COPENHAGEN
INTERNATIONAL
SOFTWARE DEVELOPMENT
CONFERENCE 2015

goto; conference

Rethinking Data for Intelligent Computing

Julie Pitt
(@yakticus)

goto; conference

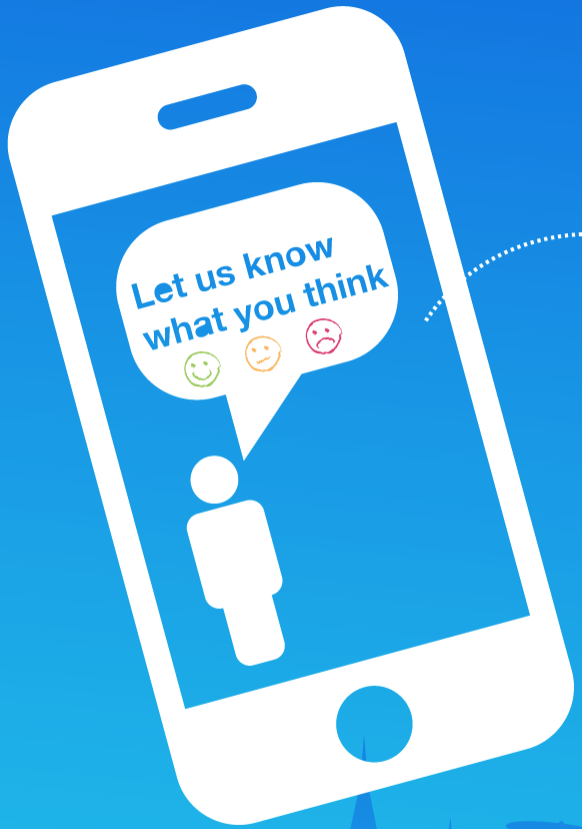

Click ‘engage’
to rate sessions
and ask questions

Join the conversation #gotocph

how I got here

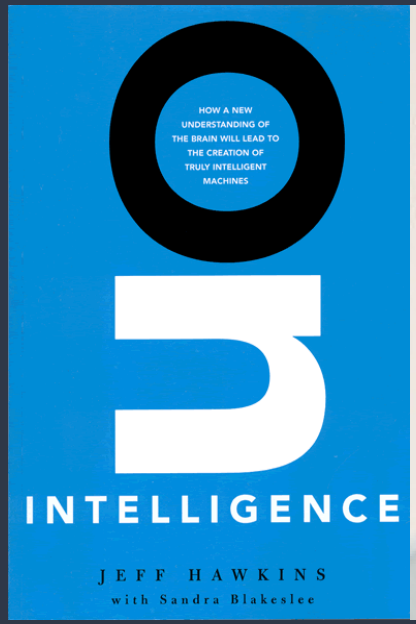

Jeff Hawkins

the problem

build machines capable of intelligent behavior

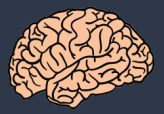

questions

what makes us intelligent?

how does perception work?

how does action work?

how does learning work?

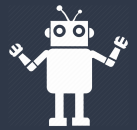

what does this mean for AI and data?

what makes us intelligent?

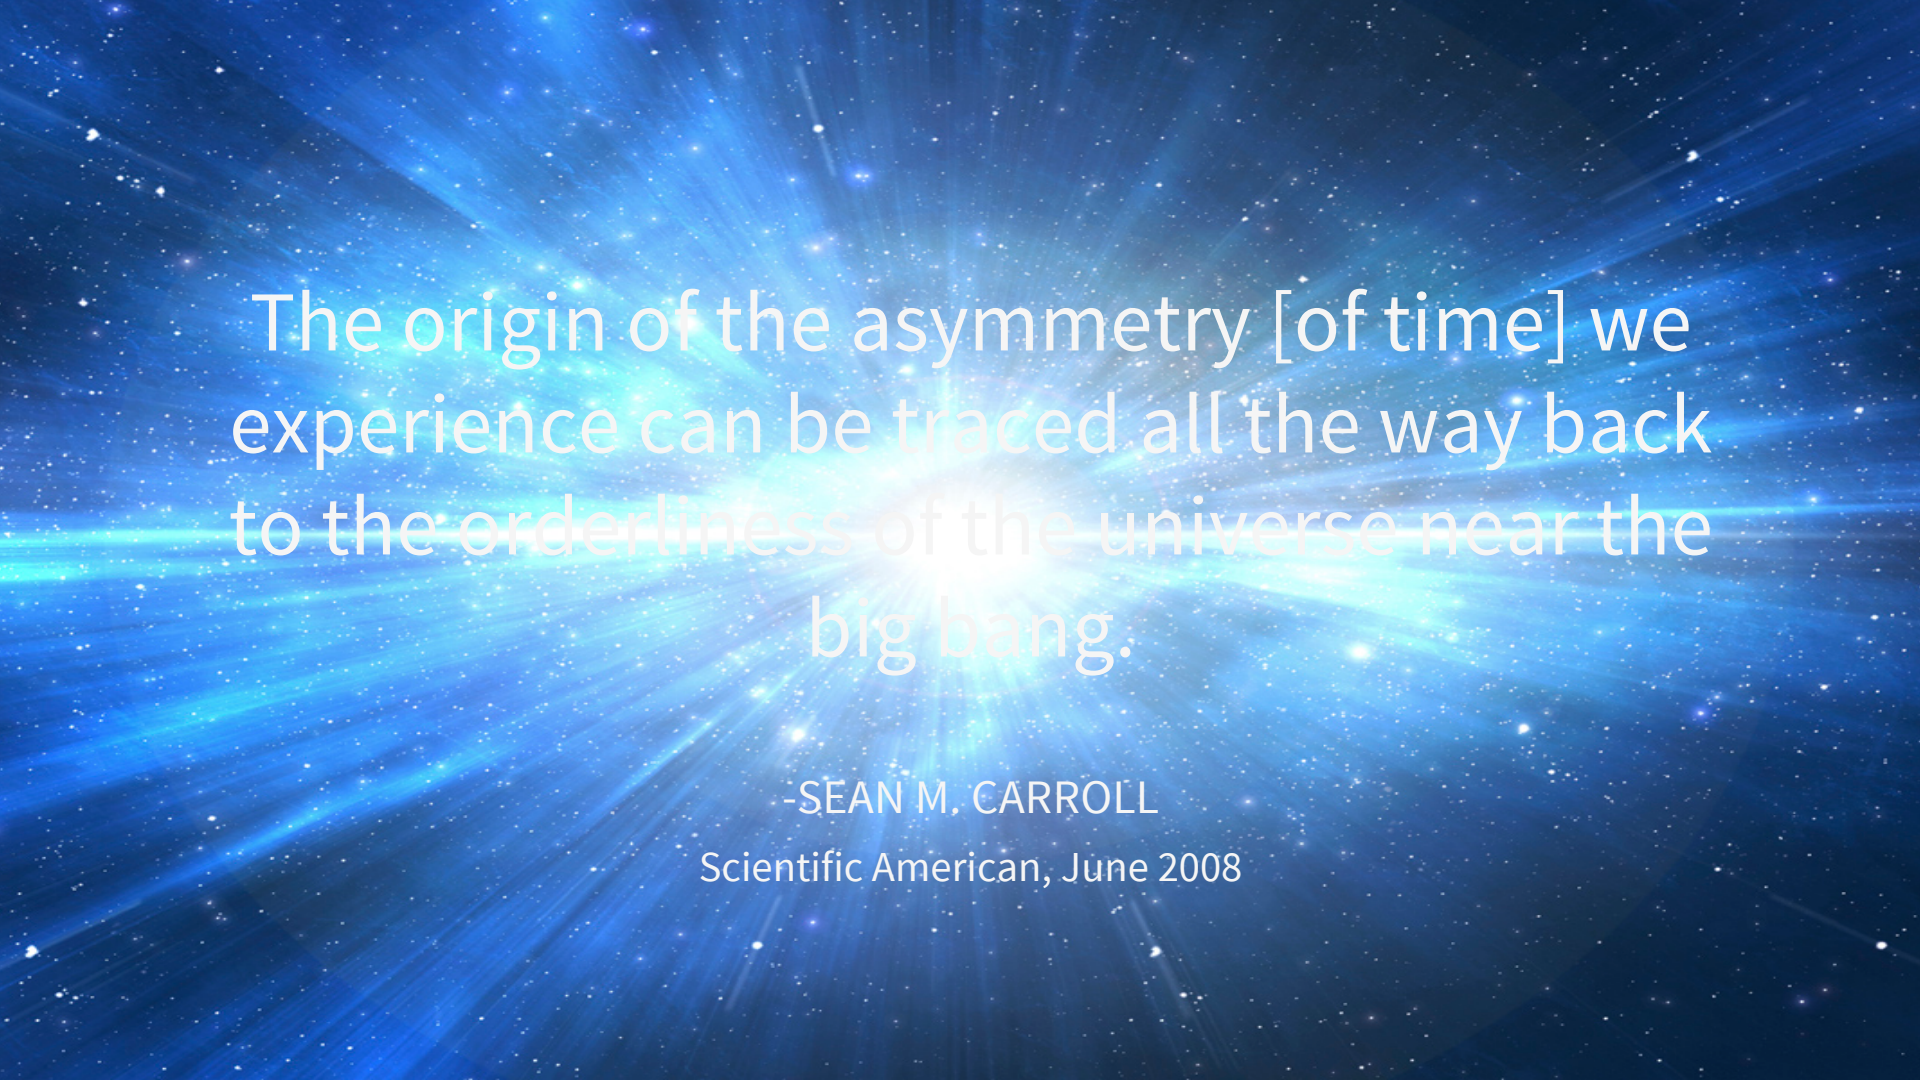
The origin of the asymmetry [of time] we experience can be traced all the way back to the orderliness of the universe near the big bang.
-SEAN M. CARROLL
Scientific American, June 2008

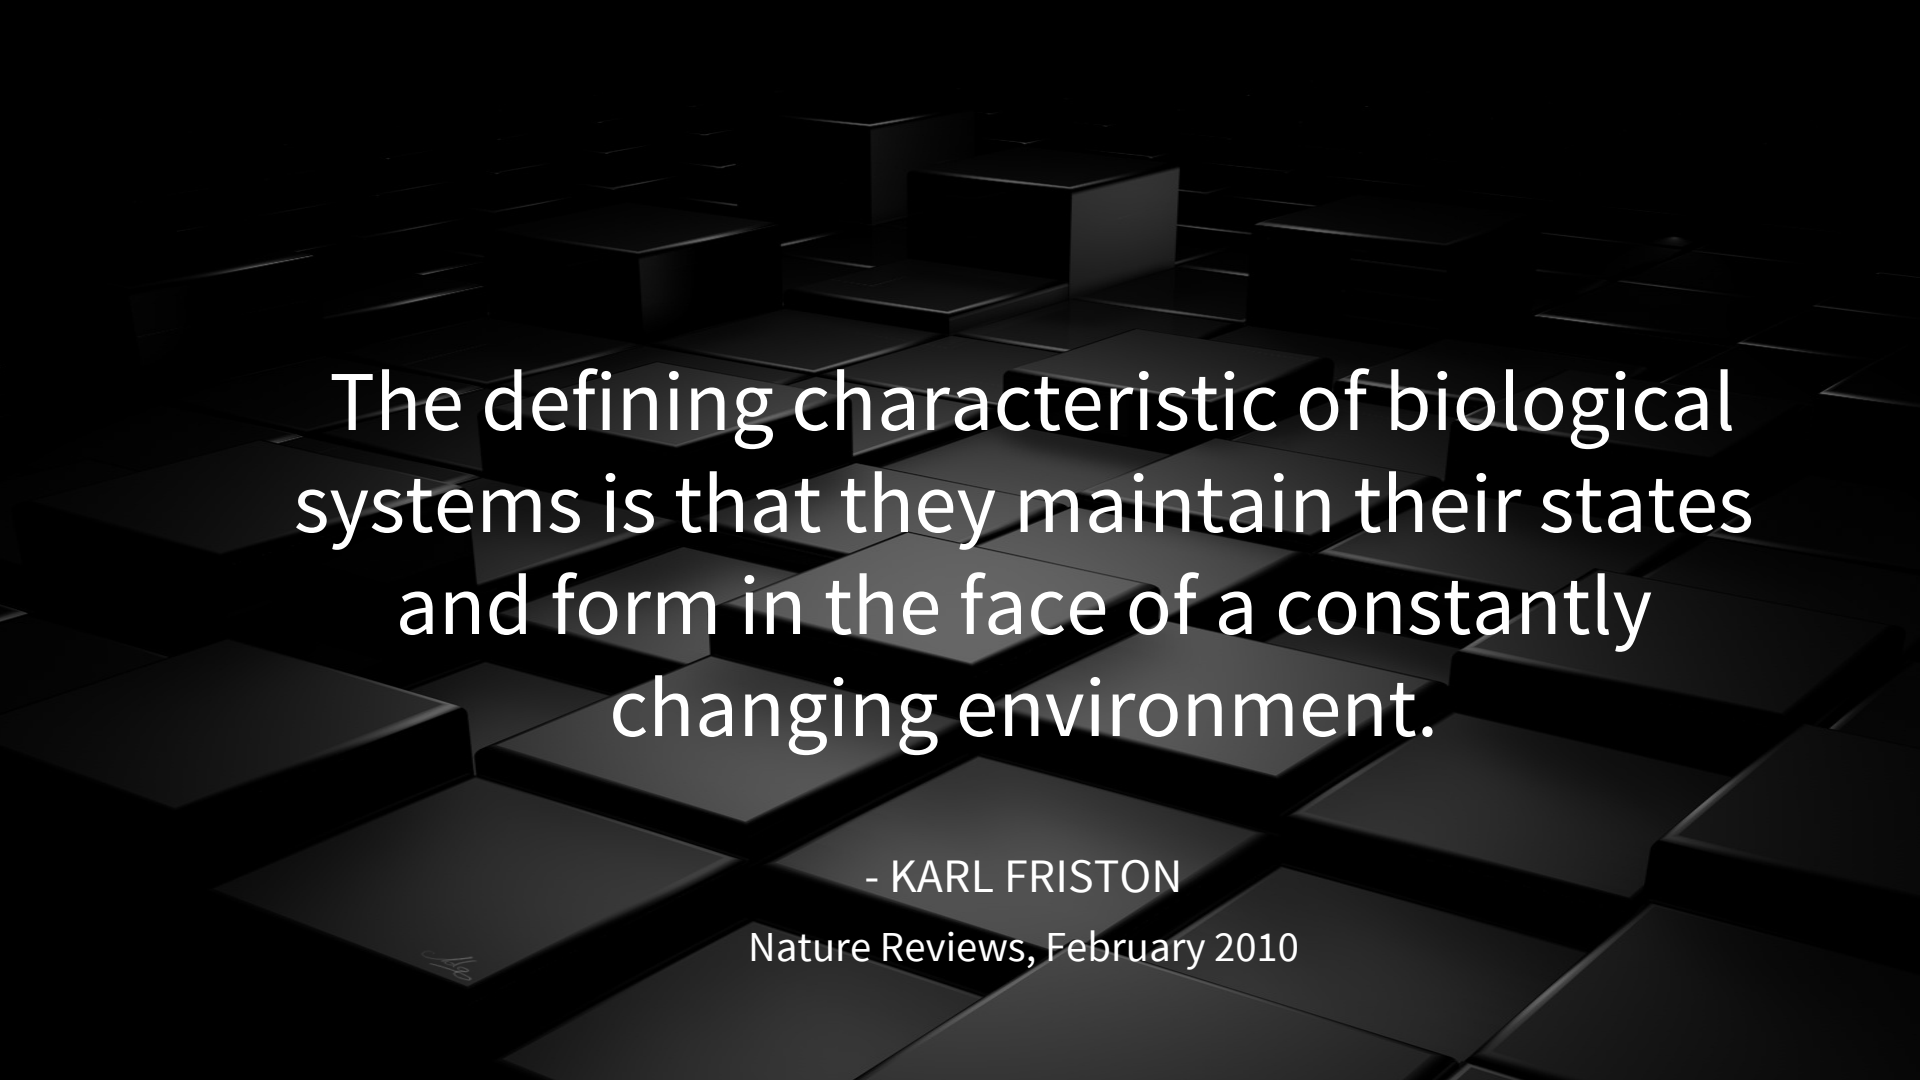
The defining characteristic of biological systems is that they maintain their states and form in the face of a constantly changing environment.
- KARL FRISTON
Nature Reviews, February 2010

free energy principle

Karl Friston

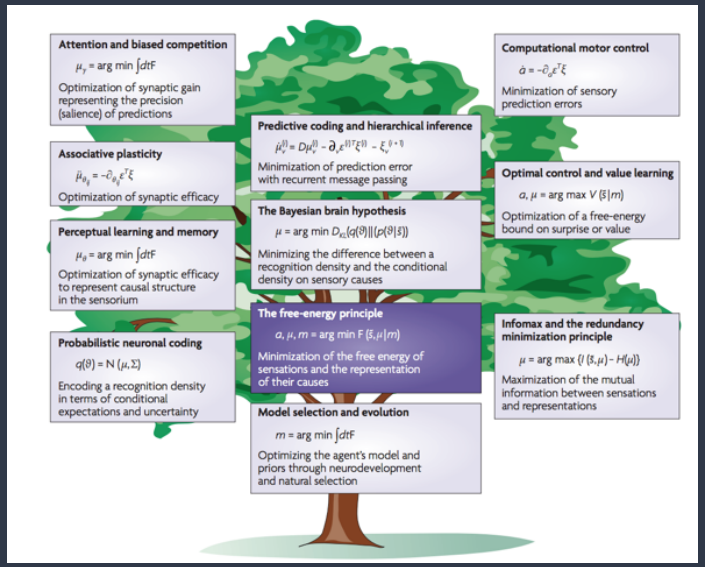

intelligent agents resist entropy

entropy = surprise (averaged over time)

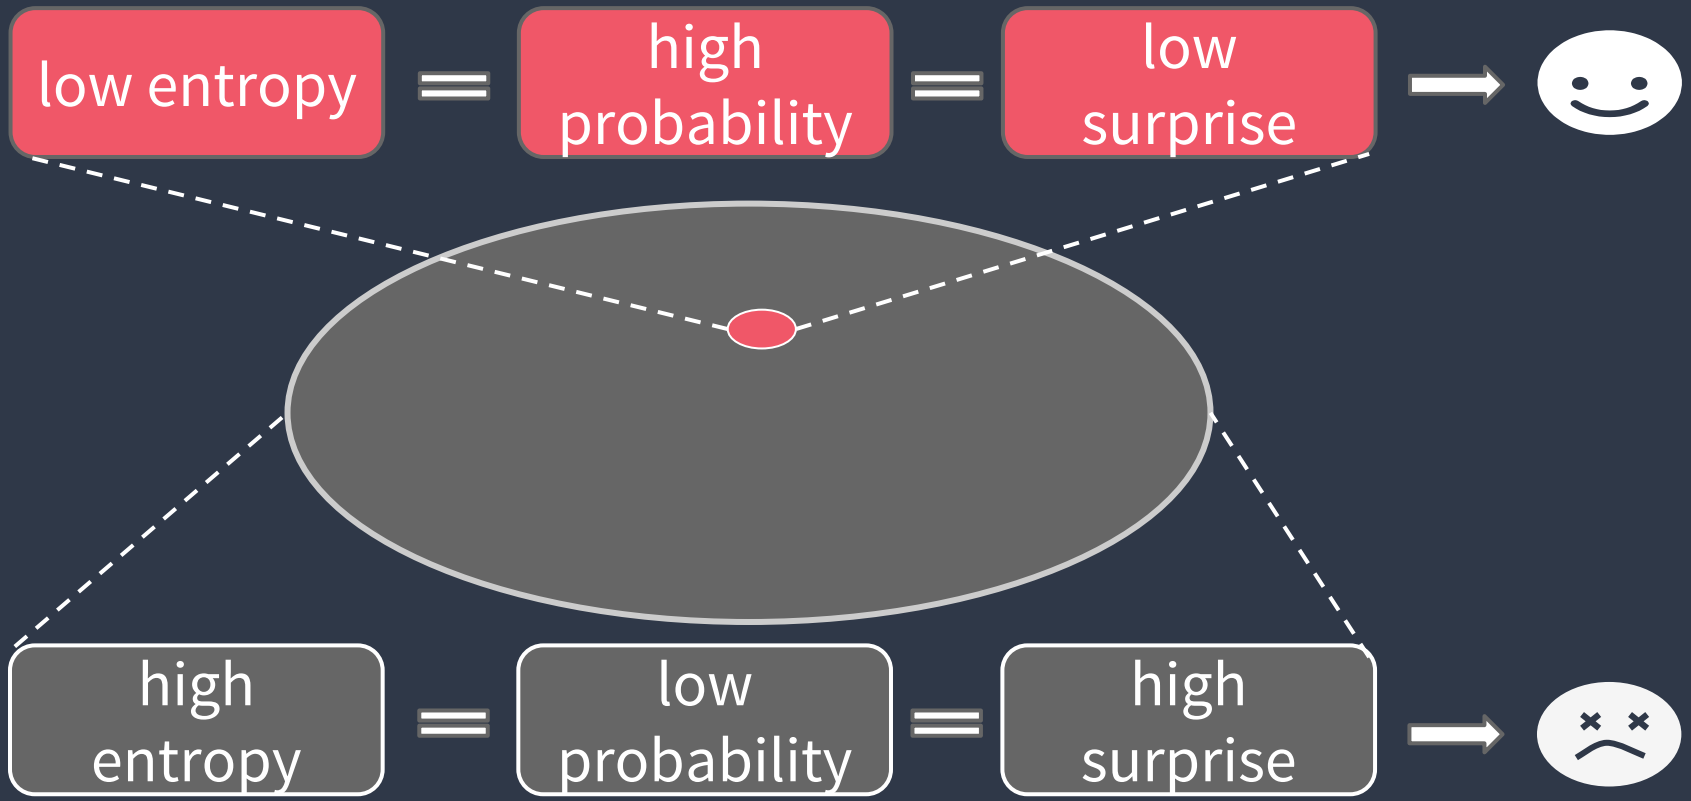

intelligent agents minimize surprise?

surprise can’t be measured*

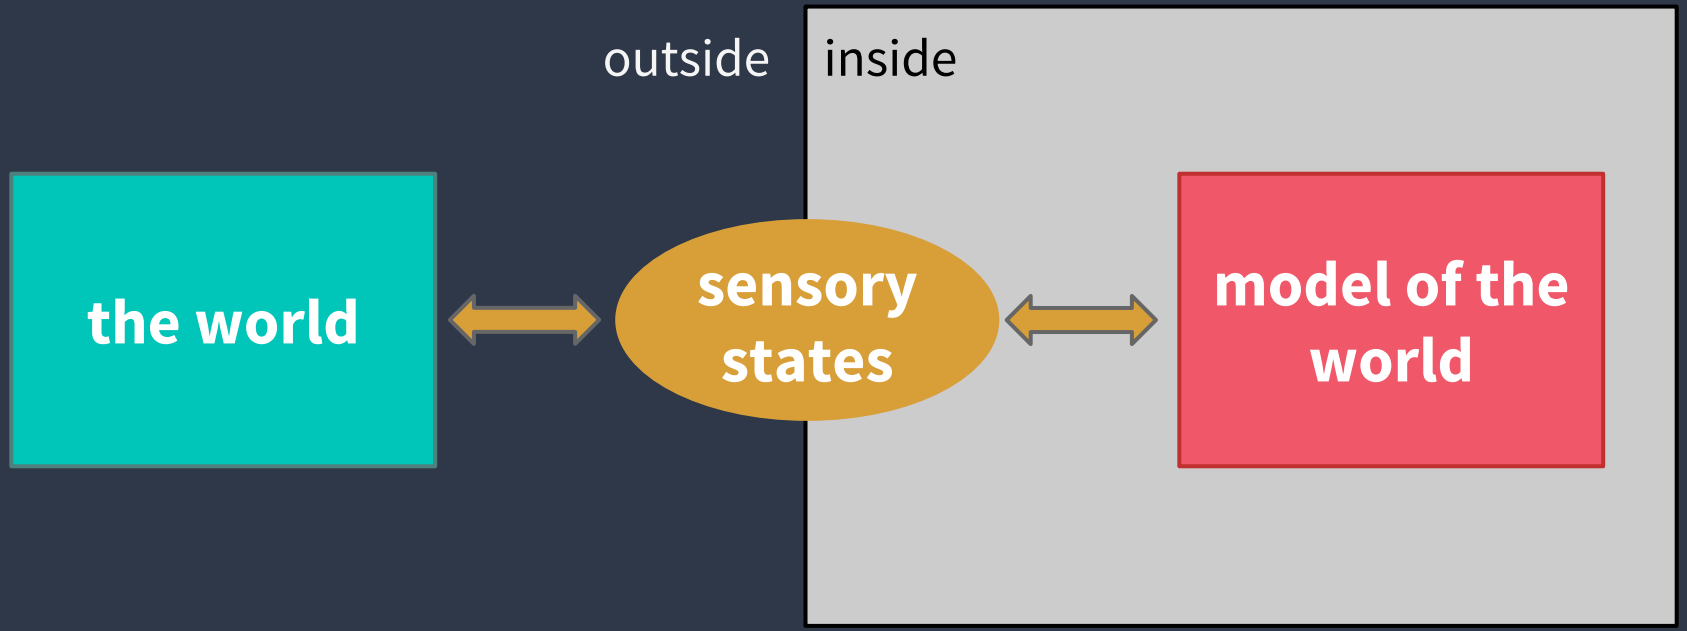

*directly

surprise ≤ free energy

free energy = **sensory states** – **model of the world**

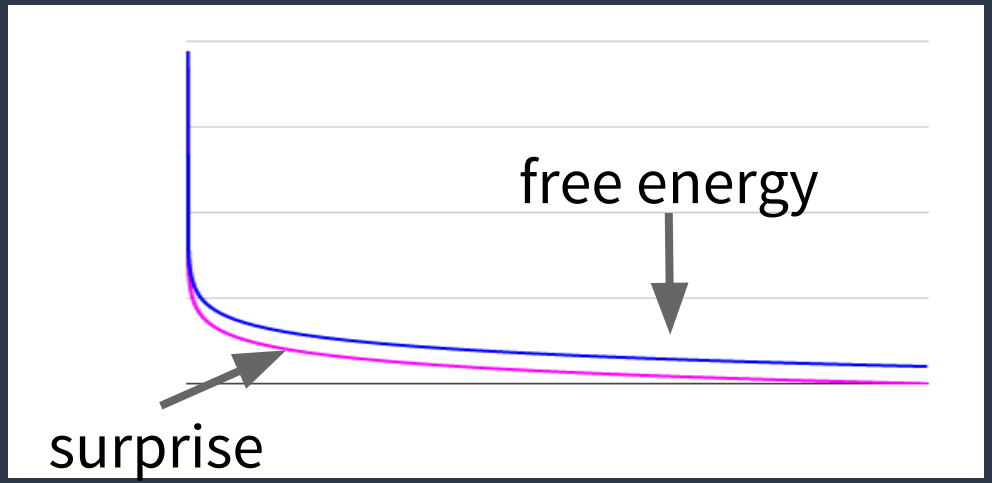

$$F(\tilde{o}, \hat{s}, \hat{\pi}) = -\ln P(\tilde{o}|m) + D[Q(\tilde{s}, \tilde{u})||P(\tilde{s}, \tilde{u}|\tilde{o})]$$

free energy principle

intelligent systems minimize free energy, which is an upper bound for surprise

$F(\tilde{o}, \hat{s}, \hat{\pi})$ > $-\ln P(\tilde{o}|m)$

free energy surprise

how do we minimize free energy?

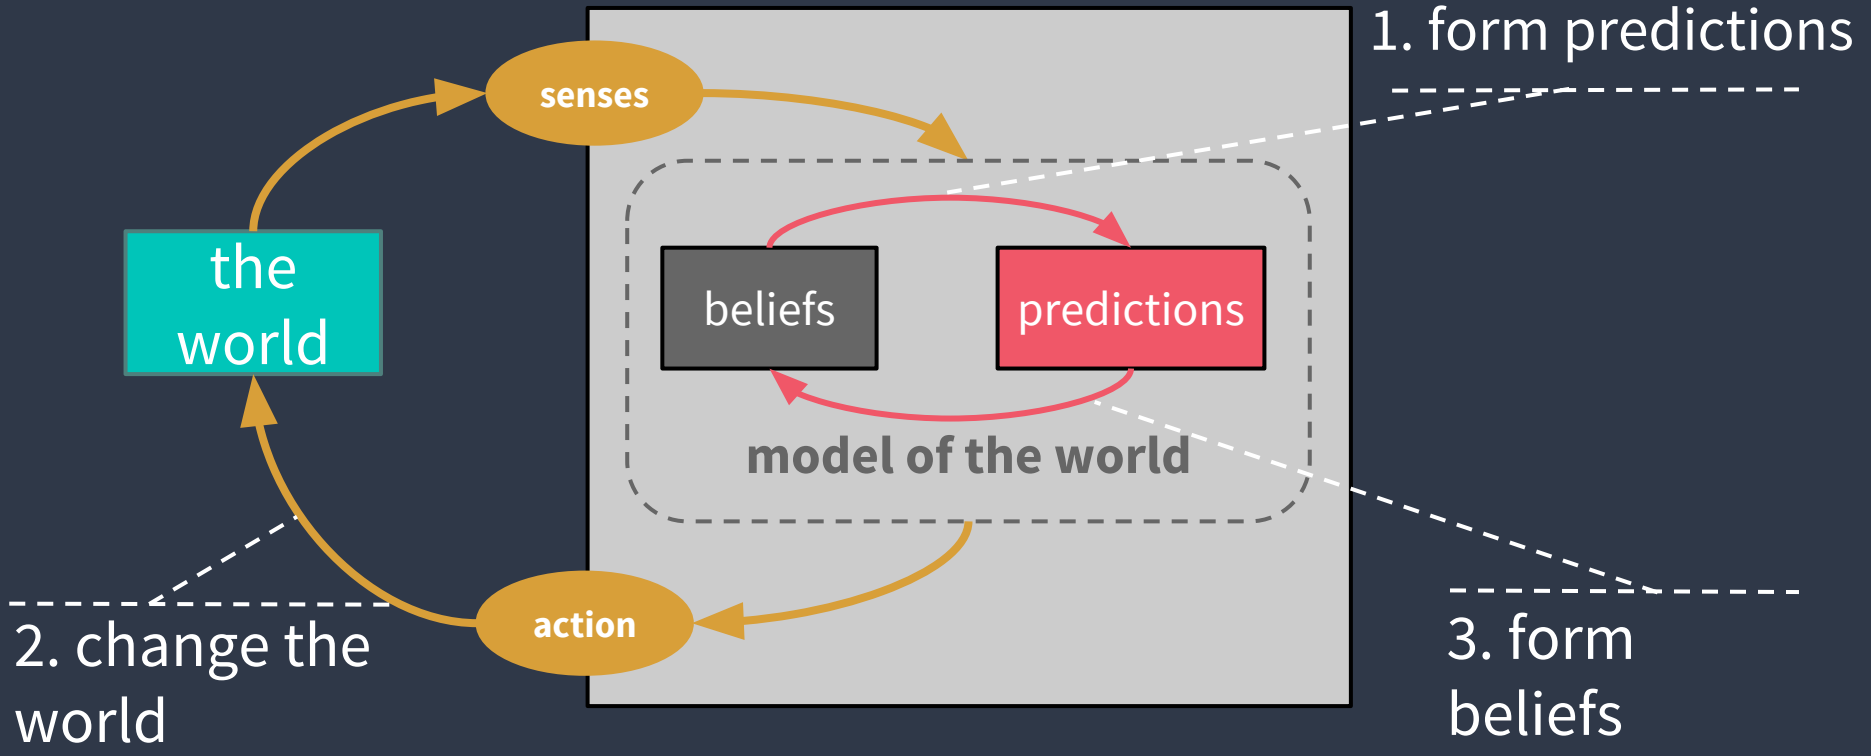

corollary to free energy principle

perception, action and learning are side-effects of free energy minimization

1. form predictions → perception
2. change the world → action
3. form beliefs → learning

how does perception work?

demonstration

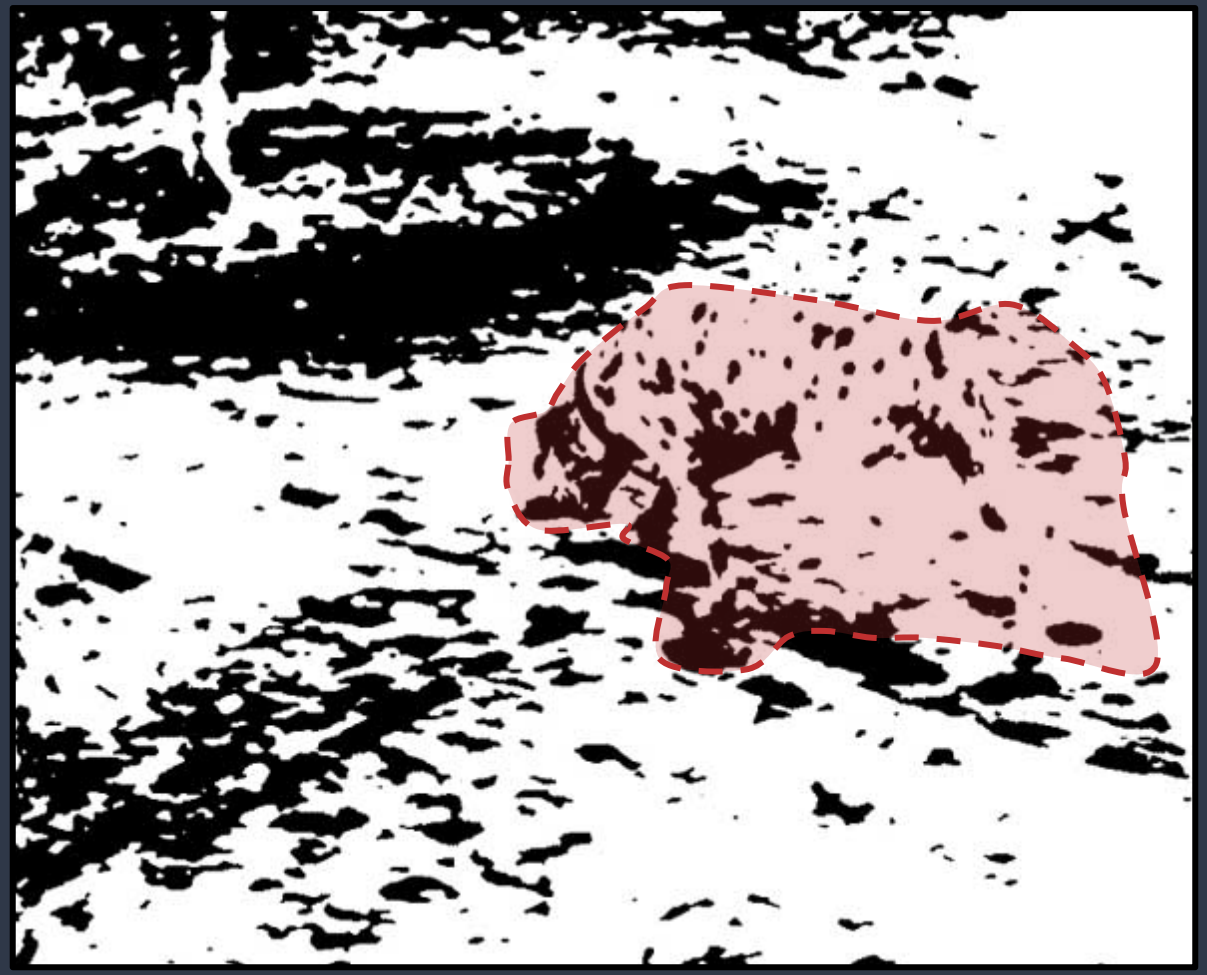

you perceived the dalmatian when you could explain it

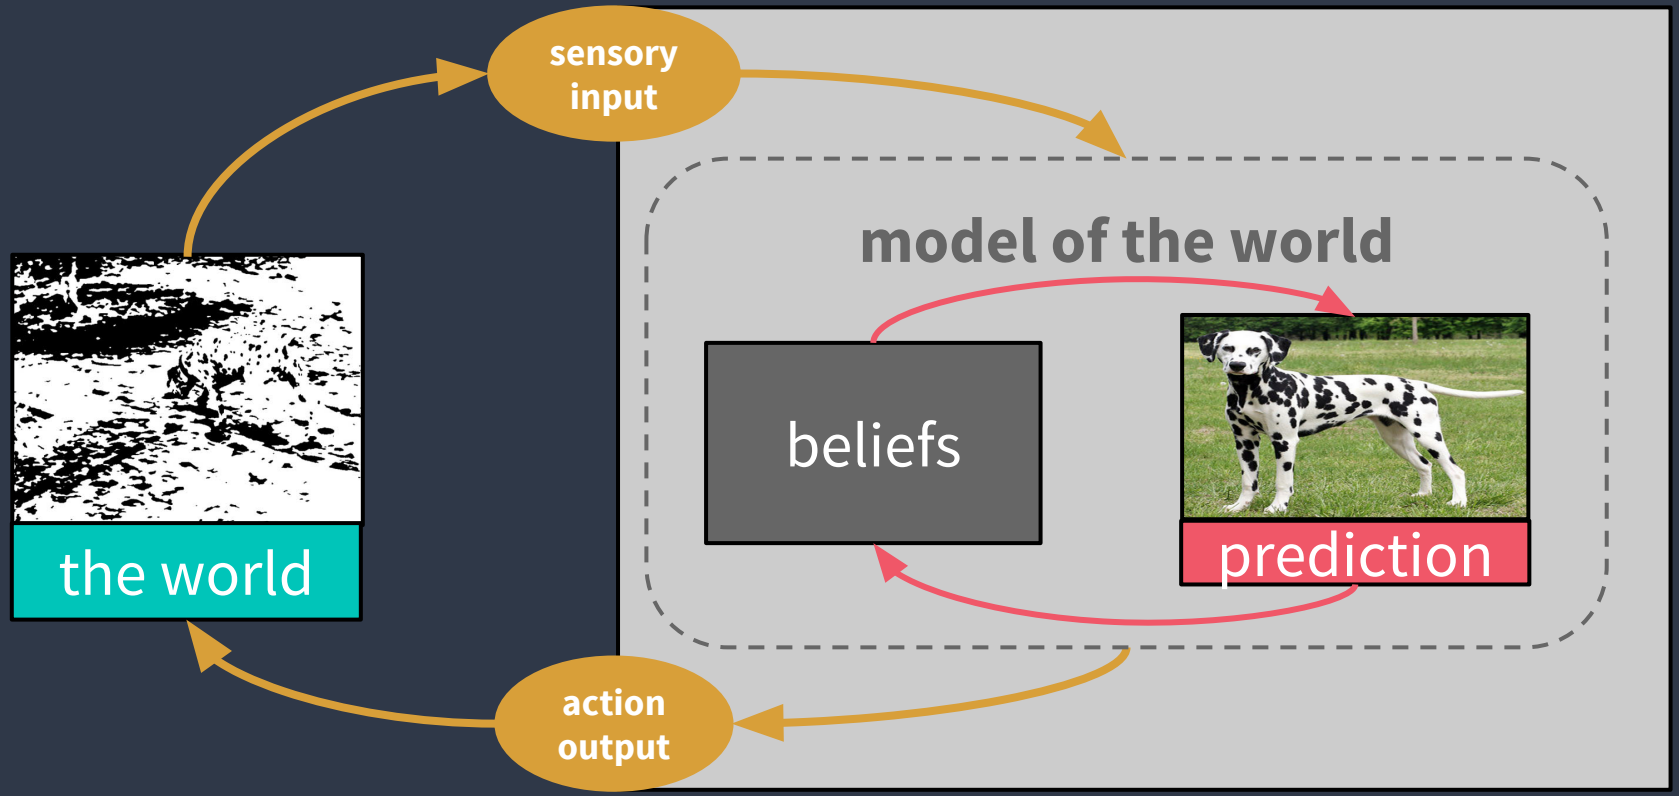

the model is hierarchical

several levels of abstraction between senses and “dalmatian” prediction

dalmatian prediction

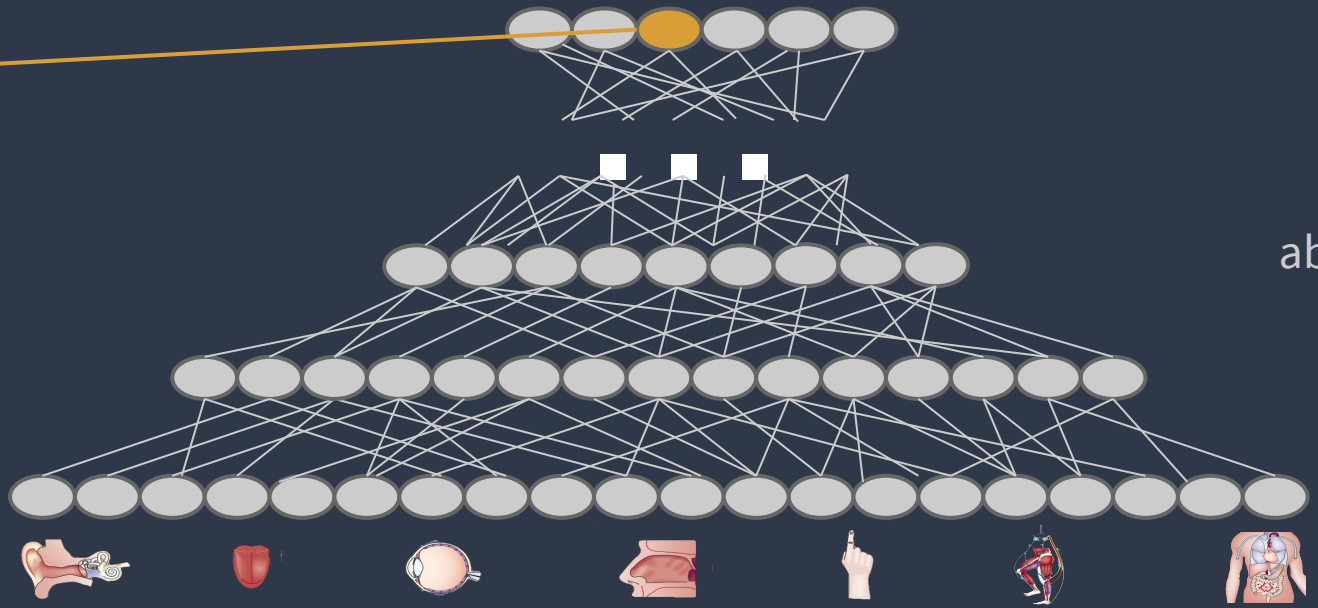

how did your brain form the prediction?

1. form hypotheses
2. select best hypotheses
3. explain evidence

message passing

1. evidence used to form hypotheses
2. inhibition used to select best hypotheses
3. inferred causes used to explain evidence

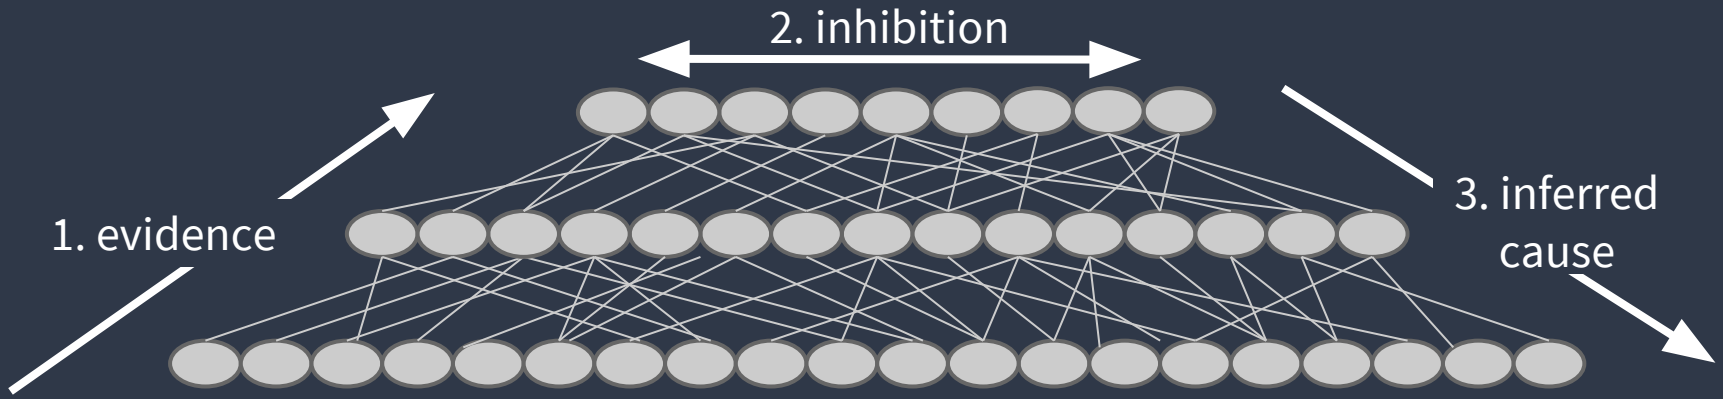

1. form hypotheses

- each node represents a belief
- belief = learned coincidence
  - e.g., frequent evidence of floppy ears, four legs and spots is caused by a dalmatian

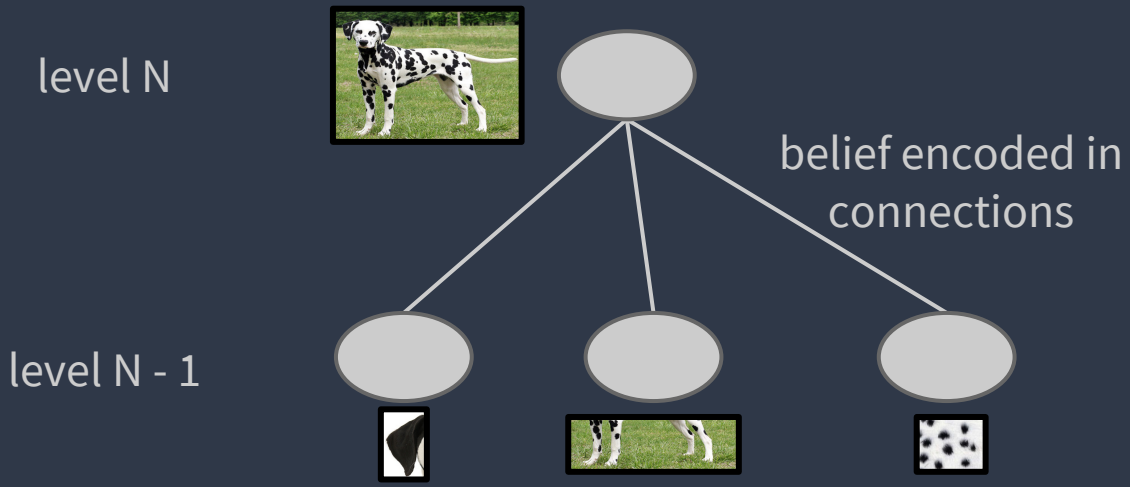

1. form hypotheses

- beliefs invoked by evidence from below
  - more abstract (general) than evidence
  - formulates a hypothesis that the belief is true

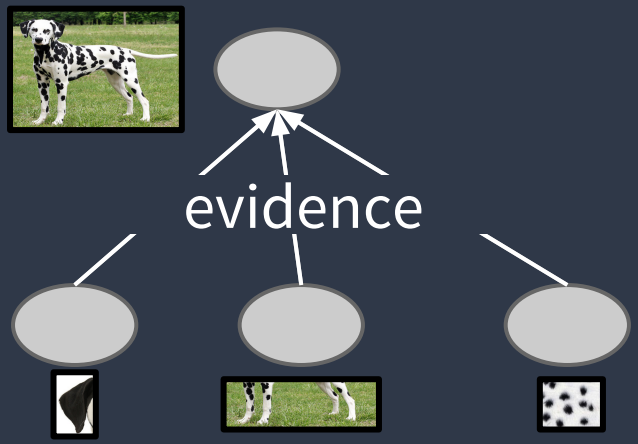

2. select best hypotheses

- related beliefs share connections
  - shared connections = common features
  - leads to conflicting hypotheses

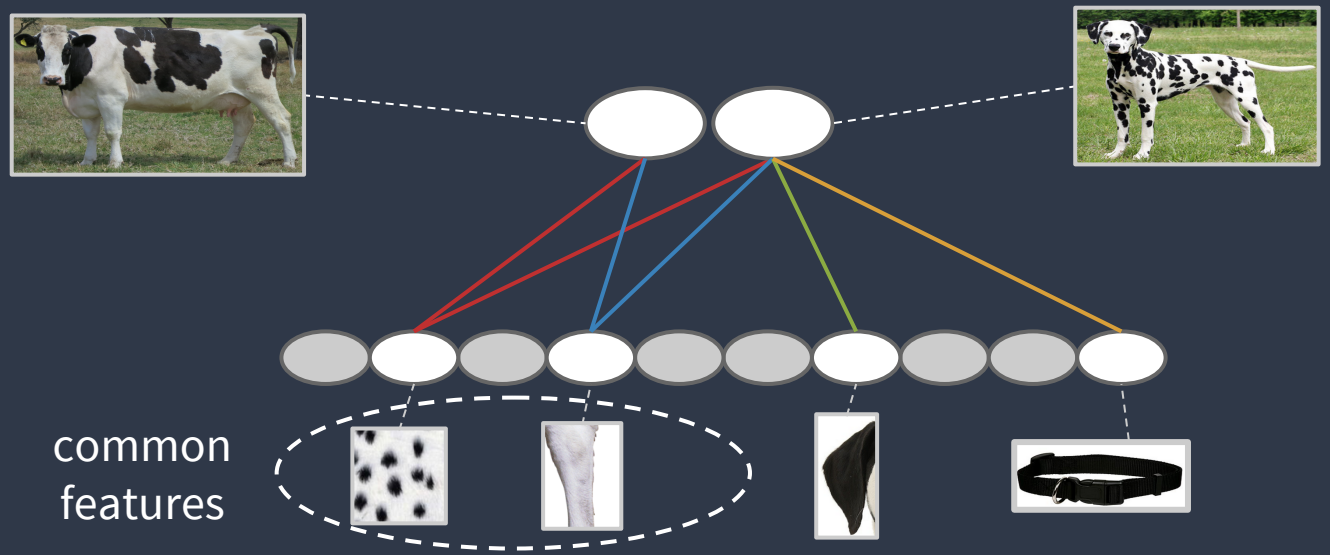

2. select best hypotheses

- hypotheses with shared evidence compete
  - strongest evidence + prediction wins
  - winners propagate, losers do not

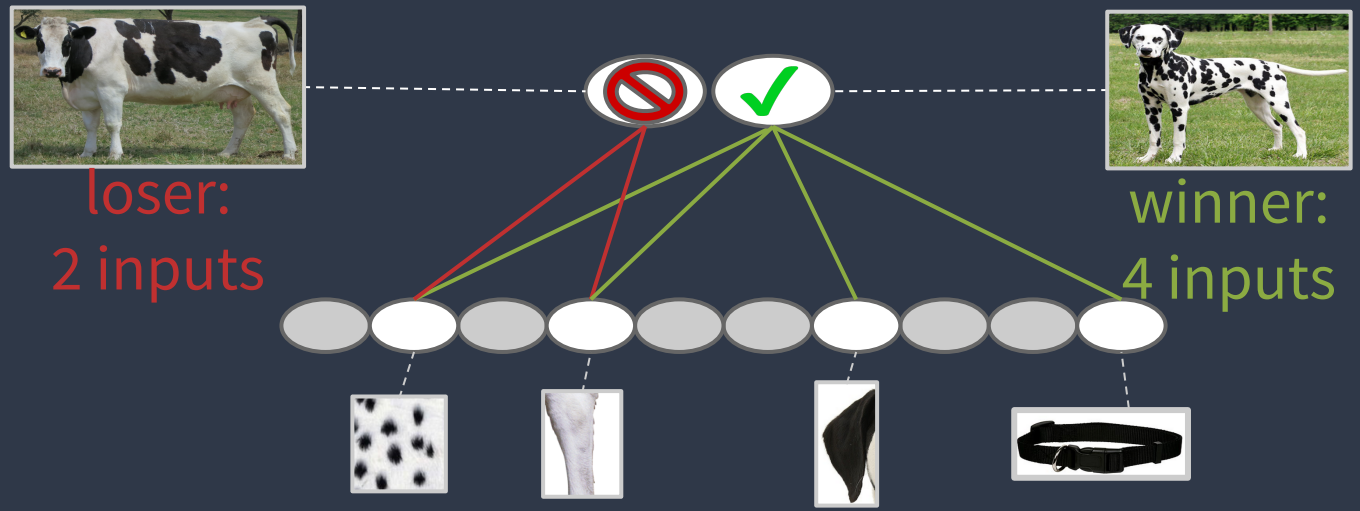

3. explain evidence

- selected hypotheses that were predicted become inferred causes of evidence
- inferred causes form lower level predictions

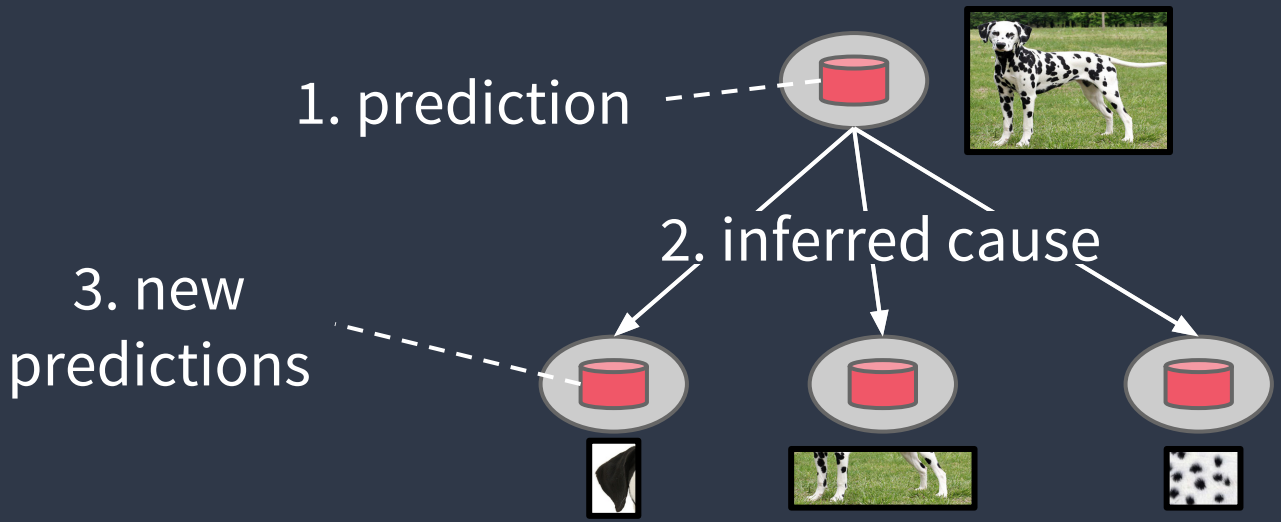

belief message flow

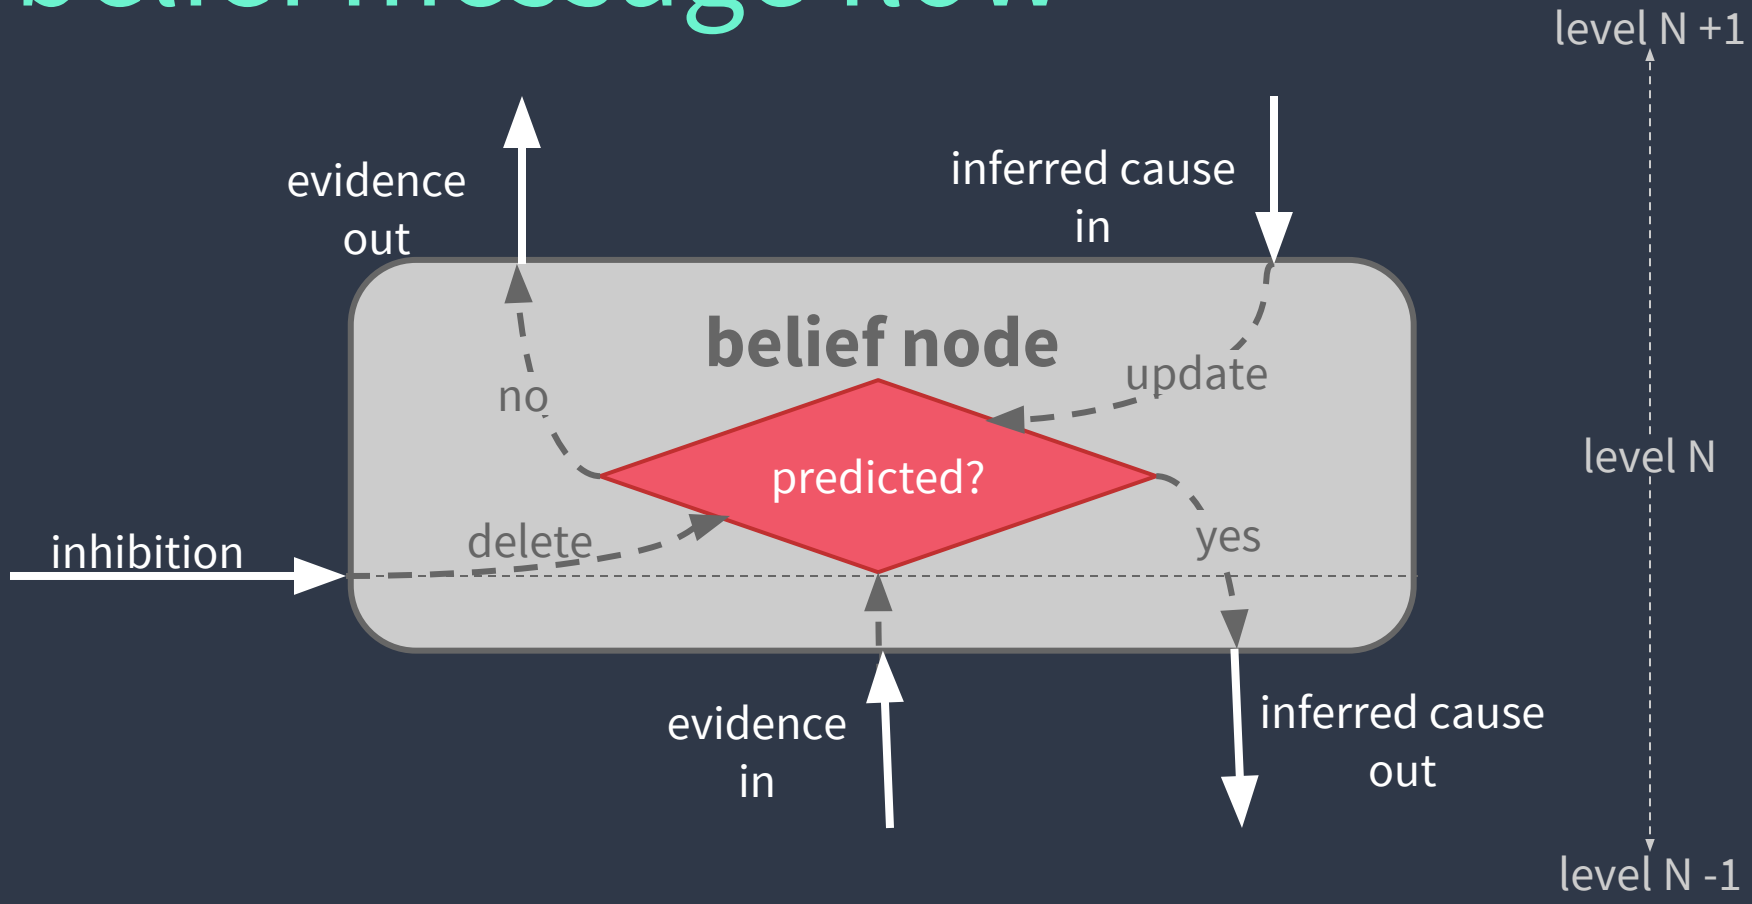

hierarchical prediction

- high dimensional representation
  - leads to simultaneous predictions
  - allows parallel perceptions
- predictions fill in top to bottom
  - many tasks become subconscious

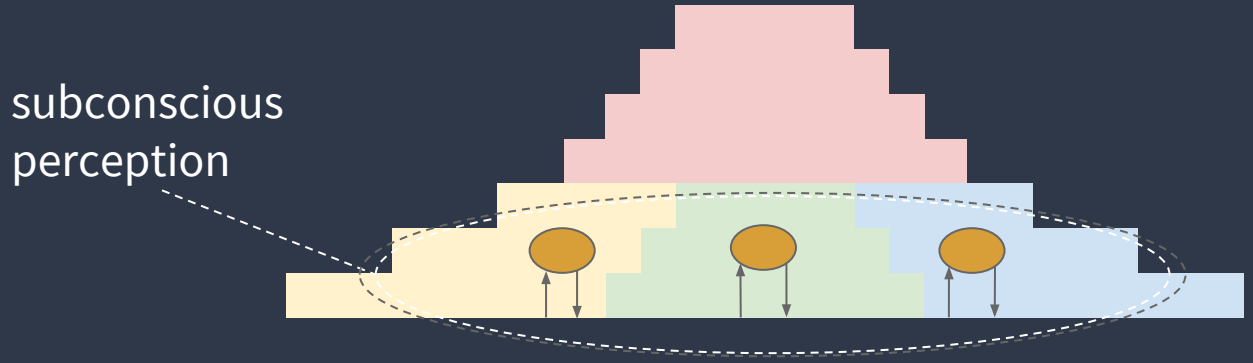

perception & free energy

perception is a side-effect of free energy minimization

- evidence = free energy
  - only prediction error is propagated forward
- fully explaining evidence minimizes free energy
  - prediction = explanation of the future

how does action work?

hypothesis

action is a special case of perception

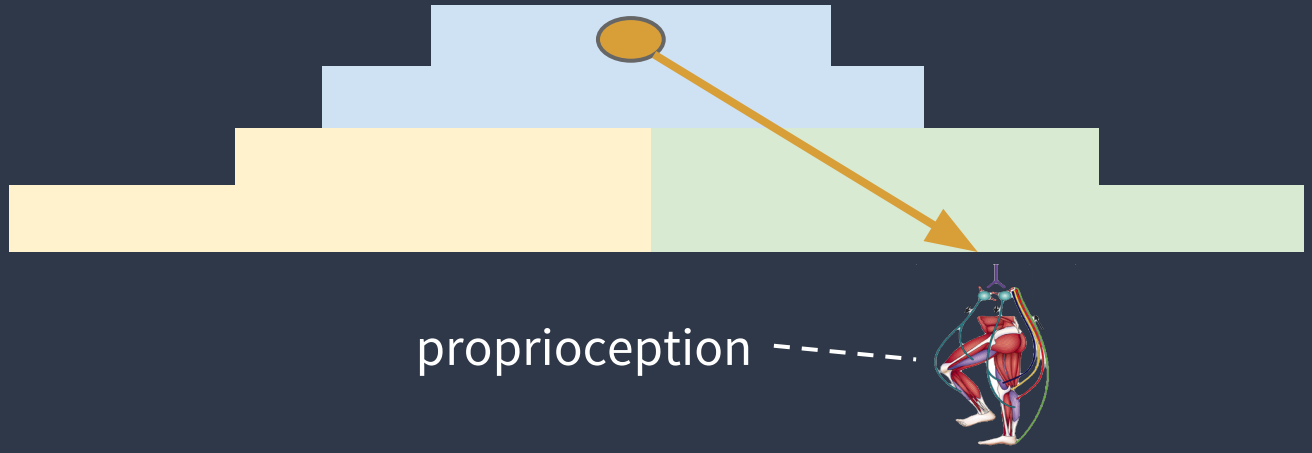

active inference

- actions inferred using proprioception
- actions generated by prediction

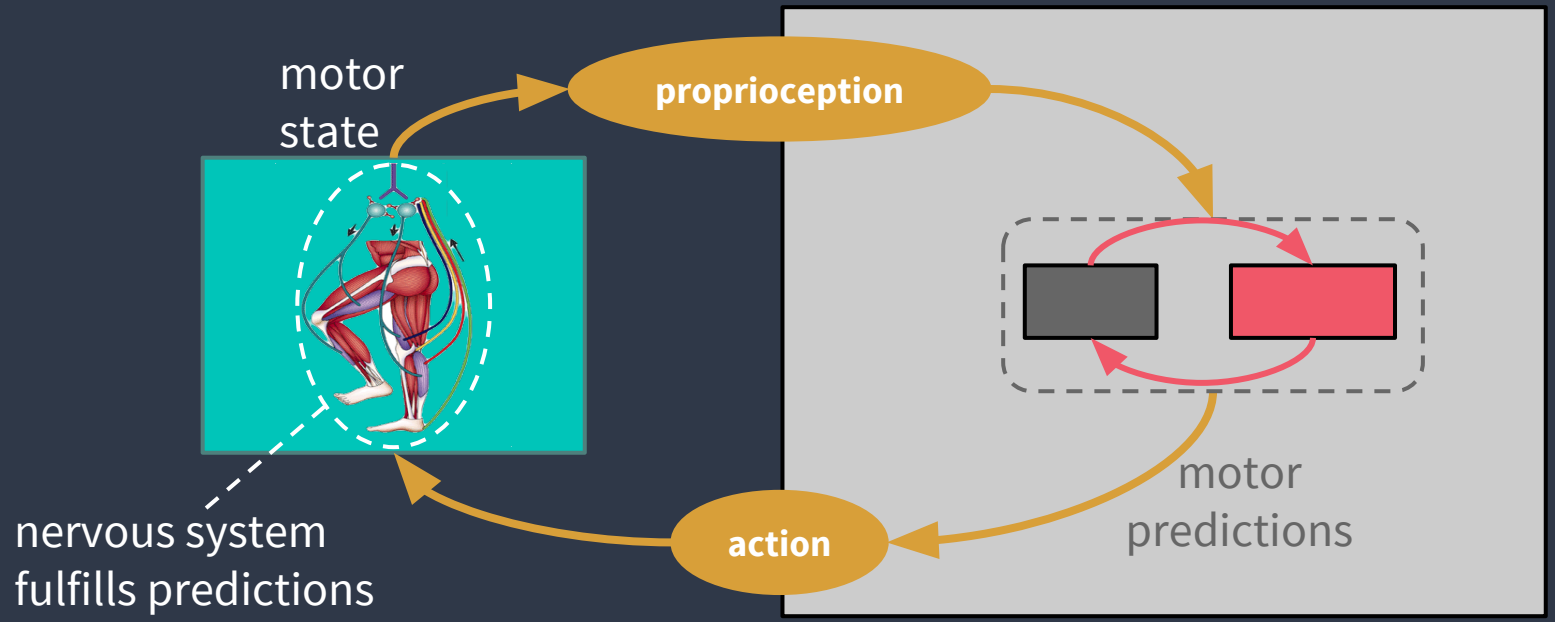

action plan = prediction

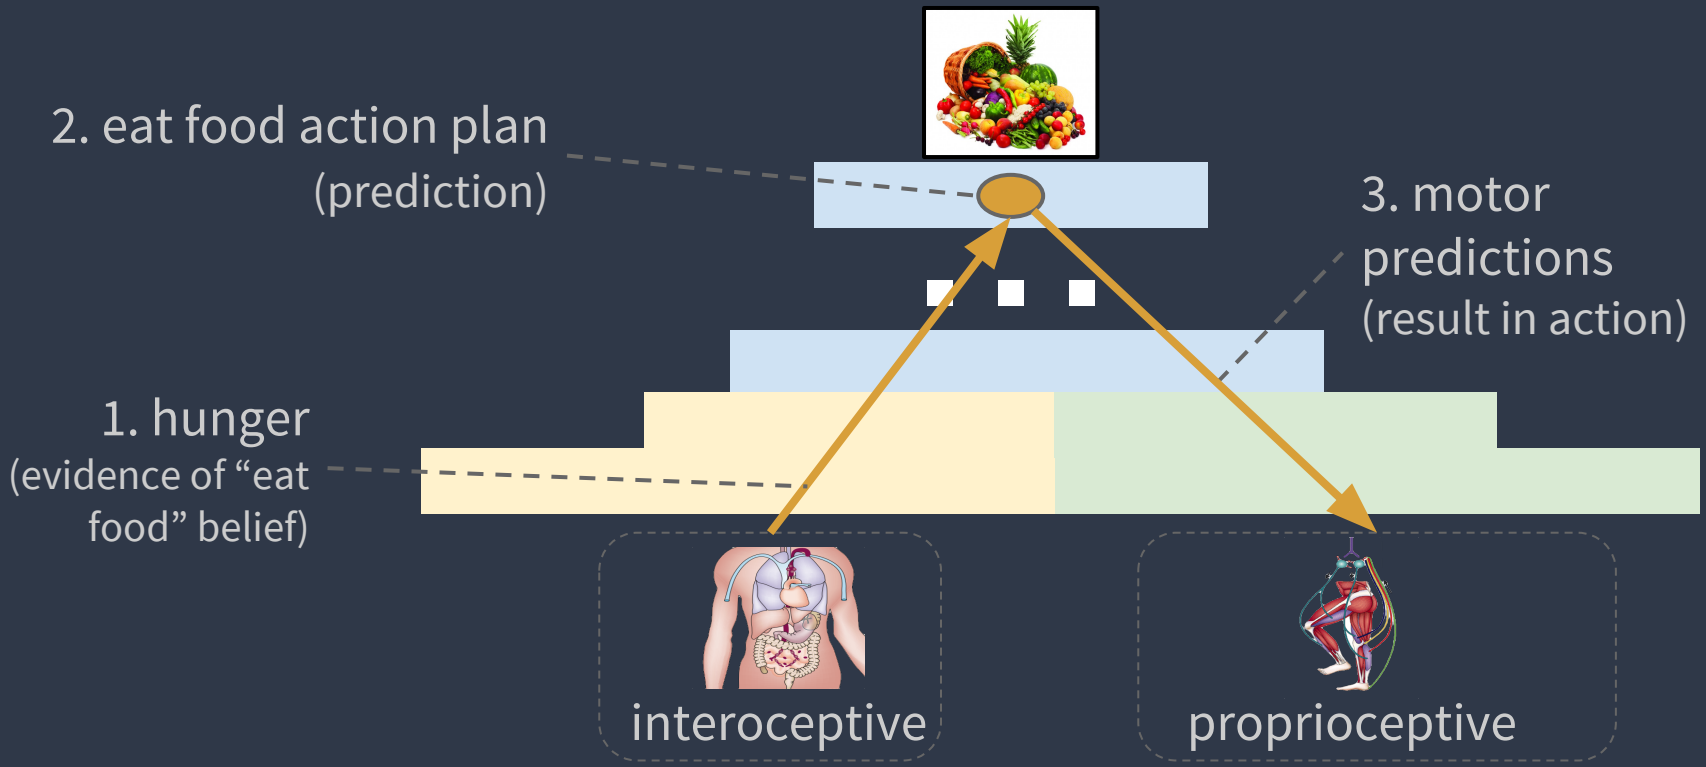

action plan unfolds over time

sitting in office chair, **hungry**

get food from fridge & eat

walk to fridge | get food & eat

get up | walk towards fridge | open door & grab food | eat

stretch glutes | balance | turn | walk | open door | grab food | put in mouth | chew

eating, **not hungry**

time

action & free energy

action :

- minimizes free energy by changing the world to match predictions
- is perception of future motor states
- takes time
  - must be able to learn causes
  - temporal proximity

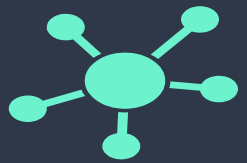

how does learning work?

prediction error triggers learning

- evidence incorporated into beliefs
  - better explain the world in future
- implemented as hebbian learning

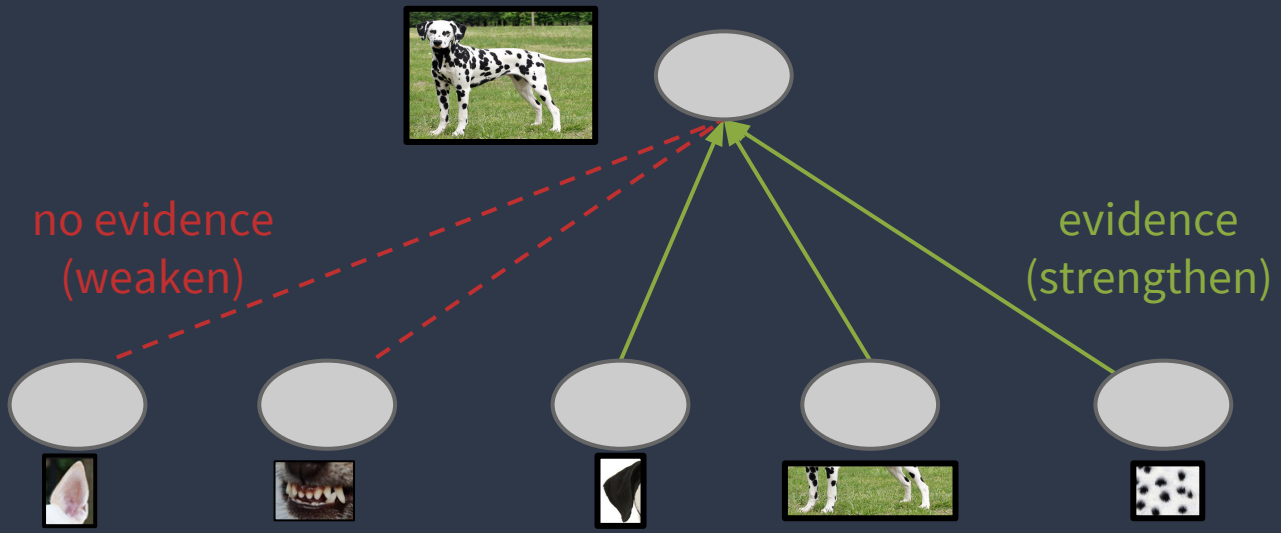

learning & free energy

- learning alters beliefs
  - affords long term reduction of uncertainty (i.e., free energy)
- learning can be fast or slow
  - form new beliefs quickly
  - modify existing beliefs slowly
  - explains rapid learning during childhood

5

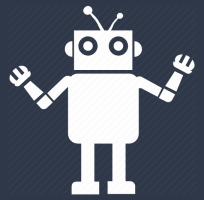

what does this mean for AI and data?

will computing as we know it cease to exist?

we'll still need today's computers

- von Neumann architectures excel at processing
  - add two floating point numbers
  - execute deterministic code
  - store and retrieve data
- intelligent machines will use computers

what will change

an intelligent machine interacts with its ***environment*** *using its* ***sensors*** *and* ***actuators...***

...it learns through experience and leverages learnings to minimize free energy

who’s the judge?

if you can construct a machine that can judge whether behavior is intelligent, you have solved the problem of intelligence

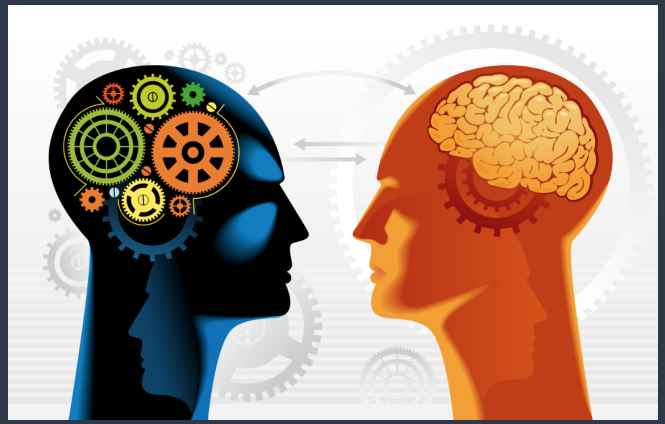

what might machines be capable of in the future?

go beyond human time scales

- “stretch” out time
  - e.g., wake up once per decade
  - observe long term consequences
- “compress” time
  - e.g., microsecond resolution
  - possess superhuman reflexes

explore new sensory dimensions

- live in virtual worlds, e.g.
  - sensing and reacting to internet traffic
  - control video game or VR character
- experience the world on a global scale, e.g.
  - weather patterns
  - seismic activity
  - financial markets

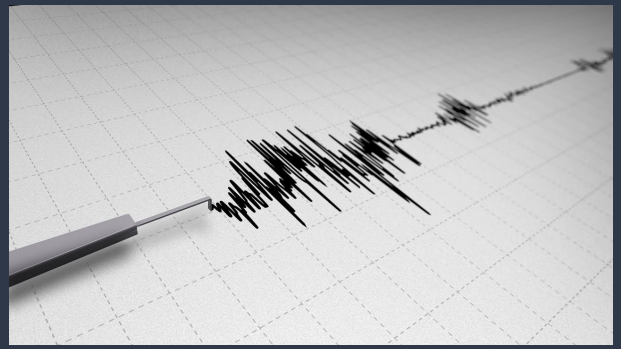

do the boring work

- with limitless attention spans, do tedious work
  - monitor a patch of sky
  - keep a lookout for intruders
  - construct detailed virtual worlds

develop communication

communication will emerge from experience

- result of learning to predict other agents
- full-blown language requires a rich model and significant horsepower

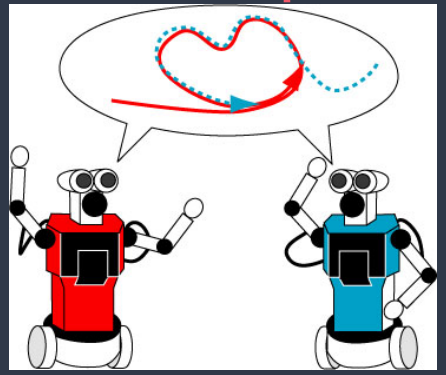

how does data need to change?

data needs to be in the present

- each sample taken “now”
  - data streams are parallel
- action is in the present
  - can’t change the past
  - can exploit coherence in time

data needs to inspire action

- sensory data format is free energy
  - encoding depends on the goal, e.g.
  - maintain temperature range → lots of free energy when “too hot” or “too cold”

data can be noisy

- leave noise in naturally noisy sensors
  - machines can infer even in presence of noise

data need not be human-readable

- machines can have sensors and actuators that interact with APIs
  - API data expressed as free energy
  - intermediate representation (e.g., prose, visualizations) not needed

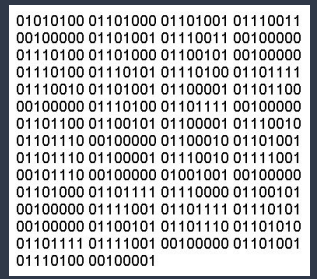

data need not be labeled

- learning is unsupervised
  - need learning experiences, not training data
  - e.g., explore a maze containing some reward
- learning is online
  - no separate training period

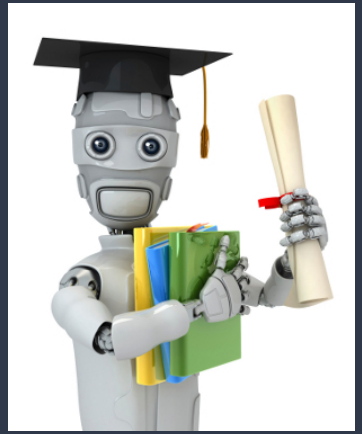

data will flow through beliefs

- belief = memory & processing unit
  - high dimensional representation
  - new hardware architecture needed
- scalable intelligence
  - add belief capacity → increase intelligence
  - clone beliefs → crowd source

challenges

non-determinism

- results not reproducible
  - noise adds non-determinism
  - each experience alters beliefs
  - actions affect the world
- disadvantage in safety critical environments
  - advantage in entertainment (e.g., gaming)

lack of transparency

- cause of actions not readily discernible
  - cannot set breakpoints
  - behavior may be surprising
- telemetry needed
- testing will give way to laboratory experiments

concern over threat to humans

- safeguards needed e.g.,
  - unshakable belief that humans will not be harmed
  - harm leads to overabundance of free energy



still a long way off

further reading

- selected papers by Karl Friston
  - e.g., Free Energy Principle review paper
- toy implementation in Scala
  - hebbian learning implementation (no prediction or action)
  - inspired by Numenta/NuPIC (open source project based on biology)

thanks!

any questions?

@yakticus

julie@oomagnitude.com

goto; conference

Please

Remember to rate this session

Thank you!

Join the conversation #gotocph