




Information Theory for Intelligent Information Processing

Sander G. van Dijk

Adaptive Systems Research Group
University of Hertfordshire


October 15, 2009





"I visualize a time when we will be to robots what dogs are to humans, and I'm rooting for the machines."

Claude E. Shannon



"Information is the resolution of uncertainty."

Claude E. Shannon

C.E. Shannon: A Mathematical Theory of Communication. Bell System Technical Journal **27**
(1948) 379-423

Outline

- 1 Information Theory and Agents
- 2 Empowerment
- 3 Relevant Information
- 4 Hierarchical Behaviours


Outline

- 1 Information Theory and Agents
- 2 Empowerment
- 3 Relevant Information
- 4 Hierarchical Behaviours

Entropy

- How much choice involved in selection of event
- How uncertain are we of the outcome

$$H(X) = - \sum_{x \in \mathcal{X}} p(x) \log p(x)$$




Mutual Information

- Mutual dependence between two random variables
- The amount of uncertainty in Y which is removed by knowing X

$$\begin{aligned}
 I(X; Y) &= H(Y) - H(Y|X) \\
 &= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \\
 &= \sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y|x) \log \frac{p(y|x)}{\sum_{x \in \mathcal{X}} p(y|x)p(x)}
 \end{aligned}$$


Channels



Channel capacity: maximum amount of information the received signal can contain about the transmitted signal.

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

Sensor-Actuator Loops



$$p(r_{t+1}|r_t, a_t)$$

$$p(a_t|s_t)$$

Outline


- 1 Information Theory and Agents
- 2 Empowerment**
- 3 Relevant Information
- 4 Hierarchical Behaviours

Empowerment


- *Perceivable* amount of influence or *control* over the world
- How much *information* can an agent “inject” via a sequence of *actions* into its *sensor*?

A. S. Klyubin, D. Polani and C. L. Nehaniv: Empowerment: A Universal Agent-Centric Measure of Control. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation 1* (2005) 128-135

Empowerment





Empowerment




Empowerment

$$A_t^n = (A_t, A_{t+1}, \dots, A_{t+n})$$


$$\mathfrak{E}_t = C(p(s_{t+n})) = \max_{p(a_t^n)} I(A_t^n; S_{t+n})$$

Figure: $n = 1$ Figure: $n = 10$

Empowerment-Maximizing Control




Empowerment-Maximizing Control



Outline

- 1 Information Theory and Agents
- 2 Empowerment
- 3 Relevant Information**
- 4 Hierarchical Behaviours

Relevant Information



“Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem.”

Claude E. Shannon

Relevant Information

- Not all available information has meaning to an agent
- How much information is actually *relevant*?

	a_1	a_2	a_3	a_4
s_1	2	1	0	0
s_2	1	2	0	0
s_3	0	0	2	1
s_4	0	0	1	2

Polani, D., Nehaniv, C., Martinetz, T. and Kim, J. T.: Relevant Information in Optimized Persistence vs. Progeny Strategies. In M.Rocha, L., Bedau, M., Floreano, D., Goldstone, R., Vespignani, A., and Jaeger, L., editors, *Proc. Artificial Life X* (2006)

Relevant Information


- Amount of information “extracted” from sensor for action selection:
 $I(S; A)$
- Minimize required relevant information while maintaining optimal behavior:

$$I(S; A^*) = \min_{p(a|s): p(a|s) \Rightarrow a \text{ optimal for } s} I(S; A)$$

- Trade-off utility/information:

$$\min_{p(a|s)} (I(S; A) - \beta \mathbb{E}[U(s, a)])$$

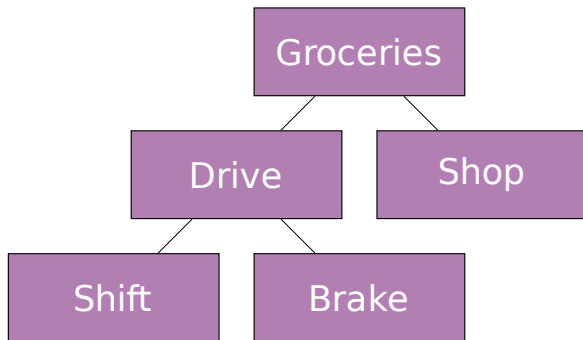
Relevant Information



Outline

- 1 Information Theory and Agents
- 2 Empowerment
- 3 Relevant Information
- 4 Hierarchical Behaviours**

Hierarchical Behaviour



Goal

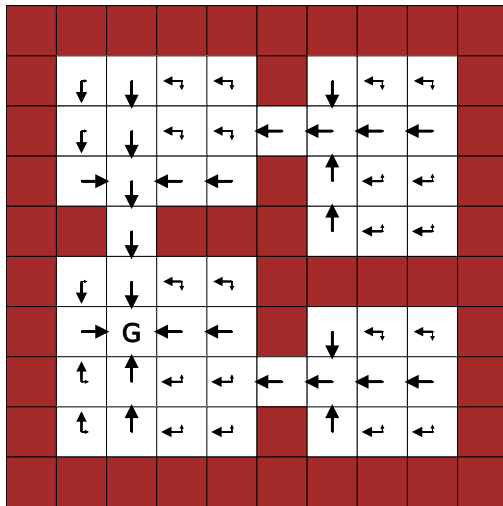
Systematic rather than *heuristic* understanding

Hypothesis


Benefits are grounded in their effect on the way an agent processes information:

- Divide the burden of information processing over different layers
- Information about the current context is retained in more abstract behaviours at higher levels

Grid-world




Flat Policy




$$\mathfrak{J}_{flat} = I(S; A)$$

Layered Policy




$$\mathfrak{J}_{layered} = I(S; O) + I(S; A|O)$$

Policy with Memory



$$\mathfrak{J}_{memory} = \frac{1}{2} [I(A_t; S_t) + I(A_{t+1}; S_{t+1} | A_t)]$$

Hierarchical Policy



$$\mathfrak{J}_{\text{hierarchical}} = \frac{1}{2} [I(S_t; O_t) + I(S_t; A_t | O_t) + I(S_{t+1}; A_{t+1} | O_t)]$$

Results

Policy	\mathcal{J}^*
Flat	1
Layered	$0.65 + 0.77 = 1.42$
Memory	$\frac{1}{2}(1.04 + 0.85) = 0.95$
Hierarchical	$\frac{1}{2}(0.6 + 0.94 + 0.35) = 0.95$

Discussion

- Novel quantitative insight in effects of hierarchical structuring on necessary information intake
- Reduction at both parts of the structure and whole system
- Future work:
 - Determine the optimality of existing models of organization
 - Guide search for optimal behaviour organisation in organisms and artificial agents

Thank you!

Questions?