Artificial Cognitive Systems

David Vernon
Carnegie Mellon University Africa

vernon@cmu.edu
www.vernon.eu

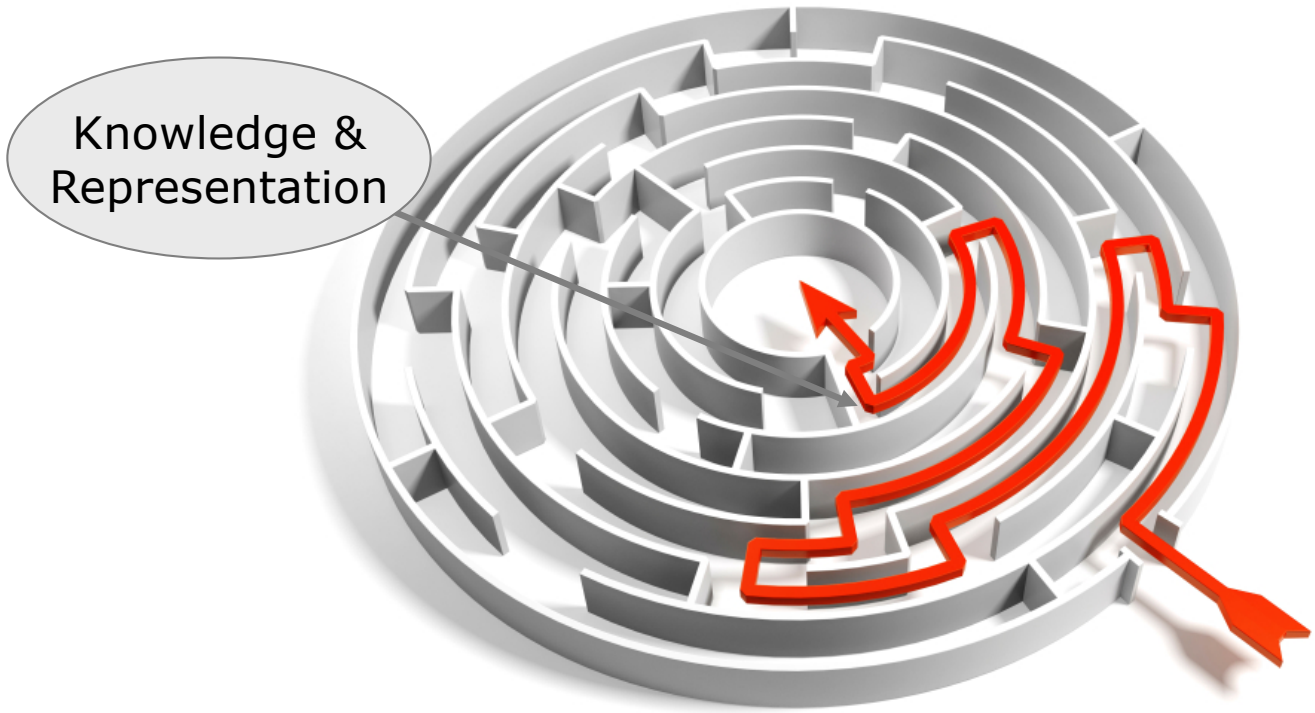
Knowledge &
Representation

Topic Overview

- The duality of memory and knowledge
- Representation and anti-representation
- The symbol grounding problem
- Joint perceptuo-motor representations
- Acquiring and sharing knowledge

Knowledge and Representation

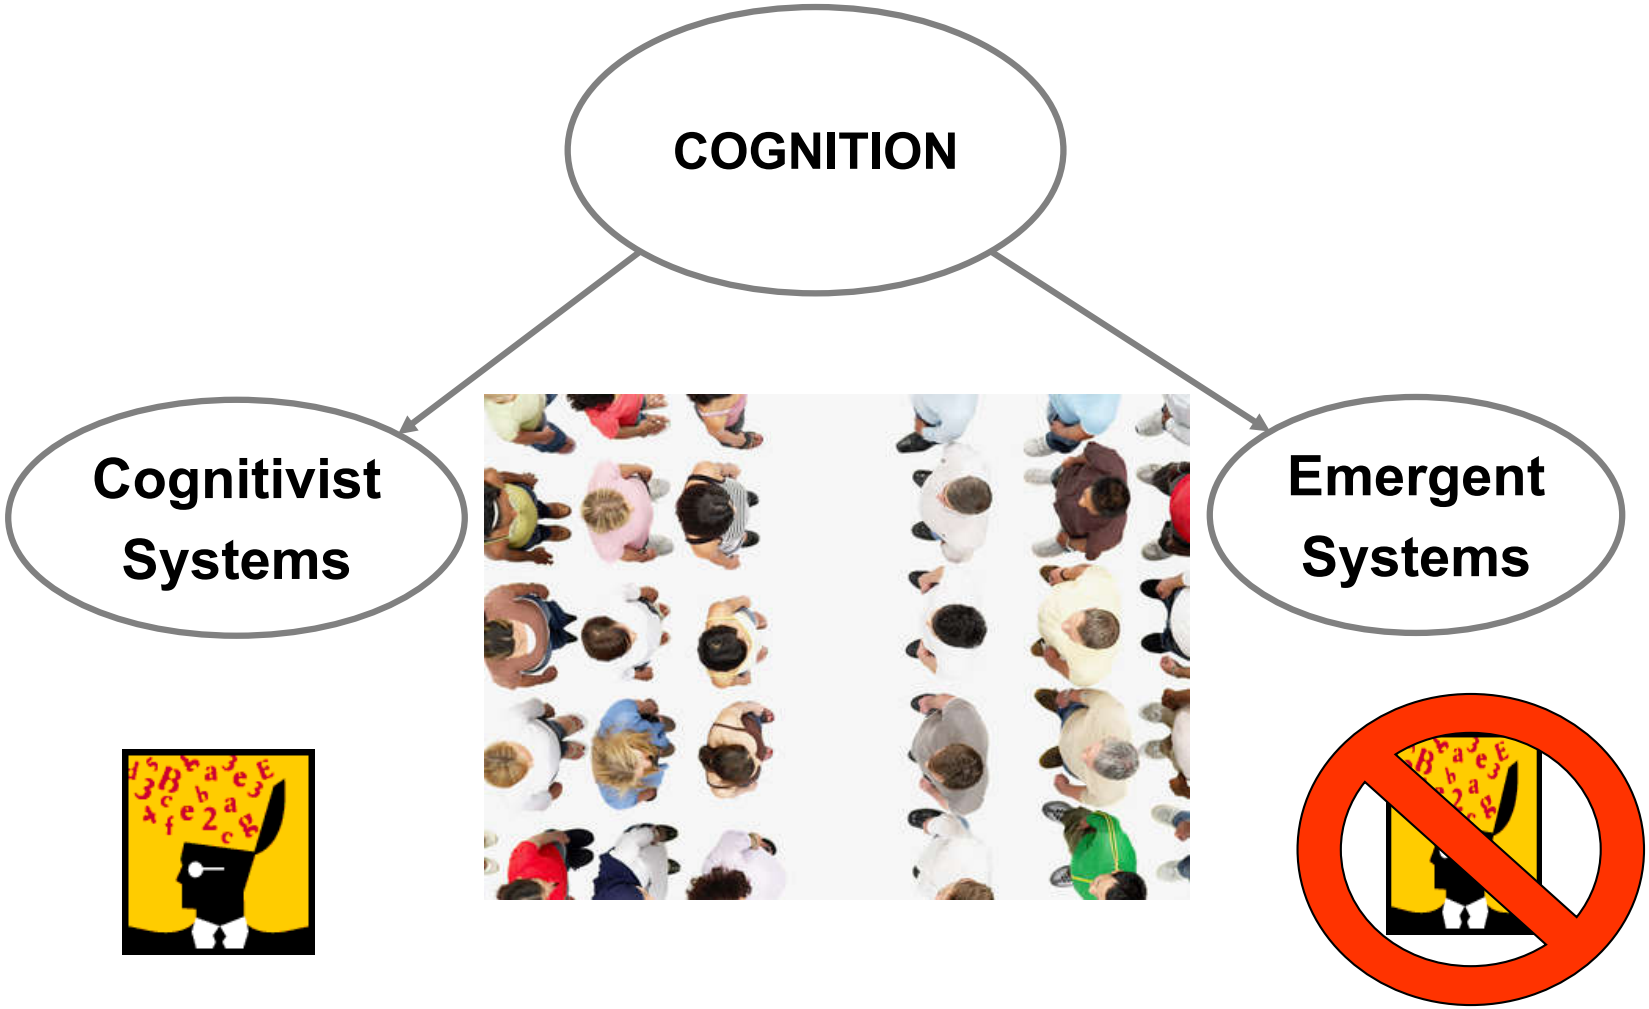

The Duality of Memory and Knowledge

Memory and knowledge are intimately related

- Declarative, procedural, episodic, and semantic knowledge
- But there are hidden assumptions about the nature of that knowledge

The Duality of Memory and Knowledge

Cognitivist paradigm

- Two tacit assumptions

1. Knowledge is the content that complements the cognitive architecture

2. Very often, knowledge, even procedural knowledge, is assumed to be symbolic

The Duality of Memory and Knowledge

Emergent paradigm

- Memory viewed as both content and process: as a mechanism for prediction and recollection
- Knowledge is the manifestation of that process: it is what emerges when memory works effectively
- Knowledge and memory as complementary aspects of the same thing
- Representation of that knowledge: symbolic, non-symbolic (sub-symbolic)

Representation and Anti-Representation

- Representation *vs.* non-representation debate
- *Replacement hypothesis* on embodiment
  - Cognitive system does not have to represent anything because all the information a cognitive system needs is already immediately accessible as a consequence of its non-stop real-time sensorimotor interaction with the world around it
- Arguments against:
  - examples are not "representation hungry"
  - they do not involve situations where the cognitive agent has to act on the basis of knowledge which is *not* presently available to it

Representation and Anti-Representation

Representation and Sharing Knowledge

- Cognitivist approach
  - Representation of the world is a direct one-to-one mapping between an (symbolic) internal state and its counterpart in the real world
  - This mapping is established by perceptual processes
  - Assumes that the things we perceive in the world ***are*** just as we perceive them
  - **All other cognitive systems perceive the world in the same way**
  - **Sharing knowledge among cognitive systems poses no problems in principle**

Representation and Anti-Representation

Representation and Sharing Knowledge

- Emergent approach
  - Perceptions and understanding are fundamentally linked to the manner in which you interact with that reality
  - Your perceptions, and therefore your representations of what it is you are perceiving, are shaped by your actions and the range of possible actions you can perform
  - It is **not** possible for a human designer to implant knowledge directly into an artificial cognitive system
  - Knowledge must be acquired by an embodied cognitive agent by learning

Representation and Anti-Representation

What Qualifies as a Representation?

- One view: any stable state of a cognitive system, and of its memory in particular, that correlates with events in the world is a representation

- Not sufficient condition according to some experts in the field

Representation and Anti-Representation

What Qualifies as a Representation?

- To qualify as a representation these states — these "stand-ins" for the things in the world that are not immediately accessible to the cognitive agent — must also
  - be used for some purpose or function
  - be generally available for such use by the cognitive system
- A representation must play an active causal role in generating the system's behaviour

Representation and Anti-Representation

Weak and Strong Representation

- Weak representations correspond to events that are currently accessible by our senses

- Strong representation correspond to those that are not (e.g. objects that are out of sight or that we saw previously)

  - Required in circumstances where the events to be represented might

    - no longer be present
    - might not even exist
    - might be counter-factual … the opposite of affairs as they appear to be

Representation and Anti-Representation

Radical Constructivism

- The constructive aspect of enactivism is referred to as
  - **Constructivism** [Riegler 2005]
  - **Radical constructivism** [von Glaserfeld 1996, 1996]
- *Radical:* emphasizes that the principles of constructivism have to be applied at every level we chose to describe a cognitive system

Representation and Anti-Representation

Radical Constructivism

- (Radical) constructivism rejects representationalism
  - In the sense that representationalism assumes an external world to which cognitive agents have direct access and can represent
- Constructivism does allow for knowledge
  - The result of an active process of construction whereby the cognitive agent determines through its structural coupling with its environment what matters for its survival and what doesn't
  - Sense-making (enaction)
  - Model generation (computational modelling)

The Symbol Grounding Problem

- Assuming a cognitive system has some form of symbolic representation of the world around it
  - i.e. some set of tokens that denote objects in the agent's world
- **How does the representation, i.e. the symbols, acquire meaning?**

  How do purely symbolic representations acquire semantic content?
- These apparently innocent questions are made difficult by the fact that (physical) symbol systems are governed by purely syntactic processes

The Symbol Grounding Problem

- Physical symbol systems
  - Atomic symbols
  - Strings of symbols
  - Symbol-based rules that define the manipulation and recombination of symbols and strings of symbols
- Defined in terms that make no reference to what these symbols mean

The Symbol Grounding Problem

- But they are all "semantically interpretable"
  - the syntax can be assigned a semantic meaning so that symbols and strings of symbols can represent objects, events, or concepts, and describe them or stand in for them
- The problem is how to assign this meaning
- This is the **symbol grounding problem** [Harnad 1990]

The Symbol Grounding Problem

- Symbolic representations have to be grounded bottom-up in non-symbolic representations of two kinds:

- **Iconic** representations

  - Derived directly from sensory data (e.g. visual imagery, motor imagery)
  - Allow you to discriminate between different objects

- **Categorical** representations

  - Based on the output of both learned and innate processes that detect invariant features of object & event categories from these sensory data (e.g. object affordances, repeated behaviours ... food is nice; dogs bite)

The Symbol Grounding Problem

- Higher-order symbolic representations can then be derived from these elementary symbols
- Both types of representation are non-symbolic
- **A non-symbolic process is required to learn the invariances and thereby form the categories**

The Symbol Grounding Problem

- Usually, we use some form of connectionist approach
  - Create the mapping
  - Form the categorical representation
- As a consequence, according to this argument, a grounded symbol system is a *hybrid system*: a combination of symbolic and emergent approaches (h not H)

The Symbol Grounding Problem

- A restricted form is the **symbol anchoring problem** [Coradeschi and Saffiotti 2003]

- Concerned only with artificial systems

The Symbol Grounding Problem

- Establish a relationship between a symbolic label denoting some object and the sensory perception of that object
- Maintaining that relationship over extended periods of time, even when that object cannot be seen
- Only concerned only with grounding physical objects
  - Doesn't address the grounding of abstract concepts (war or peace)

The Symbol Grounding Problem

- Not everyone agrees with this view of the symbol grounding problem
- An alternative viewpoint
  - Internal symbolic representations are the result of ontogenetic development
  - They are are **tethered** to the world through sensory perception rather than being *grounded*
  - **Symbol tethering** is also known as **Symbol attachment**
  - [Sloman and Chappell 2005]

The Symbol Grounding Problem

An alternative viewpoint

- Symbols don't derive directly from the sensory data (as they do in symbol grounding), they derive from development
  - The process of developing new items of knowledge that are specific to the embodiment of the agent in question

The Symbol Grounding Problem

An alternative viewpoint

- **Symbol grounding**: a cognitivist approach
- **Symbol tethering** is more neutral: makes no strong claims about
  - the relationship between world and representation
  - the necessary uniqueness of these representations
- Representations that *denote* objects *vs.* those that *connote* objects
  - absolute designation vs.
  - convenient association

Joint Perceptuo-motor Representations

- Mental imagery – internal simulation – comprises both visual imagery (or, better still, perceptual imagery) and motor imagery

- These two forms of imagery are tightly entwined

  - the **simulation of perception and covert action** both involve elements of visual and motor imagery

  - Neuro-scientific evidence for the interdependence of perception & action

Joint Perceptuo-motor Representations

Sensory-motor Theory and **Ideo-motor Theory** [Stock & Stock 2004]

- **Sensory-motor action planning**
  - Treats actions as reactive responses to sensory stimuli
  - **Assumes that perception and action use distinct and separate representational frameworks**
  - Builds on the classic uni-directional approach to perception, proceeding stage by stage from stimulus to percept and then to response
  - **Doesn't allow the resultant (or intended) action to impact on the related sensory perception**

Joint Perceptuo-motor Representations

Sensory-motor Theory and **Ideo-motor Theory** [Stock & Stock 2004]

- **Ideo-motor action planning**
  - Treats action as the result of **internally-generated goals**
  - **The selection and control of a particular goal-directed movement depends on the anticipation of the sensory consequence of accomplishing the intended action**
  - The agent images (e.g. through **internal simulation**) the desired outcome and selects the appropriate actions in order to achieve it

Joint Perceptuo-motor Representations

Sensory-motor Theory and **Ideo-motor Theory** [Stock & Stock 2004]

- **Ideo-motor action planning**
  - There is an important difference between the concrete movements comprising an action and the higher-order goals of an action
  - Actors do not voluntarily pre-select the exact movements required to achieve a desired goal
  - **Instead, they select prospectively-guided intention-directed goal-focussed action**

    with the specific movements being adaptively controlled as the action is executed

Joint Perceptuo-motor Representations

Sensory-motor Theory and **Ideo-motor Theory** [Stock & Stock 2004]

- **Ideo-motor action planning**
  - How can the goal, achieved through action, cause the action in the first place?
  - How can the later outcome affect the earlier action?
  - **Prospection**! It is the anticipated goal state, not the achieved goal state, that impacts on the associated planned action
  - **Goal-directed action is a centre-piece of ideo-motor theory**
  - Also referred to as the **goal trigger hypothesis** [Hommel et al. 2001]

Joint Perceptuo-motor Representations

Sensory-motor Theory and **Ideo-motor Theory** [Stock & Stock 2004]

- **Ideo-motor action planning**
  - **Perception and action share a common representational framework**

Joint Perceptuo-motor Representations

The Theory of Event Coding (TEC) [Hommel et al. 2001]

- Concerned with perceptual features but not with how those features are extracted or computed
- Concerned with preparing actions — action planning — but not with the final execution of those actions and the adaptive control of various parts of the agent's body
- Perception, attention, intention, and action all work with a common representation
- Action depends on both external and internal causes

Joint Perceptuo-motor Representations

The Theory of Event Coding (TEC) [Hommel et al. 2001]

Feature codes

Simple (e.g. colour, shape, moving to the left, falling)

Complex (e.g. as an affordance)

Possibly learned

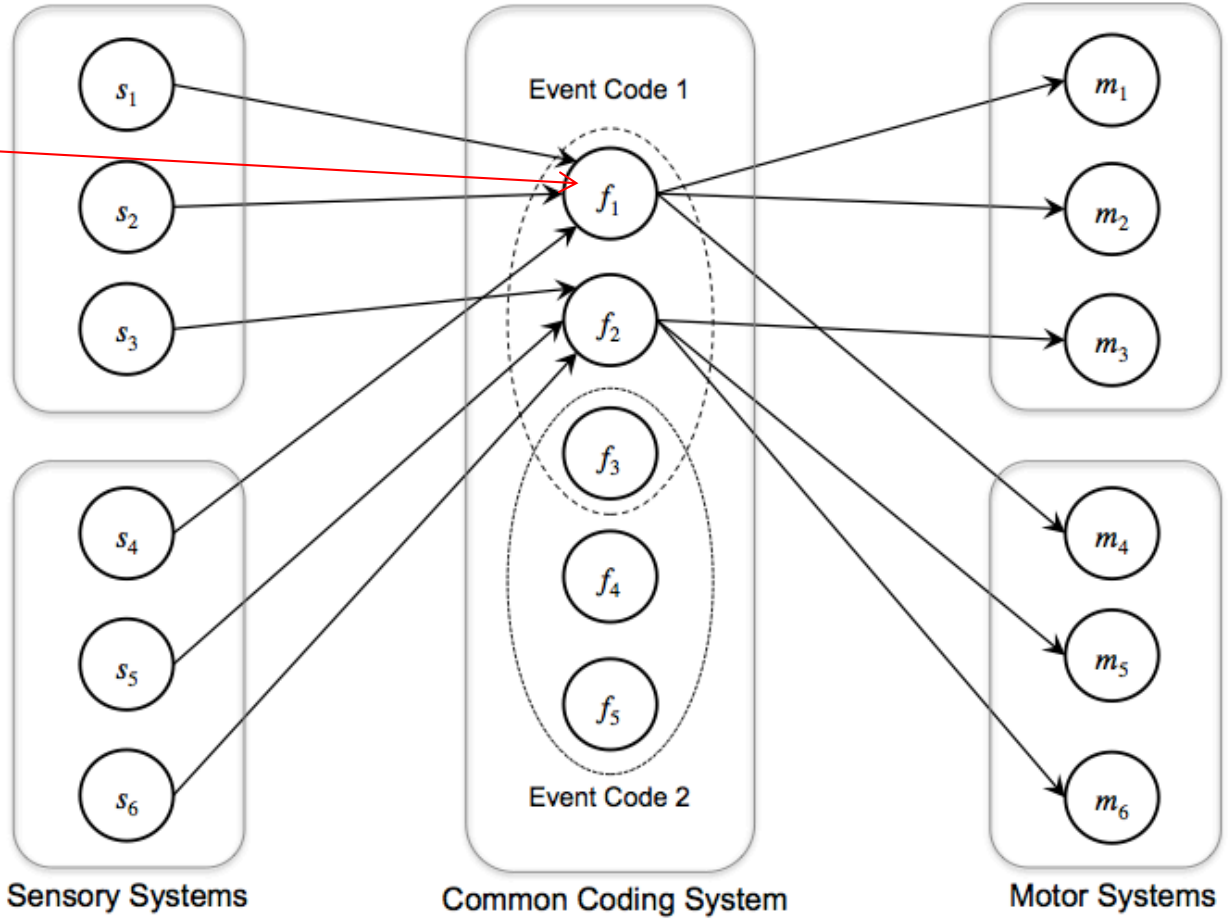

Joint Perceptuo-motor Representations

The Theory of Event Coding (TEC) [Hommel et al. 2001]

- Feature codes associated with an event are activated both when the event is perceived and when it is planned
- Features can be elements of many event codes
  - the activation of a given feature effectively primes, i.e. predisposes, all the other events of which this feature is a component
- The features that make up an event are bound together: integrated into some event code
- The nature of the binding isn't specified in TEC
  - the effect of binding is a form of event code suppression
  - one event inhibits other events that share some of the event codes features

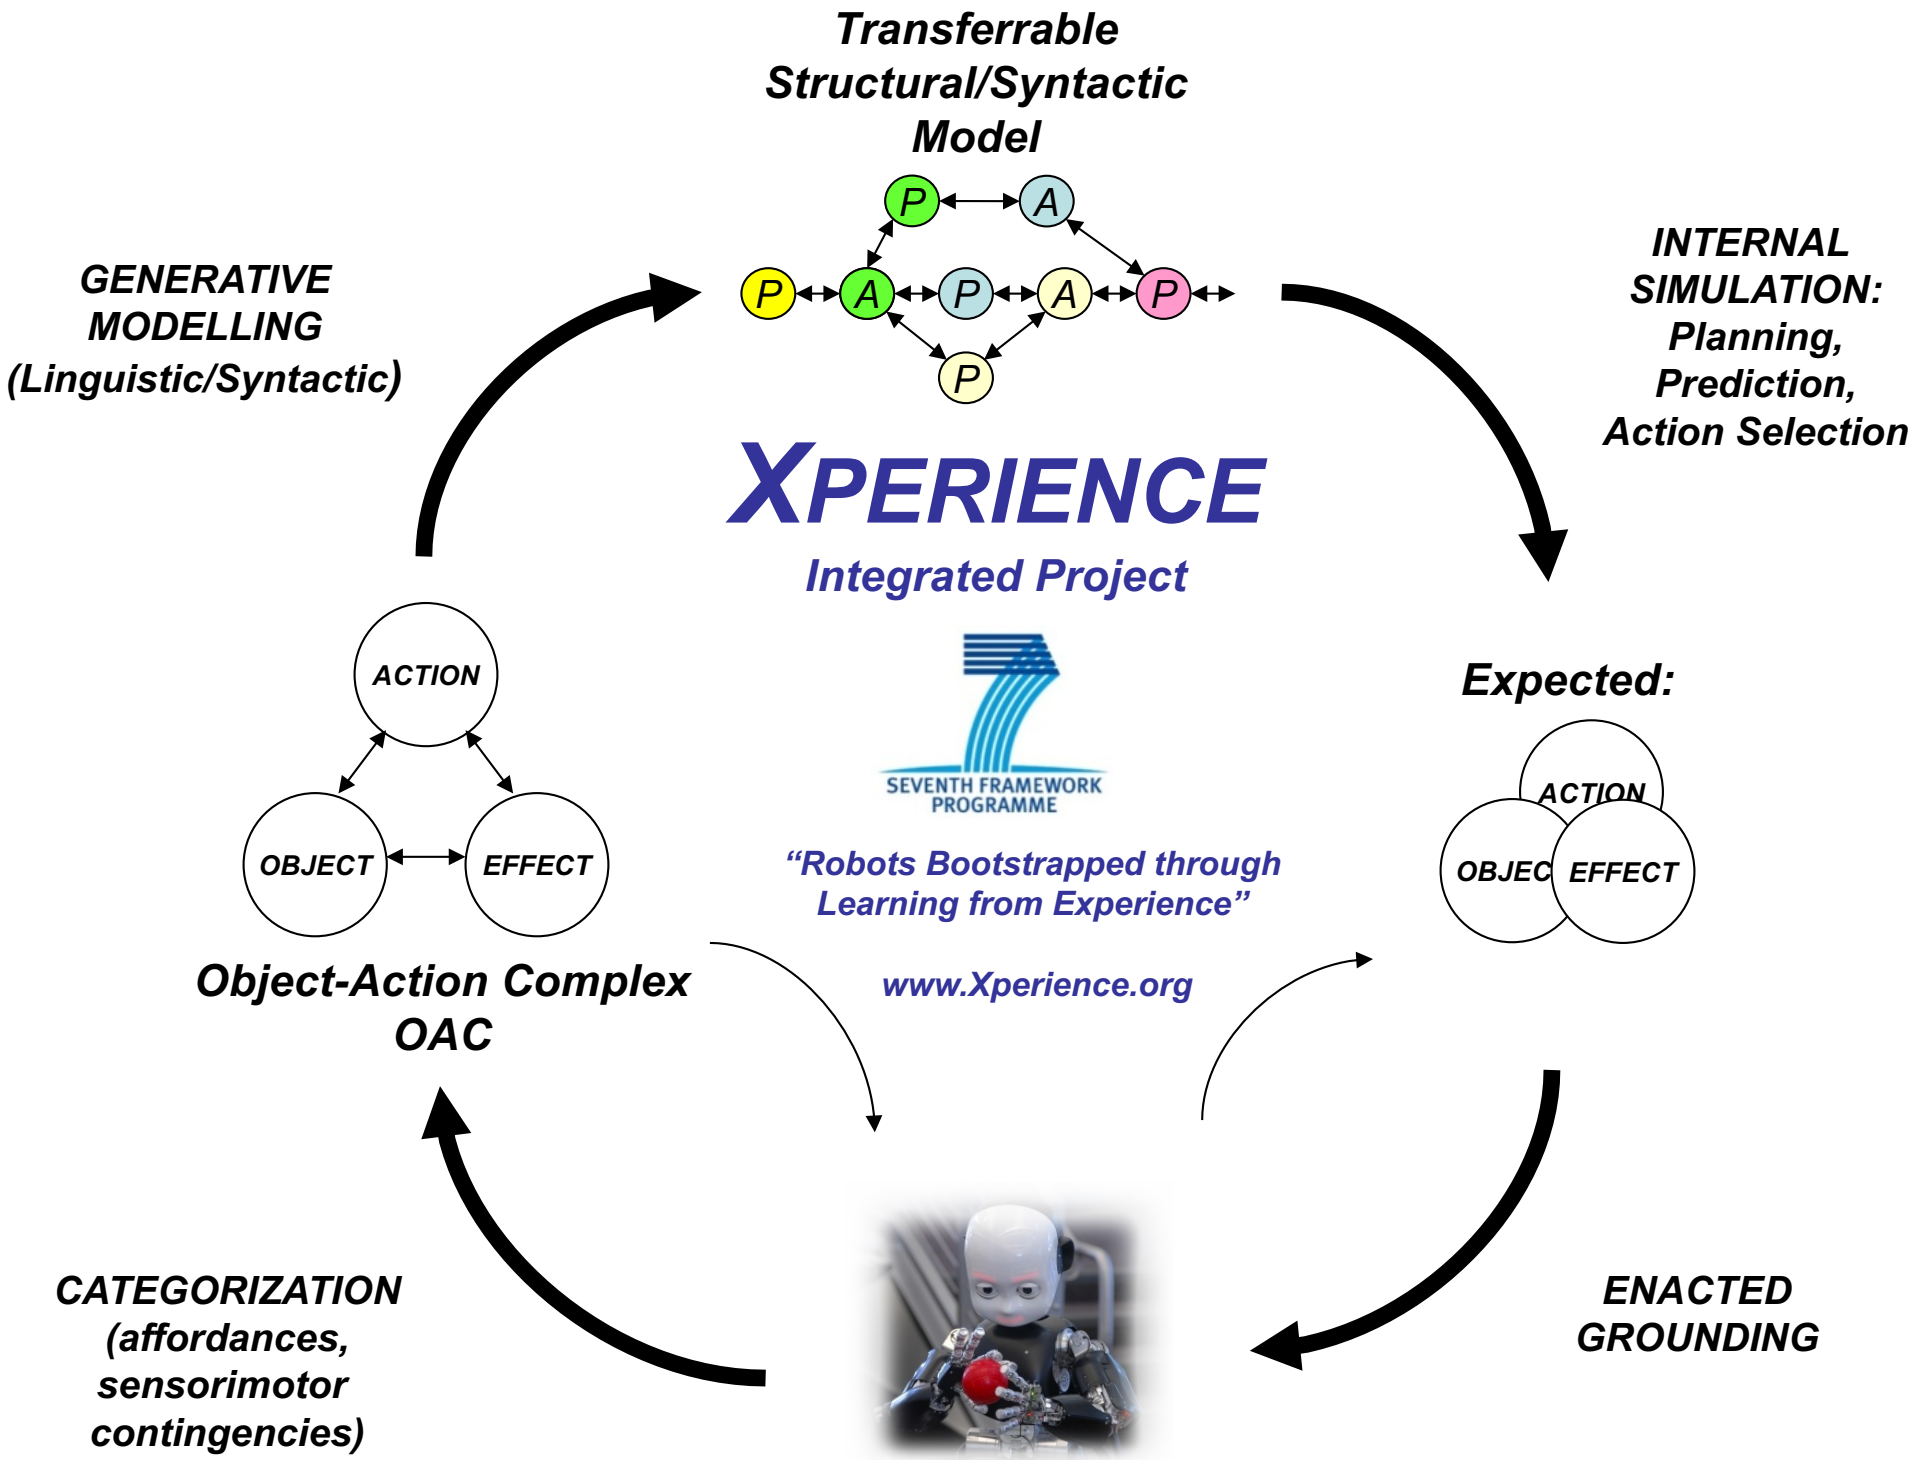
Transferrable Structural/Syntactic Model
P
A
P
A
P
P
GENERATIVE MODELLING (Linguistic/Syntactic)
INTERNAL SIMULATION: Planning, Prediction, Action Selection
XPERIENCE
Integrated Project
SEVENTH FRAMEWORK PROGRAMME
"Robots Bootstrapped through Learning from Experience"
www.Xperience.org
ACTION
OBJECT
EFFECT
Object-Action Complex OAC
Expected:
ACTION
OBJEC
EFFECT
CATEGORIZATION (affordances, sensorimotor contingencies)
ENACTED GROUNDING

Joint Perceptuo-motor Representations

Object-Action Complex, or OAC [Kruger 2011]

- An OAC is a triple, i.e. a unit with three components: (E, T, M)
  - E is an "execution specification"; think of it as an action
  - T is a function that predicts how the attributes that characterize the current state of the agent's world will change if the execution specification is executed
    - Think of T as a prediction of how the agent's perceptions will change as a result of carrying out the actions given by E. S is just the space of all possible perceptions of the agent
  - M is a statistical measure of the success of the OAC's past predictions

Joint Perceptuo-motor Representations

Object-Action Complex, or OAC [Kruger 2011]

- An OAC: a predictor that links current perceived states and future predicted perceived states that would result from carrying out that action
- An OAC models an agent's interaction with the world as it executes some motor program (low-level control program CP in the OAC literature)
- For example, an OAC might encode how to grasp a object or push an object into a given position and orientation (usually referred to as the object pose)
- OACs can be learned and executed, and they can be combined into more complex representations of actions and their perceptual consequences.

Recommended Reading

Vernon, D. *Artificial Cognitive Systems – A Primer*, MIT Press, 2014; Chapter 8.