



Cognitive Game Characters

Science and Techniques behind **MarioAI**
Fabian Schrod, University of Tübingen



- **MarioAI** is a Super Mario Bros. clone, where characters are controlled by a **cognitive architecture**, that implements...
 - a **motivational** behavioral system
 - **learning** schematic world knowledge
 - **reasoning and planning** of abstract action sequences
 - **natural language** comprehension and generation
 - **social interactions** with and learning about other intelligent agents
- MarioAI was (mainly) developed at the University of Tübingen in several software engineering courses since 2012, about 45 people involved (currently: 17 new students)




What is this talk about?

1) Motivation

Impressions about MarioAI

A cognitive science perspective on game characters

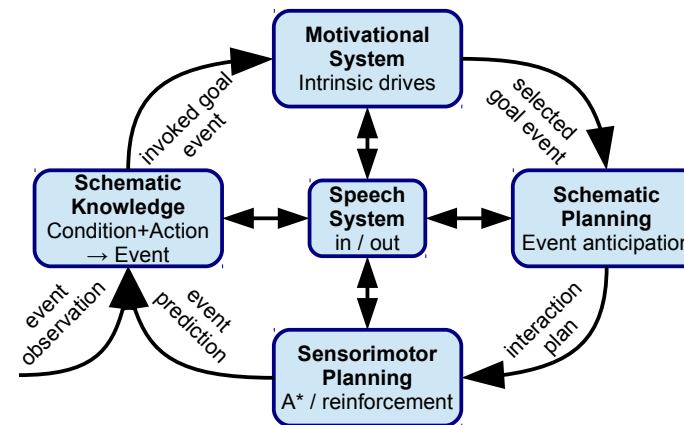
Aspects of cognition



2) About MarioAI

Science


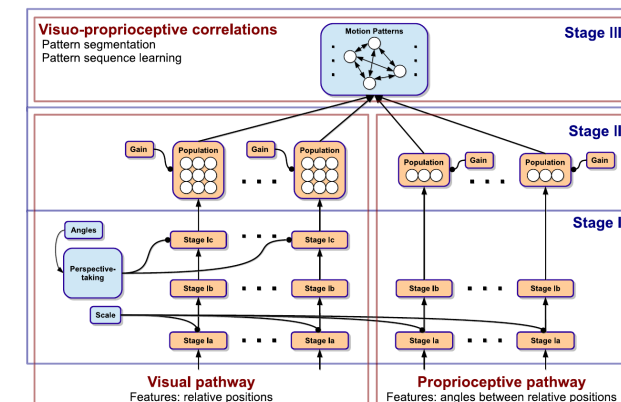
Methodologies



3) Future Work and Opportunities

Social AI

(Deep) neural networks



1.1) Motivation: Impressions

- In 2015 we published the video “Mario Lives!” about MarioAI
- It went somewhat viral...
 - 1 million views in a few days
 - World wide online media coverage (200+ online articles), etc...
 - Tons of feedback in social media
 - Peoples Choice Award of the AAAI
- Recently we published the follow-up “Mario Becomes Social!” about the current state of the project, including multiple agents with some “social AI”.
- Here are some impressions...



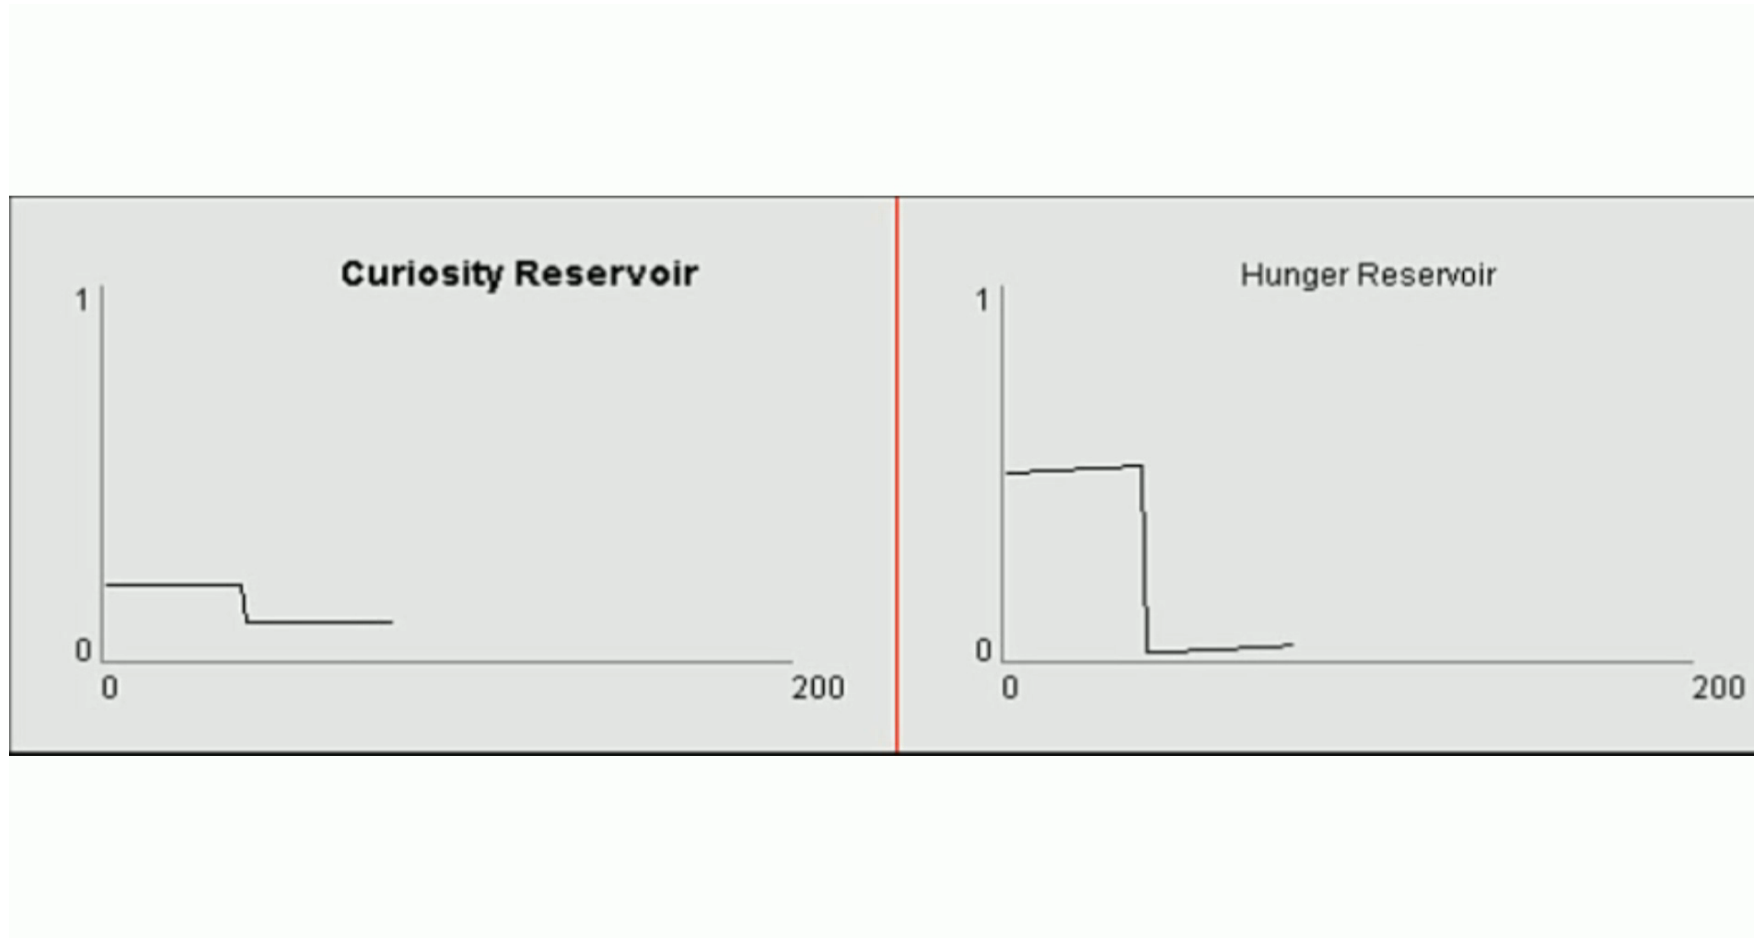
1.1) Motivation: Impressions

- Natural language: How does Mario feel about being 'self-aware'?



1.1) Motivation: Impressions

- Self-motivated behavior:



1.1) Motivation: Impressions

- Learning and observation: Multi-agent interactions and “social AI”



Quote:

"Researchers at Germany's University of Tuebingen are the naive harbingers of doom, and the Super Mario artificial intelligence they developed is their omen."
(joystick.com)



Quote:

"Hopefully the researchers haven't shown future AI robots how to work together to overthrow humanity." (qz.com)



1.2) Motivation: A Perspective on Games

- Why did that possibly happen?
 - A lot of phantasy from the audience...
 - Contemporary discussion about and public attention for the dangers of AI
 - Huge familiarity with the Super Mario games
 - Certainly some kind of humor in this connection...




- The impression that Mario has now become an **autonomous, intelligent, and intentional creature**




1.2) Motivation: A Perspective on Games

- One thing involved in this impression: Our ability to **project onto others**, even robots, game characters, or inanimate things




1.2) Motivation: A Perspective on Games

- One thing involved in this impression: Our ability to **project onto others**, even robots, game characters, or inanimate things



1.2) Motivation: A Perspective on Games

- One thing involved in this impression: Our ability to **project onto others**, even robots, game characters, or even inanimate things



(Heider & Simmel, 1944)



- Also when we play games, we take roles, and we project roles.
 - We want to believe that game characters are somewhat living, motivated beings that have feelings and goals.
 - More generally, we are able to imagine and simulate hypothetical situations, and infer and emulate other's intentions and emotions.
- **Perspective-taking**
 - is the ability to “step into another person's shoes” to understand their position or point of view,
 - is a key aspect of social intelligence,
 - **works better/faster when you identify with the observed person (the model).**
- Apart from appearance, **believability** of a game character is mediated by human-like, versatile behavior.



1.3) Motivation: Aspects of Cognition

- Thus, to develop believable game agents, we should understand **human cognition!**
- What is cognition?
 - Cognition is "the mental action or process of acquiring **knowledge and understanding** through **thought, experience**, and the **senses**." (oxford dictionary)
 - Cognition (lat. cognoscere ‚recognize‘, ‚experience‘, ‚get-to-know‘) is about **information structuring** by a **behaviorally-controlled** system. (wikipedia)
- So, cognition is about **behavior** and **information**, but...
 - where does our behavior come from?
 - how do we structure and use information (for conceptualization, abstraction)?
 - how are we able to talk about all that?
- Cognitive science tries to answer these questions in an interdisciplinary approach...



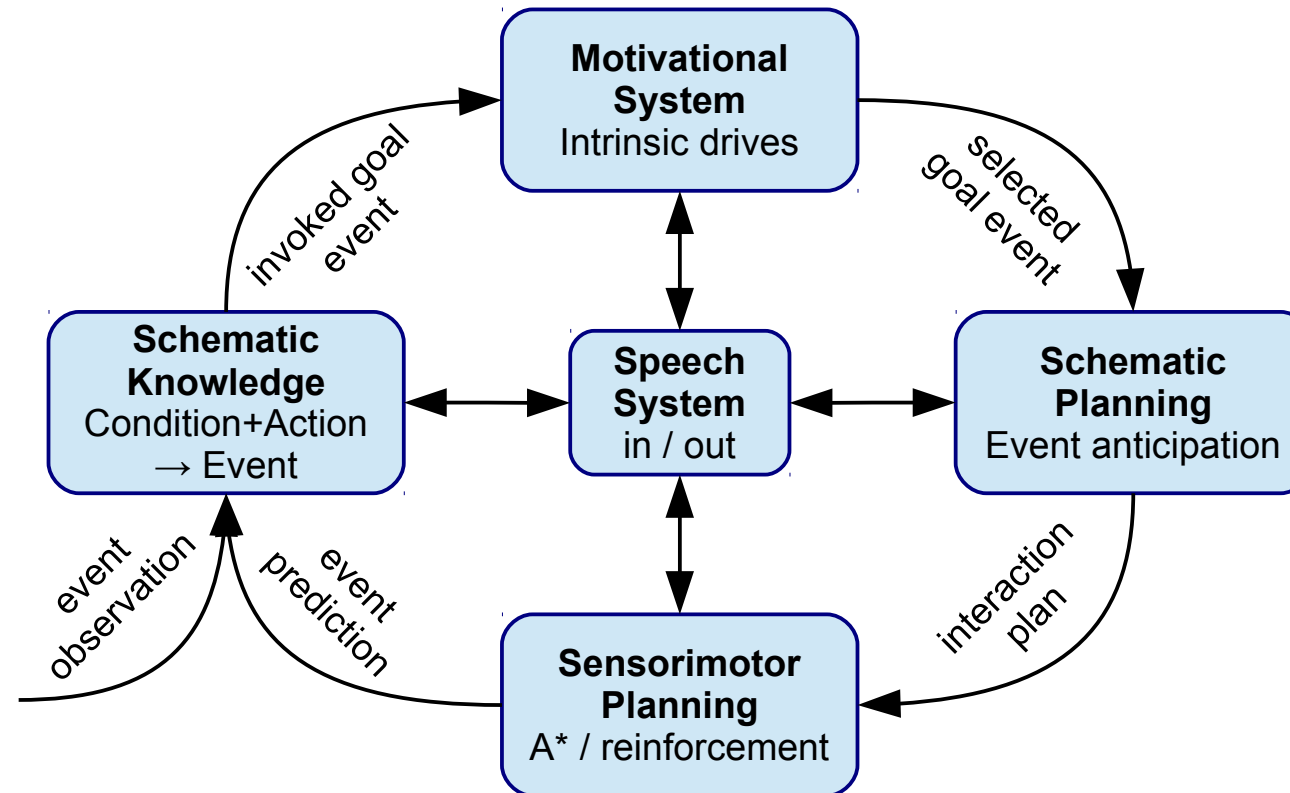
1.3) Motivation: Aspects of Cognition

- Here are some (recent, relevant, assumed) insights from cognitive science:
 - Cognition is **embodied**.
 - Experiences are gathered through and shaped by our own body.
 - Innate motivations are the driving force of behavior.
 - Projecting our embodied codes is probably the key to understand others.
 - Cognition means to **predict**
 - ...which sensory information correlates with other information.
 - ...what is likely to happen next.
 - Cognition means perceiving and causing **events**.
 - The brain plans in terms behaviorally relevant event codes.
 - Cognition means to **simulate and imagine**
 - ...situations, hypotheses and perspectives.
 - Human cognition is highly **interactive and social!**
 - Communication and culture are fundamentals of our thinking.



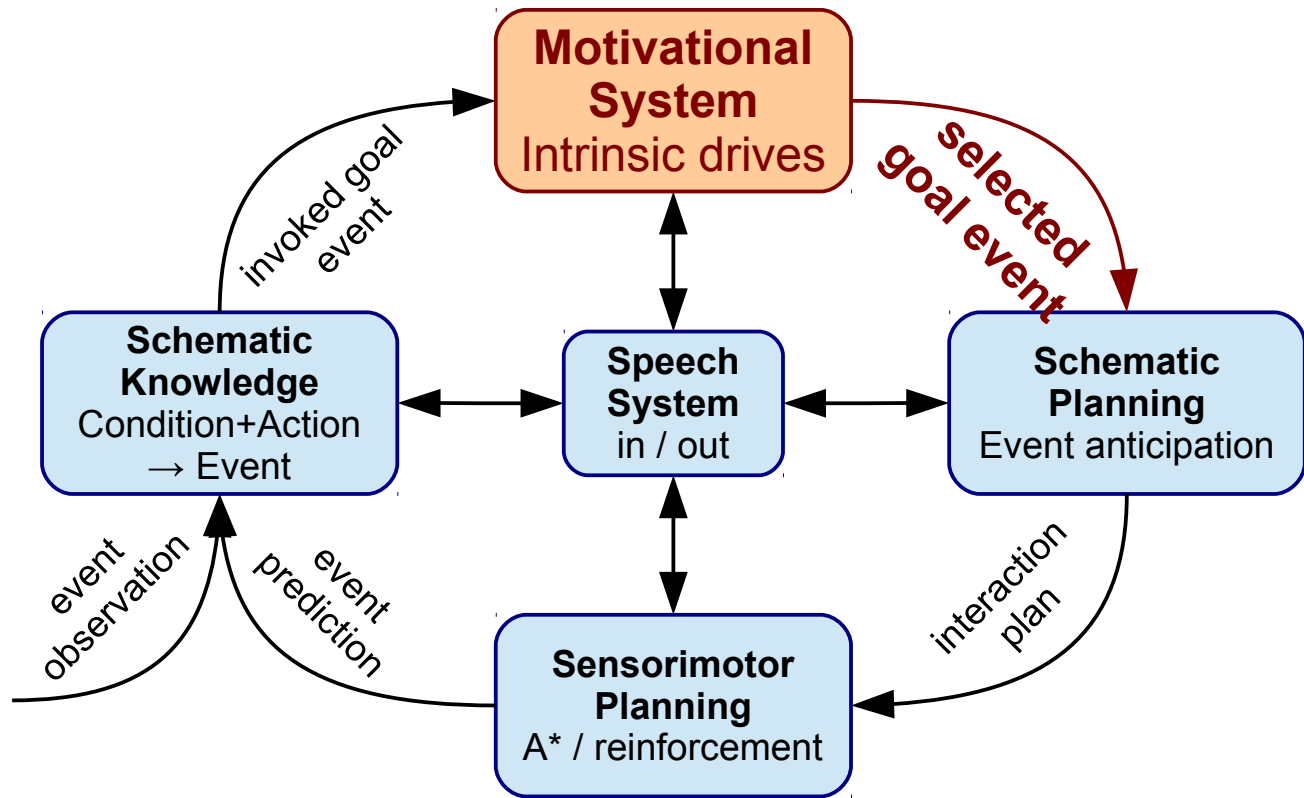
2) MarioAI: Overview

- In MarioAI, we model these processes by means of a cognitive processing loop.
- Simplified, it looks like this:



2.1) MarioAI: Motivations and Goals

- **Motivations result in drives** that maintain an internal homeostasis.
 - Homeostasis: Property of a system in which variables are regulated such that internal conditions remain stable.
- **Drives result in goals** → no goals without motivations!



I want to get healthy!

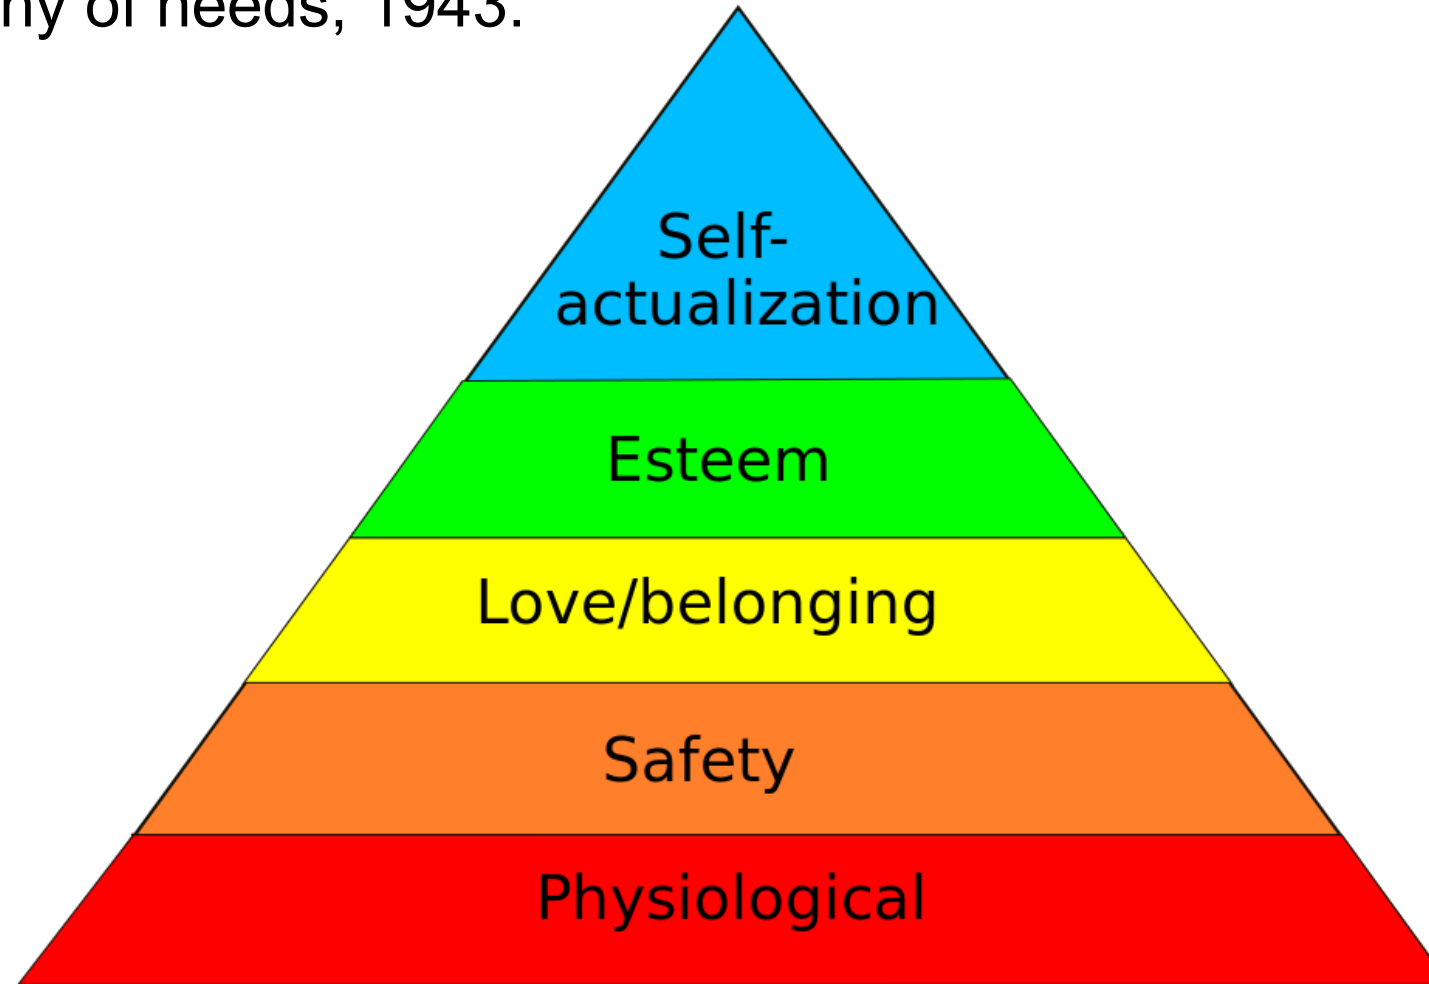


- The brain wants to keep certain signals (neurotransmitters, hormonal concentration) on a certain level.
- Typical example for a drive: **hunger!**
 - **Long-term** eating behavior driven by hormone **Leptin**, which is generated by fat-storing cells
 - **Short-term** eating behavior driven by peptide **Ghrelin**, which is generated the more the emptier the stomach.
 - **Dopamine** release is associated with
 - **rewards** (the feeling of satisfaction when eating while being hungry), and
 - **learning** (which food we like).



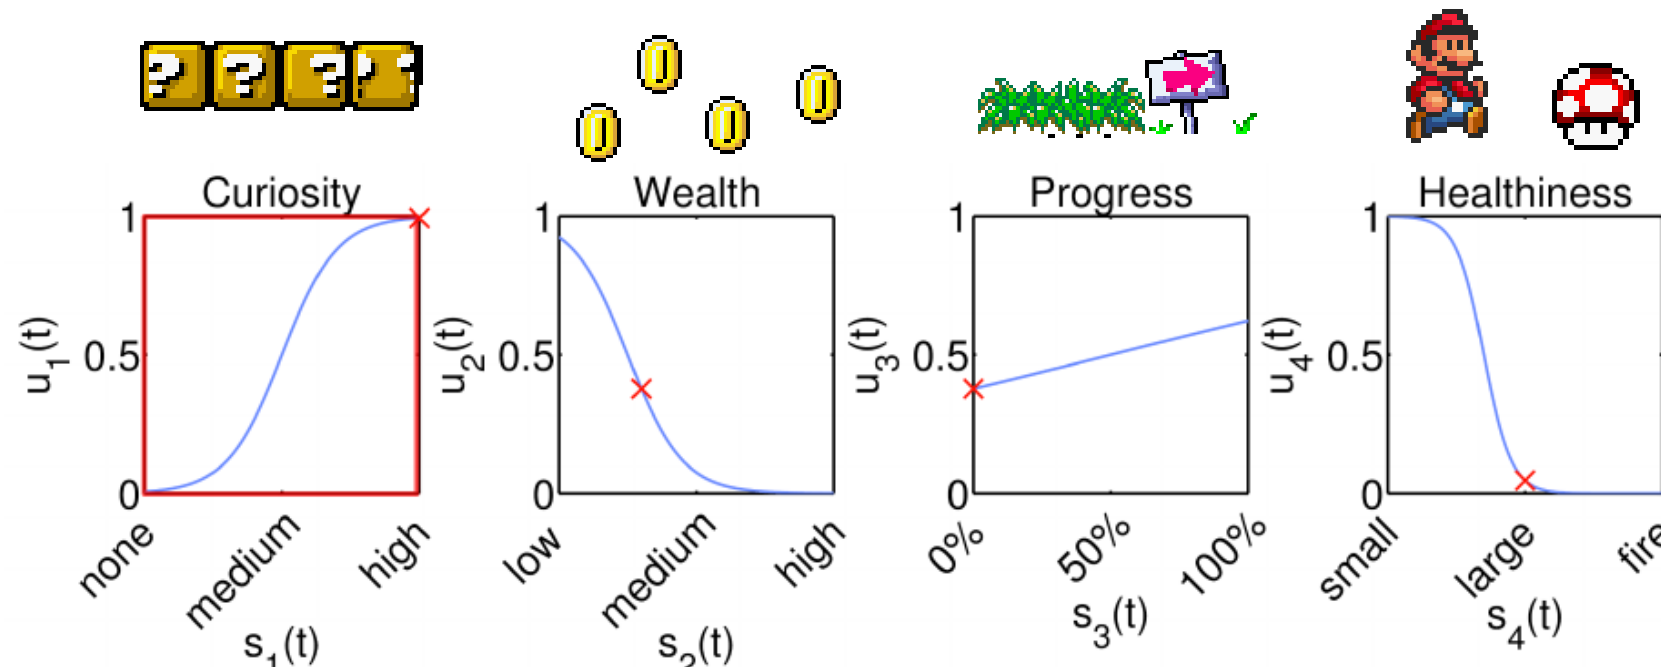
2.1) MarioAI: Motivations and Goals

- Motivations can be much more complex.
- Maslow's hierarchy of needs, 1943:



2.1) MarioAI: Motivations and Goals

- In MarioAI, we primarily focused on four basic motivations:
 - **Curiosity** (learn about the world and interactable objects)
 - **Wealth** (gain a high number of coins)
 - **Progress** (complete the level)
 - **Healthiness** (“survival instinct”)
- These are modeled by **dynamic reservoirs** (drive as function of a reservoir state):



2.1) MarioAI: Motivations and Goals

- The drives are changed in response to specific, **motivation-relevant game events**.
- Thus, Mario strives to invoke **rewarding goal events**, while avoiding **punishments**:

Reservoir		Trigger Events	
Curiosity	(r)	KNOWLEDGE_INCREASE	(+)
		OBJECT_INTERACTION	(+)
Wealth	(c)	COIN_INCREASE	(+)
		COIN_DECREASE	(-)
Progress	(c)	RIGHT_GOAL_REACHED	(+)
		LEFT_GOAL_REACHED	(-)
Healthiness	(c)	HEALTH_INCREASE	(+)
		HEALTH_DECREASE	(-)

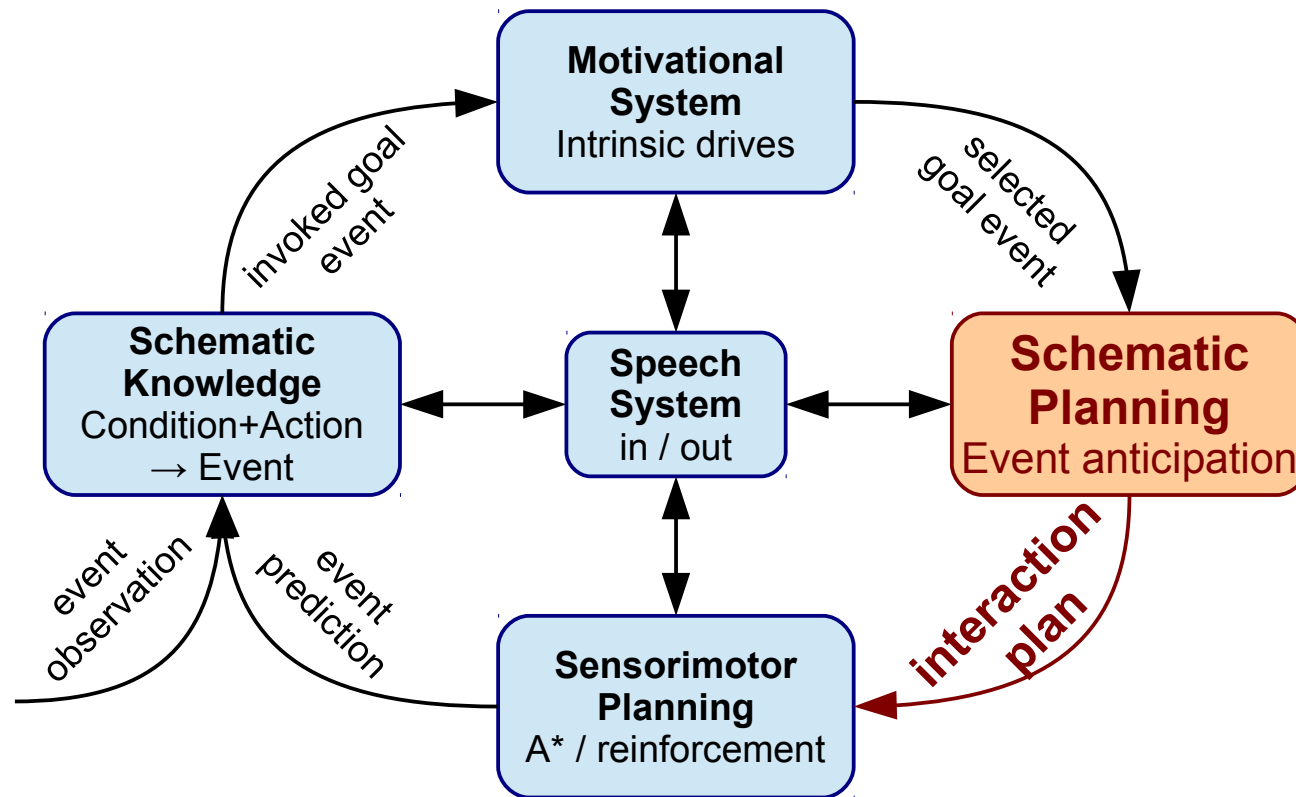
- A rewarding **goal event** is selected with a **probability** proportional to the “state” of a reservoir:

$$P(e = d) = \frac{u_d(t)}{\sum_d u_d(t)}$$



2.2) MarioAI: Schematic Planning

- A **goal event** is now selected. But how to invoke it in the game?
- The **schematic planning** module allows to plan action sequences on an abstract, conceptual level.




I want to get healthy!
How do I do that?



2.2) MarioAI: Schematic Planning

- In cognitive science, an event can be defined as a **behaviorally relevant segment of time** with **beginning** and **end**.
 - Significant changes in our perception can be considered as beginning or end of an event (e.g. touching an object, or higher level: a new year).
- We can perceive but also **cause** events. This is assumed to be involved in **planning**.
 - In the brain, neural pathways between the striatum and the frontal lobe have been associated with planning.
 - Frontal lobe: “organ of civilization”, situation-dependent behavior
 - Striatum: executive part of the brain, motivation driven selection and suppression of action patterns
- However, to cause events, we need to know about the **circumstances** that allow us to do so...



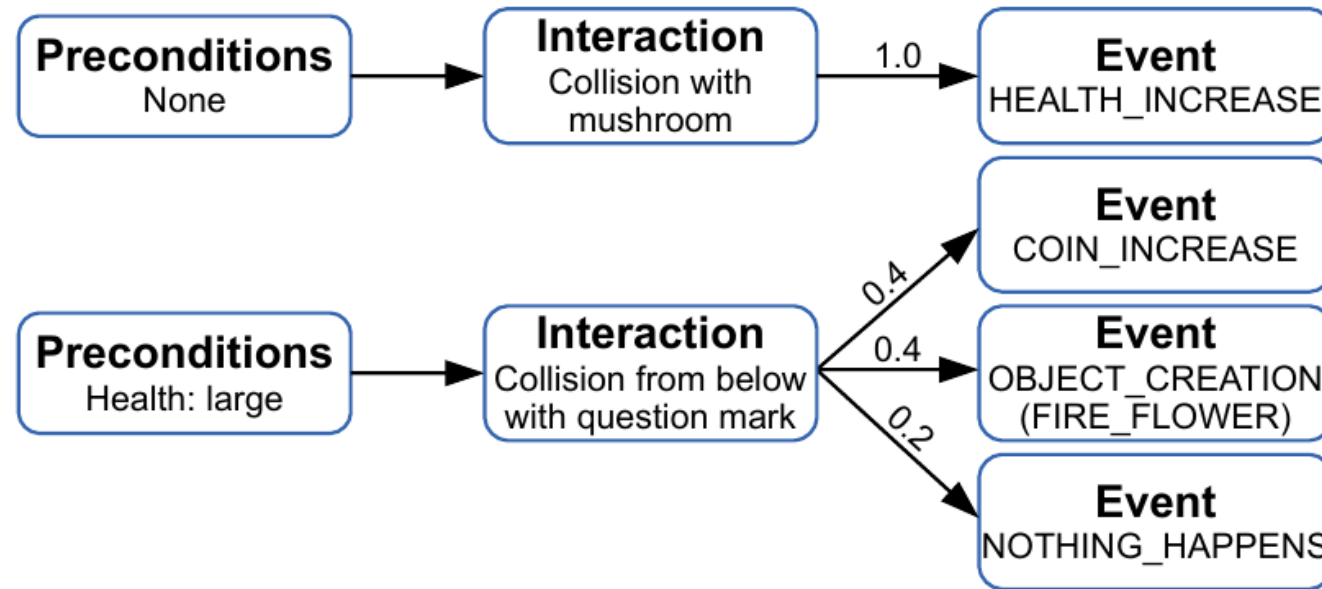
2.2) MarioAI: Schematic Planning

- In psychology, a **scheme** is defined as a mental knowledge structure that provides generalized information about **specific objects** or **concepts**.
- In MarioAI, schematic knowledge consists of a mapping from
 - **preconditions** (mainly character attributes) and
 - **object interactions** (mainly directional collisions) to
 - **event probabilities** (all effects and their likelihood).



2.2) MarioAI: Schematic Planning

- Assume the following schematic knowledge base is given:

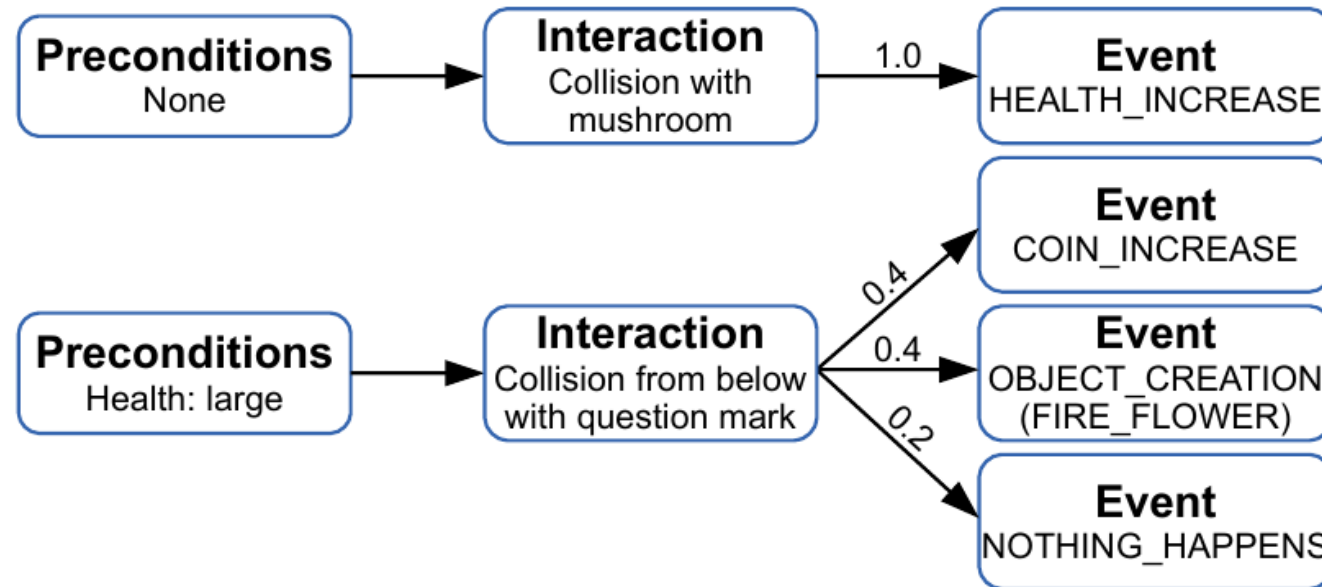


Mushrooms are always good for my health!



2.2) MarioAI: Schematic Planning

- Assume the following schematic knowledge base is given:





Hm.. if my health condition is 'large', I might get a coin when colliding with a question mark block.



2.2) MarioAI: Schematic Planning

- Event goals without preconditions are easy to translate to concrete interactions.
- Given the knowledge, the event “HEALTH_INCREASE” can directly be converted to a collision with a mushroom:



I'll look for a mushroom!

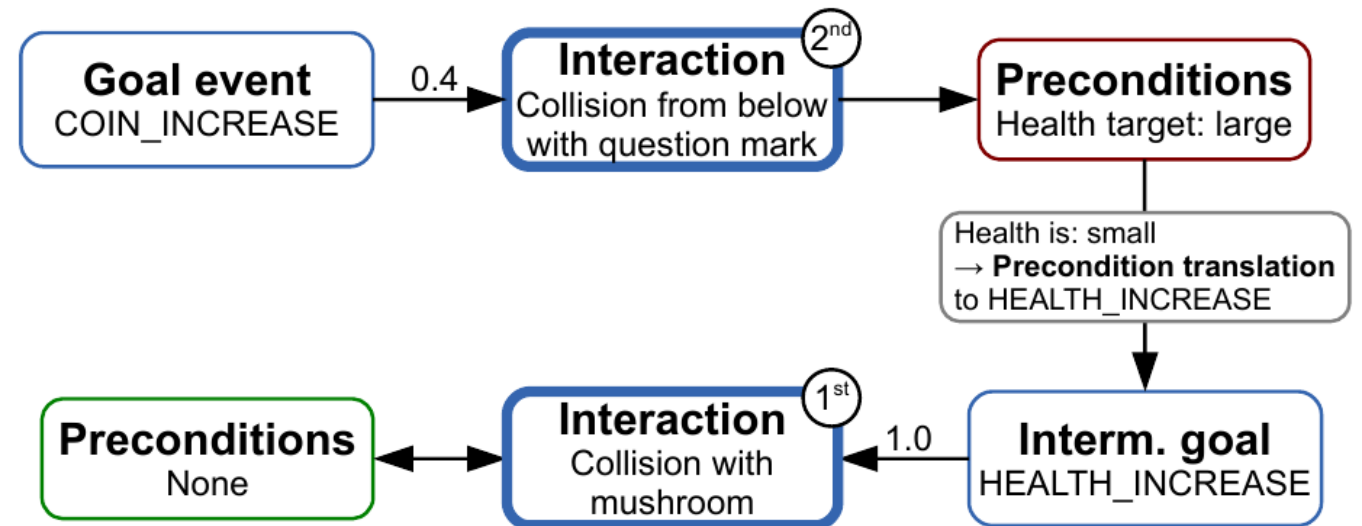
- This extends to goals for which the preconditions are **fulfilled**.



2.2) MarioAI: Schematic Planning

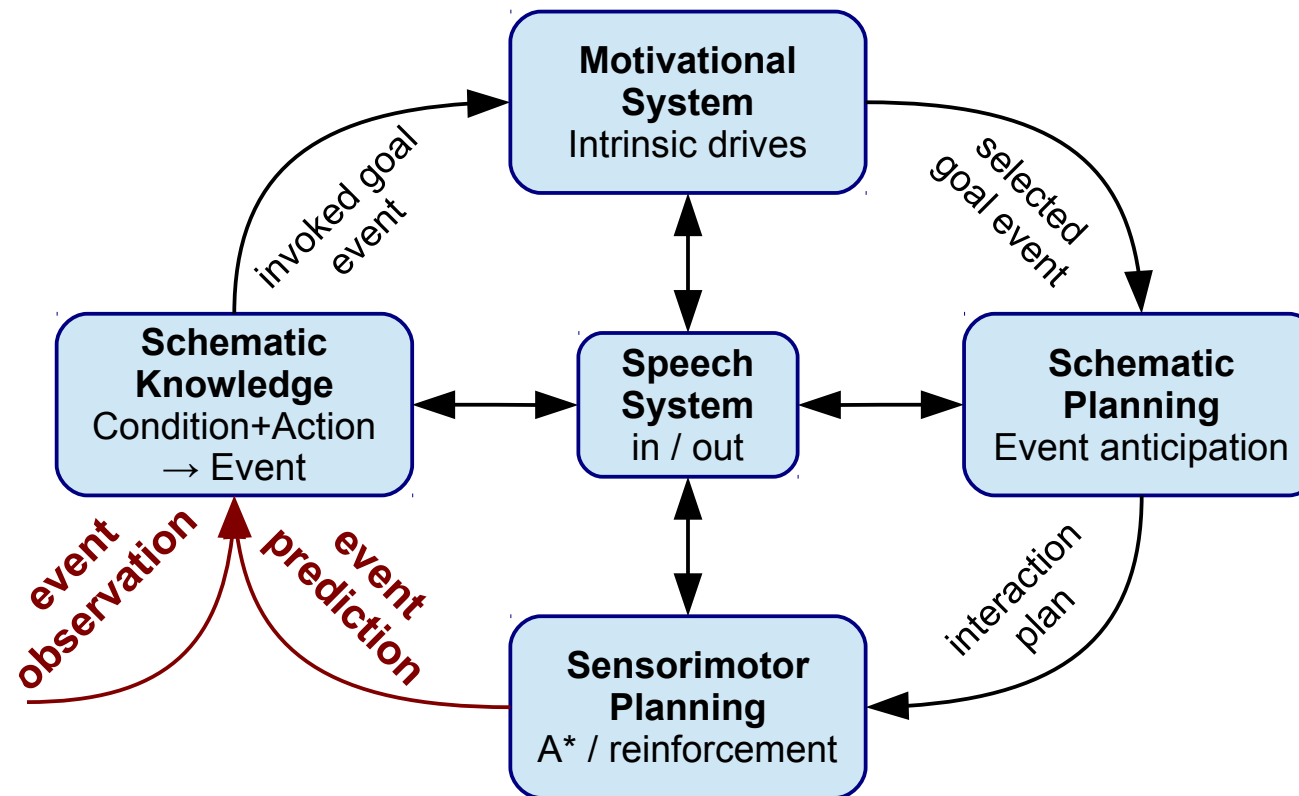
- What if current preconditions are **not fulfilled**?
- Mario simulates a **sequence of interactions** that will **most likely fulfill the preconditions** for the desired goal and invoke the event.
- Technically, this is done by Dijkstra graph search in the schematic knowledge base.
 - First expands (=simulates) nodes (=interactions) with low costs (=high probability of success) until a path from the goal event to the current condition is found.

- Hypothetical example:
Coin increase when being “small”:



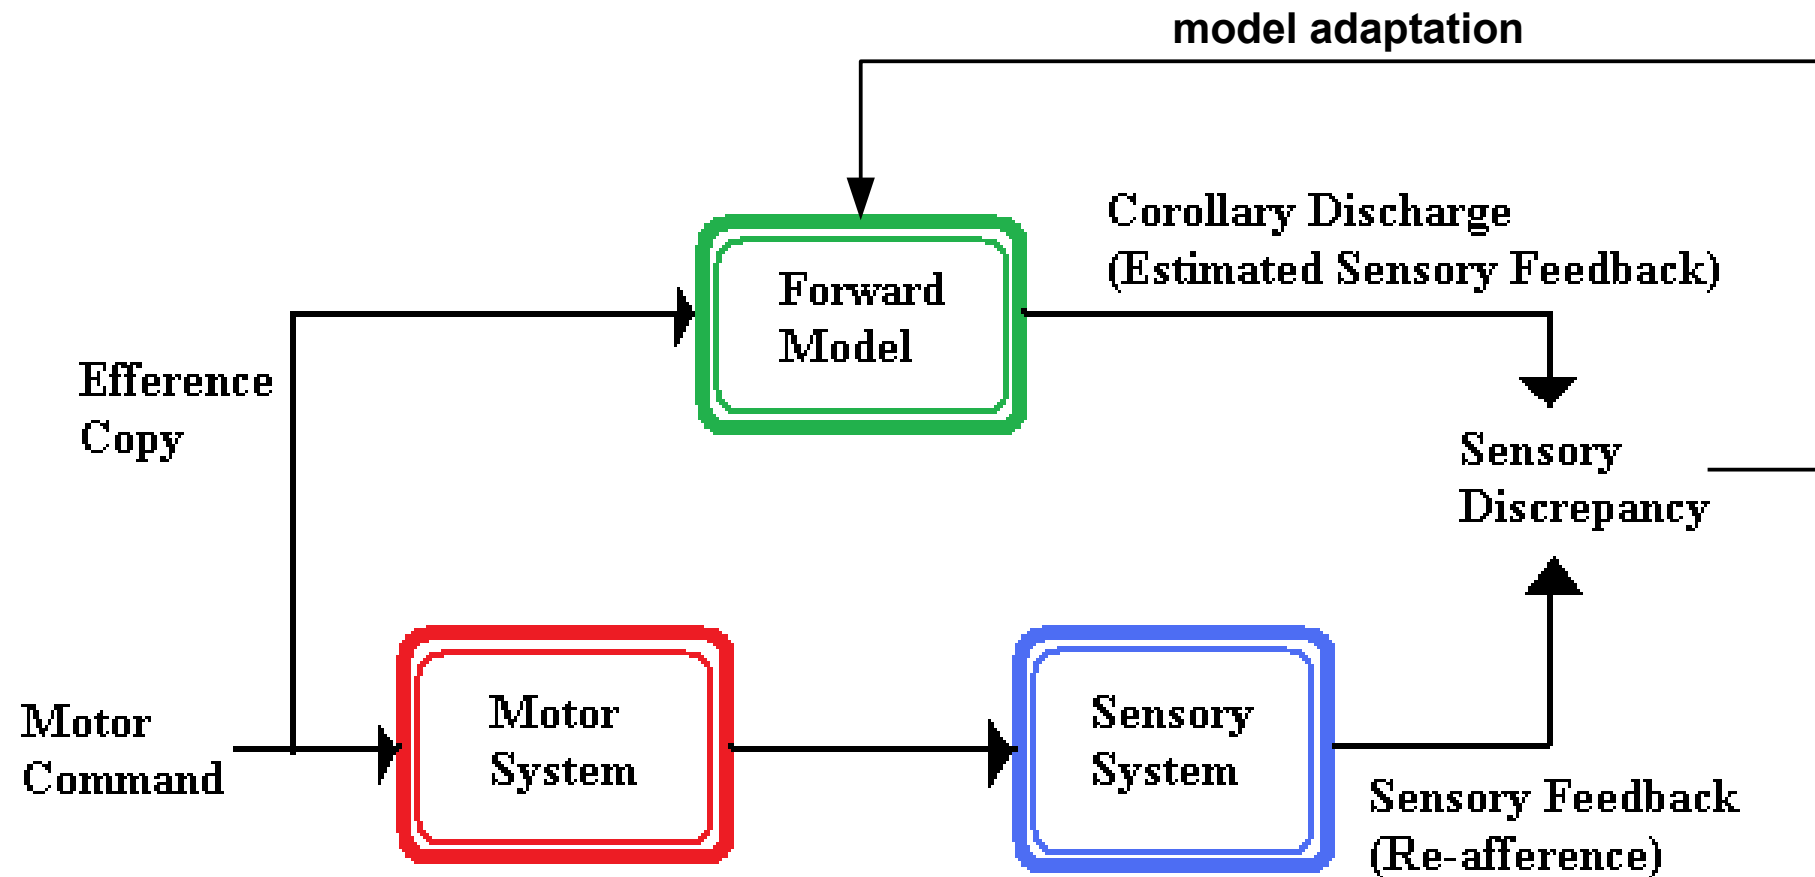
2.3) MarioAI: Learning Schematic Knowledge

- How can such schematic knowledge be **learned**?
- Learning about events can be done by comparing **forward predictions** with the **actual observations** to improve the forward model.



2.3) MarioAI: Learning Schematic Knowledge


- Example from psychology: **Reafference principle** is used to develop a forward model of the effects of own body movements.



2.3) MarioAI: Learning Schematic Knowledge

- Analogously, Mario **continuously predicts** what outcomes (events) his actions will have, given the current (world or bodily) conditions.
- When an occurring event is **unexpected**, a precondition (p) interaction (i) event (e) entry is added to the schematic knowledge (or merged with existing knowledge).
- Bayesian statistics:

$$P(e_a(t)|\vec{p}_b(t), i_c(t)) = \frac{\sum_t e_a(t) \cdot \vec{p}_b(t) \cdot i_c(t)}{\sum_t \vec{p}_b(t) \cdot i_c(t)}$$

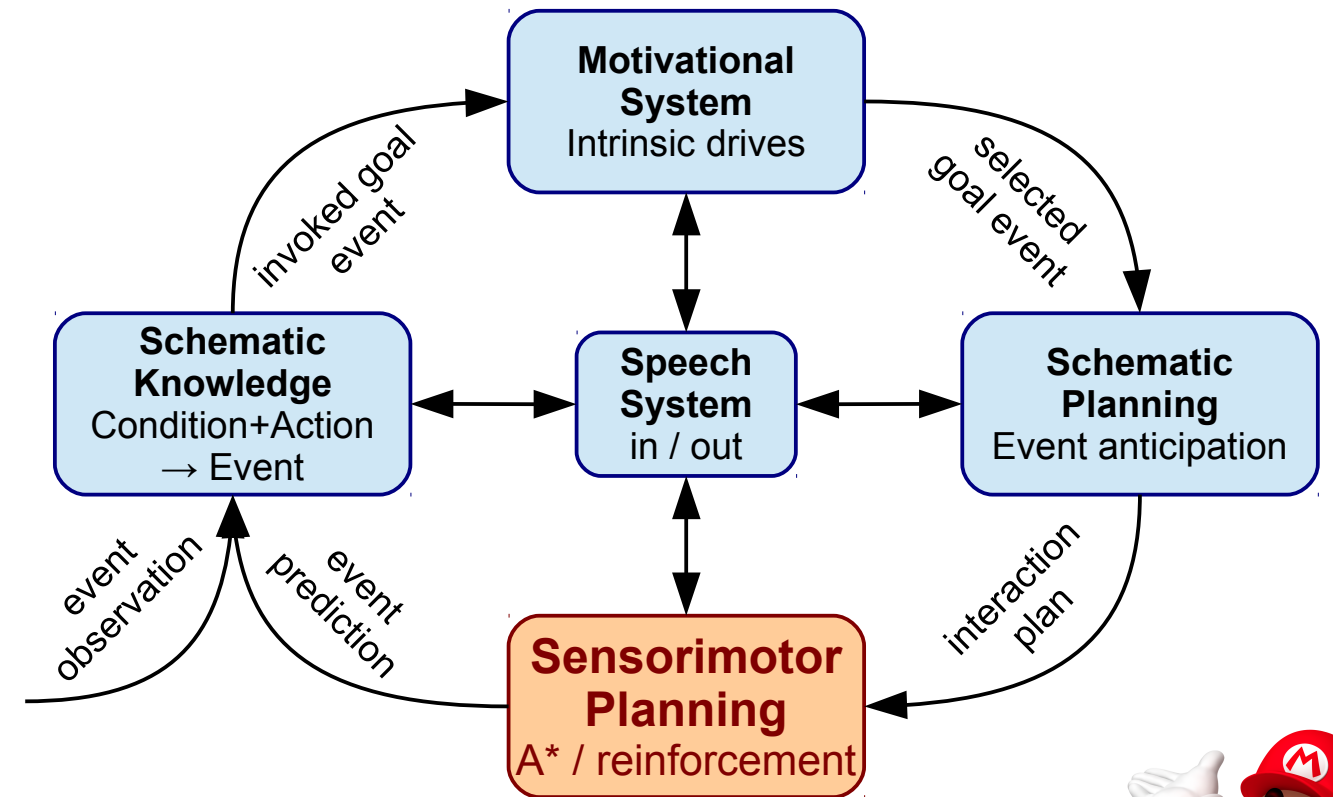



- Note: Determining the **relevance** of preconditions for interactions is not as simple...




2.4) MarioAI: Action Control

- After schematic planning, a **sequence of environmental interactions** that will probably invoke a desired event is available.
- But, what motor commands (here: **keystrokes**) to execute to get at these interactions?
- For this, Mario needs **sensorimotor** planning capabilities.



2.4) MarioAI: Action Control

- Interesting facts: Sensorimotor models develop very early in life, already in the womb!
- What is a newborn capable of?
- **Selective orientation:**
 - Infant reacts with orienting and suckling when touched on the cheek.
 - They do not so, however, when their own hand / arm touches the cheek.
- **Head-eye coordination** is functional (rather well) to follow a stimulus.
- **Hand-to-mouth motion:** mouth opens in anticipation of the hand / the finger.



2.4) MarioAI: Action Control

MarioAI implements **two sensorimotor forward models**.

1) For **path planning**: A deterministic, offline forward simulation of the local game world.

- A* graph search calculates a short path from the current position to the goal position (similar to schematic search)

- First expands nodes (=simulated game states) that will likely lead to the goal with **low cost**.
- Cost estimate is based on **potential field heuristics**.

- Works quite well for deterministic environments...



2) For **object/agent interaction**: Probabilistic, online reinforcement learning

- can handle **uncertain situations** and **plan on-the-fly**,
- is able to actually **learn** from scratch (related to reafference)!

- “TGNG” algor. inspired by hippocampal place cells of rats.


- It learns a **graph structure**, where

- **Nodes** encode **sensory situations**, incl.

- Near obstacles, relative position of relevant objects,
- directional velocity of both agent and target, etc...


- **Edges** encode

- **motor commands** that have been observed to transit from one sensory state to another,
- **transition probabilities** for successor states (Markov process).



2.4) MarioAI: Action Control

- How does this implement sensorimotor **planning**?
- Rewards are **propagated** from goals states over edges and nodes to the current sensory state!
- Example:



R: Reward state (goal interaction)

P: Punishment state (e.g. collision with enemy)

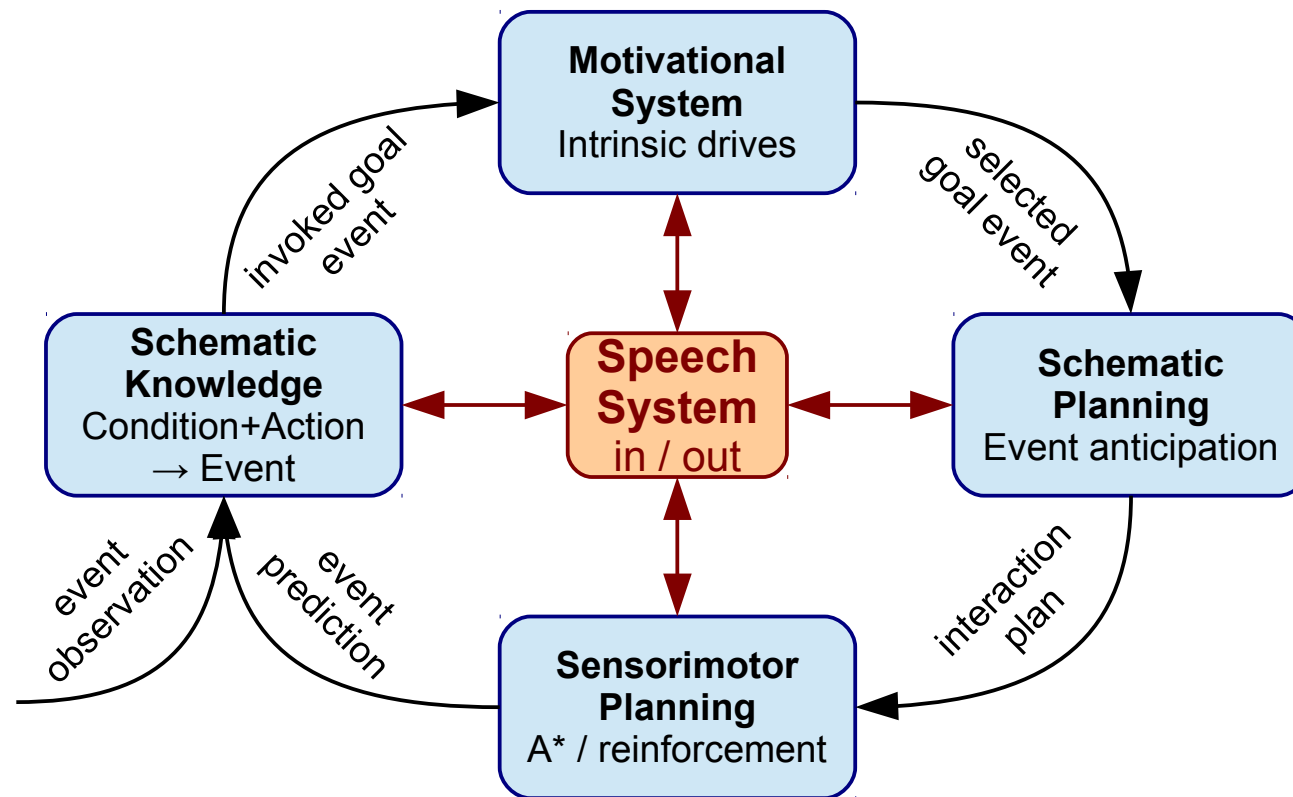
W: Winner node / observed state
(In which situation am I?)

The agent will chose motor command m_2 !



2.5) MarioAI: Speech Comprehension

- Now, Mario can learn how to play the game autonomously!
- Does it make sense to have a game playing itself? ... Not exactly ...
- Via the **speech system**, the user may intervene in all of the processes mentioned beforehand.



2.5) MarioAI: Speech Comprehension


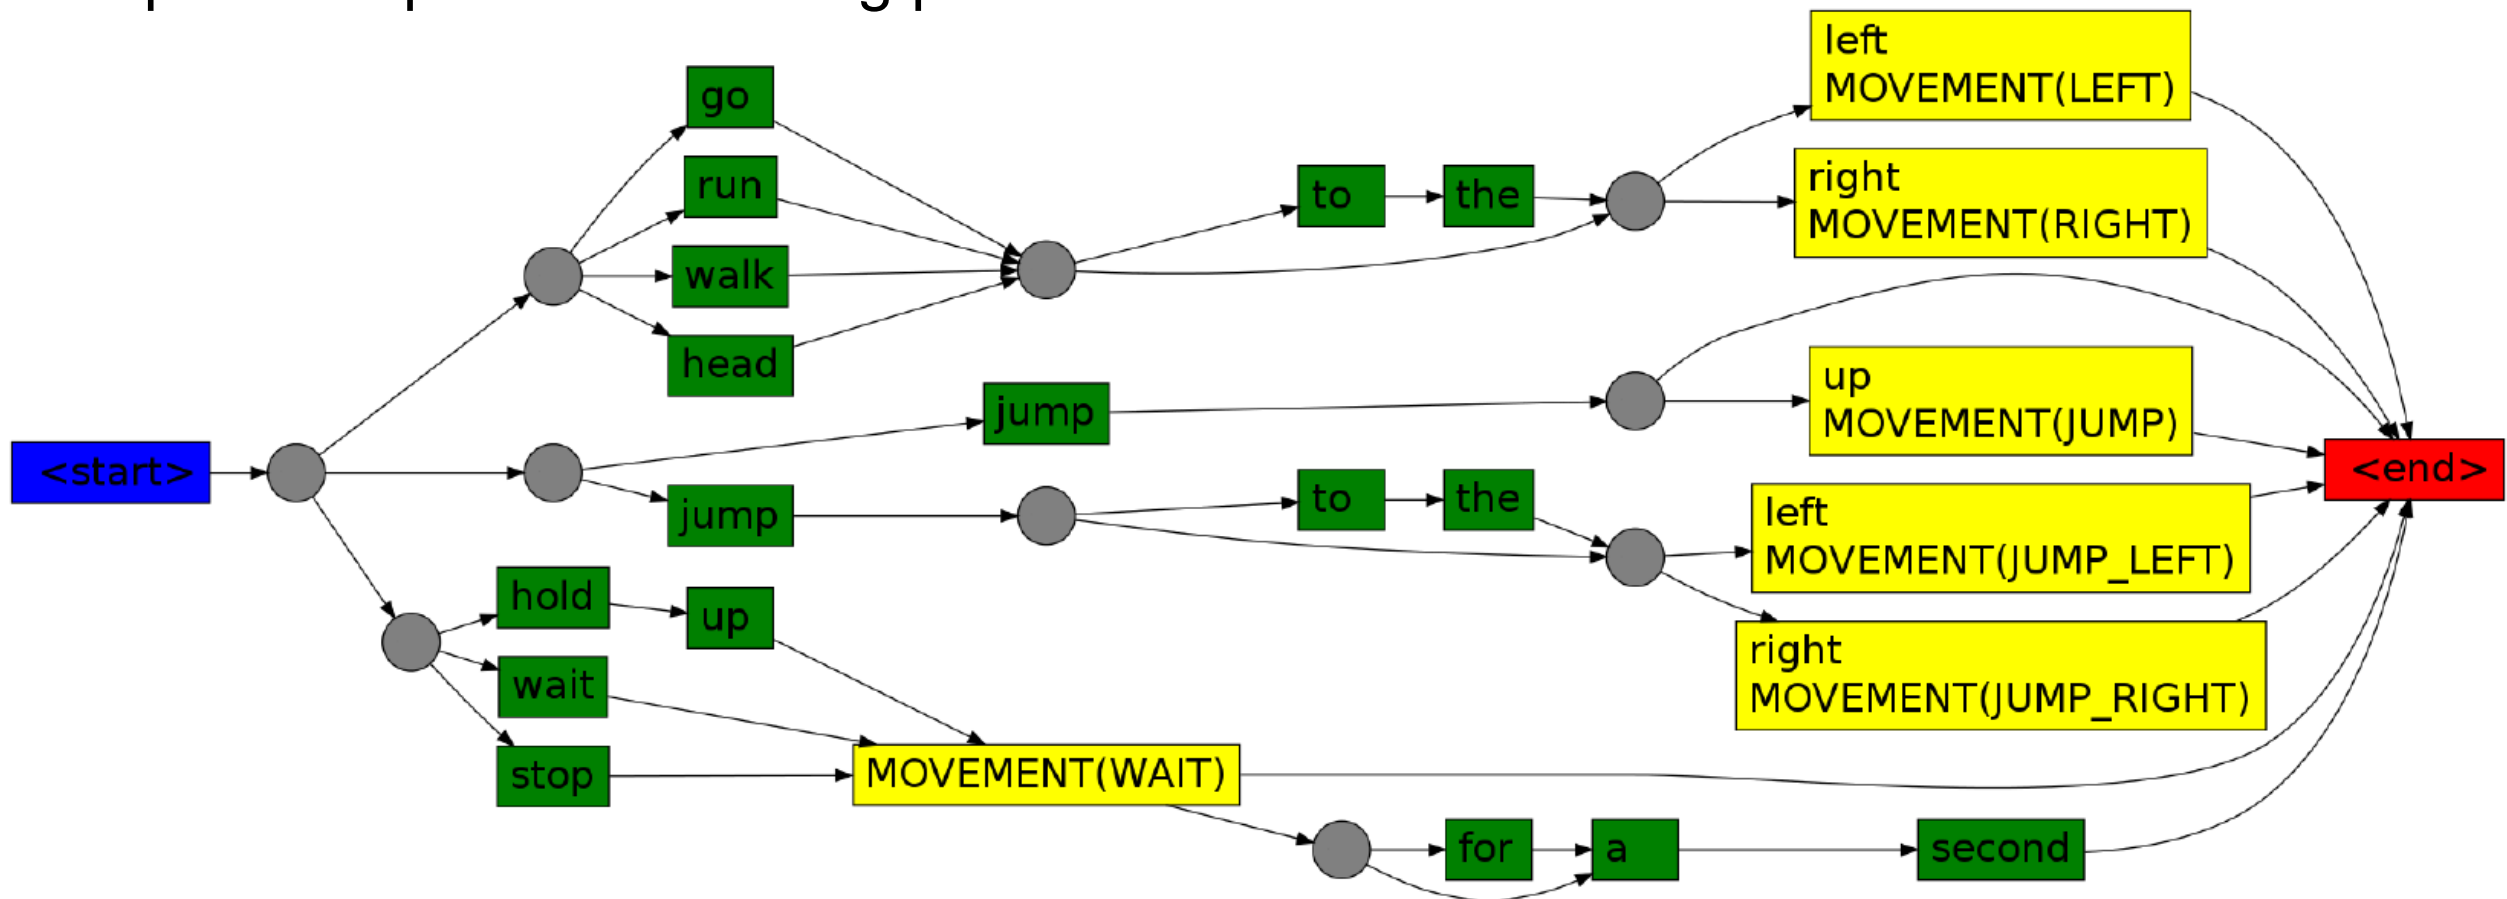
- In humans, **word conceptualizations** can develop when systematically abstracting from behaviorally-grounded, sensorimotor encodings in an **event-oriented** manner.
- In MarioAI, the **schema knowledge** is thus considered as **semantic structure**.
- **Natural language** processing becomes possible (comprehension and generation).
- Technology:
 - **CMU Sphinx-4** (Carnegie Mellon University) for speech recognition.
 - eSpeak (Terminator voice) for speech production.

Powered By
 CMUSphinx



2.5) MarioAI: Speech Comprehension

- Sphinx-4 allows the specification of **tagged context-free grammars** (in Java Speech Grammar Format, JSGF) that specify/restrict the set of utterances that will be recognized.
- Simple example for a resulting parse tree:



2.5) MarioAI: Speech Comprehension

- **Semantics** is identified by **parsing the set of tags** that comes with a recognized sentence.
- Each tag A(B) defines a **keyword A** and **parameters B**.
- **Command type** tag, e.g.:
 - GOAL(...) - instruct Mario to invoke an event or interact with an object
 - RESERVOIR(...) - influence Mario's motivations
 - KNOWLEDGE(...) - Add or query schematic knowledge
 - MOVEMENT(...) - instruct simple movement commands
- **Precondition, interactions and events** tags, e.g.:
 - EVENT(HEALTH_INCREASE)
 - INTERACTION(COLLISION, ABOVE), etc...



2.6) MarioAI: Speech Production

- Then, how can schematic structures be converted to speech?
- Mario uses an tagged output grammar very similar to the input grammar.
 - The **output grammar is inverted** to get a mapping from tags to sentences!
 - Results in a variety of utterances per semantic.
- For example the tag set

ACTOR(MARIO),
INTERACTION(COLLISION, BELOW),
TARGET(QUESTION_MARK),
EVENT(COIN_INCREASE, 0.7)


may lead to the generation of the sentence.....

If I knock my head
with a question mark,
then I will probably
get another coin!



2.6) MarioAI: Social Abilities

- Furthermore, MarioAI implements some **social AI** components:
 - Learning from and about other players
 - Passing goals to agents that are able to accomplish a task
 - Exchanging knowledge between players
 - Teampay scenarios (branch)
- These follow (mainly) from the described techniques.



3) Future Work and Opportunities

- Summarizing, Mario has become a self-motivated, acting, learning, and conversing game agent.
 - Note: The integrated AI components (per se) are not the final answer to all questions. Most of them are kept rather simple (for teaching and illustration).
- Future game agents may model more complex behavioral and learning capabilities, e.g.
 - hierarchical planning and motivations,
 - improved conceptualization abilities (beyond object interactions),
 - understanding spatial relations,
 - episodic memory,
 - language learning, etc.
- Behavioral versatility of scripted AI is inherently limited. **Integrated AI** by means of **cognitive systems** may write their own stories some day!
- **Quo vadis?** So, where are we going? And how?




3) Future Work and Opportunities

- Maybe most important, “integrated” ability for human-like behavior: **Sociability!**
- Humans are able to socially interpret the behavior also of artificial beings!
- Why shouldn't we turn the tables? Game characters could try to understand us and learn from and about us.
 - Think about artificial agents that are able to take others' (also your) perspectives.
 - Theory of mind: “What does my character think or know about me?”
- How can an AI understand e.g. a movement, facial expression or a gesture?
 - If real-world data (e.g. raw visual streams) are to be processed, abstract symbolic representations are not available!
 - Deep neural networks can robustly handle massive amounts of these data and learn codes that may be suitable for interpretation...



3) Future Work and Opportunities

- Example: Embodied, generative, deep neural network (deepRBM) that models the developmental fundamentals of mirror neurons and associated capabilities, like




- **embodied learning of body scheme**
- **spatial perspective-taking**
- recognition of facial expressions
- recognizing point-light motion
- **simulation of movements**
- imitation learning

(Schrodt et al. 2014/15/16)



3) Future Work and Opportunities

- Embodied learning of own bodily movements,
- Inference of the spatial perspective and reenactment of observed movements:



3) Future Work and Opportunities

- 'Imagination' of movements:



- **Open question:** How to generalize the power of deep learning to “ill-posed problems“ (such as life itself), instead of building even more capable expert systems?
 - **Take a look at the human mind!**



Thank you for your attention!

Questions?

Contact:

tobias-fabian.schrodt@uni-tuebingen.de

martin.butz@uni-tuebingen.de

