Norman and Shallice (1986)

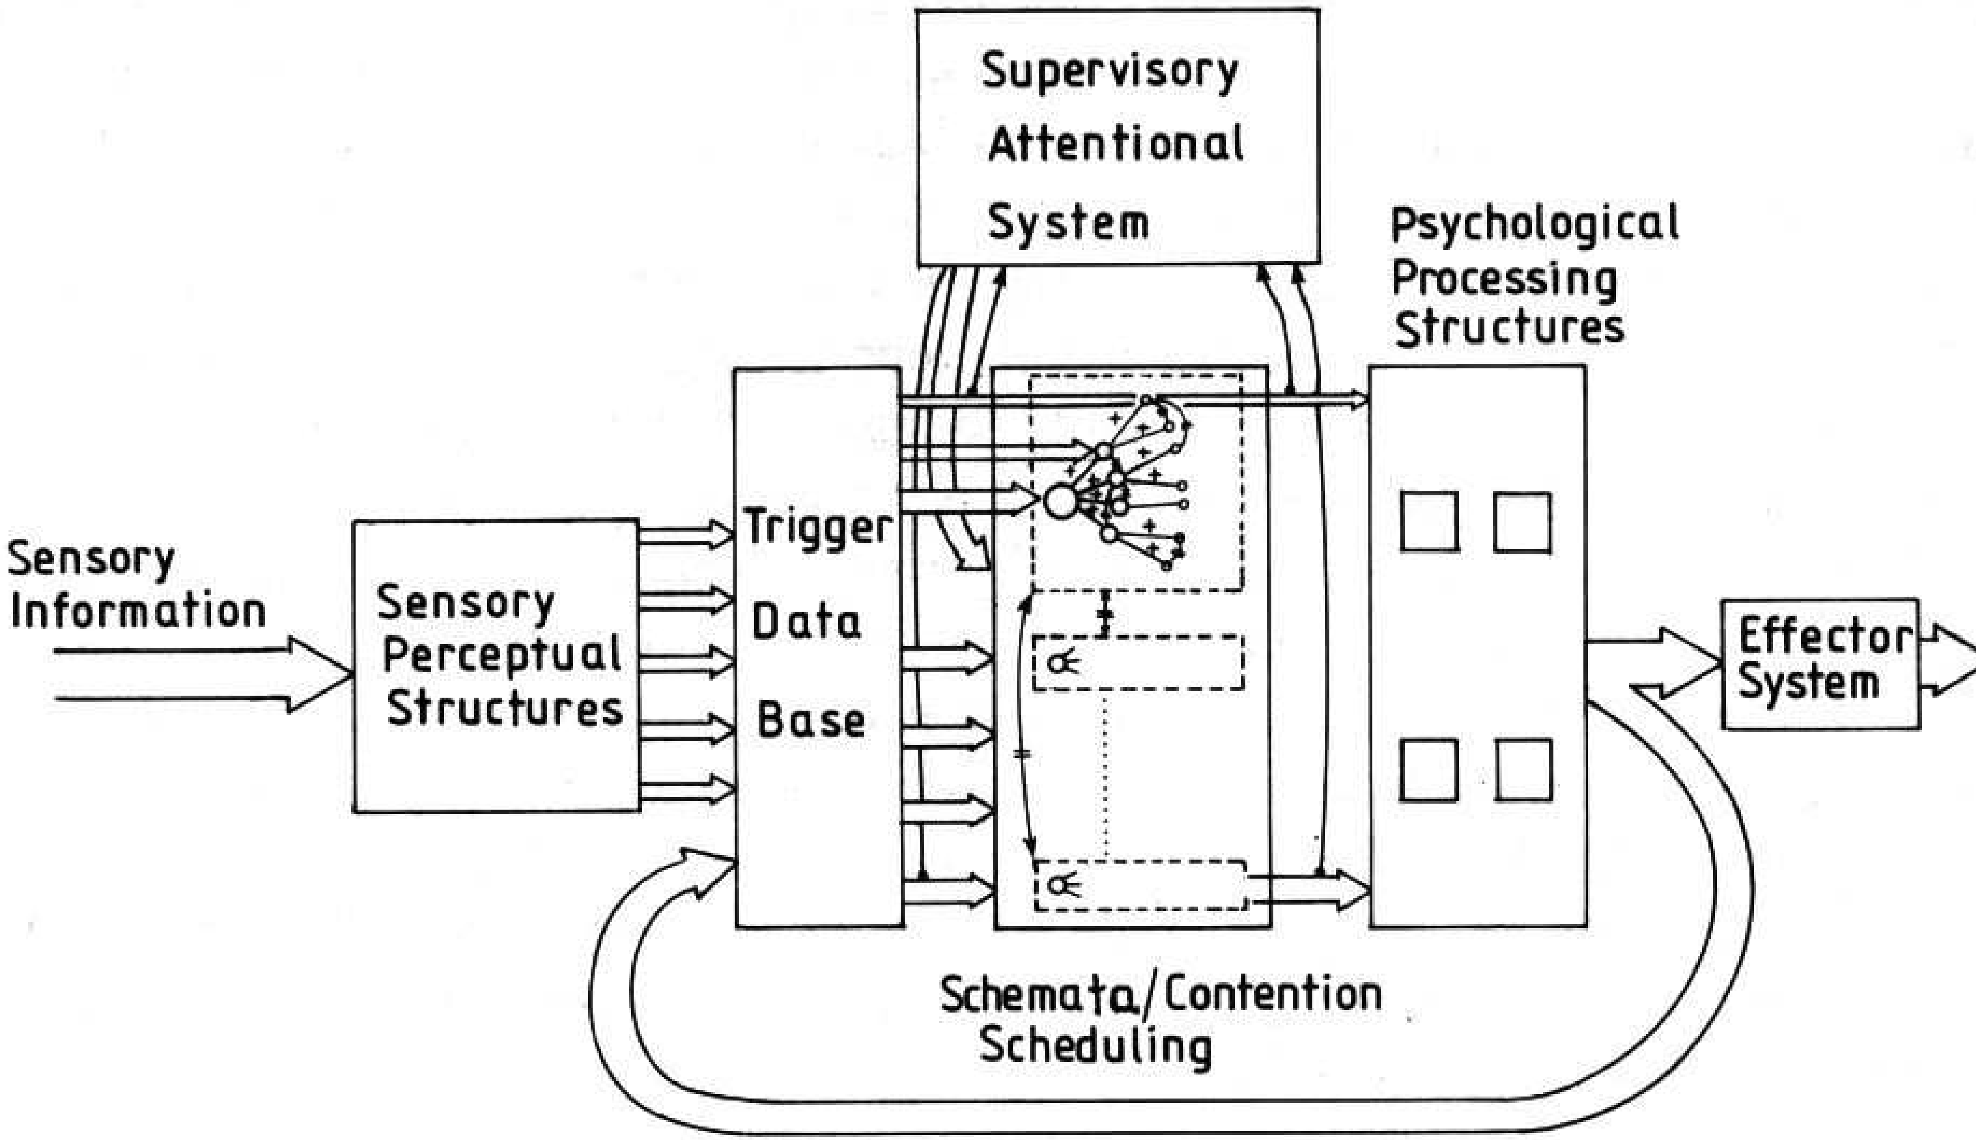

Hierarchical structure in routine sequential action

Miller, Gelanter and Pribram (1960), Estes (1972), Rumelhart and Norman (1982)
Shank and Abelson (1977), MacKay (1985, 1987), Fuster (1989), Grafman (1995)
Forde and Humphreys (1998), **Cooper and Shallice (2000)**

Knowledge

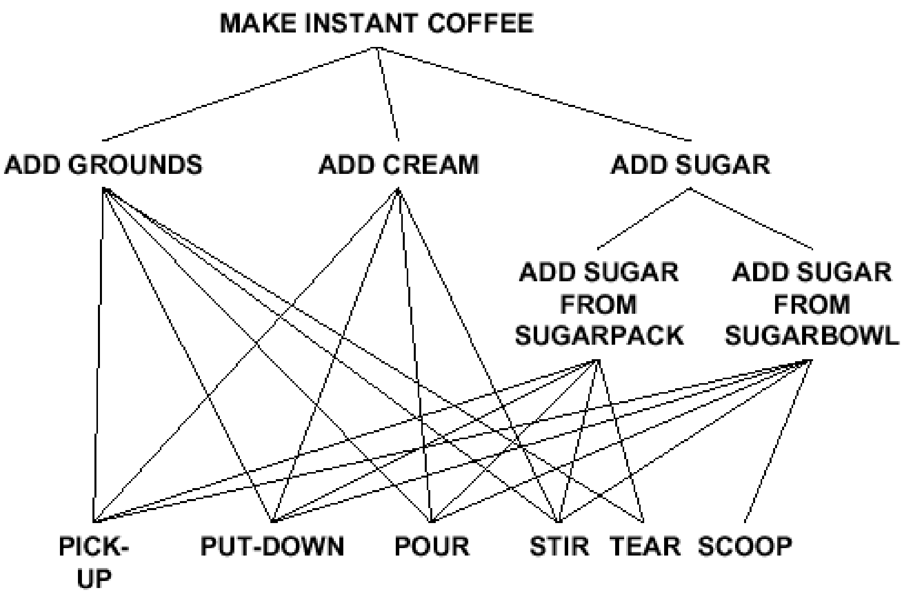

Representations/Processes?

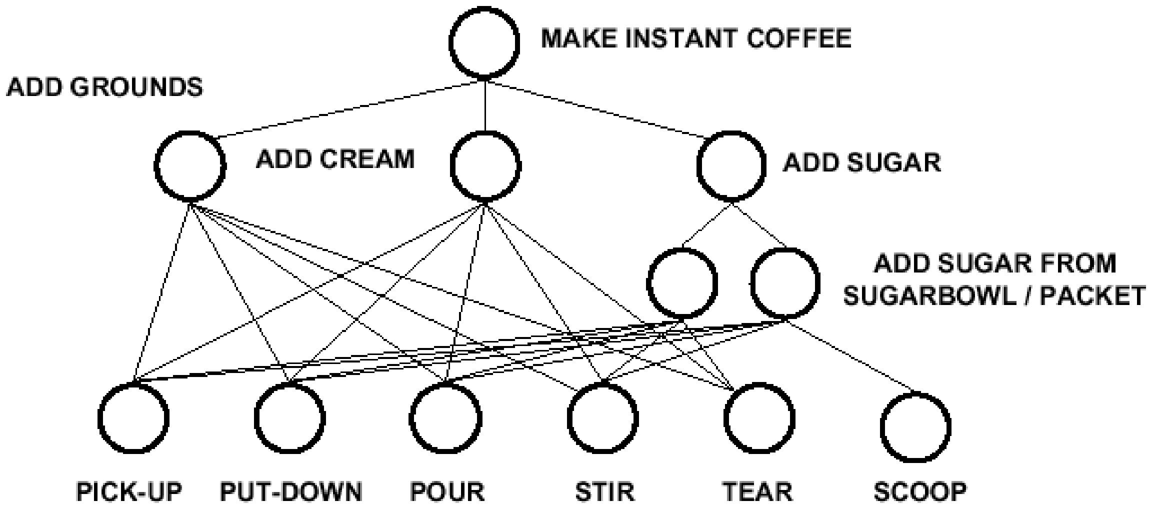

Problems

- Weak learning theory for when (and how) to elaborate vs. add schemas
- No intrinsic sequencing mechanism

Hierarchical structure in conceptual knowledge

Quillian (1967)
Collins and Quillian (1969)
Anderson (1983, 1990);
Anderson and Lebiere (1998)

Hinton (1981)
Rumelhart and Todd (1993)
McClelland, McNaughton and O'Reilly (1995)
Rogers and McClelland (2004)

Knowledge

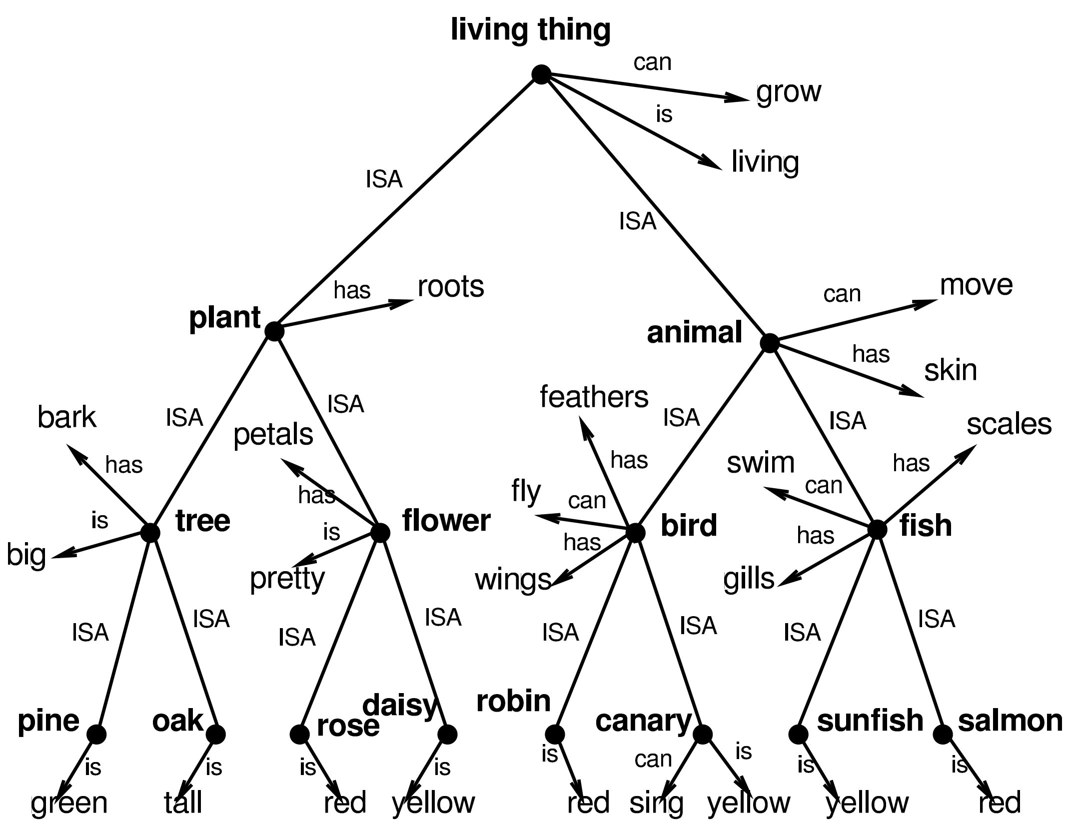

Representations

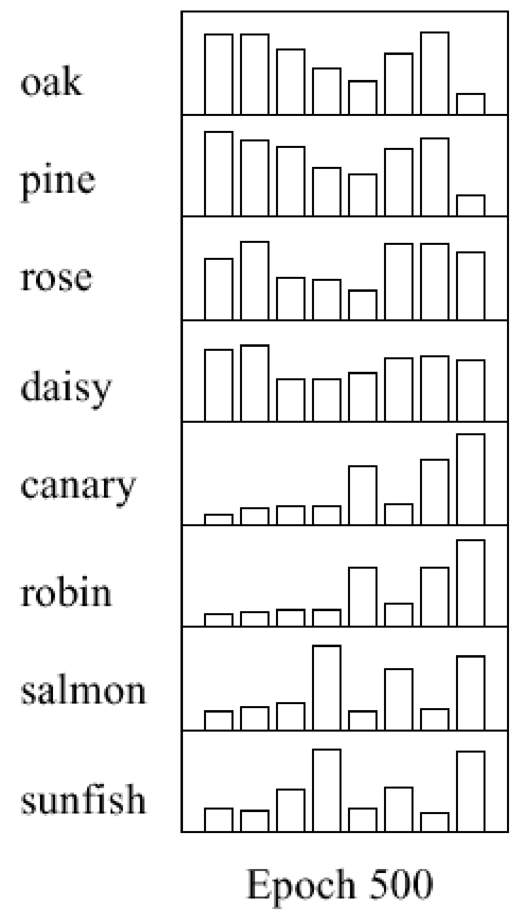

Similarities

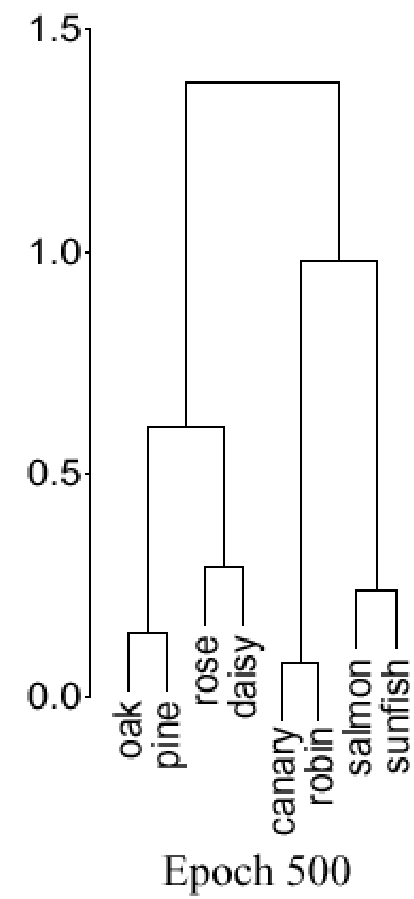

Hierarchical structure in sentence processing

"Beth told Bill that the cow left the pasture."

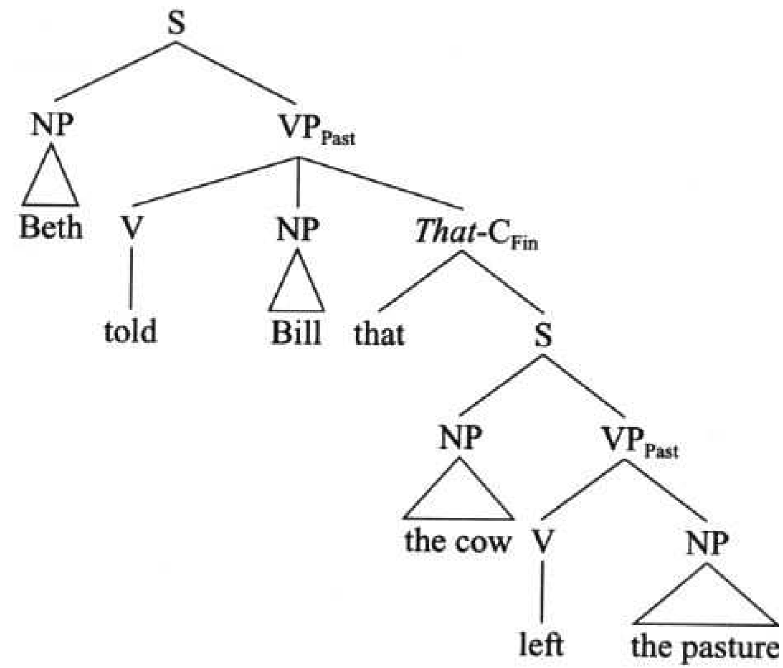

Elman (1991, 1993)

- Simple recurrent network trained to predict words in pseudo-English sentences

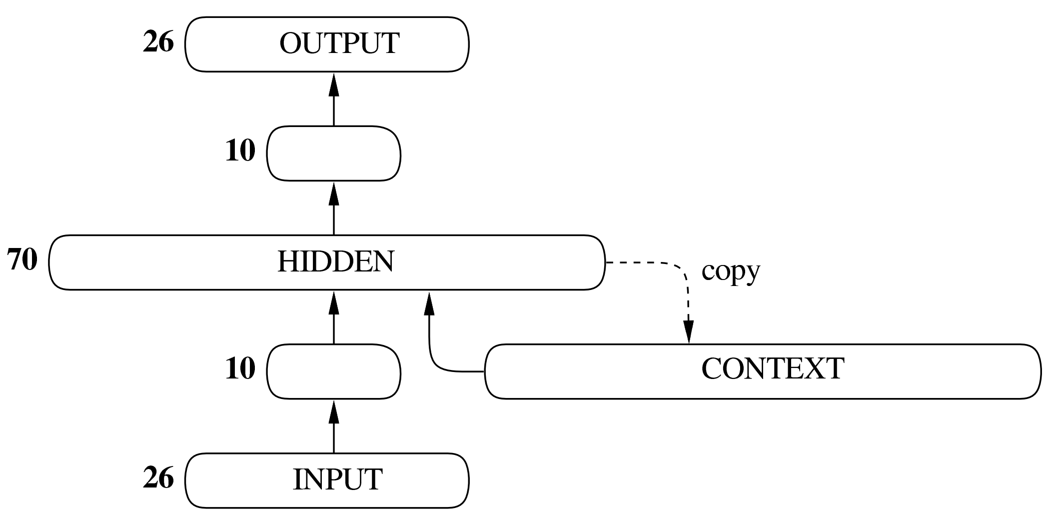

S → NP VI . | NP VT NP .
NP → N | N RC
RC → who VI | who VT NP | who NP VT
N → boy | girl | cat | dog | Mary | John | boys | girls | cats | dogs
VI → barks | sings | walks | bites | eats | bark | sing | walk | bite | eat
VT → chases | feeds | walks | bites | eats | chase | feed | walk | bite | eat

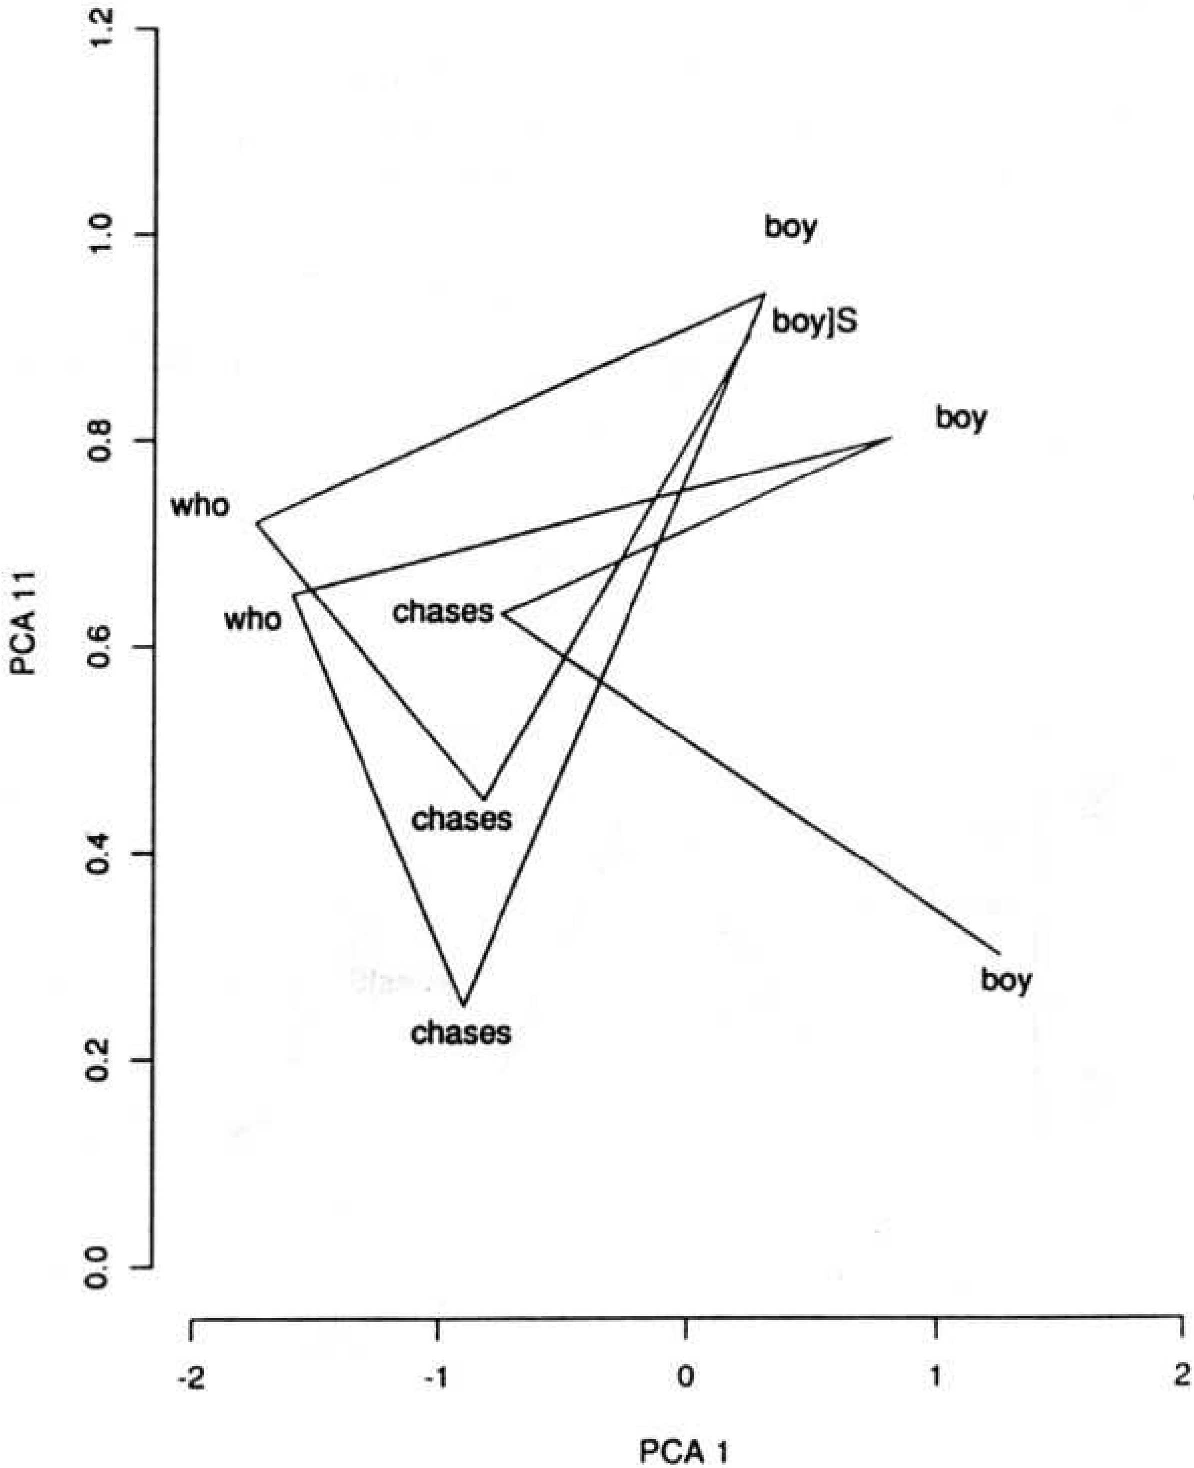

Boy chases boy who chases boy who chases boy .

Principal Components Analysis (PCA) of network's internal representations

- Largest amount of variance (PC-1) reflects **word class** (noun, verb, function word)
- Separate dimension of variation (PC-11) encodes **syntactic role** (agent/patient) for nouns and **level of embedding** for verbs

A distributed connectionist approach to sequential action

Botvinick and Plaut (2004, *Psych. Rev.*)

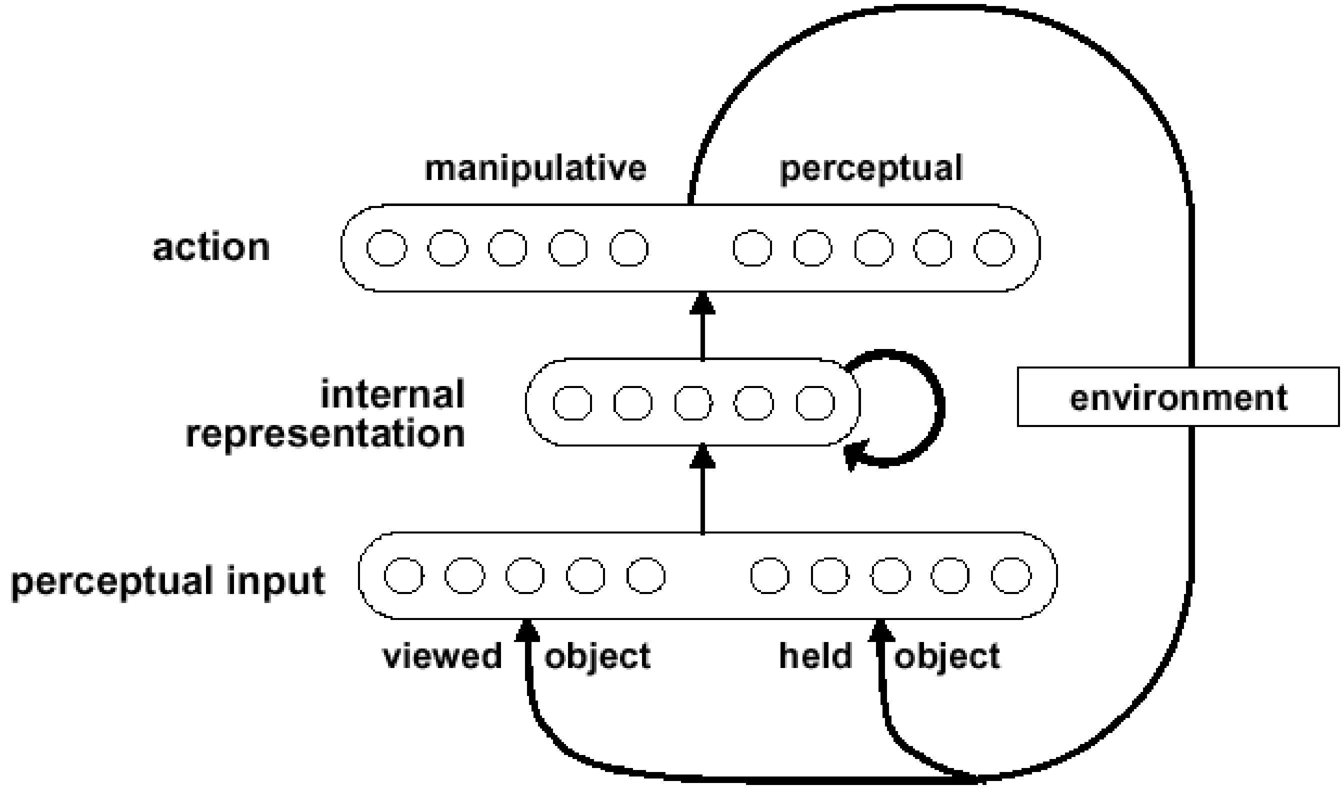

- Simple recurrent network that maps perceptual inputs *and internal representations of task context* onto actions
  - Input codes currently viewed/held objects; output is manipulative or perceptual action
  - Trained by observing skilled coffee- and tea-making, and on general *affordances*
  - Tested by applying most strongly activated action to environment and repeating

Task structure: Making instant coffee (and tea)

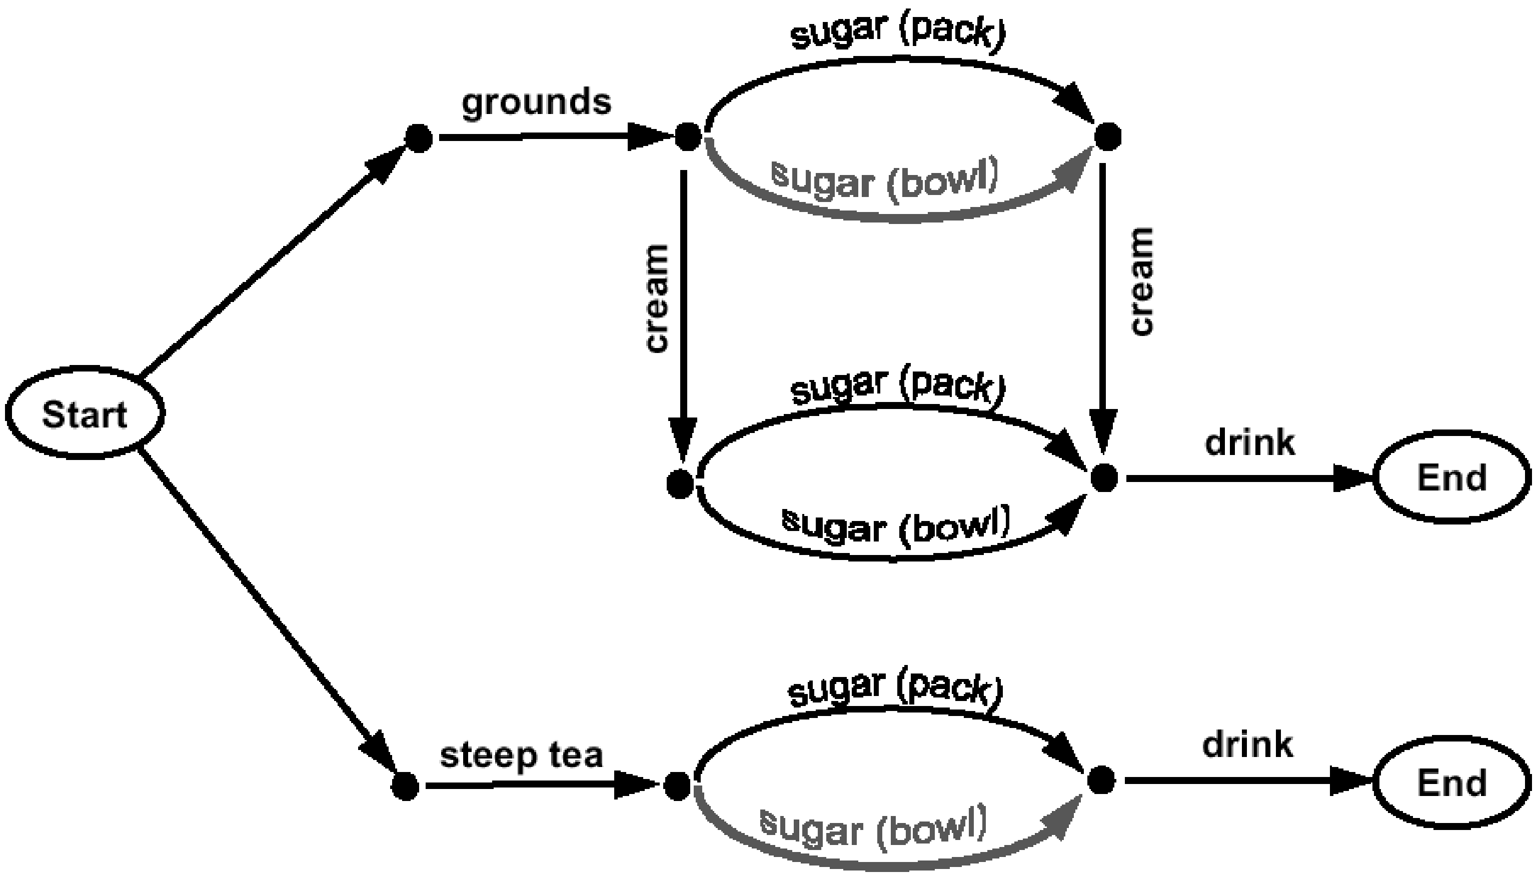

- Hierarchically structured
- Actions/subtasks may appear in multiple contexts
- Environmental cues alone sometimes insufficient to guide action selection
- Subtasks may be disjoint or executed in variable order

	Step	Fixated object	Held object	Action
	1	cup, 1-handle, clear-liquid	nothing	fixate-coffee-pack
	2	packet, foil, untorn	nothing	pick-up
	3	packet, foil, untorn	packet, foil, untorn	pull-open
	4	packet, foil, torn	packet, foil, torn	fixate-cup
Grounds	5	cup, 1-handle, clear-liquid	packet, foil, torn	pour
	6	cup, 1-handle, brown-liquid	packet-foil-torn	fixate-spoon
	7	spoon	packet, foil, torn	put-down
	8	spoon	nothing	pick-up
	9	spoon	spoon	fixate-cup
	10	cup, 1-handle, brown-liquid	spoon	stir
	11	cup, 1-handle, brown-liquid	spoon	fixate-sugar
	12	cup, 2-handles, lid	spoon	put-down
	13	cup, 2-handles, lid	nothing	pull-off
Sugar	14	cup, 2-handles, sugar	lid	fixate-spoon
	15	spoon	lid	put-down
(bowl)	16	spoon	nothing	pick-up
	17	spoon	spoon	fixate-sugarbowl
	18	cup, 2-handles, sugar	spoon	scoop
	19	cup, 2-handles, sugar	spoon-sugar	fixate-cup
	20	cup, 1-handle, brown-liquid	spoon-sugar	pour
	21	cup, 1-handle, brown-liquid	spoon	stir
	22	cup, 1-handle, brown-liquid	spoon	fixate-carton
	23	carton, closed	spoon	put-down
	24	carton, closed	nothing	pick-up
	25	carton, closed	carton, closed	peel-open
Cream	26	carton, open	carton, open	fixate-cup
	27	cup, 1-handle, brown-liquid	carton-open	pour
	28	cup, 1-handle, light-, brown-liquid	carton-open	fixate-spoon
	29	spoon	carton-open	put-down
	30	spoon	nothing	pick-up
	31	spoon	spoon	fixate-cup
	32	cup, 1-handle, light-, brown-liquid	spoon	stir
	33	cup, 1-handle, light-, brown-liquid	spoon	put-down
	34	cup, 1-handle, light-, brown-liquid	nothing	pick-up
Drink	35	cup, 1-handle, light-, brown-liquid	cup, 1-handle, light-, brown-liquid	sip
	36	cup, 1-handle, light-, brown-liquid	cup, 1-handle, light-, brown-liquid	sip
	37	cup, 1-handle, empty	cup, 1-handle, empty	say-done

Adding coffee grounds: Detail

Step	Fixated object	Held object	Action
1	cup, 1-handle, clear-liquid	nothing	fixate-coffee-pack
2	packet, foil, untorn	nothing	pick-up
3	packet, foil, untorn	packet, foil, untorn	pull-open
4	packet, foil, torn	packet, foil, torn	fixate-cup
5	cup, 1-handle, clear-liquid	packet, foil, torn	pour
6	cup, 1-handle, brown-liquid	packet-foil-torn	fixate-spoon
7	spoon	packet, foil, torn	put-down
8	spoon	nothing	pick-up
9	spoon	spoon	fixate-cup
10	cup, 1-handle, brown-liquid	spoon	stir

•

•

•

•

Neural network models of serial order. . . [retain] a dynamics dependent on chaining. . . . They also seem to us unlikely to be prone to the kinds of serial order errors discussed below [omissions, transpositions].

—Houghton and Hartley (1995, *Psyche*, p. 5)

Recurrent networks lack "temporal competence. . . the intrinsic dynamics that would enable them to progress autonomously through a sequence."

—Brown, Preece and Hulme (2000, *Psych. Review*, p. 133)

The principal difficulty in obtaining [omission and other sequence errors] within recurrent networks appears to arise from the lack of any separate representation of hierarchical relations (i.e., source/component schema relationships) and order information (i.e., the relative ordering of component schemas within a single source schema). It is thus difficult for order information to be disrupted without disruption to hierarchical relations.

—Cooper and Shallice (2000, *Cog. Neuropsych.*, p. 329)

Acquisition

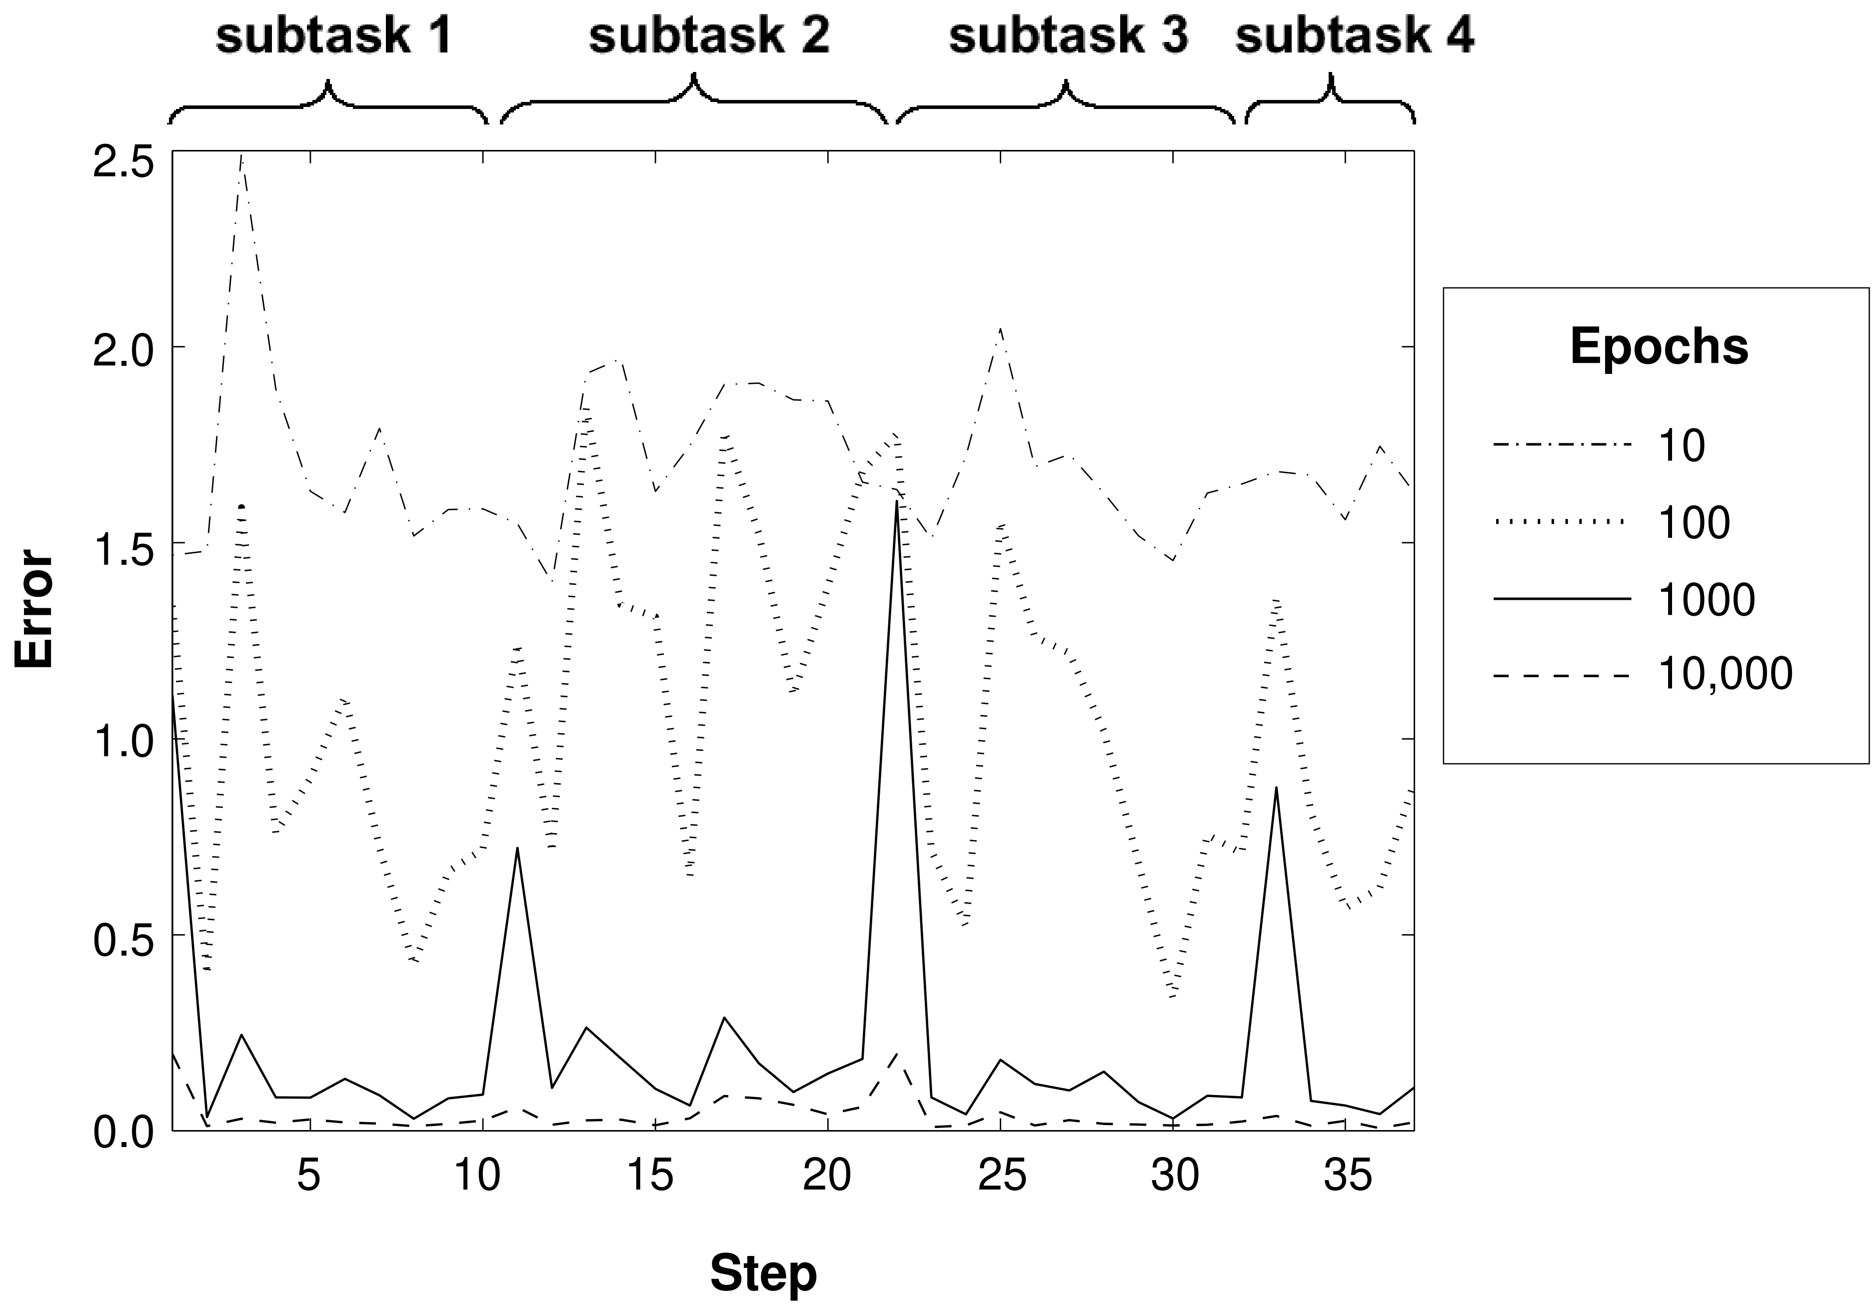

Normal performance

Proportions of testing trials

With coffee instruction

GROUNDS $\longrightarrow$ SUGAR (PACK) $\longrightarrow$ CREAM $\longrightarrow$ DRINK	0.35
GROUNDS $\longrightarrow$ SUGAR (BOWL) $\longrightarrow$ CREAM $\longrightarrow$ DRINK	0.37
GROUNDS $\longrightarrow$ CREAM $\longrightarrow$ SUGAR (PACK) $\longrightarrow$ DRINK	0.14
GROUNDS $\longrightarrow$ CREAM $\longrightarrow$ SUGAR (BOWL) $\longrightarrow$ DRINK	0.14

With tea instruction

TEABAG $\longrightarrow$ SUGAR (PACK) $\longrightarrow$ DRINK	0.46
TEABAG $\longrightarrow$ SUGAR (BOWL) $\longrightarrow$ DRINK	0.54

With no instruction

GROUNDS $\longrightarrow$ SUGAR (PACK) $\longrightarrow$ CREAM $\longrightarrow$ DRINK	0.15
GROUNDS $\longrightarrow$ SUGAR (BOWL) $\longrightarrow$ CREAM $\longrightarrow$ DRINK	0.18
GROUNDS $\longrightarrow$ CREAM $\longrightarrow$ SUGAR (PACK) $\longrightarrow$ DRINK	0.12
GROUNDS $\longrightarrow$ CREAM $\longrightarrow$ SUGAR (BOWL) $\longrightarrow$ DRINK	0.10
TEABAG $\longrightarrow$ SUGAR (PACK) $\longrightarrow$ DRINK	0.20
TEABAG $\longrightarrow$ SUGAR (BOWL) $\longrightarrow$ DRINK	0.25

Normal performance: Task context

How are different task contexts maintained across identical subtasks?

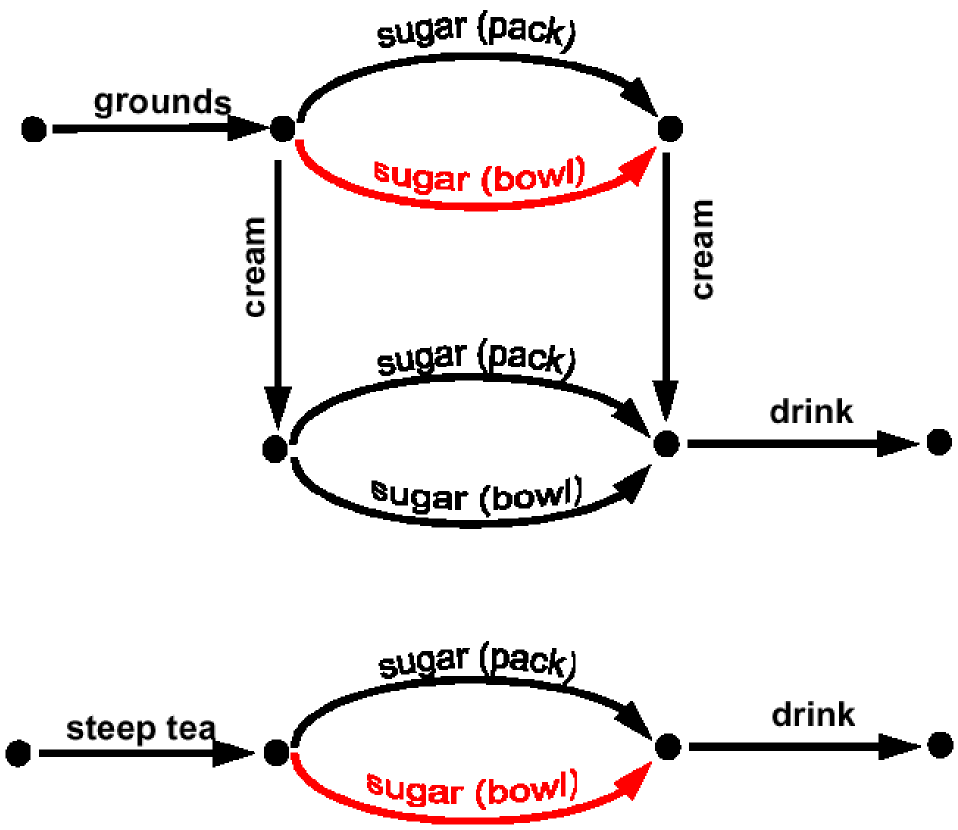

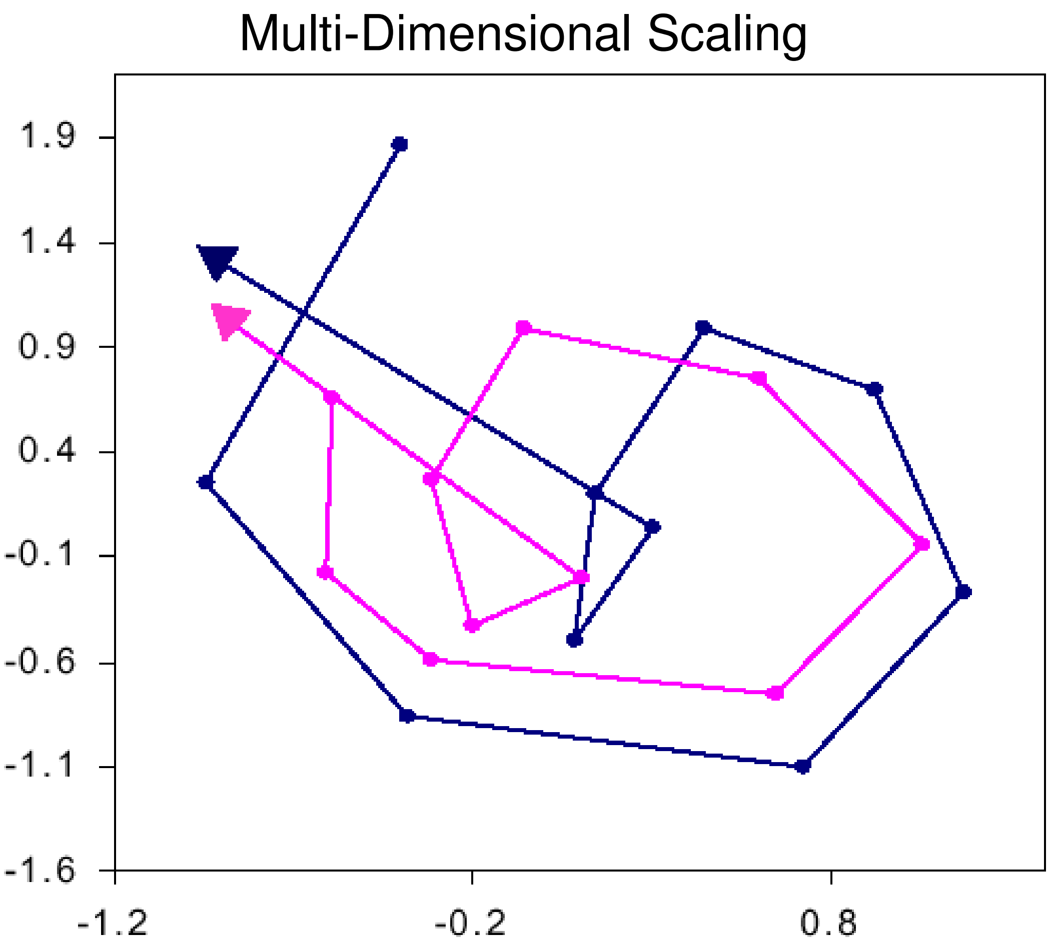

Slips of action (Reason, 1990)

Distraction: Distort context activations with mild-to-moderate noise

Sample of behavior:

pick-up coffee-pack
pull-open coffee-pack
pour coffee-pack into cup
put-down coffee-pack
pick-up spoon
stir cup
put-down spoon

grounds

pick-up sugar-pack
tear-open sugar-pack
pour sugar-pack into cup
put-down sugar-pack
pick-up spoon
stir cup
put-down spoon

sugar (pack)

cream omitted → **pick-up cup***
sip cup
sip cup
say-done

drink

Slips of action

Errors occur at **decision points** (boundaries between subtasks)

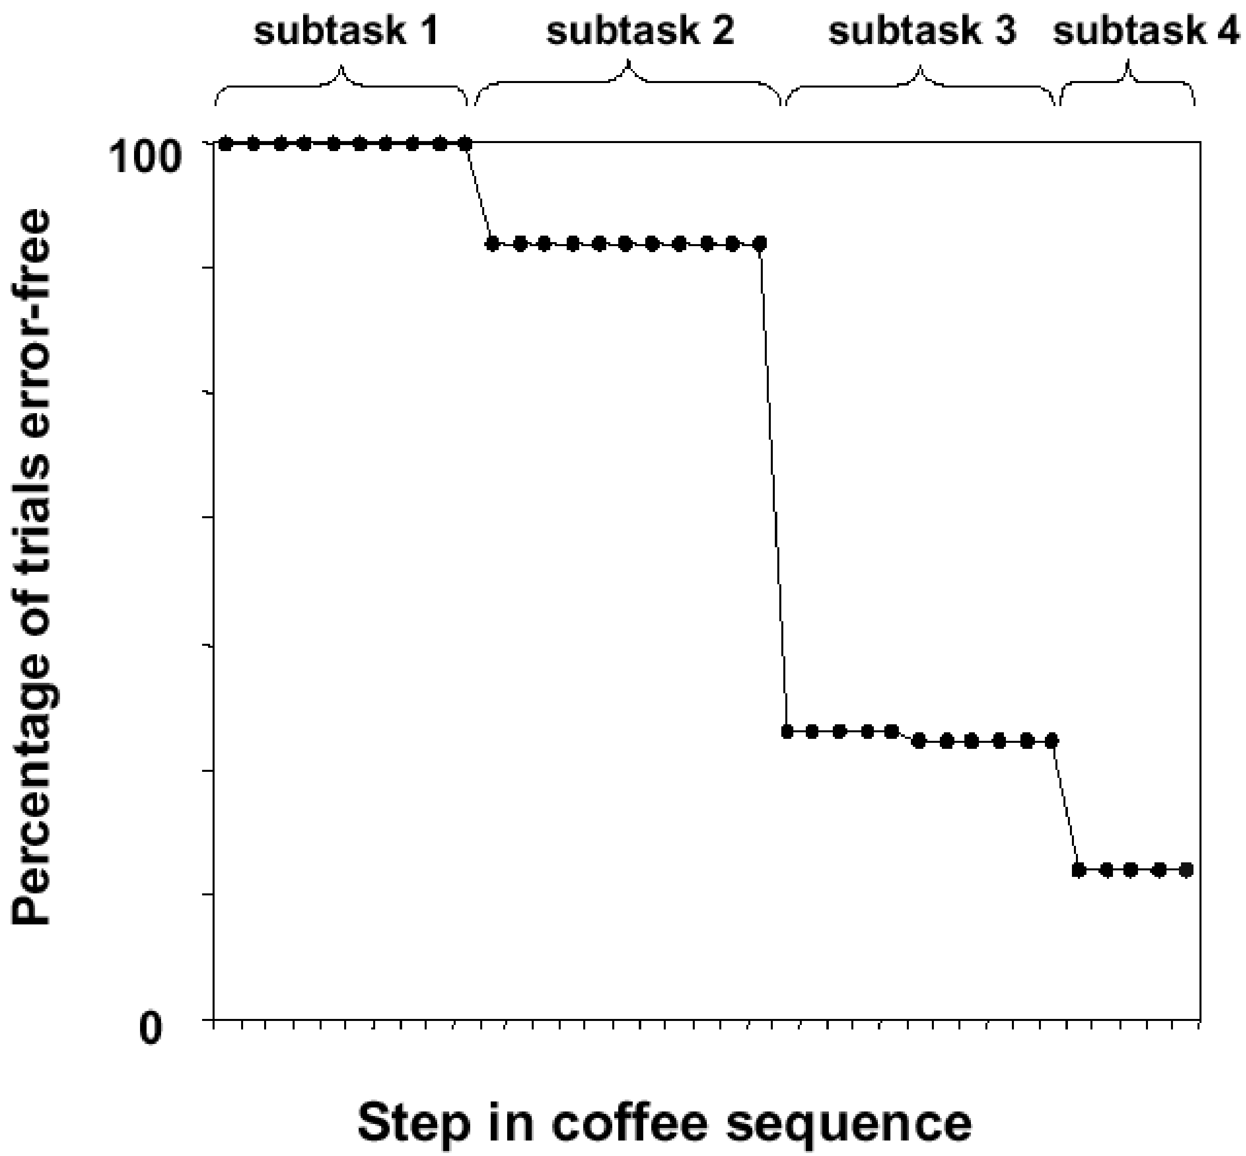

Slips of action

Errors take the form of **displaced but intact subtask sequences**

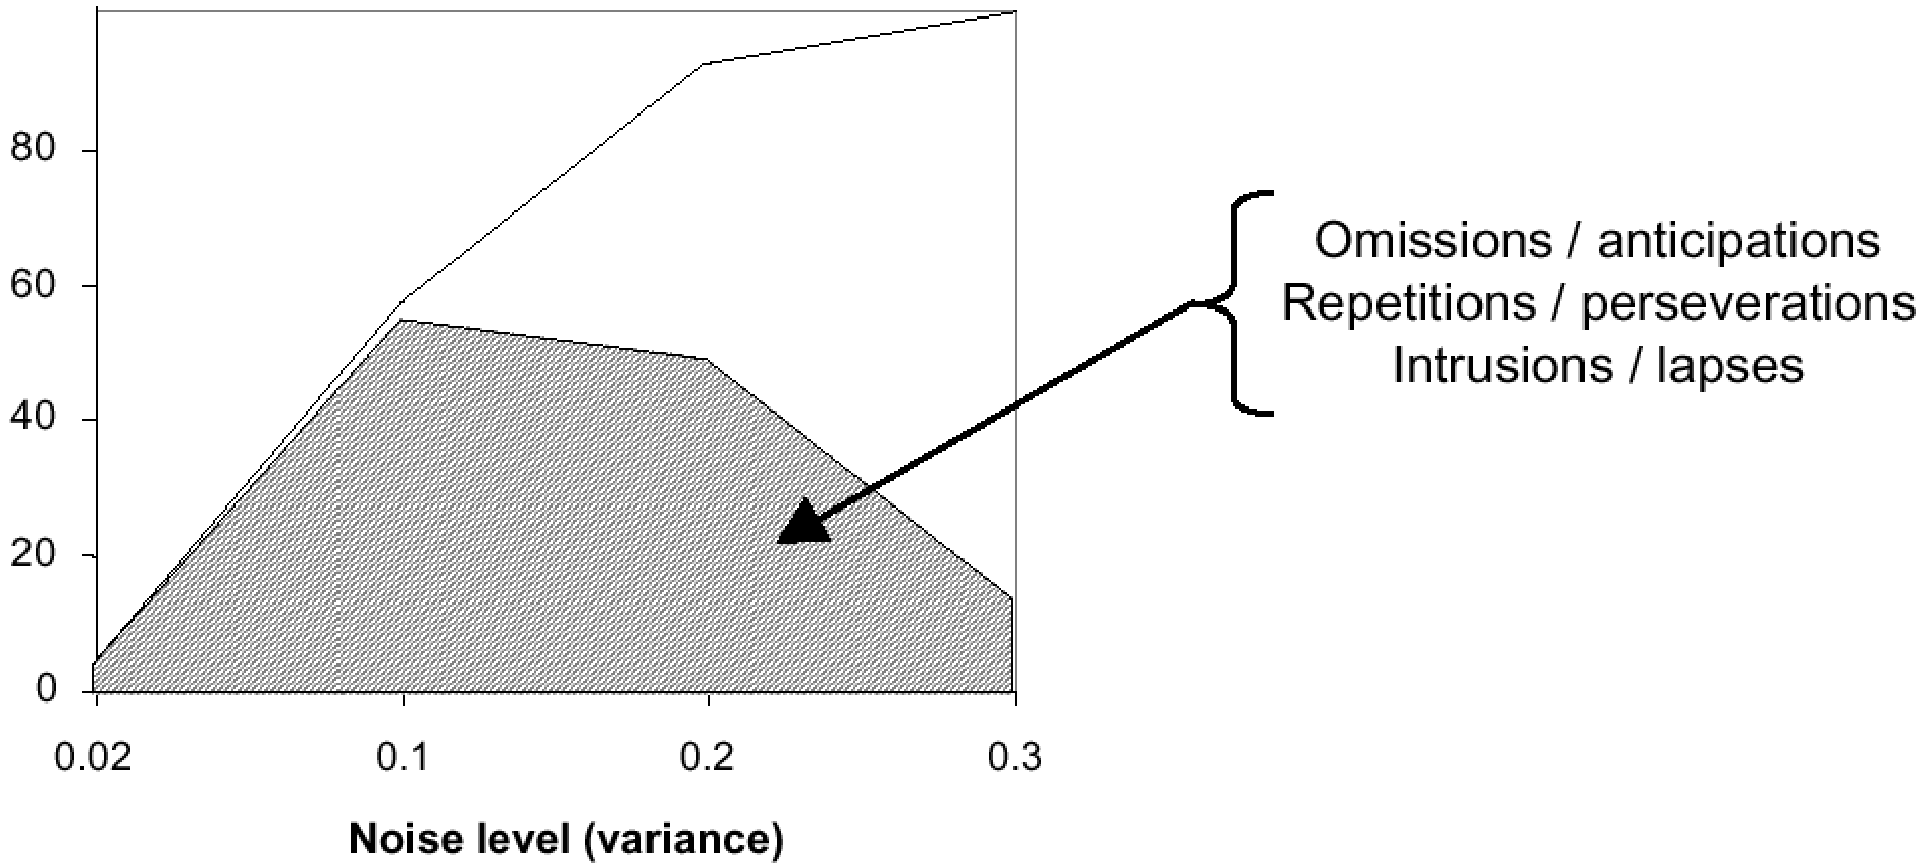

Slips of action

Lapses typically involve shift **from less frequent to more frequent task**

STEEP-TEA $\Rightarrow$ ADD-SUGAR $\Rightarrow$ ADD-CREAM*

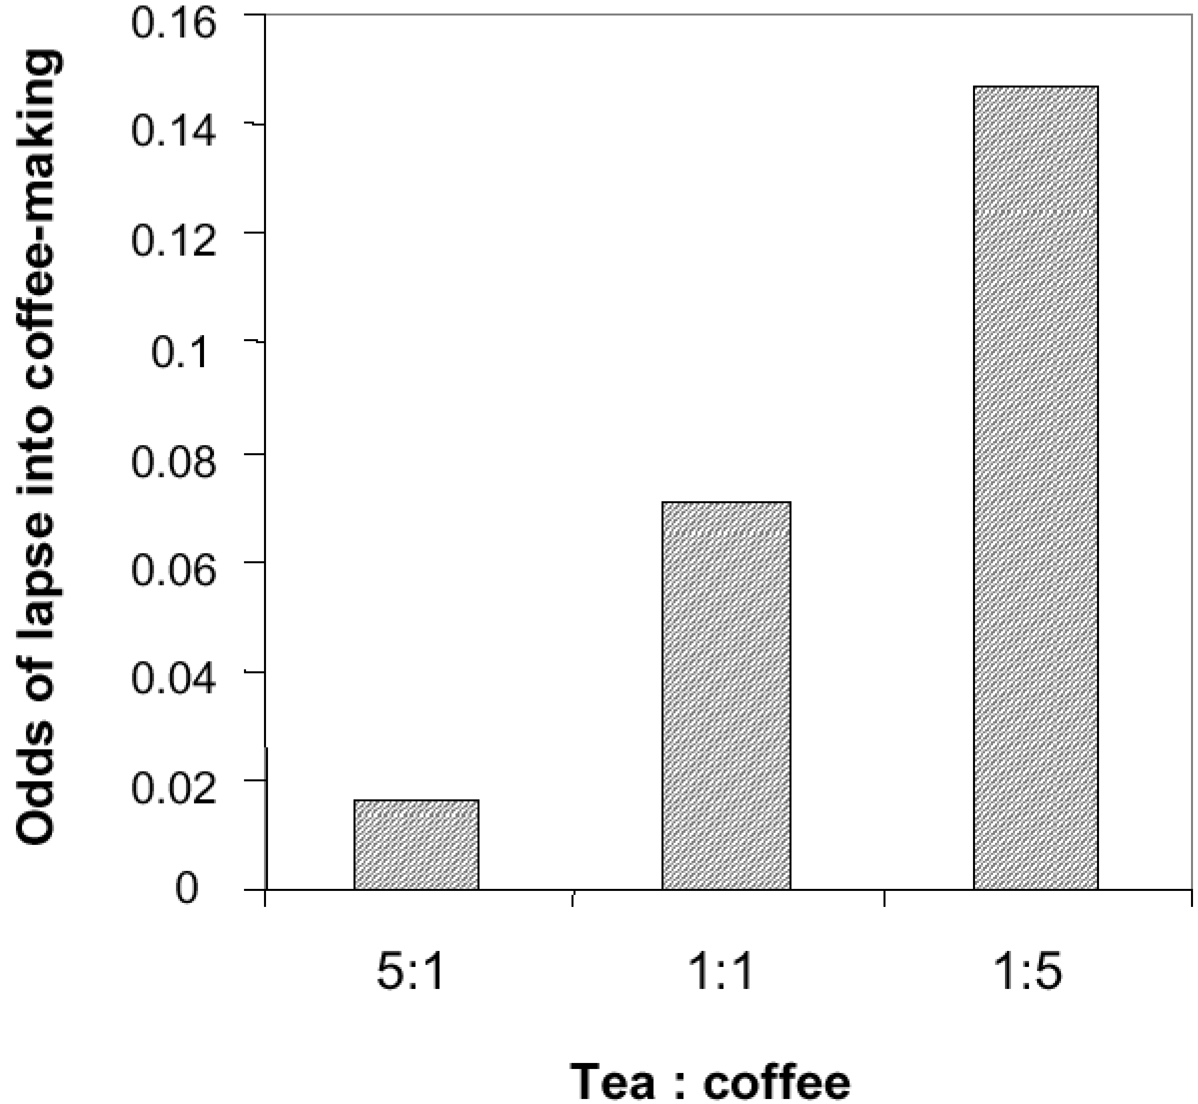

Prediction: Timing of distraction

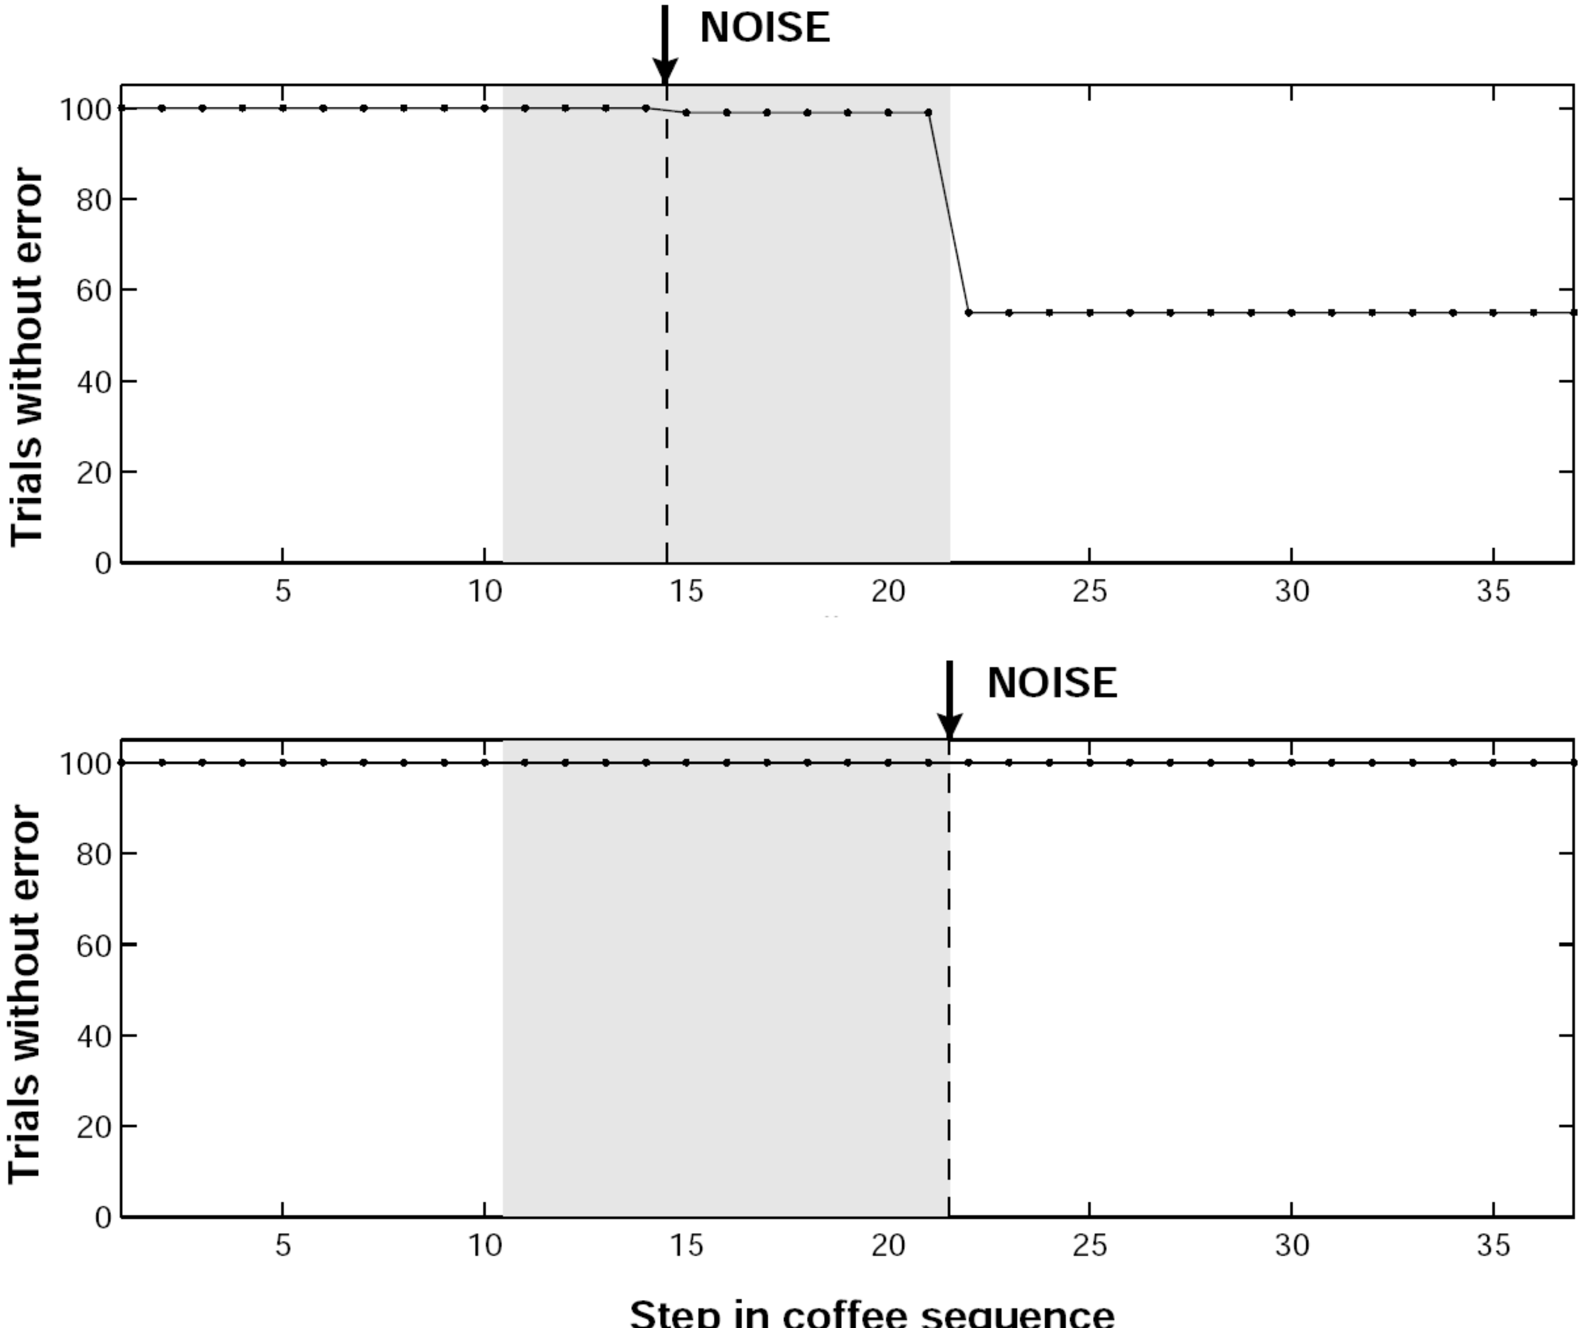

Action Disorganization Syndrome (Schwartz et al., 1991)

Neural damage: Distort context activations with severe noise

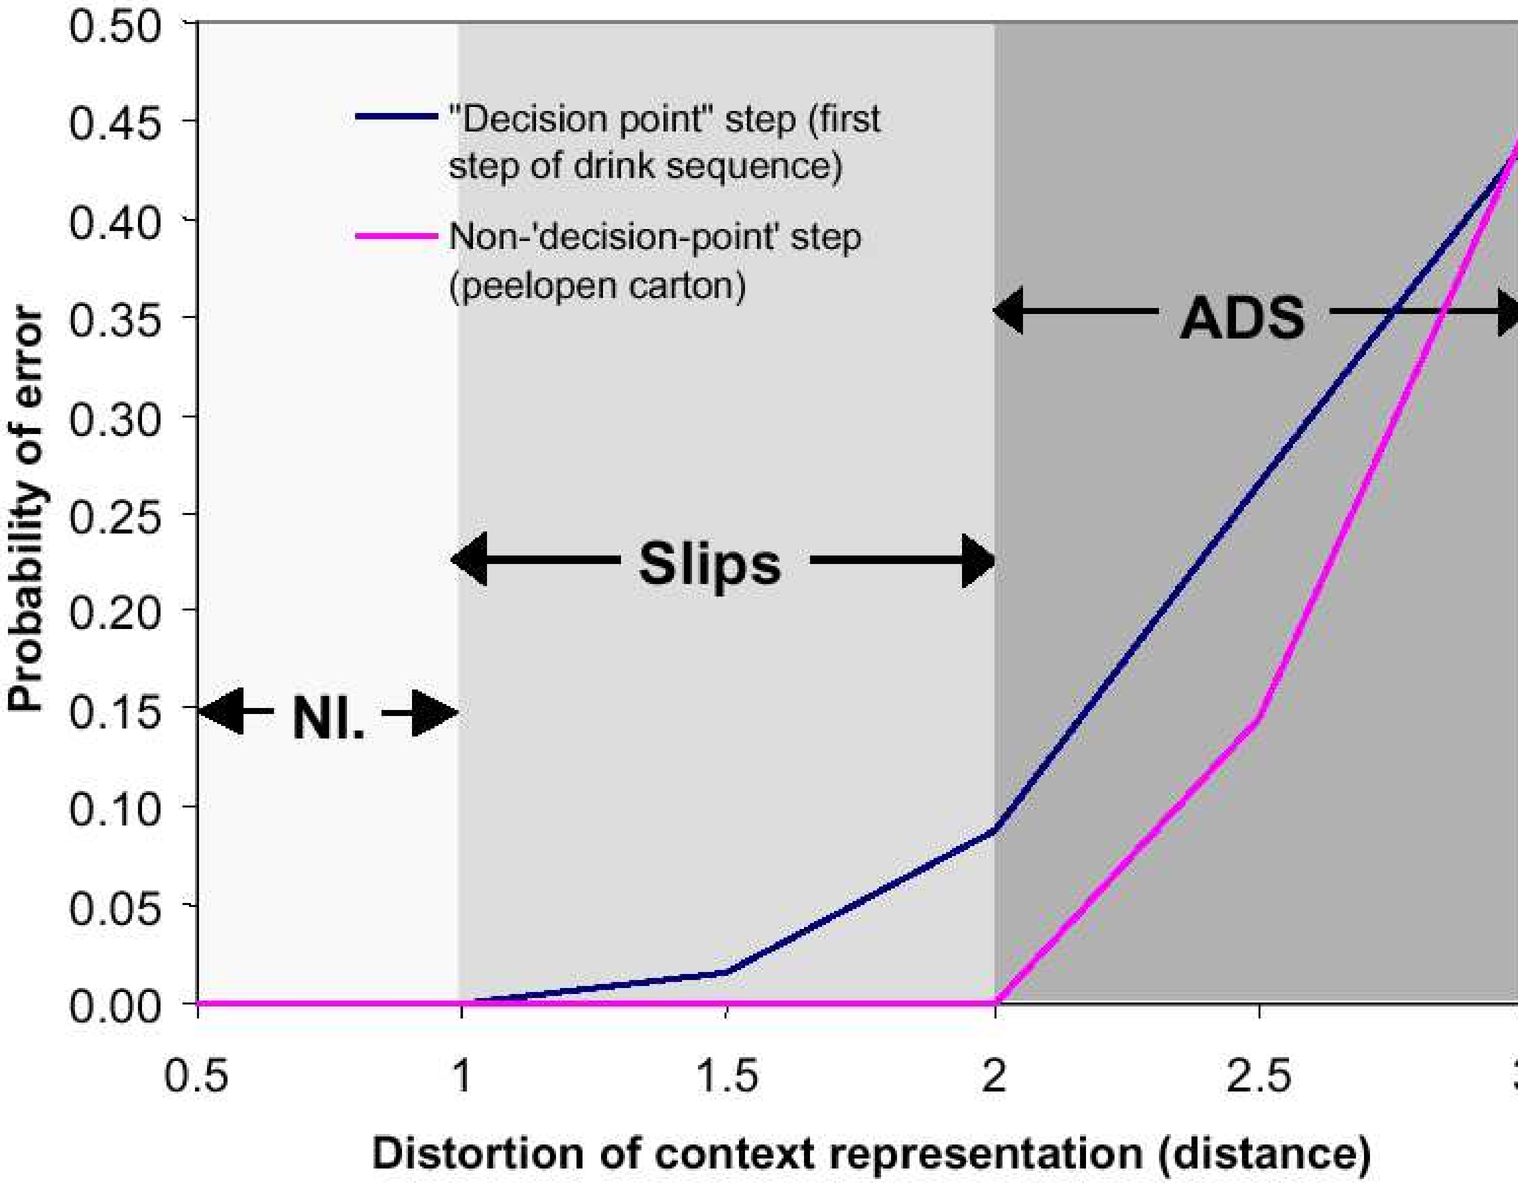

Action Disorganization Syndrome

Sample of behavior:

pick-up coffee-pack
pull-open coffee-pack
put-down coffee-pack*
pick-up coffee-pack
disrupted subtask
pour coffee-pack into cup
put-down coffee-pack
pick-up spoon
stir cup
put-down spoon
pick-up sugar-pack
tear-open sugar-pack
subtask fragment
pour sugar-pack into cup
put-down sugar-pack
pick-up cup*
put-down cup
sugar repeated
pull-off sugarbowl lid
put-down lid
pick-up spoon
subtask fragment
scoop sugarbowl with spoon
put-down spoon*
pick-up cup*
sip cup
cream omitted
sip cup
say-done

Action Disorganization Syndrome

Decrease in **independent actions** with decreasing ADS severity

Schwartz et al. (1991, *Cog. Neuropsych.*)

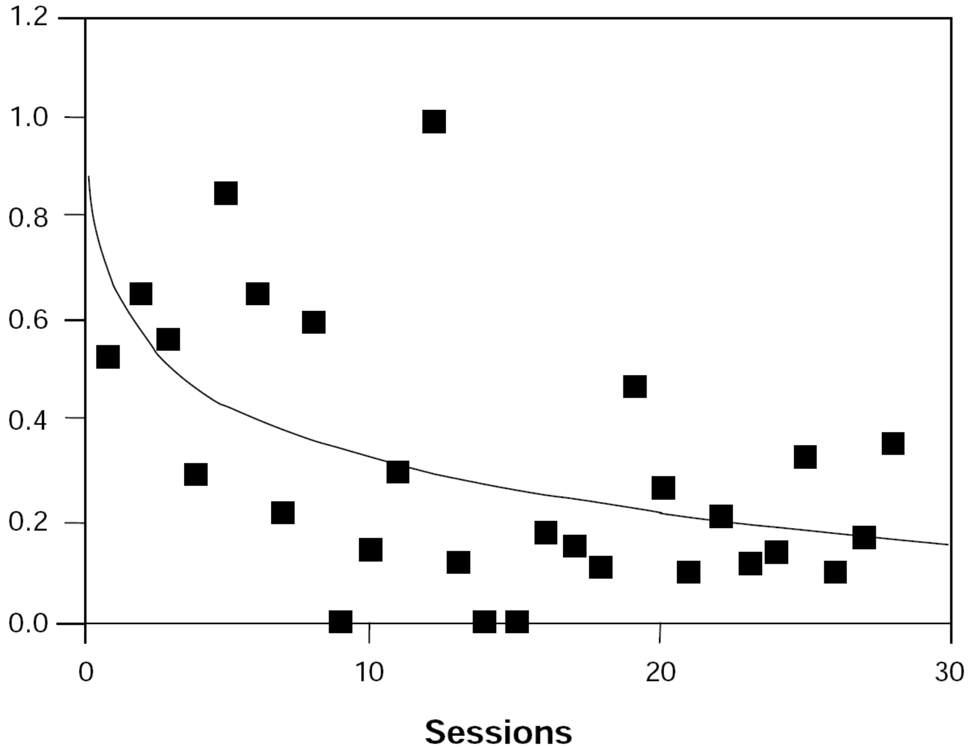

Model

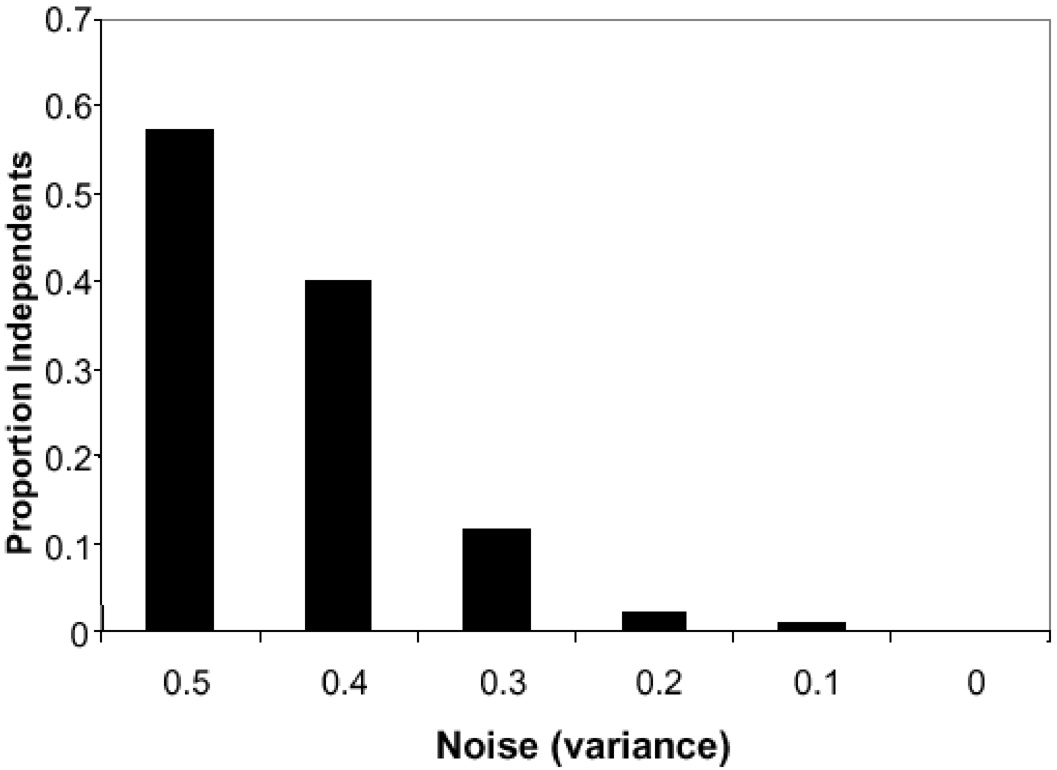

Action Disorganization Syndrome

Patients who make more errors commit a higher proportion of **omission errors**

Schwartz et al. (1998, *Neuropsych.*)

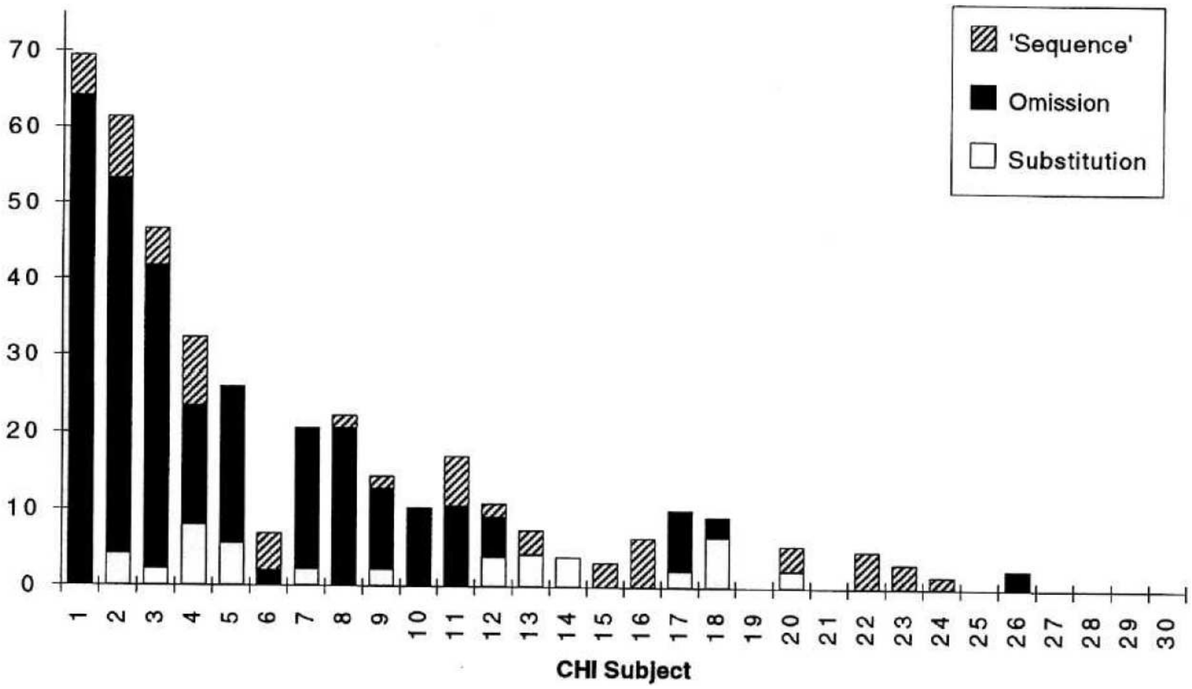

Model

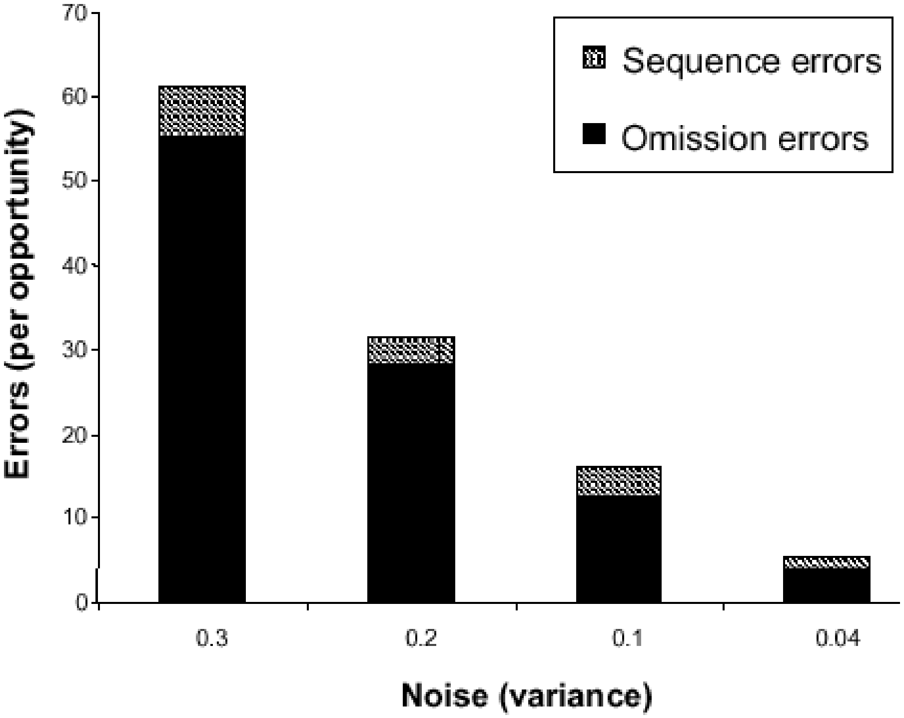

Recurrent networks and "chaining"

- Recurrent networks thought to depend on item-item associations (chaining)
- Incompatible with findings from **immediate serial recall**

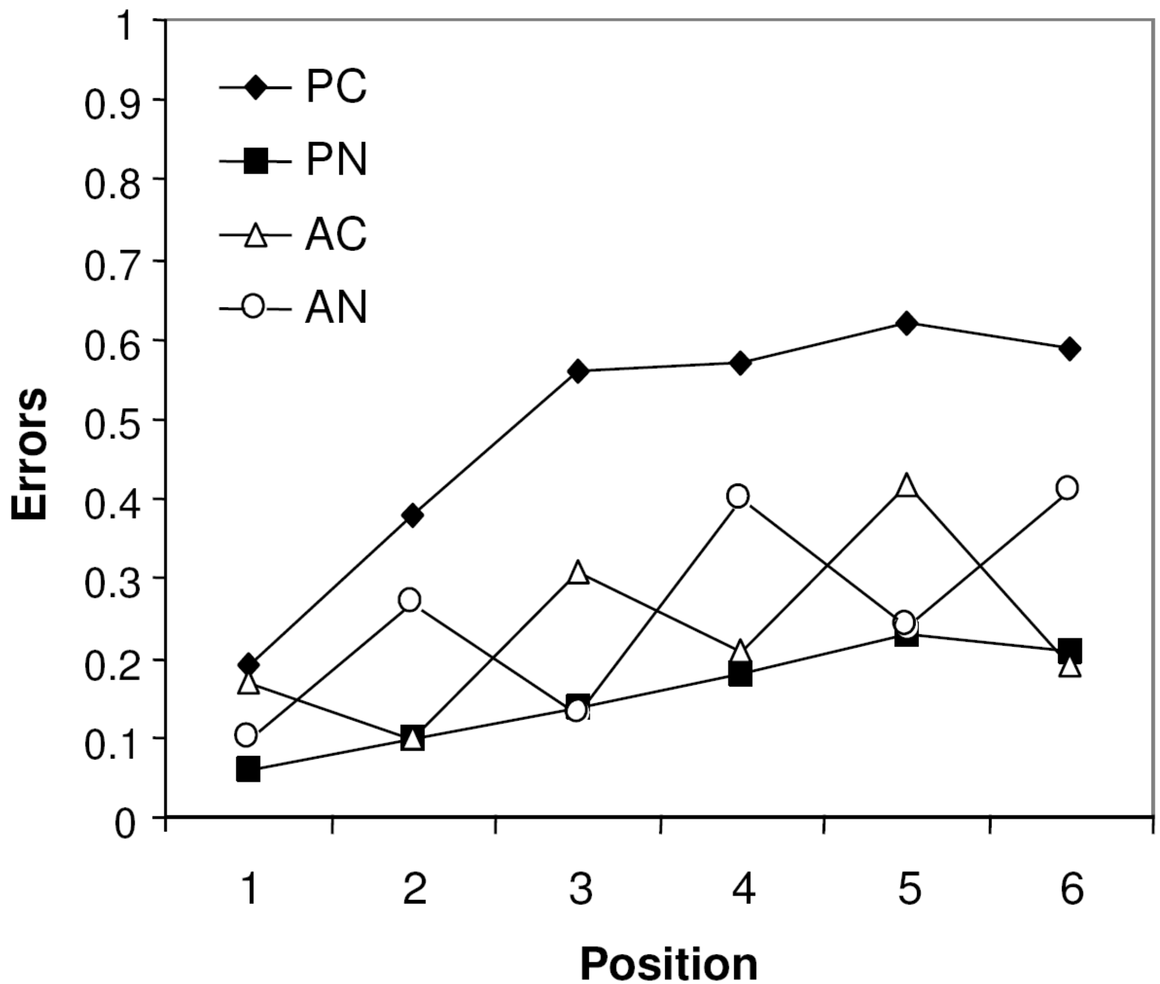

Pure/Alternating
Confusable/Nonconfusable

Example AC list:

B R **D** Q **P** L

Baddeley (1968, *QJEP*); Henson, Norris, Page, & Baddeley (1996, *QJEP*)

Context-association models of serial recall

Burgess and Hitch (1992)
Henson (1996, 1998)
Houghton (1990)
Brown, Preece, and Hulme (2000)

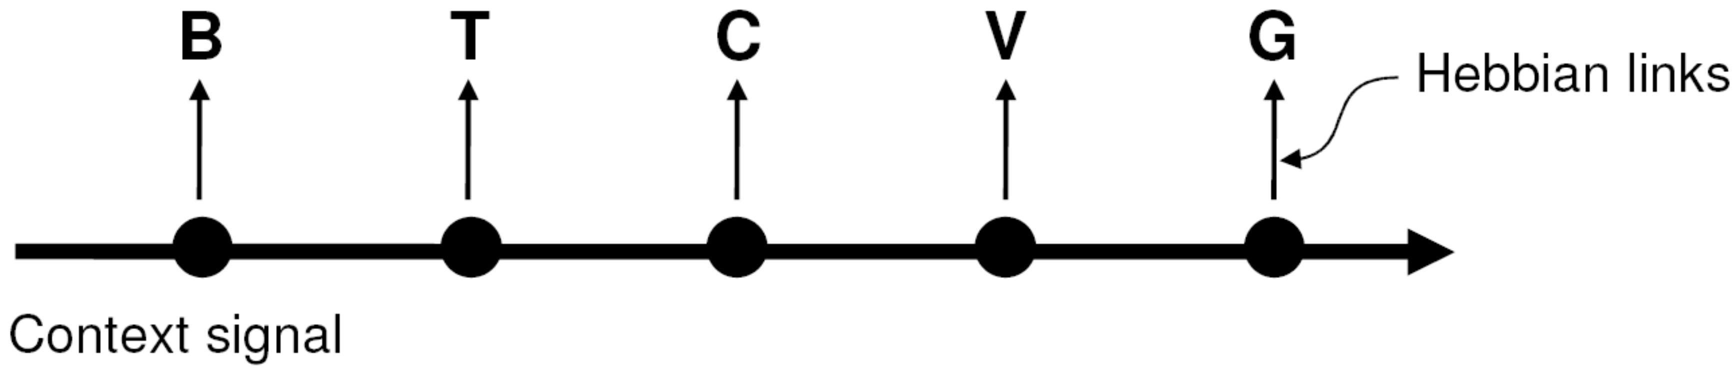

"Interactions between short- and long-term memory pose problems for most models of serial recall."

—Henson (1998, *Cog. Psych.*)

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

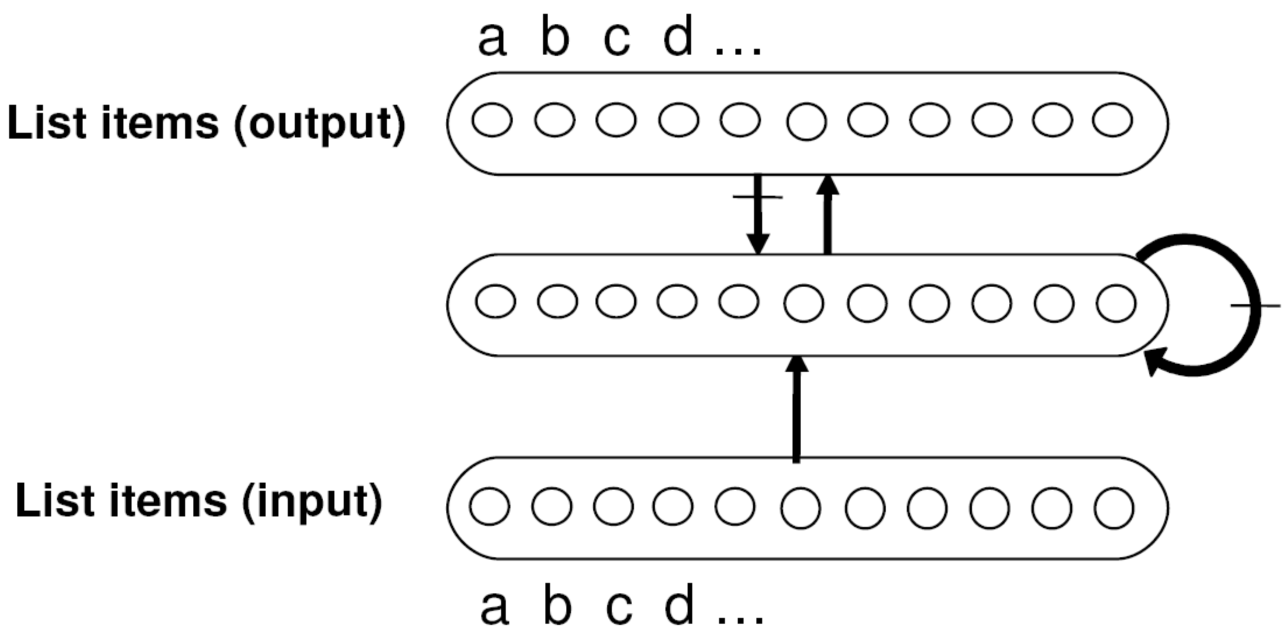

- Trained and tested on ISR (list lengths 1-9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

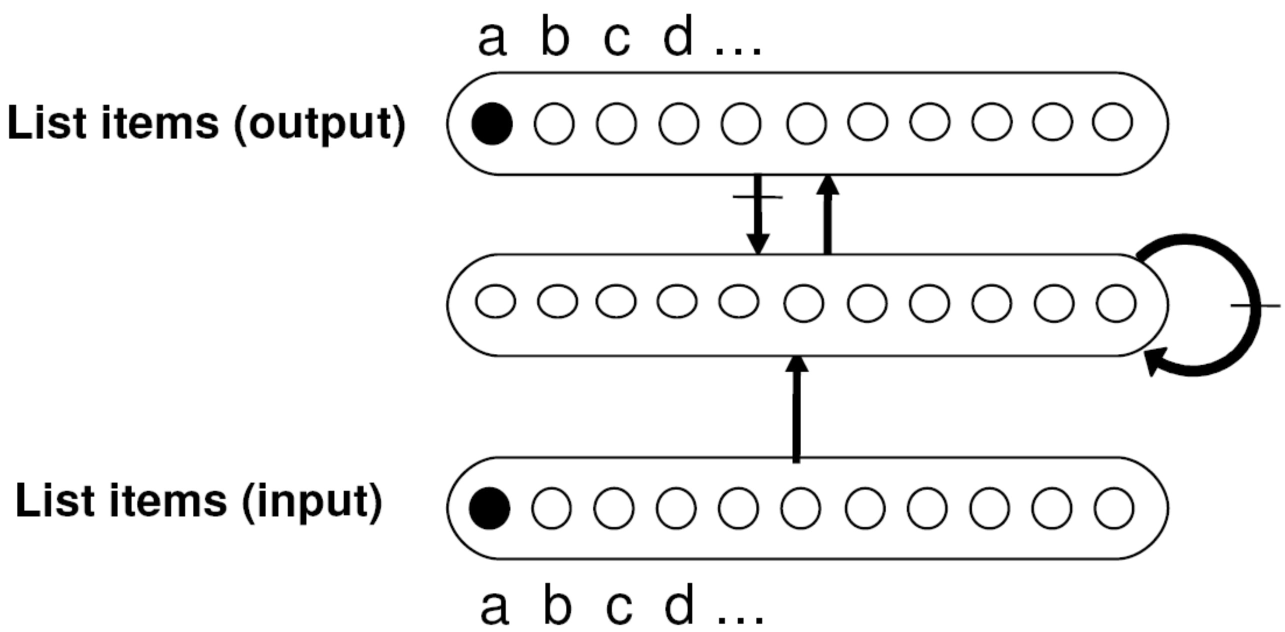

- Trained and tested on ISR (list lengths 1–9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

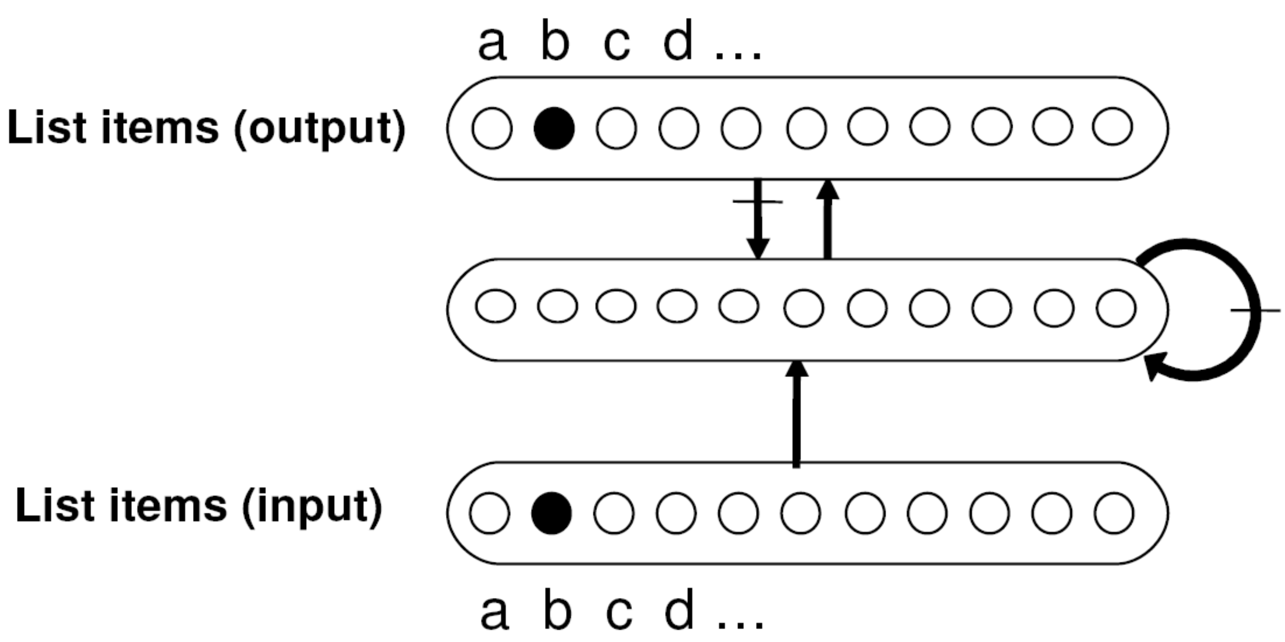

- Trained and tested on ISR (list lengths 1-9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

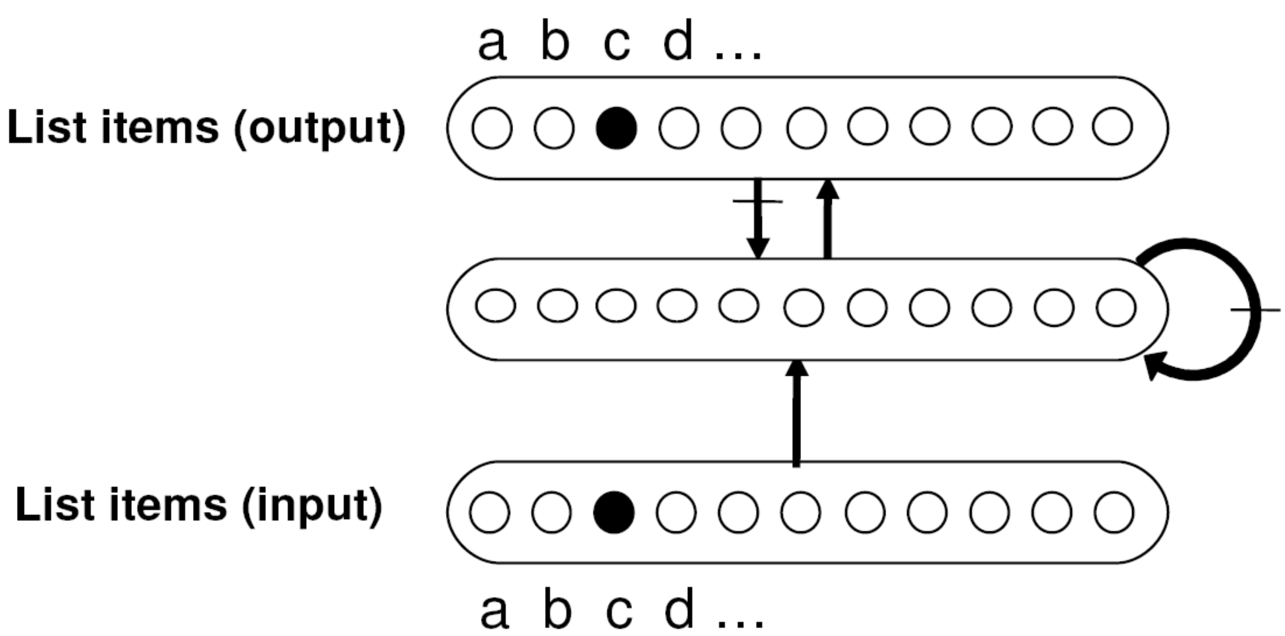

- Trained and tested on ISR (list lengths 1-9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

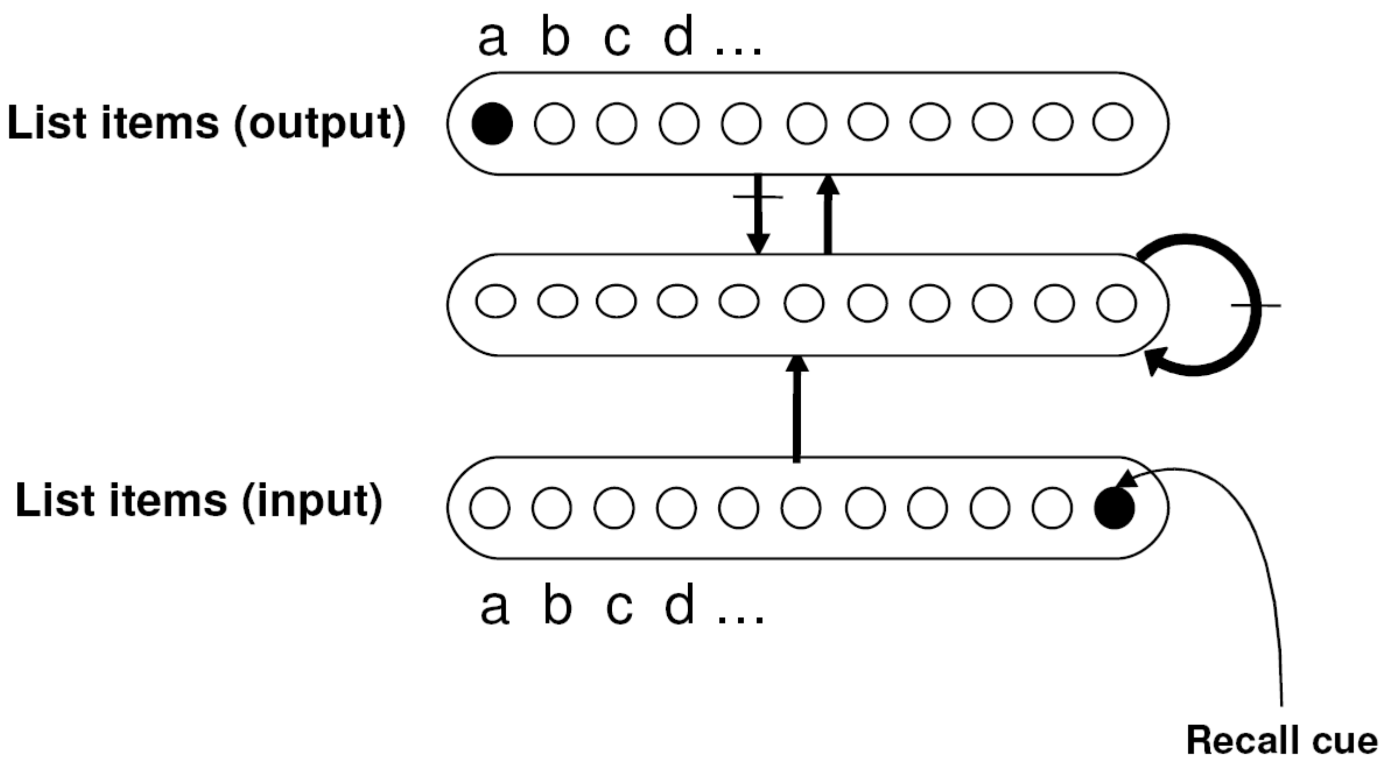

- Trained and tested on ISR (list lengths 1-9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

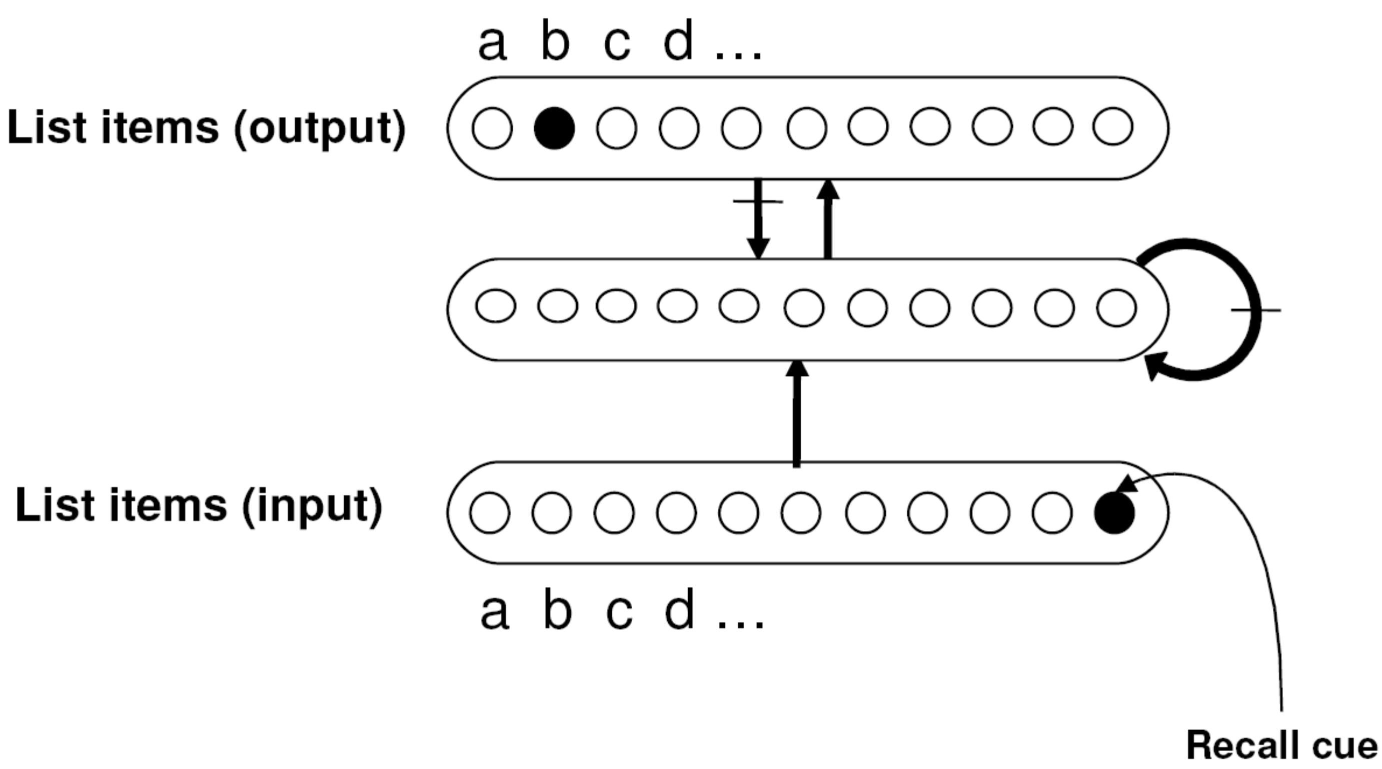

- Trained and tested on ISR (list lengths 1-9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

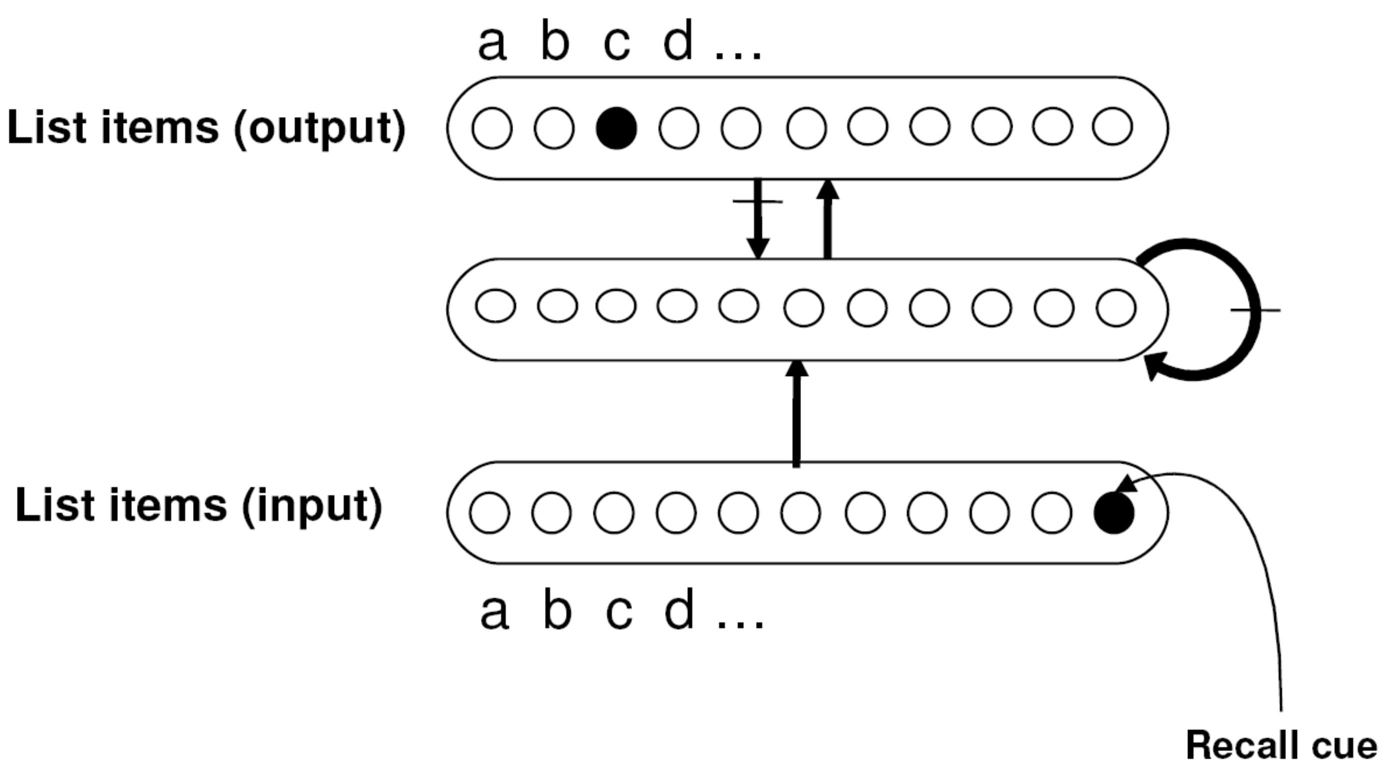

- Trained and tested on ISR (list lengths 1-9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

A recurrent network model of immediate serial recall

Botvinick and Plaut (2006, *Psych. Rev.*)

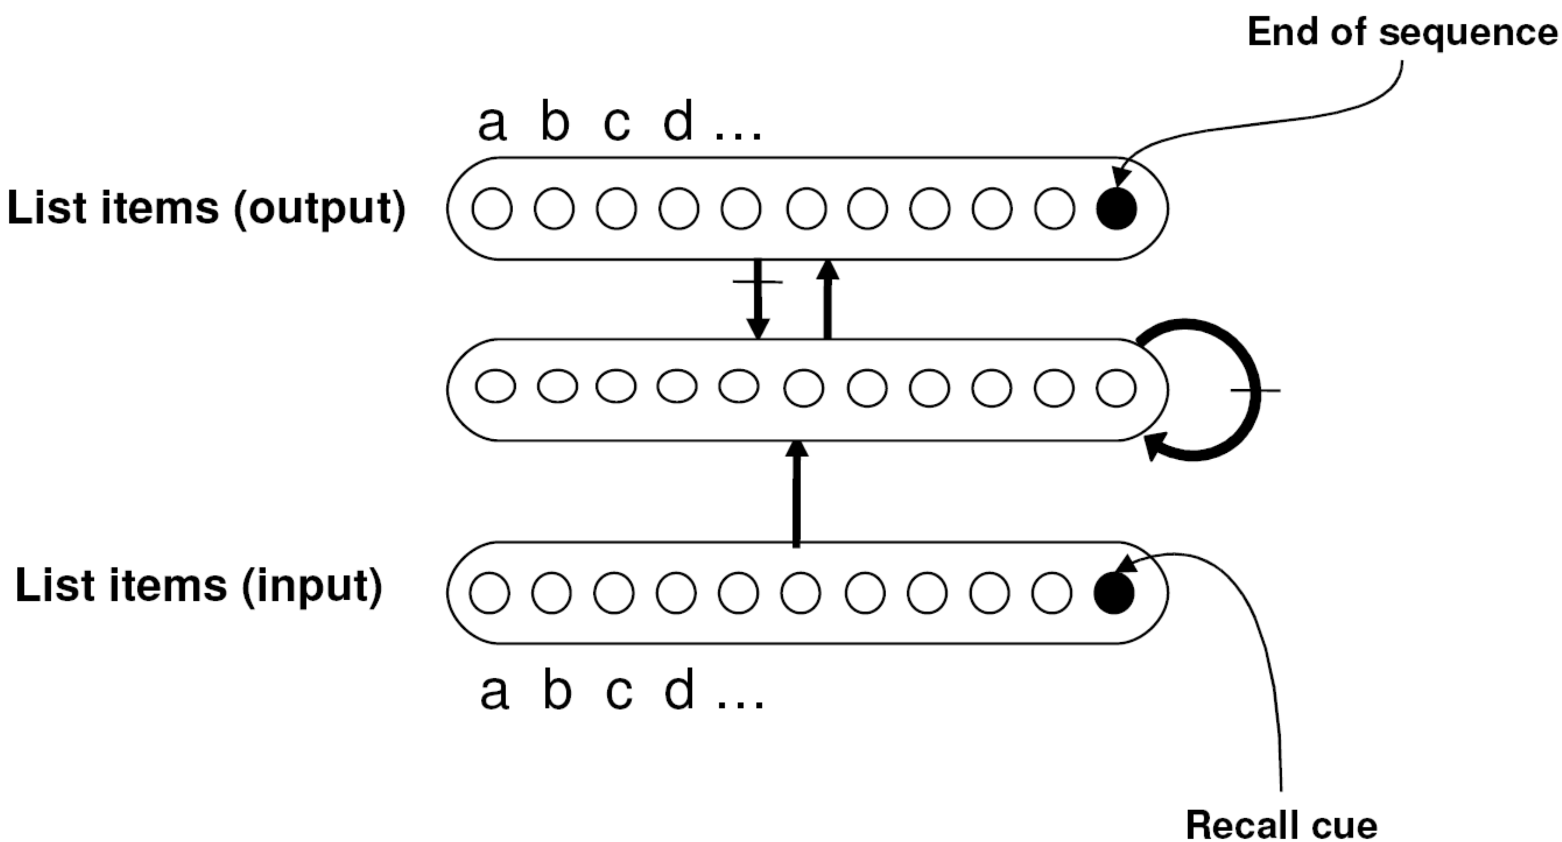

- Trained and tested on ISR (list lengths 1-9); proxy for *language* learning
- Weights not allowed to change during testing
- Three versions: localist inputs; inputs with similarities; bigram frequencies

Results: List length

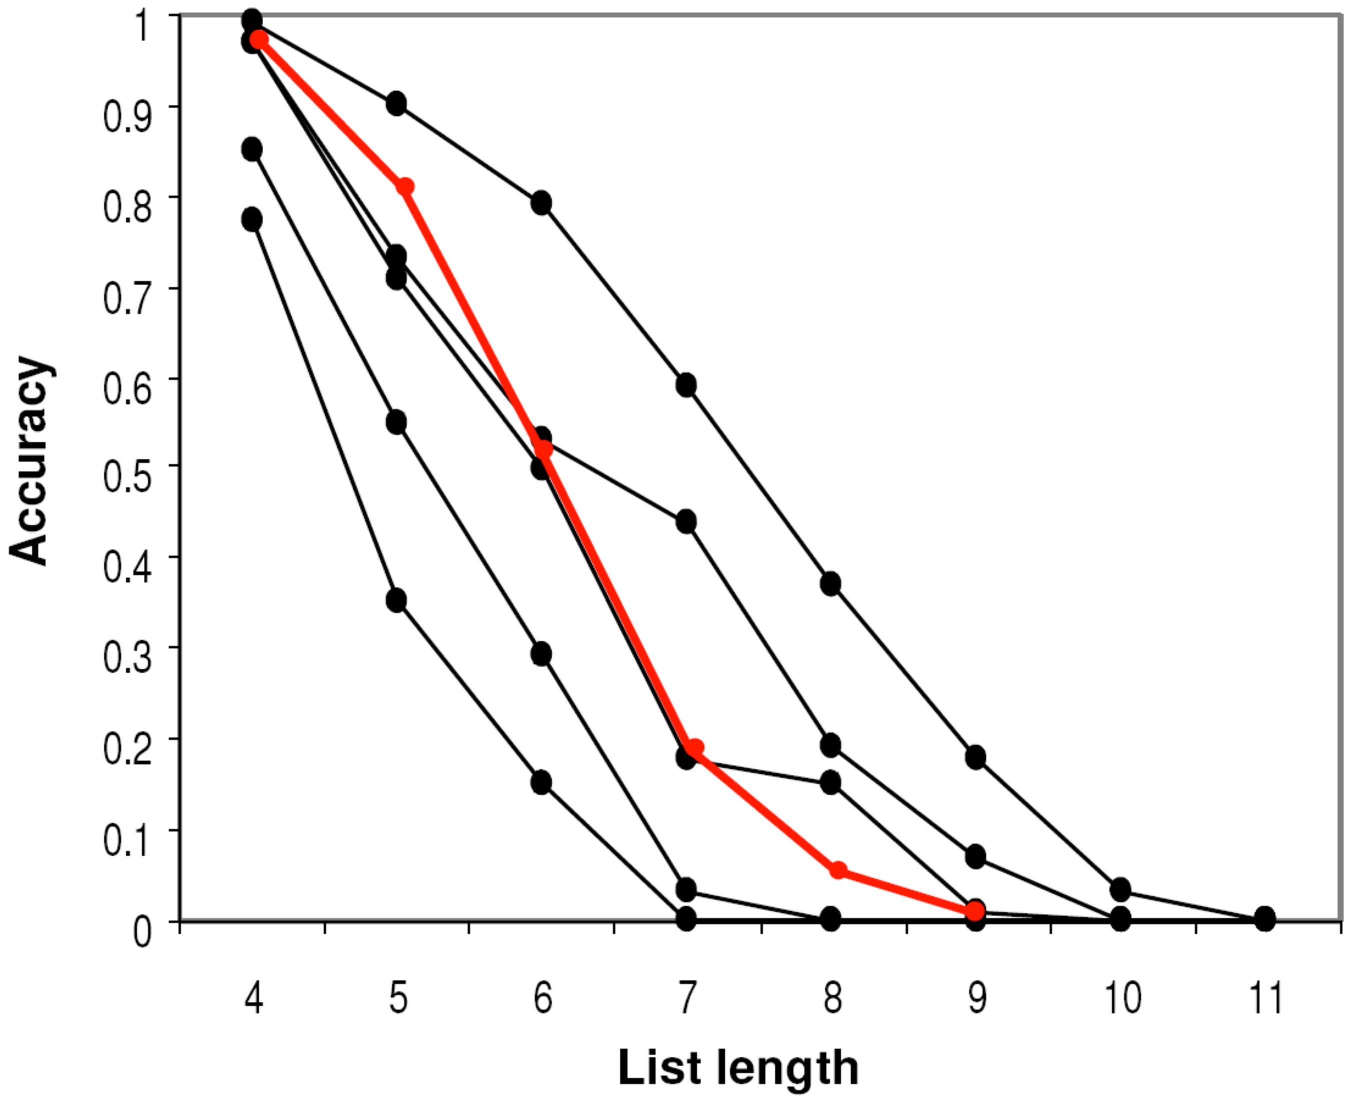

Crannel and Parrish (1957, J. Psychol.)

Results: Primacy, recency, transpositions

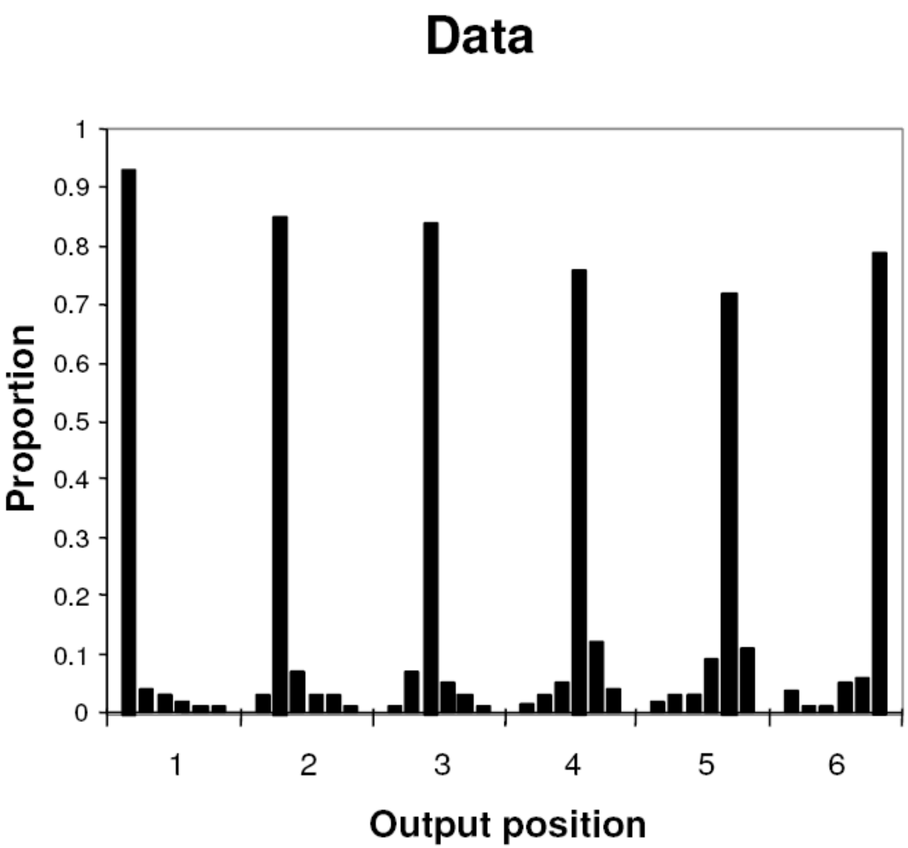

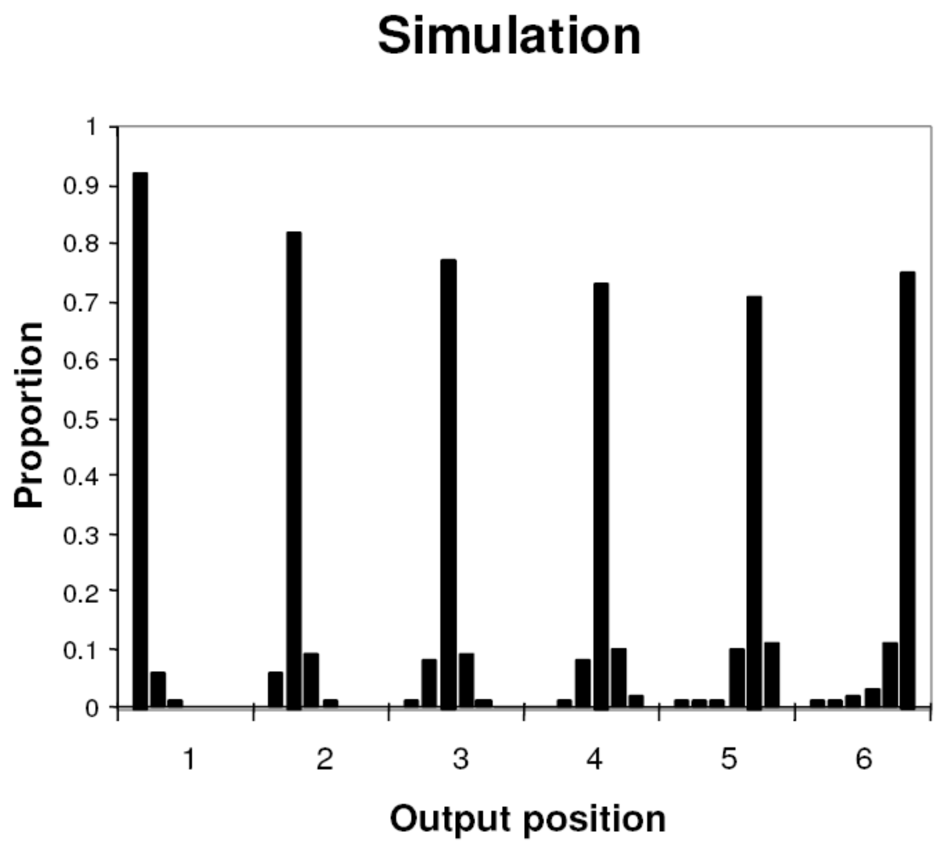

Henson, Norris, Page and Baddeley (1996, *QJEP*)

Results: Bigram frequencies

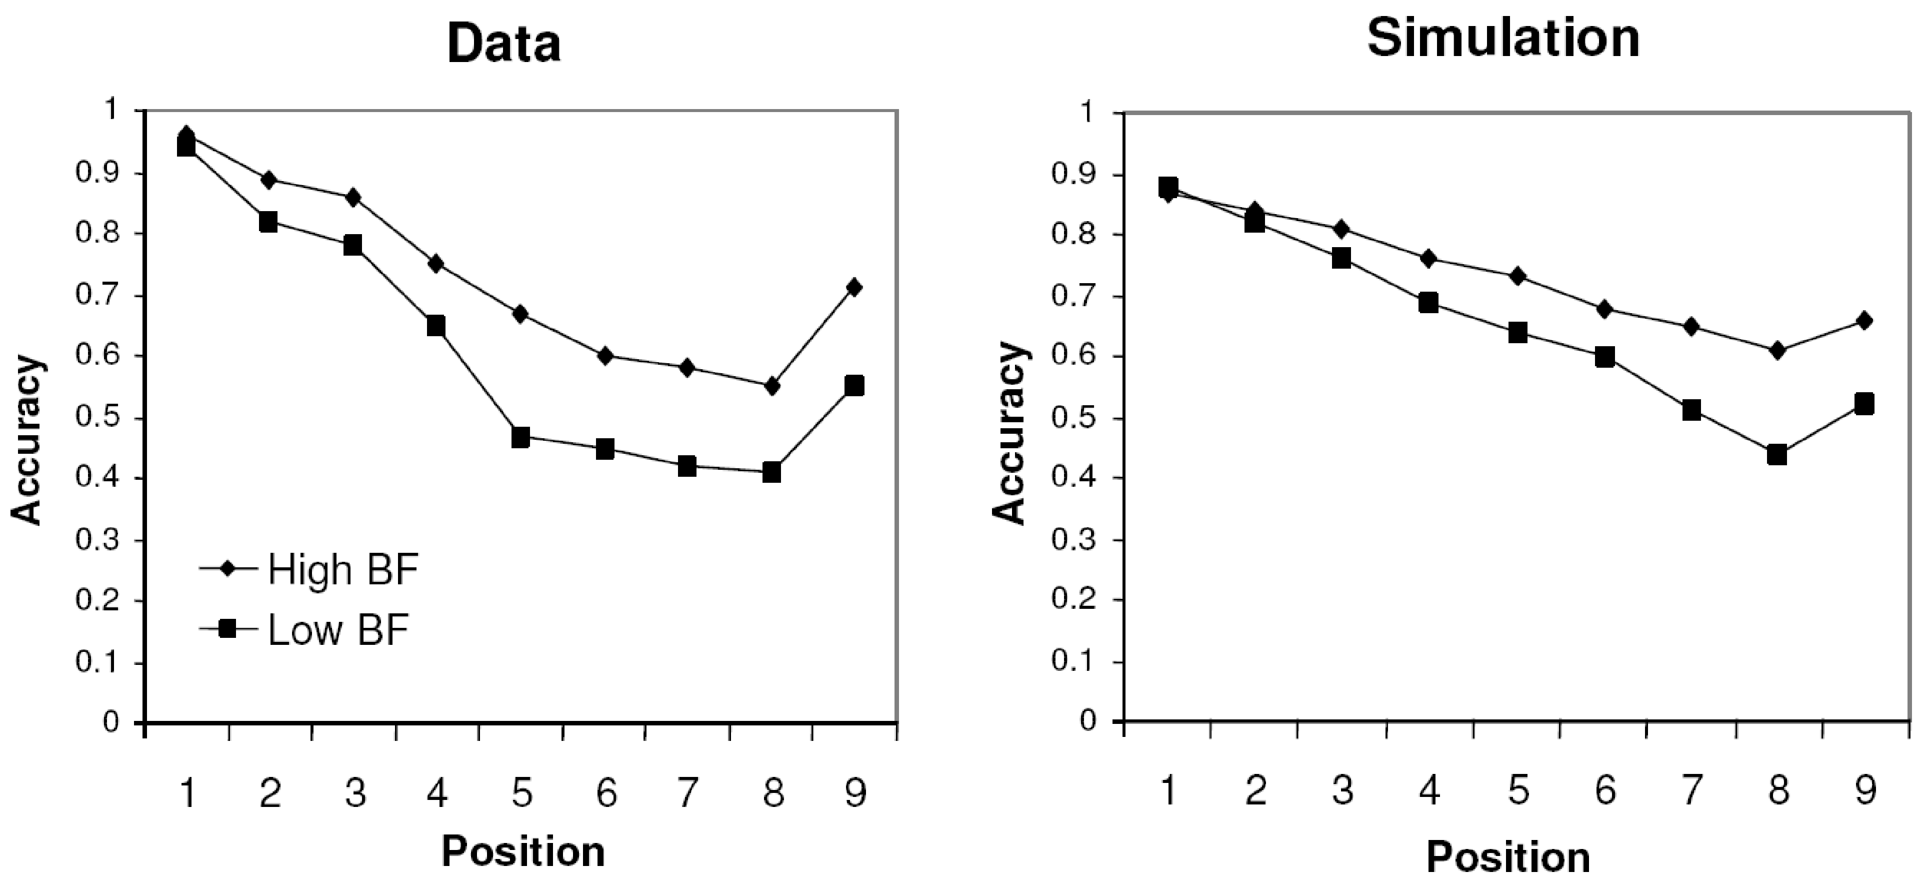

Kantowitz, Ornstein, and Schwartz (1972, J. Exp. Psych.)

Results: "Sawtooth" pattern

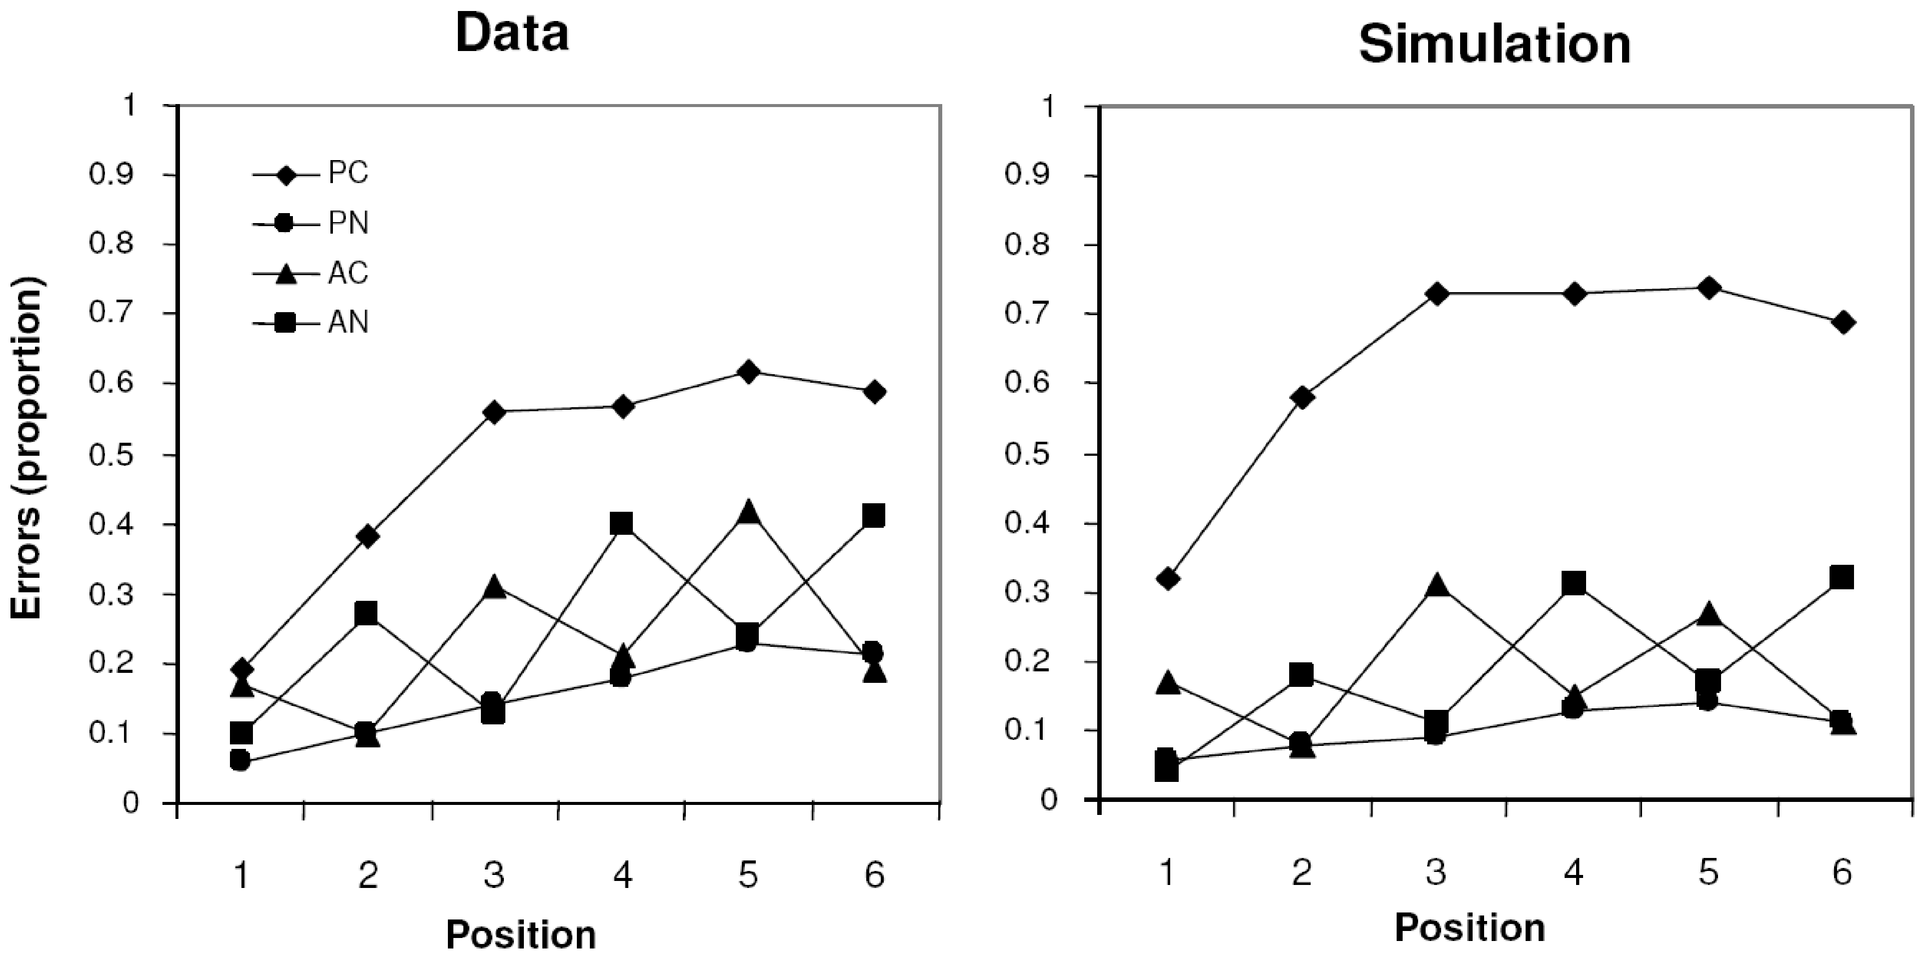

Baddeley (1968, *QJEP*)

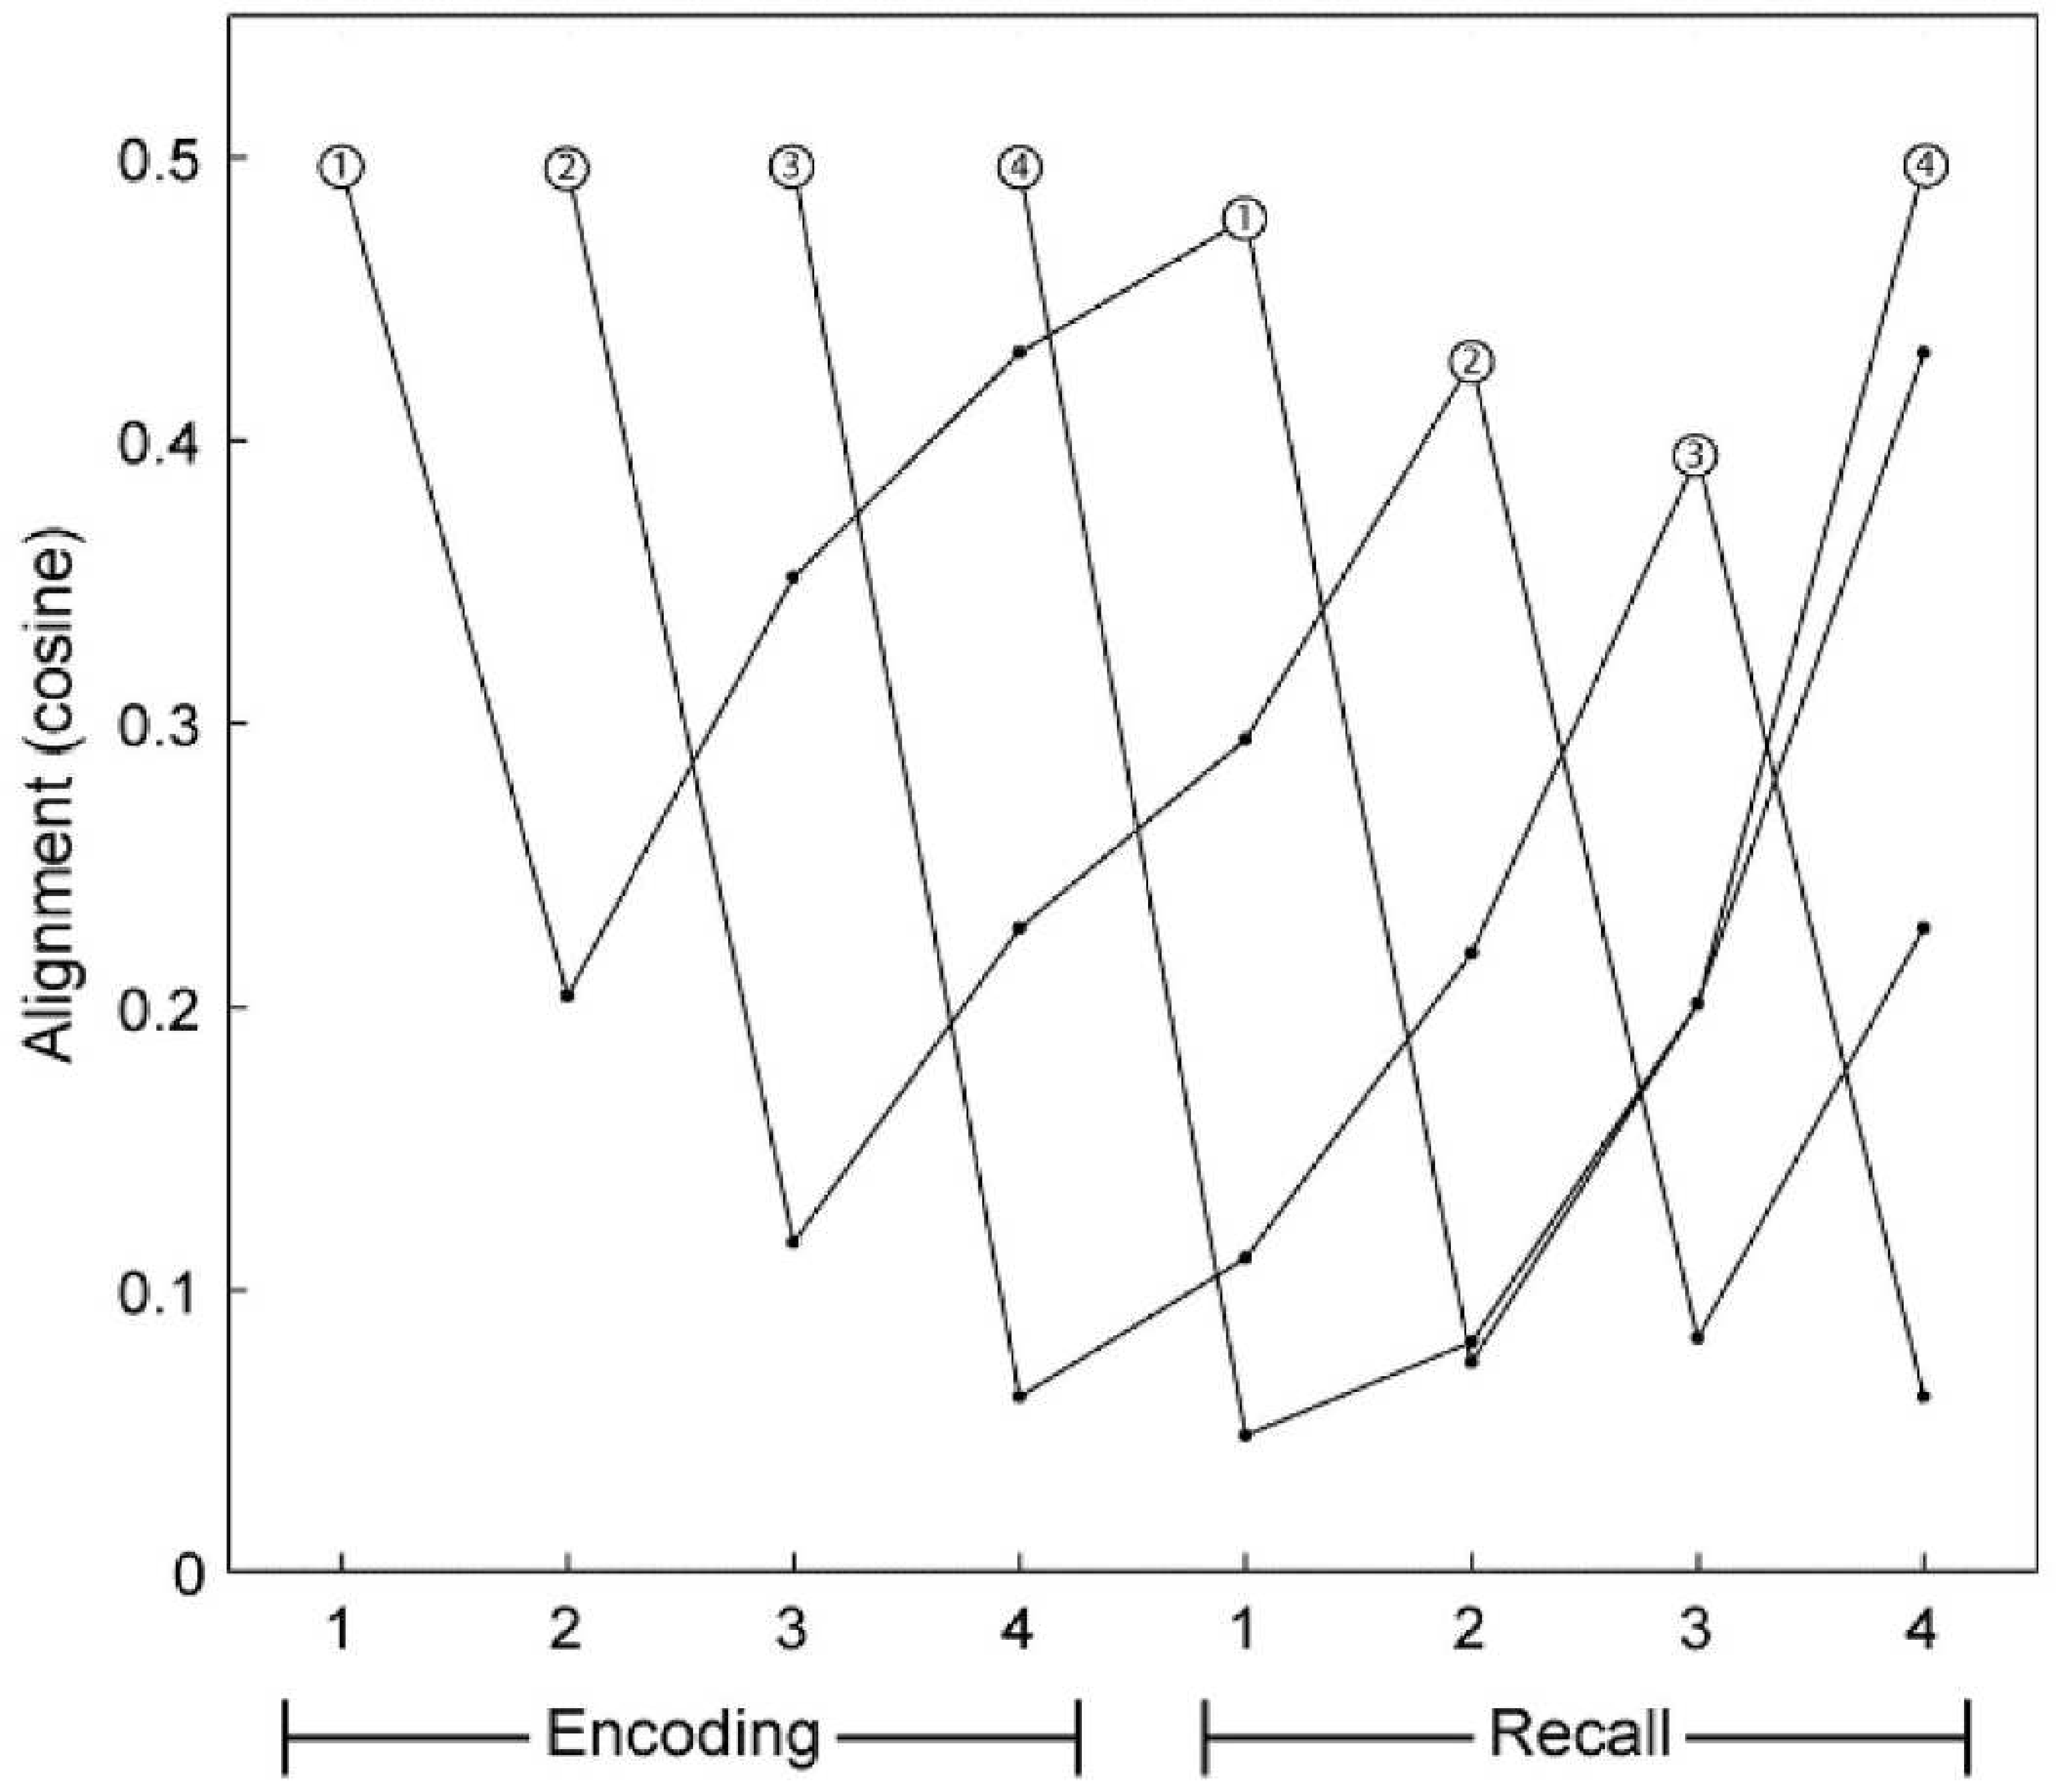
Alignment (cosine)
0.5
0.4
0.3
0.2
0.1
0
1
2
3
4
1
2
3
4
1
2
3
4
1
2
3
4
Encoding
Recall

Analysis: Representational similarity

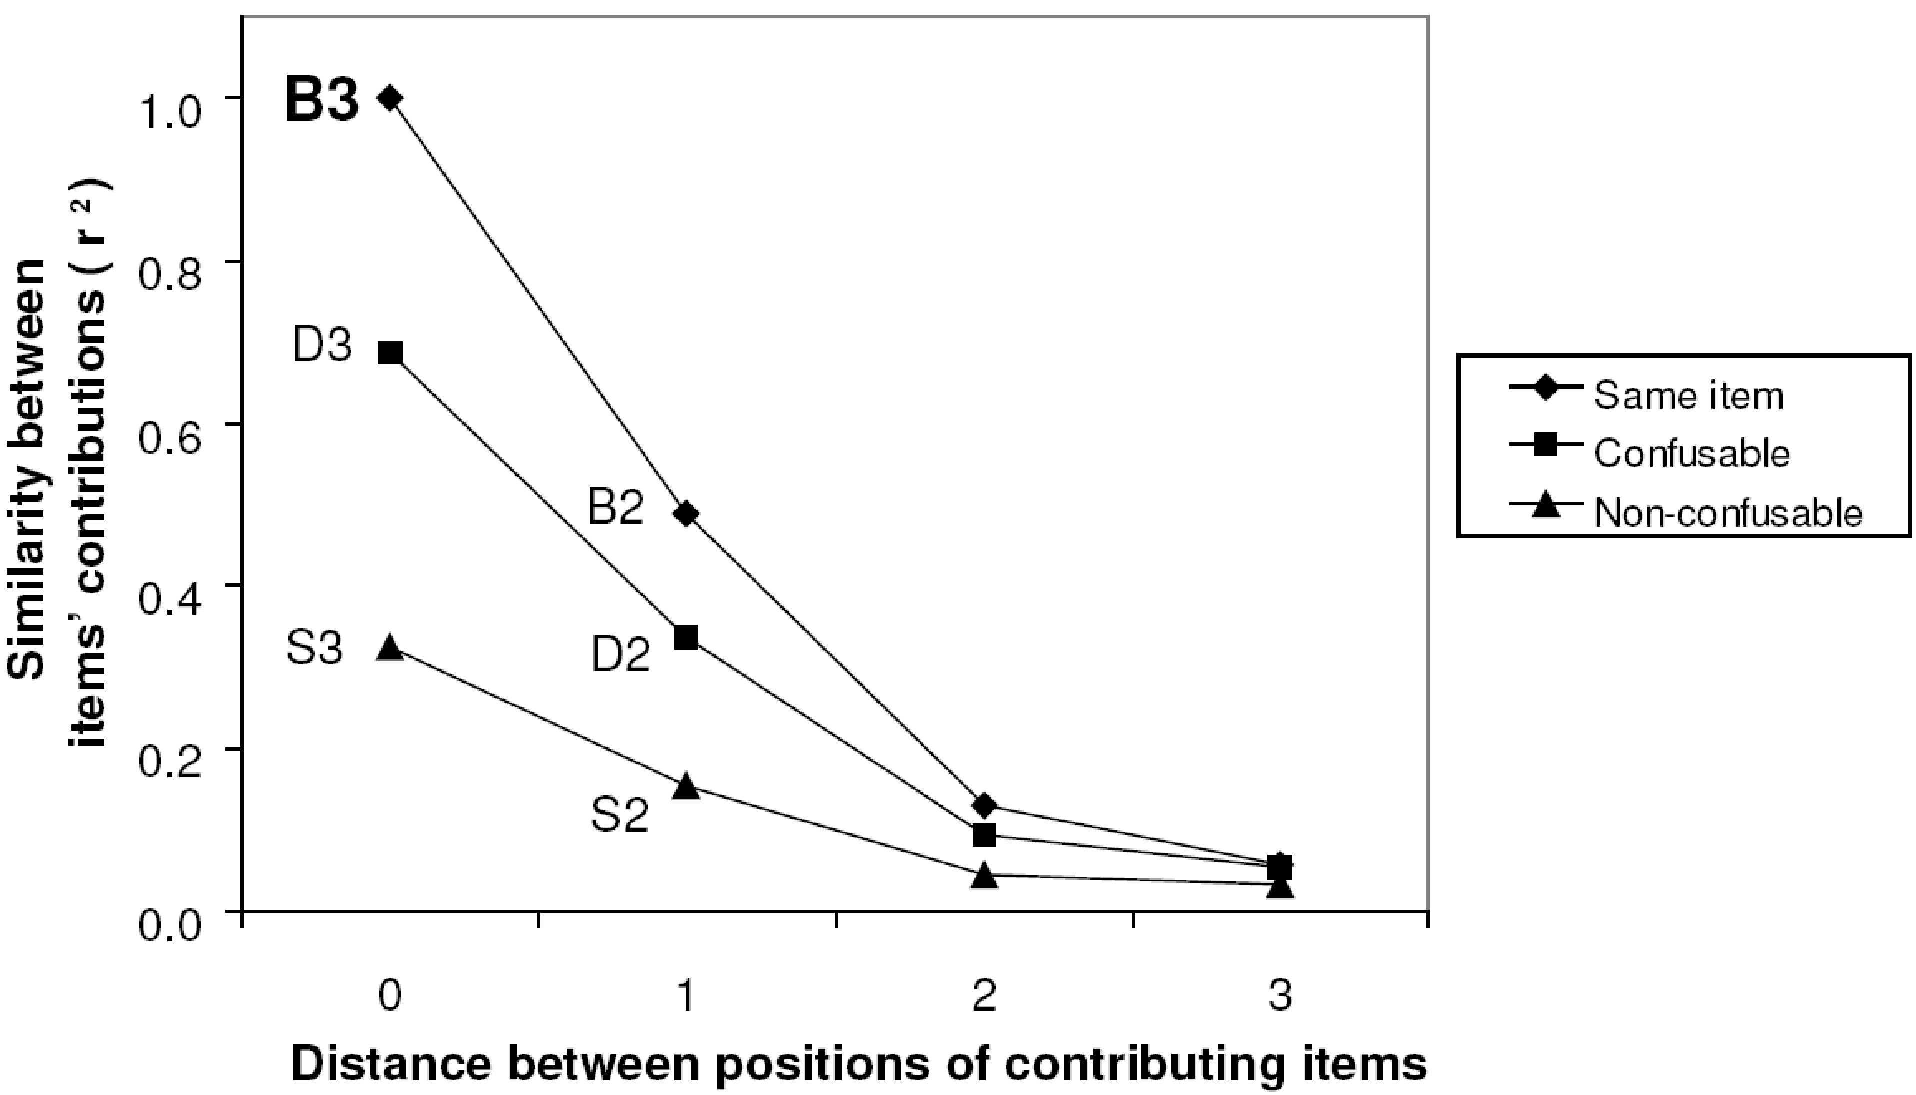

Conclusions

- Distributed recurrent networks can learn to exhibit hierarchically organized behavior without (structurally) hierarchically organized representations
  - Schemas are emergent functional properties of a system mapping perception to action
  - Sequential knowledge shaped by experience with task domains

- Sequential knowledge in recurrent networks need not *rely* on item-item associations (chaining)
  - Networks *are* sensitive to statistical structure of training environment (including item-item associations) when tested on this structure

- Explicit computational modeling can play a critical role in fully understanding the implications of theoretical claims
  - Intuitions about the computational properties of neural (and neural-like) systems can be misleading