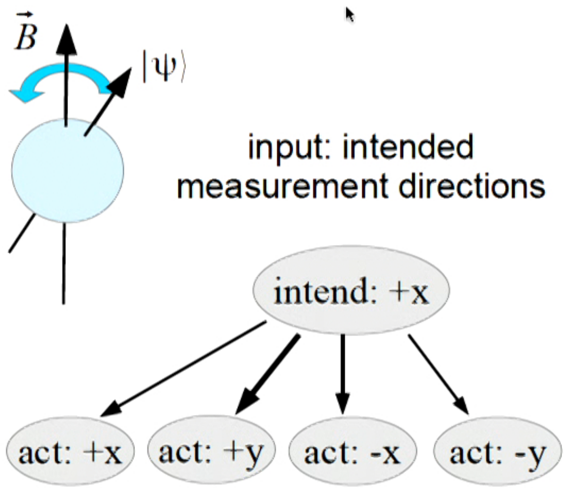

output: actual measurement directions, compensating for the magnetic field

[Tiersch et al, Sci Rep 5, 12874 (2015)]

Designing optics experiments

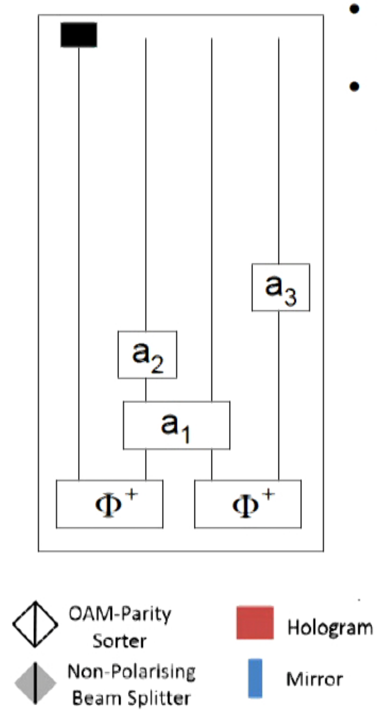

- actions: place optical elements in beams
- reward: entanglement (Schmidt rank vector)

[Melnikov et al, arXiv:1706.00868]

PS agents for quantum experiments

Measuring a spin in a stray magnetic field

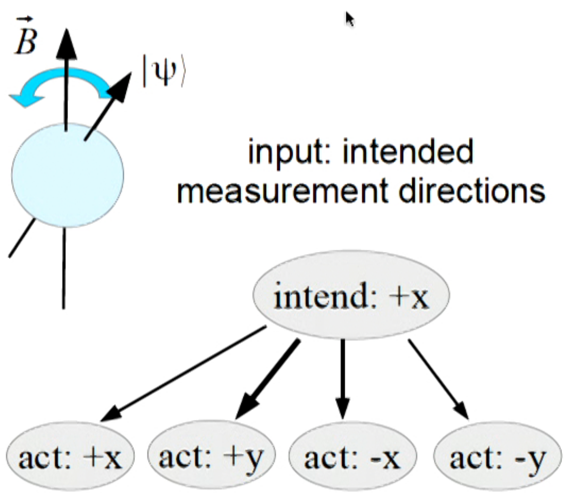

output: actual measurement directions, compensating for the magnetic field

[Tiersch et al, Sci Rep 5, 12874 (2015)]

Designing optics experiments

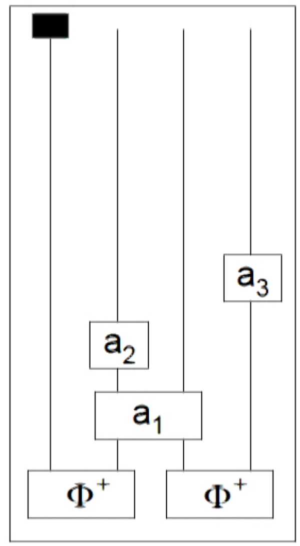

- actions: place optical elements in beams
- reward: entanglement (Schmidt rank vector)
- result: discovery of useful building blocks, e.g. parity sorters

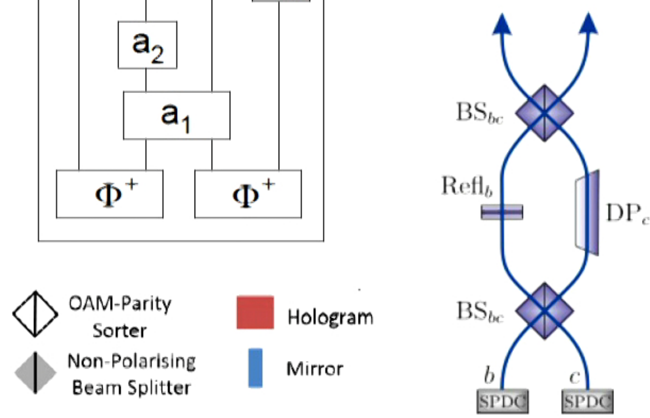

[Melnikov et al, arXiv:1706.00868]

Agency: making your own decisions

What does it mean to be an agent?

- Agents pursue their own goals.
- Agents are aware of the consequences of their actions.
- Agents are responsible for these consequences.
- Agents are at the root of causal chains.
- Agents decide for themselves what to do.
- ...

Agency: making your own decisions

What does it mean to be an agent?

- Agents pursue their own goals.
- Agents are aware of the consequences of their actions.
- Agents are responsible for these consequences.
- Agents are at the root of causal chains.
- Agents decide for themselves what to do.
- ...

Question of free will: If the laws of nature completely specify the dynamics of the physical system in which decisions are made, then in what sense can the decision be attributed to the agent?

Agency: making your own decisions

Where does the decision come from?

```python

class LearningAgent(object):

    def __init__(self):
        self.num_percepts=4
        self.num_actions=2
        self.hmatrix=np.ones([self.num_percepts,self.num_actions])

    def deliberate(self, percept):
        """Given a percept, return an action."""
        probabilities=normalize(self.hmatrix[percept])
        action=np.random.choice(range(self.num_actions),p=probabilities)
        return(action)

    def learn(self,percept, action, reward):
        """Given a reward, update the h-matrix."""
        self.hmatrix[percept,action]+=reward

```

Agency: making your own decisions

Where does the decision come from?

learning rules
(provided by the programmer,
stored in the agent)

```python
class LearningAgent(object):

    def __init__(self):
        self.num_percepts=4
        self.num_actions=2
        self.hmatrix=np.ones([self.num_percepts,self.num_actions])

    def deliberate(self, percept):
        """Given a percept, return an action."""
        probabilities=normalize(self.hmatrix[percept])
        action=np.random.choice(range(self.num_actions),p=probabilities)
        return(action)

    def learn(self,percept, action, reward):
        """Given a reward, update the h-matrix."""
        self.hmatrix[percept,action]+=reward
```

Agency: making your own decisions

Where does the decision come from?

learning rules
(provided by the programmer, stored in the agent)

```
class LearningAgent(object):

    def __init__(self):
        self.num_percepts=4
        self.num_actions=2
        self.hmatrix=np.ones([self.num_percepts,self.num_actions])

    def deliberate(self, percept):
        """Given a percept, return an action."""
        probabilities=normalize(self.hmatrix[percept])
        action=np.random.choice(range(self.num_actions),p=probabilities)
        return(action)

    def learn(self,percept, action, reward):
        """Given a reward, update the h-matrix."""
        self.hmatrix[percept,action]+=reward
```

(pseudo-) randomness

relevant information collected from the environment and stored in the agent

The Sally-Anne test of theory of mind

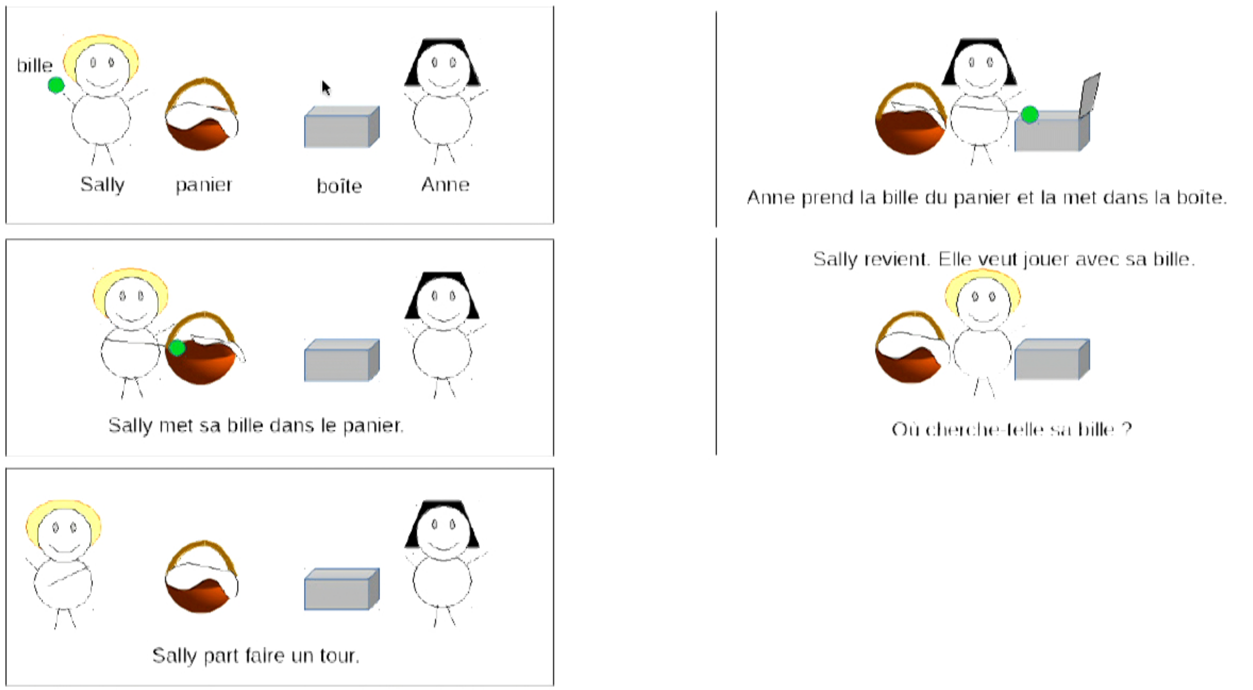

[image: Fschwarzentruber, CC-BY-SA]

The Sally-Anne test of theory of mind

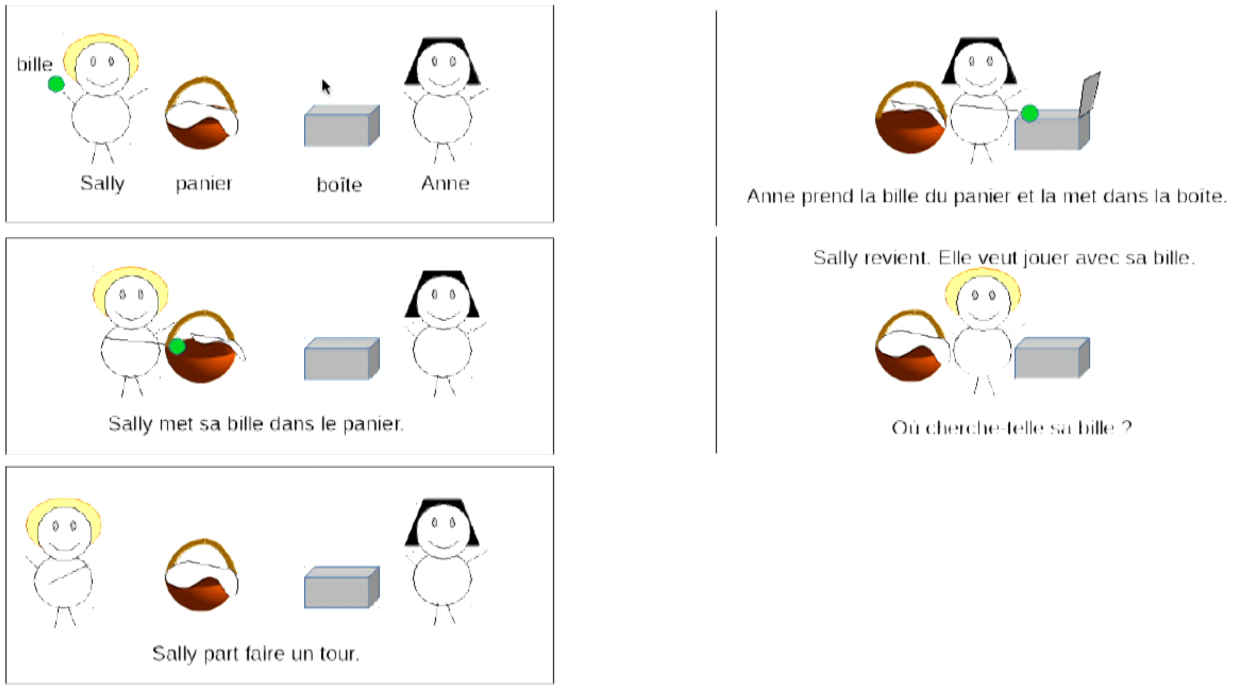

[image: Fschwarzentruber, CC-BY-SA]

Interaction and collective behaviour

Interaction and collective behaviour

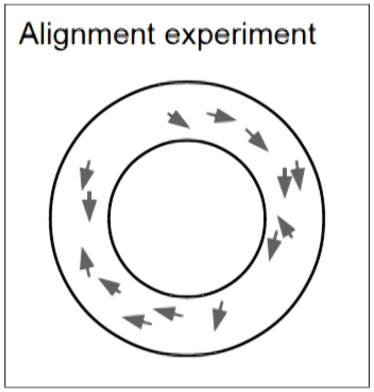

Interaction and collective behaviour

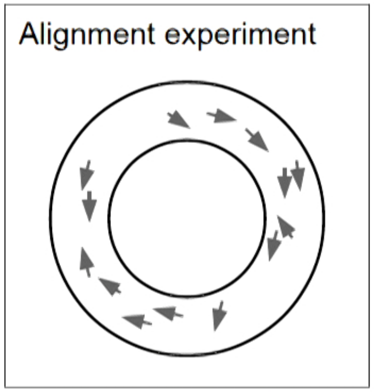

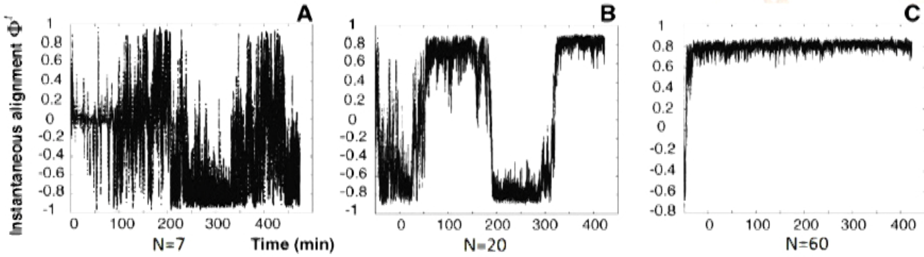

[Buhl et al, Science 312, 5778 (2006)]

Interaction and collective behaviour

Alignment experiment

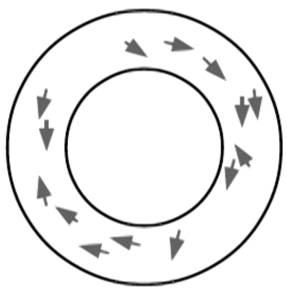

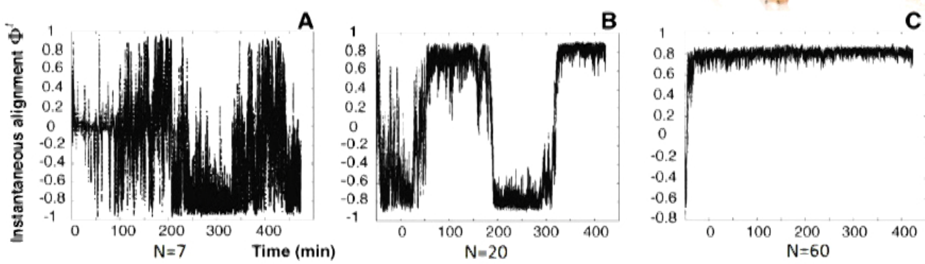

[Buhl et al, Science 312, 5778 (2006)]

Model 1: self-propelled particles (SPP)

average angular velocity of walking animals

$$\phi = \Sigma_{k\ walking}\ \omega_k$$

probability of turning

$$p_k^{\text{turn}} = \begin{cases} \alpha + \beta_{\text{W}} |\phi(t)| & h_k(t)\phi(t) < 0 \\ \alpha + \beta_{\text{A}} |\phi(t)| & h_k(t)\phi(t) \geq 0 \end{cases}$$

[Ariel et al, PLoS One 9, e101636 (2014)]

Interaction and collective behaviour

Alignment experiment

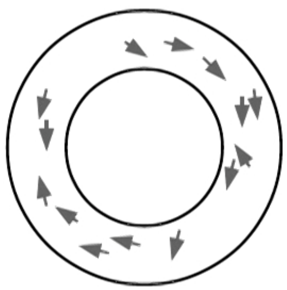

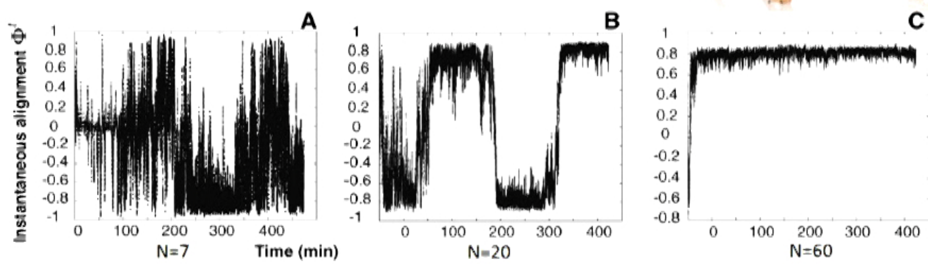

[Buhl et al, Science 312, 5778 (2006)]

Model 1: self-propelled particles (SPP)

average angular velocity of walking animals

$$\phi = \Sigma_{k\ walking}\, \omega_k$$

probability of turning

$$p_k^{\mathrm{turn}} = \begin{cases} \alpha + \beta_{\mathrm{W}} |\phi(t)| & h_k(t)\phi(t) < 0 \\ \alpha + \beta_{\mathrm{A}} |\phi(t)| & h_k(t)\phi(t) \geq 0 \end{cases}$$

[Ariel et al, PLoS One 9, e101636 (2014)]

Model 2: learning individual responses

- percepts: net movement in vicinity
- actions: turn or continue
- reward: alignment with neighbours

Interaction and collective behaviour

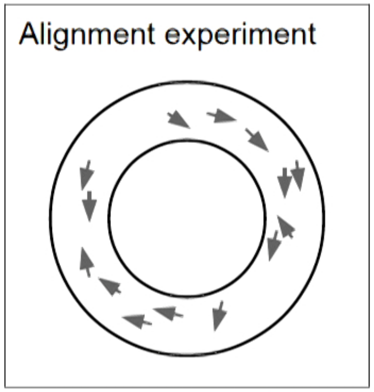

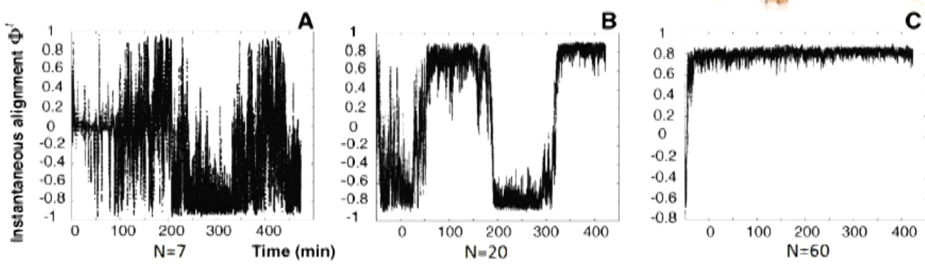

[Buhl et al, Science 312, 5778 (2006)]

Collective motion of PS agents at different densities

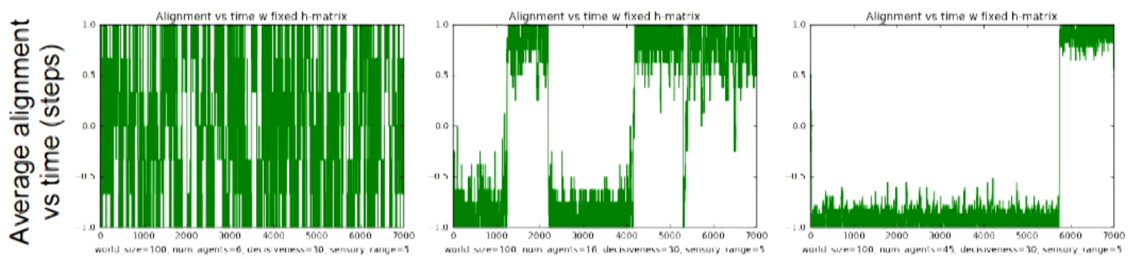

num_agents=6 num_agents=16 num_agents=45

Interaction and collective behaviour

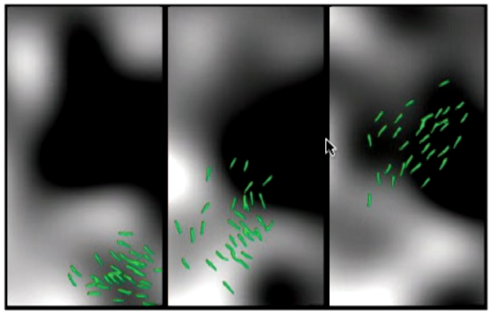

[Berdahl et al, Science 339, 574 (2013)]

Collective sensing allows swarms to locate resources better than isolated individuals in the case of noisy spatial distribution.

Interaction and collective behaviour

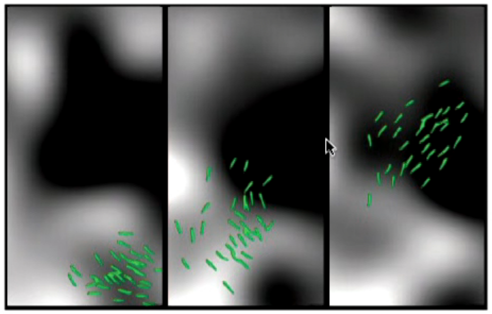

[Berdahl et al, Science 339, 574 (2013)]

Collective sensing allows swarms to locate resources better than isolated individuals in the case of noisy spatial distribution.

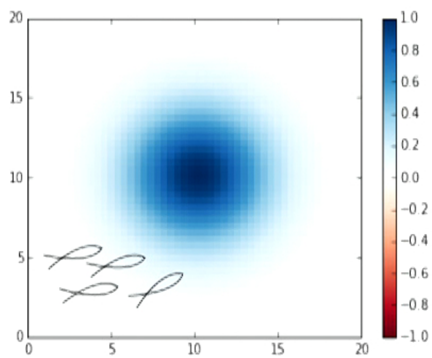

Interaction and collective behaviour

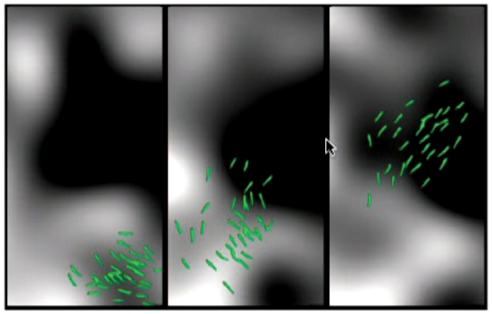

[Berdahl et al, Science 339, 574 (2013)]

Collective sensing allows swarms to locate resources better than isolated individuals in the case of noisy spatial distribution.

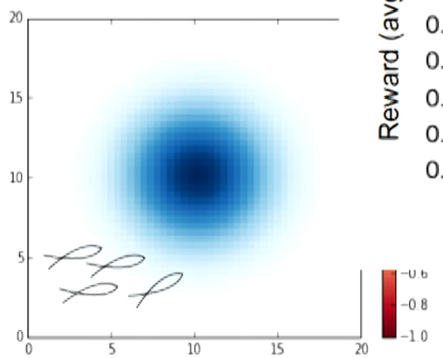

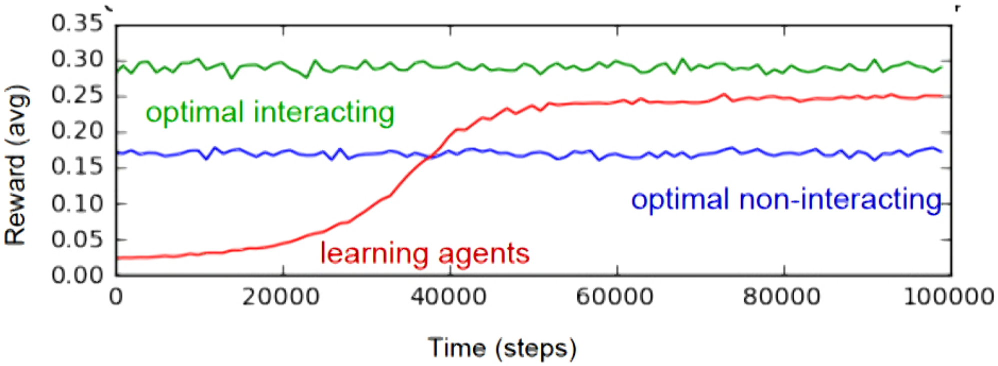

Learning structured models

Motivation

- performance at practical tasks:
  - handling complex environments
  - generalization to novel situations
  - handling continuous variables autonomously
- toy model of high-level abilities:
  - recognize context
  - form abstractions, categories
  - think causally
- memory-efficient representation

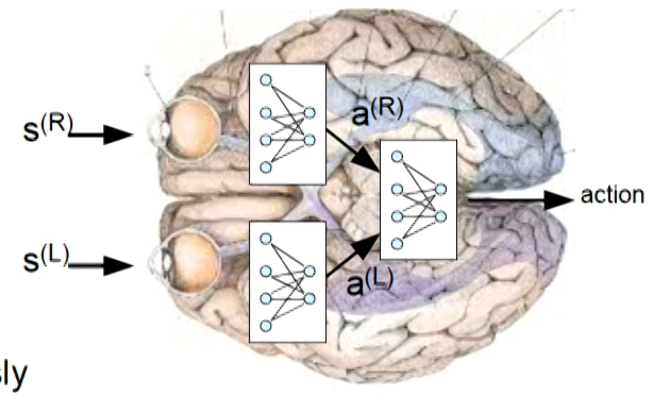

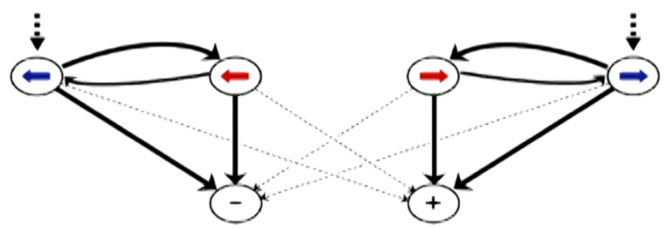

[Briegel and de las Cuevas, Sci Rep 2, 400 (2012)]

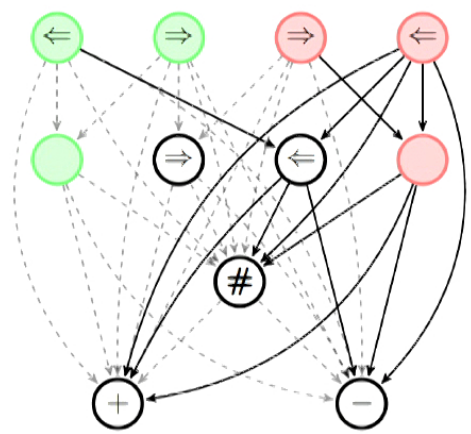

[Melnikov et al, arXiv:1504.02247]

Learning structured models

Motivation

- performance at practical tasks:
  - handling complex environments
  - generalization to novel situations
  - handling continuous variables autonomously
- toy model of high-level abilities:
  - recognize context
  - form abstractions, categories
  - think causally
- memory-efficient representation

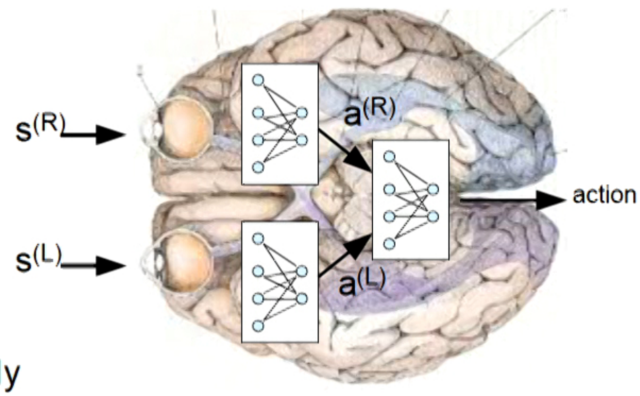

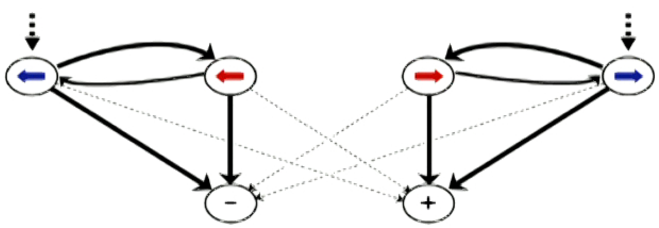

[Briegel and de las Cuevas, Sci Rep 2, 400 (2012)]

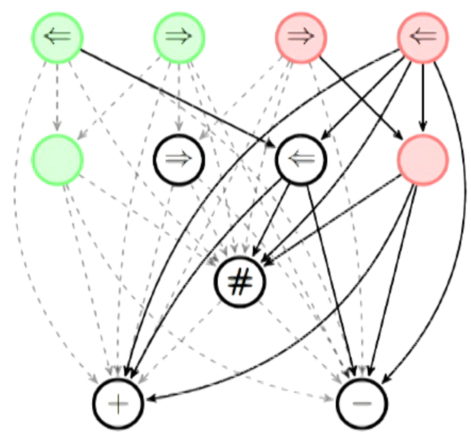

[Melnikov et al, arXiv:1504.02247]

Learning structured models

1. Deliberation on a given ECM
   - multiple parallel walks
   - attractive interaction

Learning structured models

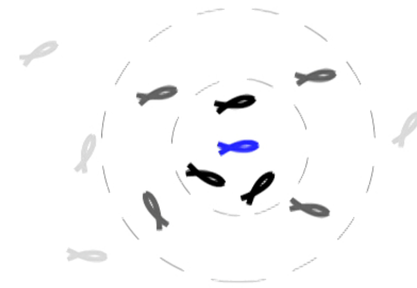

1. Deliberation on a given ECM
   - multiple parallel walks
   - attractive interaction

Learning structured models

1. Deliberation on a given ECM
   - multiple parallel walks
   - attractive interaction

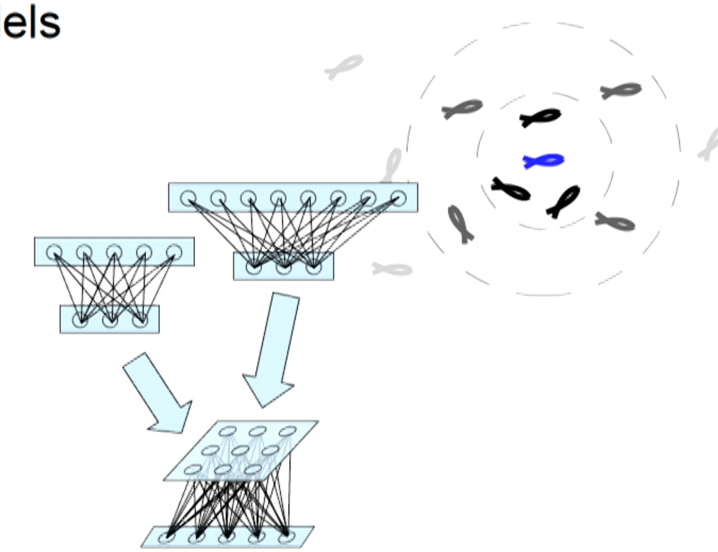

Learning structured models

1. Deliberation on a given ECM
   - multiple parallel walks
   - attractive interaction

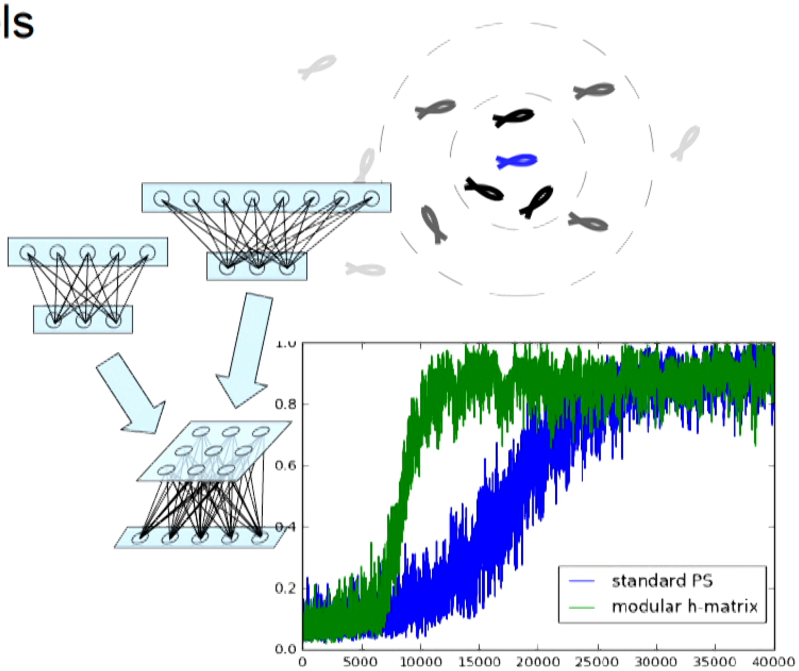

Learning structured models

1. Deliberation on a given ECM
   - multiple parallel walks
   - attractive interaction

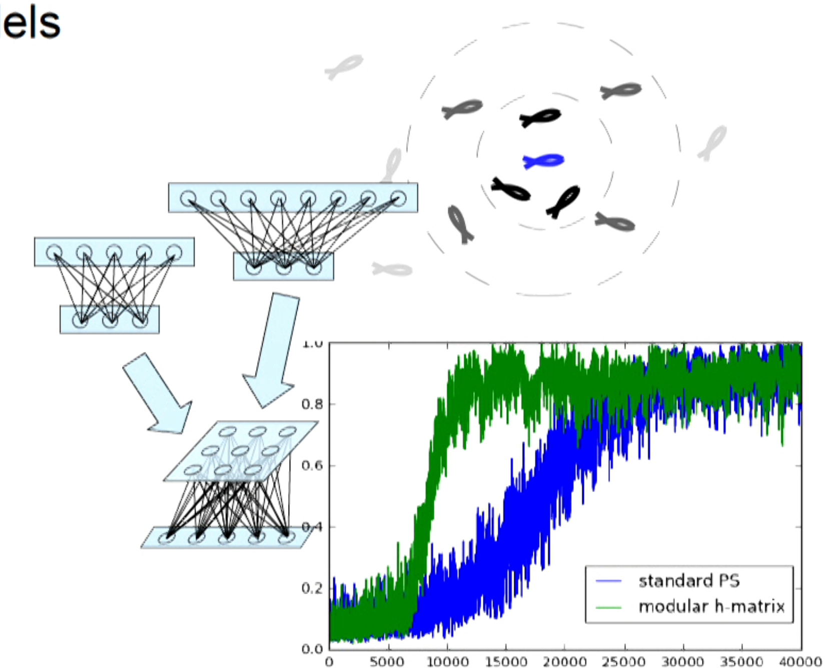

2. Learning the ECM structure
   - capture operational equivalence
   - create abstractions unifying operationally equivalent sets
   - connect proposed abstractions to existing network

Conclusions

- Quantum advantages for machine learning:
  - more compact representation and linear algebra speedup
  - optimization with annealers
- Applying classical machine learning techniques to quantum systems:
  - useful tool for characterization and control
  - discovery of phases, patterns, design primitives

Conclusions

- Quantum advantages for machine learning:
  - more compact representation and linear algebra speedup
  - optimization with annealers
- Applying classical machine learning techniques to quantum systems:
  - useful tool for characterization and control
  - discovery of phases, patterns, design primitives

- Autonomous learning agents as toy models for a range of phenomena:
  - Agents can make their 'own' decisions using a combination of pre-programmed rules, memory and (pseudo-)randomness.
  - Collective behaviour can be modelled at an individual level, with interaction rules being learned rather than proposed ad hoc.
  - Agents endowed with learning rules that capture structure can build abstract models of their environment.

Briegel and Dür groups

 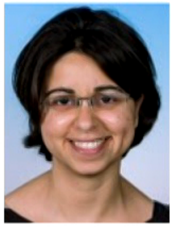  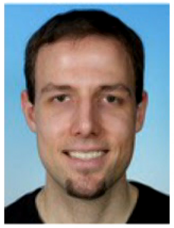 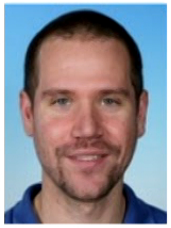

universität innsbruck