The Binding Problem

- *Massively Parallel Brain*
- *Unitary Conscious Experience*
- *Many Variations and Proposals*
- *Our focus: The Variable Binding Problem*

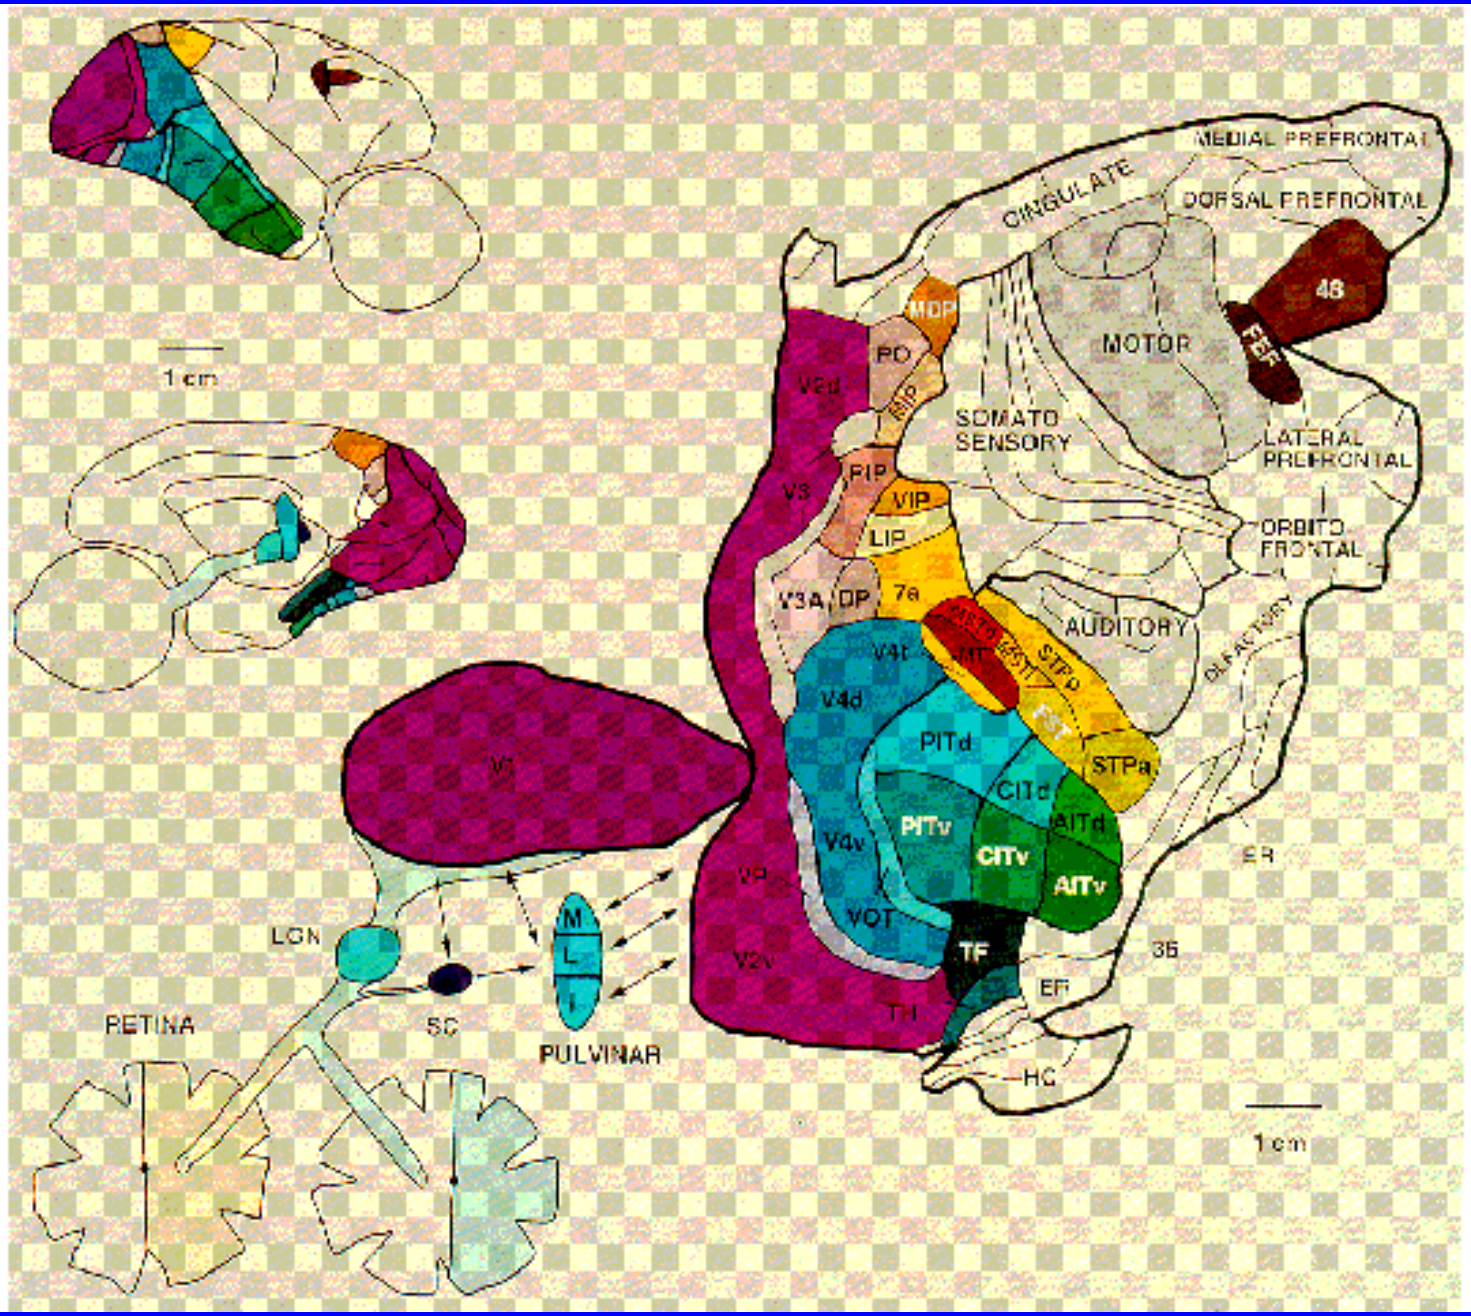

Problem

- **Binding problem**
  - **In vision**
    - You do not exchange the colors of the shapes below

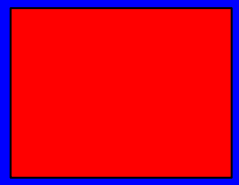

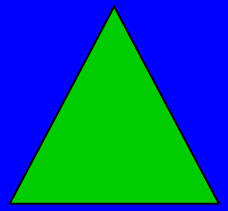

  - **In behavior**
    - Grasp motion depends on object to grasp
  - **In inference**
    - Human(x) -> Mortal(x)
    - Must bind a variable to x

Automatic Inference

- Inference needed for many tasks
  - Reference resolution
  - General language understanding
  - Planning
- Humans do this quickly and without conscious thought
  - Automatically
  - No real intuition of how we do it

Other Solutions in Inference

- Brute-force enumeration
  - Does not scale to depth of human knowledge
- Signature propagation (direct reference)
  - Difficult to pass enough information to directly reference each object
  - Unifying two bindings (e.g. reference resolution) is difficult
- Temporal synchrony example (SHRUTI)
  - Little biological evidence

SHRUTI

- SHRUTI does inference by connections between simple computation nodes
- Nodes are small groups of neurons
- Nodes firing in sync reference the same object

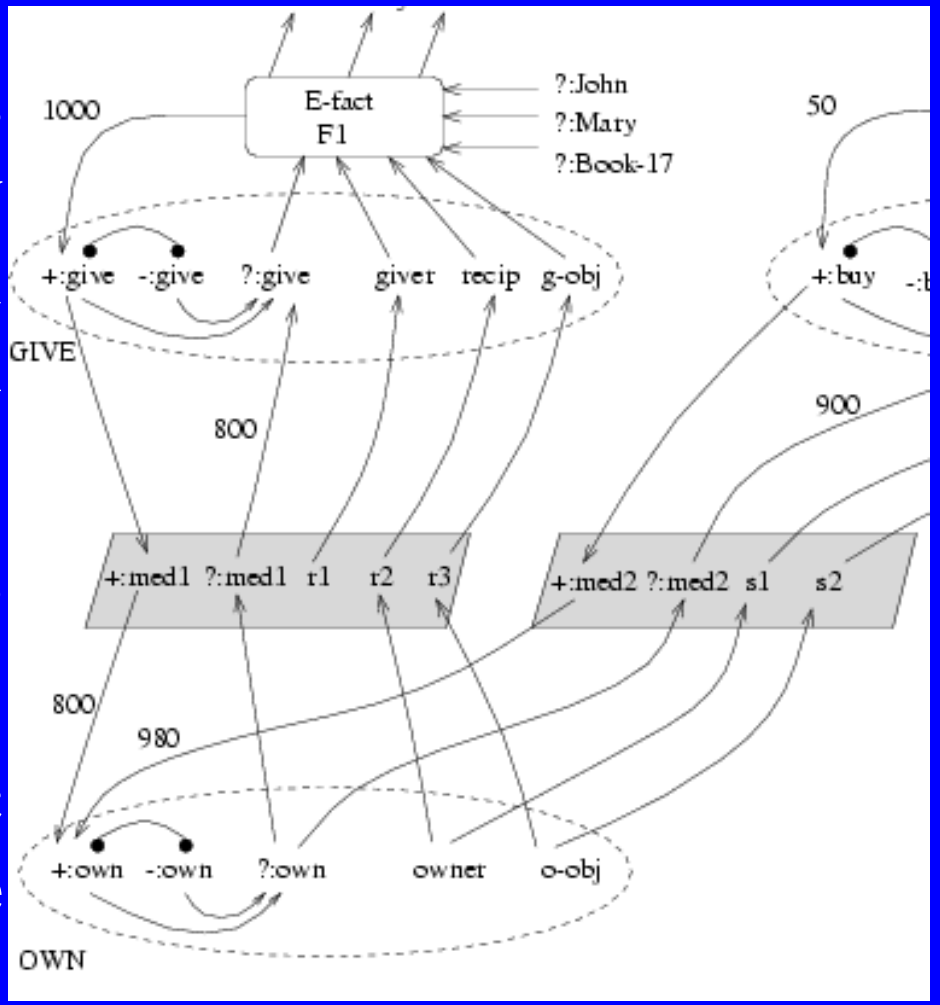

shruti

A Neurally Plausible model of Reasoning

Lokendra Shastri
International Computer Science Institute
Berkeley, CA 94704

Five levels of Neural Theory of Language

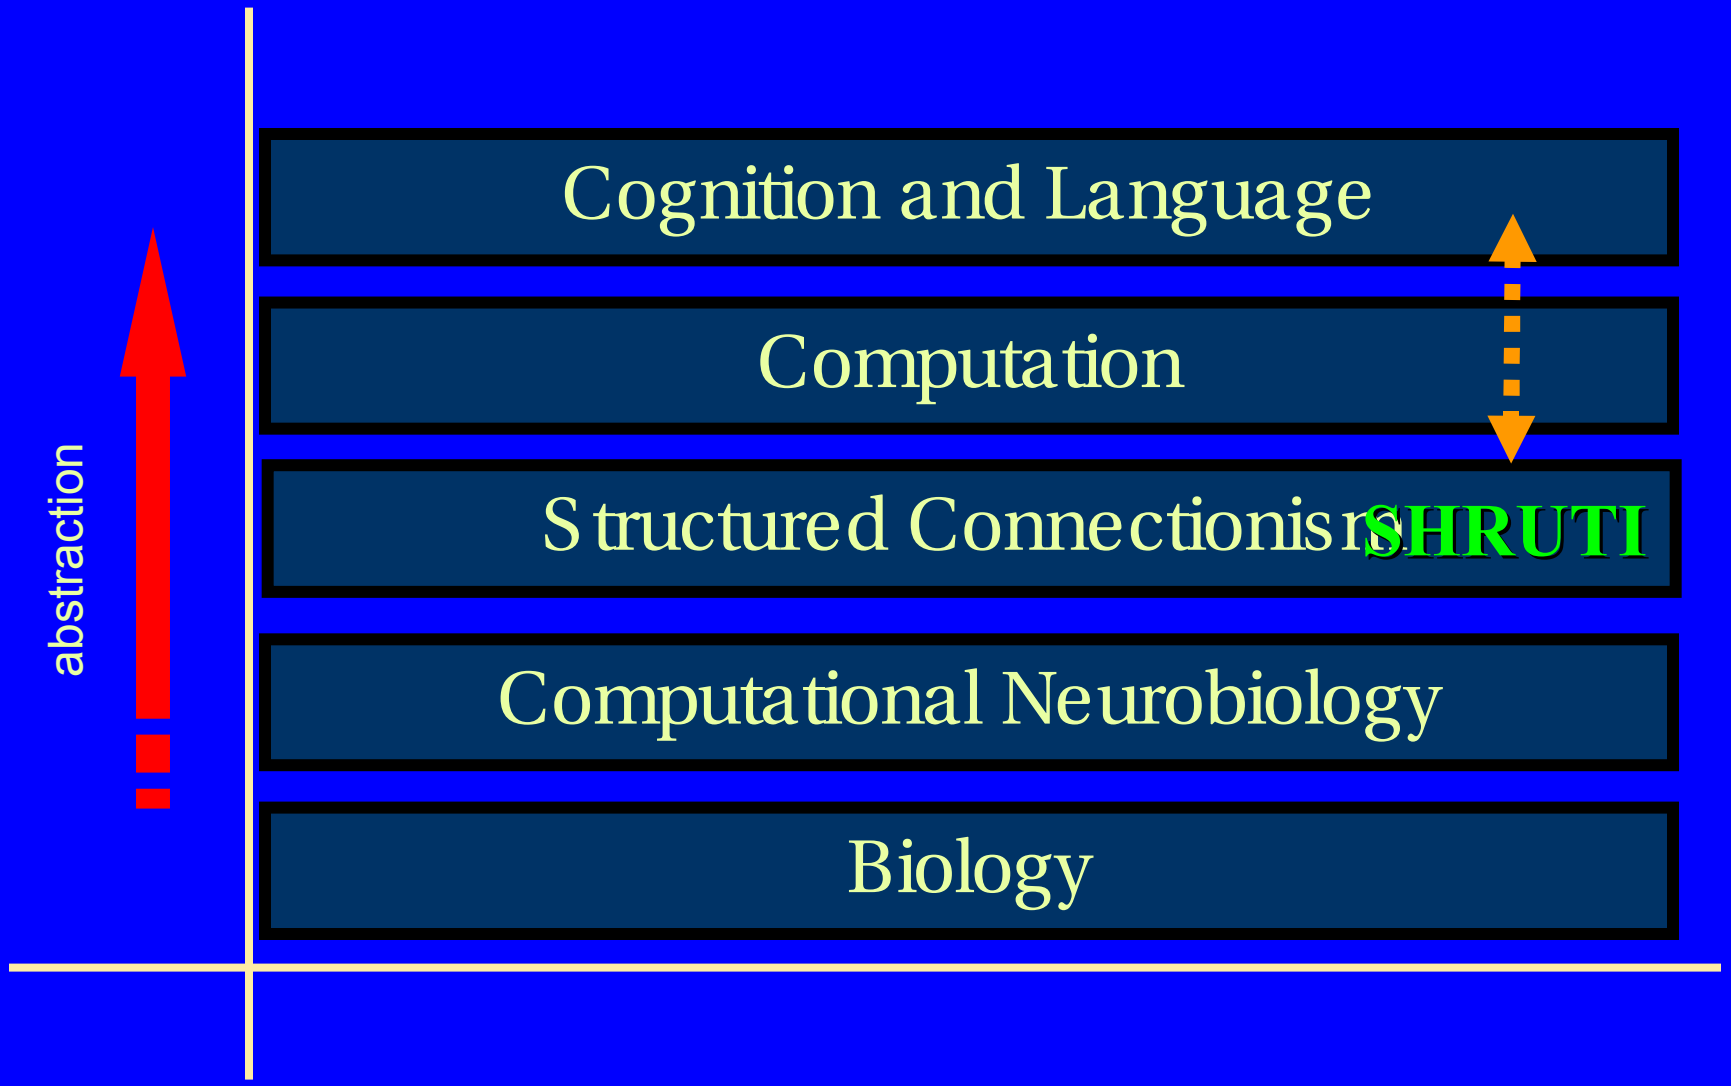

"John fell in the hallway. Tom had cleaned it. He got hurt."

⇒ *Tom had cleaned the hallway.*

⇒ *The hallway floor was wet.*

⇒ *John slipped and fell on the wet floor.*

⇒ *John got hurt as a result of the fall.*

*such inferences establish **referential** and **causal coherence.***

Reflexive Reasoning

- *Ubiquitous*
- *Automatic, effortless*
- *Extremely fast --- almost a reflex response of our cognitive apparatus*

Reflexive Reasoning

Not all reasoning is reflexive

*Contrast with **reflective** reasoning*

deliberate

involves explicit consideration of alternatives

require props (paper and pencil)

e.g., solving logic puzzles … differential equations

How fast is reflexive reasoning?

- **We understand language at the rate of 150-400 words per minute**

⇒ ***Reflexive inferences required for establishing inferential and causal coherence are drawn within a few hundred milliseconds***

How can a system of slow and simple neuron-like elements

- *encode a large body of semantic and episodic knowledge and yet*
- *perform a wide range of inferences within a few hundred milliseconds?*

Characterization of reflexive reasoning?

- *What can and cannot be inferred via reflexive processes?*

Shruti

http://www.icsi.berkeley.edu/~shastri/shruti

- Lokendra Shastri
- V. Ajjanagadde (Penn, ex-graduate student)
- Carter Wendelken (UCB, ex-graduate student)
- D. Mani (Penn, ex-graduate student)
- D.J. Grannes (UCB, ex-graduate student)
- Jerry Hobbs, USC/ISI (abductive reasoning)
- Marvin Cohen, CTI (metacognition; belief and utility)
- Bryan Thompson, CTI (metacognition; belief and utility)

Reflexive Reasoning
representational and processing issues

- *Activation-based (dynamic) representation of events and situations (relational instances)*

Dynamic representation of relational instances

"John gave Mary a book"

giver:	***John***
recipient:	***Mary***
given-object:	***a-book***

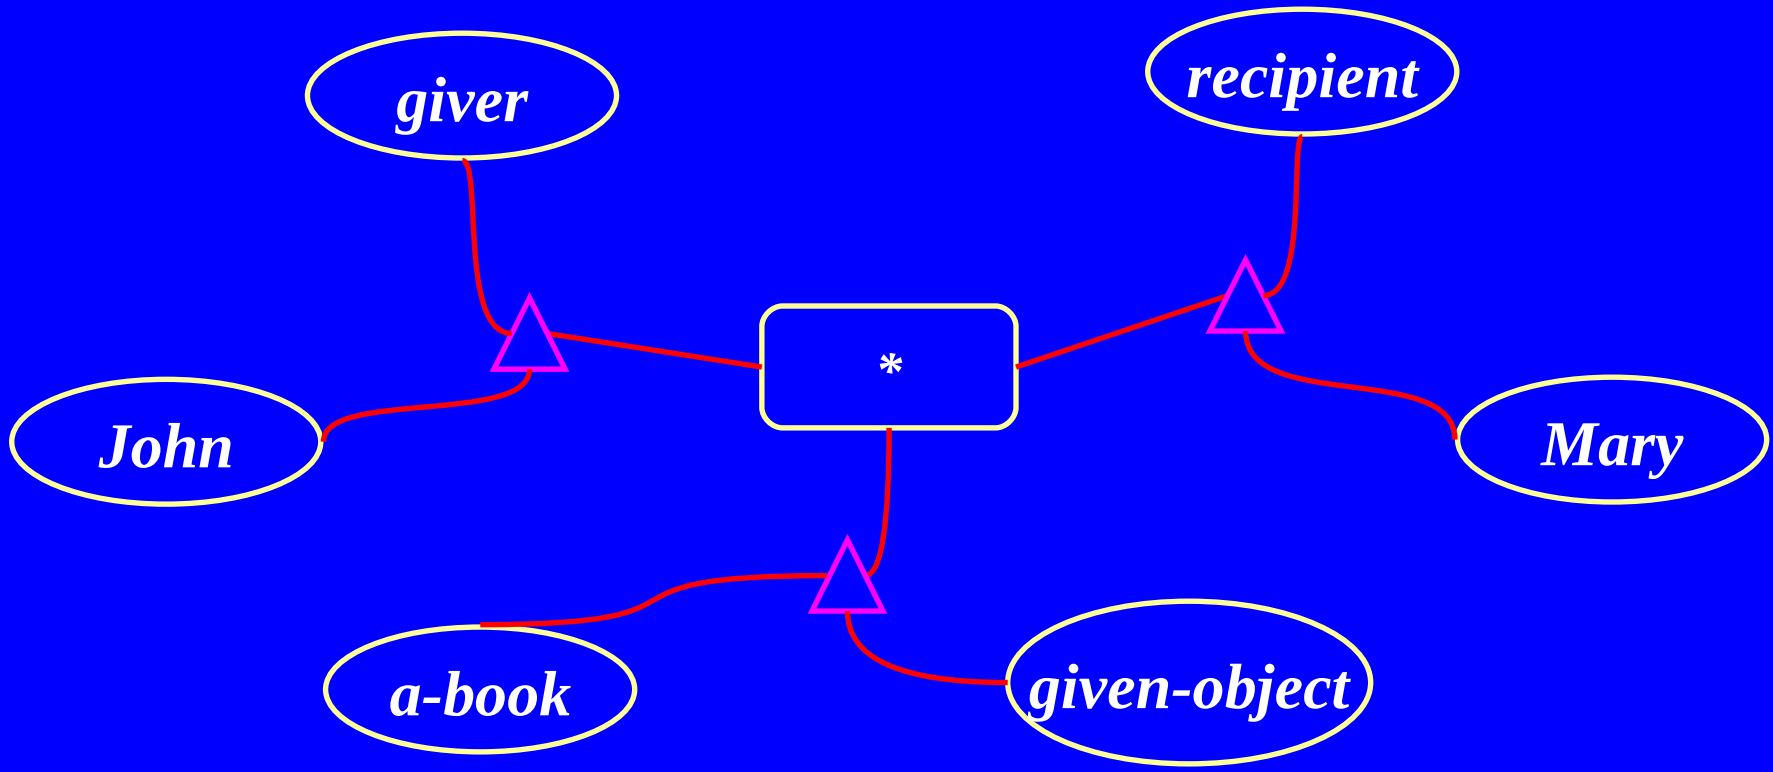

Reflexive Reasoning

Requires compatible neural mechanisms for:

- *Expressing dynamic bindings*
- *Systematically propagating dynamic bindings*
- *Computing coherent explanations and predictions*
  - *evidence combination*
  - *instantiation and unification of entities*

All of the above must happen rapidly

Learning

- *one-shot learning of events and situations (episodic memory)*
- *gradual/incremental learning of concepts, relations, schemas, and causal structures*

Relation focal-cluster

FALL

Entity, category and relation focal-clusters

FALL

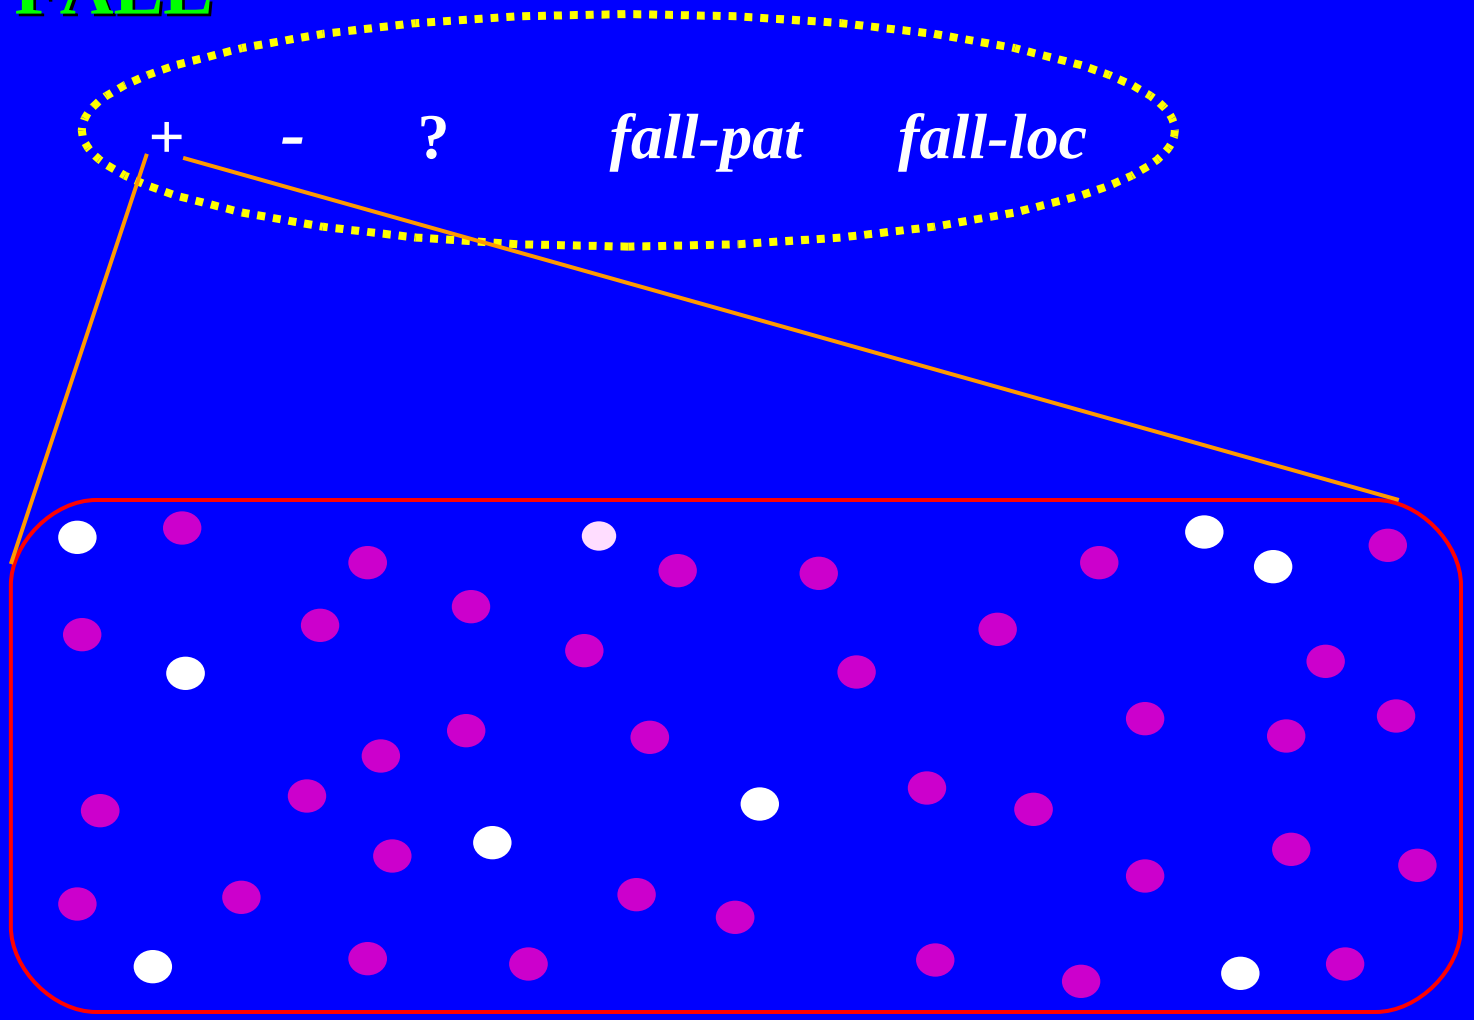

Entity, category and relation focal-clusters

FALL

Functional nodes in a focal-cluster [collector (+/-), enabler (?), and role nodes] may be situated in different brain regions

Focal-cluster of a relational schema

focal-clusters of perceptual schemas and sensory representations associated with fall

focal-clusters of other relational schemas causally related to fall

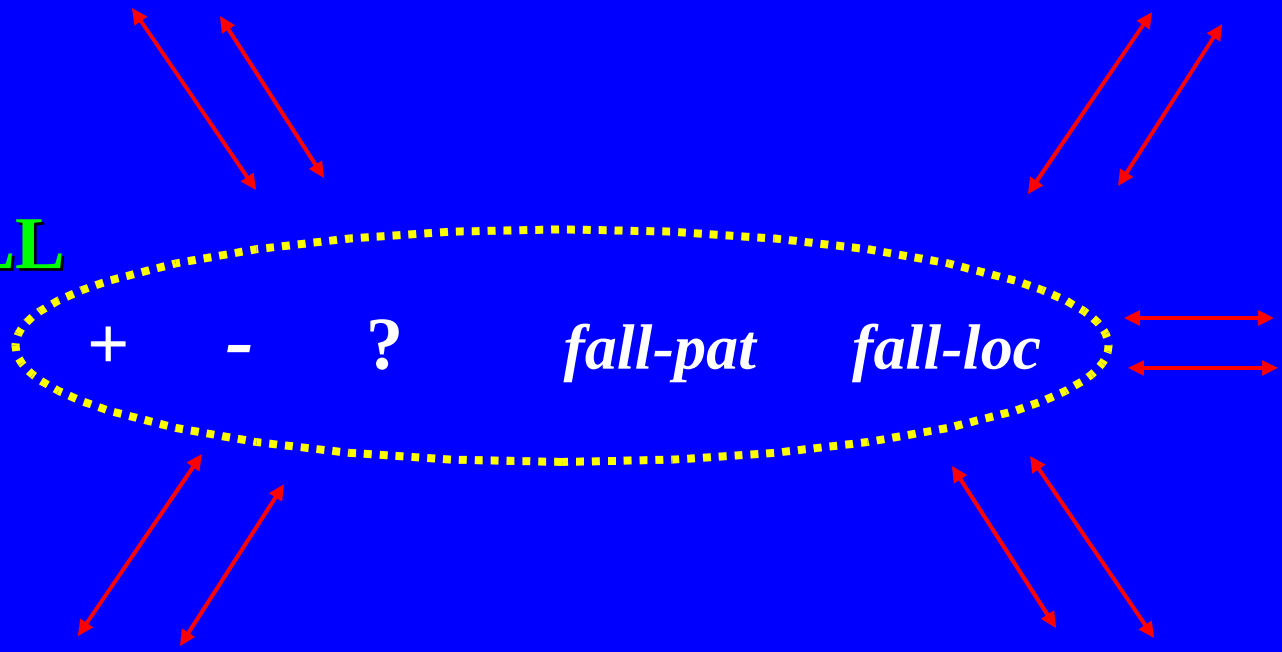

episodic memories of fall events

focal-clusters of motor schemas associated with fall

focal-clusters of lexical know-ledge associated with fall

Focal-clusters

Nodes in the fall focal-cluster become active when

- ***perceiving a fall event***
- ***remembering a fall event***
- ***understanding a sentence about a fall event***
- ***experiencing a fall event***

A focal-cluster is like a "supra-mirror" cluster

Focal-cluster of an entity

focal-clusters of perceptual schemas and sensory representations associated with John

focal-clusters of other entities and categories semantically related to John

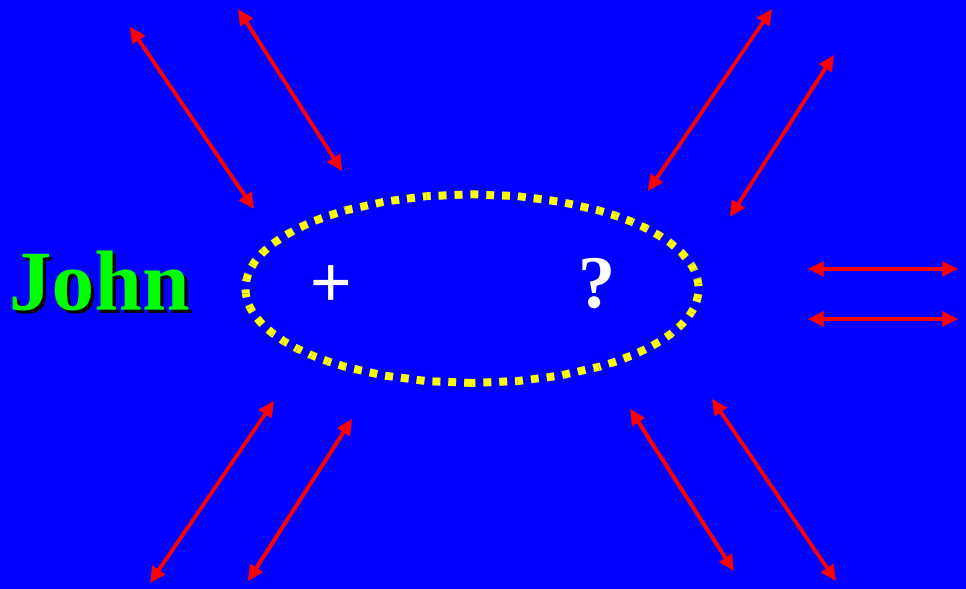

episodic memories where John is one of the role-fillers

focal-clusters of motor schemas associated with John

focal-clusters of lexical know-ledge associated with John

"John fell in the hallway"

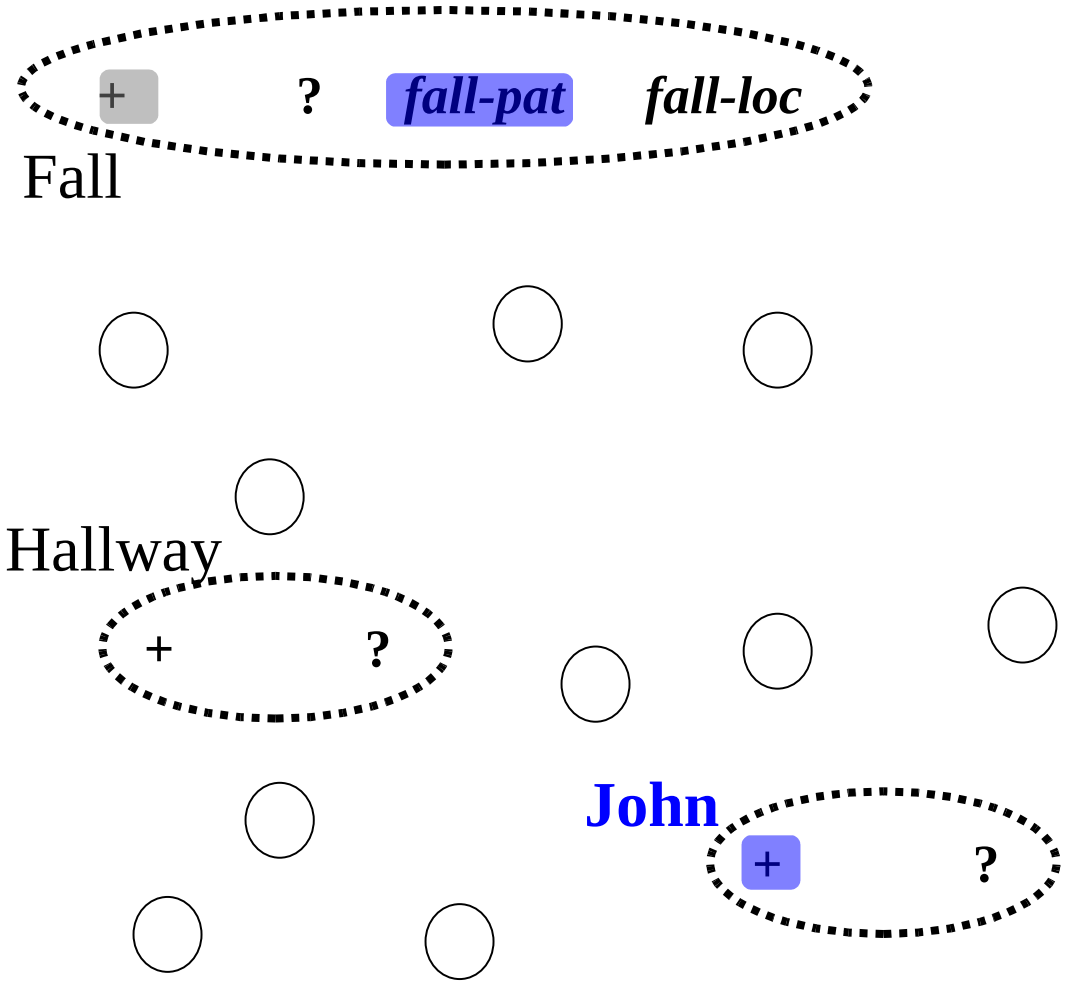

"John fell in the hallway"

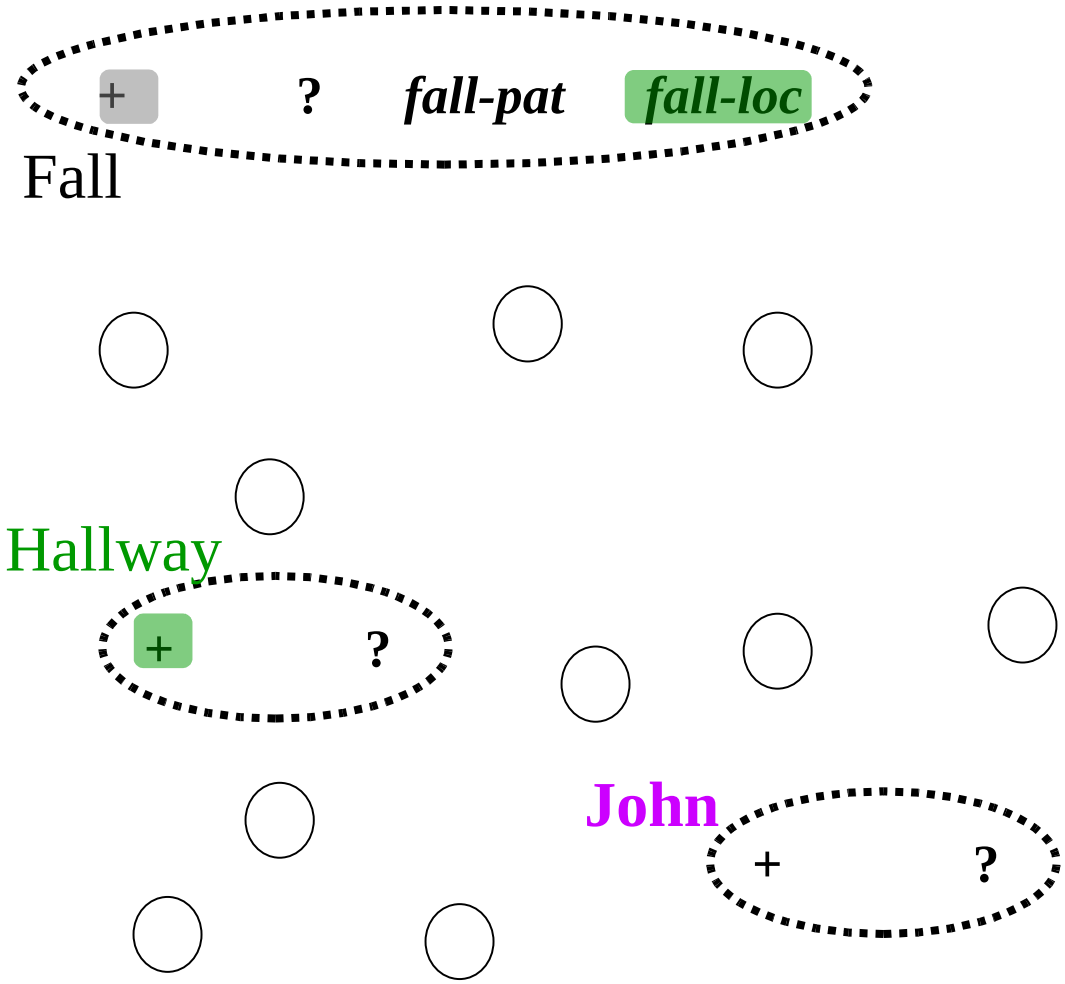

"John fell in the hallway"

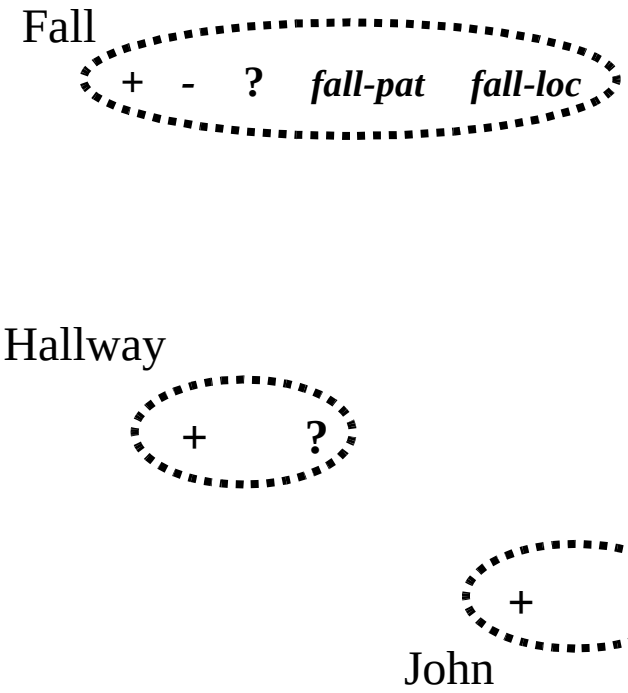

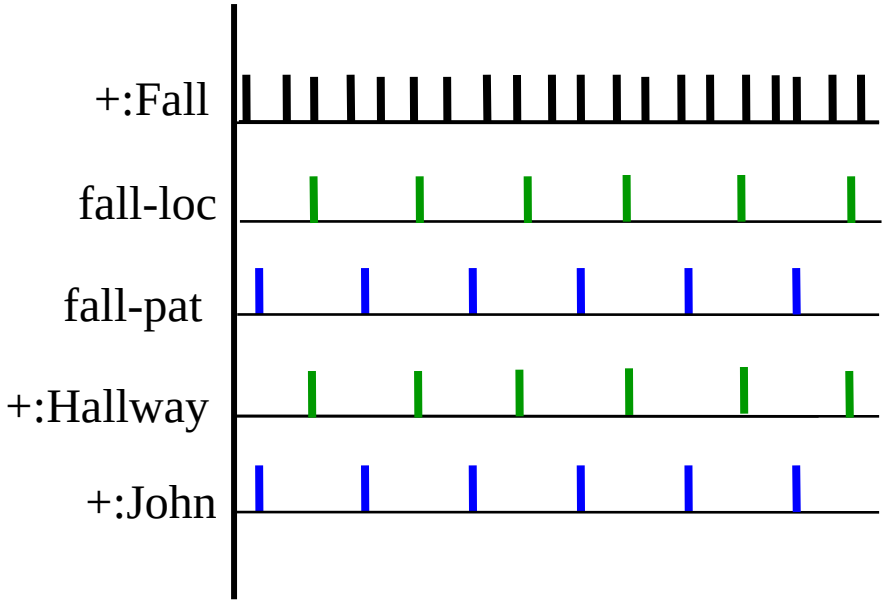

Encoding "slip => fall" in Shruti

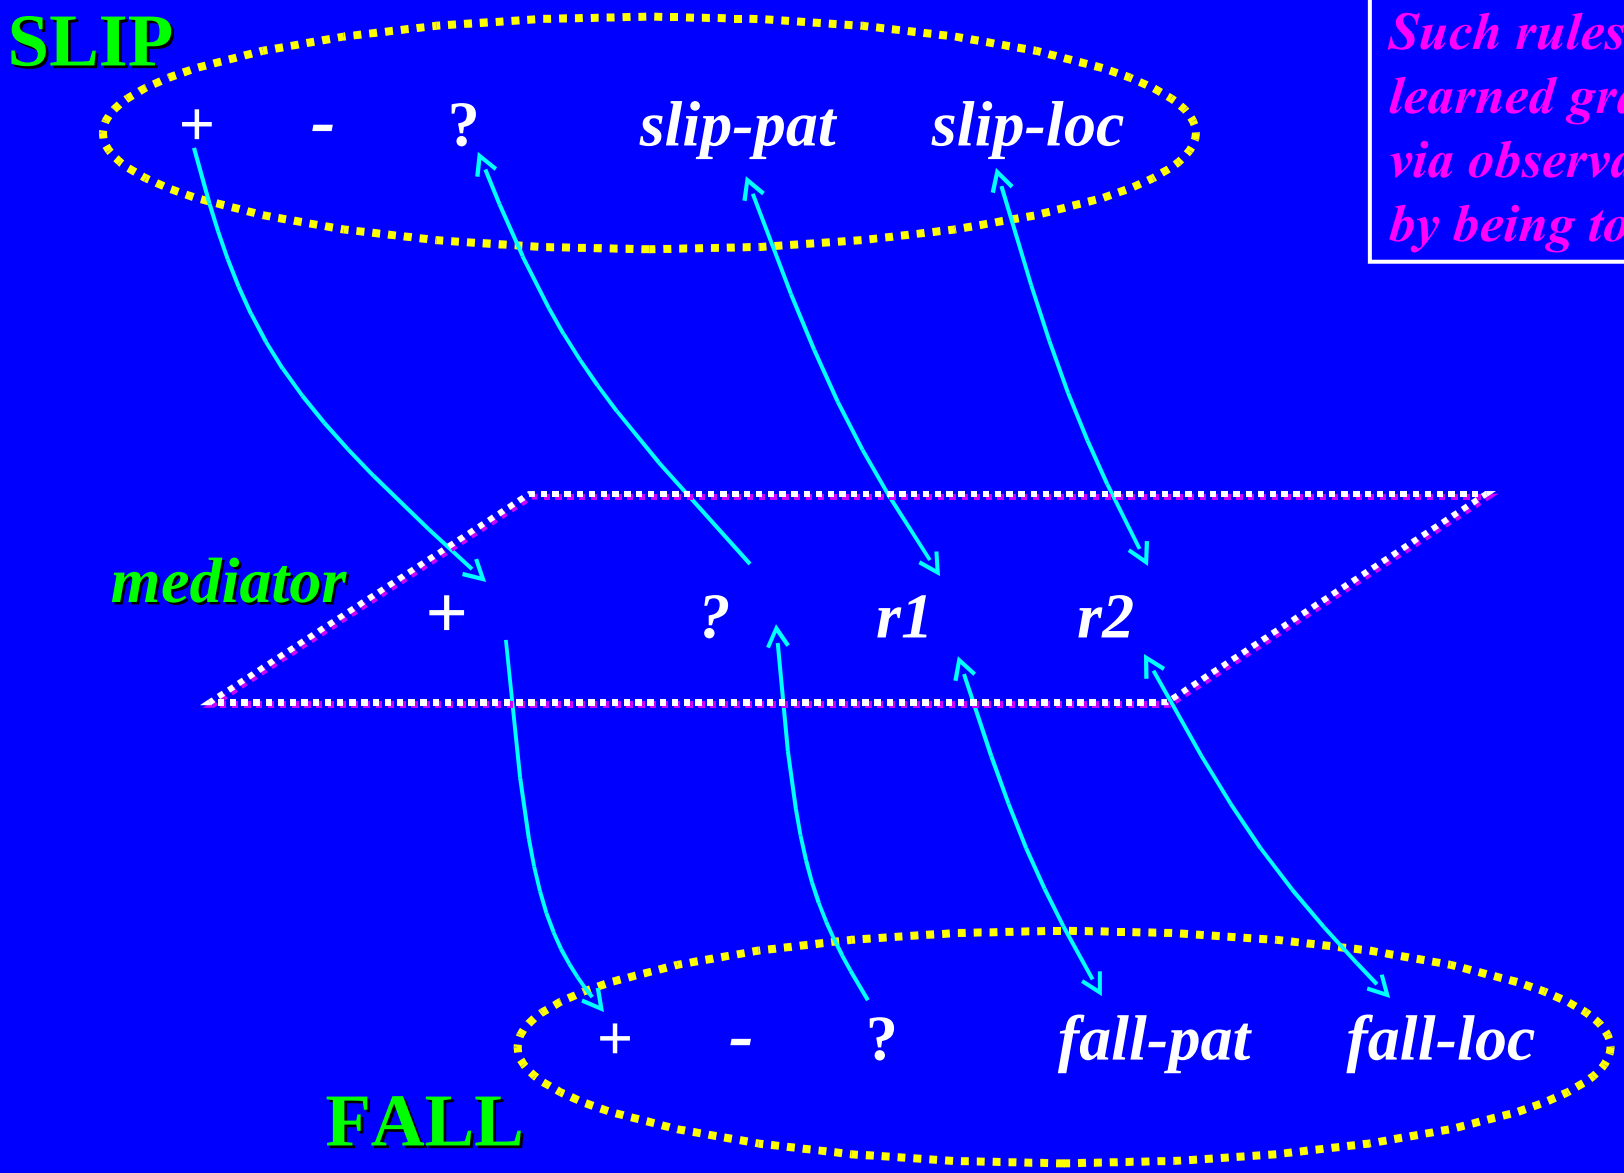

Such rules are learned gradually via observations, by being told …

"John slipped in the hallway" → "John fell in the hallway"

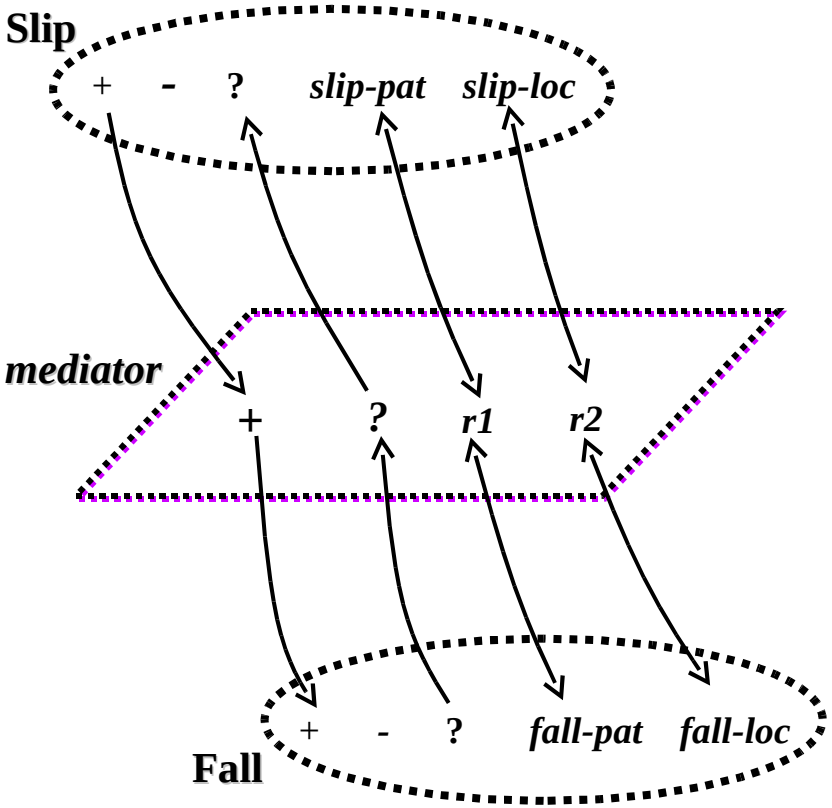

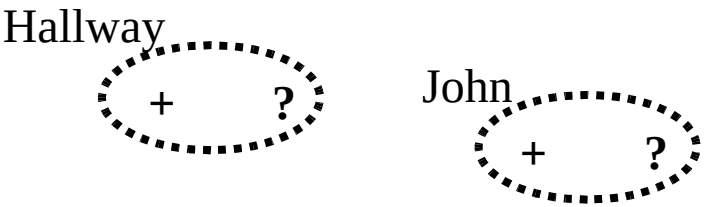

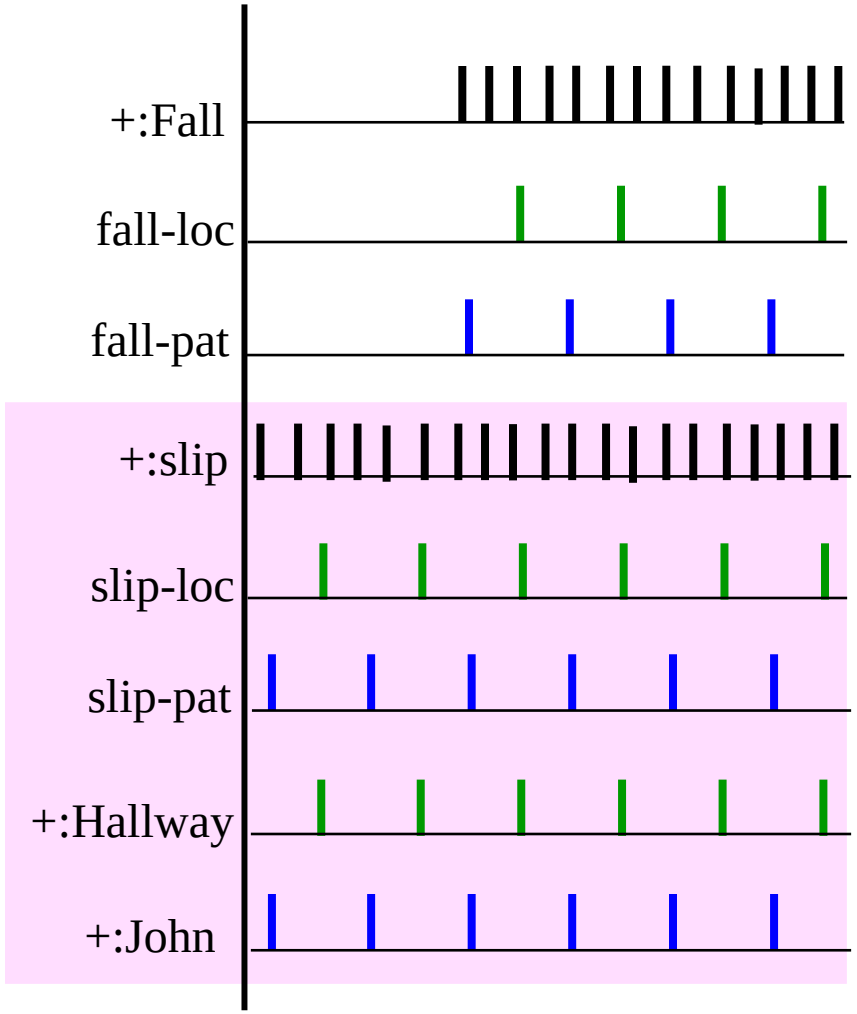

A Metaphor for Reasoning

- *An episode of reflexive reasoning is a transient propagation of rhythmic activity*
- *Each entity involved in this reasoning episode is a phase in this rhythmic activity*
- *Bindings are synchronous firings of cell clusters*
- *Rules are interconnections between cell-clusters that support propagation of synchronous activity*

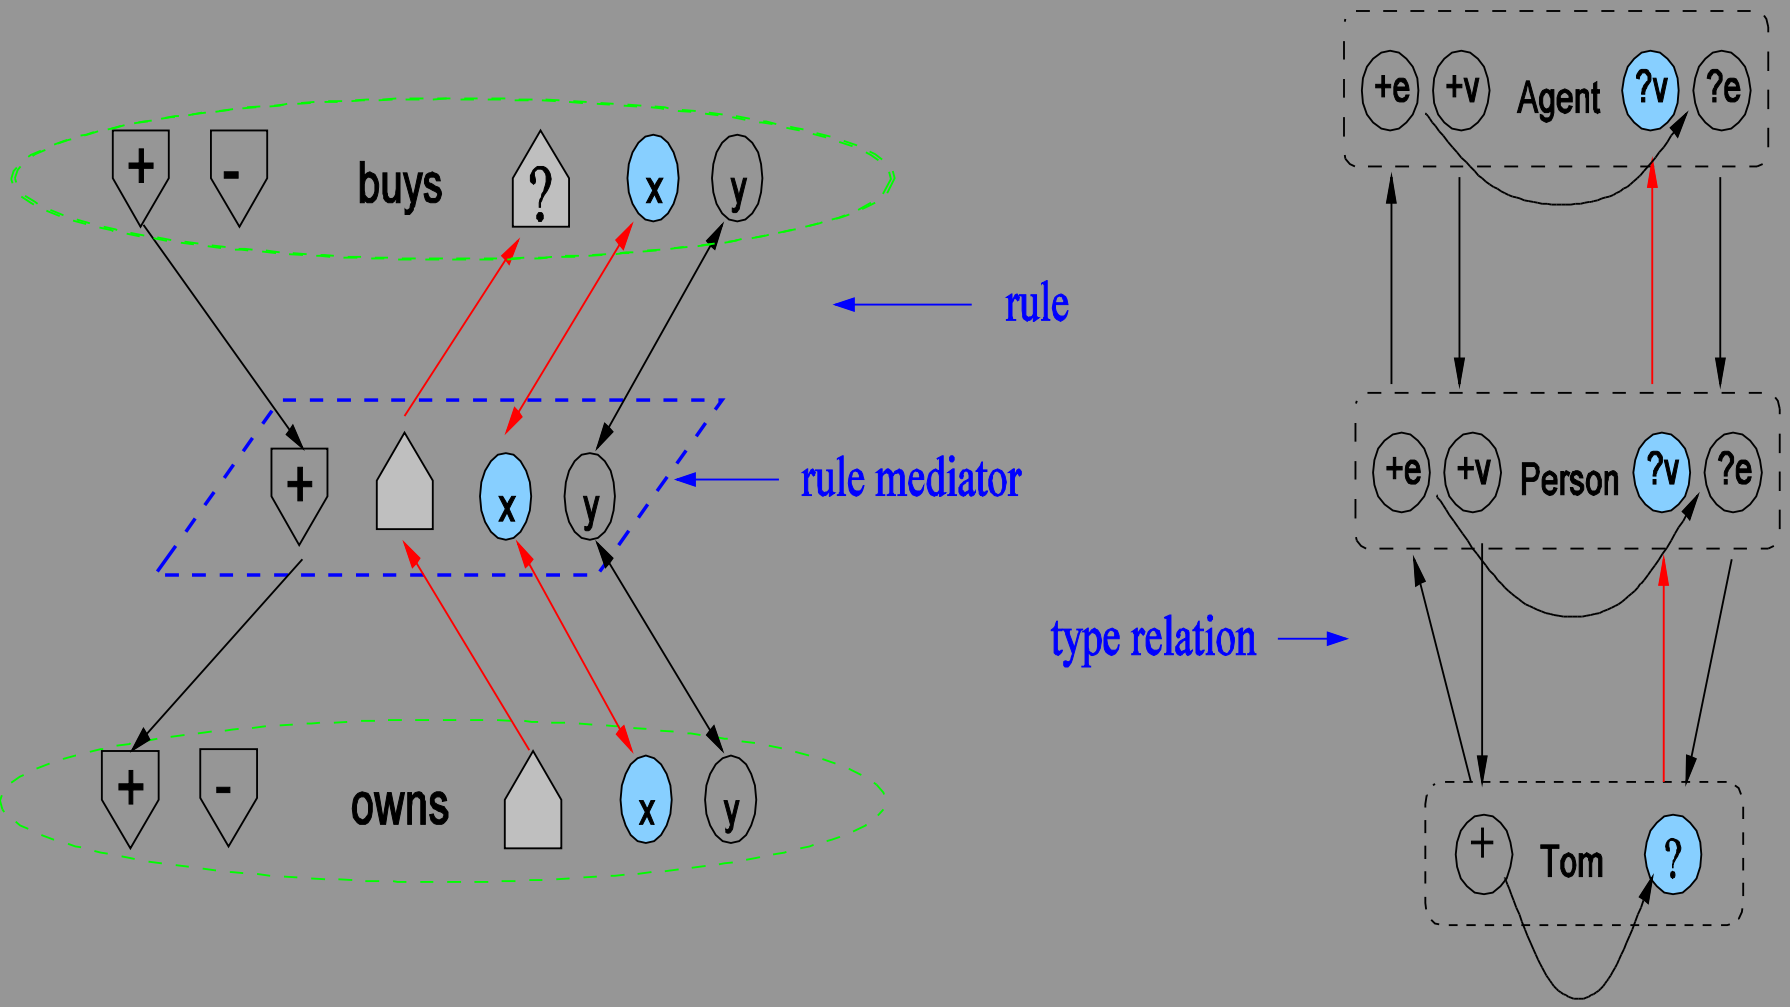

Focal-clusters with intra-cluster links

FALL

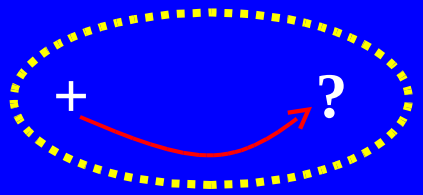

John

Shruti always seeks explanations

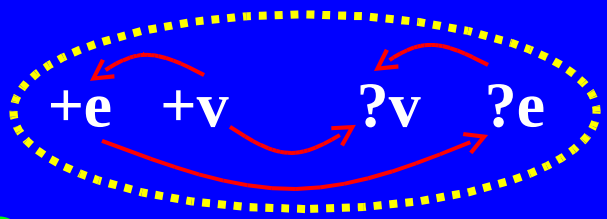

Person

Encoding "slip => fall" in Shruti

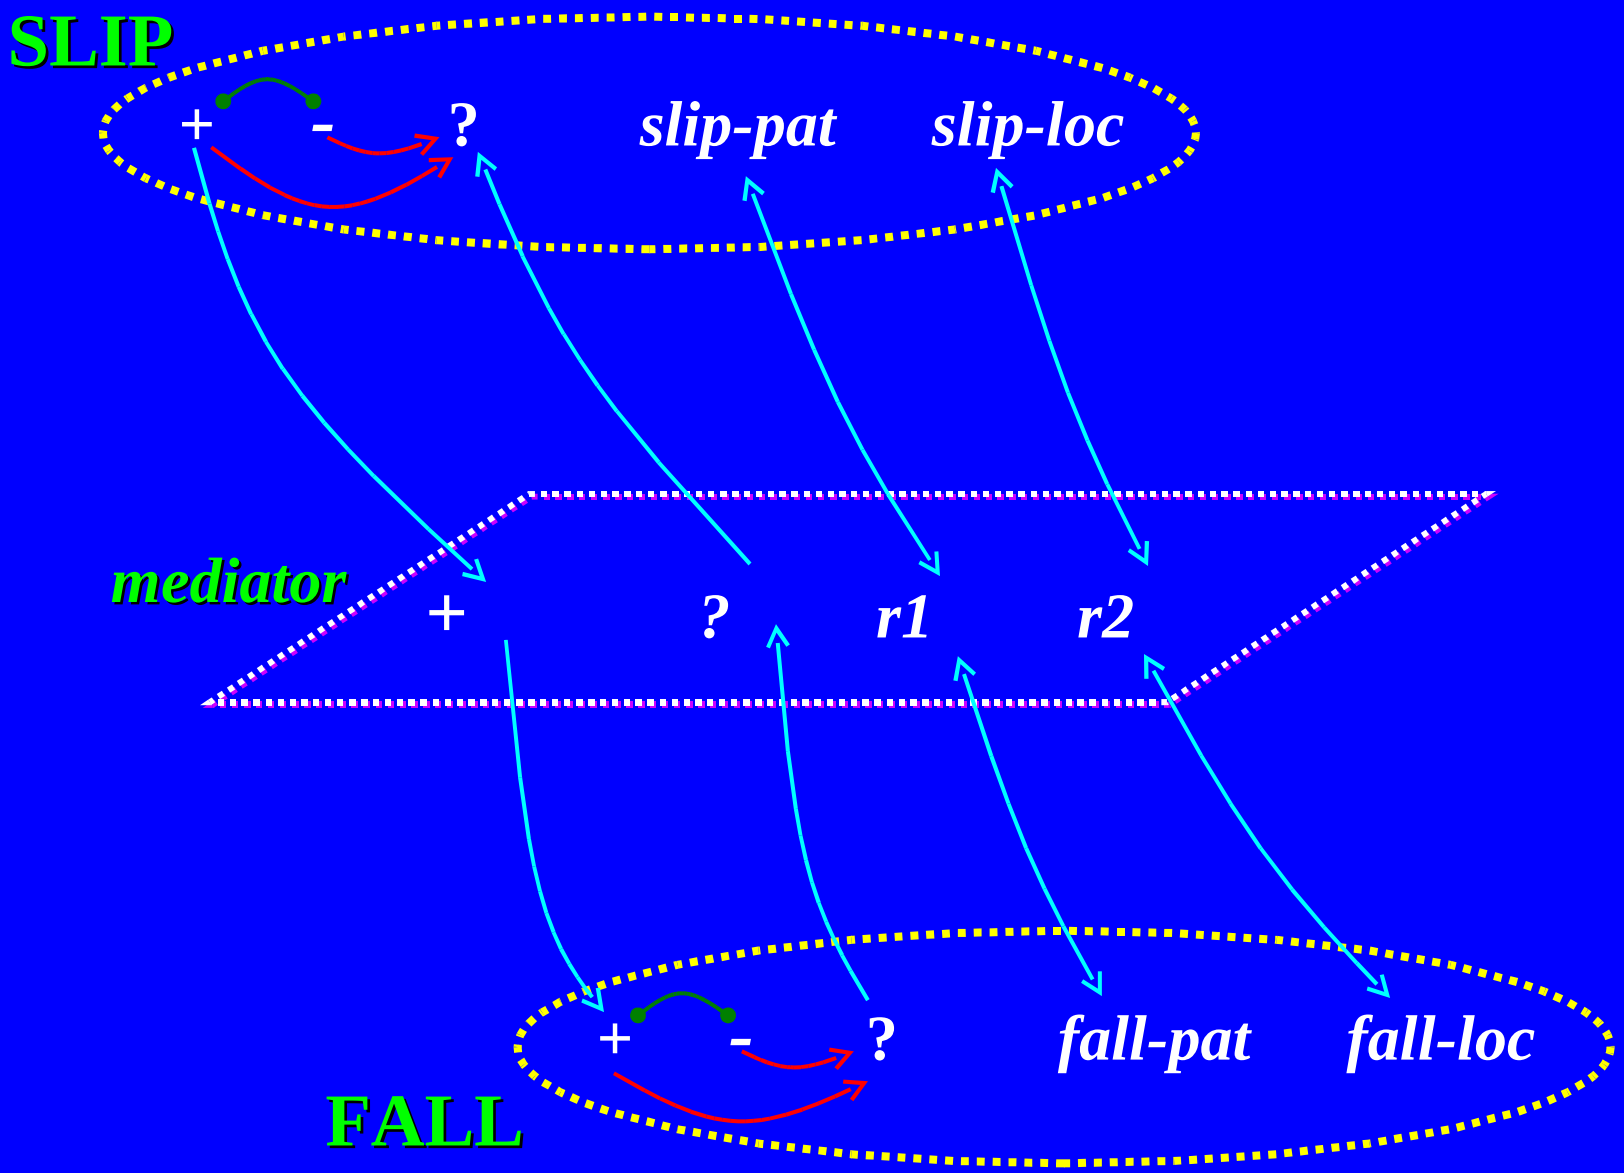

Linking focal-clusters of types and entities

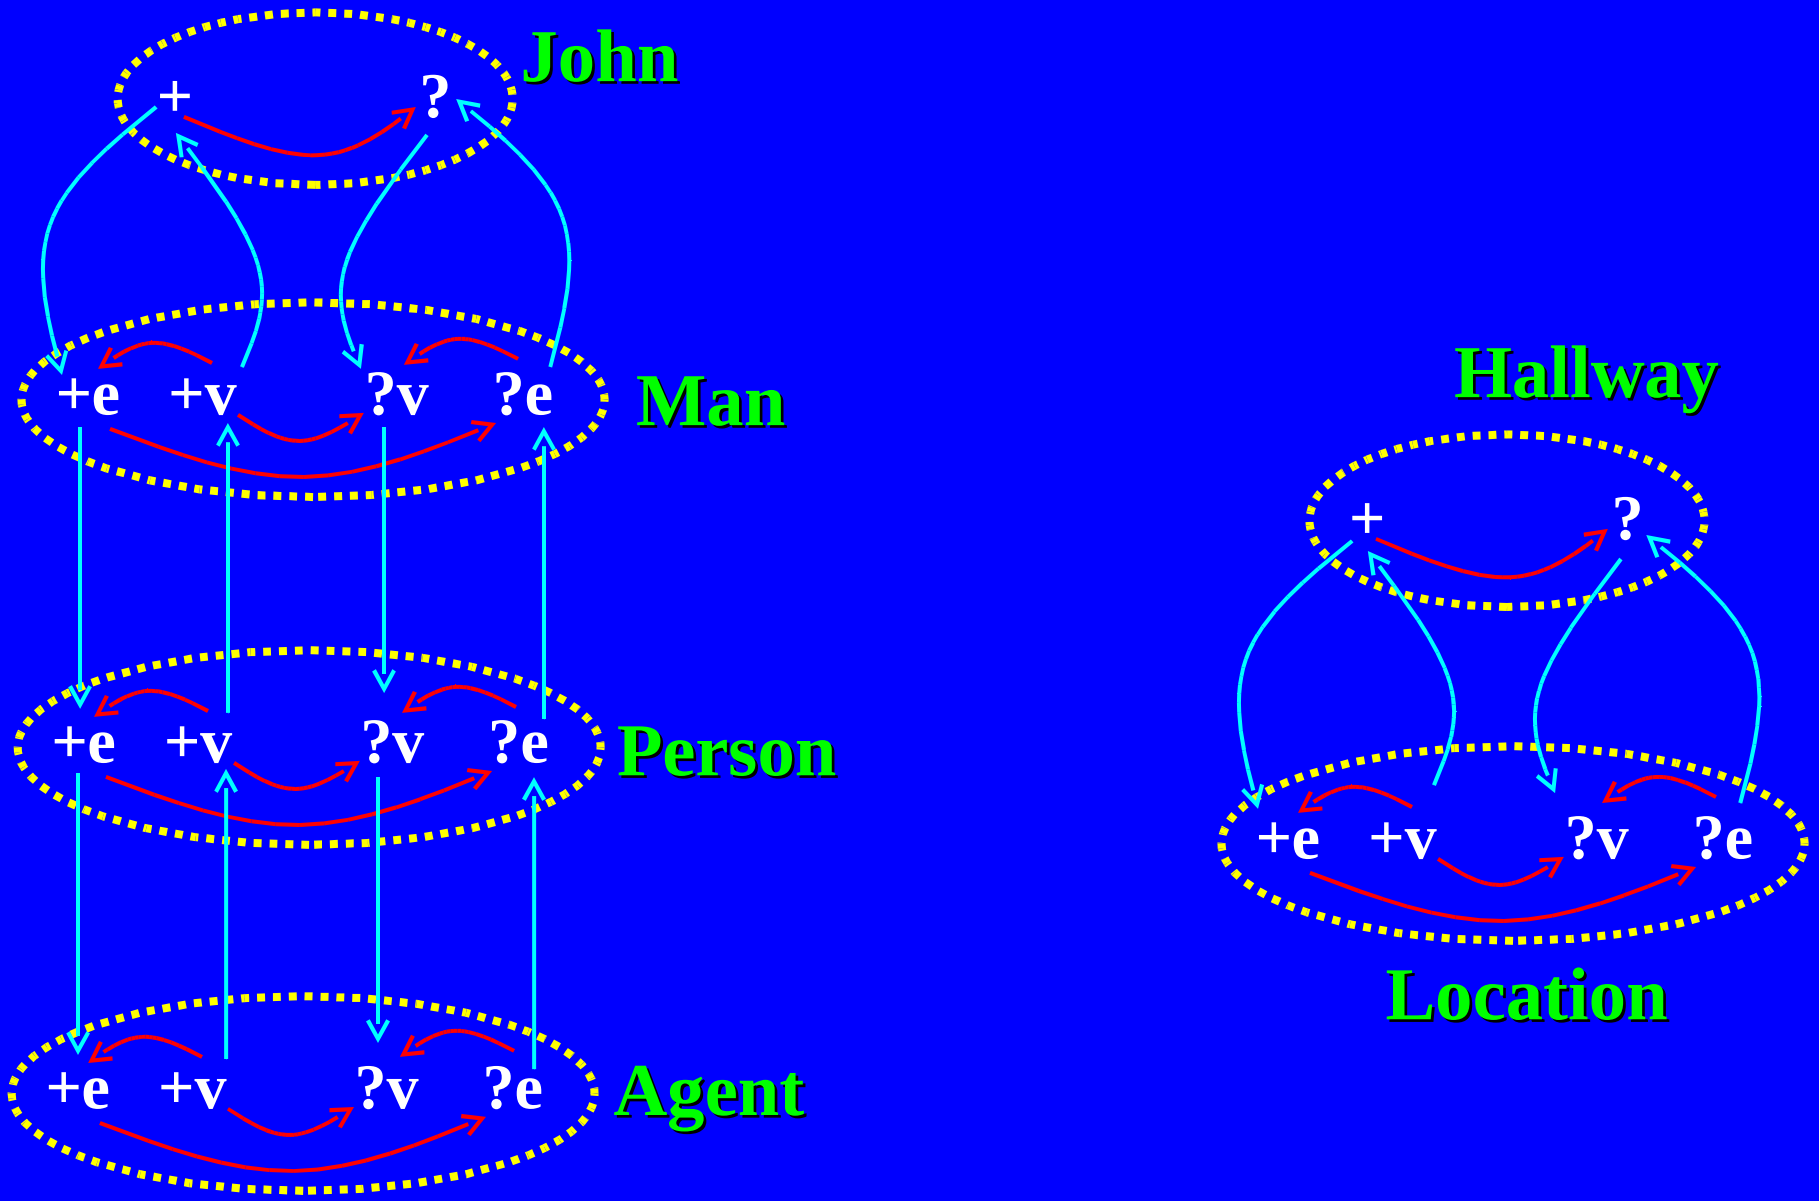

Focal-clusters and context-sensitive priors (T-facts)

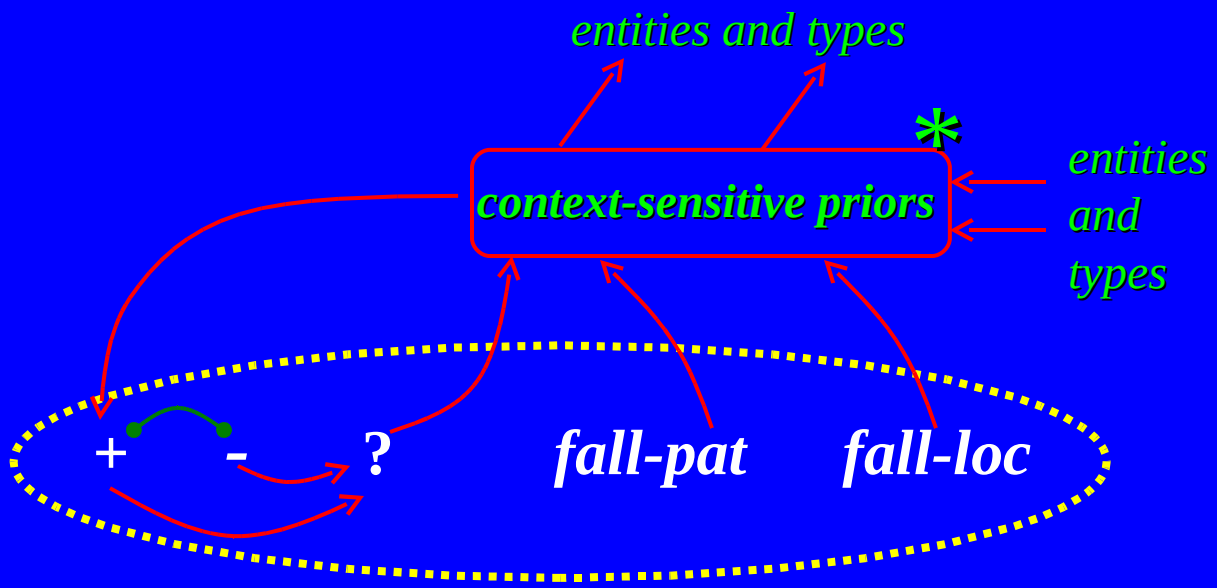

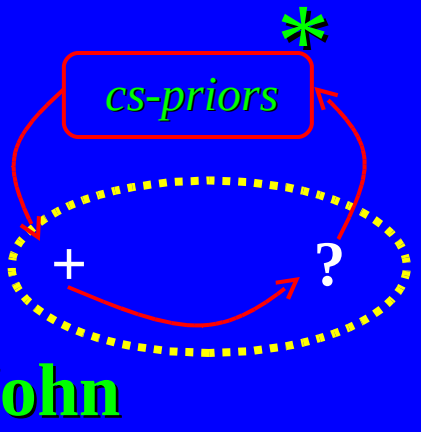

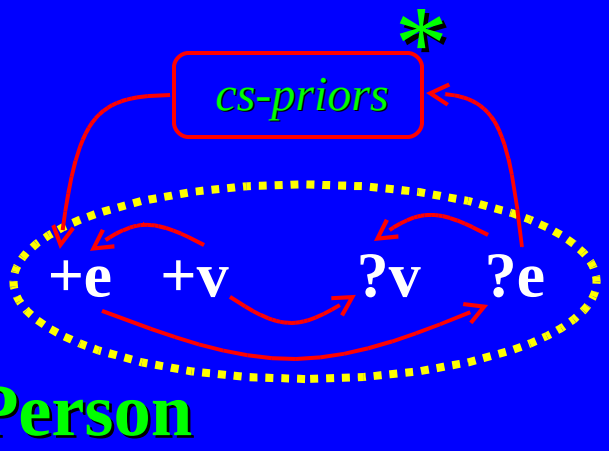

* *cortical circuits*

Focal-clusters and episodic memories (E-facts)

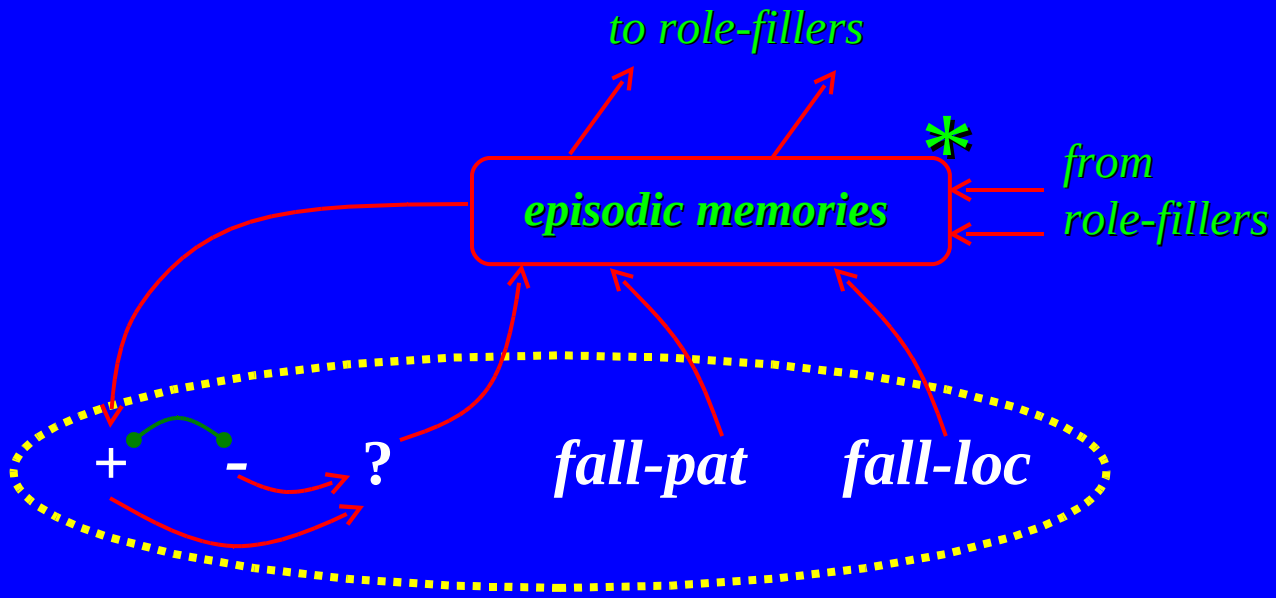

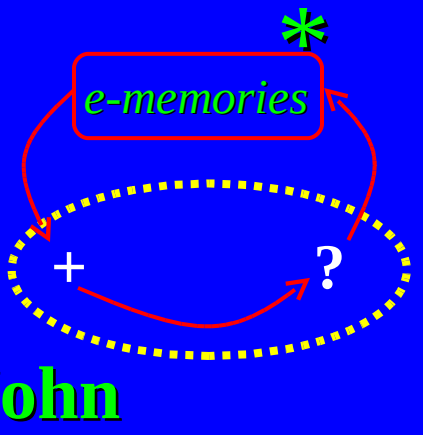

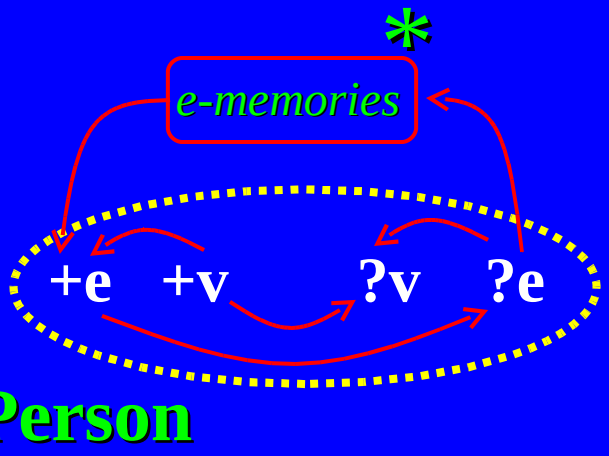

* *hippocampal circuits*

Explaining away in Shruti

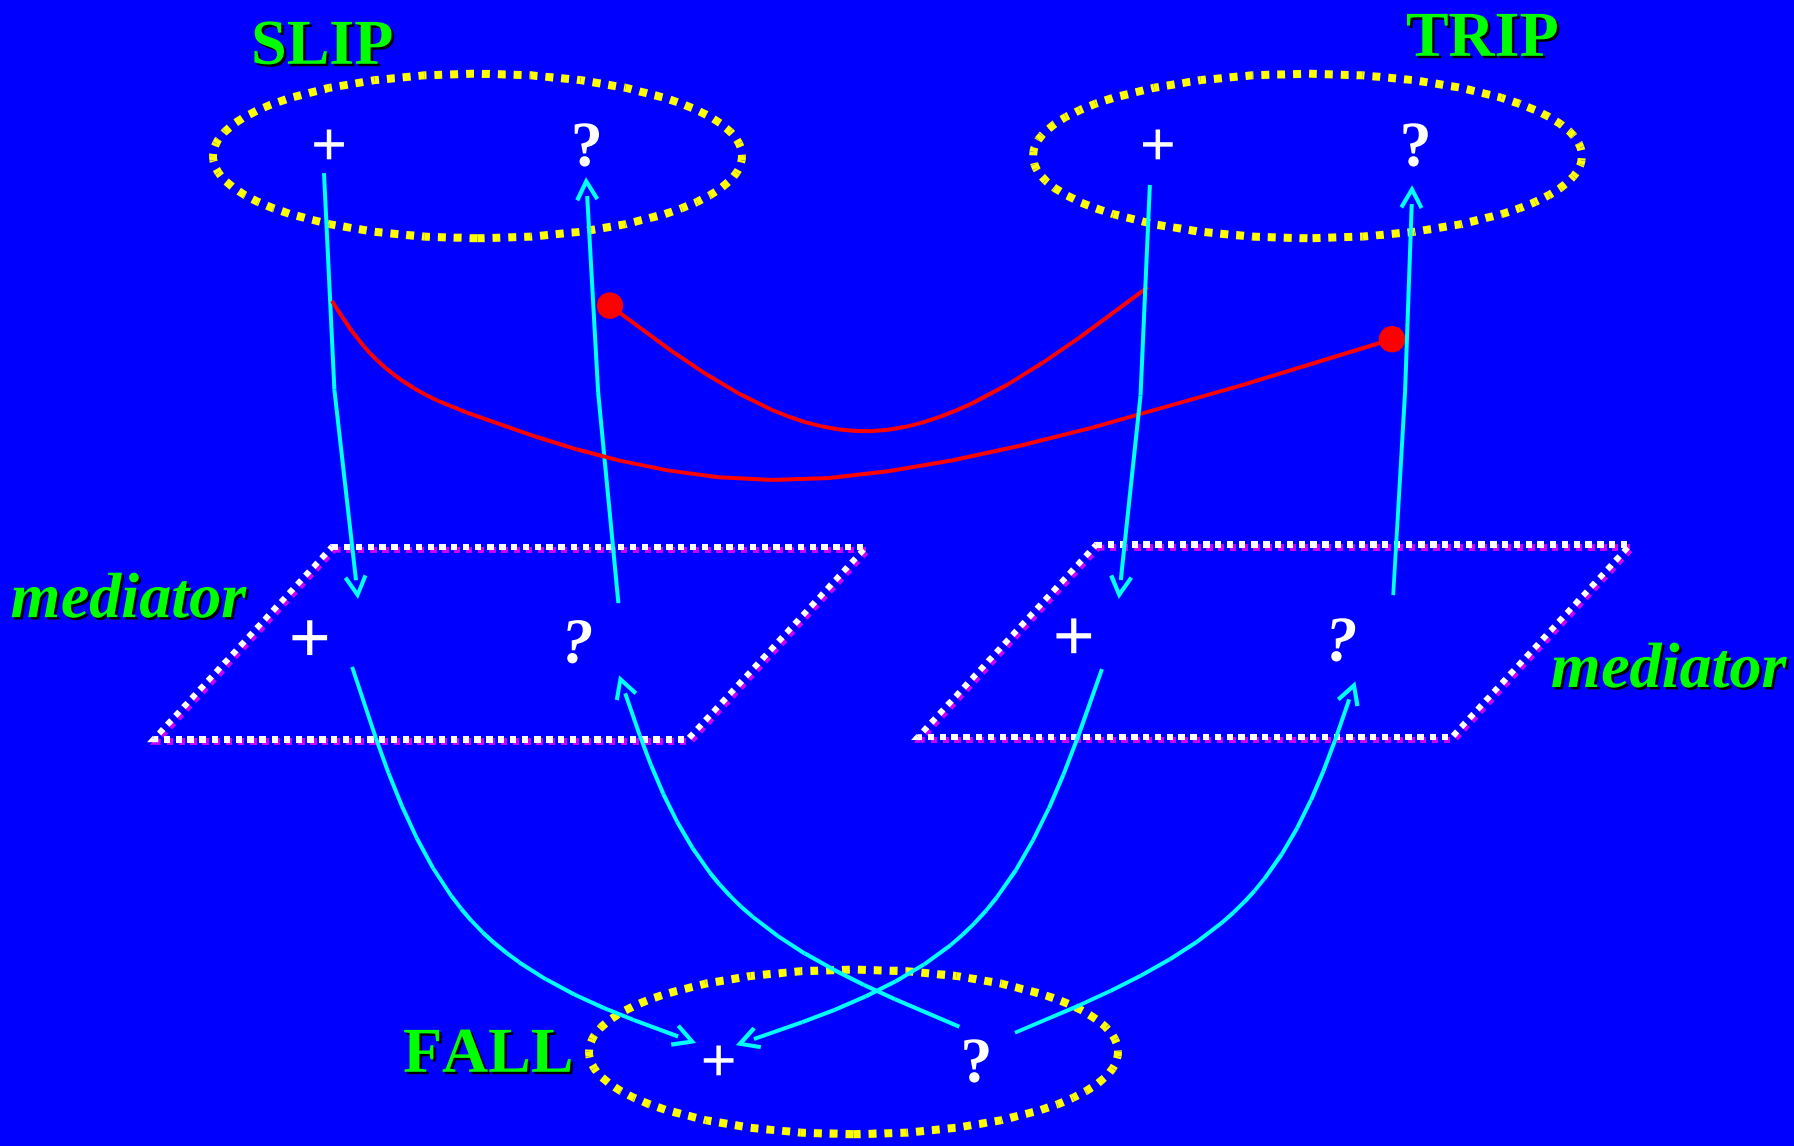

Other features of Shruti

- *Mutual inhibition between collectors of incompatible entities*
- *Merging of phases -- unification*
- *Instantiation of new entities*
- *Structured priming*

Unification in Shruti : merging of phases

The activity in focal-clusters of two entity or relational instances will synchronize if there is evidence that the two instances are the same

R1: ***Is there an entity A of type T filling role r in situation P?***
(Did a man fall in the hallway?)

R2: ***Entity B of type T is filling role r in situation P.***
(John fell in the hallway.)

In such a situation, the firing of A and B will synchronize.

Consequently, A and B will unify, and so will the relational instances involving A and B.

Entity instantiation in Shruti

If Shruti encodes the rule-like knowledge:

x:Agent y:Location fall(x,y) => hurt(x)

it automatically posits the existence of a location where John fell in response to the dynamic instantiation of *hurt(x)*

Encoding "fall => hurt" in Shruti

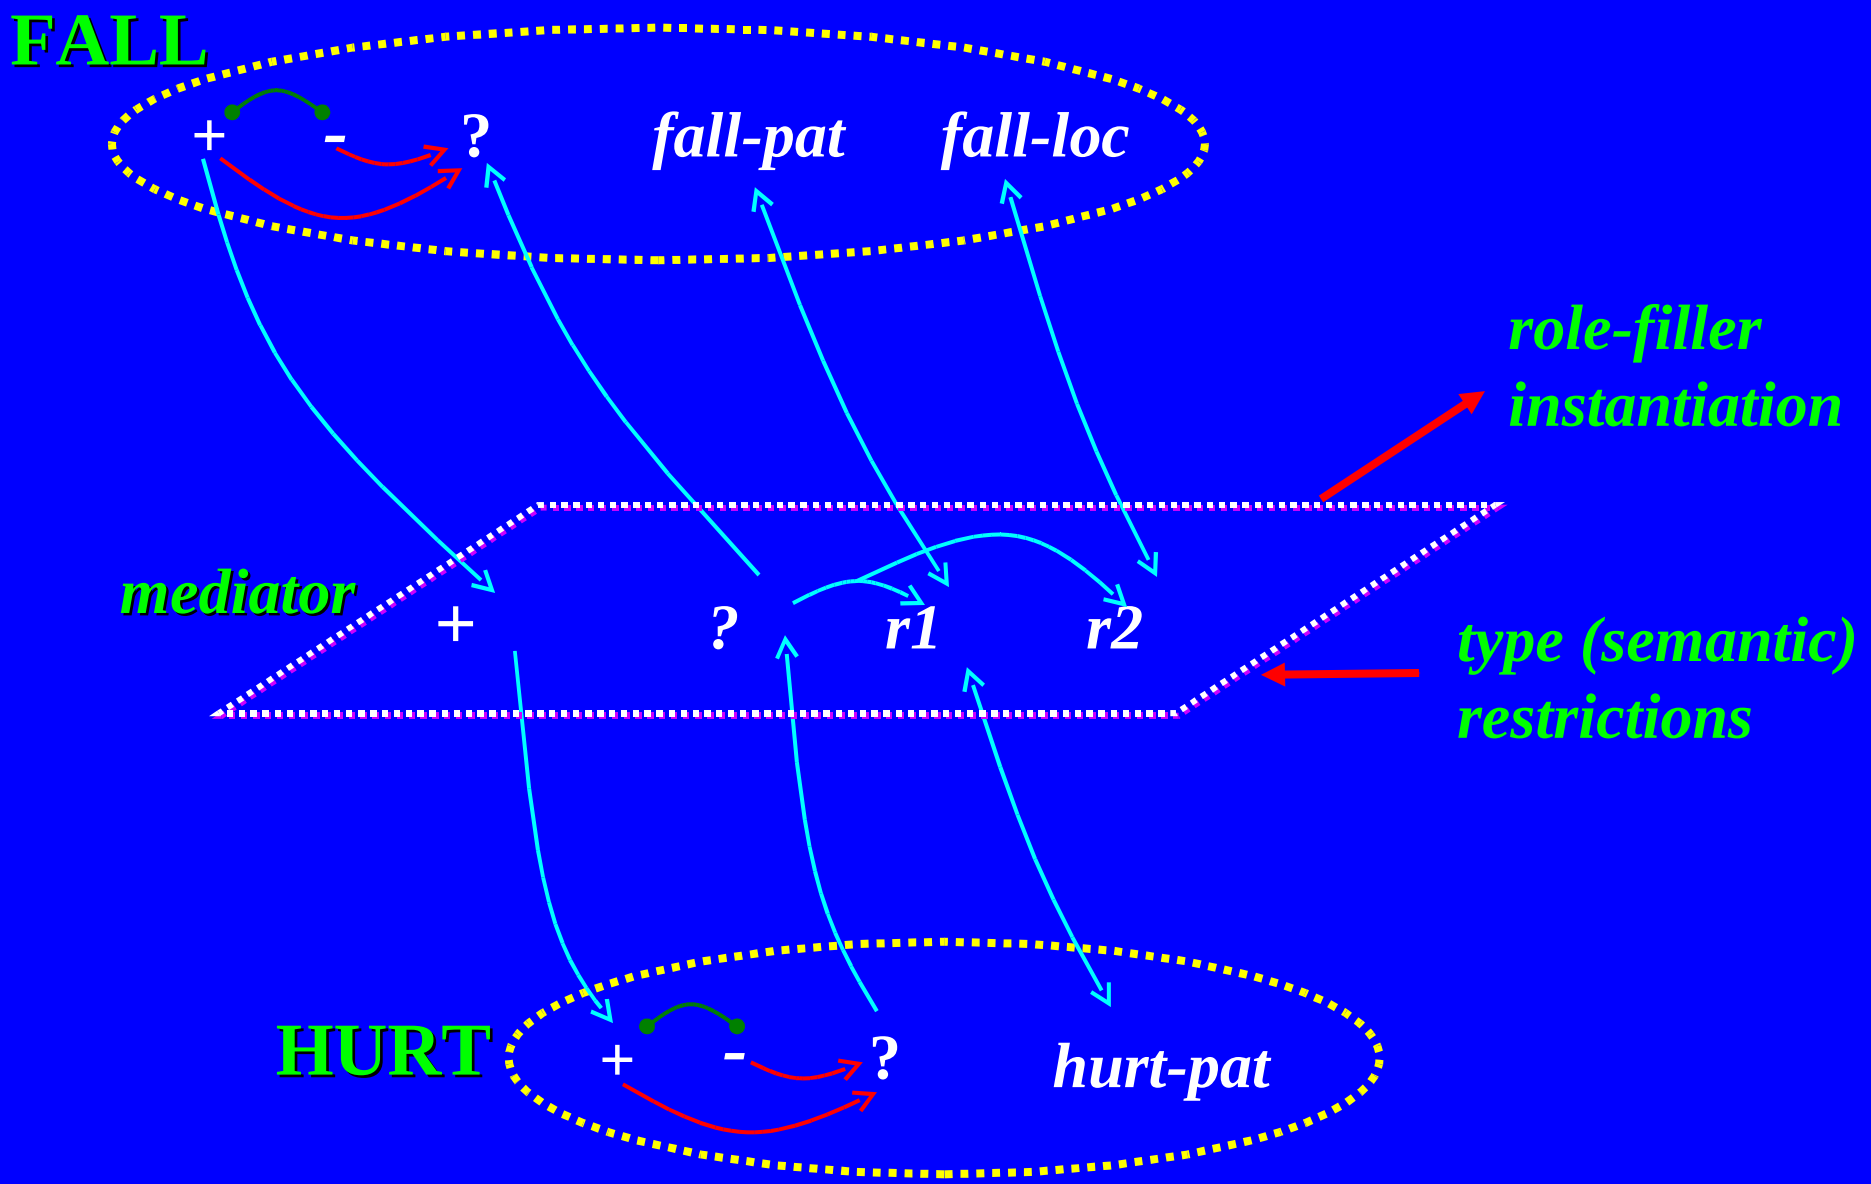

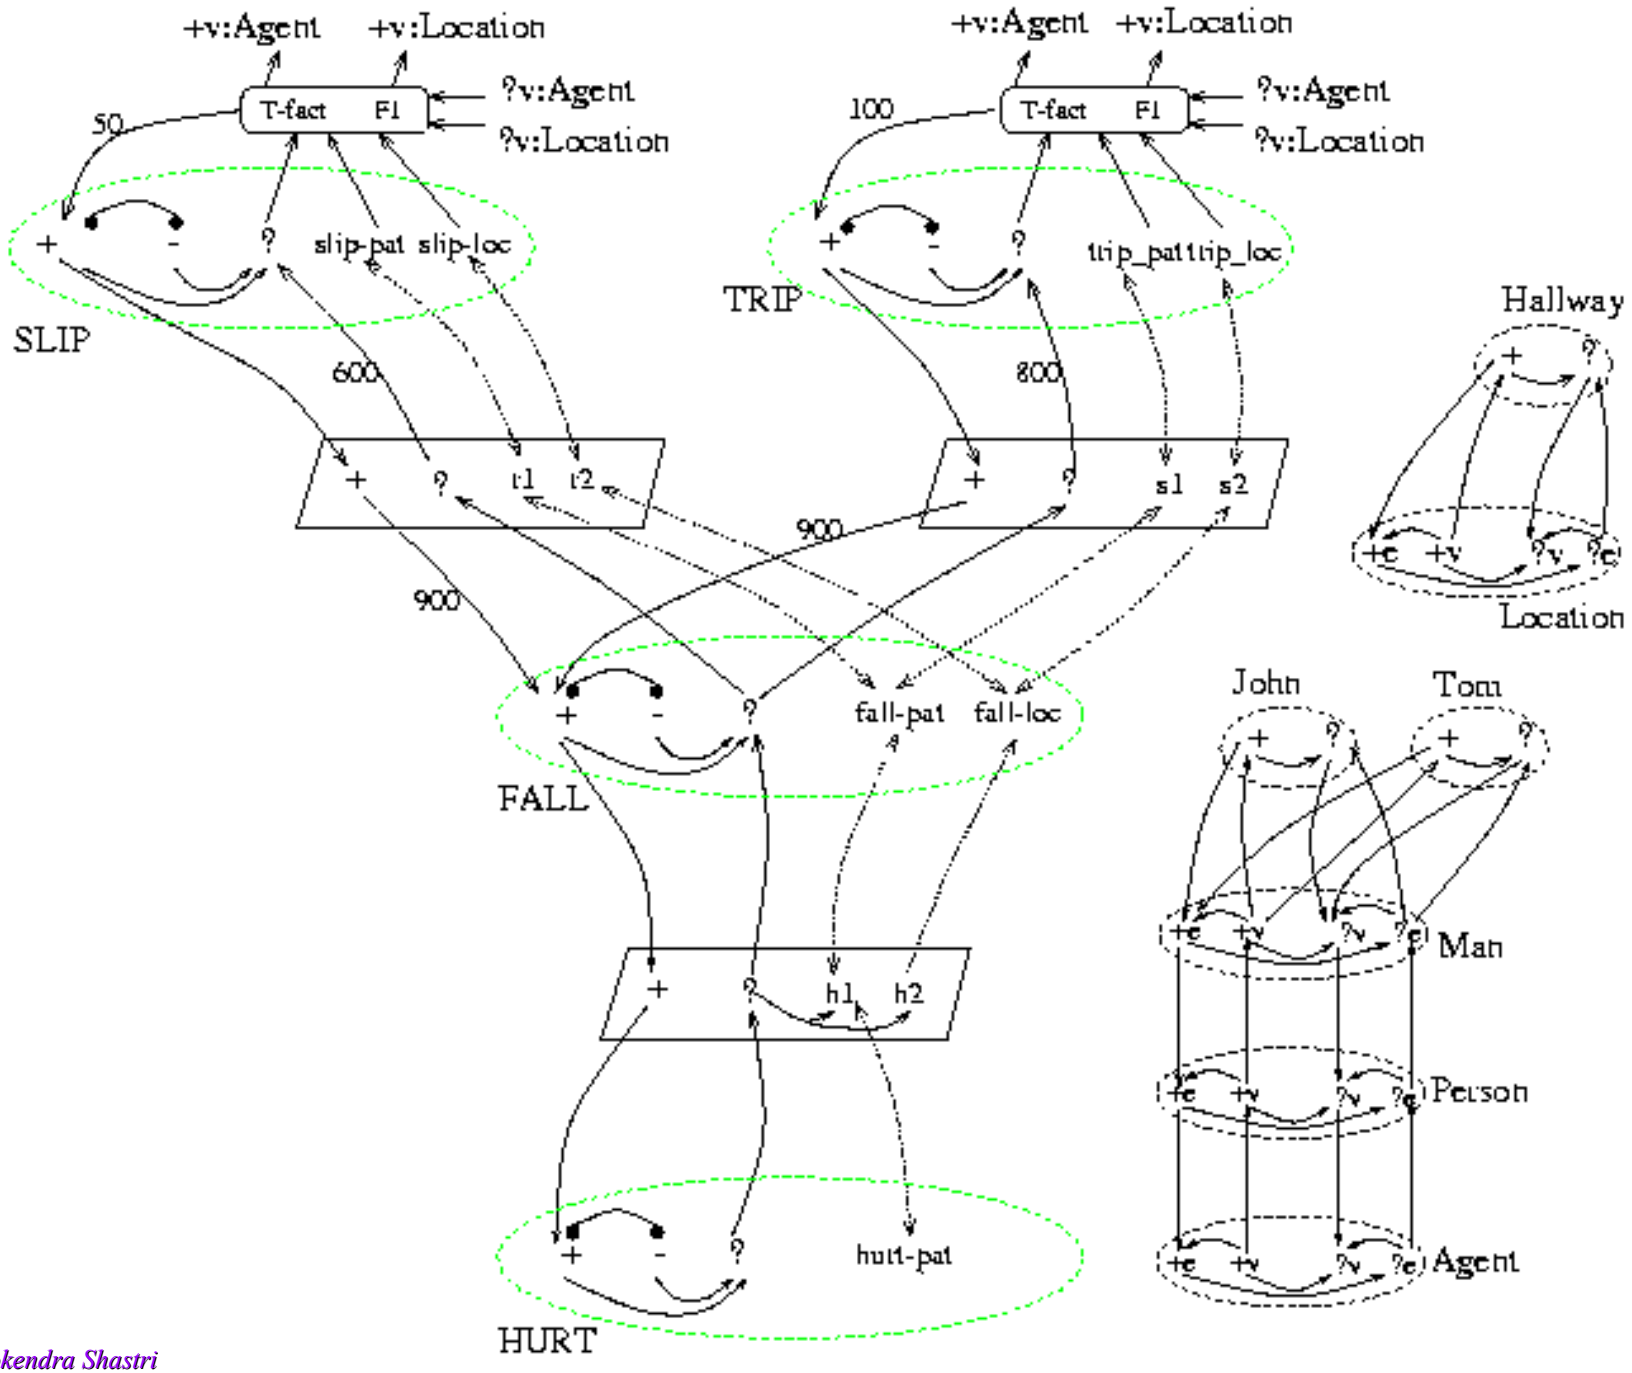

The activation trace of +:slip and +:trip

"John fell in the hallway.
Tom had cleaned it.
He got hurt."

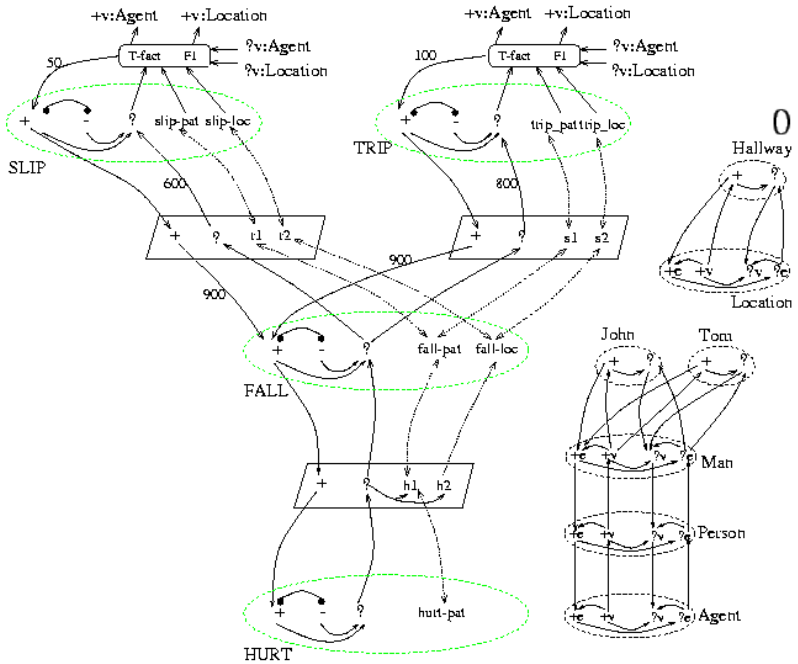

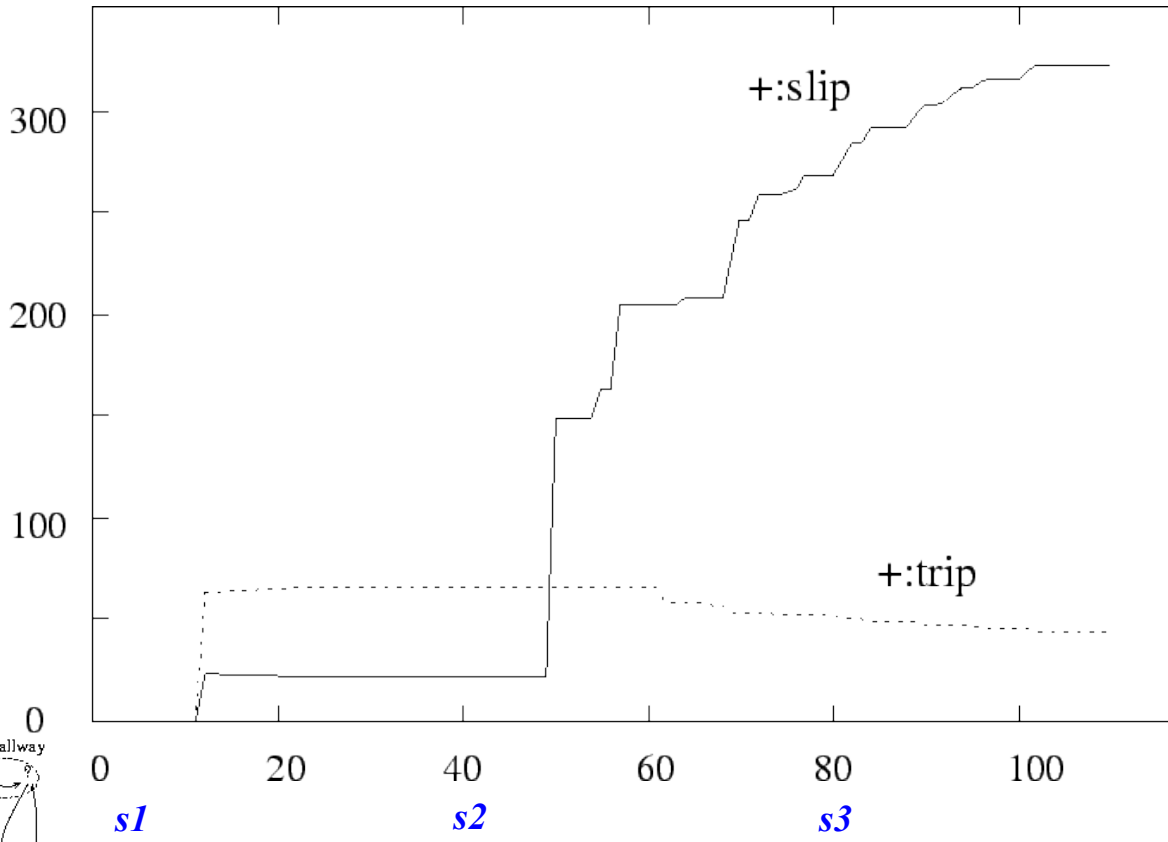

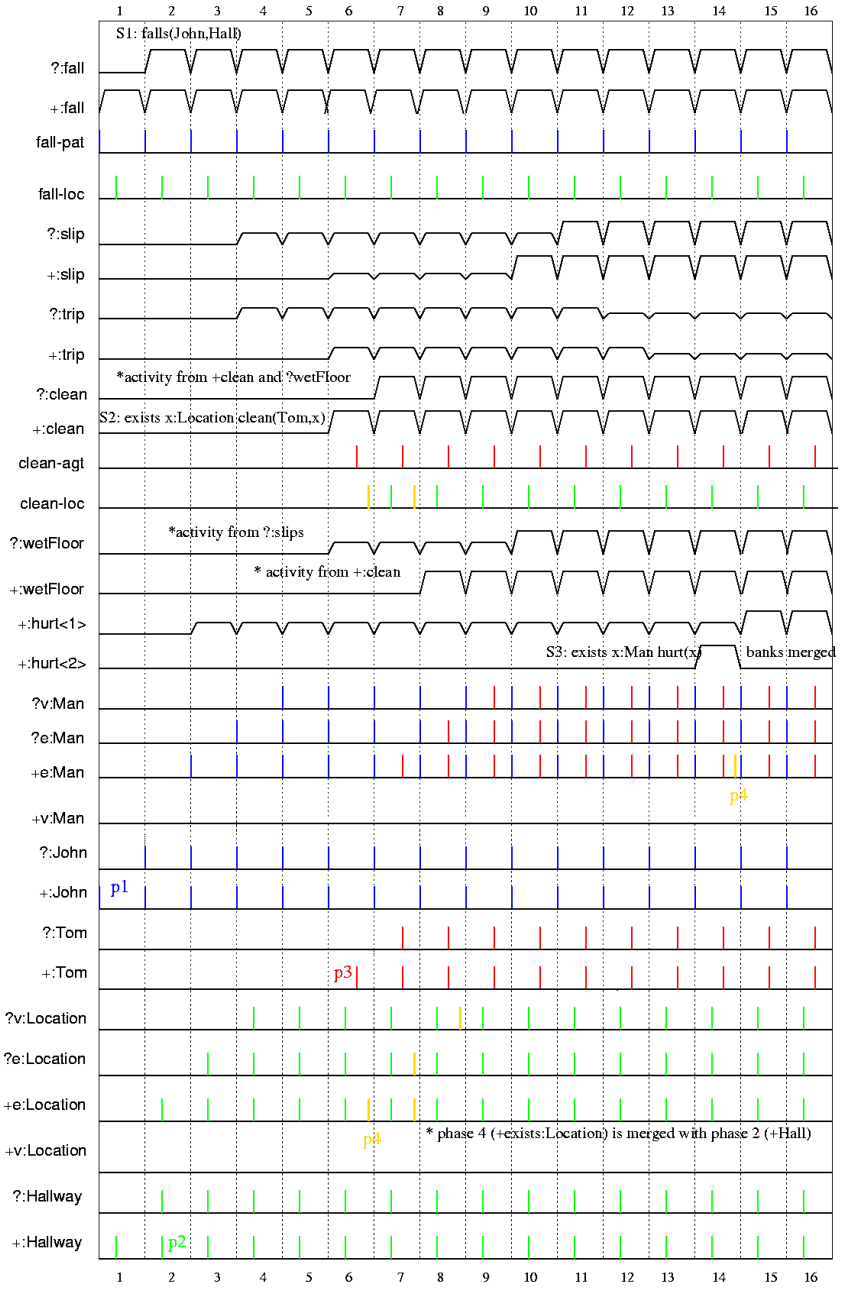

A Metaphor for Reasoning

- *An episode of reflexive reasoning is a transient propagation of rhythmic activity*
- *Each entity involved in this reasoning episode is a phase in this rhythmic activity*
- *Bindings are synchronous firings of cell clusters*
- *Rules are interconnections between cell-clusters that support context-sensitive propagation of activity*
- *Unification corresponds to merging of phases*
- *A stable inference (explanation/answer) corresponds to reverberatory activity around closed loops*

Support for Shruti

- *Neurophysiological evidence: transient synchronization of cell firing might encode dynamic bindings*
- *Makes plausible predictions about working memory limitations*
- *Speed of inference satisfies performance requirements of language understanding*
- *Representational assumptions are compatible with a biologically realistic model of episodic memory*

Neurophysiological evidence for synchrony

- *Synchronous activity found in anesthetized cat as well as in anesthetized and awake monkey.*
- *Spatially distributed cells exhibit synchronous activity if they represent information about the same object.*
- *Synchronous activity occurs in the gamma band (25--60Hz) (maximum period of about 40 msec.)*
- *frequency drifts by 5-10Hz, but synchronization stays stable for 100-300 msec*
- *In humans EEG and MEG signals exhibit power spectrum shifts consistent with synchronization of cell ensembles*
  - *orienting or investigatory behavior; delayed-match-to- sample task; visuo-spatial working memory task*

Predictions: constraints on reflexive inference

- *gamma band activity (25-60Hz) underlies dynamic bindings (the maximum period ~40 msec.)*
- *allowable jitter in synchronous firing 3 msec. lead/lag.*

⇒ **only a small number of distinct conceptual entities can participate in an episode of reasoning**

7 +/- 2 *(40 divided by 6)*

as the number of entities increases beyond five, their activity starts overlapping, leading to cross-talk

Note: Not a limit on the number of co-active bindings!

Predictions: Constraints on reflexive reasoning

1. *A large number of relational instances (facts) can be co-active, and numerous rules can fire in parallel, but*
2. *only a small number of distinct entities can serve as role-fillers in this activity*
3. *only a small number of instances of the same predicate can be co-active at the same time*
4. *the depth of inference is bounded – systematic reasoning via binding propagation degrades to a mere spreading of activation beyond a certain depth.*

2 and 3 specify limits on Shruti's working memory

Massively Parallel Inference

- *if gamma band activity underlies propagation of bindings*
- *each binding propagation step takes ca. 25 msec.*
- *inferring "John may be hurt" and "John may have slipped" from "John fell" would take only ca. 200 msec.*
- *time required to perform inference is independent of the size of the causal model*

Probabilistic interpretation of link weights

$P(C/E) = P(E/C)\ P(C)/P(E)$

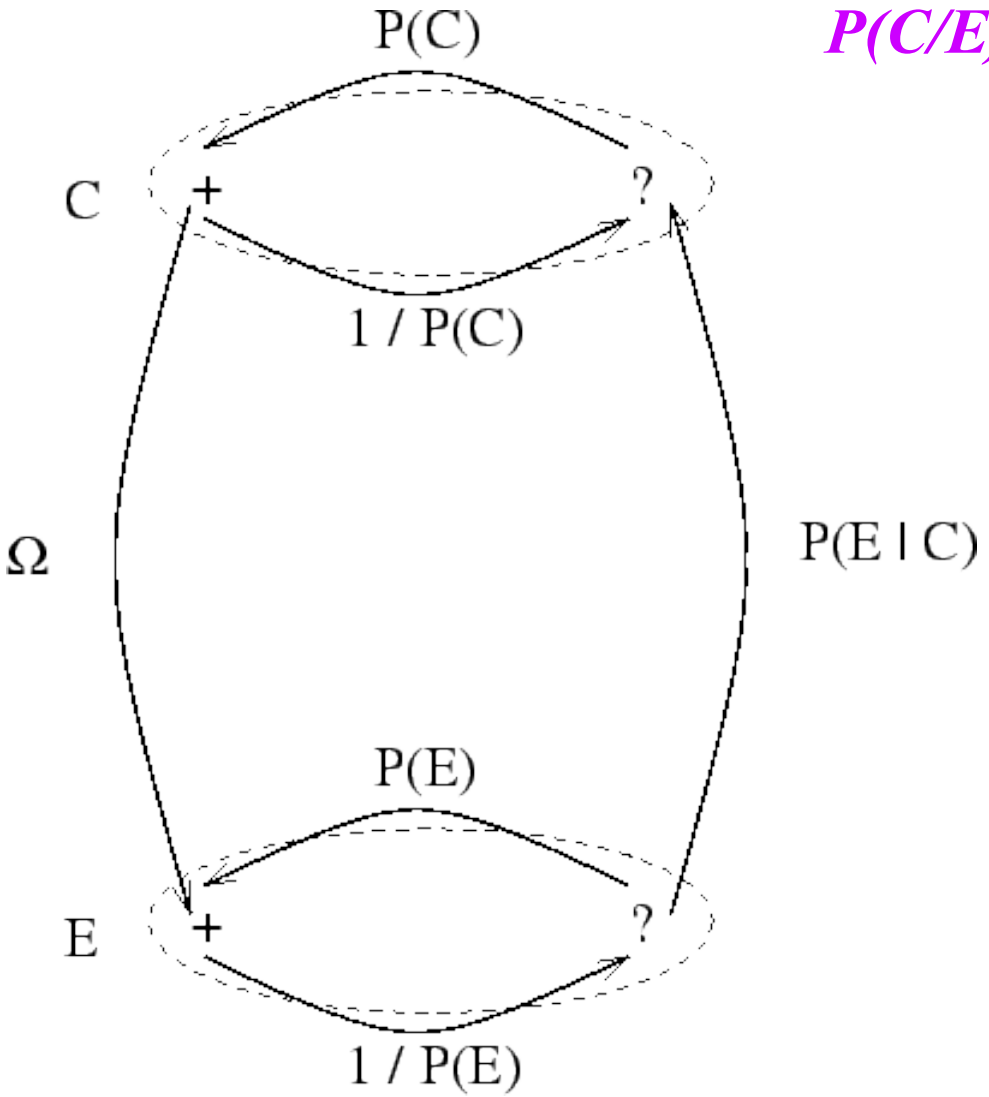

Evidence Combination

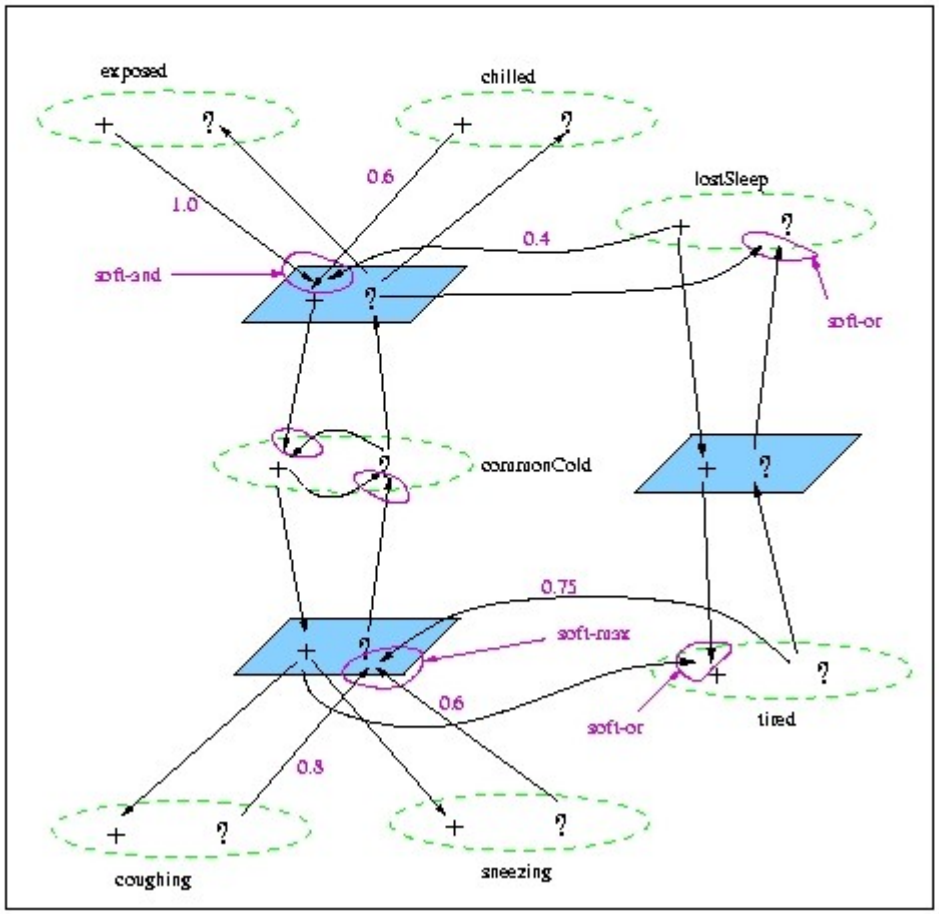

Noisy-Or

$1 - \mathrm{PROD}(1 - X_i W_i)$

Soft-Or

Or / And

$k \to \infty$ Max

Soft-Max $k > 1$

$$\left(\frac{\mathrm{SUM}(X_i^k W_i)}{\mathrm{SUM}(W_i)} \right)^{1/k}$$

Weighted Average

Soft-Min $k < 1$

$k \to 0$ Min

Soft-And

1 - SoftOr

Noisy-And

$\mathrm{PROD}(1 - (1 - X_i) W_i)$

Encoding X-schema

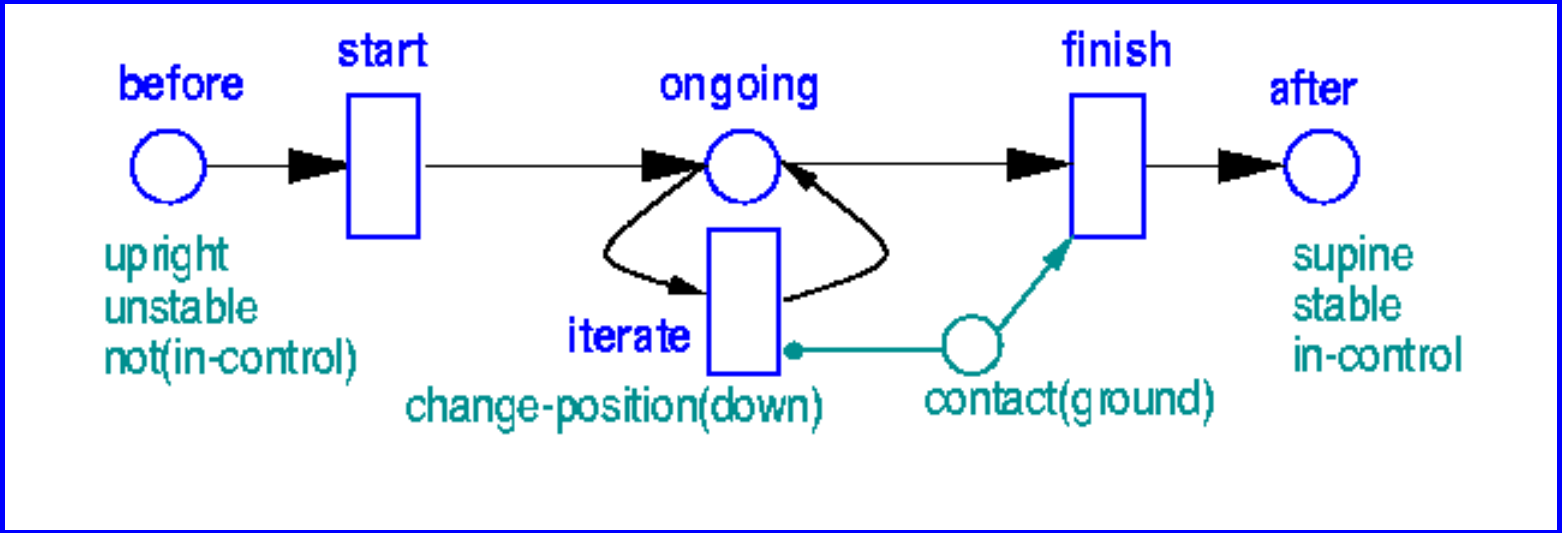

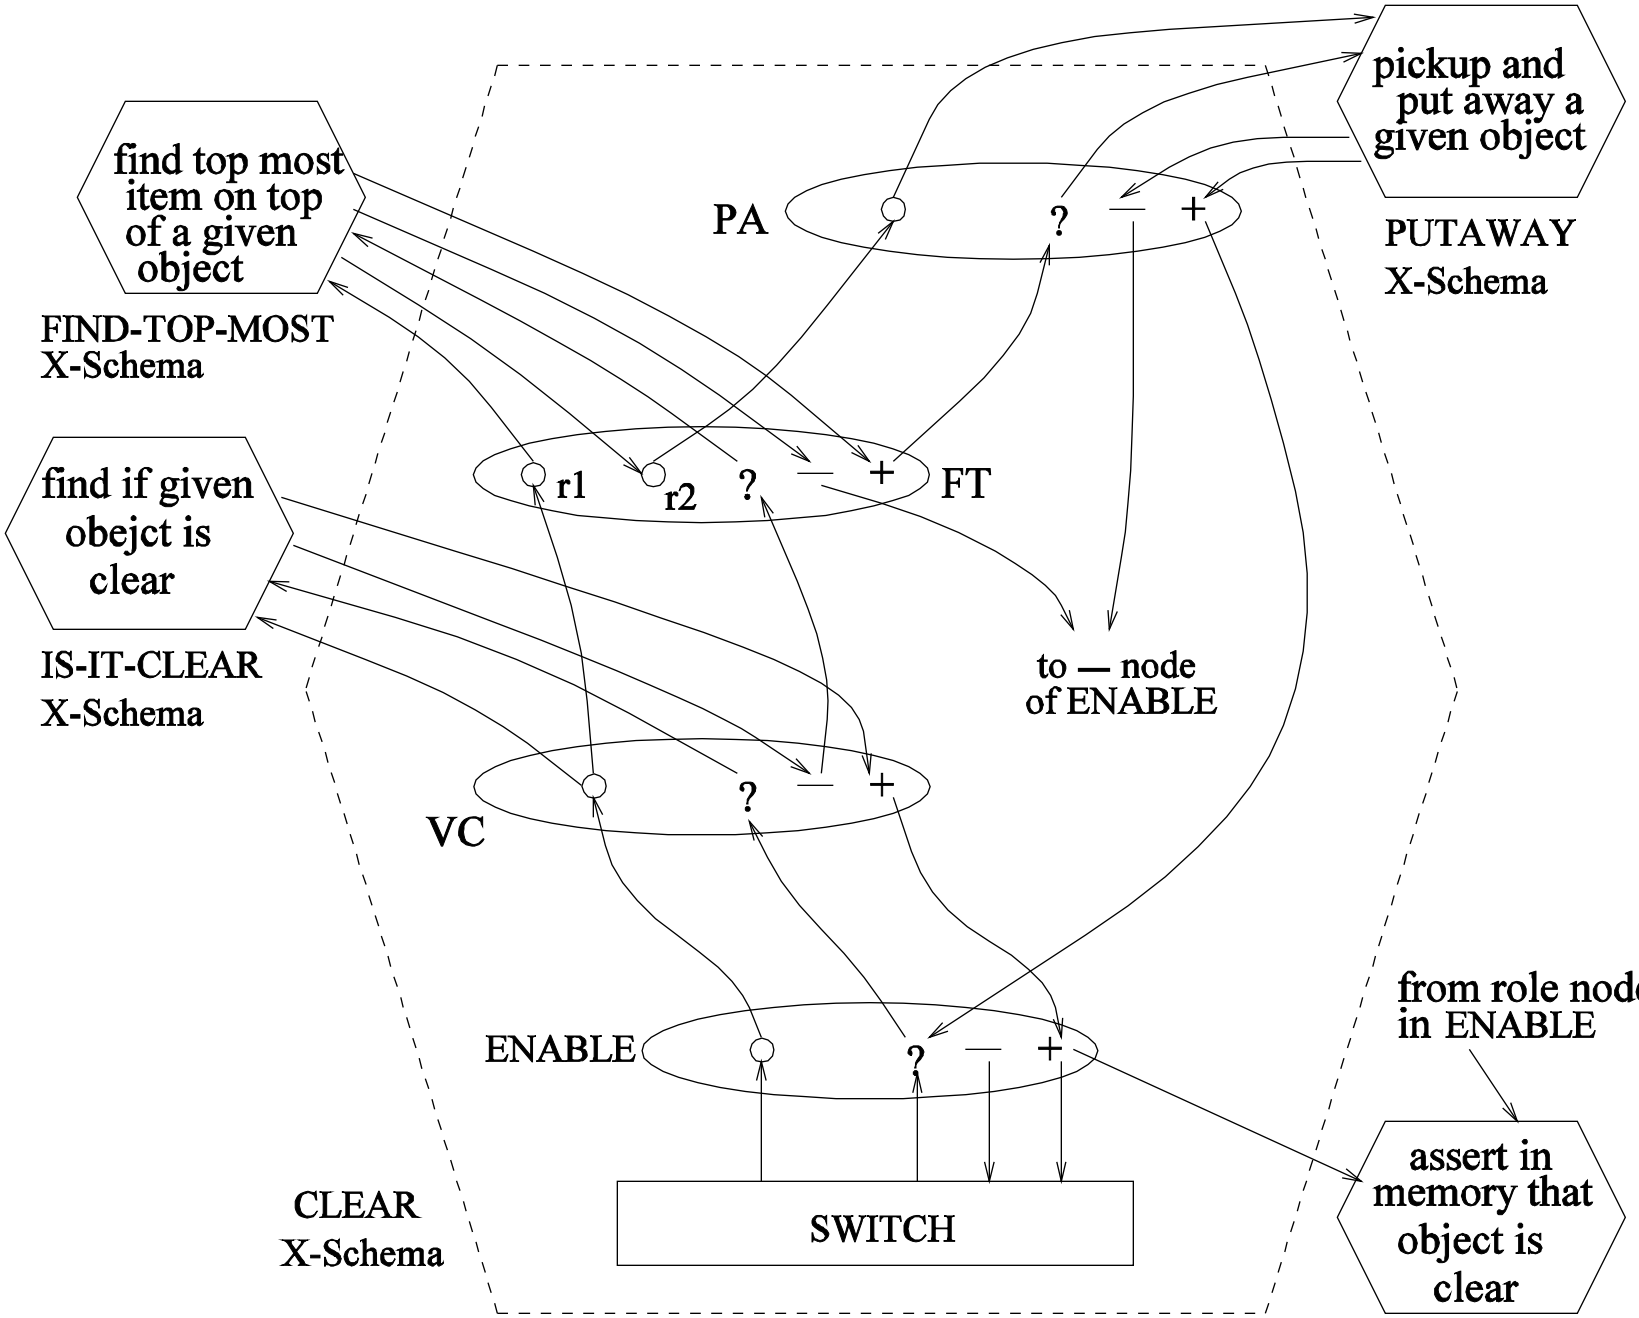

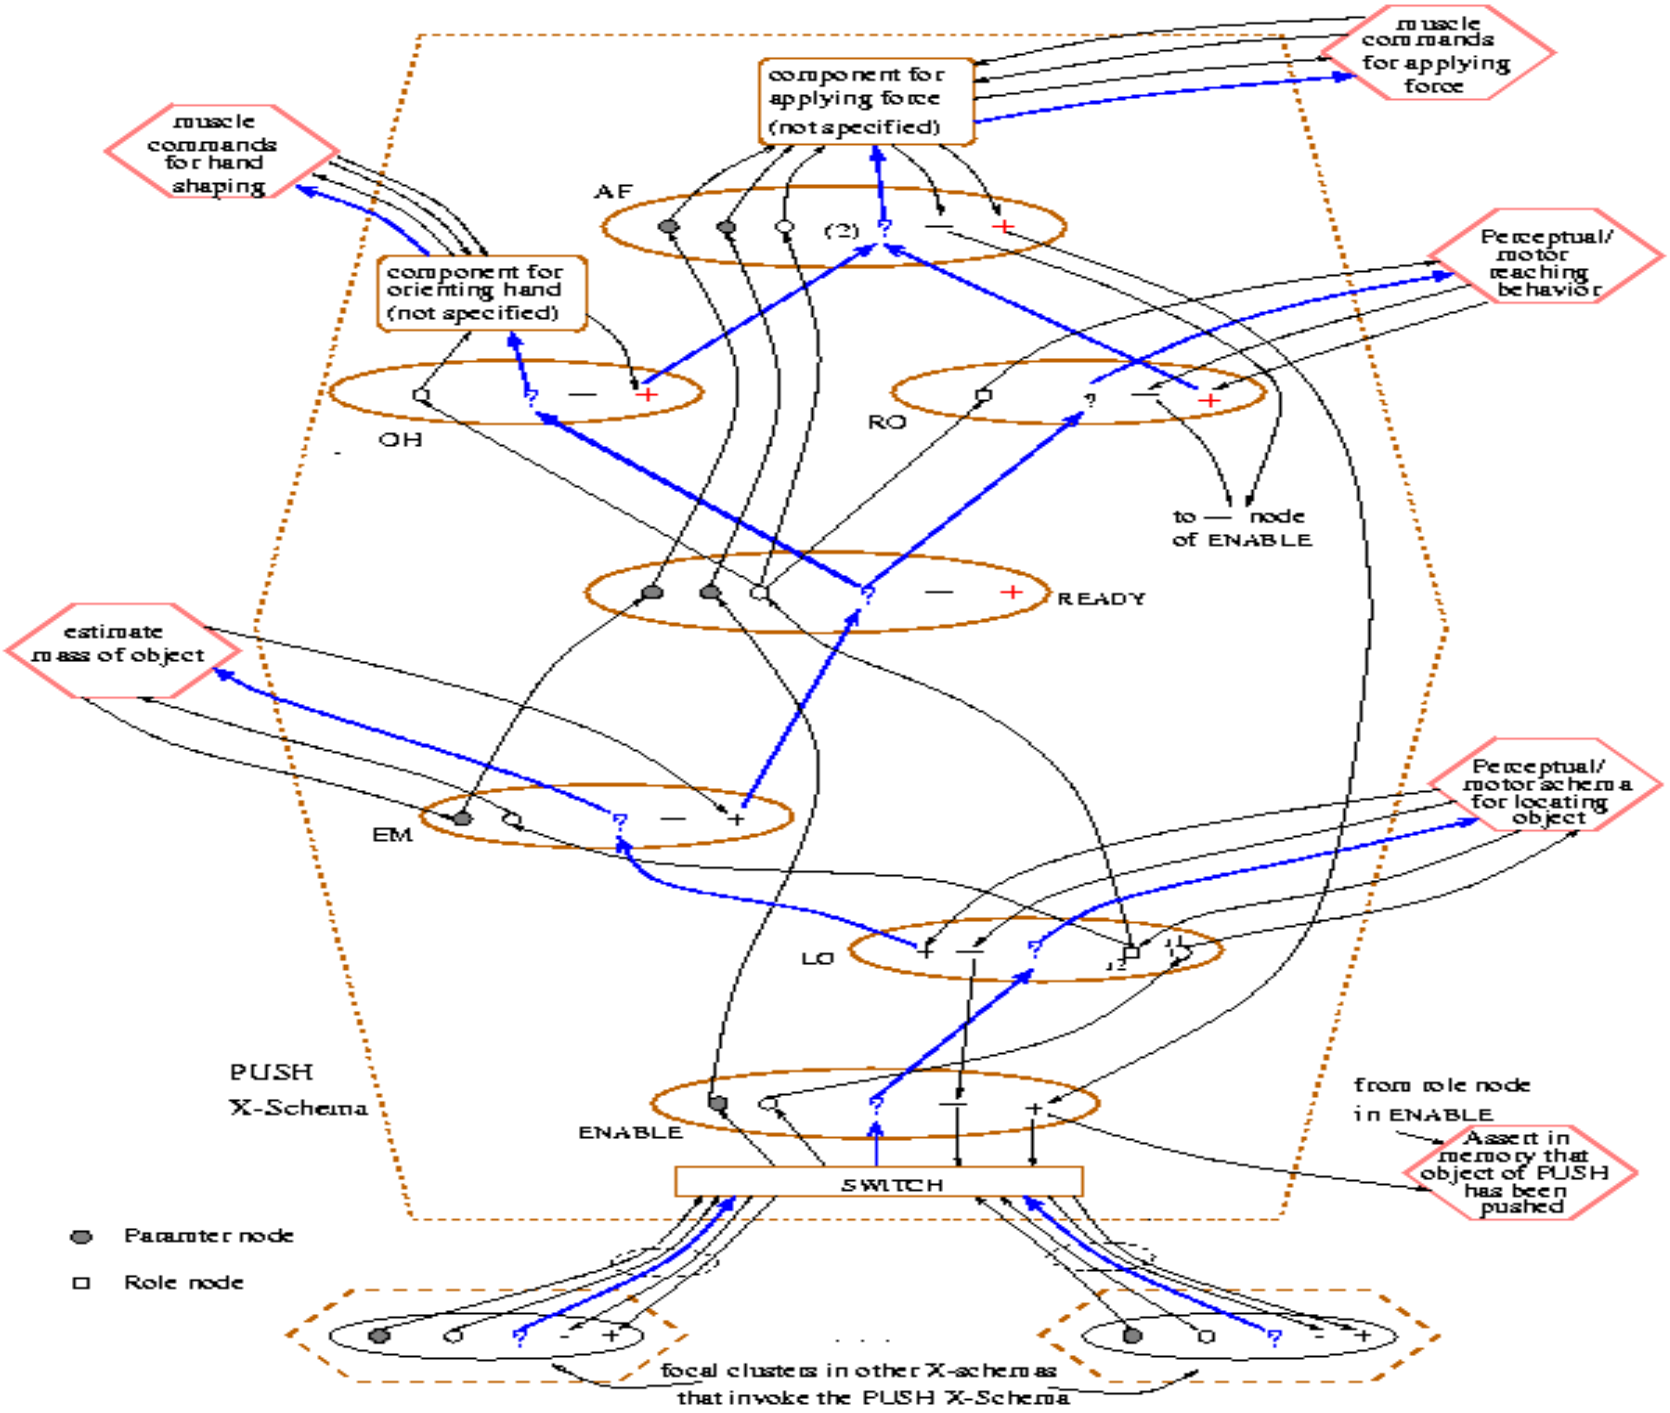
muscle commands for applying force
component for applying force (not specified)
muscle commands for hand shaping
AF
(2)
component for orienting hand (not specified)
Perceptual/ motor reaching behavior
OH
RO
to — node of ENABLE
READY
estimate mass of object
EM
Perceptual/ motor schema for locating object
LO
PUSH
X-Schema
ENABLE
from role node in ENABLE
Assert in memory that object of PUSH has been pushed
SWITCH
Parameter node
Role node
focal clusters in other X-schemas that invoke the PUSH X-Schema

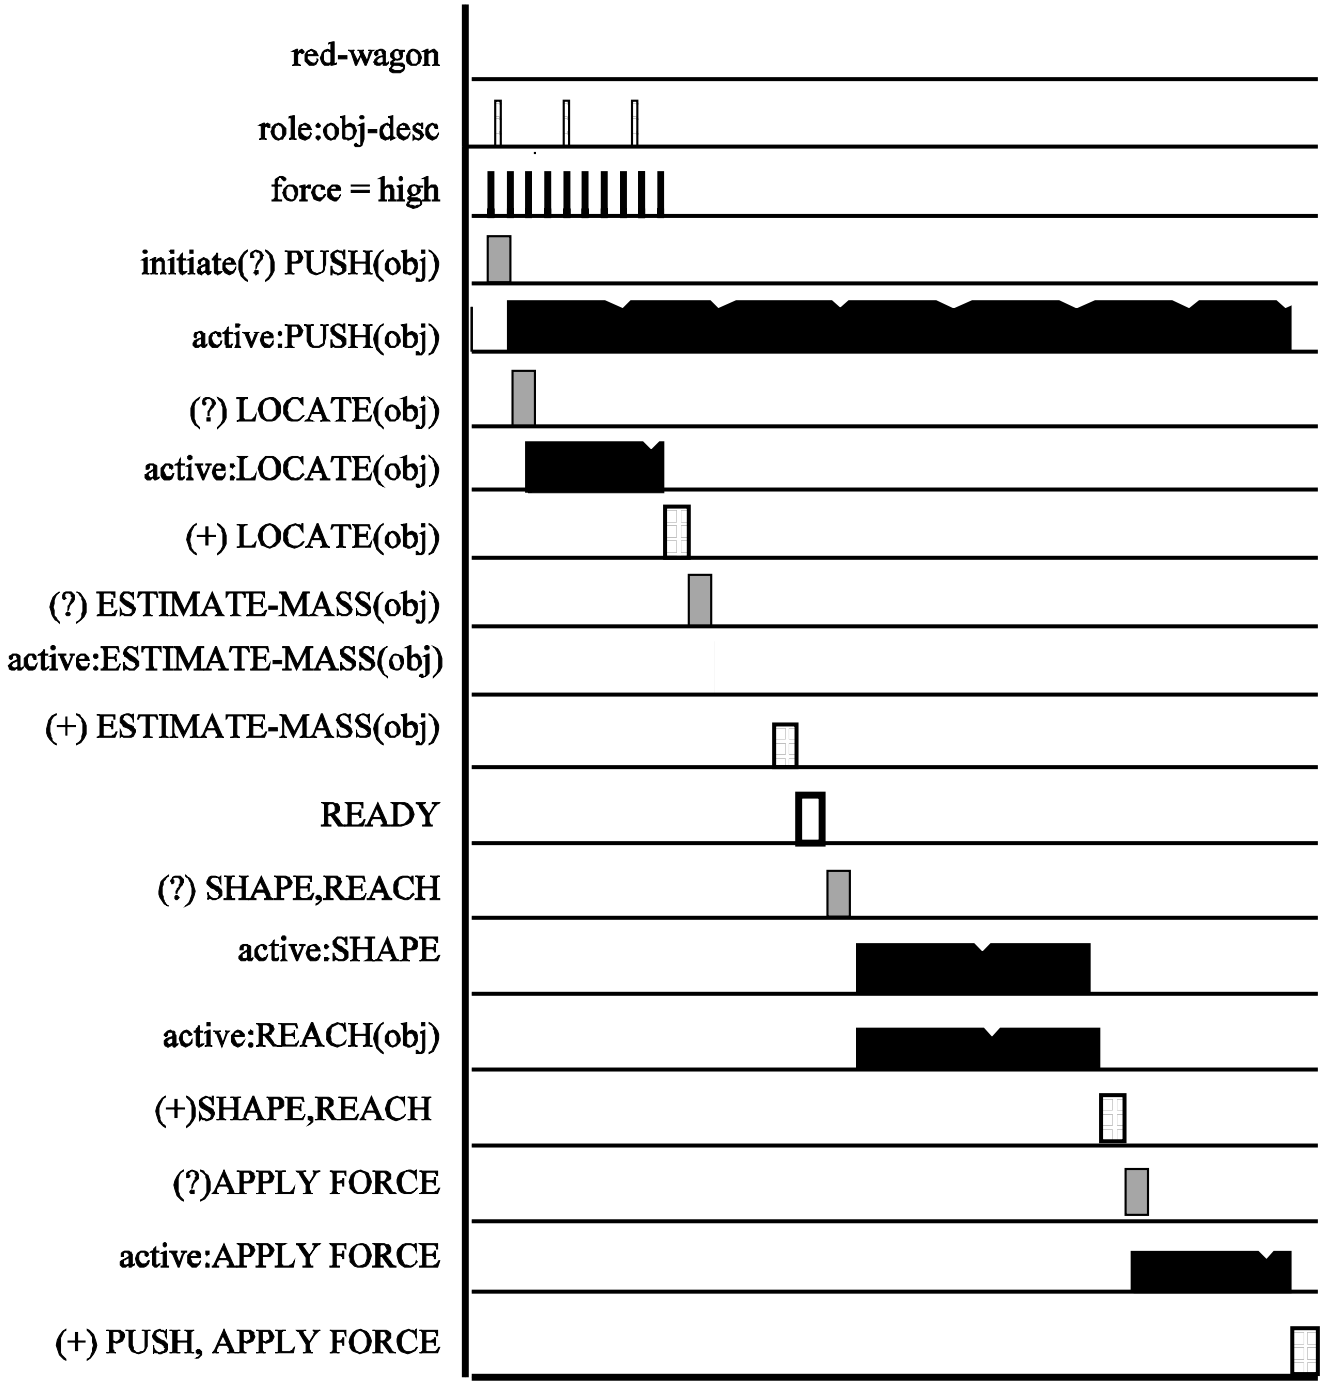

Proposed Alternative Solution

- **Indirect references**
  - **Pass short signatures, "fluents"**
    - Functionally similar to SHRUTI's time slices
  - **Central "binder" maps fluents to objects**
    - In SHRUTI, the objects fired in that time slice
  - **Connections need to be more complicated than in SHRUTI**
    - Fluents are passed through at least 3 bits
    - But temporal synchrony is not required

Components of the System

- Object references
  - Fluents
  - Binder
- Short term storage
  - Predicate state
- Long term storage
  - Facts, mediators, what predicates exist
- Inference
  - Mediators
- Types
  - Ontology

Fluents:

- Roles are just patterns of activation 3-4 bits

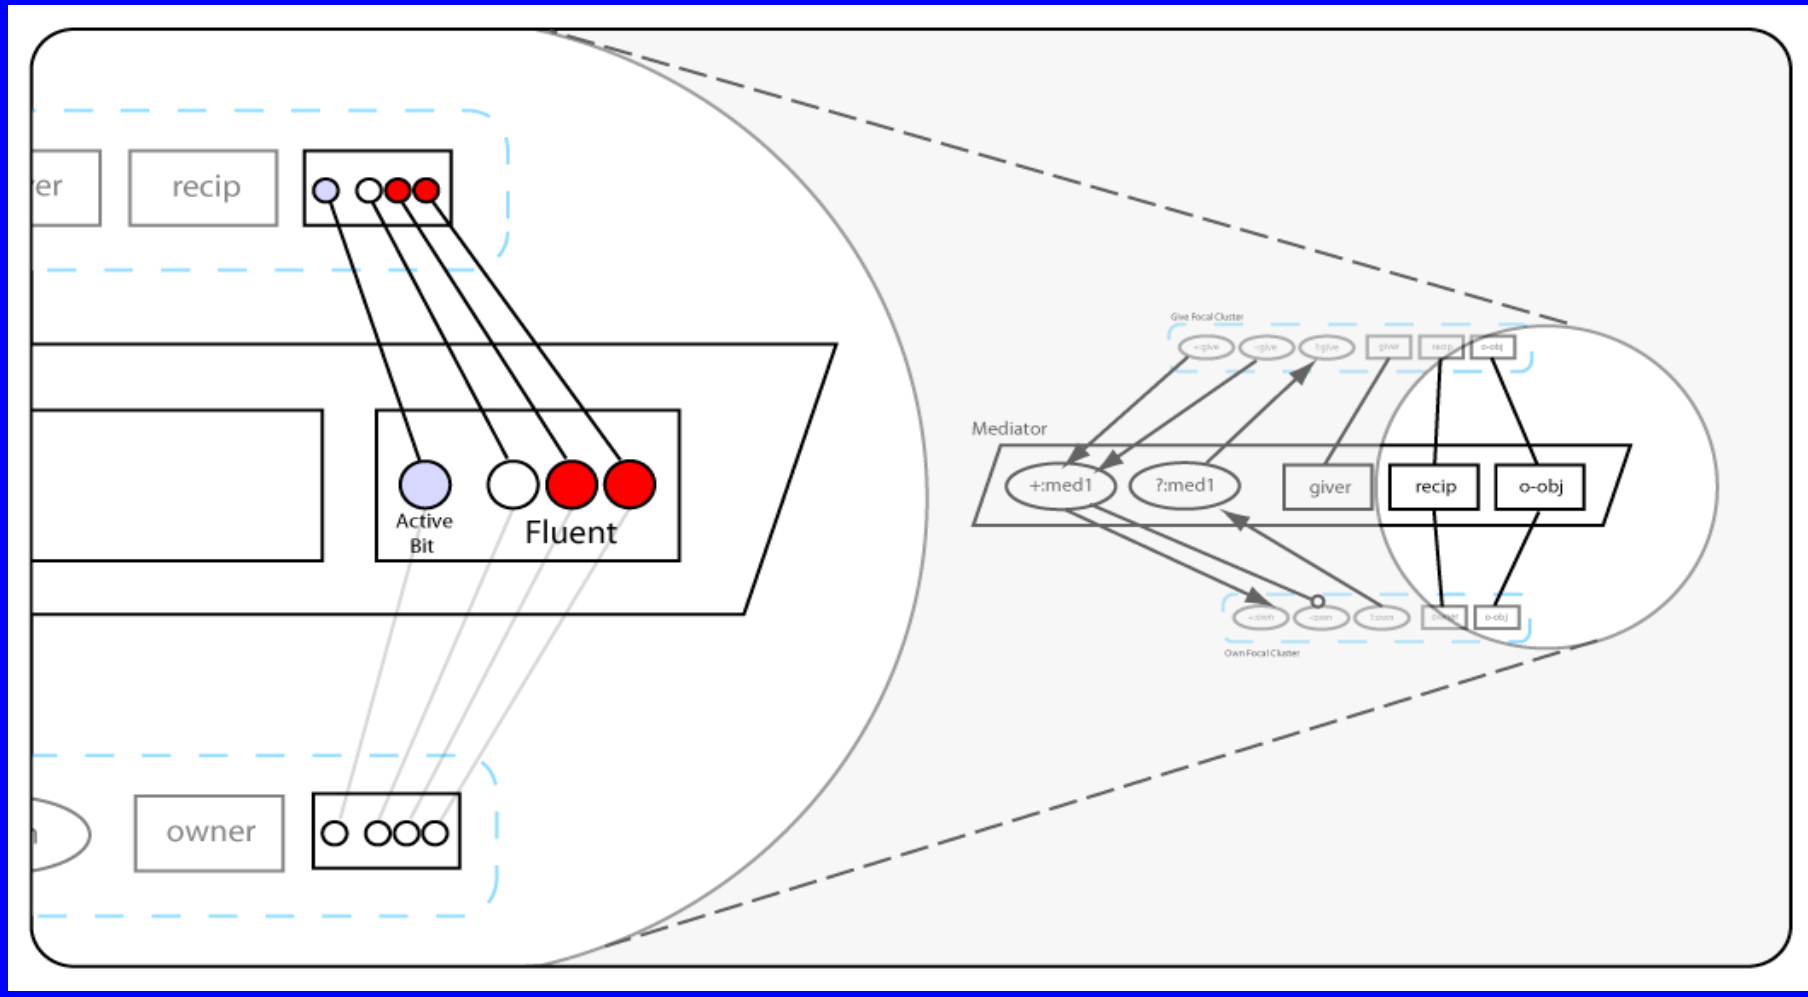

Binder:

- What does the pattern mean?
  - The binder gives fluent patterns meaning

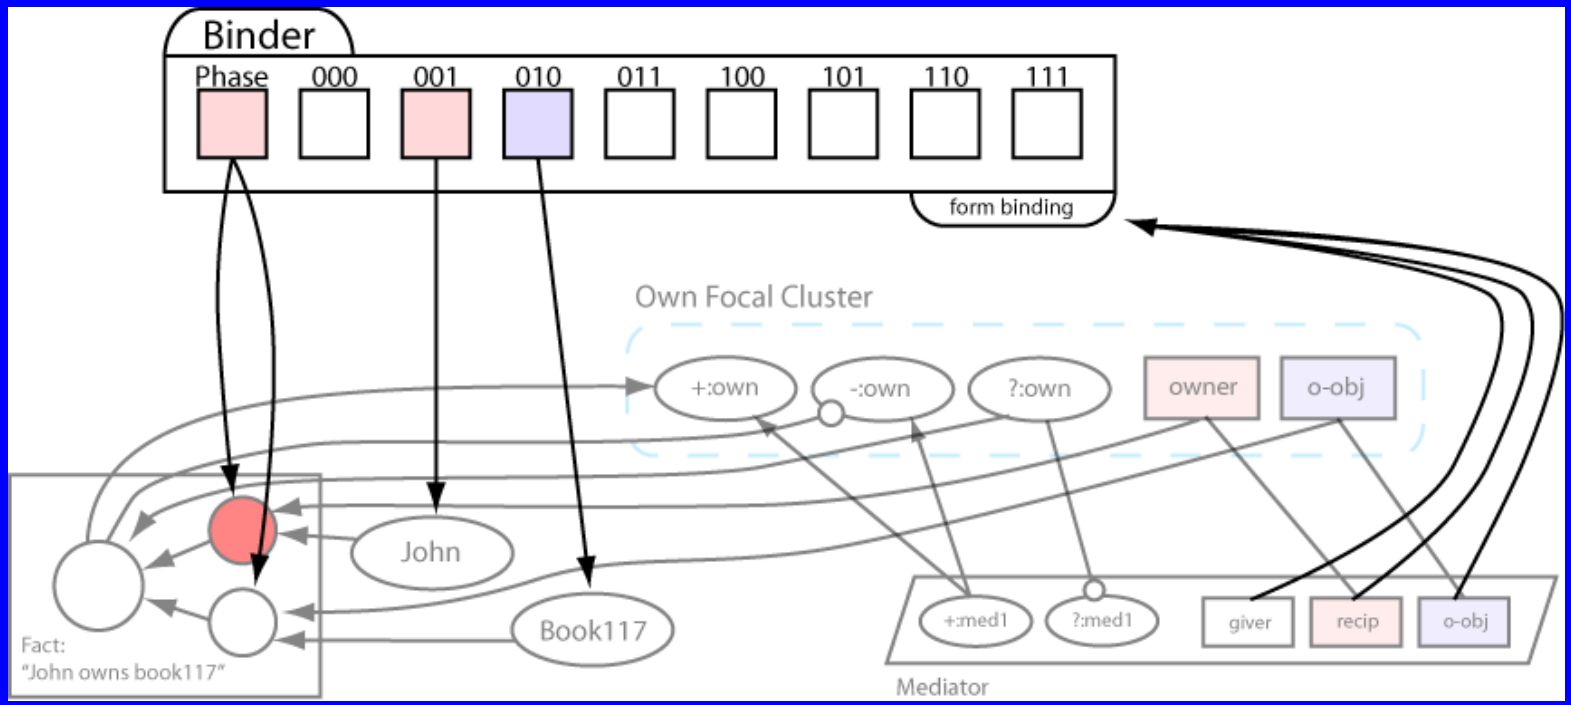

Predicates:

- **Represent short term beliefs about the world**

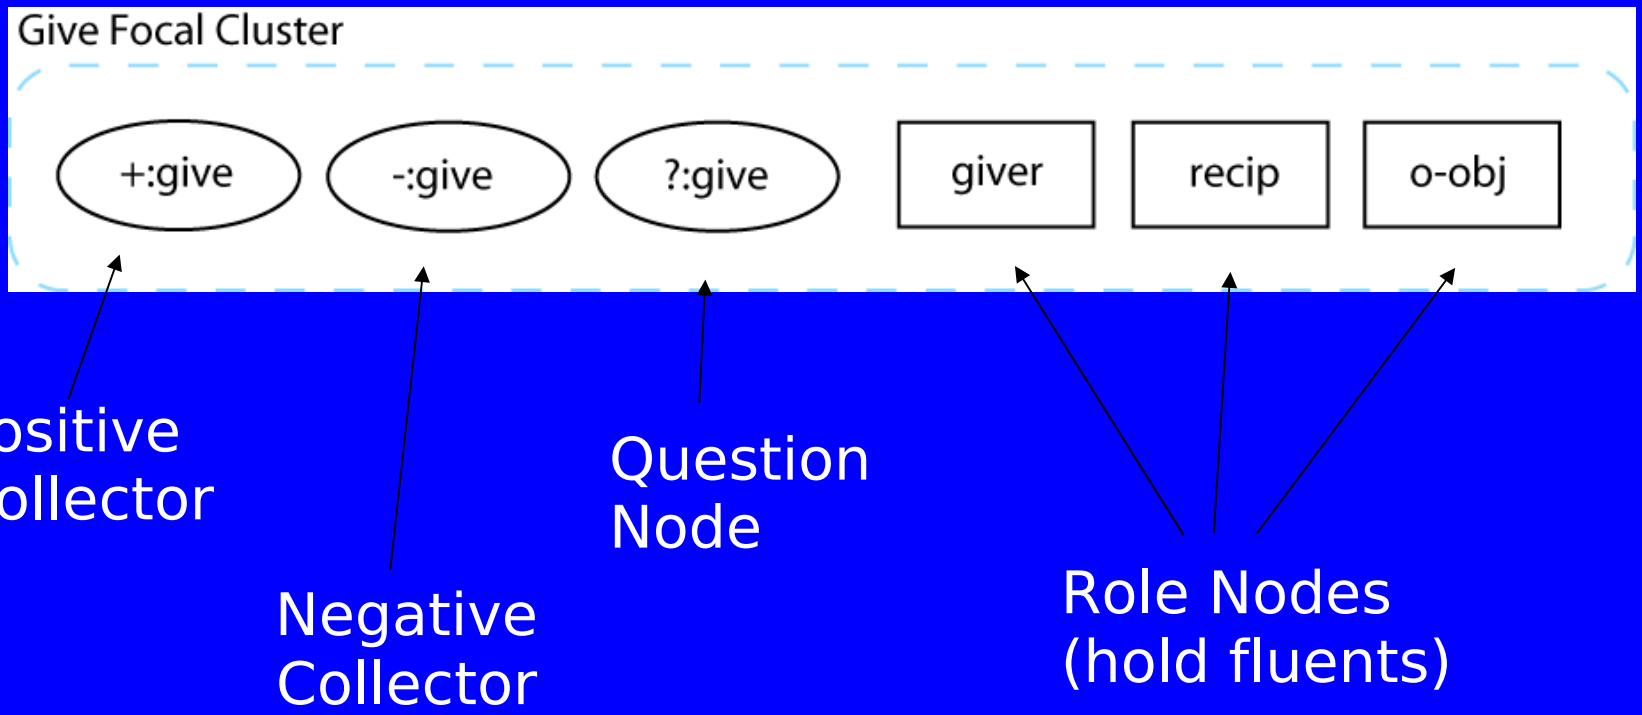

Facts:

- **Support or refute belief in a specific set of bindings of a given predicate**

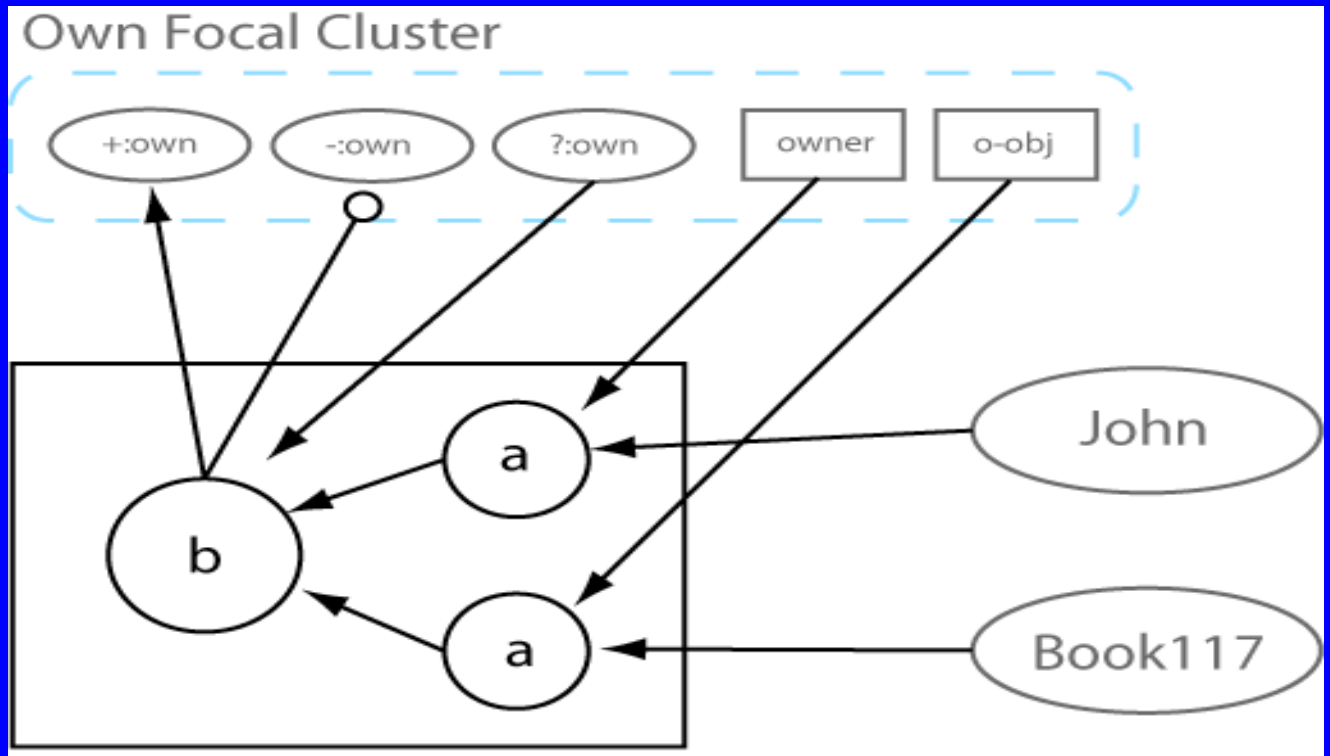

Inference:

- Connections between predicates form evidential links
  - Big(x) & CanBite(x) => Scary(x)
  - Poisonous(x) & CanBite(x) => Scary(x)
  - Strength of connections and shape of neuron response curve determines exactly what "evidence" means
- Direct connections won't work
  - Consider Big(f1) & Poisonous(f1)
  - We want to "Or" over a number of "And"s

Solution: Mediators

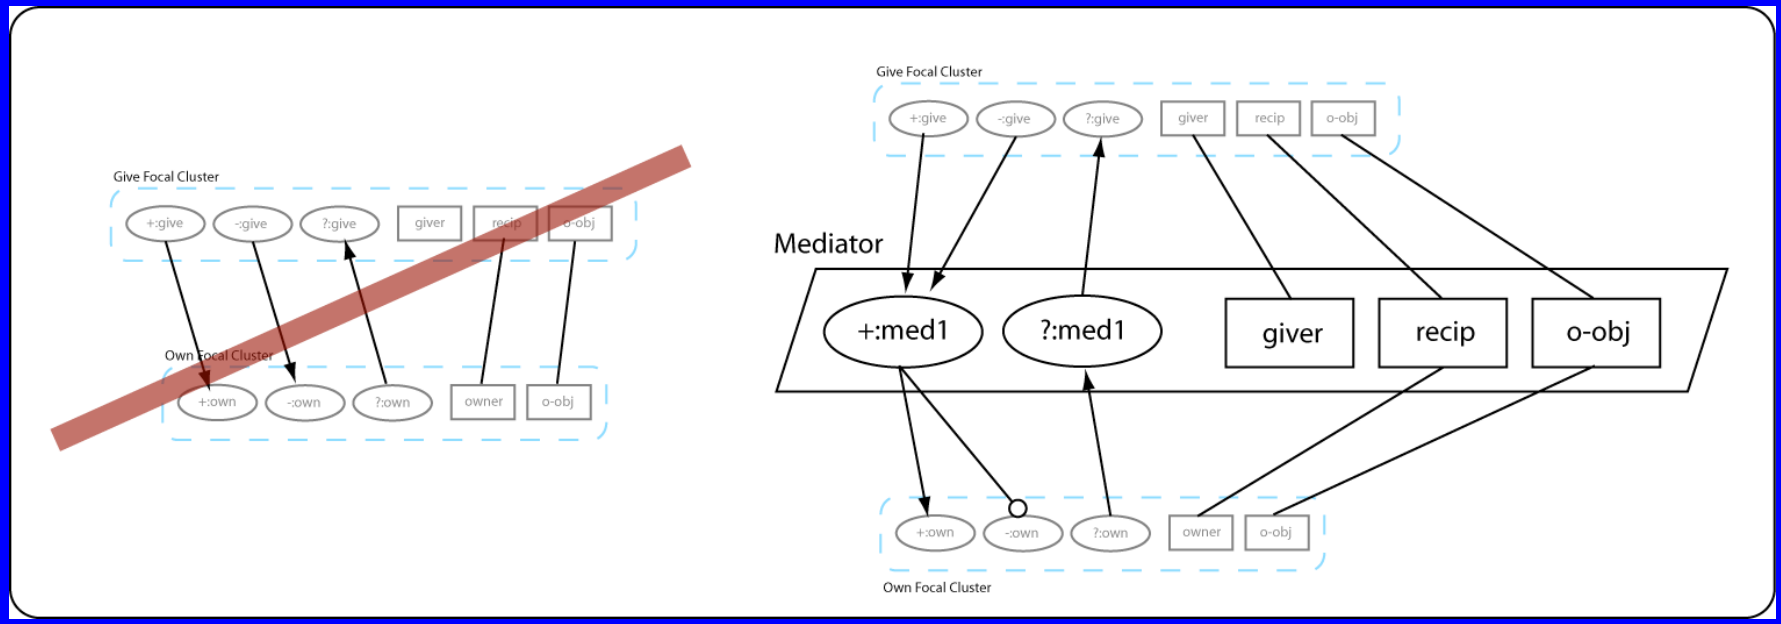

- **Multiple antecedents**
- **Role consistency**

Mediators (continued)

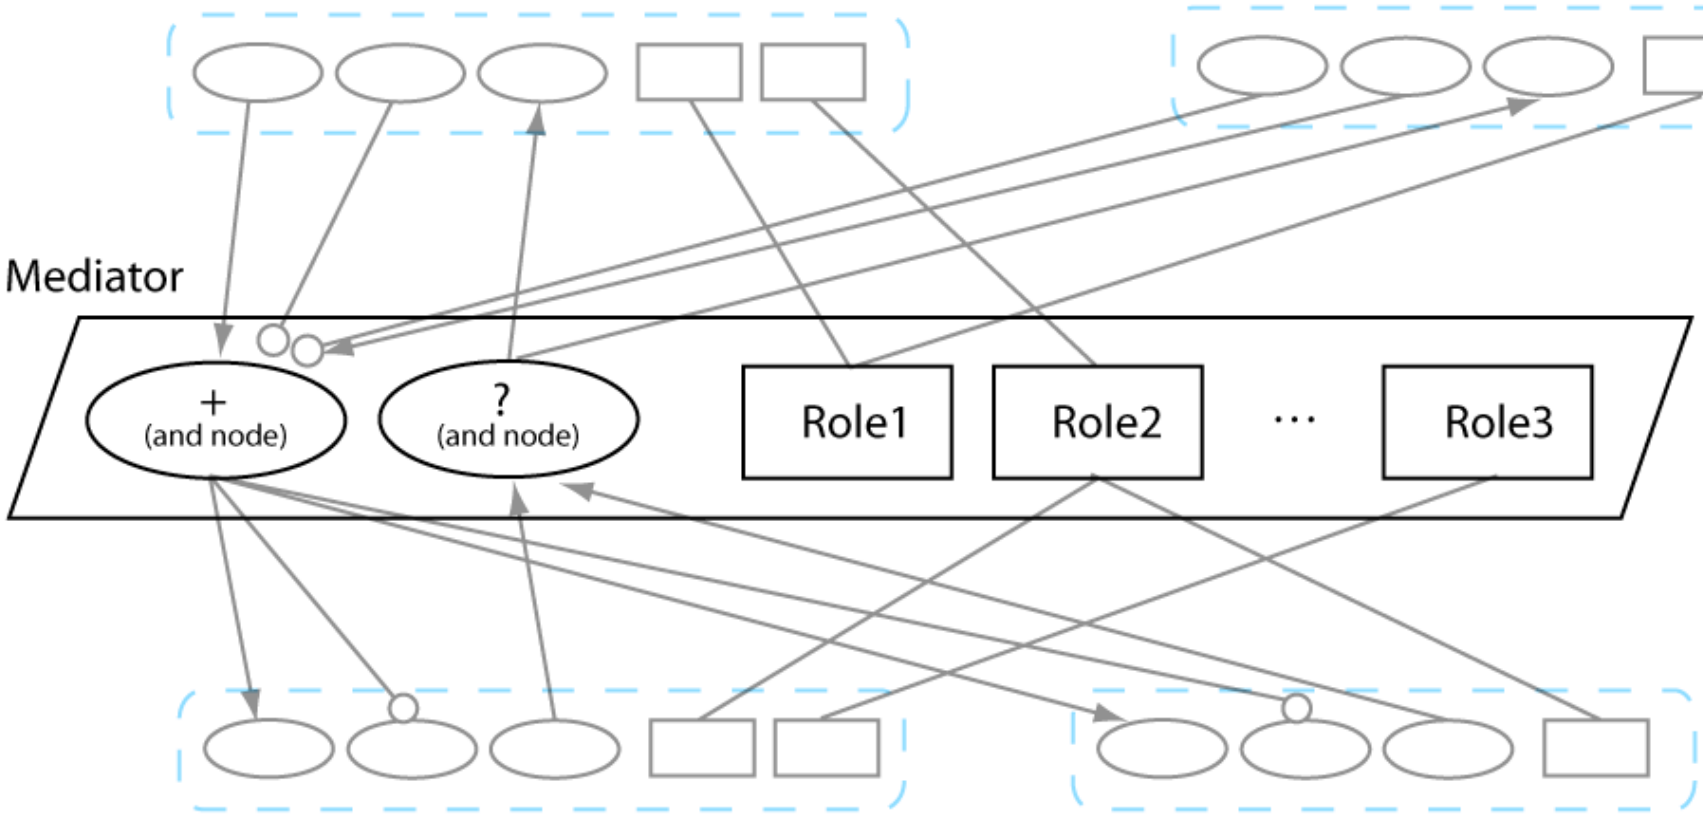

Fluents:

- Roles are just patterns of activation 3-4 bits

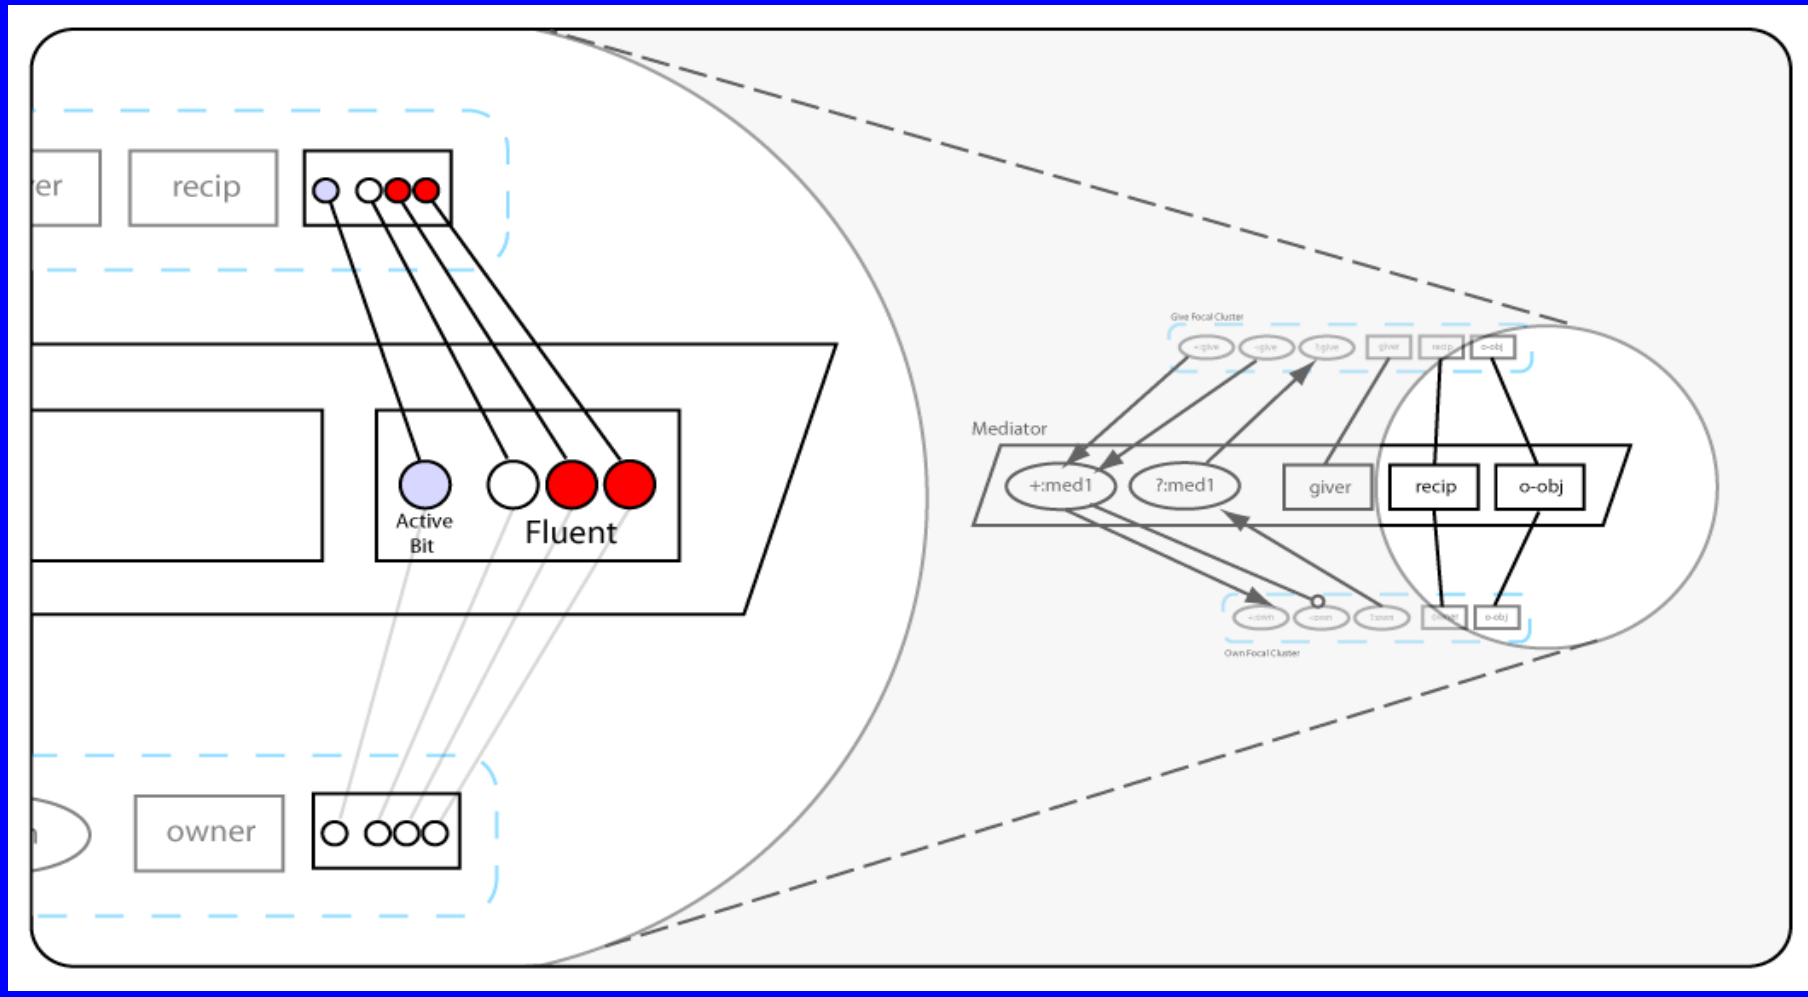

Binder:

- What does the pattern mean?
  - The binder gives fluent patterns meaning

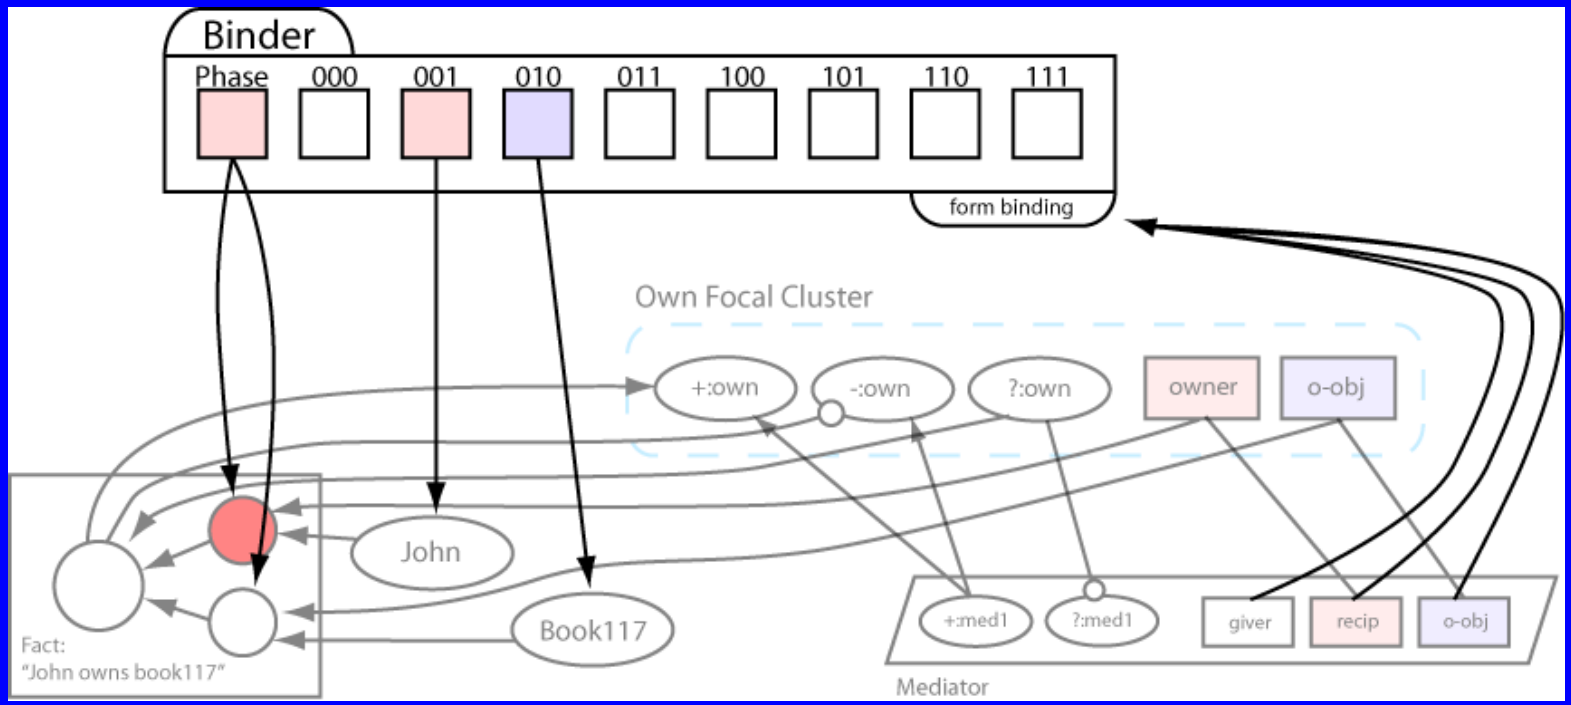

Multiple Assertions

- As described so far, the system cannot simultaneously represent Big(f1) and Big(f2)
- Solution
  - Multiple instances of predicates
  - Requires more complex connections
    - Signals must pass only between clusters with matching fluents
    - Questions must requisition an appropriate number of clusters

Multiple Assertions (detail)

- **Connections between Predicates and their evidence Mediators are easy 1-1**

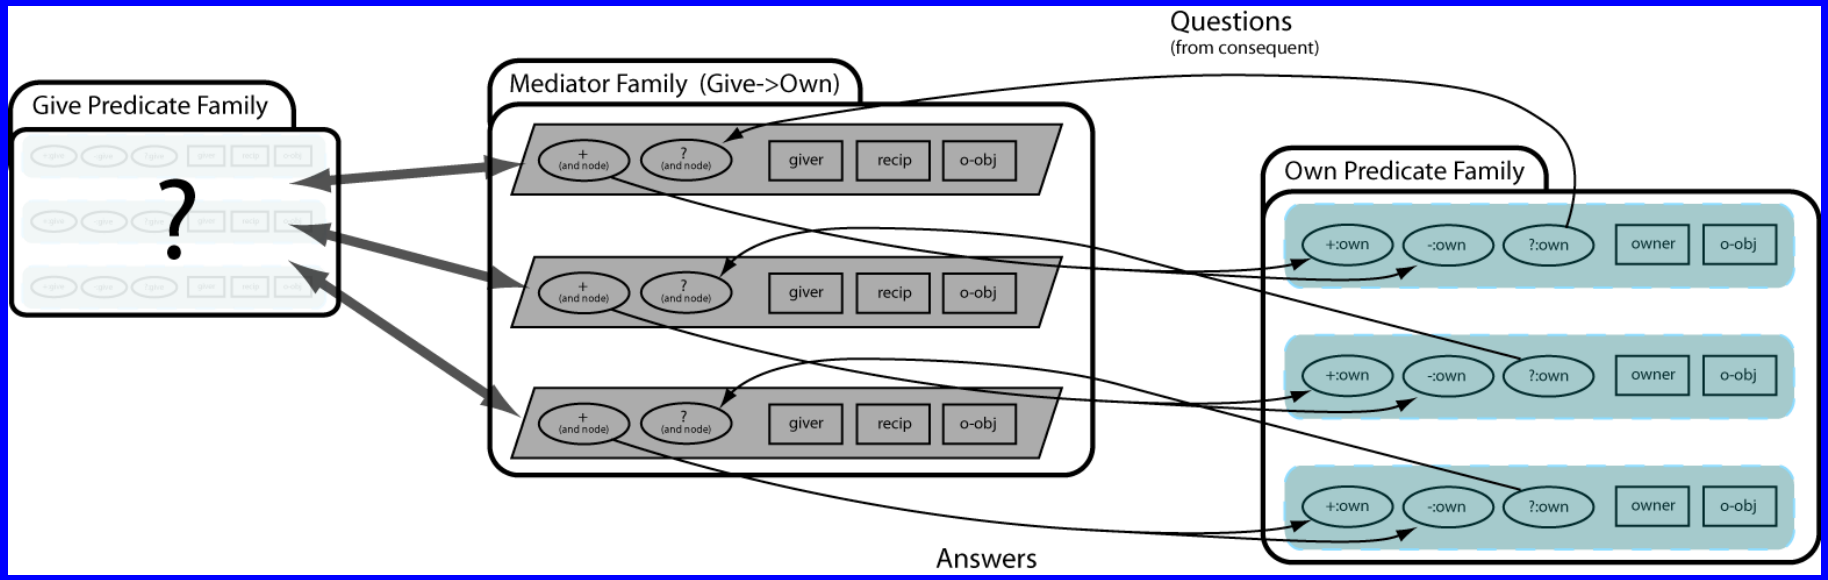

Multiple Assertions (detail)

- Connections between Predicates and their evidence Mediators are easy 1-1
- Evidential connections of Mediators and their evidence Predicates are easy
  - Just connect + and - nodes dependent on matching fluents
- Questions going between Mediators and evidence Predicates are hard
  - Add a selection network to deal with one question at a time

Components of the System

- Object references
  - Fluents
  - Binder
- Short term storage
  - Predicate state
- Long term storage
  - Facts, mediators, what predicates exist
- Inference
  - Mediators
- Types
  - Ontology

Limitations

- Size of network is linear with knowledge base
- Short-term knowledge limited by number of fluents
- Depth of inference limited in time
- Number of same assertions limited
- Inference only goes entirely correctly with ground instances (e.g. “Fido” and not “dog”)

Questions

Representing belief and utility in Shruti

- *associate utilities with states of affairs (relational instances)*
- *encode utility facts:*
  - *context sensitive memories of utilities associated with certain events or event-types*
- *propagate utility along causal structures*
- *encode actions and their consequences*

Encoding "Fall => Hurt"

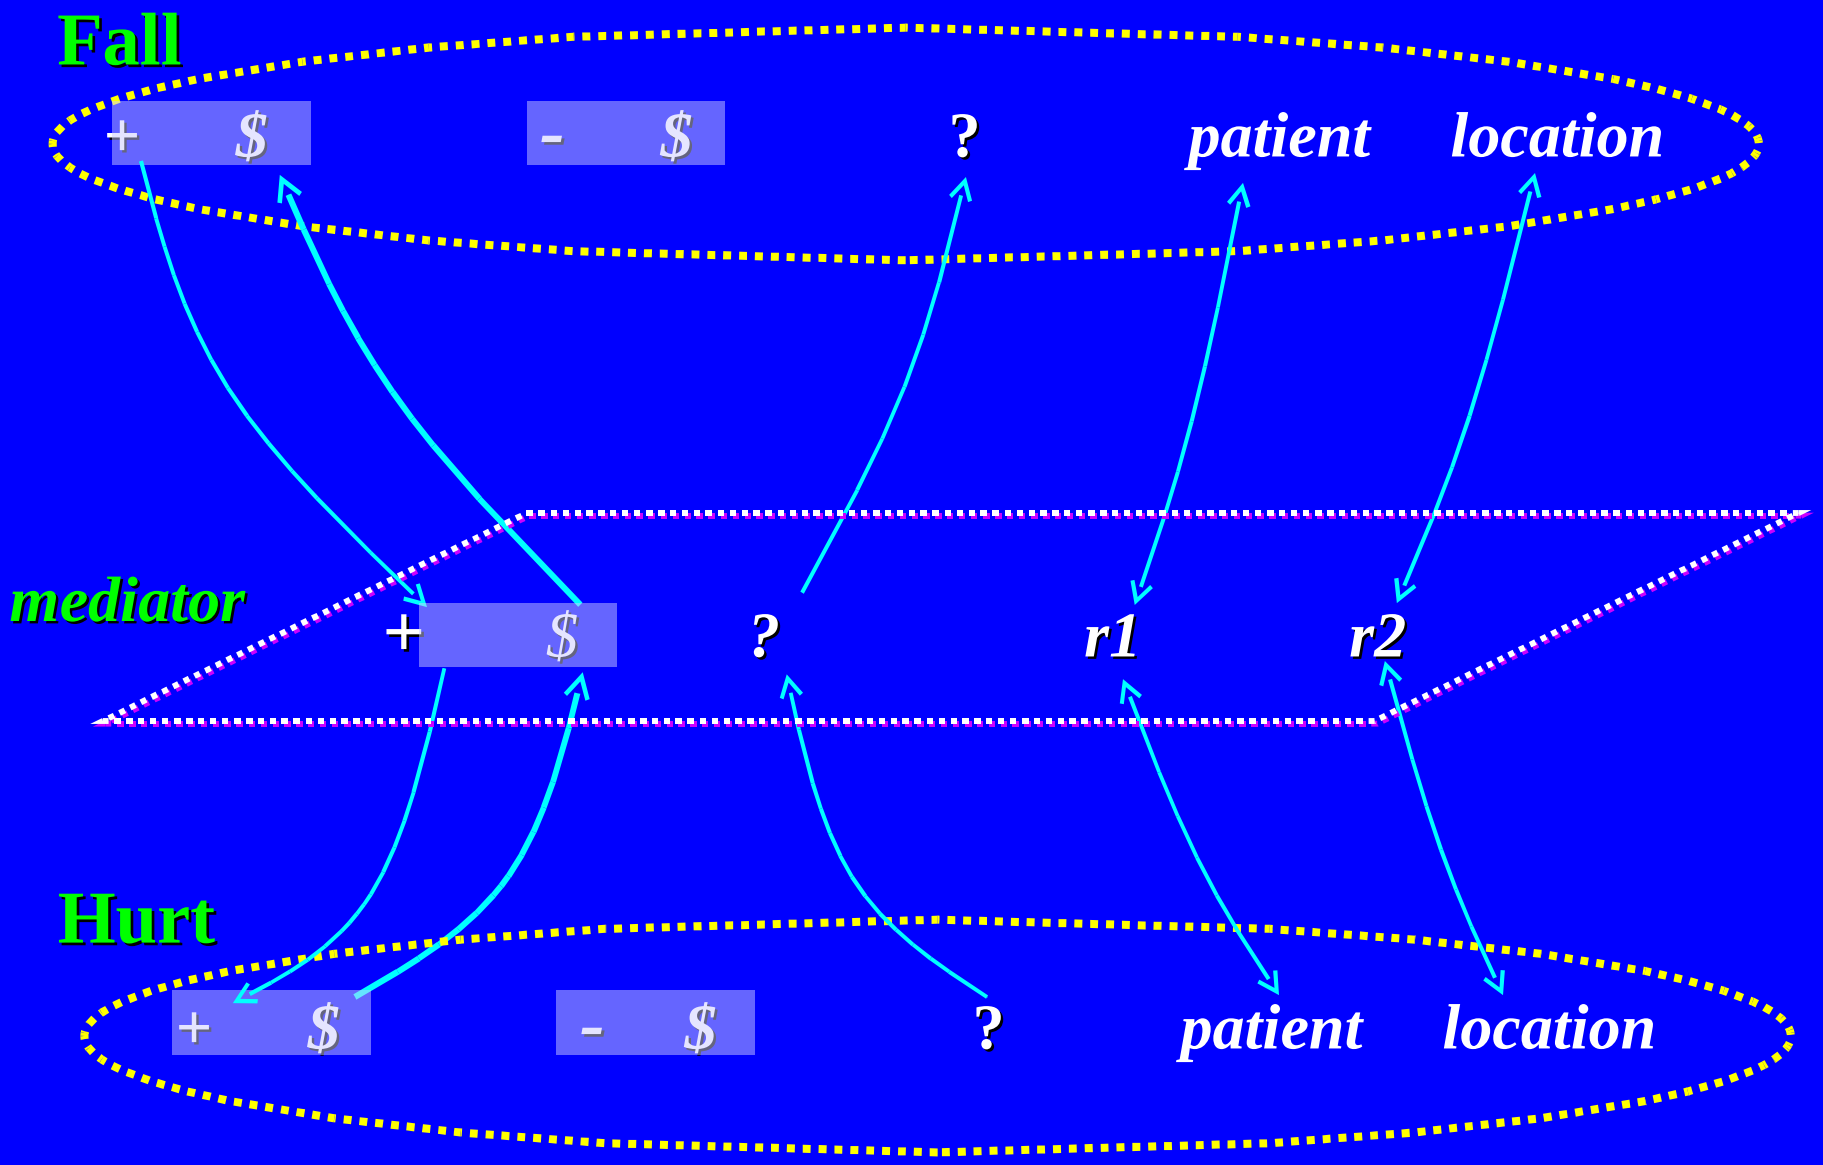

Focal-clusters augmented to encode belief and utility

Attack

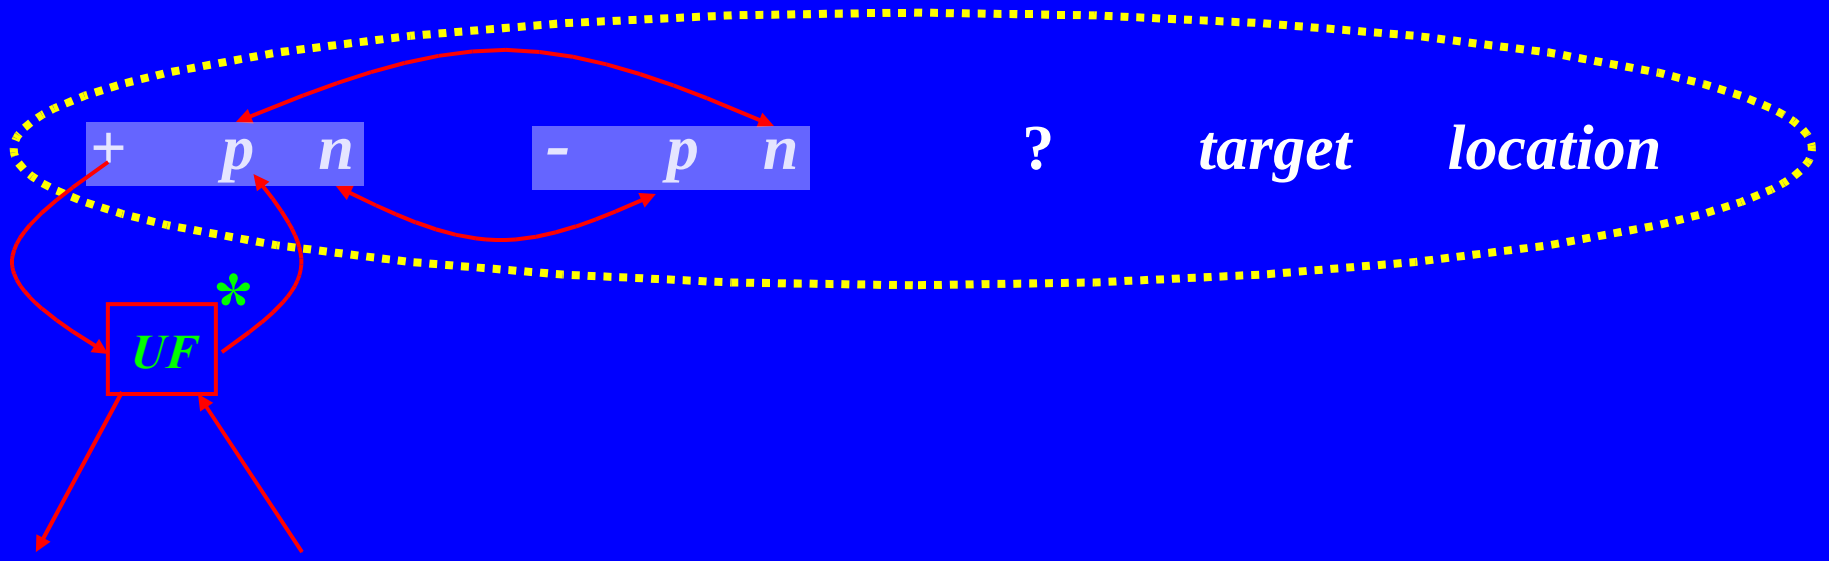

**UF: utility fact; either a specific reward fact (R-fact) or a generic value fact (V-fact)*

Behavior of augmented Shruti

Shruti reflexively

- *Makes observations*
- *Seeks explanations*
- *Makes predictions*
- *Instantiates goals*
- *Seeks plans that enhance expected future utility*
  - *identify actions that are likely to lead to desirable situations and prevent undesirable ones*

Shruti suggests how different sorts of knowledge may be encoded within neurally plausible networks

- ***Entities, types and their relationships*** ***(John is a Man)***
- ***Relational schemas/frames corresponding to action and event types*** ***(Falling, giving, …)***
- ***Causal relations between relational schemas*** ***(If you fall you can get hurt)***
- ***Taxon/Semantic facts*** ***(Children often fall)***
- ***Episodic facts*** ***(John fell in the hallway on Monday)***
- ***Utility facts*** ***(It is bad to be hurt)***

Current status of learning in Shruti

- ✓ *Episodic facts: A biologically grounded model of "one-shot" episodic memory formation*
  - *Shastri, 1997; Proceedings of CogSci 1997*
  - *_2001; Neurocomputing*
  - *_2002; Trends in Cognitive Science*
  - *_In Revision; Behavioral and Brain Science*

    (available as a Technical Report)

...current status of learning in Shruti

Work in Progress

- *Causal rules*
- *Categories*
- *Relational schemas*

Shastri and Wendelken 2003; Neurocomputing

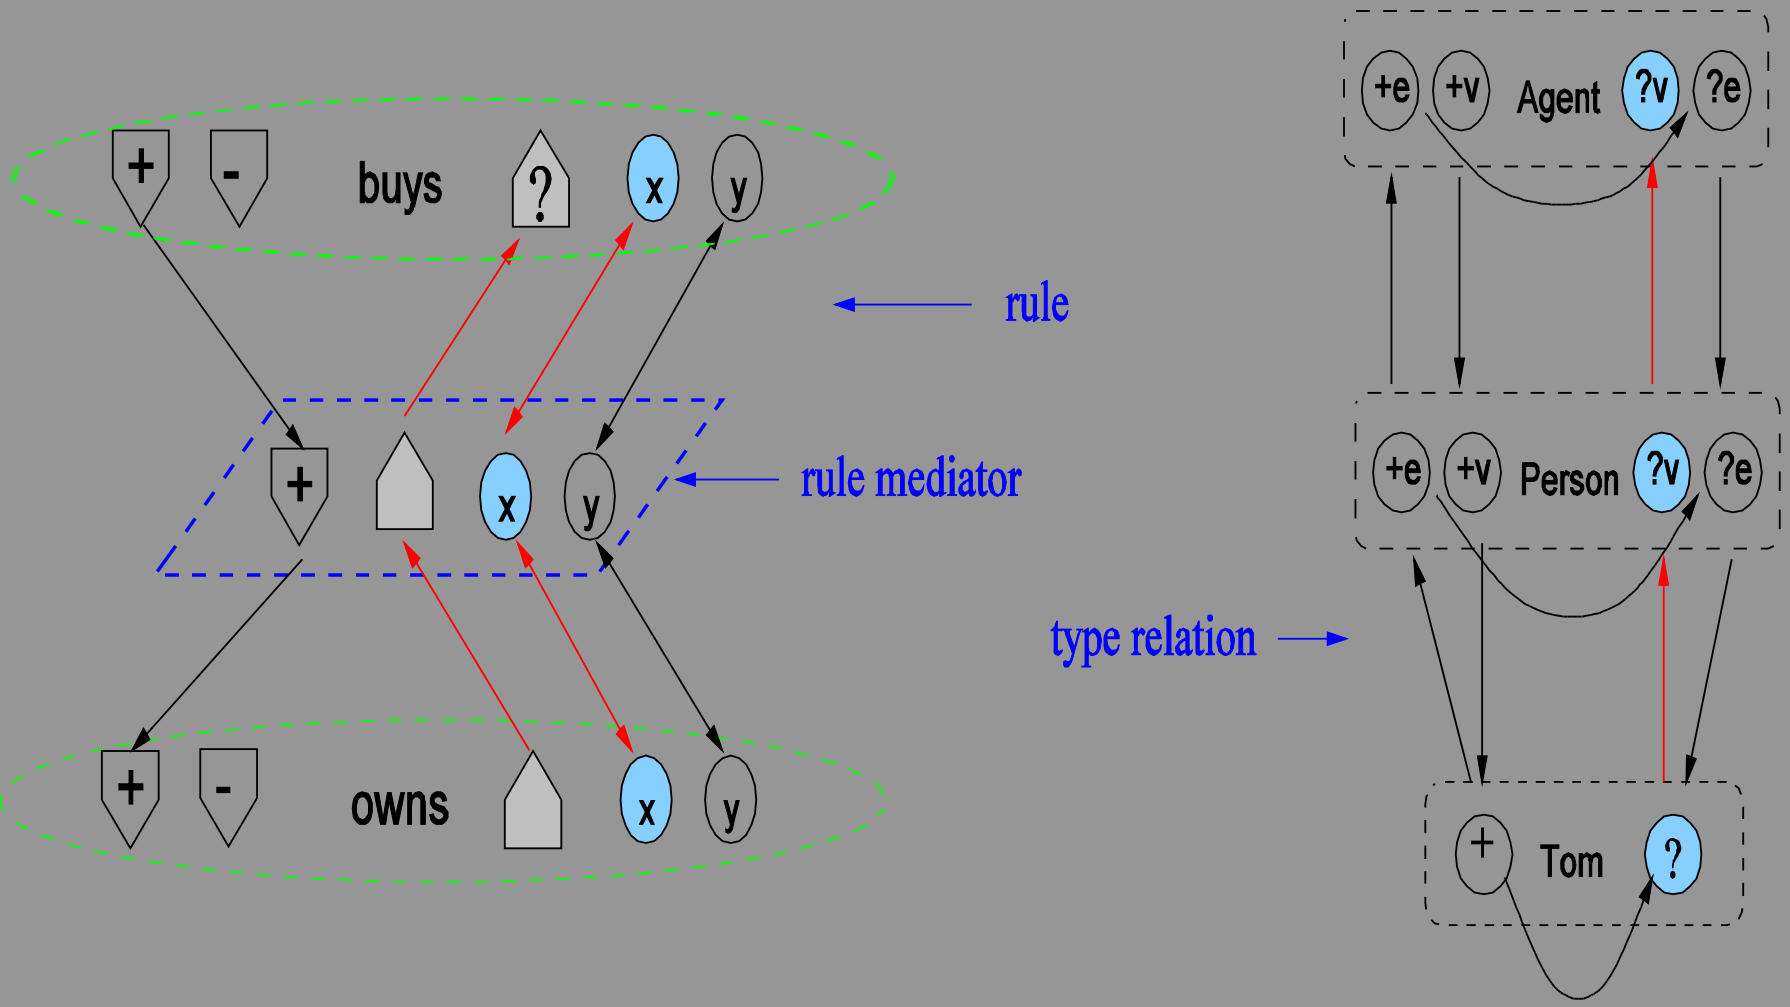

Questions