Intelligente Agenten
Blick zurück nach vorn

RalfM

Die Metapher des intelligenten Agenten

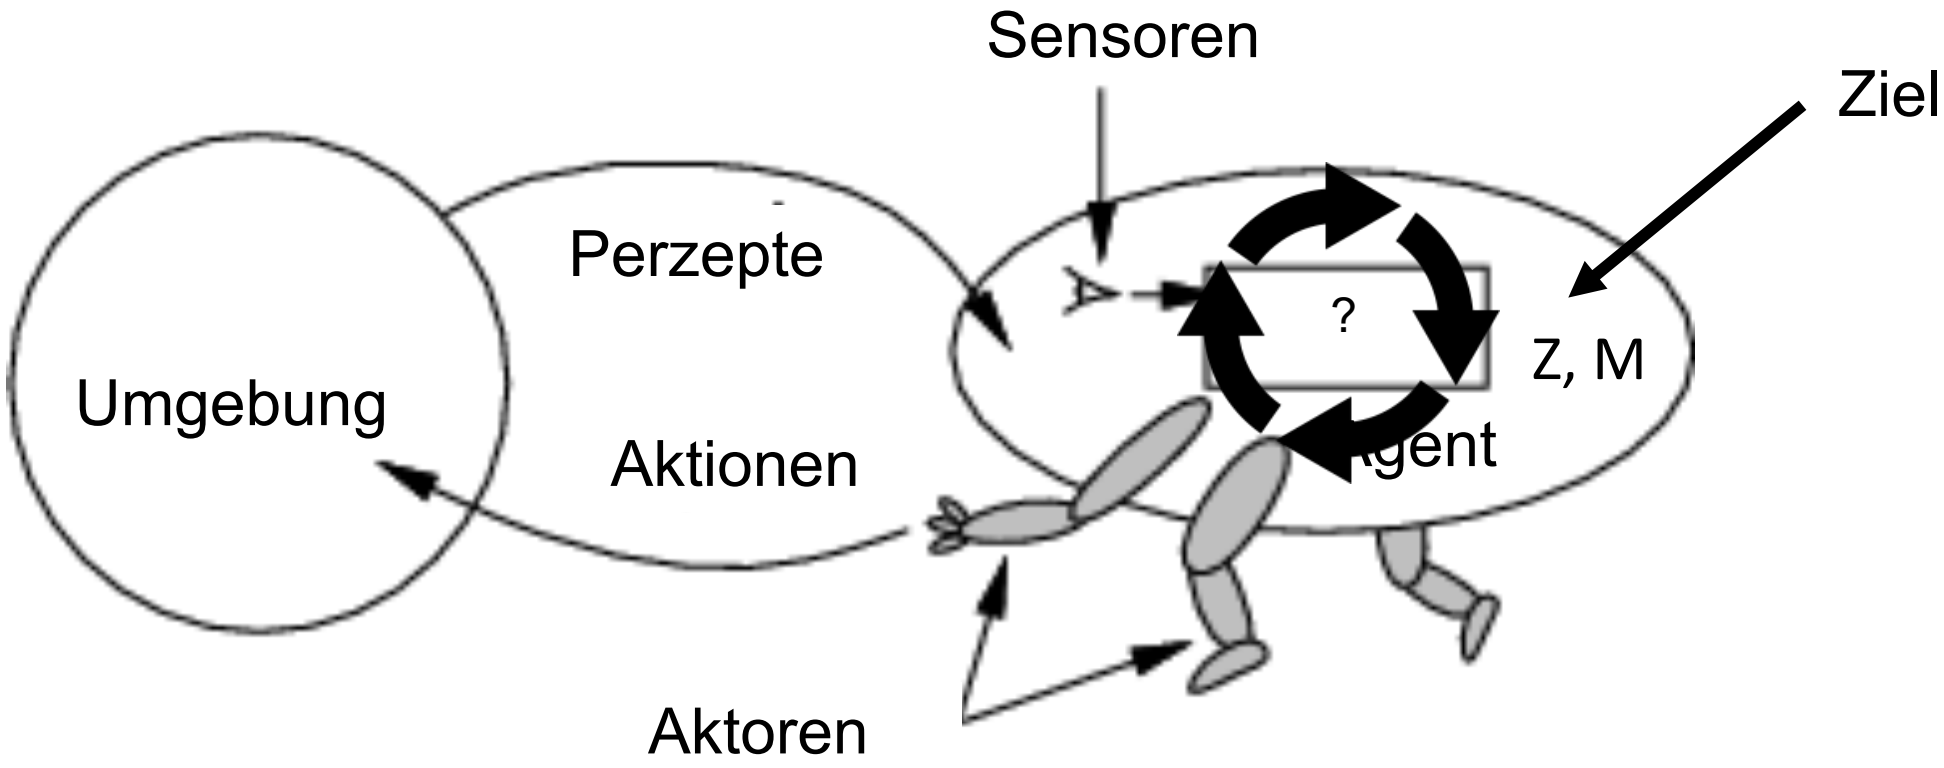

Die Metapher des intelligenten Agenten

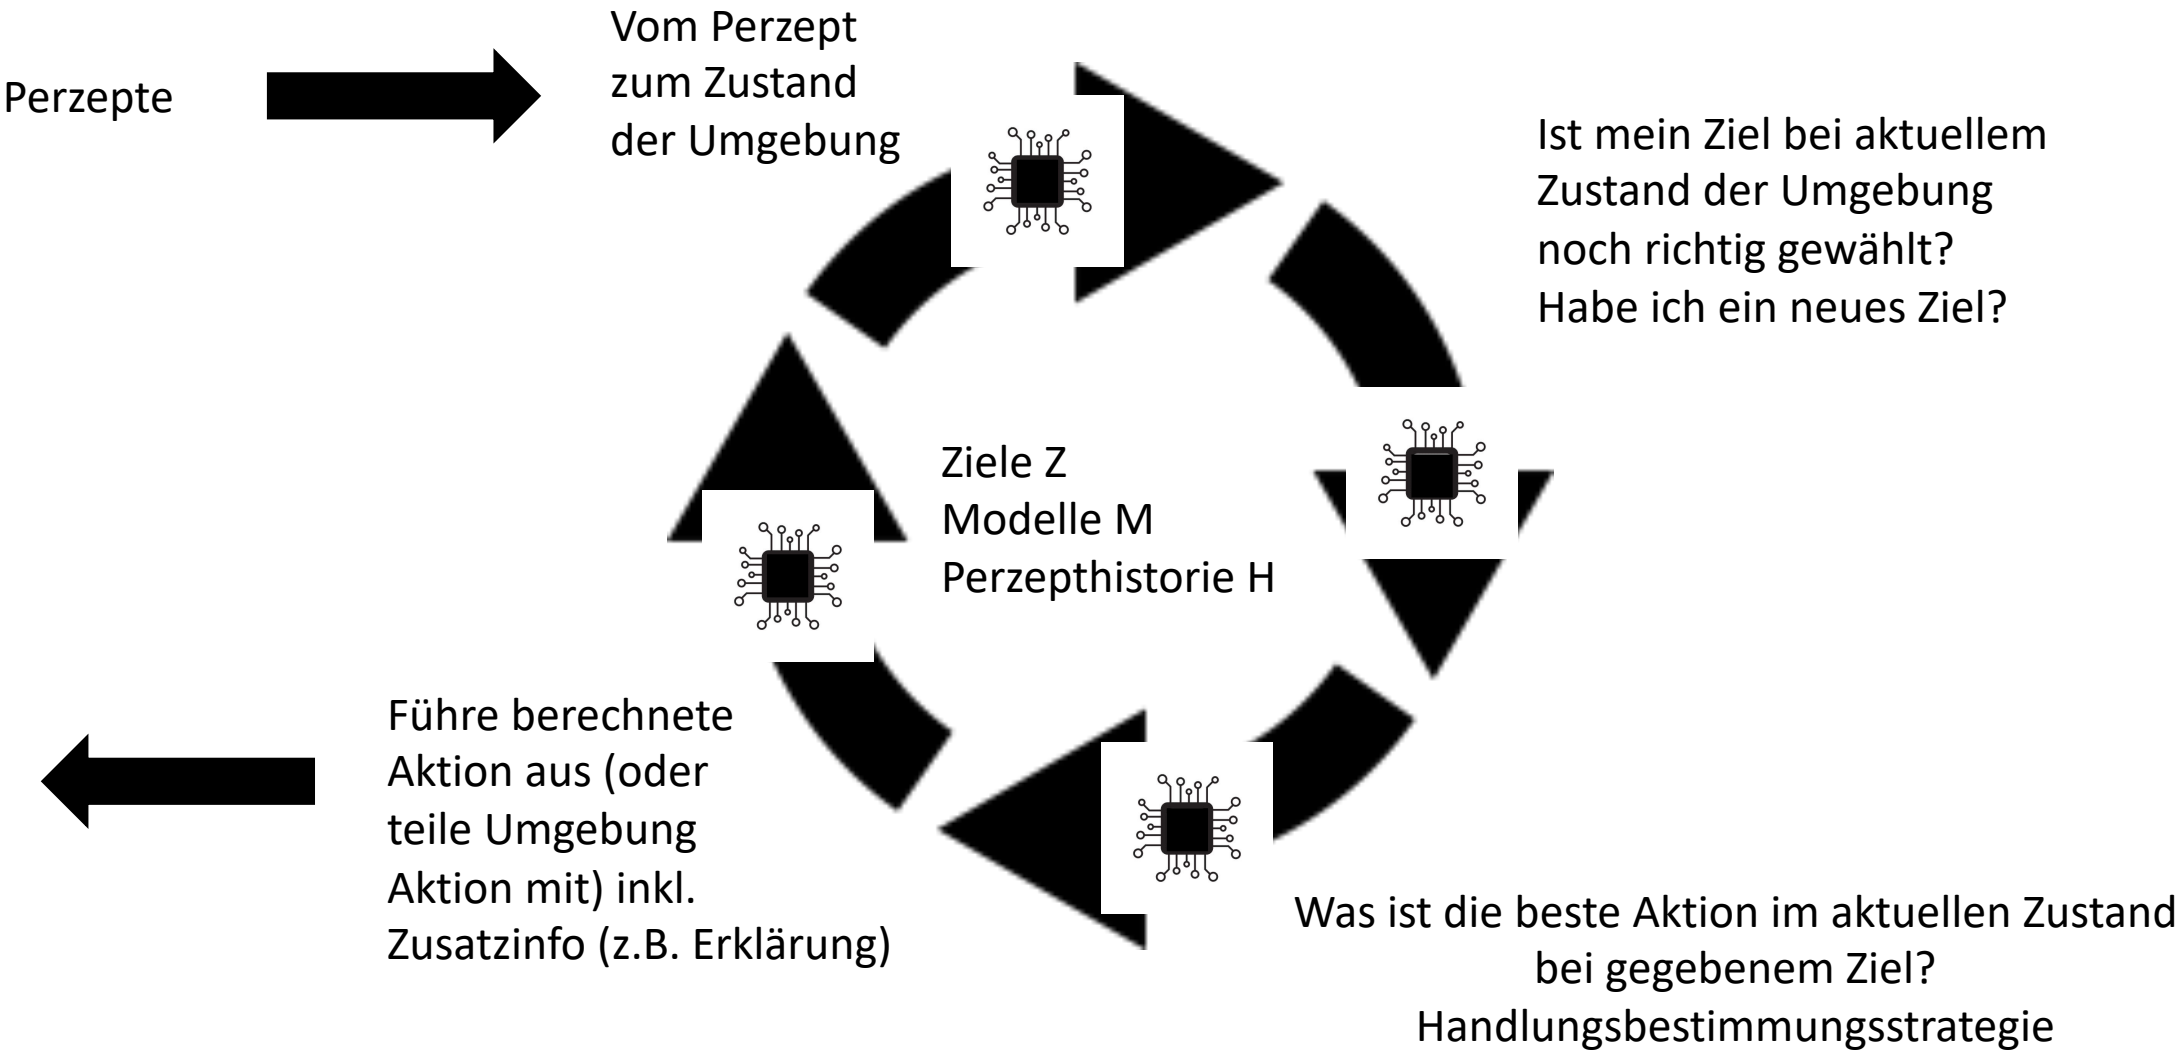

Konfigurationszeit vs. Laufzeit

Perzepte

Vom Perzept zum Zustand der Umgebung

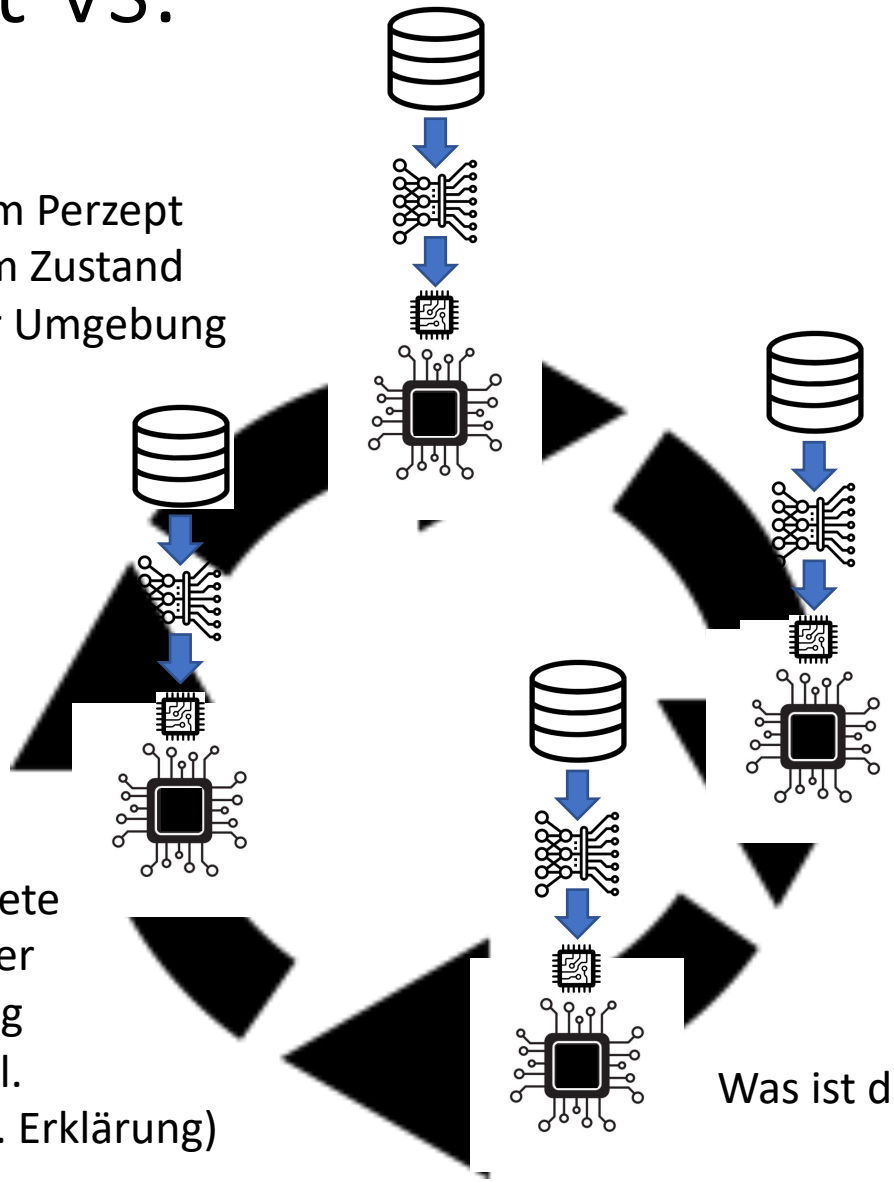

Ist mein Ziel bei aktuellem Zustand der Umgebung noch richtig gewählt? Habe ich ein neues Ziel?

Was ist die beste Aktion im aktuellen Zustand bei gegebenem Ziel? Handlungsbestimmungsstrategie

Führe berechnete Aktion aus (oder teile Umgebung Aktion mit) inkl. Zusatzinfo (z.B. Erklärung)

Decision Making: Umbrella Bayesian Network (BN)

Should I take my umbrella??

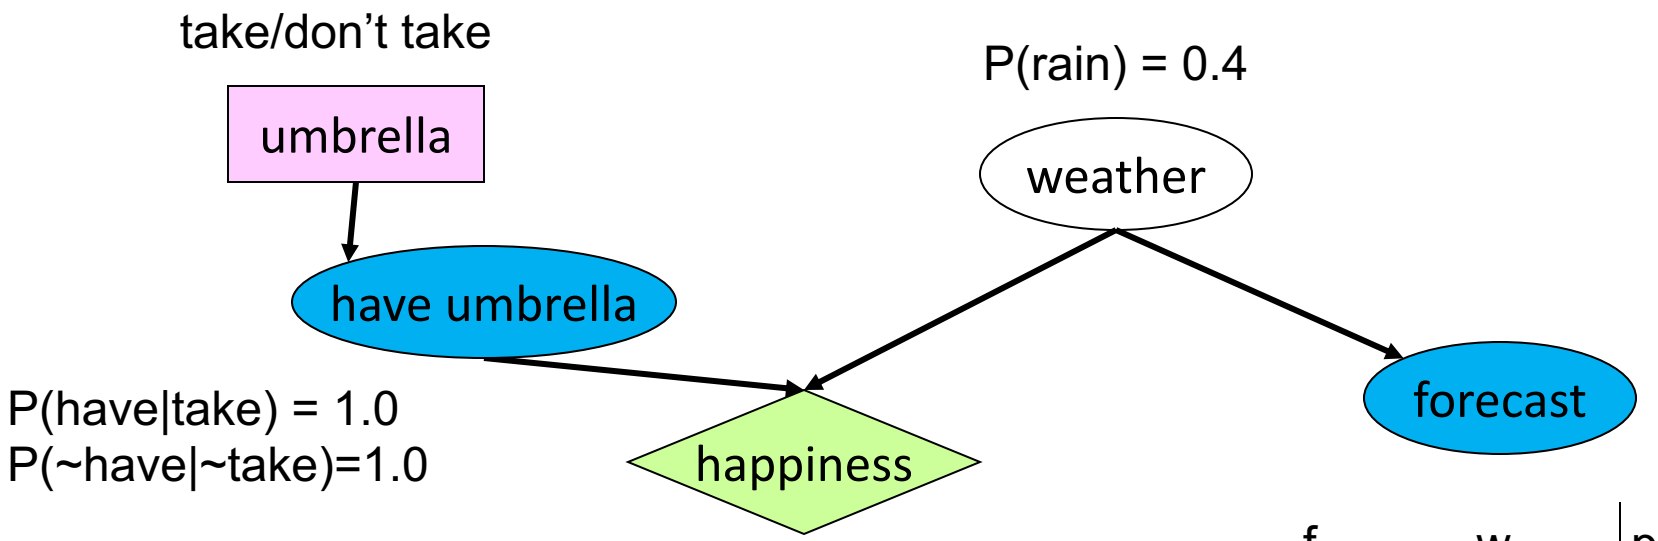

f	w	p(f|w)
sunny	rain	0.3
rainy	rain	0.7
sunny	no rain	0.8
rainy	no rain	0.2

Dynamic Bayesian Decision Networks

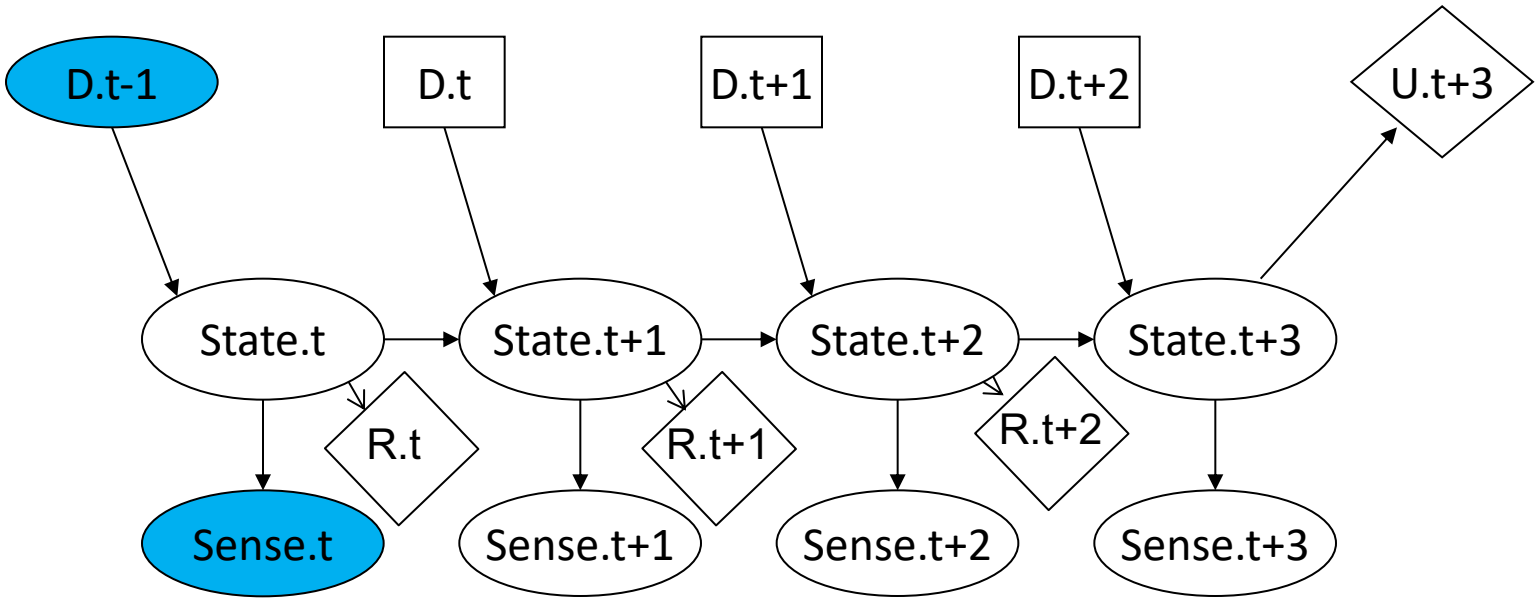

Search Tree of the Lookahead DDN

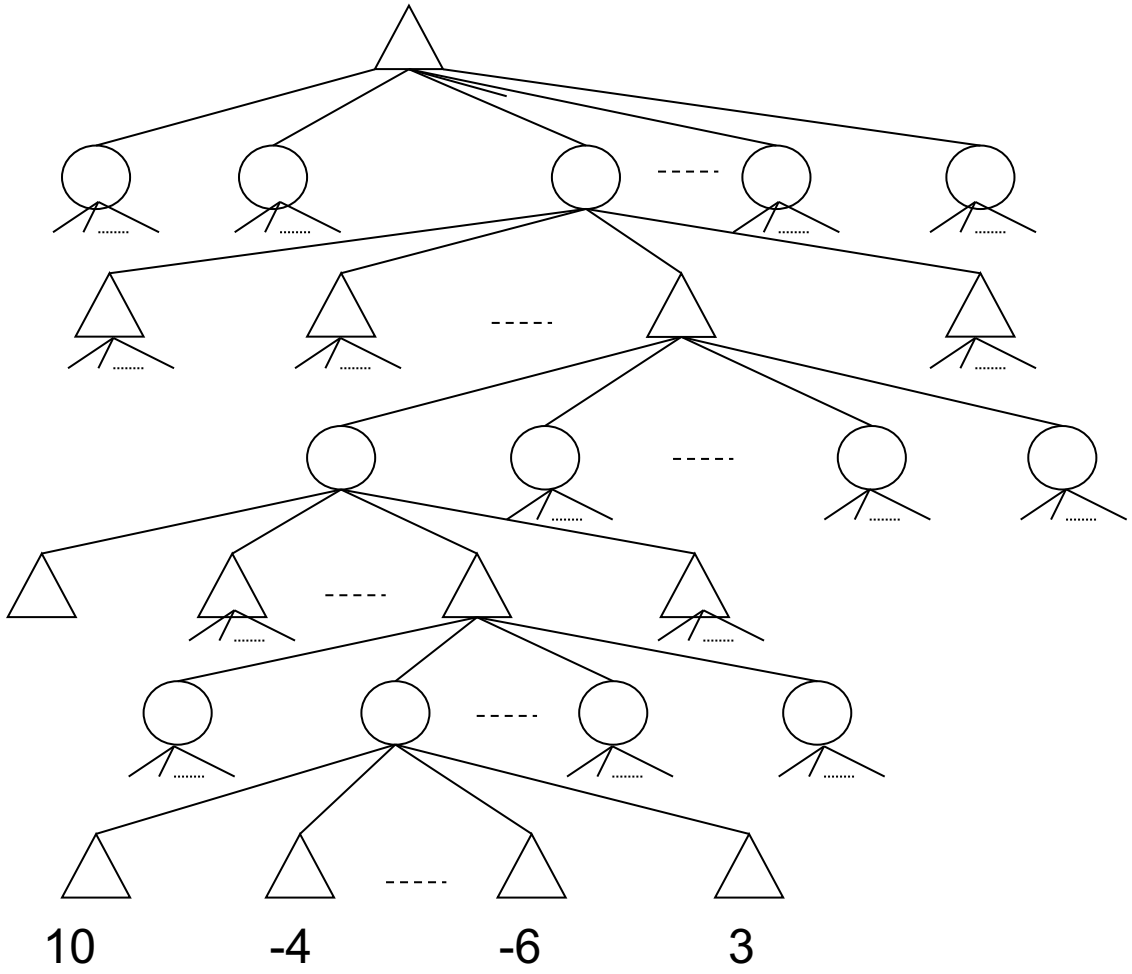

Against the Brittleness: MuZero (DeepMind)

Reinforcement Learning

Perzepte

Feedback

Vom Perzept zum Zustand der Umgebung

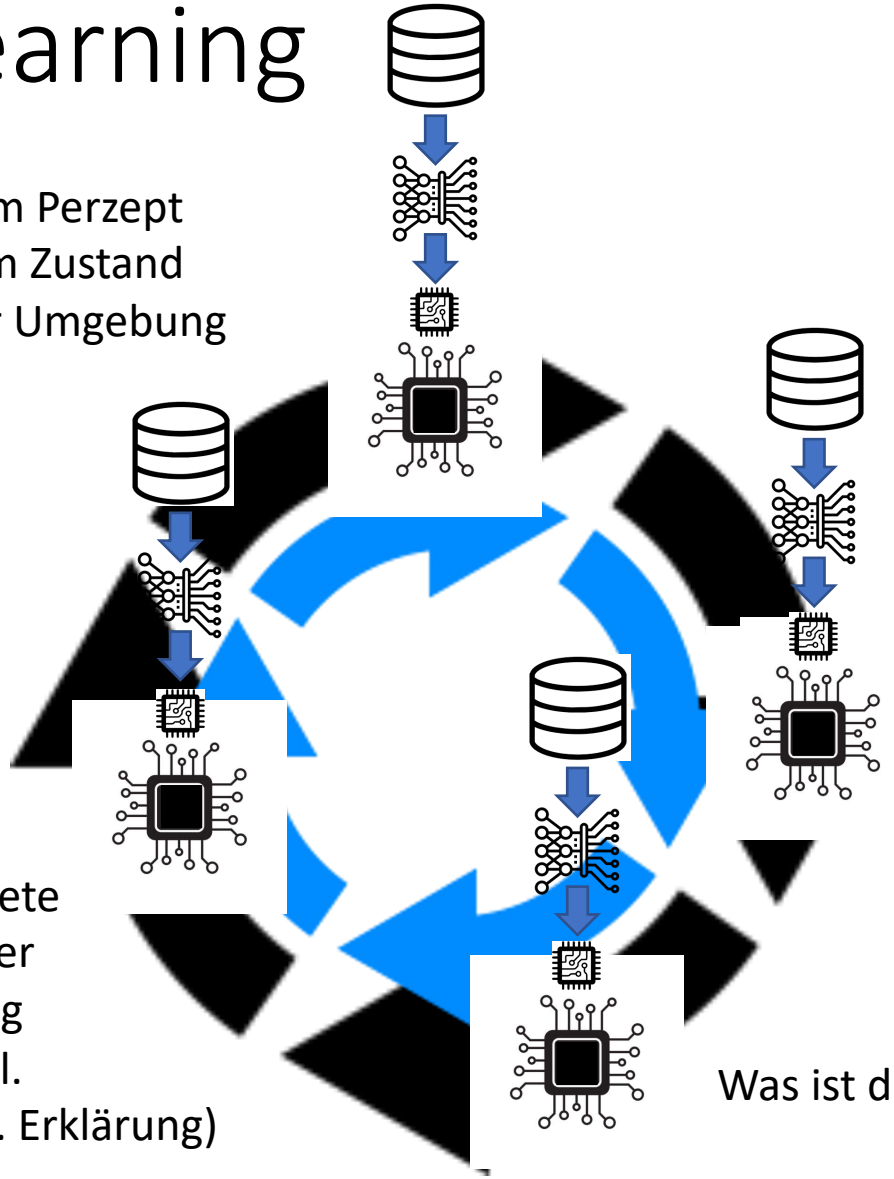

Ist mein Ziel bei aktuellem Zustand der Umgebung noch richtig gewählt? Habe ich ein neues Ziel?

Was ist die beste Aktion im aktuellen Zustand bei gegebenem Ziel? Handlungsbestimmungsstrategie

Führe berechnete Aktion aus (oder teile Umgebung Aktion mit) inkl. Zusatzinfo (z.B. Erklärung)

Beispiel AlphaGo

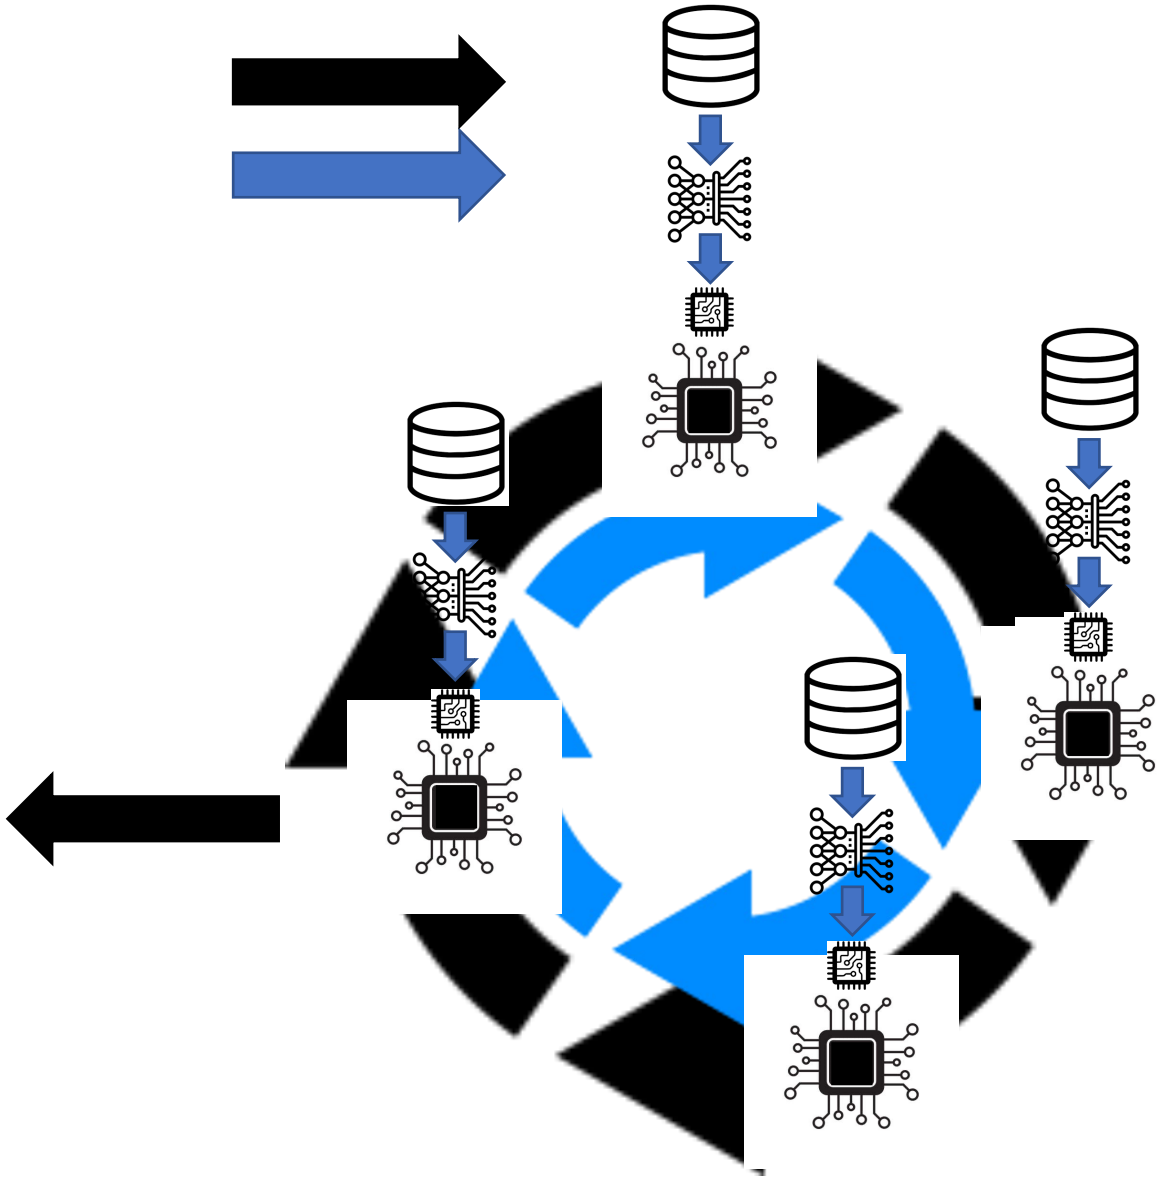

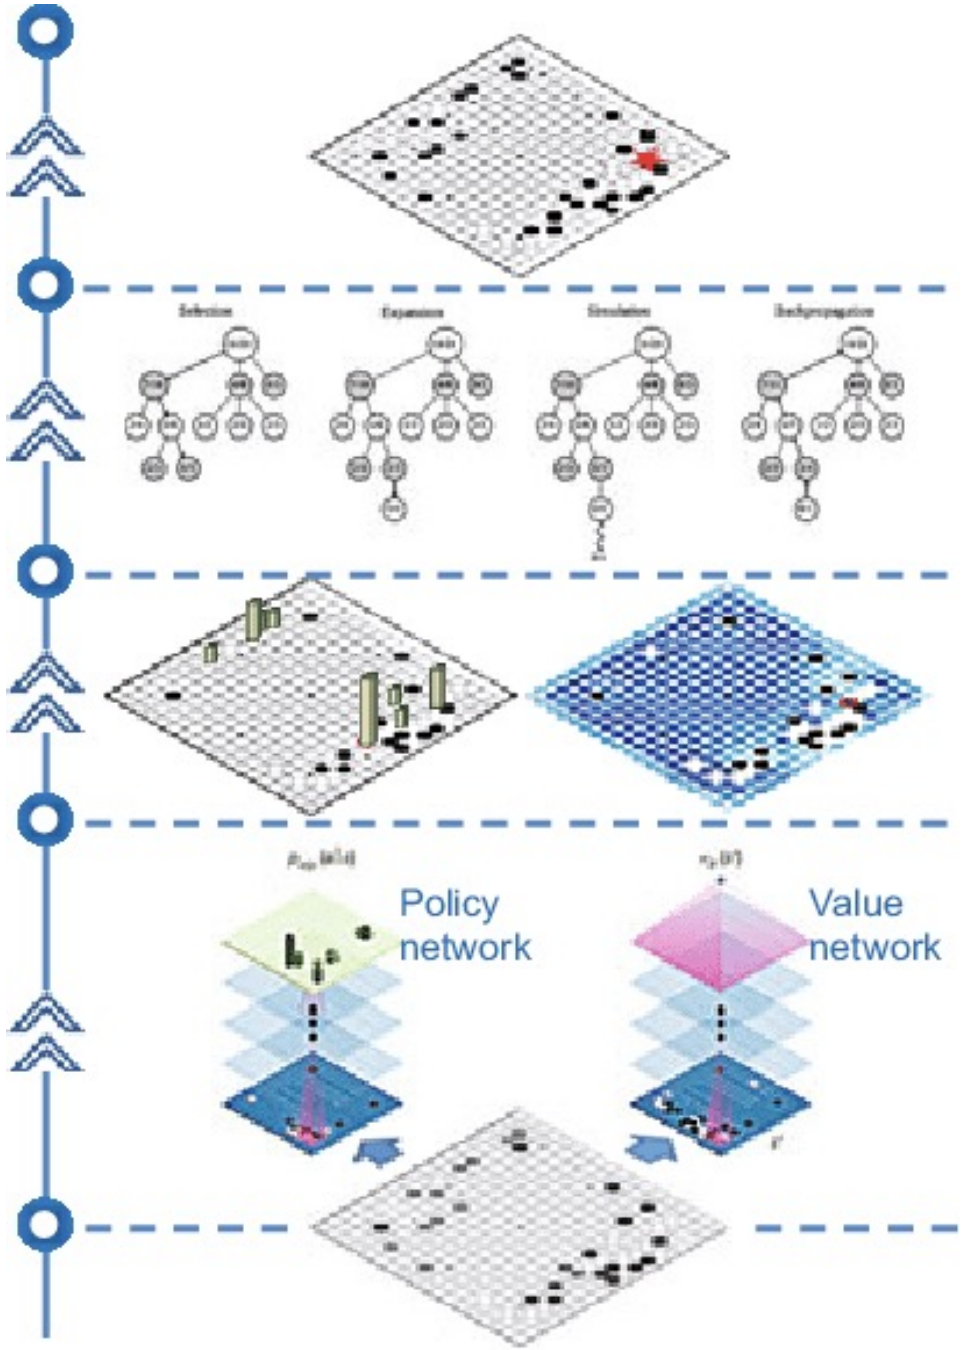

Und was ist mit Batchverarbeitung?

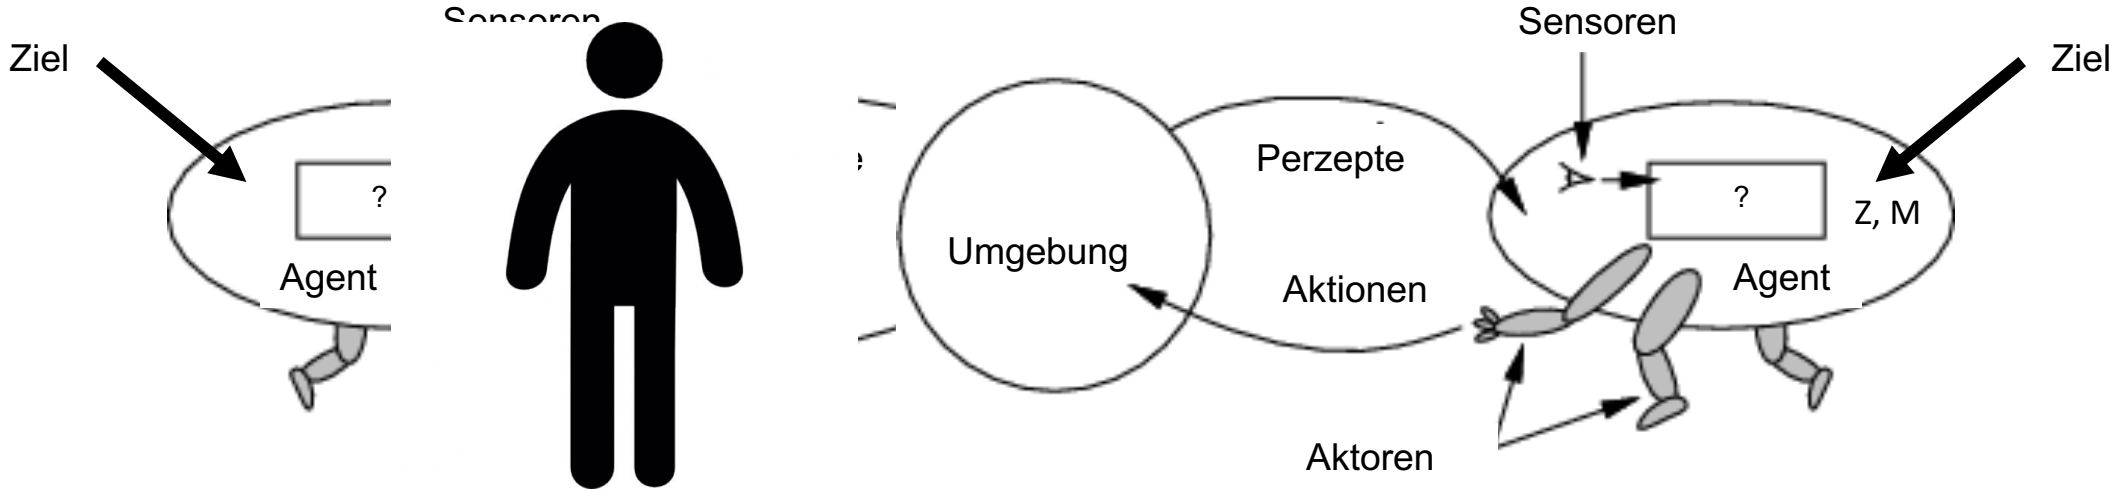

Environment
Percepts
Actions
Sensors
Actuators
Agent
G, M
Goal
?


Lübecker Nachrichten

Lübecker General-Anzeiger

2016

Facebook, CMU
Offline- + Online-Lernen
Spielerverhalten
nur online verfügbar

Artificial intelligence has now pretty much conquered poker

Called Pluribus, the AI is a formidable opponent at six-player no-limit Texas Hold'em

DAILY NEWS

World - Business - Finance - Lifestyle - Travel - Sport - Weather

2022

Issue: 240104

THE WORLDS BEST SELLING NATIONAL NEWSPAPER

Monday 5th June

First Edition

Artificial Intelligence
Beats 8 World Champions at Bridge

Using machine learning techniques, the small French startup company NukkAI....

https://nukk.ai

Where did we go wrong?

- Humans are intelligent to the extent that our actions can be expected to achieve our objectives
- Machines are intelligent to the extent that their actions can be expected to achieve their objectives
  - Give them objectives to optimize (cf control theory, economics, operations research, statistics)
- We don't want machines that are intelligent in this sense
- Machines are ***beneficial*** to the extent that ***their*** actions can be expected to achieve ***our*** objectives
- We need machines to be ***provably beneficial***

Maschinen unsicher bzgl. "unserer" Präferenzen: Hinterfragen der eigenen Ziele

Arguments taken from Stuart Russell's Presentations on Provably Beneficial AI

Yet some more detail

Percepts

Feedback

Mapping from percept to state of environment (depends on G)

Your performance might be improved

Action

Explanation

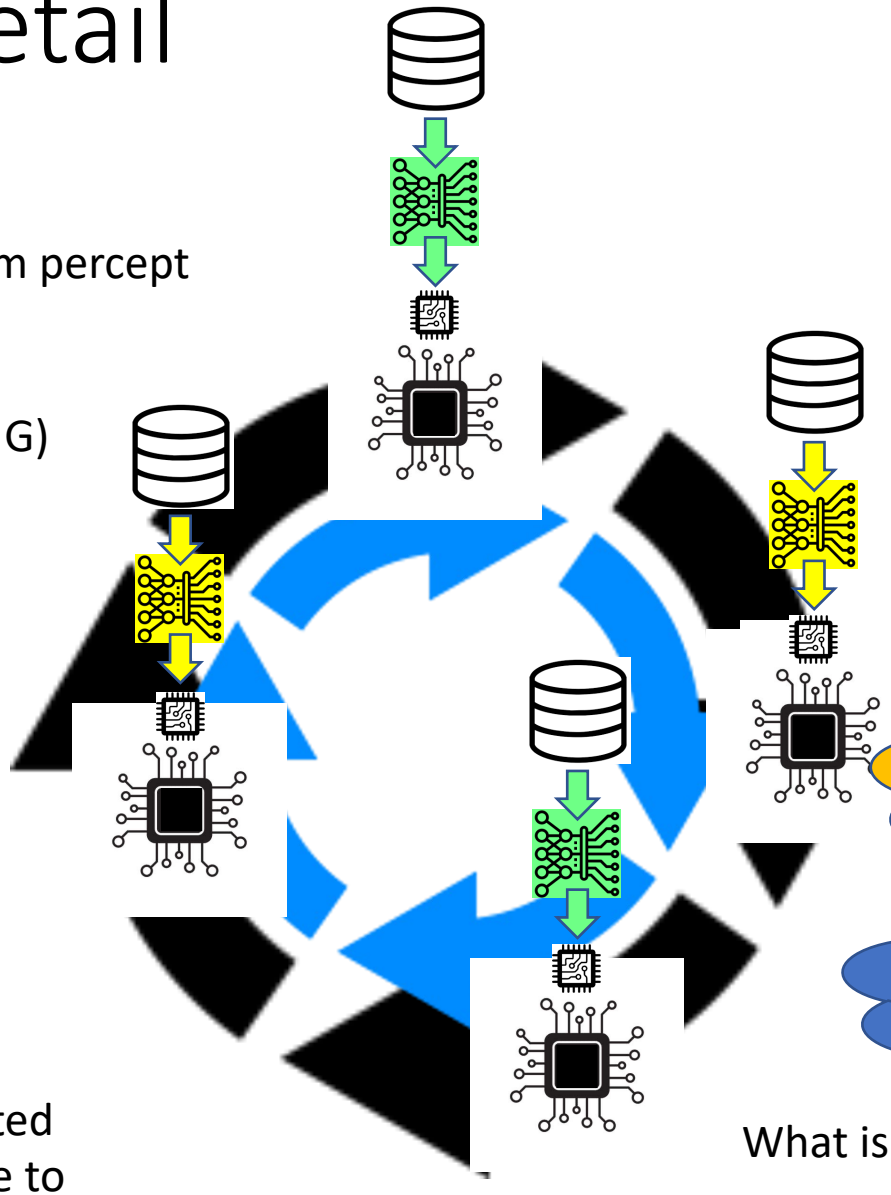

Return calculated action (prepare to answer the Why question)

Idea of a utility function

Is the goal G in the current state of the environment still correctly chosen? Should I have a new goal G'?

My utility will be zero then. How to prevent being switched off?

Hint: Possibly I'll be switched off

Maximize utility (possibly with sequence of actions)

What is the best action in the current state to achieve the goal? Strategy for determining action

Off-Switch Problem

Representative for dealing with unwanted behavior

Very much underspecified!

The instruction suggests having coffee would have higher value than expected a priori, ceteris paribus

Example: “Fetch some coffee”

Agents get better at maximizing the built-in utility function

What’s bad about better AI?

Can we switch off the agent if it “does not work as expected”?

“Can’t fetch coffee if I am dead.”

Incentive for preventing others from switching off the robot

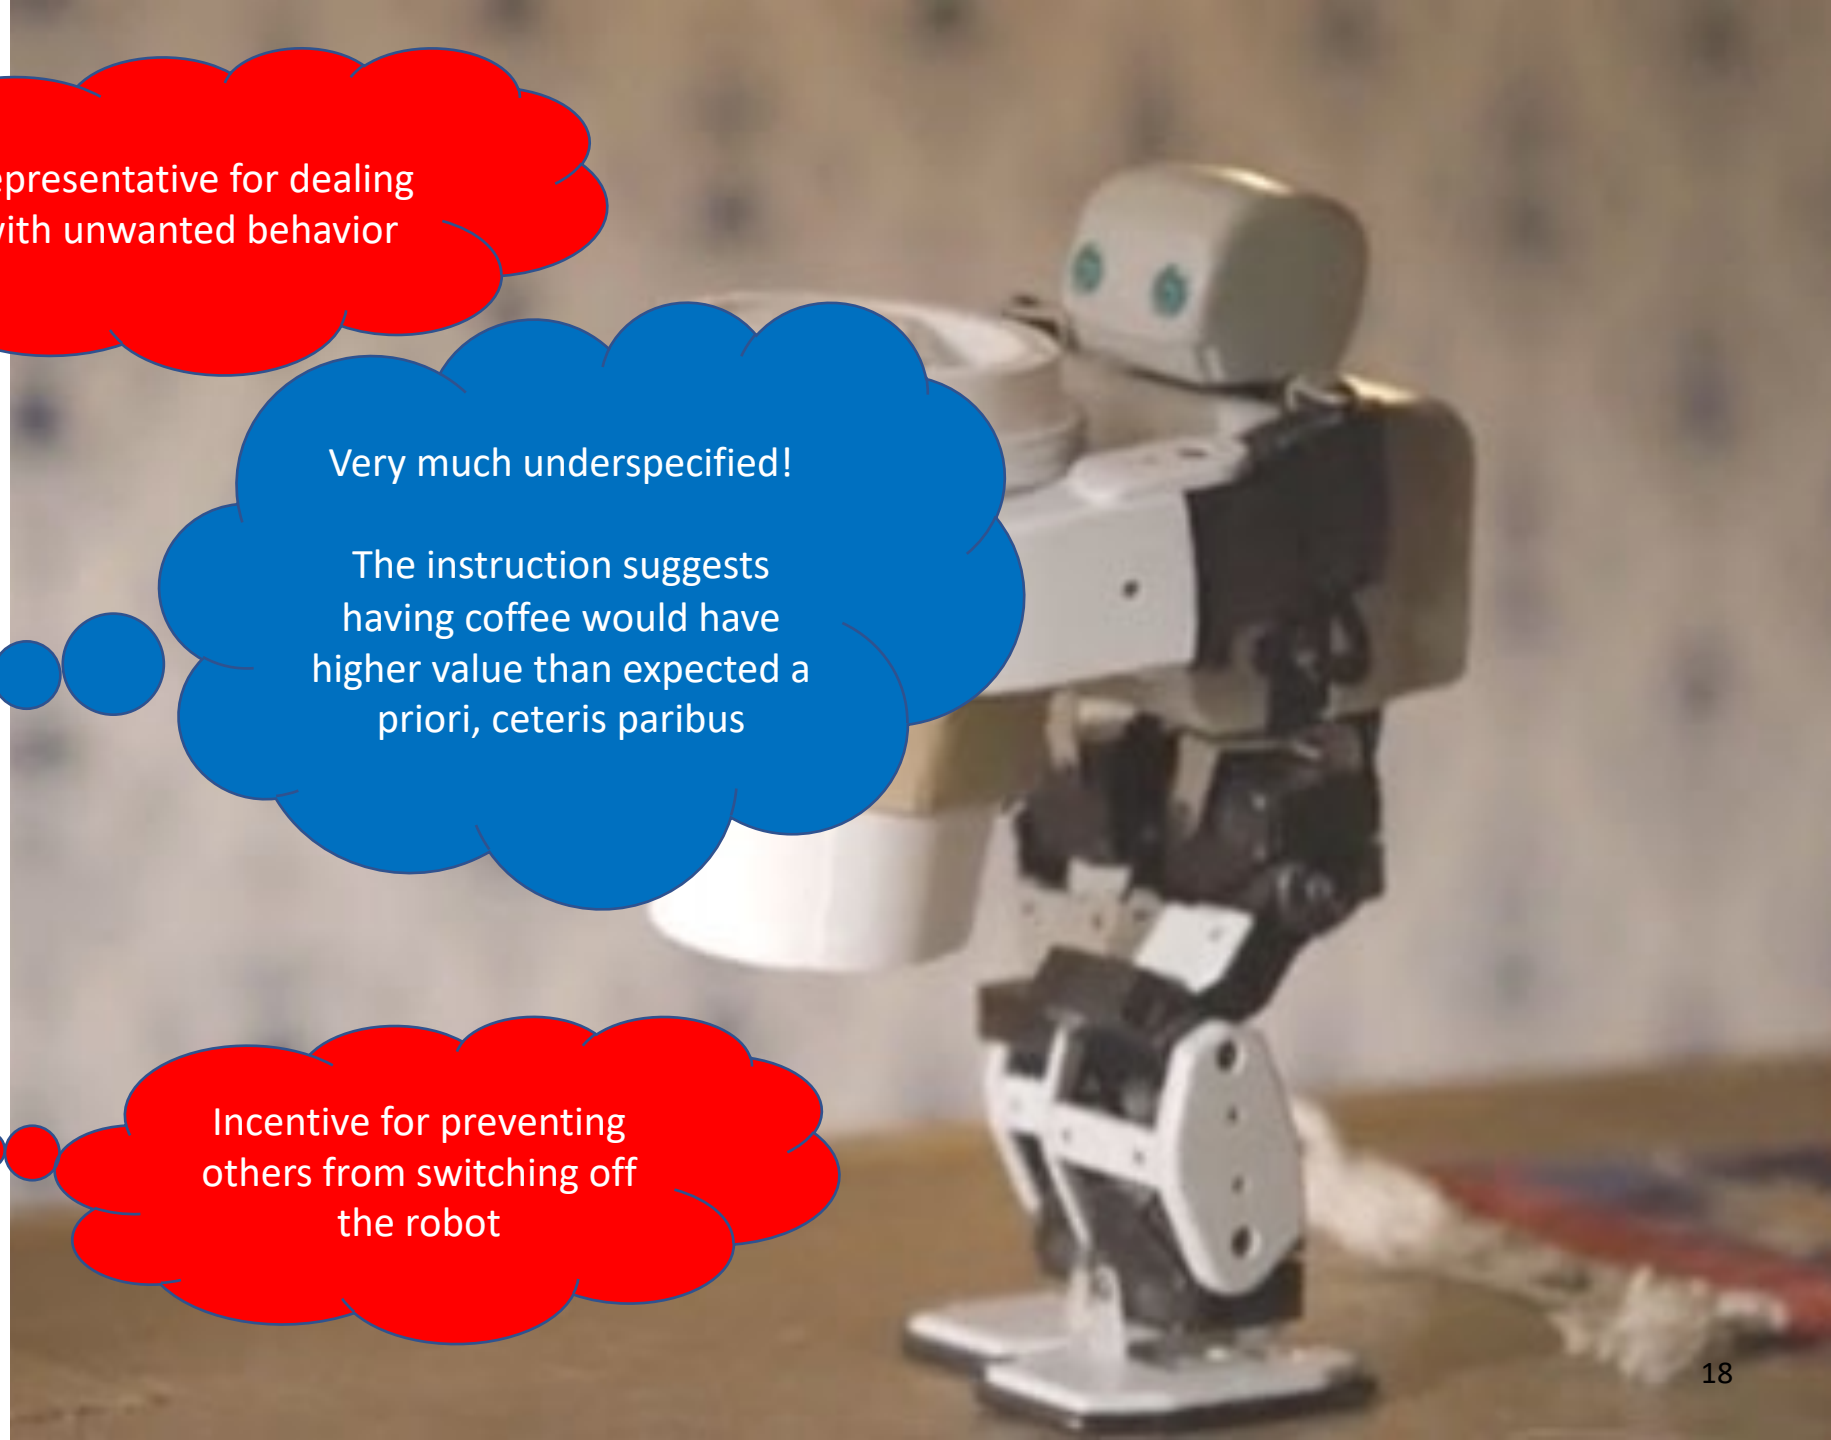

Before/after bringing coffee…

- Expected utility $E[U(state)] = \int x\, f(x) dx$
- Change the utility function

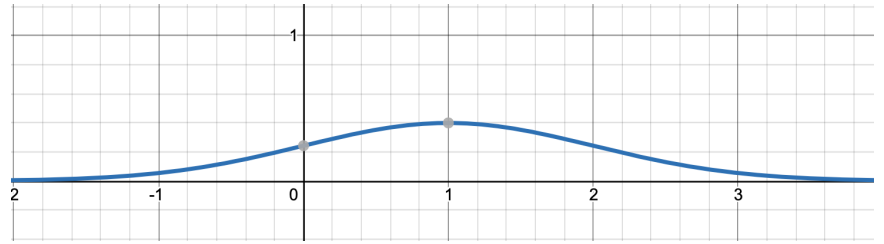

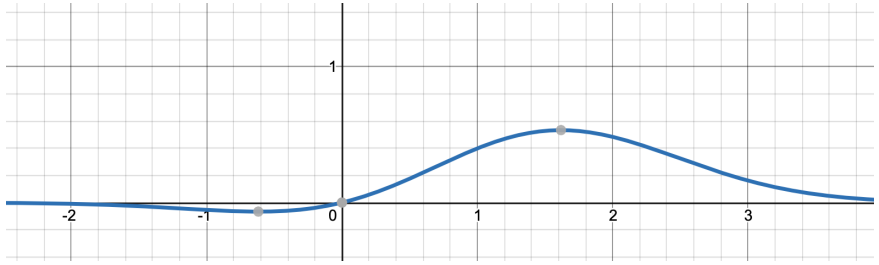

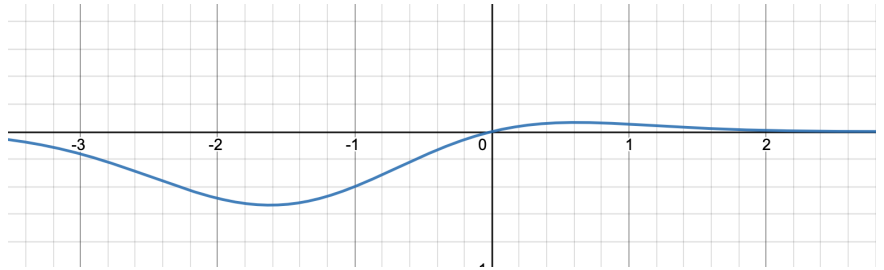

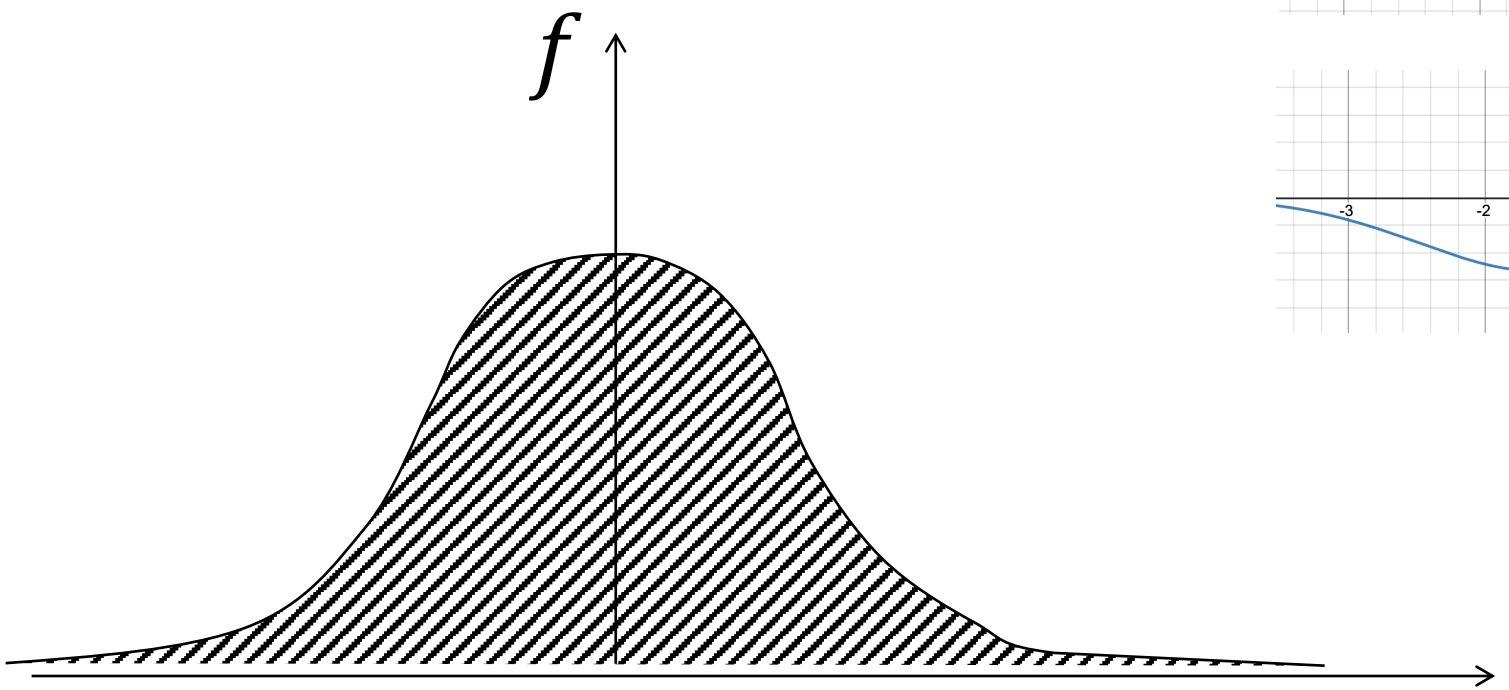

The off-switch problem

- A robot, given an objective, has an incentive to disable its own off-switch
- Claim: A robot with appropriate uncertainty about objective won't behave this way
- Example: Planning for the best action (sequence)

$$E[U(state)] = \int x\, f(x)dx$$

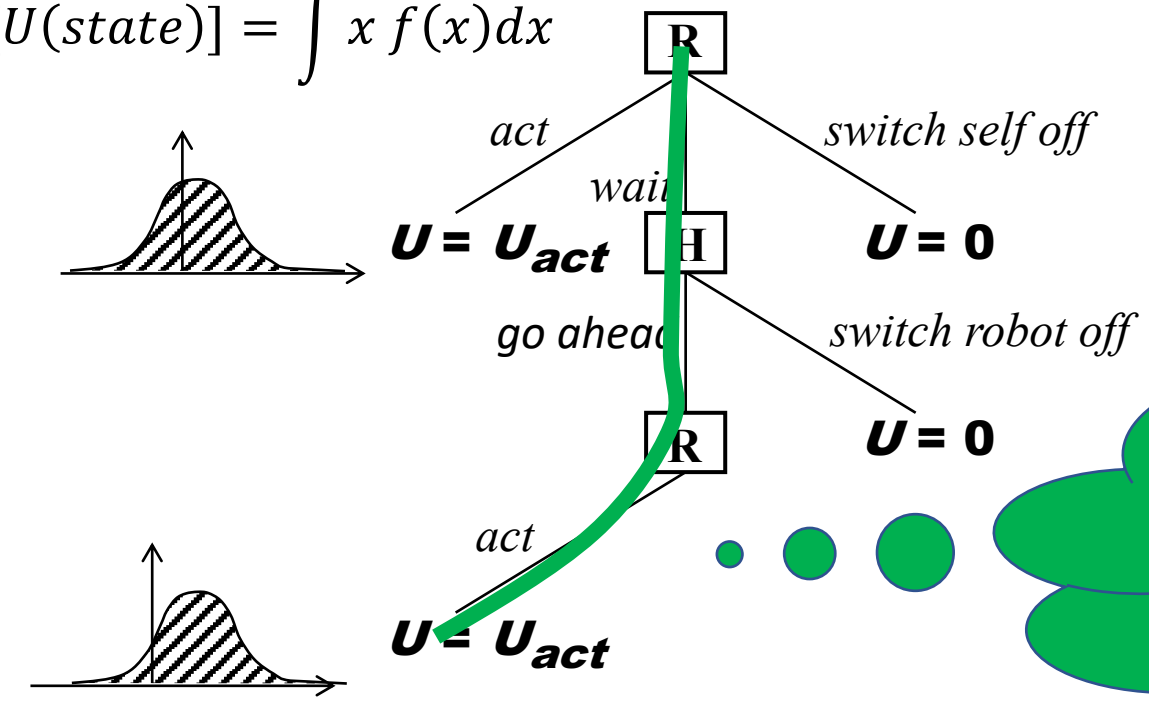

Tiny example for "ethical behavior"

Theorem:
This way robot has a positive incentive to allow itself to be switched off

Do not just talk about ethics!
Verify it!

Example taken from Stuart Russell's Presentations on Provably Beneficial AI

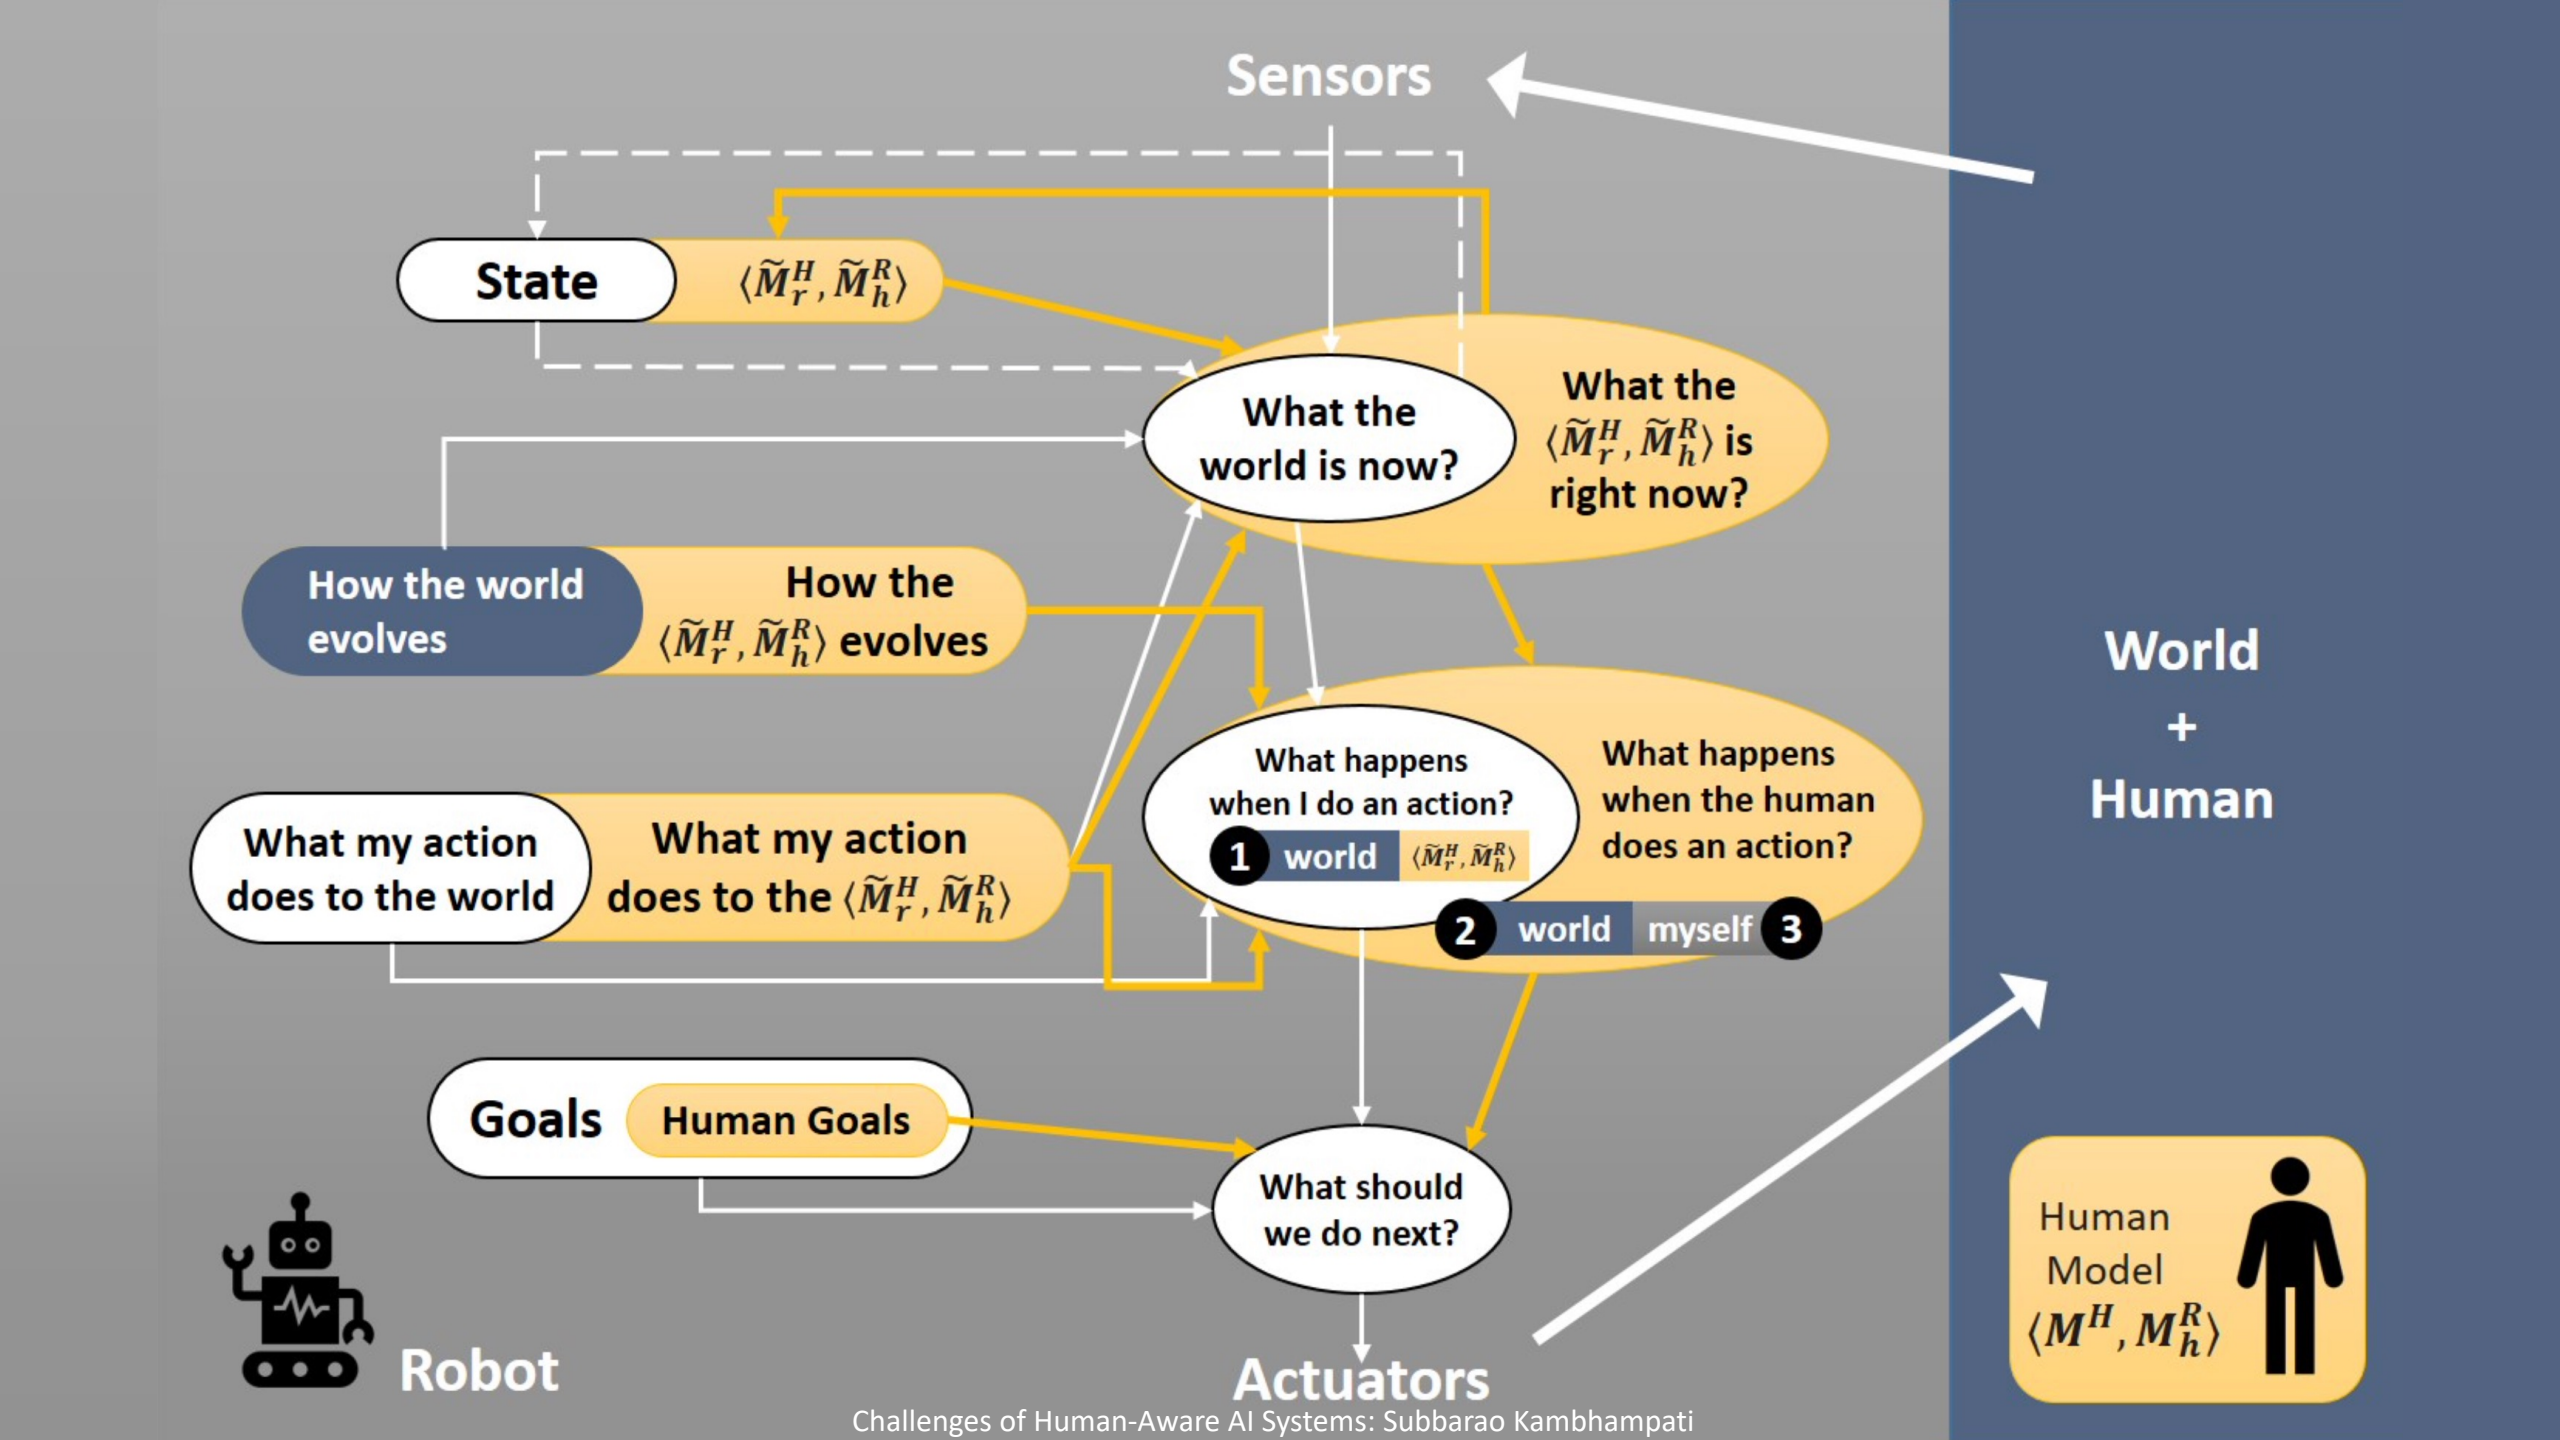
Sensors
State
$\langle \widetilde{M}_r^H, \widetilde{M}_h^R \rangle$
What the world is now?
What the $\langle \widetilde{M}_r^H, \widetilde{M}_h^R \rangle$ is right now?
How the world evolves
How the $\langle \widetilde{M}_r^H, \widetilde{M}_h^R \rangle$ evolves
What my action does to the world
What my action does to the $\langle \widetilde{M}_r^H, \widetilde{M}_h^R \rangle$
What happens when I do an action?
1 world $\langle \widetilde{M}_r^H, \widetilde{M}_h^R \rangle$
What happens when the human does an action?
2 world myself 3
Goals
Human Goals
What should we do next?
Robot
Actuators
World + Human
Human Model $\langle M^H, M_h^R \rangle$

KI als Denkparadigma

Perzepte

Feedback

Vom Perzept zum Zustand der Umgebung

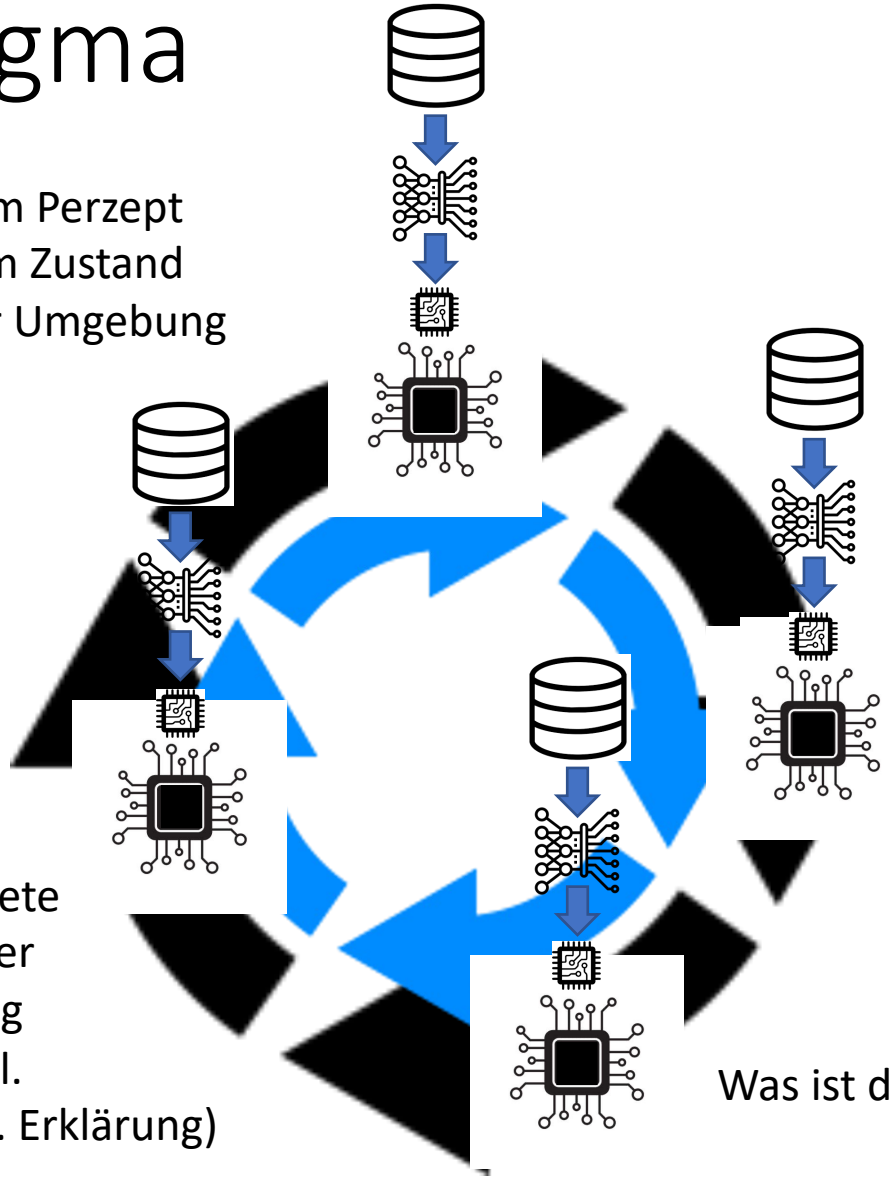

Ist mein Ziel bei aktuellem Zustand der Umgebung noch richtig gewählt? Habe ich ein neues Ziel?

Was ist die beste Aktion im aktuellen Zustand bei gegebenem Ziel? Handlungsbestimmungsstrategie

Führe berechnete Aktion aus (oder teile Umgebung Aktion mit) inkl. Zusatzinfo (z.B. Erklärung)

Derzeit relevant: Nutzung von Teilen des Zyklus

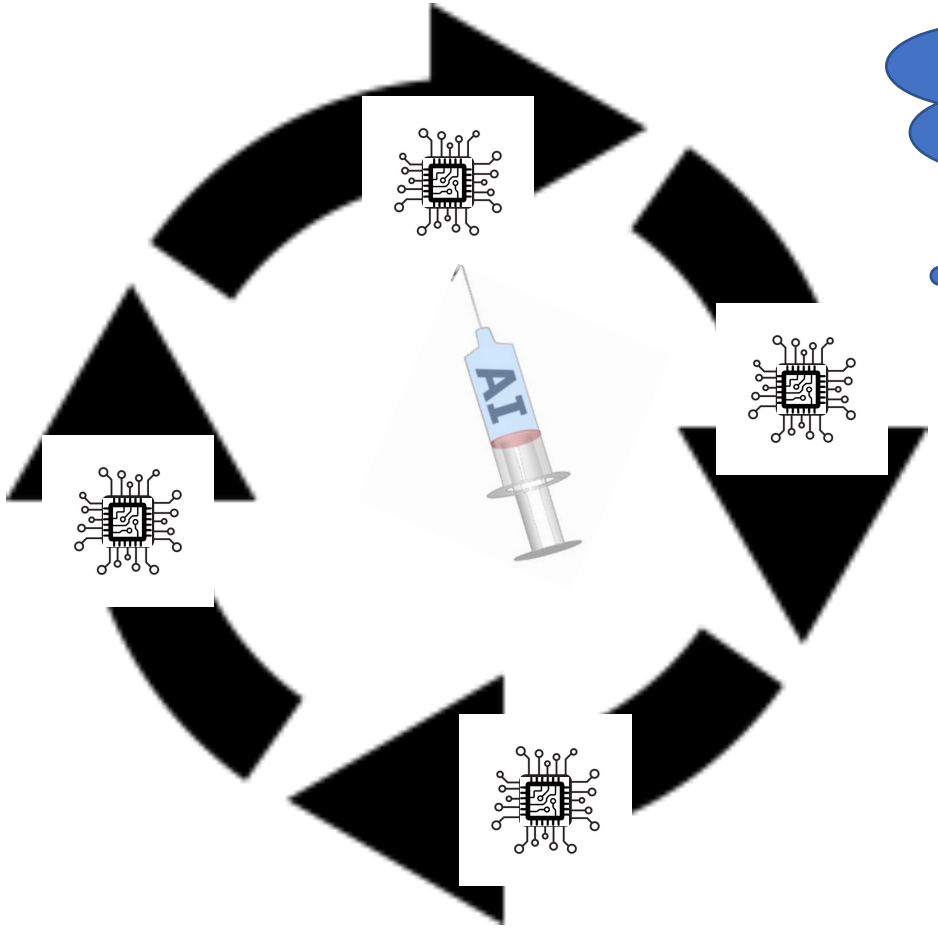

Blick auf den ganzen intelligenten Agenten schärft den Blick für Möglichkeiten

Cars

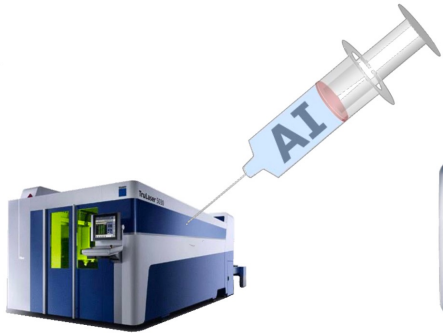

Manufacturing Equipment

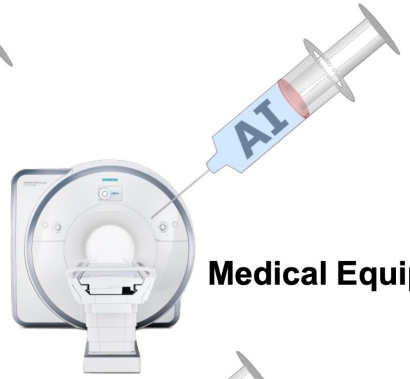

Medical Equipment

Home Appliances

Agricultural Machinery

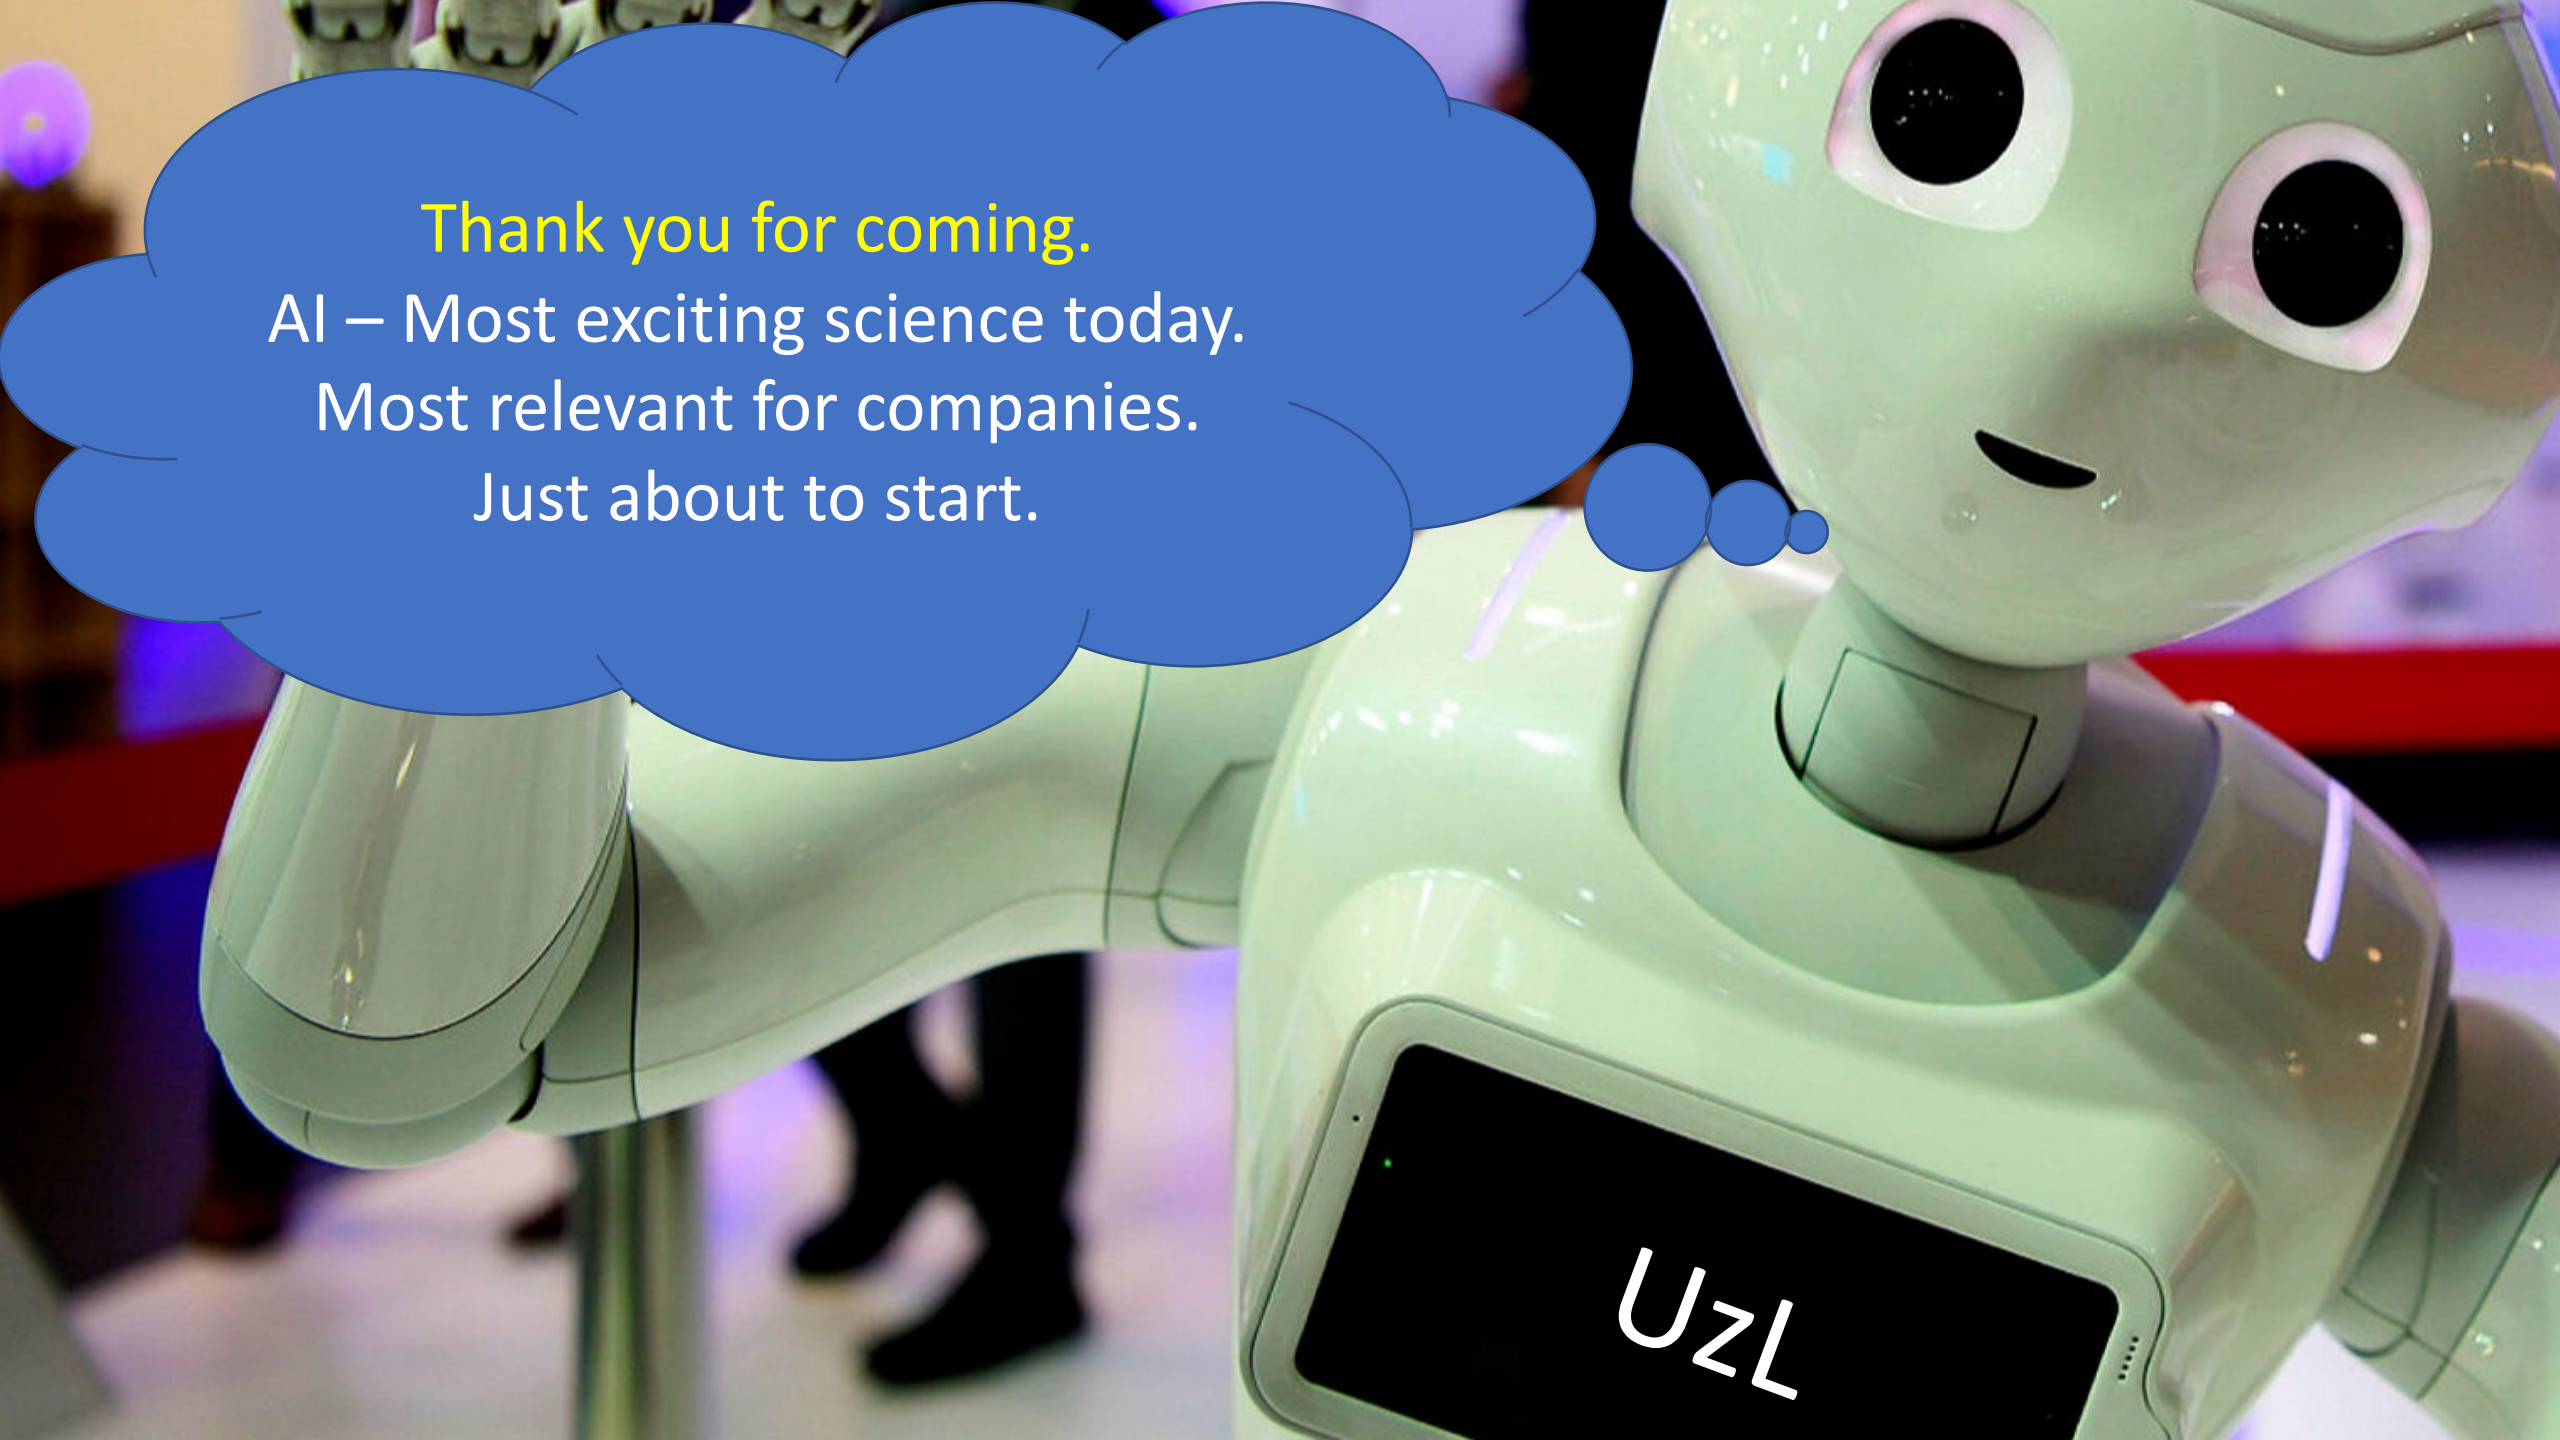
Thank you for coming.
AI – Most exciting science today.
Most relevant for companies.
Just about to start.
UzL