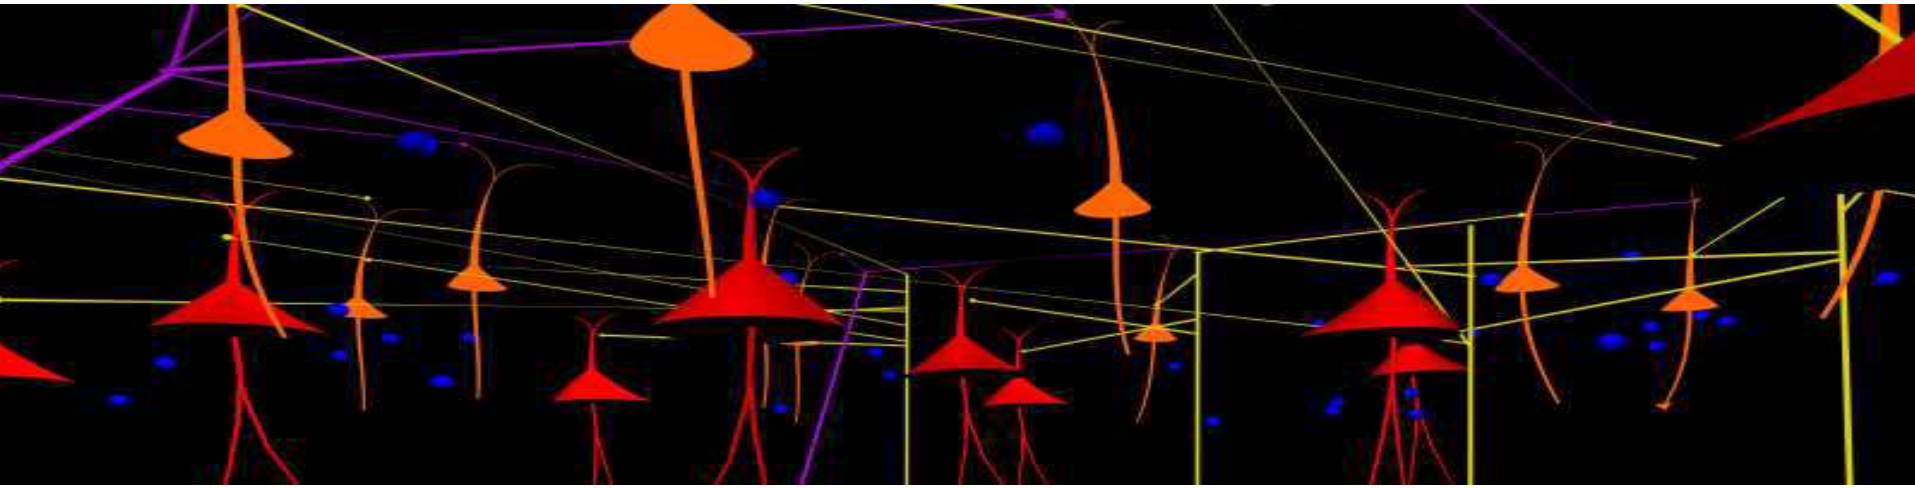

Artificial Brain Technology

Robert Hecht-Nielsen
Principal Investigator
r@hnc.com

University of California, San Diego
and
HNC Software Inc.
5935 Cornerstone Court
San Diego, CA 92121

Functions of Mammalian Cerebral Cortex and Thalamus (Thalamocortex)

- Object (perceptual, motor, conceptual, thought process, abstract, etc.) representation
- Learning and storing knowledge about objects
- Thinking
- Participation (with subcortical structures) in control of movement and thought

How Does Cerebral Cortex Represent Objects?

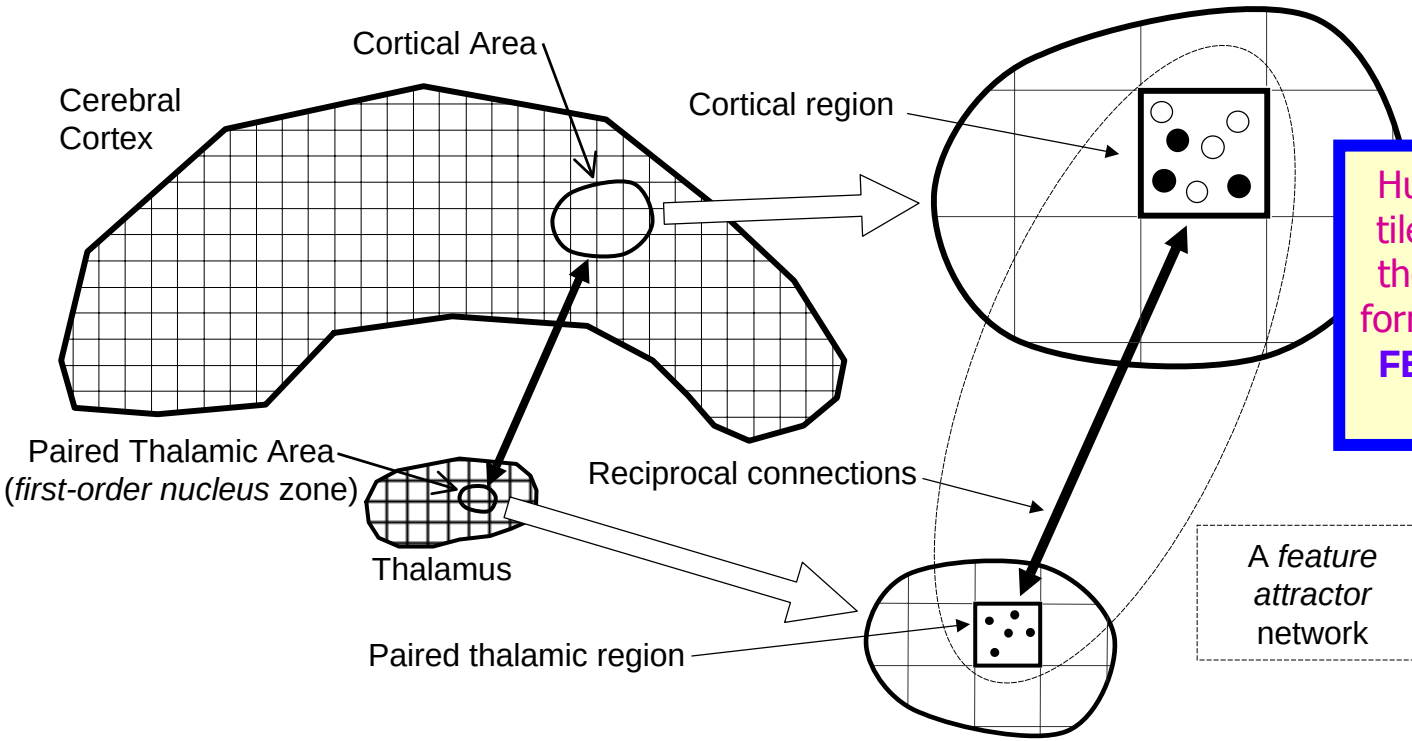

Human cerebral cortex is tiled with over a hundred thousand *regions* -- each forming the upper half of a **FEATURE ATTRACTOR** neuronal network

Each cortical region becomes equipped with a *lexicon of standardized tokens* (frozen early in learning) which are used to describe one specific *attribute* of an object

$$\{\Phi, T_1, T_2, \ldots, T_K\}$$

Neurons can randomly die and it still works.

OBJECT ATTRIBUTES ARE FUNDAMENTALLY SYMBOLIC

How Does Cerebral Cortex Represent Objects?

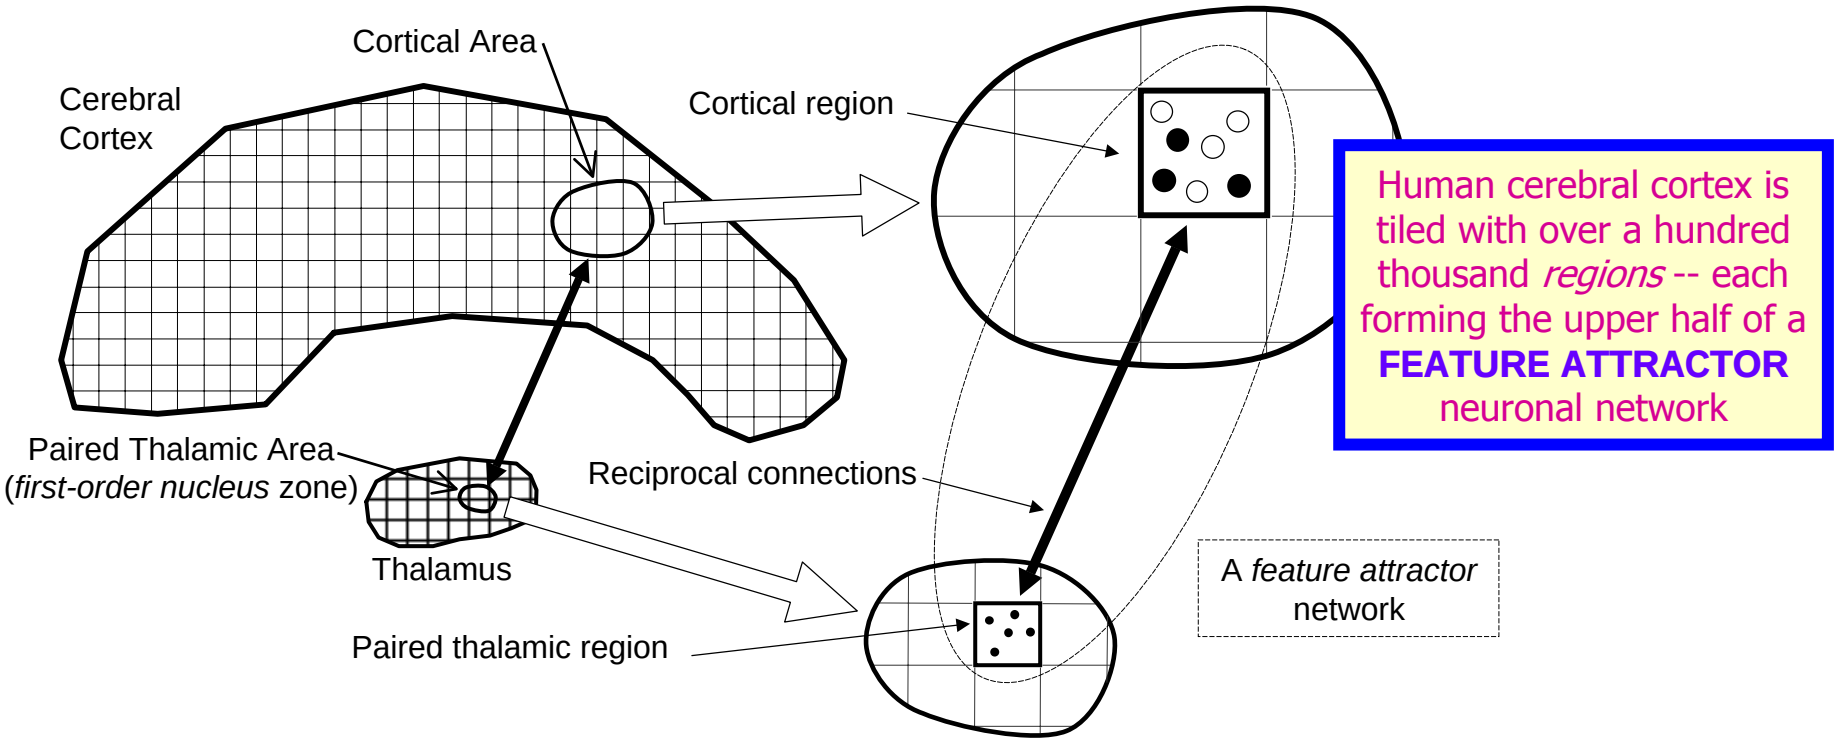

Each cortical region has a fixed *lexicon* of neuron tokens (each containing a few hundred neurons out of over 100,000). After each activation of the feature attractor network (or when external expectation input is supplied to the region in advance), the region rapidly ends up expressing **one and only one** of the lexicon tokens as its *active* output.

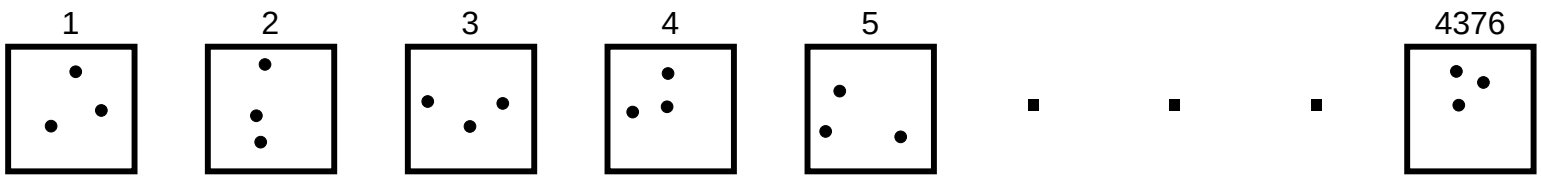

How Does Cerebral Cortex Represent Objects?

The gradual random death of its neurons has no significant effect on feature attractor performance

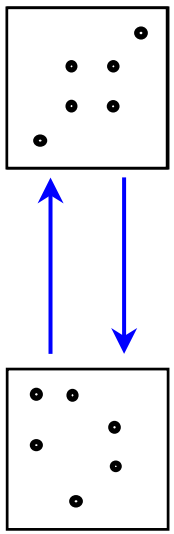

Feature Attractor Performance vs. Random Neuron Death

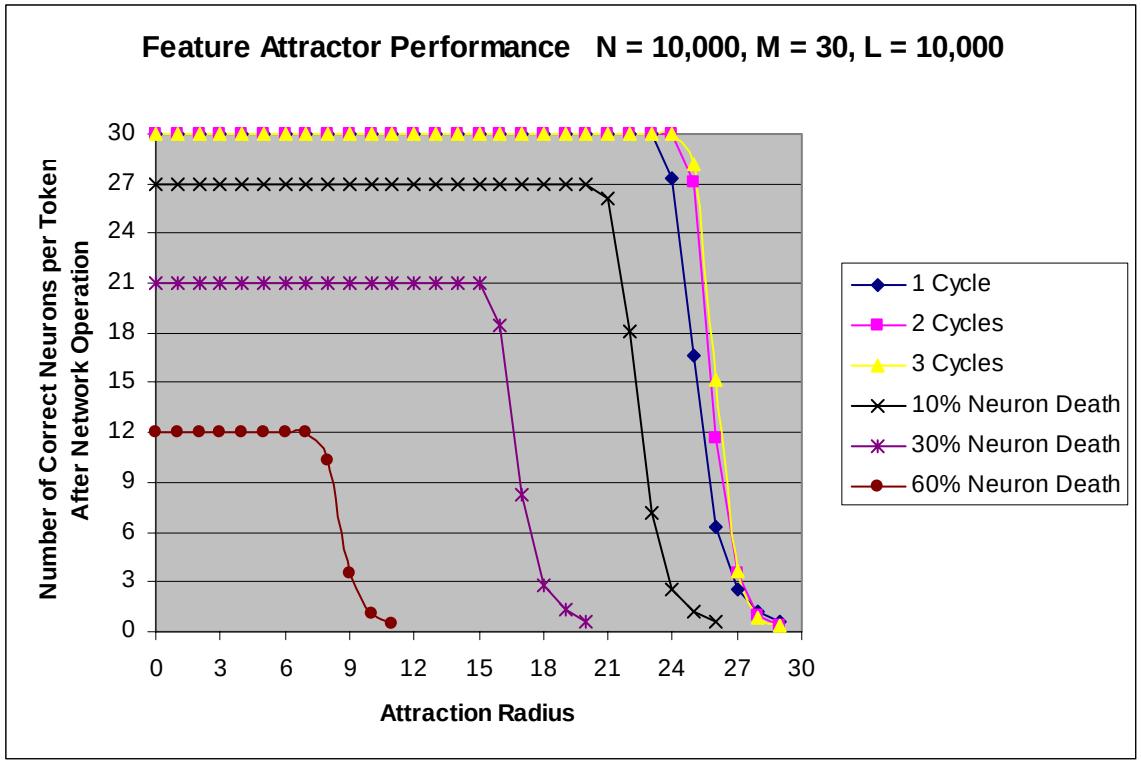

Thanks to Robert Means for this computer experiment

How Does Cerebral Cortex Represent Objects?

- Each cortical region is responsible for describing one *attribute* of an object. That description is accomplished by activating one of its lexicon tokens (or none, the *null token*, if its attribute is not applicable or not needed).

- Feature attractors can accomplish their function almost instantly if expectation is present, which significantly restricts the selection of the winning token. Otherwise, operation of the feature attractor takes about a tenth of a second.

- The 120,000 regions of the human cerebral cortex provide the mechanism for describing any object that can be considered in terms of a collection of individual attribute tokens (*attribute descriptors*). Thus, this is a *symbolic* representation.

How Does Cerebral Cortex Learn & Store Information?

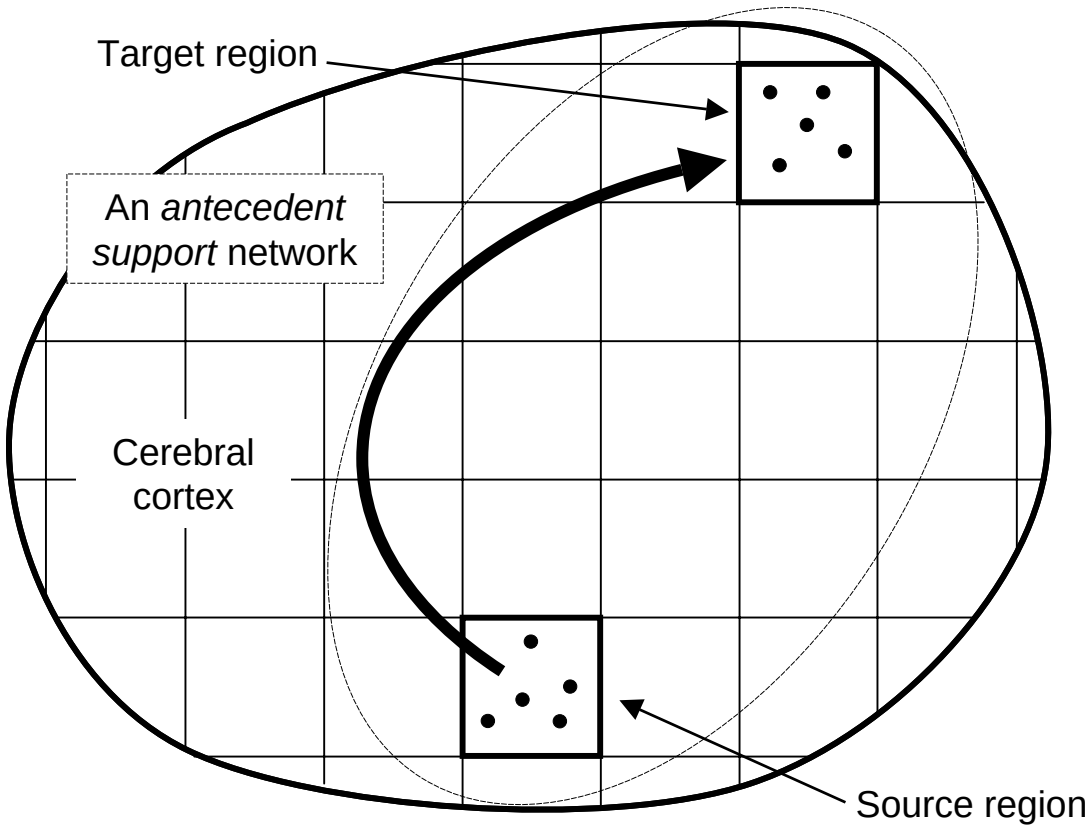

Antecedent support networks, of which there are hypothesized to be over a million in human thalamocortex, unidirectionally link each lexicon token of the network's cortical *source region* to *selected* lexicon tokens of its *target region*. The criterion for linkage selection is that the co-occurrence of source and target tokens i and j was, over some extensive set of sequential operations of the network, significantly above chance. If tokens i and j are to be linked, the weight of this linkage w(i,j) is a positive sigmoidal function of $p(i,j)/p(j) = p(i|j)$ – a quantity called the level of *antecedent support* of i, given token j.

The link strengths w(i,j) of antecedent support networks are hypothesized to constitute the only information learned and stored in mammalian thalamocortex. Each antecedent support network can be viewed as a *knowledge base* concerning unidirectional relationships between the attribute description tokens expressed on its two regions.

How Does Cerebral Cortex Learn & Store Information?

- The cerebral cortex acquires all antecedent support probabilities for token pairs which have (genetically determined) physiological links and which appear together sufficiently often over any period of N activations of the involved antecedent support network. In this sense, our knowledge is *exhaustive*.

- Once acquired, each item of knowledge is permanent (although the links very slowly weaken as a result of gradual neuron death) unless the usage of the source token changes radically, requiring the formation of many new linkages. This causes old, now disused, linkages of that source token to erode. This is how we know where the clock is in our new home and, eventually, forget where it was in our old home.

- The human cerebral cortex has a couple million antecedent support networks. Each source token is linked to, say, an average of a hundred target tokens. This implies a total human thalamocortical information storage of about 2 TBytes (= 27,892 average bits of information acquired every waking second over 30 years).

Brain Control of Brain Function

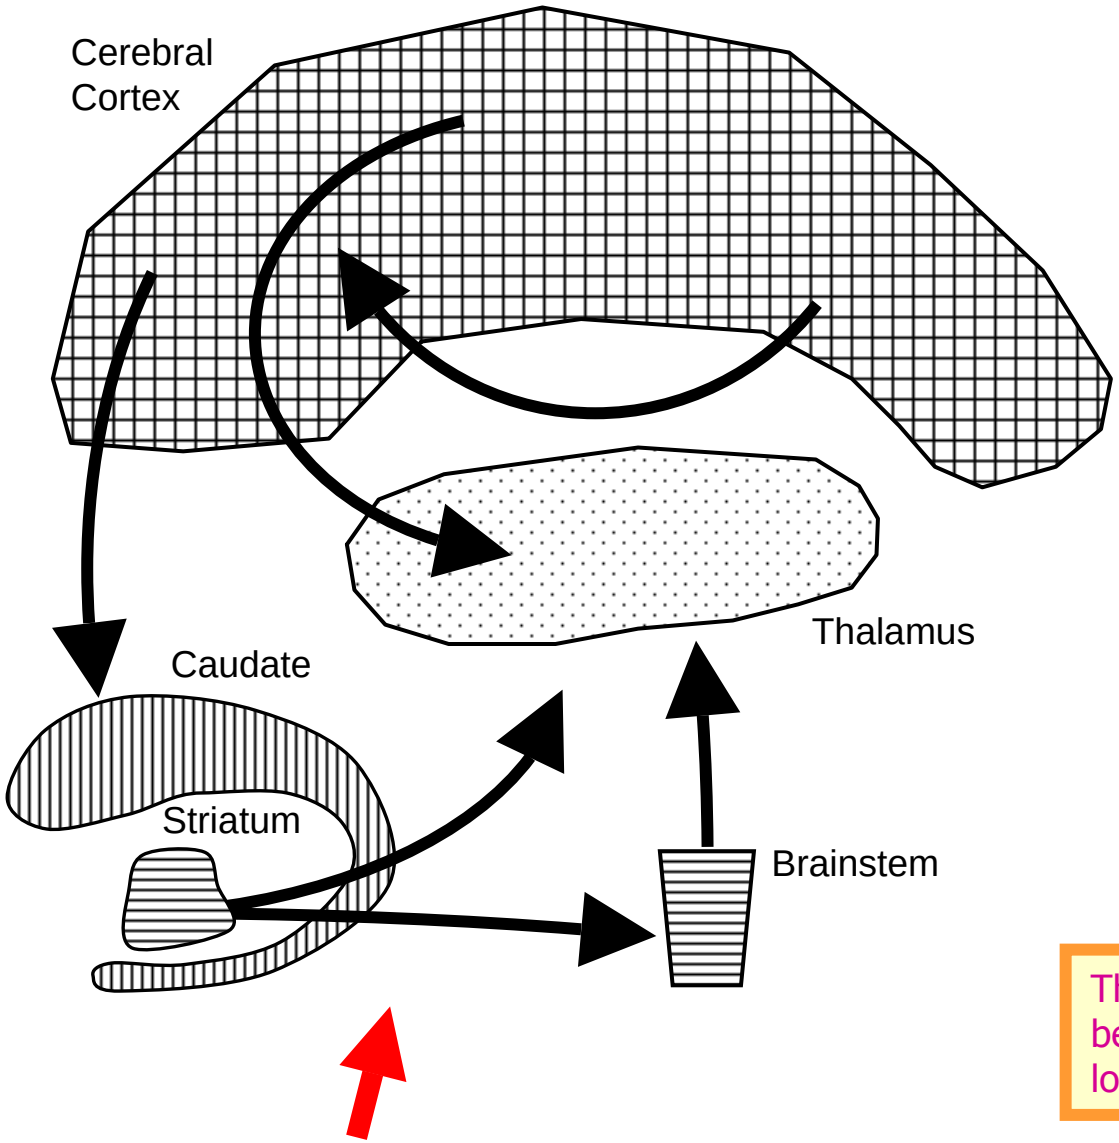

Movements and Thoughts are Siblings

The *world state* (internal and external) is constantly being expressed and delivered across cortex, causing issuance of *actions* (movements and/or thoughts) from deep cortical layers. Layer 5 projections from primary motor cortex go to spinal motor nuclei. Layer 5 projections from the rest of cortex go to subcortical nuclei which control feature attractor and antecedent support network activations (just as if they were muscles).

The brain command loop can be entrained by, and phase locked to, music and dance.

The *brain command loop* (with many parallel streams) runs continuously during wakefulness

Consensus Building: The Universal Thalamocortical Information Processing Operation

➔ **The Law of Mammalian Thought**: Choose that rational answer region token which maximally supports (in terms of p(i|j) value) the weakest item of evidence in its favor. In other words, choose that j which maximizes MIN[$p(i_1|j)$, . . . ,$p(i_k|j)$].

➔ Only answer tokens which have links from all the *assumed facts* are considered.

➔ Consensus building is hypothesized to be carried out in a fraction of a second in thalamocortex via a local neuronal dynamical behavior called *honing*.

➔ This kind of attribute co-occurrence knowledge is comprehensive (exhaustive, actually) and robust, and thus ideal for dealing with real-life situations, which are almost always novel. This is in stark contrast to Aristotelian and Bayesian rules, which are narrowly specific in their scope of application.

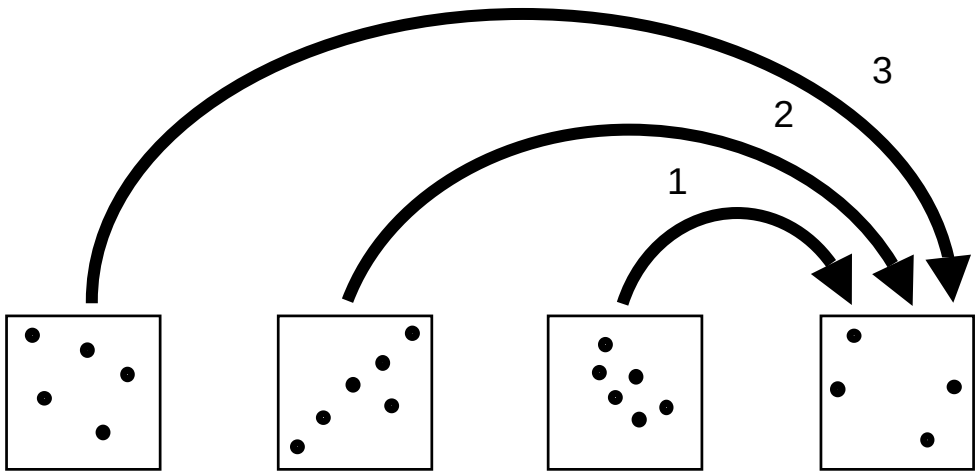

Key Attributes of the Cerebral Cortex

- Objects are described in terms of (non-unique) combinations of fixed-for-life *attribute descriptors*.
- Knowledge takes the form of pairwise rational *antecedent support* token pair linkages from one attribute descriptor to another.
- Cortical knowledge is *exhaustive* and *robust*.
- Thinking is extremely fast and yet implements a huge search which considers all rational alternatives and picks the best one.
- The 'best' answer is that one which maximally supports the weakest item of evidence in its favor. This is the ***Law of Mammalian Thought***.

Implementation of Consensus Building: *Honing*

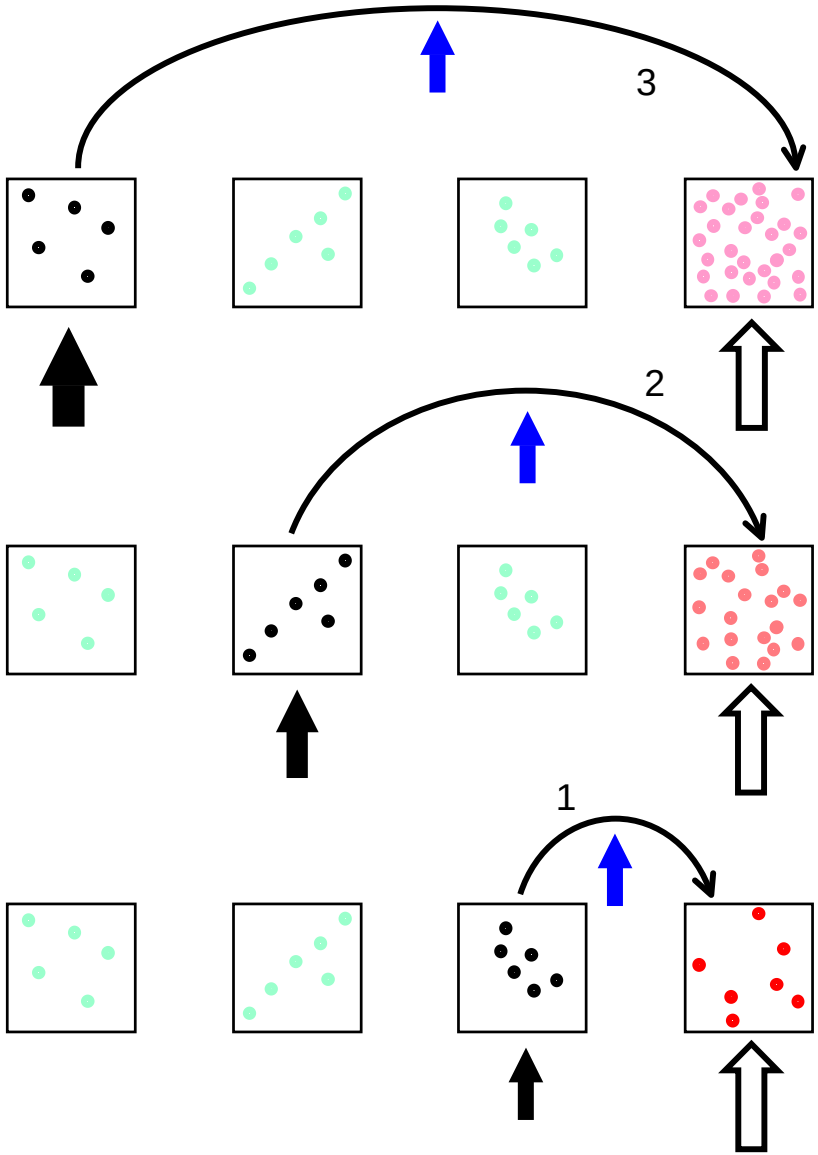

Consensus building involves *sequential activation* of groups of regions in accordance with a specified *precedence ordering*. At each stage of this sequence, each answer region is *honed*. Honing involves turning off (for the remainder of the CB operation) those neurons which do not receive significance input above a fixed minimum. Neurons which remain on have their activation level set to their new significance input value; except that this level can never increase; only decrease, during the CB process. Additional answers can be discerned by turning off answer neurons and repeating the CB process. The entire CB process typically requires only a few hundred milliseconds.

CB is a stored *thought process* (region and fascicle activation sequence) which is data independent. Such thought processes are learned, stored, and issued by cortex exactly like movements.

Consensus Building: Combining Knowledge

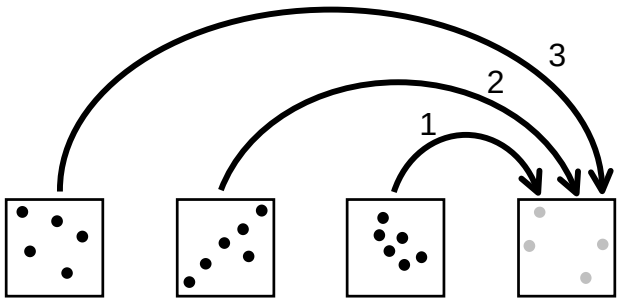

Knowledge Base 3: **driving:** influence, expired, suspended

Knowledge Base 2: **at:** university, time, A.M., P.M.

Knowledge Base 1: **top:** priority, aides, aid, executives

legend

red = phrase not in training database

blue = phrase in training database

Assumed Facts			Consensus Answers (from successive CB operations)
driving	at	top	speed
man	on	the	moon
all	of	a	sudden, common, handful
all	but	a	handful, sudden, surprising
all	over	the	map, country, place
the	future	looks	bright, promising
down	the	garden	path, mall

Computer experiment using knowledge bases established using a billion word English text corpus. A 5,000 word lexicon is employed by each region.

Consensus Building is the Holy Grail of AI.

Consensus Building: Combining Disparate Knowledge

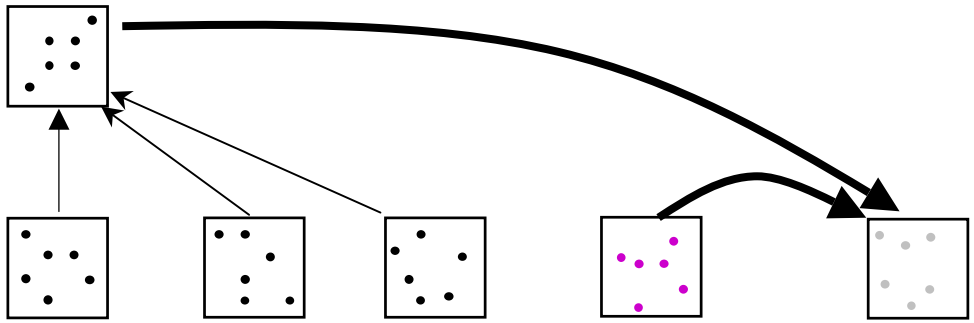

Consensus building allows context of arbitrary type and quantity to be brought to bear on segmentations, classifications, predictions, etc.

the	wishes	of	the	school	board
spoke	on	condition		of	anonymity
##XXX	degrees			below	zero
former	Prime	Minister		Margaret	Thatcher
former	President			Ronald	Reagan
according	to	the		school	administrators, records, record
the	Dow	Jones		Industrial	Average
the	Dow	Jones		Average	rose, fell, climbed
the	destruction	of		wildlife	habitat
##XXX	on	the		Richter	scale
living	in			New	York, Jersey, Delhi, Orleans
comes	in			early	afternoon, trading, September, Spring
based	on			new	evidence
according	to			new	laws, rules
earlier	in			the	day, week, evening, session
New	York			Stock	Exchange, market, index

Example illustrating a *derived* network architecture (a *hierarchical abstractor*).

Yes, this is really it. CB is the mechanism of all human thought.

Learning Word Level Synonymy

… bus drove from Chicago to …

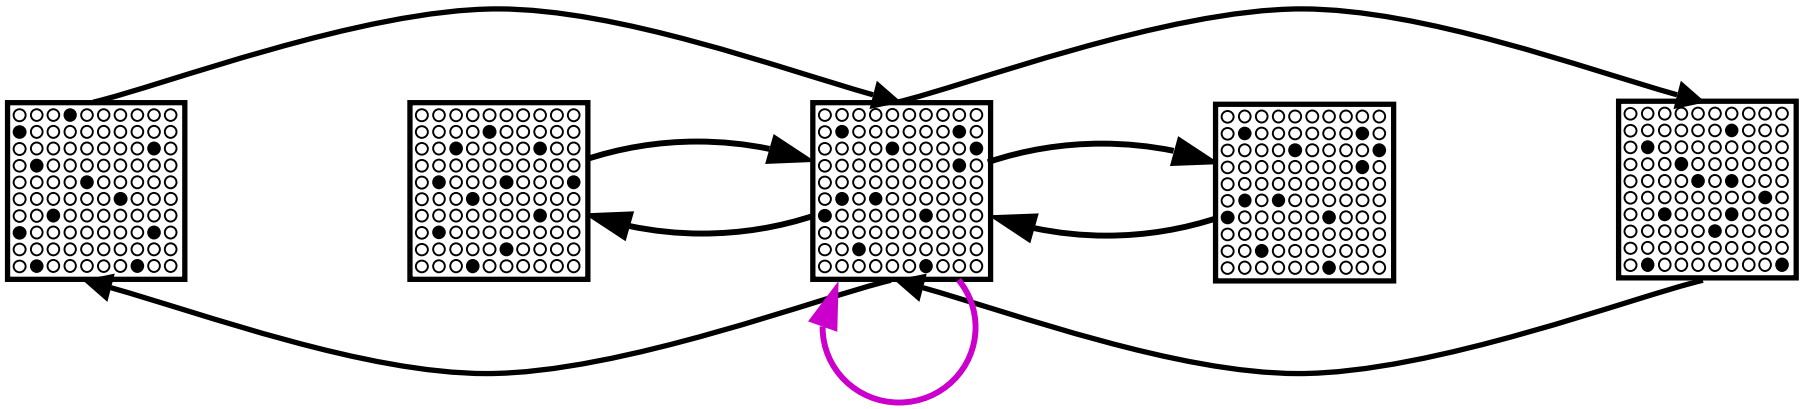

car
auto
automobile
truck

Word Synonymy Lists

John	Mary	Detroit	France
Jack	Betty	Chicago	Portugal
Ivan	Elisabeth	Sydney	China
Joseph	Sue	Seattle	Singapore
George	Sally	Phoenix	Haiti
Paul	Karen	Manila	Cuba
Peter	Jackie	Cincinnati	Spain
Daniel	Margaret	Minneapolis	Chile
Edward	Lisa	Vancouver	Japan

phone	ninety	minutes	TV	television
telephone	twenty	years	television	TV
telecommunications	thirty	days	radio	radio
communications	eighty	pounds	video	news

see	car	college	automobile	seen	president
hear	vehicle	school	automaker	got	presidents
be	truck	university	automotive	done	leader
get	jeep	graduate	auto	been	chairman

Phrase Level Synonymy Lists

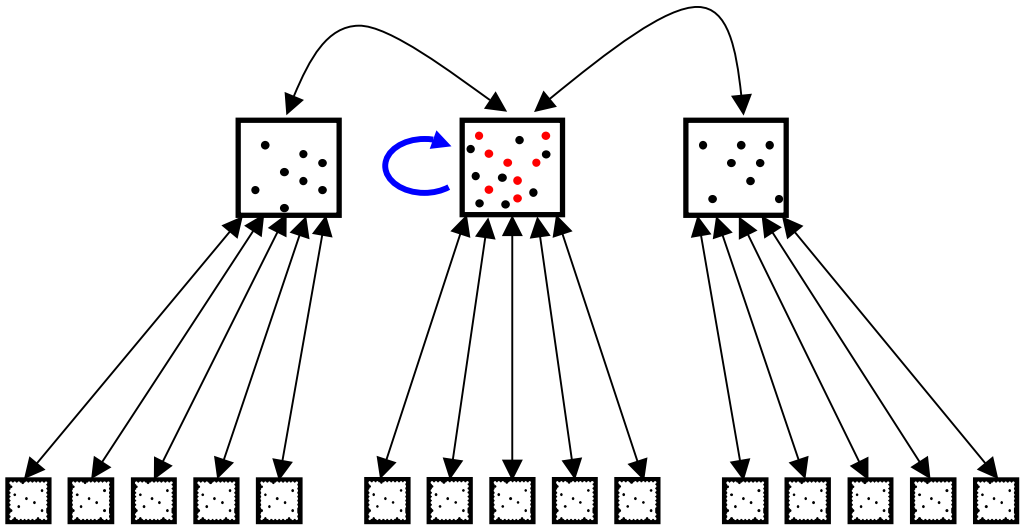

French Francs	will be	New York
German Marks	should be	Los Angeles
Italian Lira	have been	Seattle
Swiss Francs	are	New Orleans
Canadian Dollars	can be	Chicago

Antecedent Support Network Connectivity Statistics

(typical: preword-postword)

Number of Words in the Lexicon	Number of Possible Links	Number of Actual Links	Percentage Connectivity
5,000	25,000,000	6,109,422	24.4%
10,000	100,000,000	11,438,063	11.4%
30,000	900,000,000	21,724,904	2.4%

Other Simple Thought Process Examples

- Antecedent support knowledge bases contain both explicit and implicit knowledge about *word functionality* (semantics, grammar, etc.).
- For example, consider the following knowledge bases and the sentence template:

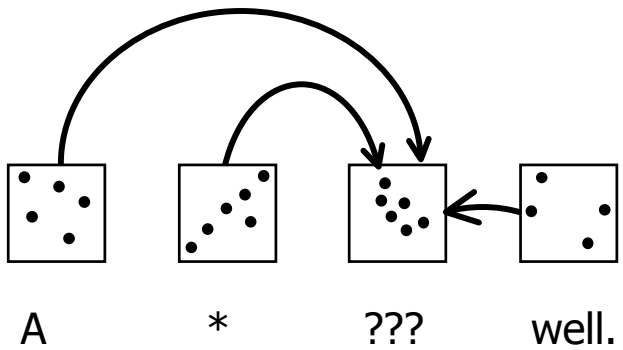

A * ??? well.

- **President:** speaks, travels, understands, defended
- **Quarterback:** throws, recruit, draw, threw
- **Car:** coat, ran, speeds, stops, handles
- **Model:** wears, sells, oh, fits, dressed
- **Person:** executed, ate, understands, feels, knows
- **Teacher:** taught, training, travels, pay, speaks

A * can ???.

- **Person:** swim, die, fly, sail, commit
- **Dog:** bite, hunt, duck, pig, eat
- **Fish:** swim, bite, feed, fish, hunt
- **President:** veto, appoint, consult, propose, testify
- **Jury:** decide, hear, spare, reasonably, consider
- **Judge** dismiss, appoint, decide, declare, impose

Word *Families:* The Emergence of Abstraction

first	football	food	talks	food	dress	coach
second	golf	meat	relations	retail	shirt	player
third	basketball	beef	negotiations	clothing	jacket	forward
fourth	soccer	milk	relationship	specialty	uniforms	owner
fifth	tennis	pork	meetings	toy	dresses	captain
sixth	hockey	dairy	agreements	merchandise	t-shirts	quarterback
seventh	athletic		links	supermarket	shirts	rookie
eighth	boxing		conversation	apparel	jackets	pitcher
ninth	rugby			jewelry	costumes	starter
				grocery	t-shirt	teammates
				liquor		defender
						trainer
						teammate
						linebacker
						goalie
						defenseman

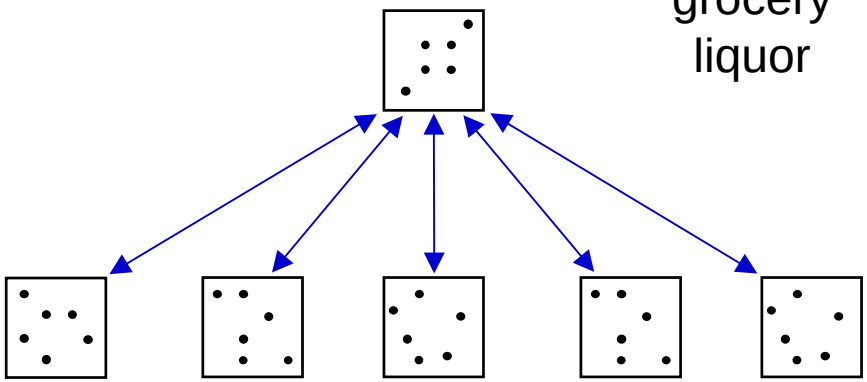

A simple automated process produced 18,339 families for a lexicon of 10,000 words. Each family is a subset of the *synonymy set* of that word. Families are like word senses, but more abstract and more useful. For example, word family matching between sentences can be used to evaluate their similarity of meaning

Families for Sentence Meaning Comparison

1: ***They*** *brought a* ***lawsuit*** *against* ***Firestone****.*

they	lawsuit	Firestone
we	case	Michelin
you	litigation	BF Goodrich
we'll	suit	Goodyear
they've	lawsuits	Bridgestone
you've	charges	
they'll	complaint	
he'll	cases	
they'd	complaints	
	suits	
	claims	

2: ***We*** *filed a* ***case*** *charging* ***Michelin*** *of wrongdoing.*

This example illustrates the vast context-exploitation power of this new approach to information processing

Families for Sentence Meaning Comparison

1: ***The safety personnel ran out to the shuttle immediately after touchdown.***

safety
medical
legal
accounting
scientific
security
regulatory
investigative
…

safety
quarterback
linebacker
defenseman
goalie
owner
coach
starter
forward
rookie
…

Similarly, 'touchdown' can be disambiguated into "landing" and "goal", respectively, leaving wholly distinct sentence representations (and therefore small similarity measure).

2: ***The team's safety scored a touchdown.***

Test Sentence Pairs

- No common words, similar meaning
  - **we filed a suit charging microsoft with illegal behavior**
  - **they submitted a case accusing apple of unauthorized conduct**
    - SIMILARITY = 0.995
- Some common words, dissimilar meaning
  - **i was right to avoid a suit against apple**
  - **on my right was a man in a suit drinking apple juice**
    - SIMILARITY = 0.171

Context-Driven Family Disambiguation

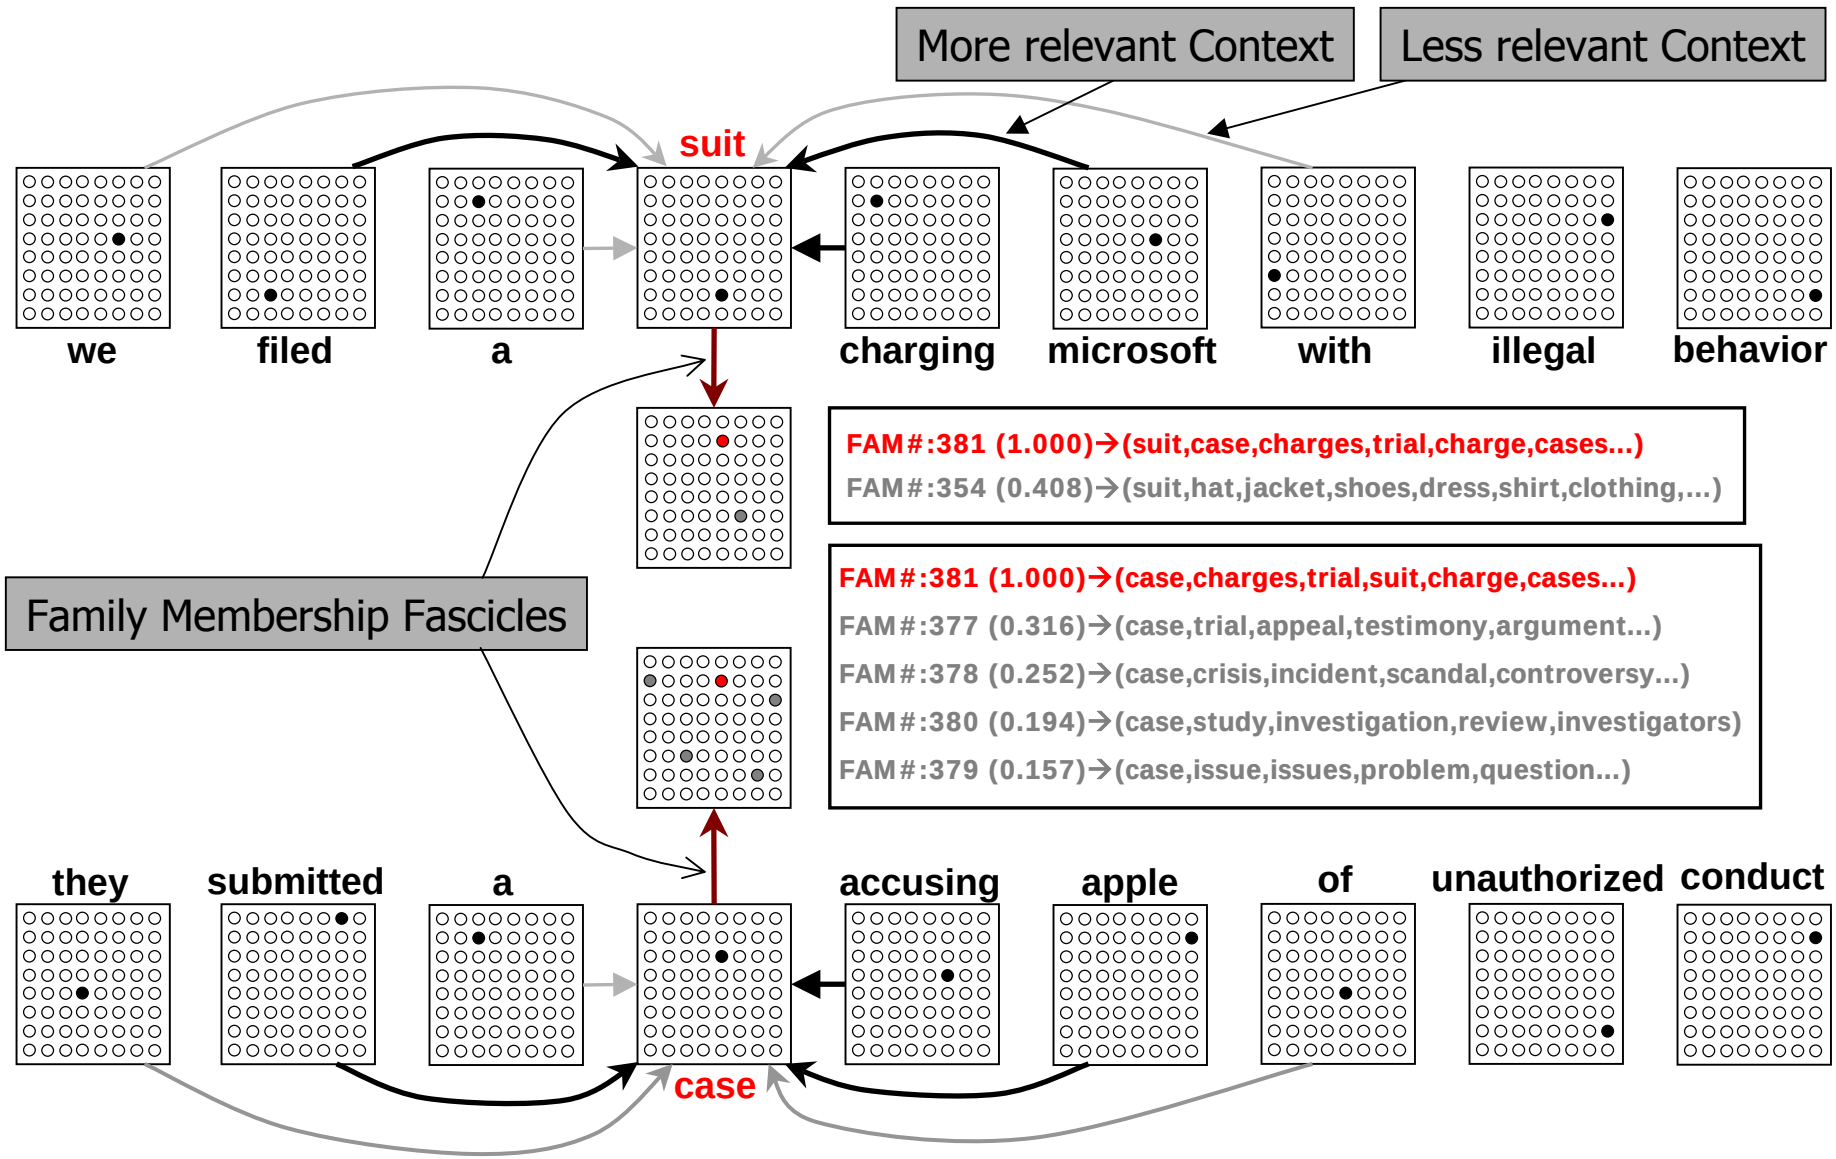

Context Driven Family Disambiguation

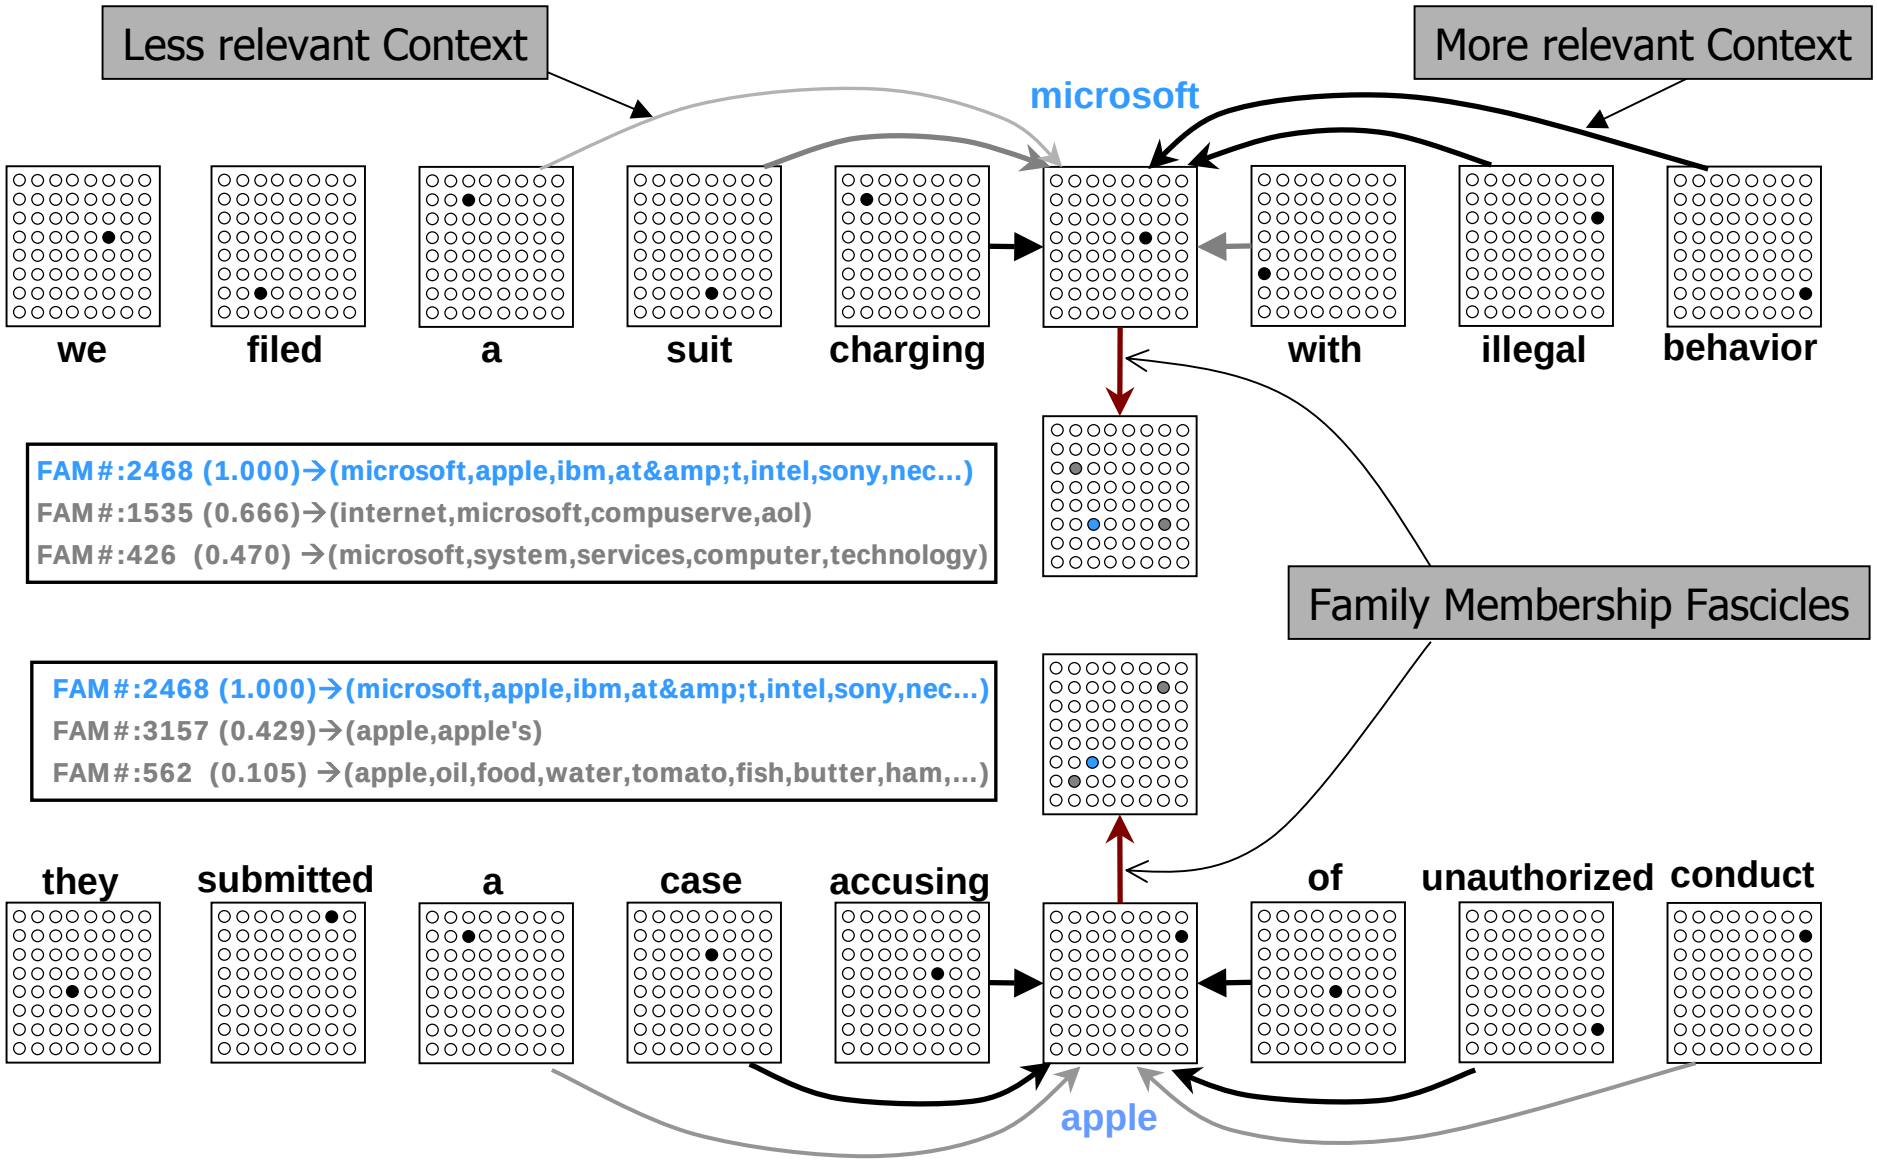

Education

- *Training* is the process of exposing a brain to a sequence of selected examples in an information environment.
- *Education* is where a brain is exposed to carefully selected examples with side information that indicates acceptable variants that can be used with each construct.
- Training depends upon the raw environmental frequency of exposure to decide what to learn.
- Education demands that everything which is presented is learned and ensures that obvious variants (not superceded by explicitly learned exceptions) will be handled similarly.

Reference

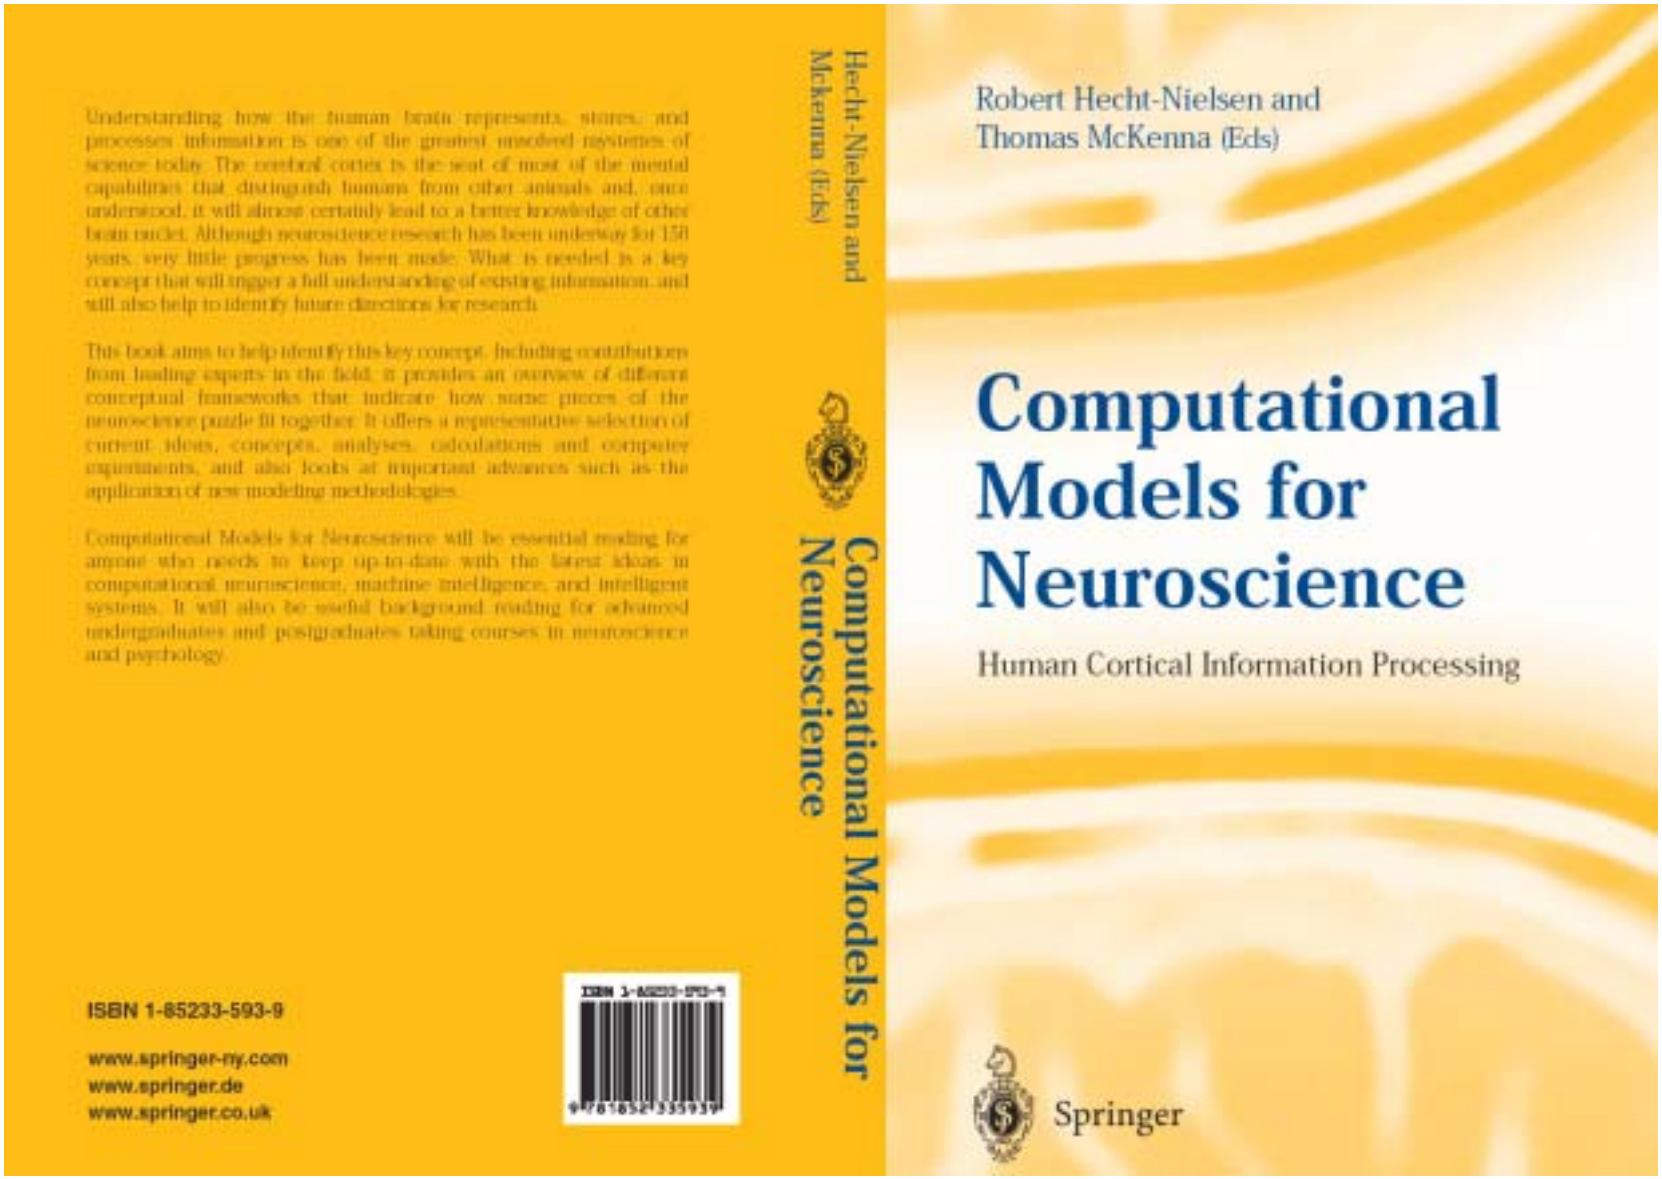

Understanding how the human brain represents, stores, and processes information is one of the greatest unsolved mysteries of science today. The cerebral cortex is the seat of most of the mental capabilities that distinguish humans from other animals and, once understood, it will almost certainly lead to a better knowledge of other brain nuclei. Although neuroscience research has been underway for 150 years, very little progress has been made. What is needed is a key concept that will trigger a full understanding of existing information and still also help to identify future directions for research.

This book aims to help identify this key concept. Including contributions from leading experts in the field, it provides an overview of different conceptual frameworks that indicate how some pieces of the neuroscience puzzle fit together. It offers a representative selection of current ideas, concepts, analyses, calculations and computer experiments, and also looks at important advances such as the application of new modeling methodologies.

Computational Models for Neuroscience will be essential reading for anyone who needs to keep up-to-date with the latest ideas in computational neuroscience, machine intelligence, and intelligent systems. It will also be useful background reading for advanced undergraduates and postgraduates taking courses in neuroscience and psychology.

ISBN 1-85233-593-9

www.springer-ny.com
www.springer.de
www.springer.co.uk

Hecht-Nielsen and Mckenna (Eds)

Computational Models for Neuroscience

Robert Hecht-Nielsen and
Thomas McKenna (Eds)

Computational Models for Neuroscience

Human Cortical Information Processing

Springer

Welcome to Our Laboratory

BRAIN 1
24 Pentium Processors
48 GBytes RAM
1 Gbit/s Optical Ethernet

HUMAN THALAMOCORTEX
25 Billion 1kHz Neurons
2 TBytes Information Storage

BRAIN 2
30 1GHz Pentium Processors
60 GBytes RAM
1 Gbit/s Optical Ethernet

The HNC Team

Other Team Members

Todd Gutschow - Project Leader and Senior Project Scientist, HNC Cortronics Project

Dr. Robert Means - Chief Technologist, HNC Cortronics Project

Kate Mark - Project Coordinator, HNC Cortronics Project

David Busby - Chief Brain Software Architect, HNC Cortronics Project

Dr. Syrus Nemat-Nasser - Scientist, HNC Cortronics Project

Rion Snow - Researcher, HNC Cortronics Project

Dr. Shailesh Kumar, Scientist, HNC Cortronics Project

Research Sponsors

HNC
ARDA
ONR

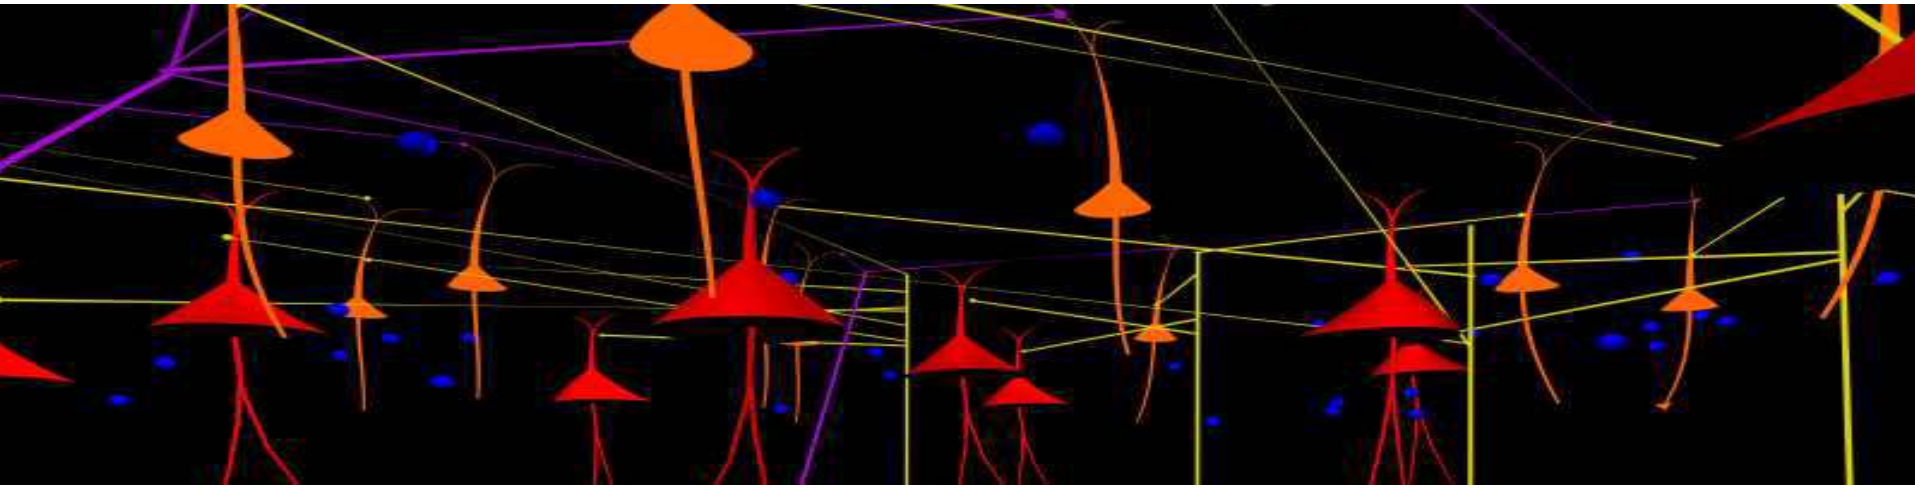

Artificial Brain Technology

Robert Hecht-Nielsen
Principal Investigator
r@hnc.com

University of California, San Diego
and
HNC Software Inc.
5935 Cornerstone Court
San Diego, CA 92121

Functions of Mammalian Cerebral Cortex and Thalamus (Thalamocortex)

- Object (perceptual, motor, conceptual, thought process, abstract, etc.) representation
- Learning and storing knowledge about objects
- Thinking
- Participation (with subcortical structures) in control of movement and thought

How Does Cerebral Cortex Represent Objects?

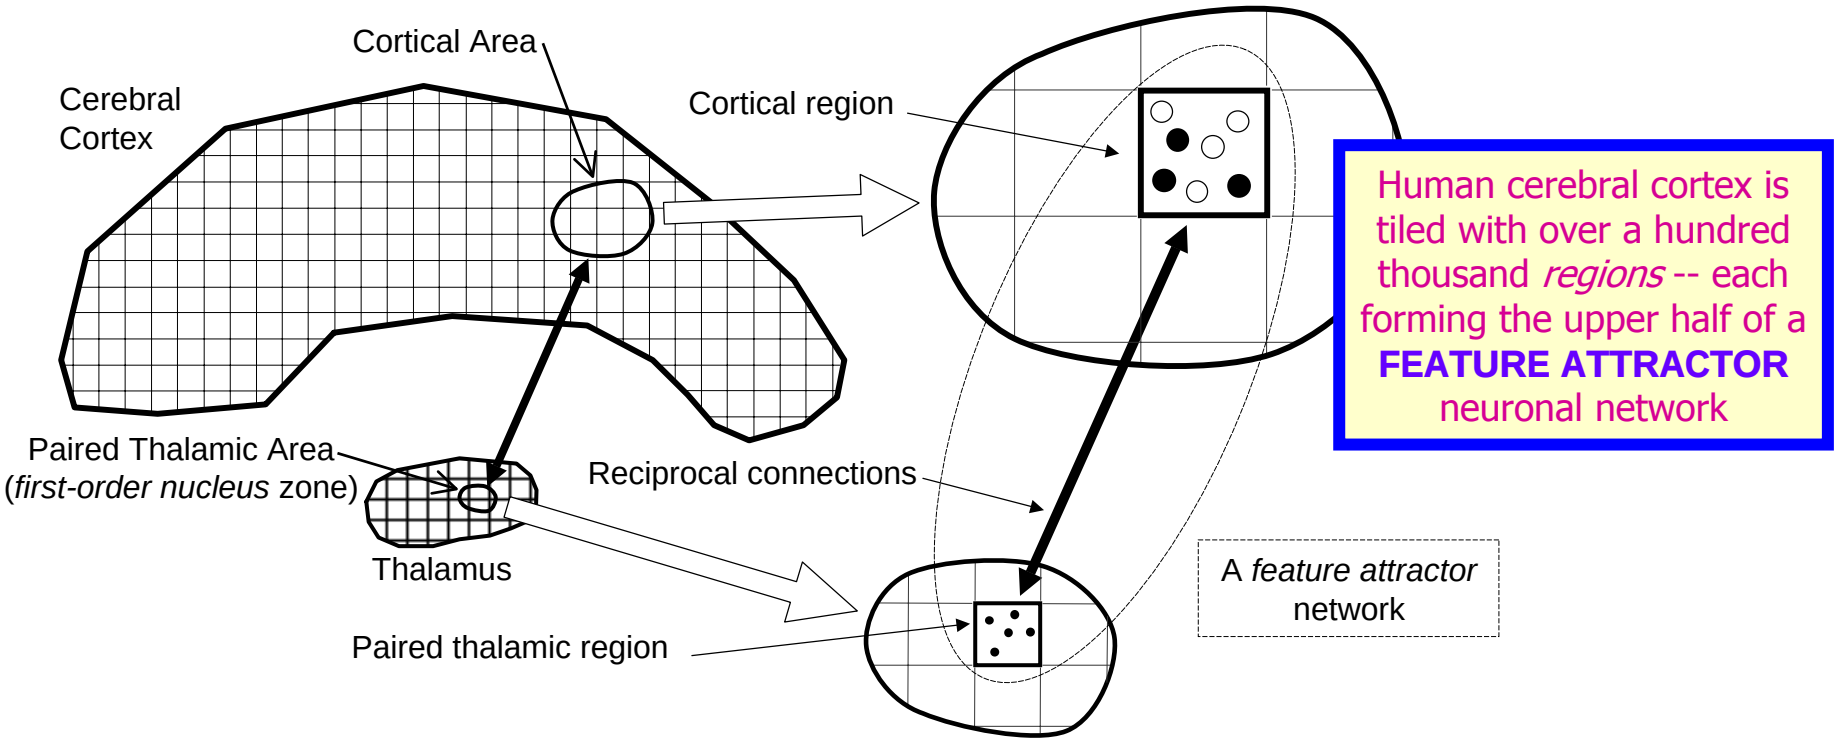

Each cortical region has a fixed *lexicon* of neuron tokens (each containing a few hundred neurons out of over 100,000). After each activation of the feature attractor network (or when external expectation input is supplied to the region in advance), the region rapidly ends up expressing **one and only one** of the lexicon tokens as its *active* output.

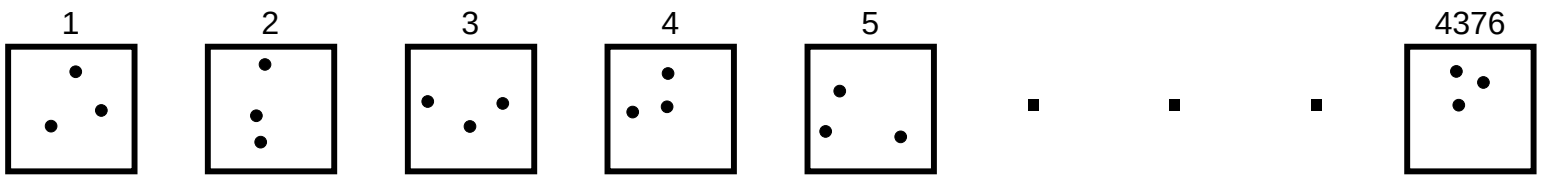

How Does Cerebral Cortex Represent Objects?

The gradual random death of its neurons has no significant effect on feature attractor performance

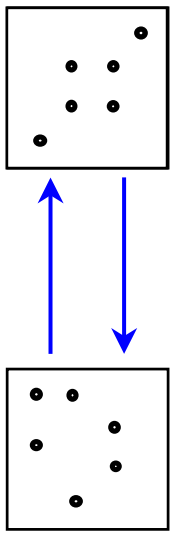

Feature Attractor Performance vs. Random Neuron Death

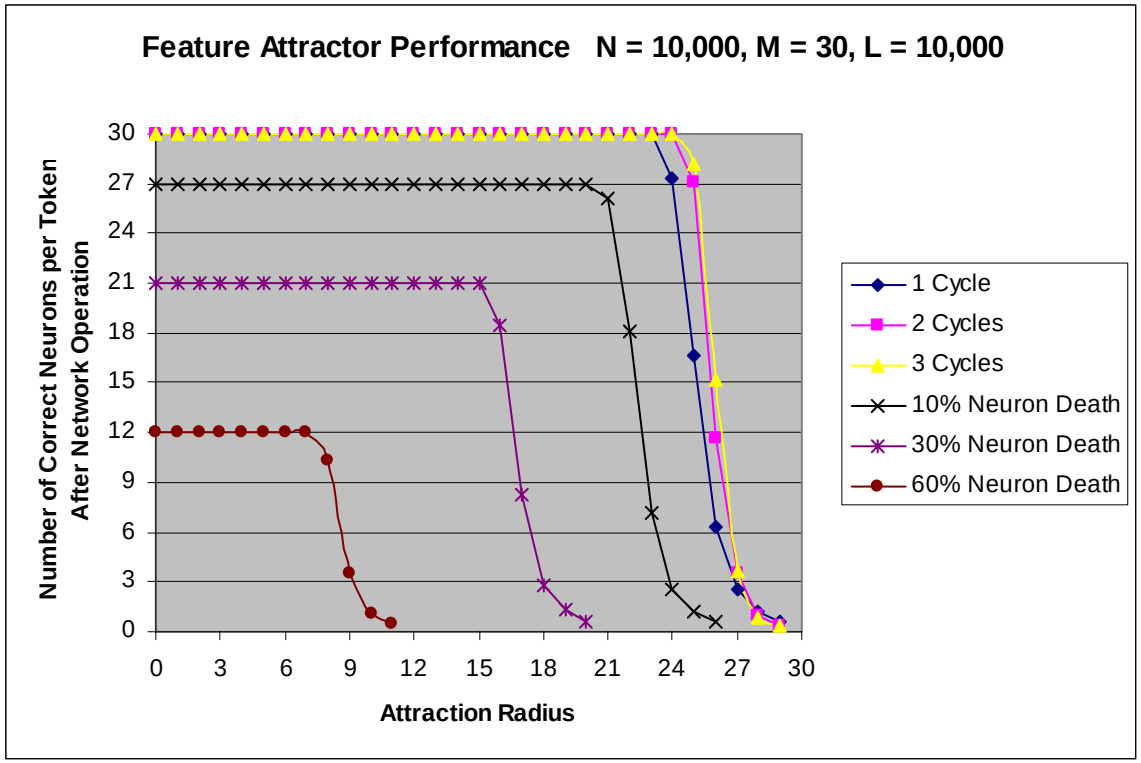

Thanks to Robert Means for this computer experiment

How Does Cerebral Cortex Learn & Store Information?

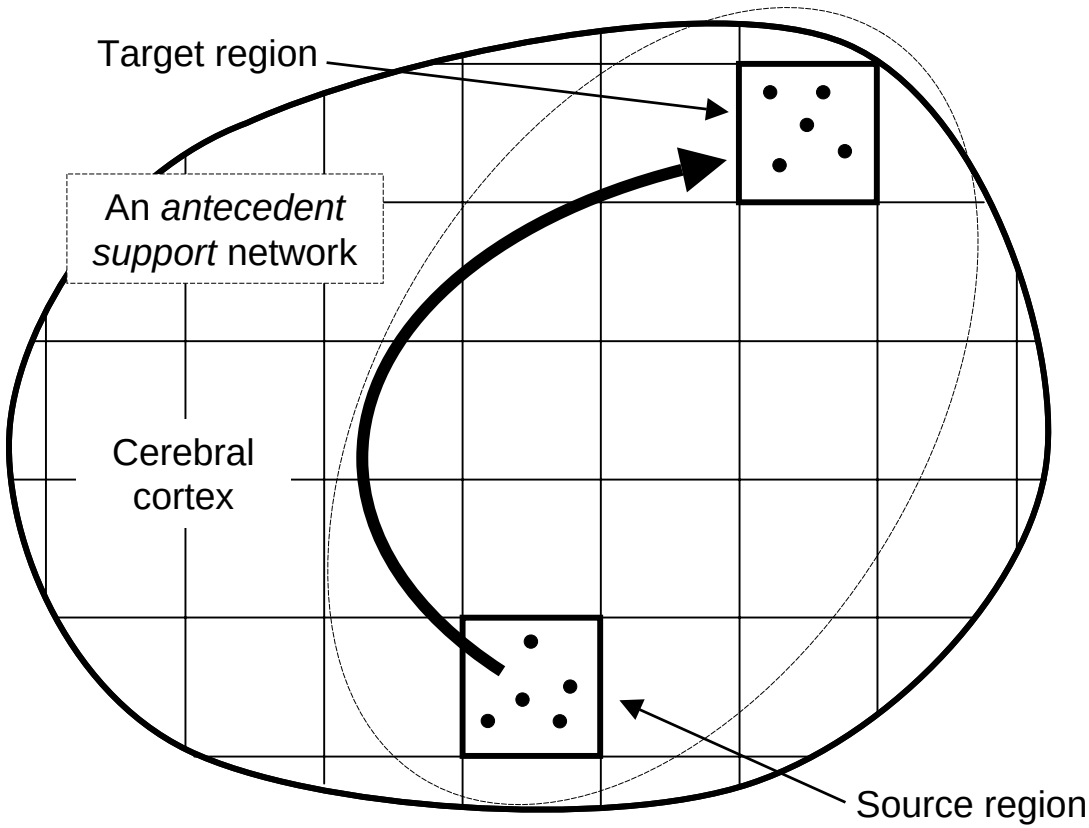

Antecedent support networks, of which there are hypothesized to be over a million in human thalamocortex, unidirectionally link each lexicon token of the network's cortical *source region* to *selected* lexicon tokens of its *target region*. The criterion for linkage selection is that the co-occurrence of source and target tokens i and j was, over some extensive set of sequential operations of the network, significantly above chance. If tokens i and j are to be linked, the weight of this linkage w(i,j) is a positive sigmoidal function of $p(i,j)/p(j) = p(i|j)$ – a quantity called the level of *antecedent support* of i, given token j.

The link strengths w(i,j) of antecedent support networks are hypothesized to constitute the only information learned and stored in mammalian thalamocortex. Each antecedent support network can be viewed as a *knowledge base* concerning unidirectional relationships between the attribute description tokens expressed on its two regions.

Brain Control of Brain Function

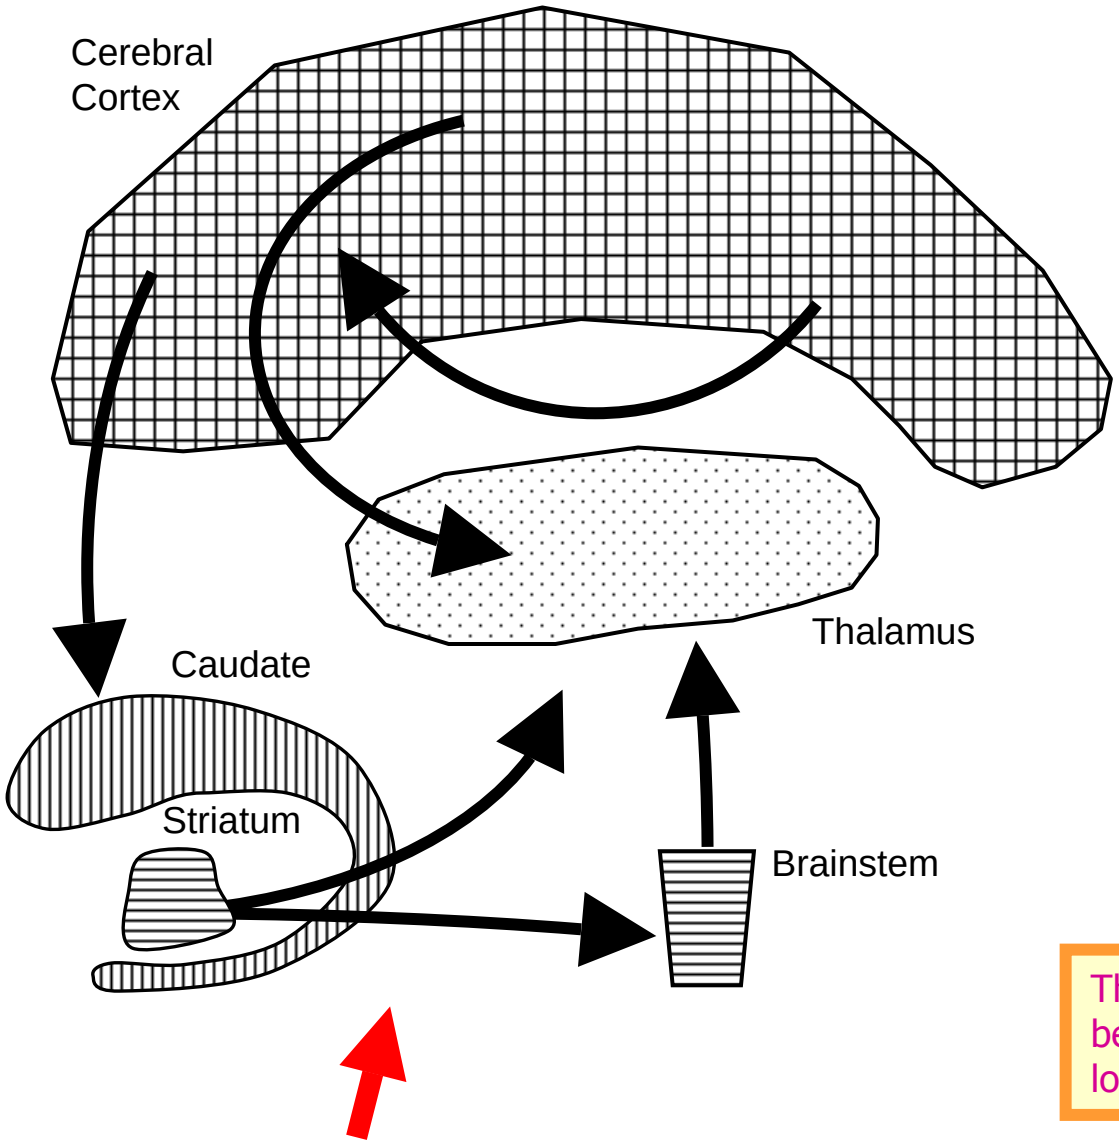

Movements and Thoughts are Siblings

The *world state* (internal and external) is constantly being expressed and delivered across cortex, causing issuance of *actions* (movements and/or thoughts) from deep cortical layers. Layer 5 projections from primary motor cortex go to spinal motor nuclei. Layer 5 projections from the rest of cortex go to subcortical nuclei which control feature attractor and antecedent support network activations (just as if they were muscles).

The brain command loop can be entrained by, and phase locked to, music and dance.

The *brain command loop* (with many parallel streams) runs continuously during wakefulness

How Thinking Works

➜ **The Law of Mammalian Thought**: Choose that rational answer region token which maximally supports (in terms of p(i|j) value) the weakest item of evidence in its favor. In other words, choose that j which maximizes MIN[$p(i_1|j)$, . . . ,$p(i_k|j)$].

➜ Only answer tokens which have links from all the *assumed facts* are considered.

➜ Consensus building is hypothesized to be carried out in a fraction of a second in thalamocortex via a local neuronal dynamical behavior called *honing*.

➜ This kind of attribute co-occurrence knowledge is comprehensive (*exhaustive*, actually) and *robust*, and thus ideal for dealing with real-life situations, which are almost always novel. This is in stark contrast to Aristotelian and Bayesian rules, which are narrowly specific in their scope of application.

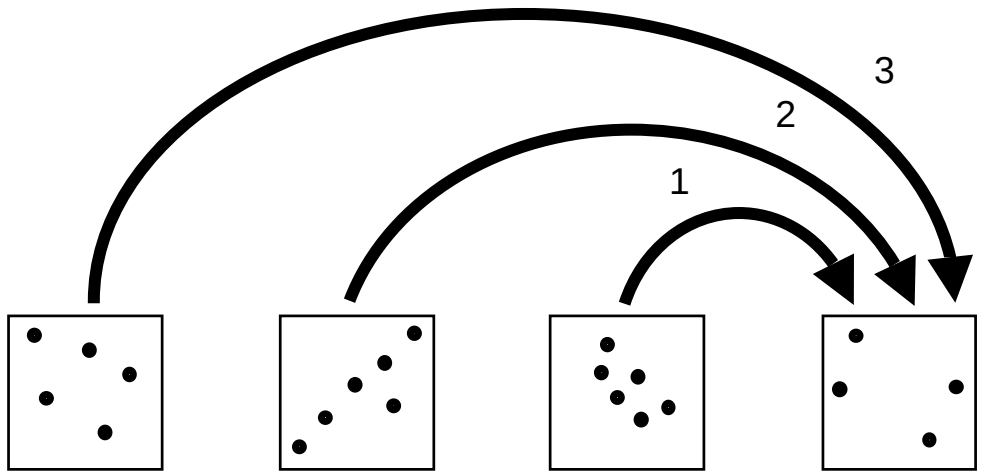

Key Attributes of the Cerebral Cortex

- Objects are described in terms of (non-unique) combinations of fixed-for-life *attribute descriptors*.
- Knowledge takes the form of pairwise rational *antecedent support* token pair linkages from one attribute descriptor to another.
- Cortical knowledge is *exhaustive* and *robust*.
- Thinking is extremely fast and yet implements a huge search which considers all rational alternatives and picks the best one.
- The 'best' answer is that one which maximally supports the weakest item of evidence in its favor. This is the ***Law of Mammalian Thought***.

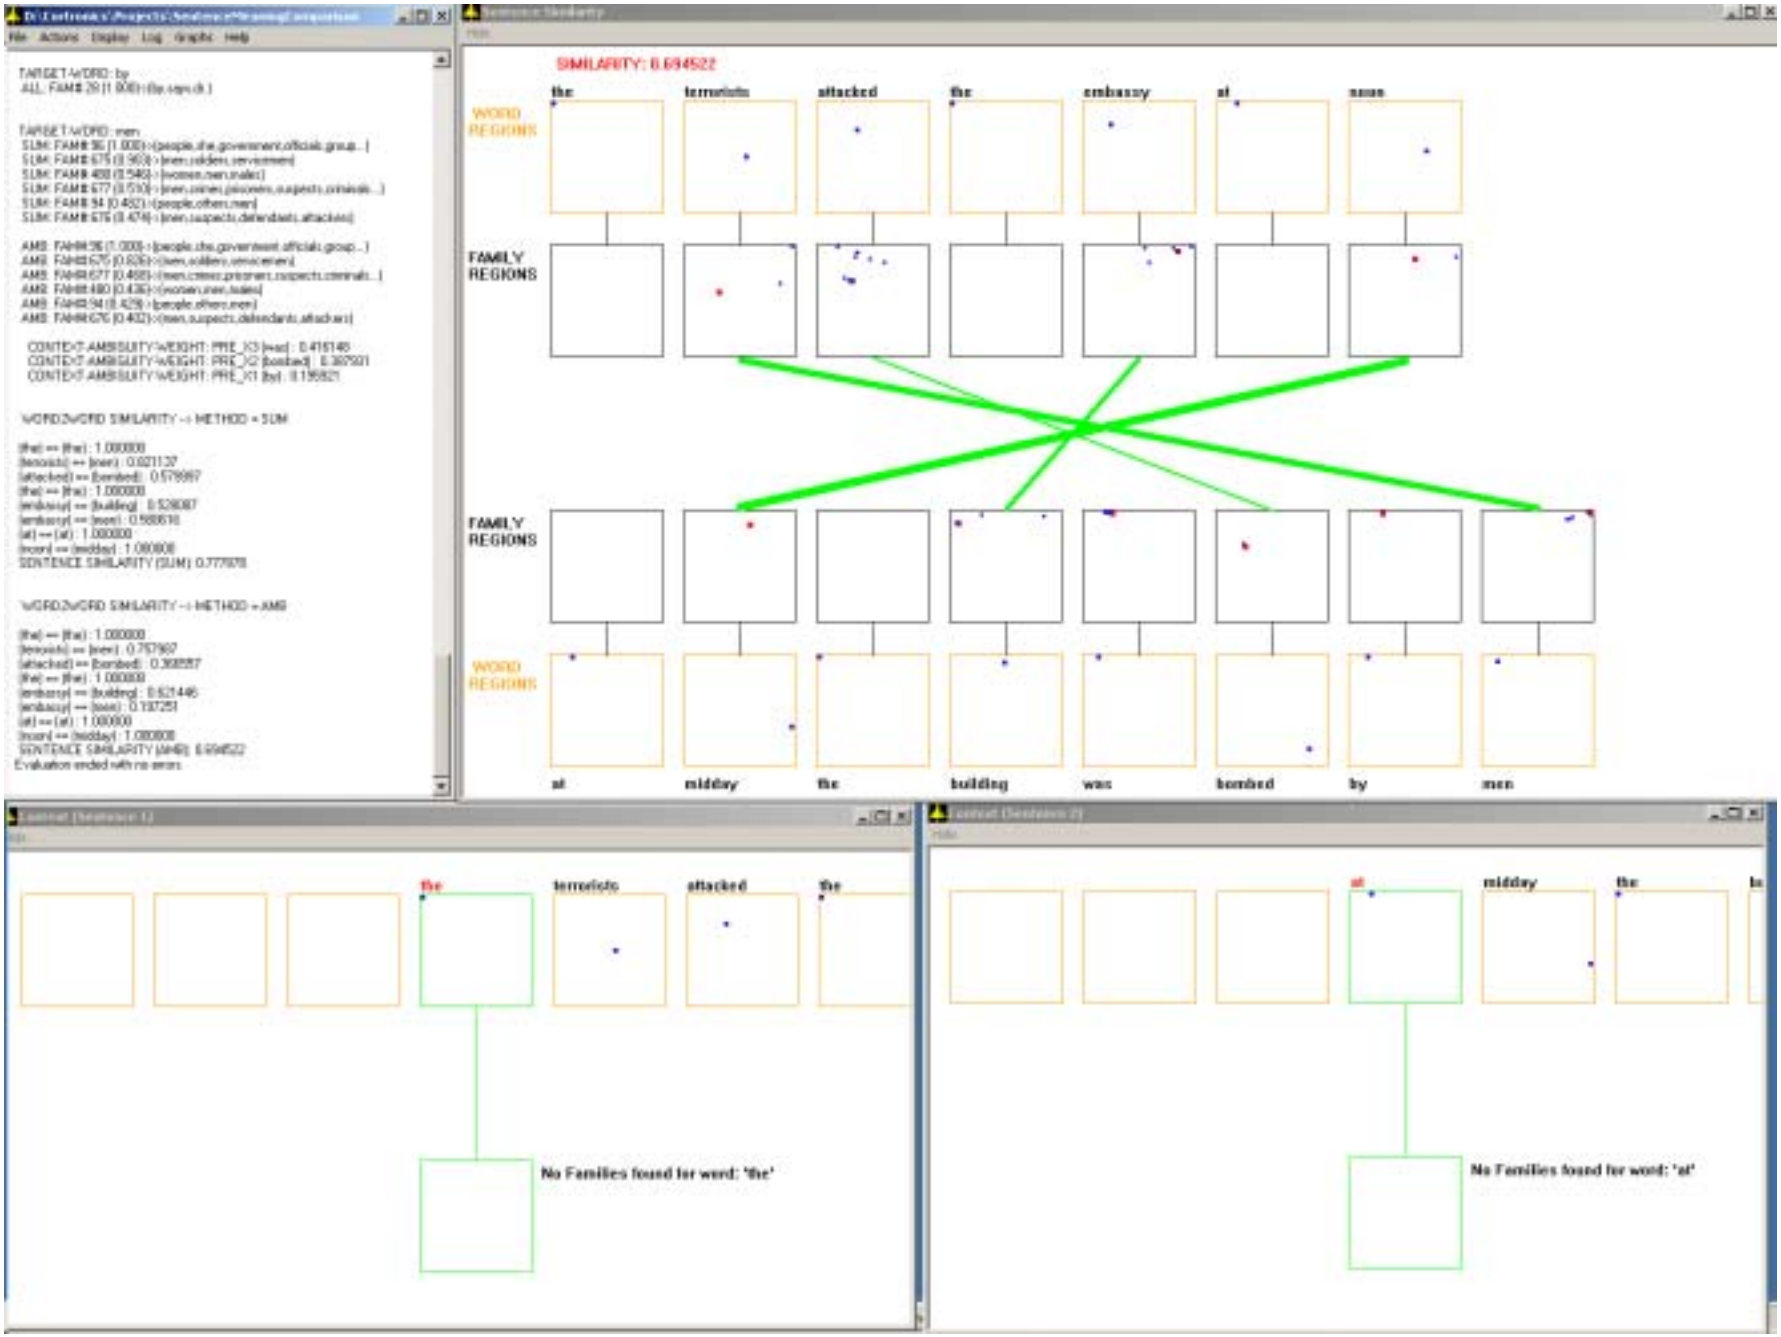
SIMILARITY: 0.694522
WORD REGIONS
FAMILY REGIONS
FAMILY REGIONS
WORD REGIONS
the
terrorists
attacked
the
embassy
al
noon
al
midday
the
building
was
bombed
by
men
No Families found for word: 'the'
No Families found for word: 'al'

The HNC Team

Other Team Members

Todd Gutschow - Project Leader and Senior Project Scientist, HNC Cortronics Project

Dr. Robert Means - Chief Technologist, HNC Cortronics Project

Kate Mark - Project Coordinator, HNC Cortronics Project

David Busby - Chief Brain Software Architect, HNC Cortronics Project

Dr. Syrus Nemat-Nasser - Scientist, HNC Cortronics Project

Rion Snow - Researcher, HNC Cortronics Project

Dr. Shailesh Kumar, Scientist, HNC Cortronics Project

Research Sponsors

HNC
ARDA
ONR

Artificial Brain Technology Demonstrations

YES!

IT THINKS

WORLD'S FIRST DEMONSTRATION OF THE

MAMMALIAN LAW OF THOUGHT

(Hecht-Nielsen, 2002)

Artificial Brain Technology Demonstrations

YES!

IT THINKS

WORLD'S FIRST DEMONSTRATION OF THE

MAMMALIAN LAW OF THOUGHT

(Hecht-Nielsen, 2002)