Three levels of description *(David Marr, 1982)*

Computational

Why do things work the way they do?
What is the goal of the computation?
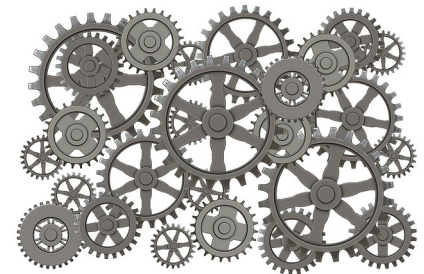What are the unifying principles?

Arithmetic:
$$a + b = c$$

Algorthmic

What representations can implement such computations?
How does the choice of representations determine the algorithm?

00001000110101+
00001000010100=
00010001001001

Implementational

How can such a system be built in hardware?
How can neurons carry out the computations?

Three levels of description *(David Marr, 1982)*

Computational

Why do things work the way they do?
What is the goal of the computation?
What are the unifying principles?

Algorthmic

What representations can implement such computations?
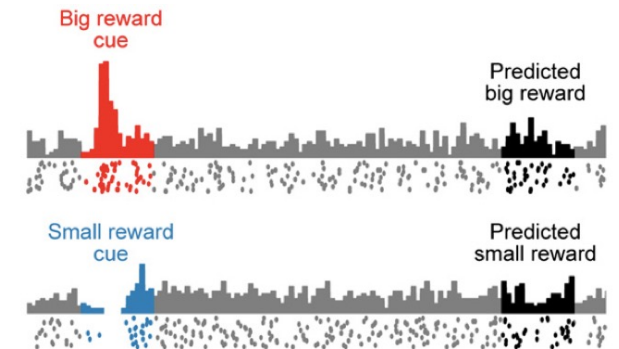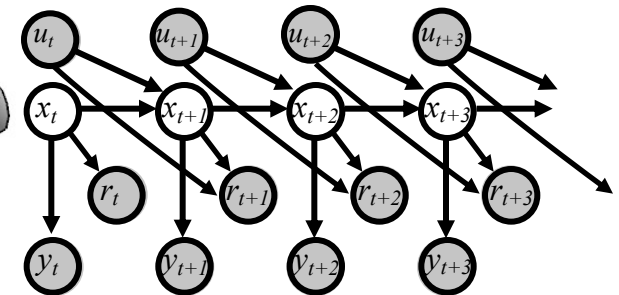How does the choice of representations determine the algorithm?

Implementational

How can such a system be built in hardware?
How can neurons carry out the computations?

maximize:

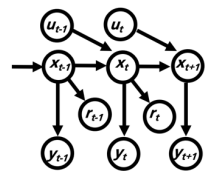$$R_t = r_{t+1} + r_{t+2} + \cdots + r_T$$

Big reward cue — Predicted big reward

Small reward cue — Predicted small reward

Three levels of description *(David Marr, 1982)*

Computational

Why do things work the way they do?
What is the goal of the computation?
What are the unifying principles?

maximize:

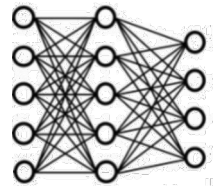$$R_t = r_{t+1} + r_{t+2} + \cdots + r_T$$

Bellman (1960)

Algorthmic

What representations can implement such computations?
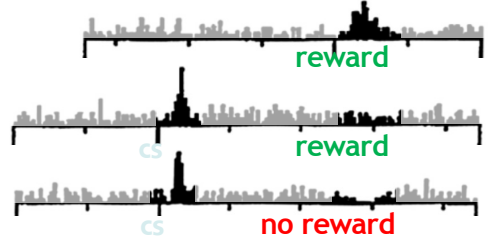How does the choice of representations determine the algorithm?

Temporal-difference-error algorithm
while *episode not terminated* **do**
A ← action given by π
Take action A, observe r_t, S_{t+1}
$\delta_t = r_t + \gamma V(S_{t+1}) - V(S_t)$
$V(S_t) \leftarrow V(S_t) + \alpha \delta_t$
$S_t \leftarrow S_{t+1}$
end

Sutton & Barto (1998)

Implementational

How can such a system be built in hardware?
How can neurons carry out the computations?

reward

reward

no reward

Schultz, Dayan, Montague (1997)