All Else Being Equal Be Empowered – Klyubin, Polani, Nahaniv

Entropy Discussion Group
26/6/2006

Joseph Lizier
PhD student
The University of Sydney – Prof. Albert Zomaya
CSIRO ICT Centre - Dr. Mikhail Prokopenko

All Else Being Equal Be Empowered – Klyubin, Polani, Nahaniv

All Else Being Equal Be Empowered

Alexander S. Klyubin, Daniel Polani, and Chrystopher L. Nehaniv

Adaptive Systems Research Group
School of Computer Science, Faculty of Engineering and Information Sciences
University of Hertfordshire, College Lane, Hatfield Herts AL10 9AB, UK
{A.Kljubin,D.Polani,C.L.Nehaniv}@herts.ac.uk

Abstract. The classical approach to using utility functions suffers from the drawback of having to design and tweak the functions on a case by case basis. Inspired by examples from the animal kingdom, social sciences and games we propose *empowerment*, a rather universal function, defined as the information-theoretic capacity of an agent's actuation channel. The concept applies to any sensorimotoric apparatus. Empowerment as a measure reflects the properties of the apparatus as long as they are observable due to the coupling of sensors and actuators via the environment.

Contents

1. Motivation
2. Derive concept of Empowerment
3. Clarify use with examples:
    1. Pushable box
    2. Maze navigation

Motivation: Generic Utility Functions

- Generic Utility Functions wanted for adaptive systems:
  - To avoid specific design using prior knowledge.
  - To help our understanding of nature ("Is there a more general principle?")
  - To be locally computable.
  - To be universally applicable.
- Examples include homeostasis
- Observations from nature: control and influence are key

The empowerment concept

- Empowerment = *"the perceived amount of influence or control the agent has over the world"*
  - Links back to perception-action loop concept:
    - "The world" is the environment
    - Perception is important
  - Agent's *potential* to change the world

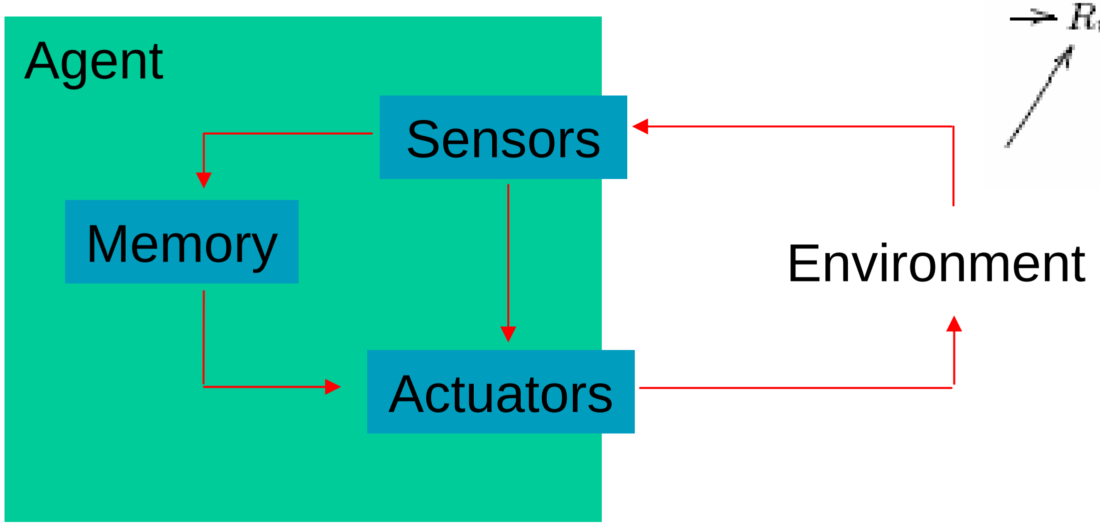

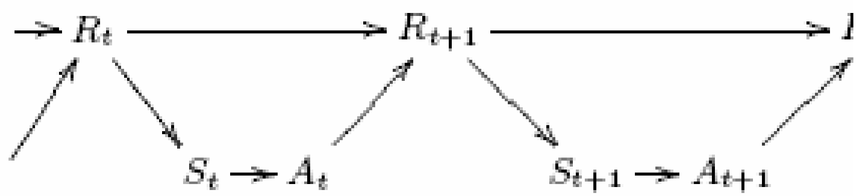

Empowerment defined – High level

- Empowerment ≡ Capacity of the actuation channel of the agent (i.e. channel from actuators to sensors via environment)
- Channel capacity (Shannon) = maximum MI for source (x) and receiver (y) over all possible distributions of transmitted signal.

$$C = \max_{p(x)} I(X;Y)$$

- Taking maximum MI over all probability distributions of the input:
  - Removes dependence on the actual input / actuation sequence
  - Removes asymmetry, makes a causal metric
  - Gives true representation of the *potential* to change the world.

CSIRO

Empowerment defined - Quantification

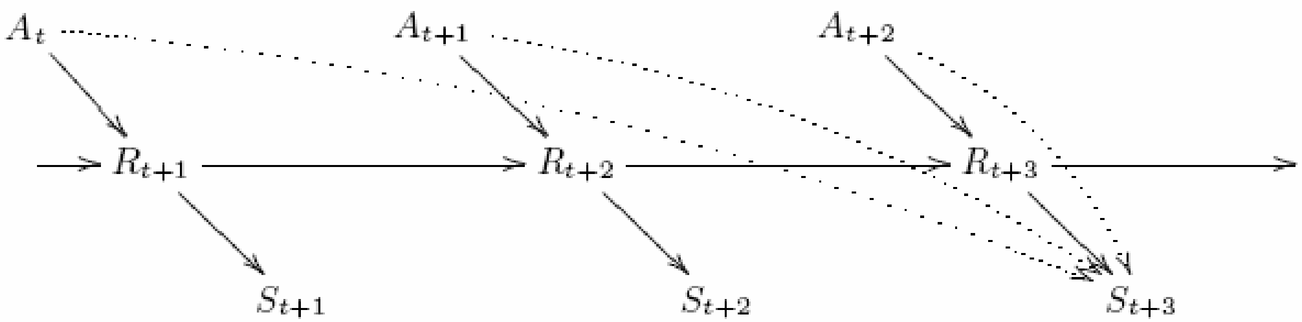

- n-step empowerment: Use n previous steps of input – source = A_t^n (vector length n):

$$\mathfrak{E} = C = \max_{p(a_t^n)} \sum_{\mathcal{A}^n, \mathcal{S}} p(s_{t+n}|a_t^n) p(a_t^n) \log_2 \frac{p(s_{t+n}|a_t^n)}{\sum_{\mathcal{A}^n} p(s_{t+n}|a_t^n) p(a_t^n)} .$$

- This rearranges to have the standard form of MI:

$$\varepsilon = \max_{p(a_t^n)} \sum_{A^n, S} p(s_{t+n}, a_t^n) \log \left[\frac{p(s_{t+n}, a_t^n)}{p(s_{t+n}) p(a_t^n)} \right]$$

  - But $p(s|a_n)$ formulation is used here to emphasise dependence of sensor inputs on actuator outputs, and maximization over $p(a_n)$.

Empowerment defined – Understanding it

- **Key point**: Distribution of p(a) to maximise the mutual information = distribution of actions to inject max info into sensors
  - **Max info into sensors → *maximum number* of *equally likely* achievable sensor states.**
  - **Deterministic environment: empowerment = log of achievable sensor states**
  - **Stochastic environment: emp will be lower**
- **Outcome**: $p(s|a_n)$ may induce equivalence classes over the set of action sequences a_n:
  - Greater diversity of sensory inputs resulting from available actions = (Intuitively) More control available to agent = (Quantitatively) Higher empowerment.
  - Lower diversity of sensory inputs resulting from available actions = (Intuitively) Less control available to agent = (Quantitatively) Lower empowerment.

Example 1 – Pushable box

- Agent wanders around 2D grid. Four situations: empowerment measured over 5 steps as function of start position of agent.

	Stationary box		Pushable box	
Agent can't perceive box	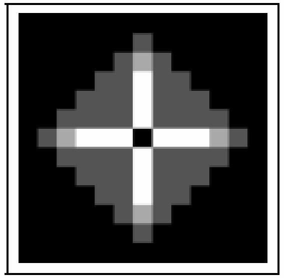 a. $\mathfrak{C} \in [5.86; 5.93]$	Max Emp away from box	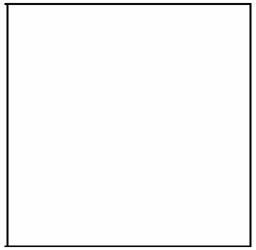 b. $\mathfrak{C} = \log_2 61 \approx 5.93$ bit	Max Emp everywhere
Agent perceives box	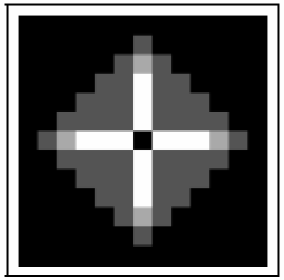 c. $\mathfrak{C} \in [5.86; 5.93]$	Max Emp away from box	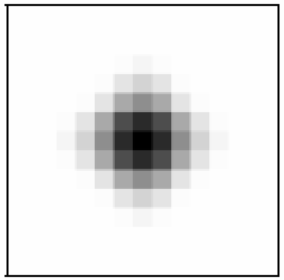 d. $\mathfrak{C} \in [5.93; 7.79]$	Max Emp close to box

- Empowerment is maximised here where agent's actions have greatest diversity of results (think of equivalence classes of actions).

Example 2 – Maze navigation

- 2D grid with obstacles = maze. Sensor = agent's absolute position.
- n-step empowerment measured as a function of position and n.

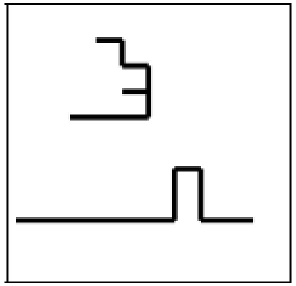

n = 1, 2, 5, 10

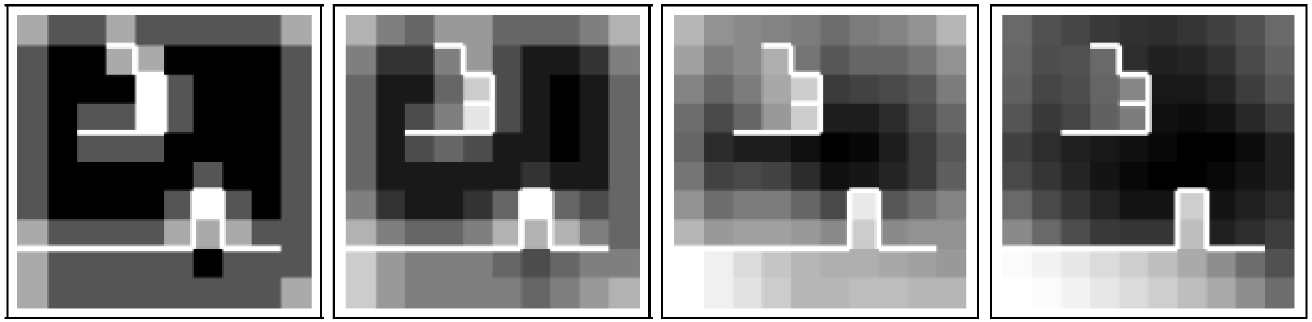

$\mathfrak{E} \in [1; 2.32]$ | $\mathfrak{E} \in [1.58; 3.70]$ | $\mathfrak{E} \in [3.46; 5.52]$ | $\mathfrak{E} \in [4.50; 6.41]$

- Empowerment max where agent can quickly reach more cells.
- In fact, direct correlation with average (shortest) path to other cells:

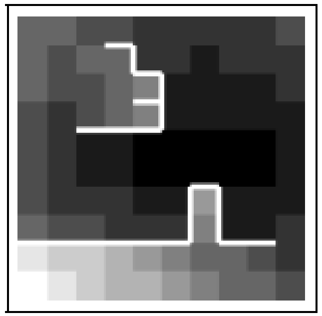

Empowerment - Summary

- Empowerment ≡ Capacity of the actuation channel of the agent (i.e. channel from actuators to sensors via environment)
- A general purpose utility function.
- Usefulness of empowerment:
  1. Universal definition
  2. Simple interpretation
  3. Easy to compute on-board an agent
  4. Useful as fitness function.