Radical neural vision for autonomous intelligence: the neural-dynamic perspective

Gregor Schöner
Theory of Cognitive Systems
Institute for Neural Computation
Ruhr-University Bochum, Germany

The old dream of AI was embodied AI

- cognitive architectures linked to sensors and action systems
- the kind of perception that symbolic AI needed was harder than thought
- the kind of motion planning and control that symbolic AI presupposed needed detailed models of everything… and highly-demanding real time computation
- background knowledge was very difficult to capture, especially for physical/perceptual processes

Progress

- motion planning through potential fields and more
- probabilistic thinking enables linking to low-level sensors
- software engineering makes design more efficient
- better computation is an enabler
- recent progress in DL may boost perception

Problems:

- background knowledge is still hard
  - => “every new task 1 PhD thesis”
  - => very little real planning, “thinking”…
- physical interaction is still very hard
  - perhaps the area with least real progress
- no real learning from experience
  - => what we have now does not scale

Organisms as models of embodied AI

- earliest analogies based on low level behaviors, reflexes, fixed action patterns
  - behavior-based robotics
- inspiration from the morphology of organisms…
- inspiration from decentralized control of locomotion, pattern generations…

Neural inspiration for perception

- ideas from ML and deep networks to train perceptual systems for given tasks with given data
  - only beginning to impact on robotics…
- but: does not deliver general "symbols"
  - perception continues to be specific to each task and environment
- difficult to automatically adapt to new tasks and environments

Hypothesis:

- insights from neuroscience and embodied cognition suggest that a new radical form of the analogy with organisms could solve core problems of embodied AI

Hypotheses

- 1) neural cognitive architectures are built on a sensory-motor foundation that incorporates background knowledge
- 2) these architectures generate classes of naturalistic behavior/naturalistic tasks (rather than solving general computational tasks)
- 3) cognition emerges from sensory-motor behaviors through increasing invariance and complexity while retaining sensory-motor grounding (rather than through symbol manipulation)
- 4) neural cognitive architectures provide the process foundation for autonomous learning from experience

The need for neural dynamics

- the forward NN capture only a small portion of (visual) cognition.. most cognition occurs without continuous sensory input
- => activation needs to be kept stable/be generated autonomously
- requires neural interaction
    - excitatory coupling among neurons belonging to same state
    - inhibitory coupling among neurons belonging to competing states

[Schöner ToPiCs (in press)]

Dynamic Field Theory (DFT)

- neural dynamics in low-dimensional spaces
- => stable activation states as units of representation

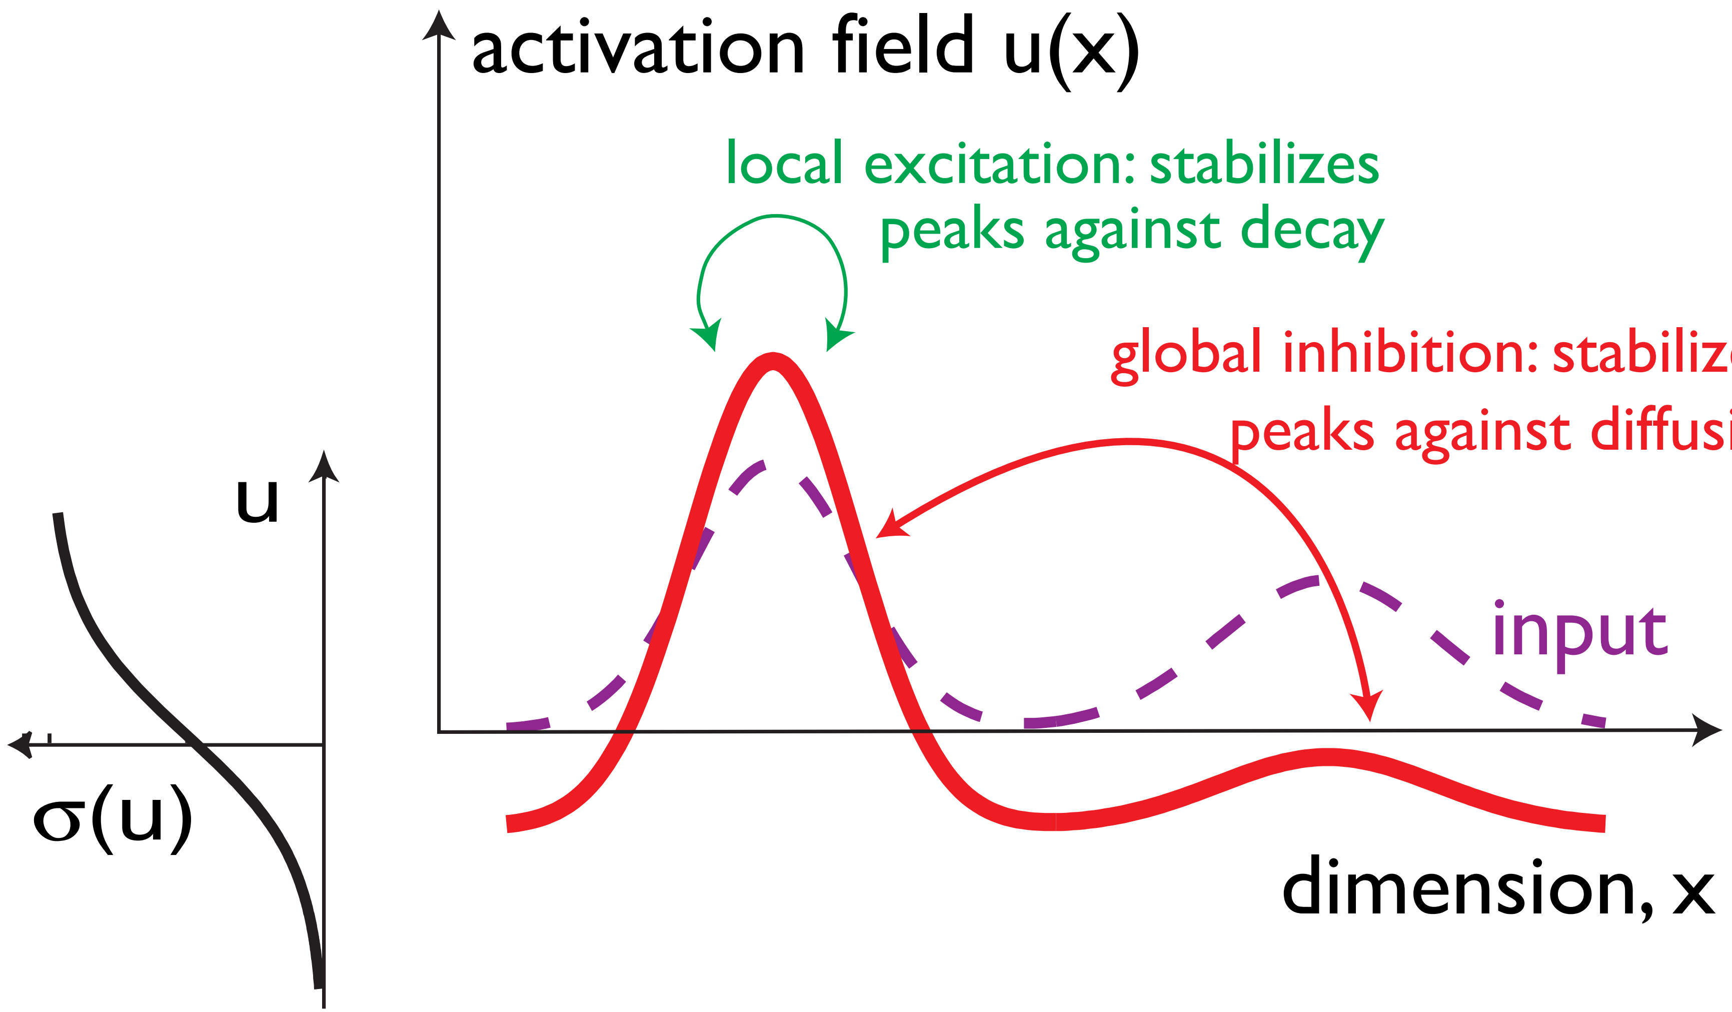

Dynamic Field Theory (DFT)

- instabilities enable elementary cognitive function: detection, selection, working memory, recall
- => simulation

www.dynamicfieldtheory.org

- online lectures
- Matlab toolbox
  Cosivina
- graphical DFT programming tool
  Cedar

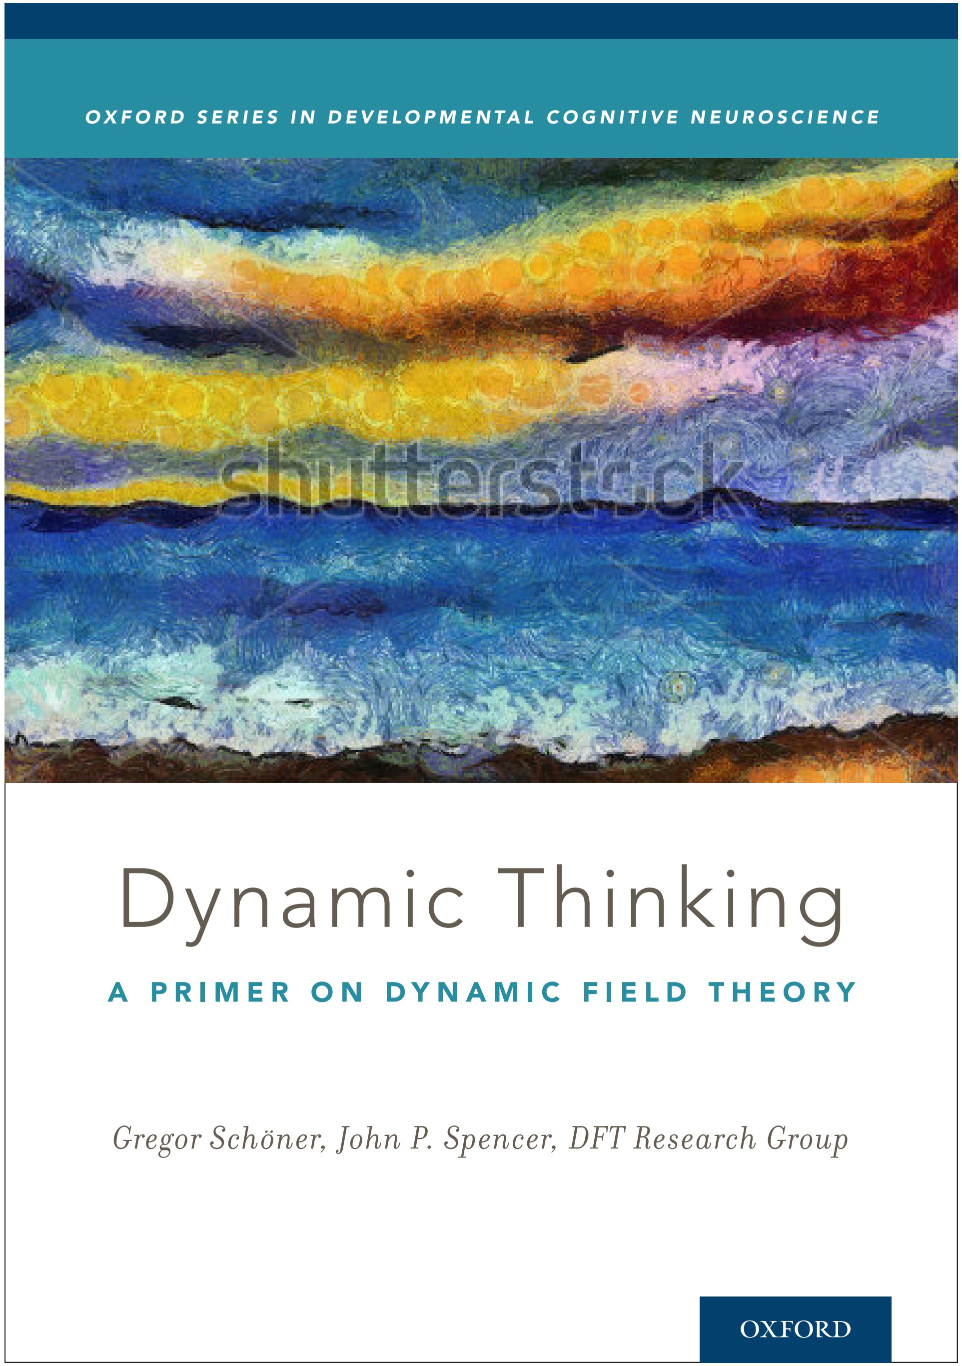

Sequences in neural dynamics

- generation of sequential transitions between units of representation is the core processing element needed to extend neural dynamics toward cognition

Sequences in neural dynamics

- intentional states
- their "condition of satisfaction" (Cos)
- may originate from sensory signals or other neural activation variables/fields inside an architecture....

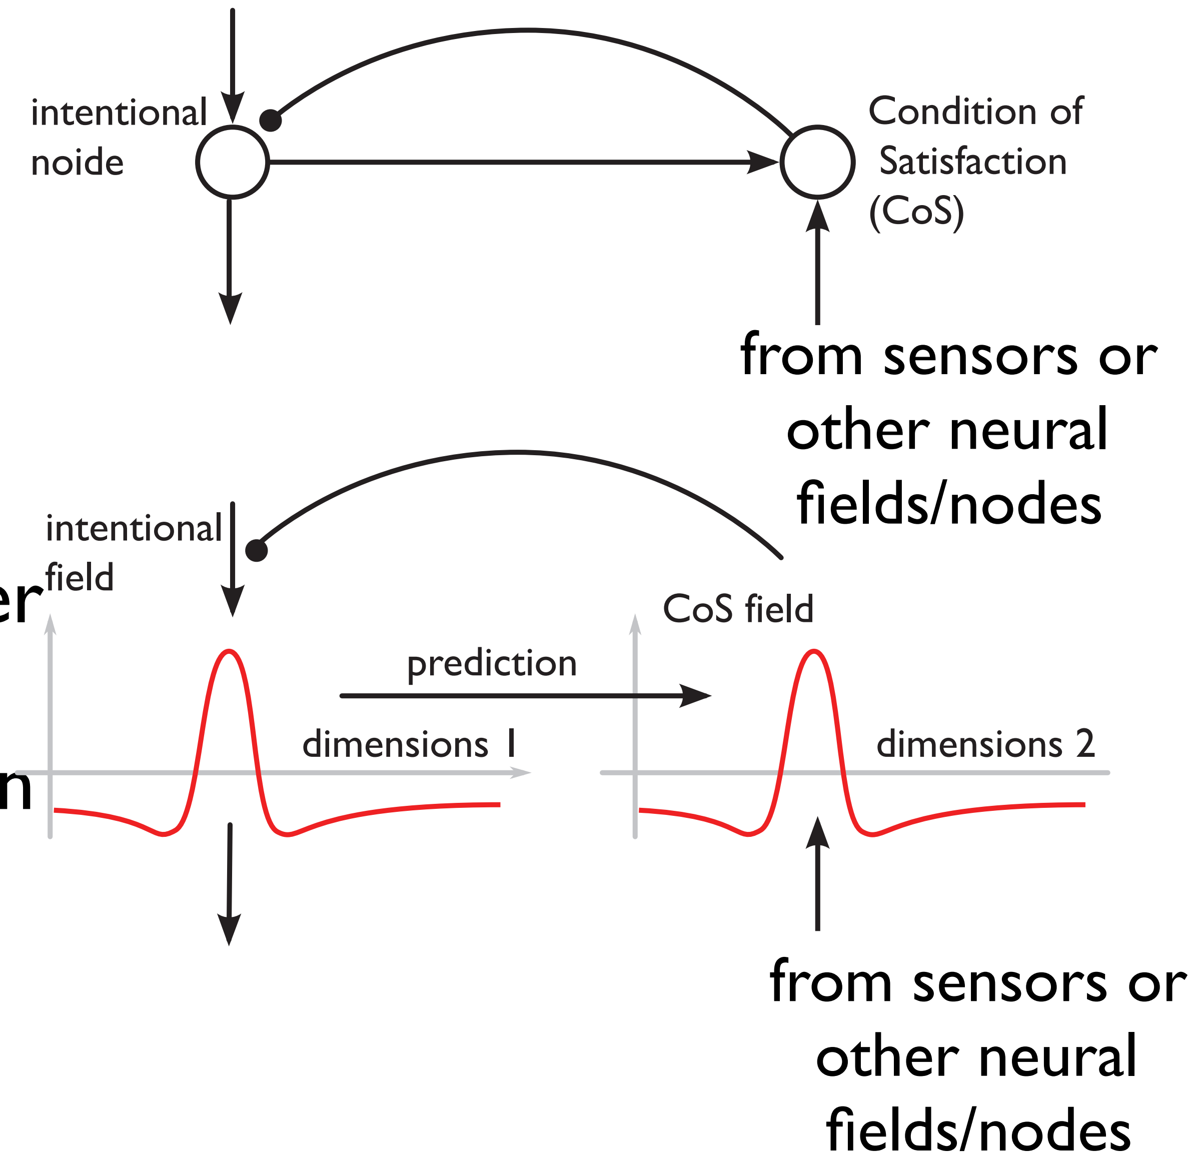

Neural dynamic architectures

- are enabled by stability which implies robustness: elements retain their function as they are coupled
- coupling fields of different dimensions offers new functions
  - cued search
  - coordinate transforms
  - peak detection

Neural dynamic architectures

- => networks of neural dynamic systems
  - not conventional cognitive architectures: no information is passed on, no symbols are manipulated..
- Examples
  - movement
  - scene representation
  - serial order
  - perceptual grounding of relations
  - mental map formation
  - intentional systems

- serial order from demonstration
- sequence of pointing movements

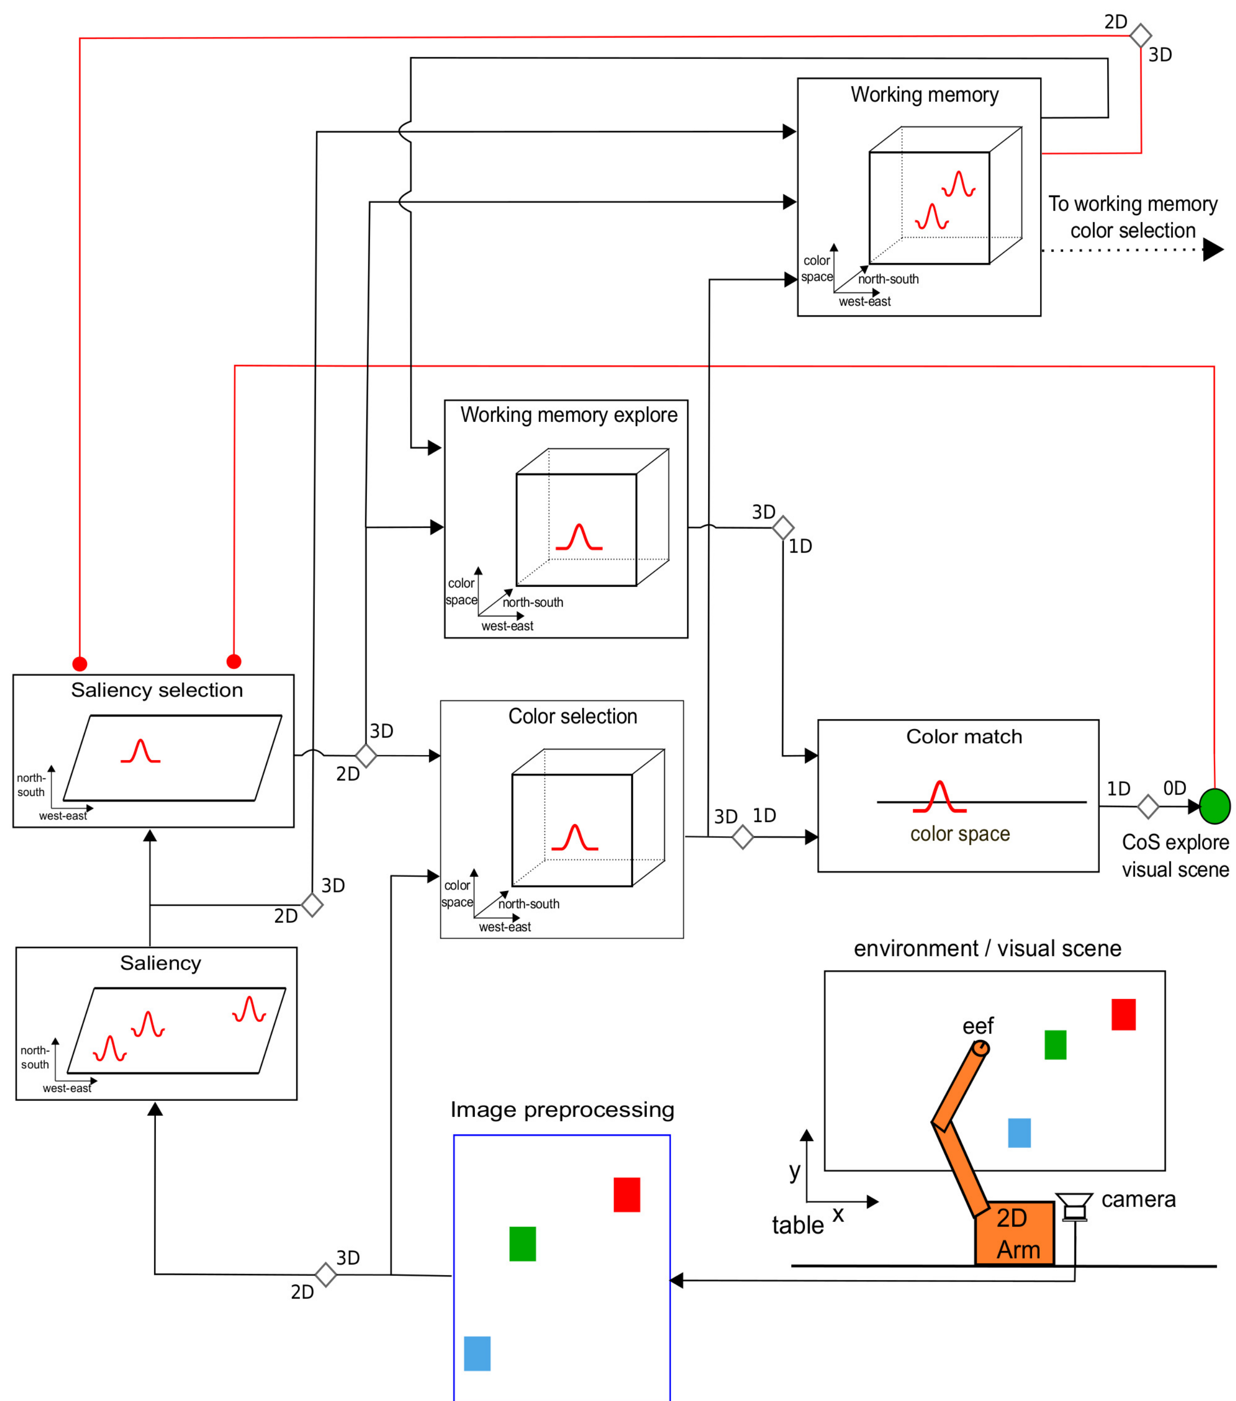

[Tekülve et al., Frontiers (under review)]

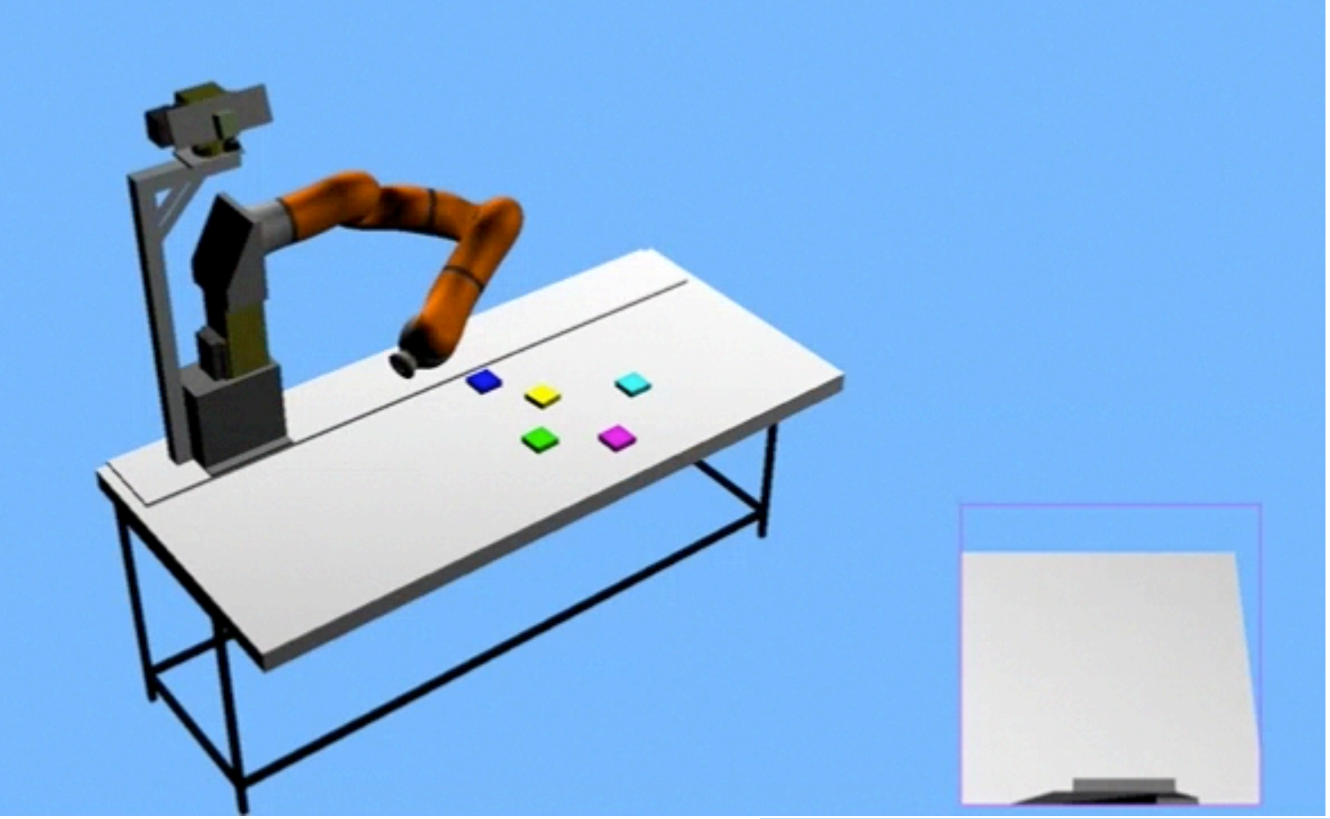

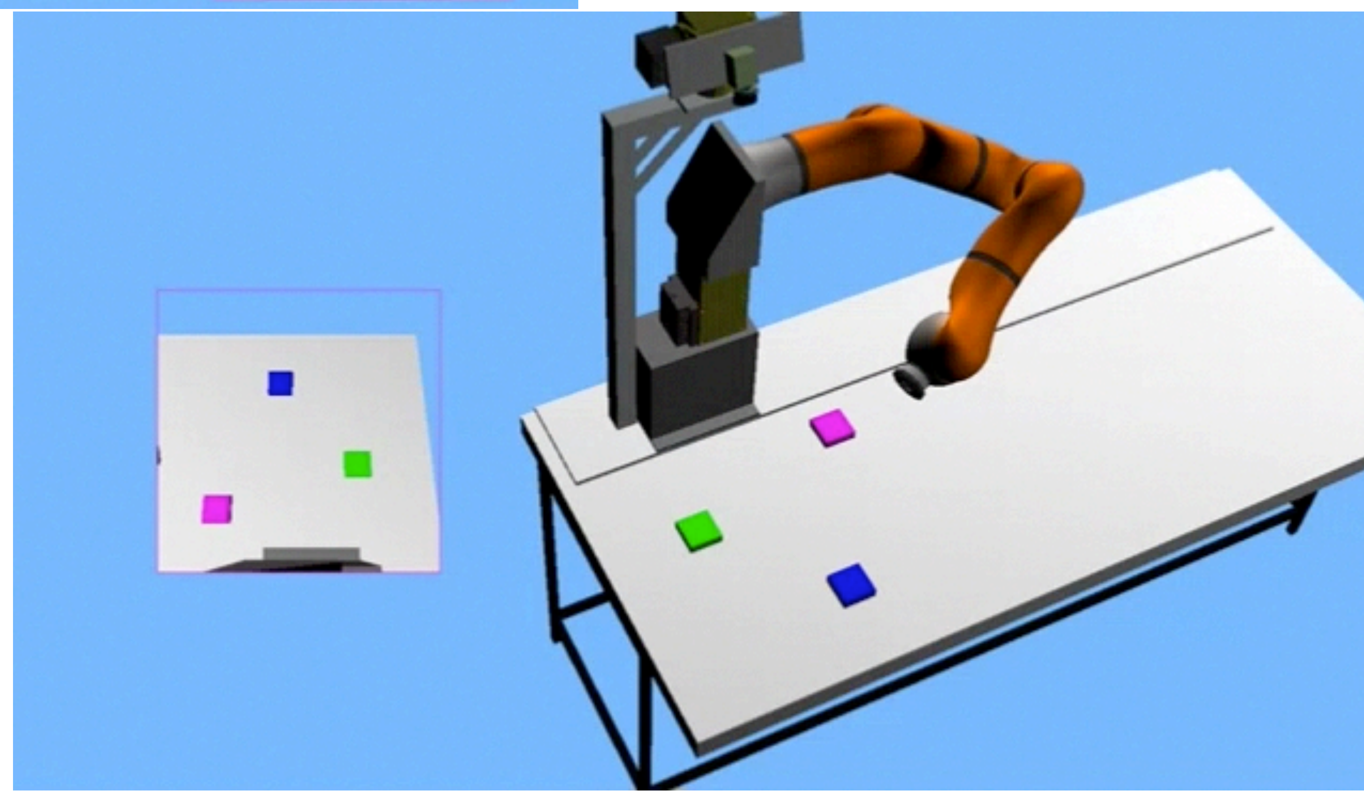

Feed-forward feature extraction

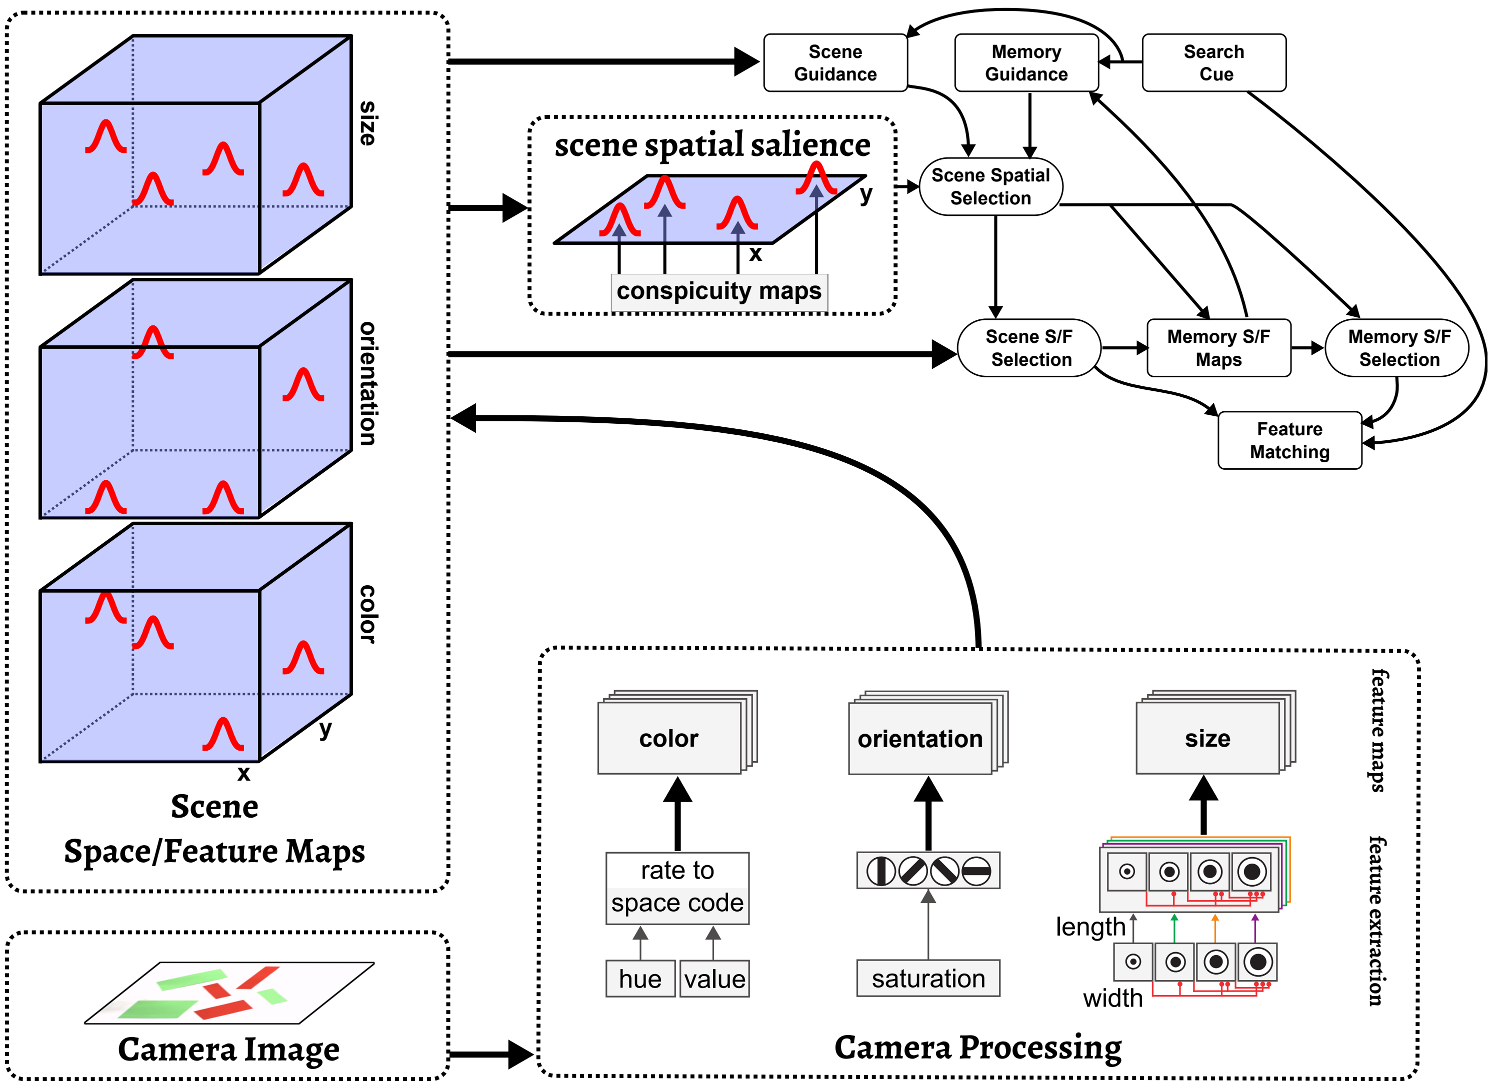

[Grieben et al (under revision)]

Scene representation and visual search

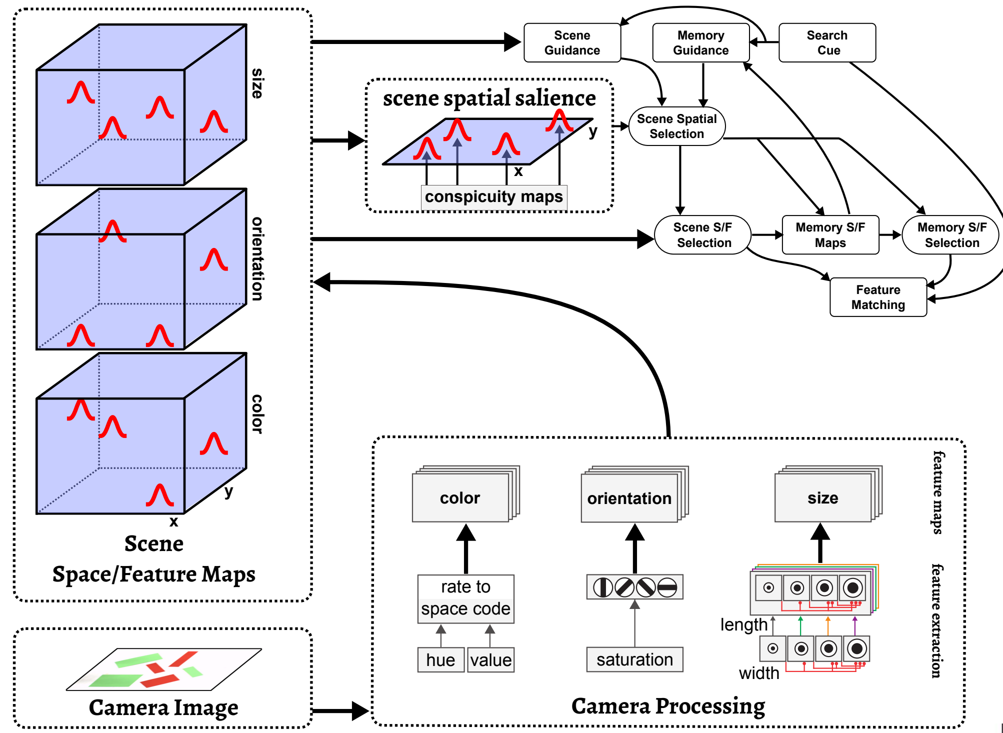

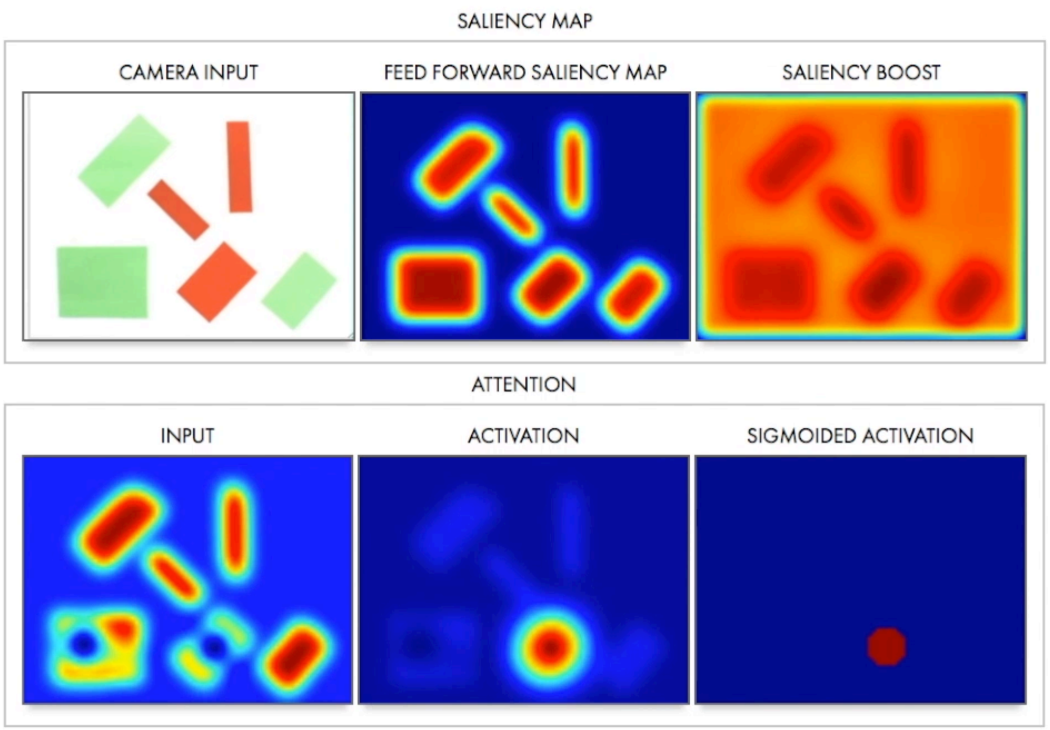

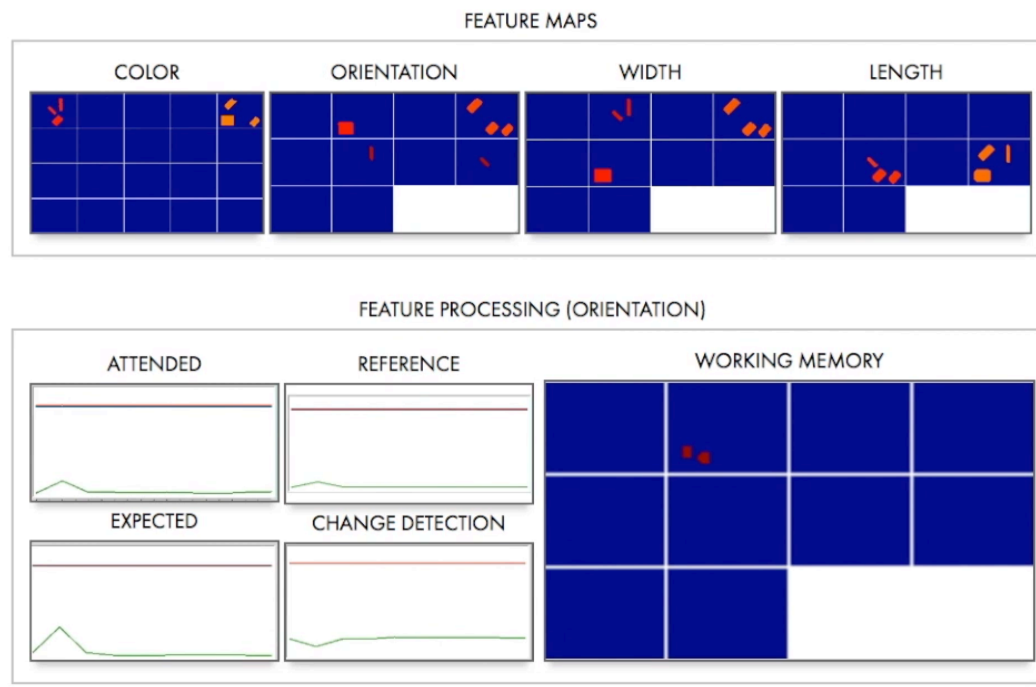

[Grieben et al., Attention, Perception, & Psychophysics (under revision)]

Perceptual grounding of relations

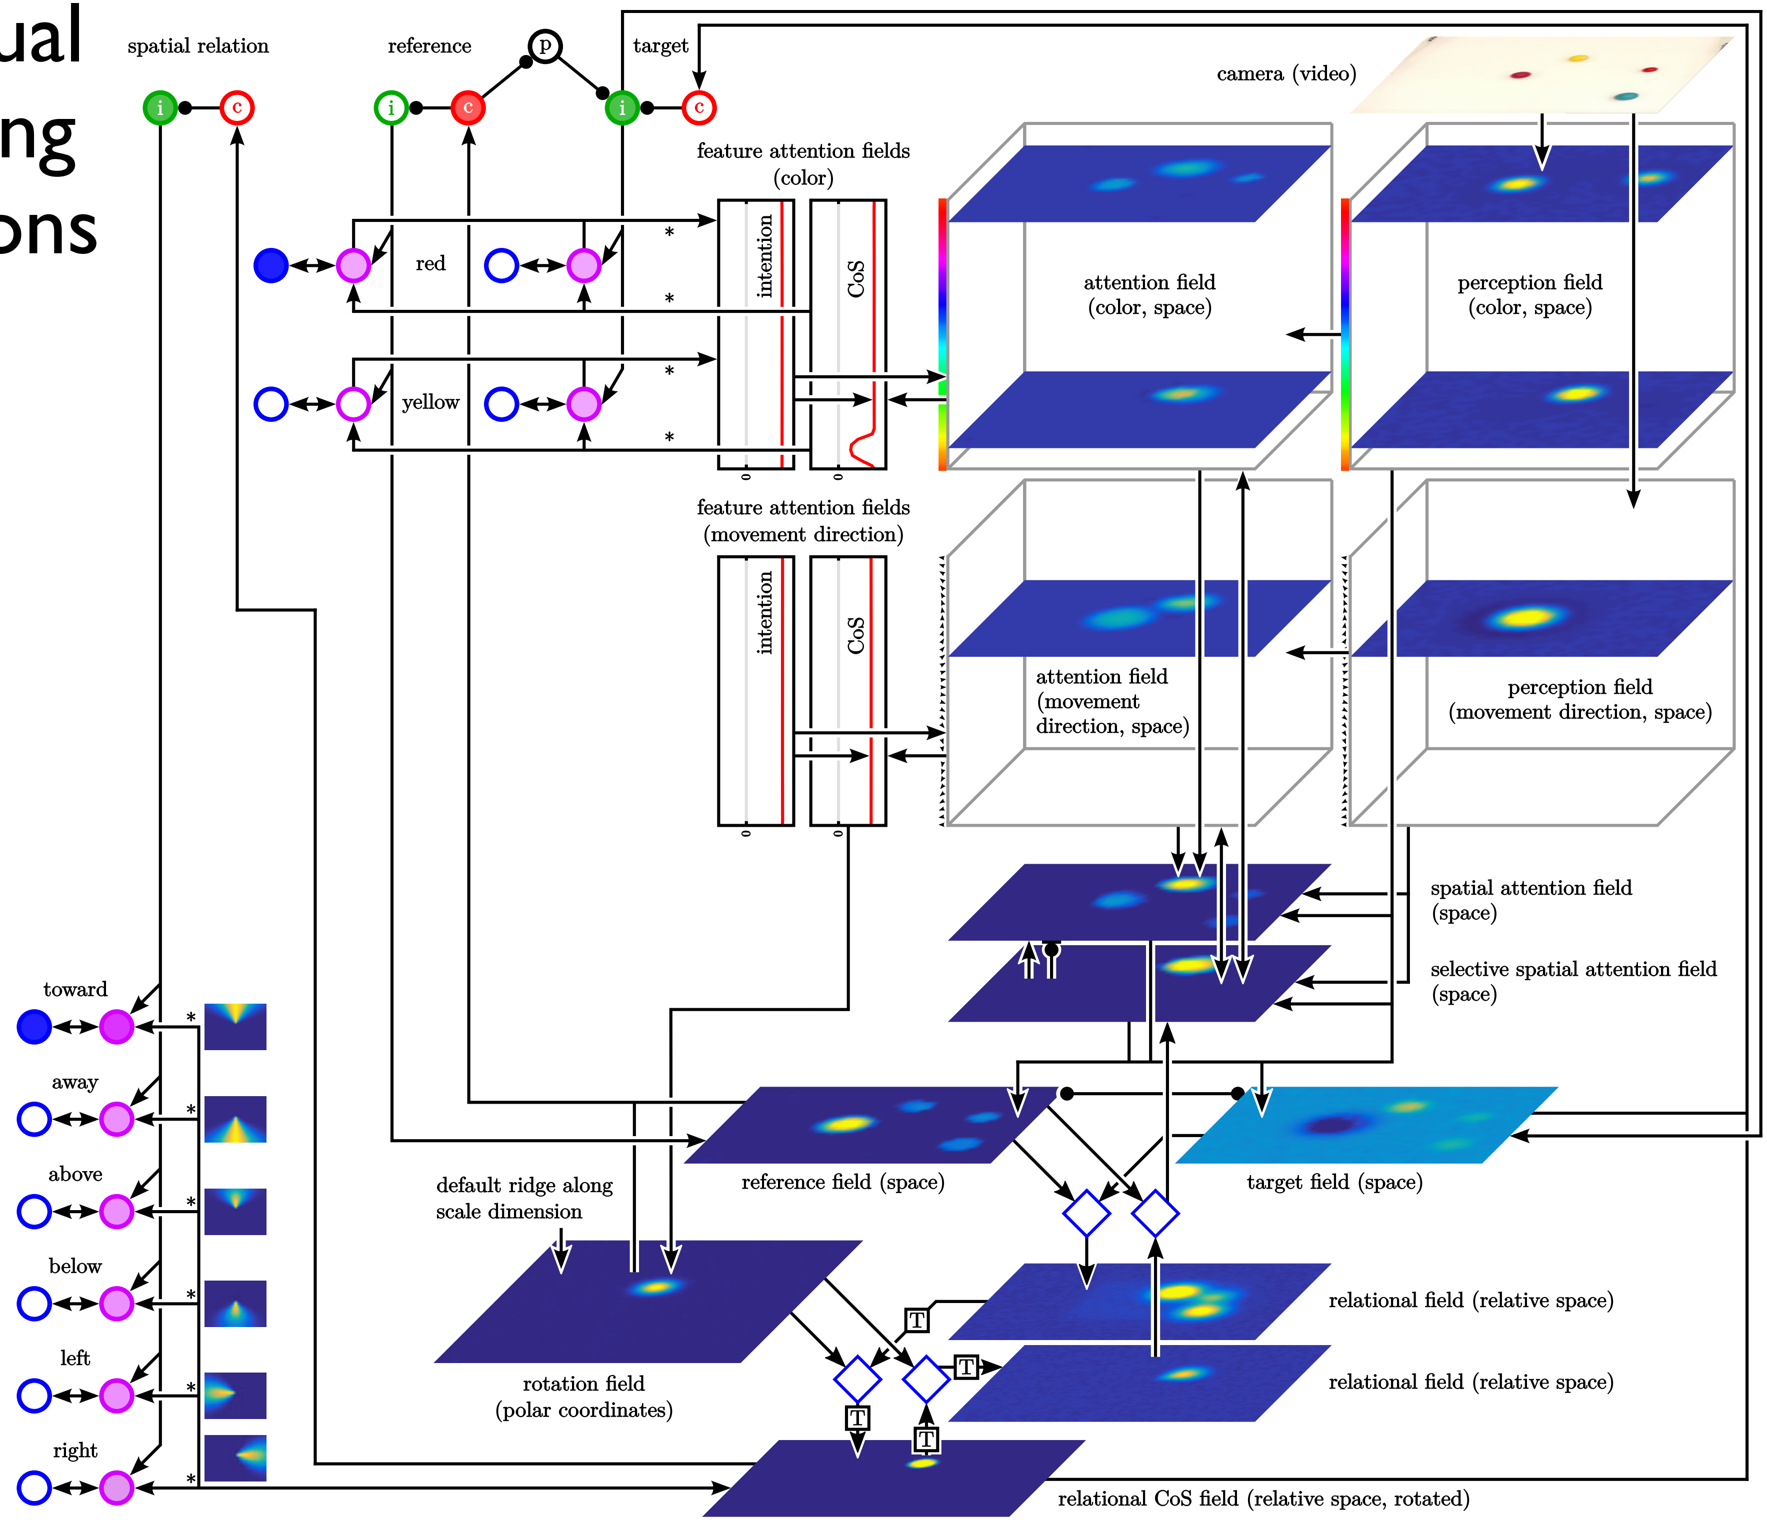

[Richter, Lins, Schöner, TopiCS (2017)]

“red to the left of green”

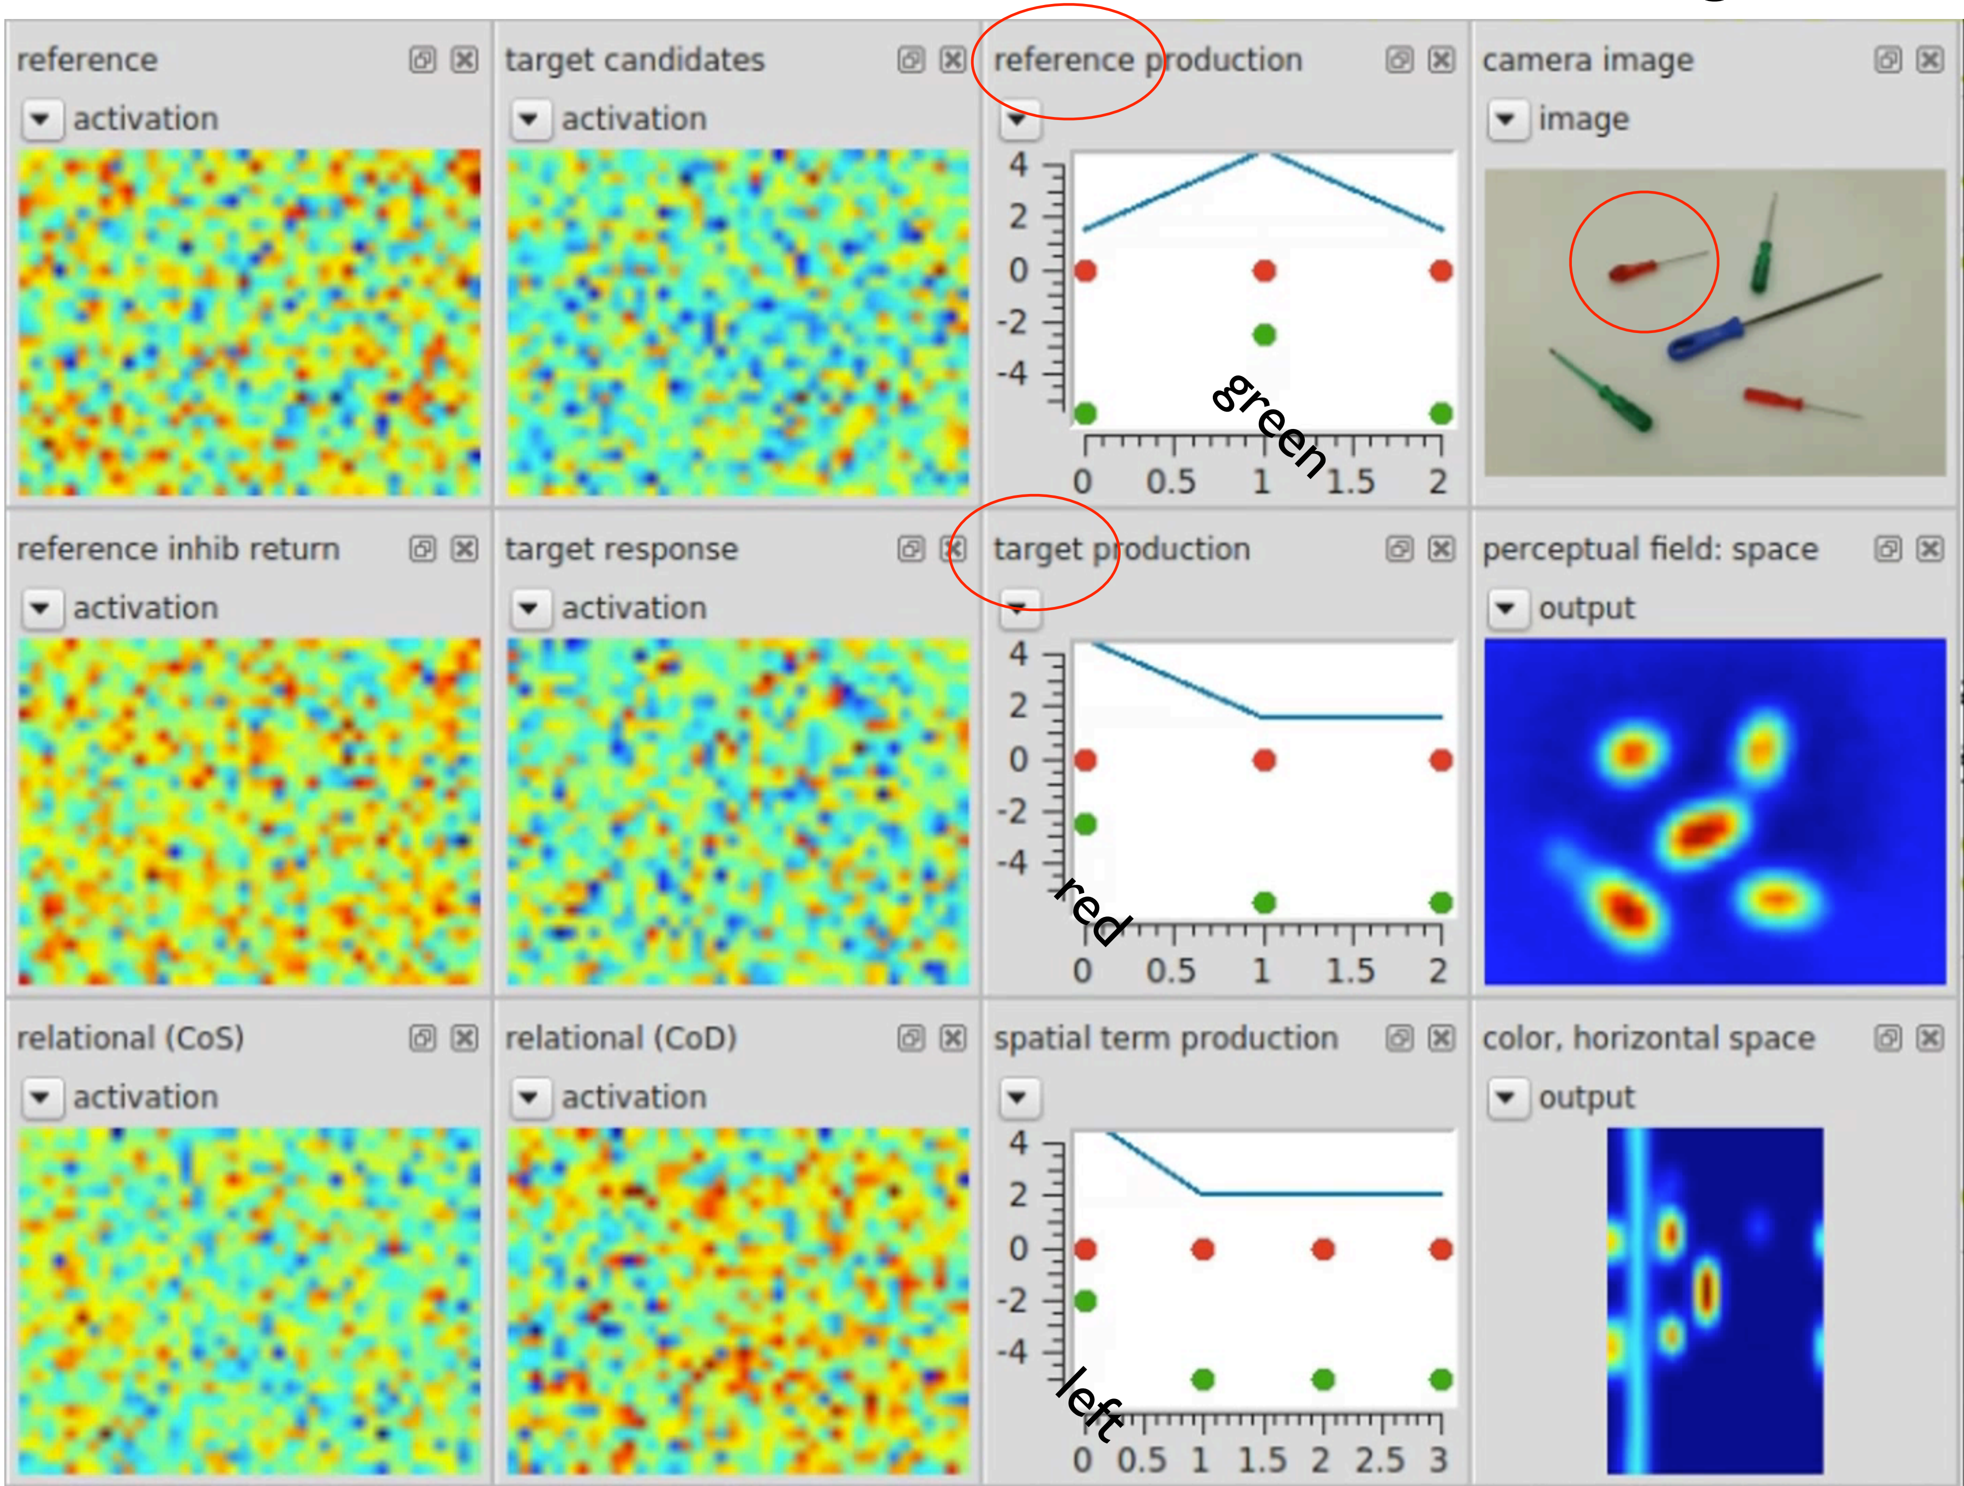

Mental mapping based on relations

"cyan above green"
"red left of green"
"blue right of red"

- "There is a cyan object above a green o

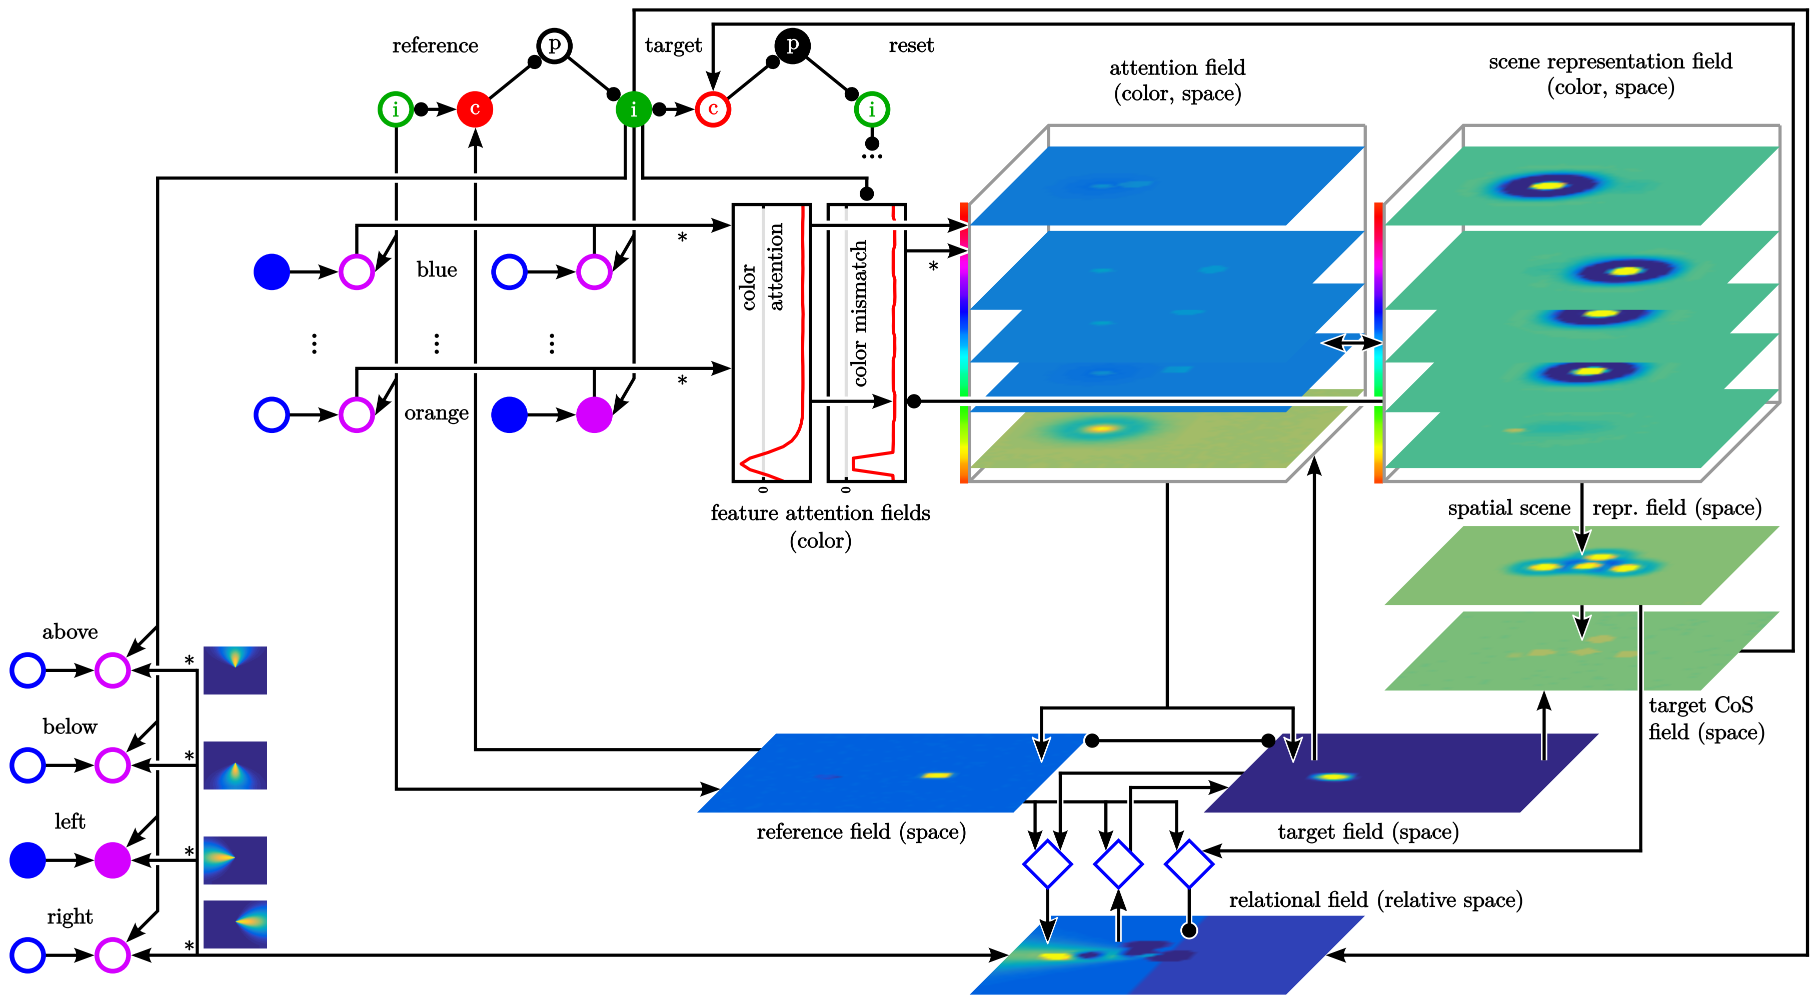

[Kounatidou, Richter, Schöner, CogSci 2018]

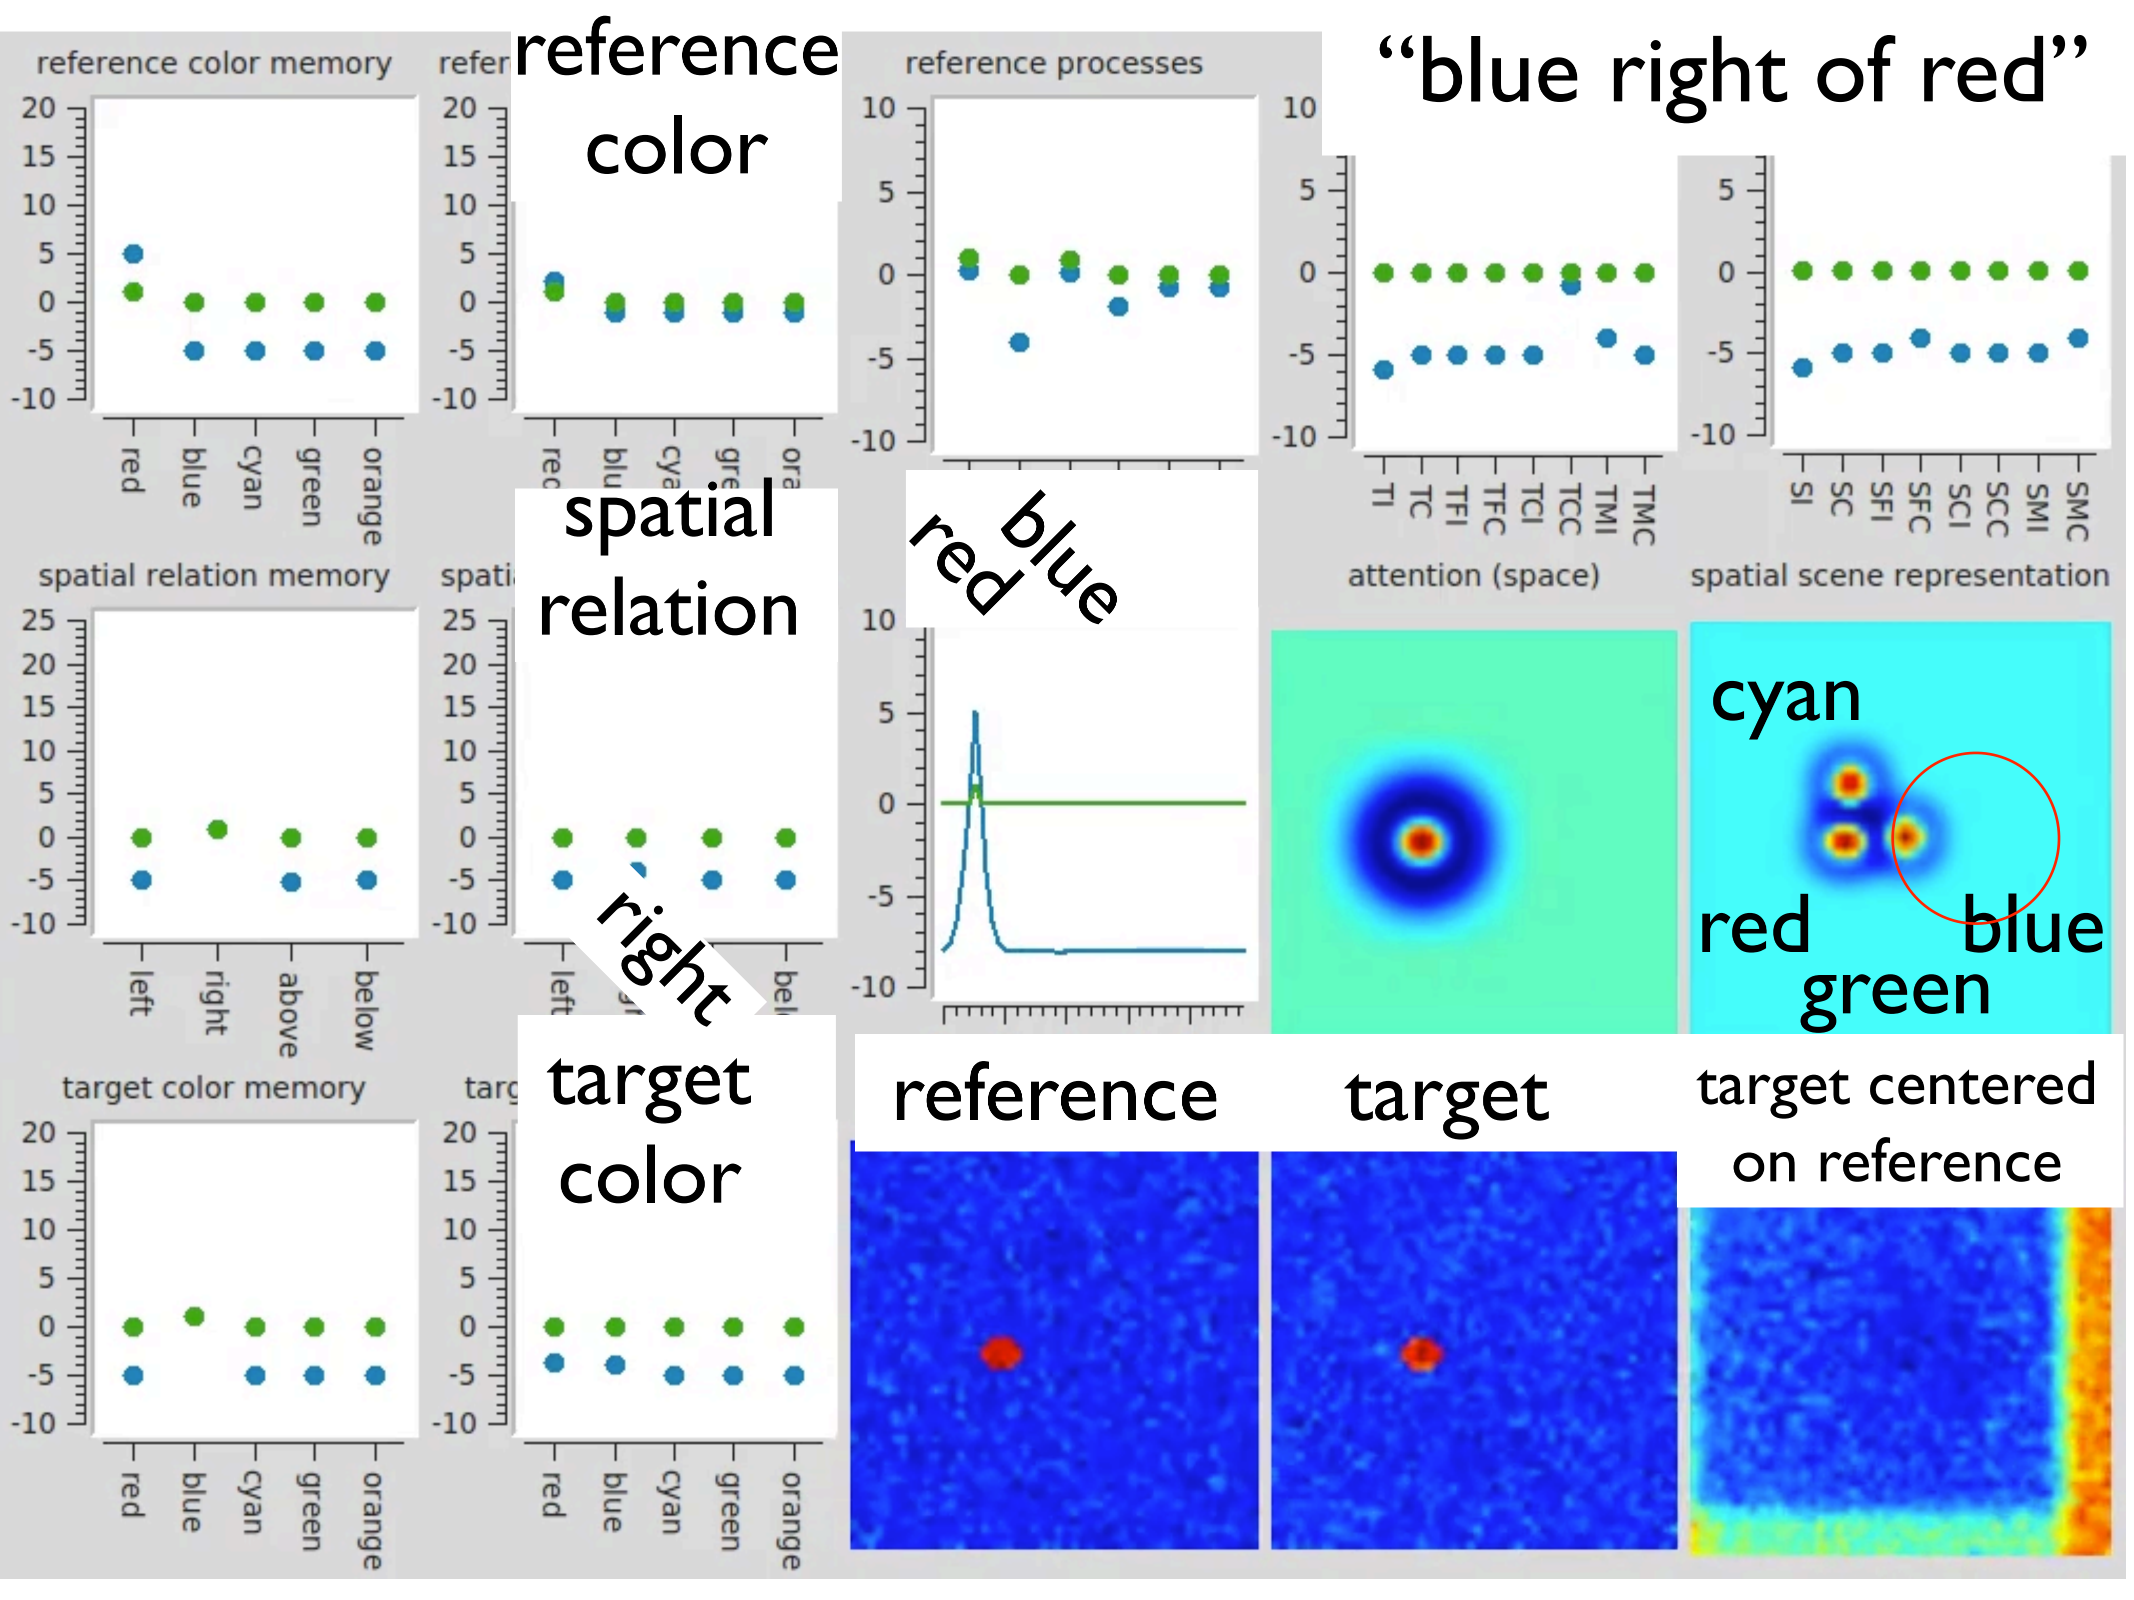

"blue right of red"
reference color memory
reference color
reference processes
spatial relation memory
spatial relation
red
blue
right
attention (space)
spatial scene representation
cyan
red
blue
green
target color memory
target color
reference
target
target centered on reference
red
blue
cyan
green
orange
left
right
above
below
TI
TC
TFI
TFC
TCI
TCC
TMI
TMC
SI
SC
SFI
SFC
SCI
SCC
SMI
SMC

Intentional agent

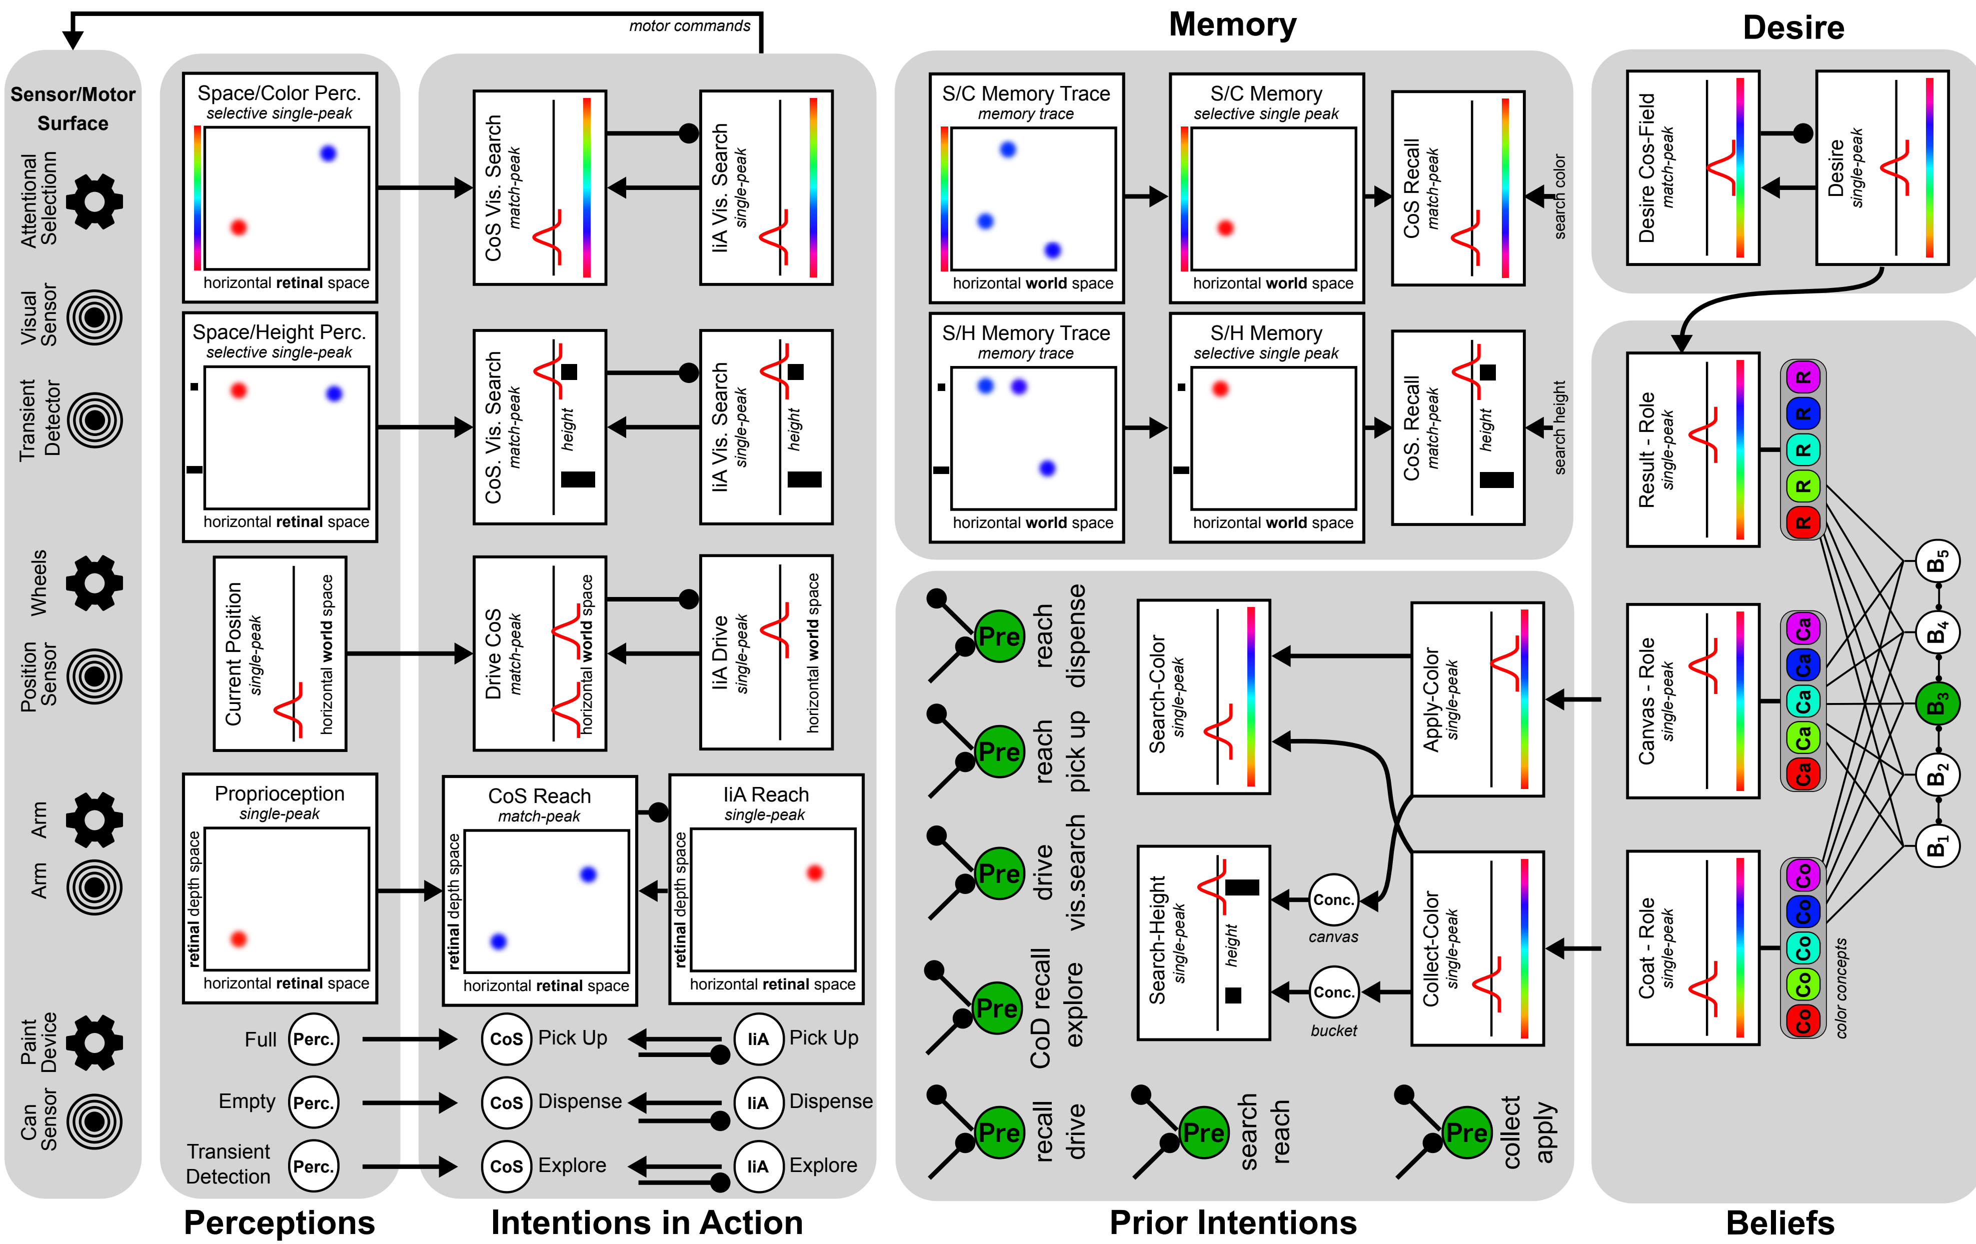

[Tekülve Schöner, CogSci 2019]

Intentional agent

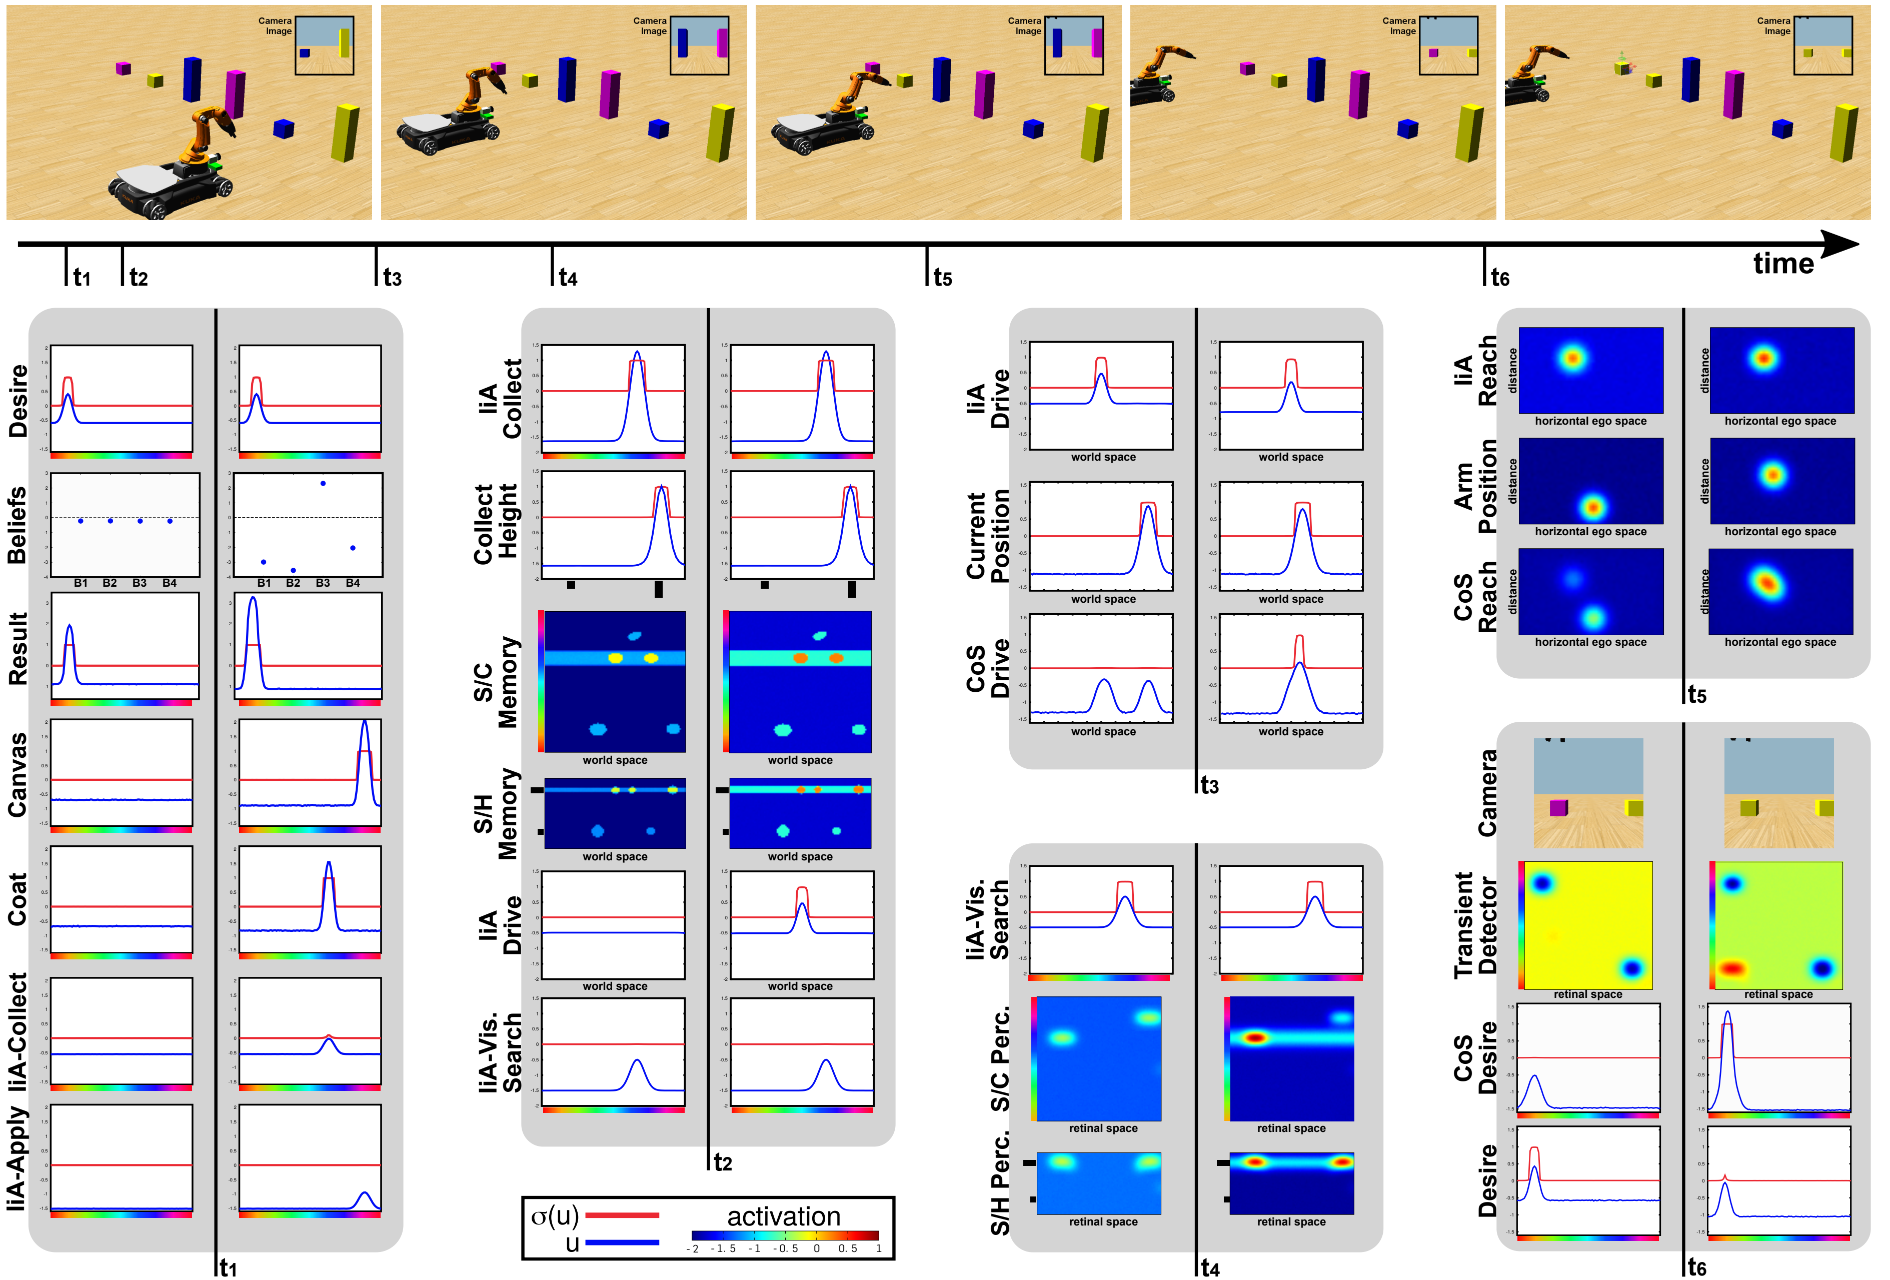

[Tekülve Schöner, CogSci 2019]

=> demonstration/interactive session

Why would this radically neural approach be useful?

- architectures represent theories of task that capture background knowledge
  - about possible actions
  - about how to reach goals by combining primitives

Why would this radically neural approach be useful?

- all elements of representation are perceptually grounded
  - open, in principle, to autonomous learning
- because behavior is built from primitives
  - no need to control arbitrary movement
  - + open, in principle, to learning physical interaction

Conclusion 1

- but: this is a research frontier rather than a ready-made theory
- capturing background knowledge in architectures …
- autonomous learning from experience …
- physical interaction …

…. the other dimension to the radical hypothesis: neuromorphics

- when all AI is on a neural chip… all AI needs to be neural…
- => Yulia Sandamirskaya's introduction

Conclusion 2

- sometimes a radical reset is needed to get out of dead-end roads..
- that is what we are contemplating