9. Language

Language Involves all of Cognition

- **Perception**: hearing & reading words
- **Attention**: picking out words, speakers from many
- **Motor**: speech, writing, etc.
- **Memory**: semantics, specific content – how do you encode plot of a book?
- **Executive Function**: maintaining context, planning speech, syntax structure…

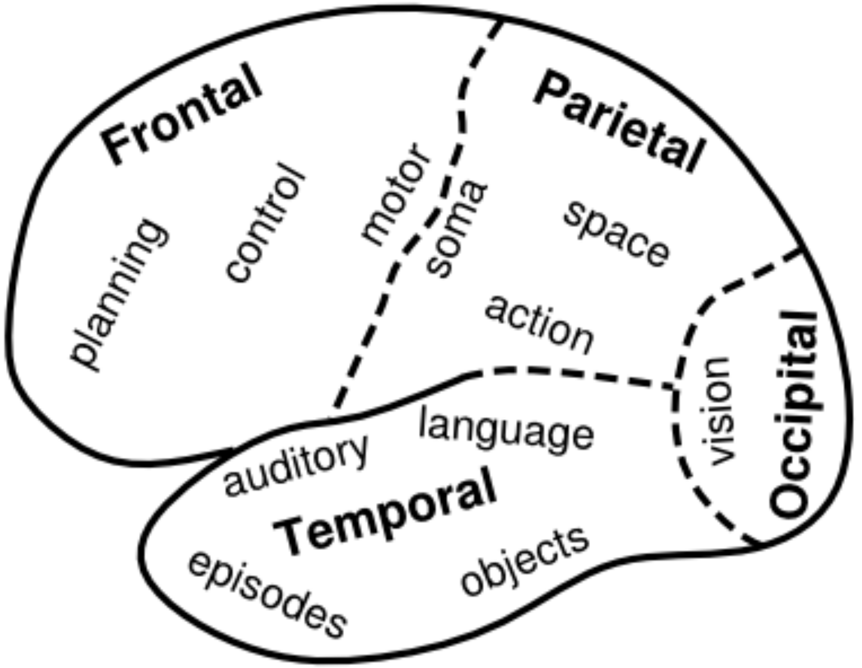

Language is Special

- Symbols
  - — thought is reduced to symbols to reconstruct thoughts
- Syntax
- Temporally-extended sequences
- Cultural transmission
- Embedded levels of structure (*generativity*):
  - — “The horse racing past the barn fell”
  - — “Isn’t it true that example-sentences that people that you know produce are more likely to be accepted?”

(slide < O’Reilly)

Language Controversies

- How special is language: just co-opting existing neural mechanisms vs. innate language modules?
  - Poverty of Stimulus (POS)
- Rules vs. regularities: Is there anything special about rule-like behavior in language?
  - Spelling to sound: Exceptions also have sub-rules
  - Over-regularization (add "-ed" = "goed") – evidence of rule system coming online?
  - Truly variable-like behavior? Generative, abstract.

What is Truly Novel?

- Pure syntax: "Adj Adj Noun Verb Adverb"
- "Colorless green ideas sleep furiously"
  - But: "Newly formed bland ideas are inexpressible in an infuriating way."
  - "It can only be the thought of verdure to come, which prompts us in the autumn to buy these dormant white lumps of vegetable matter covered by a brown papery skin, and lovingly to plant them and care for them. It is a marvel to me that under this cover they are labouring unseen at such a rate within to give us the sudden awesome beauty of spring flowering bulbs. While winter reigns the earth reposes but these colourless green ideas sleep furiously."
- "'Twas brillig, and the slithy toves…"
  - But each word has *some* overlap with real words..

(slide < O'Reilly)

Time and Language

- Distributed representation broken down and sent over a sequential channel:

The summer is a fun time for going to the beach, dancing, …

(slide < O'Reilly)

Distributed Representations of Words

I cnduo't bvleiee taht I culod aulaclty uesdtannrd waht I was rdnaieg. Unisg the icndeblire pweor of the hmuan mnid, aocdcrnig to rseecrah at Cmabrigde Uinervtisy, it dseno't mttaer in waht oderr the lterets in a wrod are, the olny irpoamtnt tihng is taht the frsit and lsat ltteer be in the rhgit pclae. The rset can be a taotl mses and you can sitll raed it whoutit a pboerlm. Tihs is bucseae the huamn mnid deos not raed ervey ltteer by istlef, but the wrod as a wlohe. Aaznmig, huh? Yaeh and I awlyas tghhuot slelinpg was ipmorantt! See if yuor fdreins can raed tihs too.

Traditional View of Language

- Language competence defined by knowledge of rules and exceptions
  - — e.g., “i before e except after c”
- Knowledge about words is stored in a central mental lexicon (dictionary)
- Each word has a lexical representation that is linked to information about its orthography, phonology, semantics

Connectionist View of Language

- Language is another set of input-output mappings (e.g., orthography to phonology, orthography to semantics)
- These mappings are trained up using the same learning algorithms used elsewhere (e.g., vision)
- The same pathways handle both rules and exceptions
- Hard to tell what is "regular" vs. "exceptional"
  - regular: clown, down
  - exception: blown… but blown goes with grown
- Distributed lexicon: Knowledge about words is embodied in reciprocal mappings between phonology, orthography, semantics – **there is no central "word representation"**

Language and Thought

- Socrates (d. 399 BCE): “that which we know we must surely be able to tell” (Plato, *Laches* 190c)
- Socrates on definition:
  - … what is that common quality, which is the same in all these cases, and which is called courage? (*Laches* 191e)
  - Well then, show me what, precisely, this ideal is, so that, with my eye on it, and using it as a standard, I can say that any action done by you or anybody else is holy if it resembles this ideal, or, if it does not, can deny that it is holy. (*Euthyphro* 6e)
  - And so of the virtues, however many and different they may be, they all have a common nature which makes them virtues. (*Meno* 72)

Formal Logic

- Originally developed by Aristotle (384–322 BCE)
- A syllogism:

  All men are mortal

  <u>Socrates is a man</u>

  ∴ Socrates is mortal

- Formal logic: the correctness of the steps depend only on their *form* (syntax), not their *meaning* (semantics):

  All M are P

  $\underline{S \text{ is } M}$

  ∴ S is P

- More reliable, because more mechanical

Calculus

- In Latin, *calculus* means pebble
- In ancient times *calculi* were used for *calculating* (as on an abacus), voting, and many other purposes
- Now, a *calculus* is:
  - — a *mechanical method* of solving problems
  - — by manipulating *discrete tokens*
  - — according to *formal rules*
- Examples: algebraic manipulation, integral & differential calculi, logical calculi

Common Features of Calculi

- Information (data) representation is:
  — formal (info. represented by arrangements)
  — finite (finite arrangements of atomic tokens)
  — definite (can determine symbols & syntax)
- Information processing (rule following) is:
  — formal (depends on arrangement, not meaning)
  — finite (finite number of rules & processing time)
  — definite (know which rules are applicable)

Thought as Calculation

- "By *ratiocination* I mean *computation*."
  — Thomas Hobbes (1588–1679)
- "Then, in case of a difference of opinion, no discussion … will be any longer necessary … It will rather be enough for them to take pen in hand, set themselves to the abacus, and … say to one another, "Let us calculate!" — Leibniz (1646–1716)
- Boole (1815–64): his goal was "to investigate the fundamental laws of those operations of mind by which reasoning is performed; to give expression to them in the symbolical language of a Calculus"

Early Investigations in Mechanized Thought

- Leibniz (1646–1716): mechanical calculation & formal inference
- Boole (1815–1864): “laws of thought”
- Jevons (1835–1882): logical abacus & logical piano ☞
- von Neumann (1903–1957): computation & the brain
- Turing (1912–1954): neural nets, artificial intelligence, “Turing test”

Symbol Grounding Problem

- If knowledge is just a network of propositions, how do symbols get their meaning?
- Think of a dictionary, which defines words in terms of other words
- If you don't know the meaning of any words, how do you know the meaning of anything?
- This is the *symbol grounding problem*
- Perhaps some words are defined in terms of images…

Symbolic AI & Cognitive Science

- Western philosophy has generally assumed knowledge and cognition can be described in terms of language-like structures (a *language of thought*)
- Symbolic AI and CogSci assumed intelligence resides in:
  - the structures of a *knowledge representation language*
  - deduction-like formal rules for their manipulation
- Knowledge represented at a *symbolic* level:
  - atomic word-level categories
  - related by sentence-like logical structures
- Semantics is reduced to syntax
  - that is, meaning is represented and manipulated formally (calculation)
- The brain doesn't matter because all computers are equally powerful
  - <u>competence</u> is relevant, but not <u>performance</u>
- Is a language of thought "the only game in town"?

Connectionist AI & Cognitive Science

- Knowledge is represented at a *subsymbolic* level:
  - in terms of minute, quantitative features
  - related by low-level, often statistical connections
  - grounded in sensorimotor interactions
- Knowledge is more akin to an image than a sentence
- The brain is very relevant because:
  - it shows us how to do connectionist information processing
  - it places tight constraints on models, because brains have to work in real time: <u>performance</u> is critical (100 step rule etc.)
  - on the other hand, e.g., symbolic approaches try to account for unlimited nesting, but connectionist approaches don’t need to

Some Questions…

- If knowledge is <u>not</u> represented in language-like rules, then how is language processed?
- What general processes are involved in reading, and how do these sometimes fail (e.g., in dyslexia)?
- How are we able to read “cat,” “yacht,” and “nust”?
- Why do children say “I goed to school” when previously they said “I went”?
- How do words come to mean anything?
- How do we go beyond words to sentences?

Outline: Language

A. Biology of Language
B. Reading and Dyslexia in the Triangle Model
C. Spelling to Sound Mappings in Word Reading
D. Latent Semantics in Word Co-occurrence
E. Syntax and Semantics in a Sentence Gestalt

A. Biology of Language

Language is Instinctual

Evidence:

- Skeletal specializations for speech in our earliest hominid ancestors
- Left hemisphere specialization evident before birth
- Ability to recognize any phoneme at birth

  — basis for learning one's native languages
- On the other hand, writing is not instinctual

Speech Output

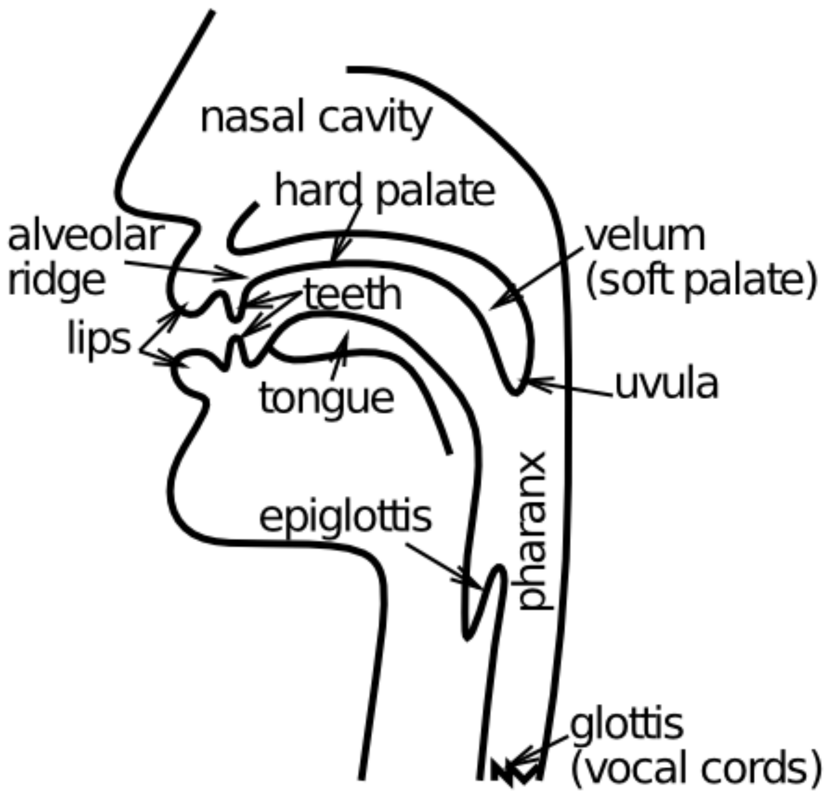

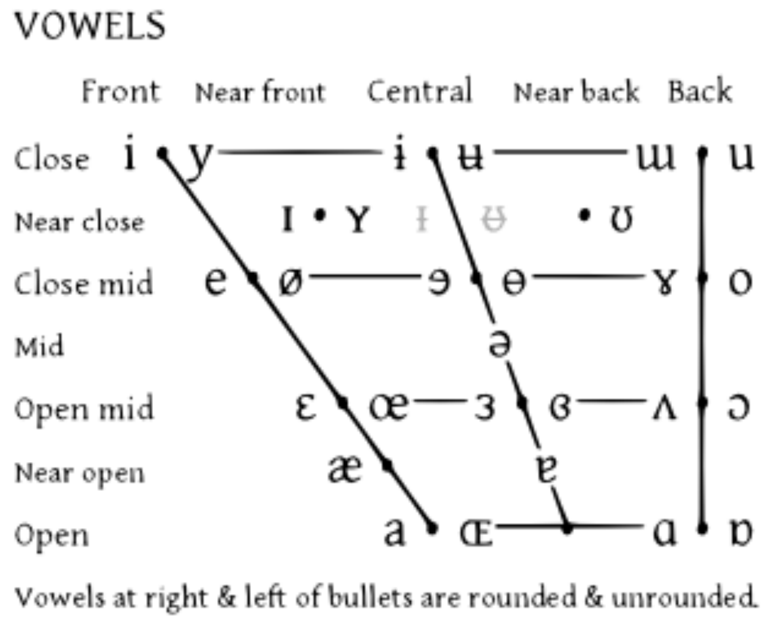

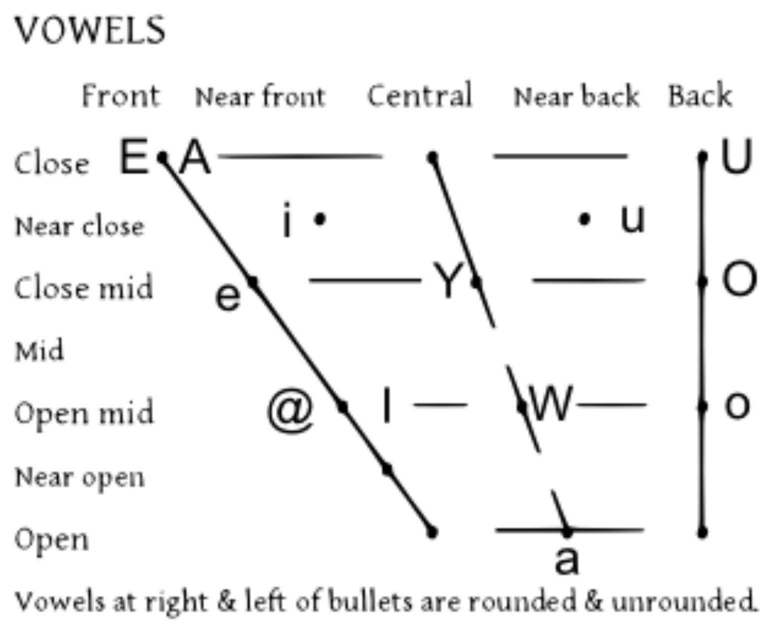

(slide < O'Reilly)

Speech Output

CONSONANTS (PULMONIC)

	LABIAL		CORONAL				DORSAL			RADICAL		LARYNGEAL
	Bilabial	Labio-dental	Dental	Alveolar	Palato-alveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Epi-glottal	Glottal
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ			
Plosive	p b	ȹ ȸ		t d		ʈ ɖ	c ɟ	k ɡ	q ɢ		ʡ	ʔ
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	ʜ ʢ	h ɦ
Approximant		ʋ		ɹ		ɻ	j	ɰ				
Trill	ʙ			r					ʀ		ᴙ	
Tap, Flap		ⱱ		ɾ		ɽ						
Lateral fricative				ɬ ɮ		ꞎ	𝼆	𝼄				
Lateral approximant				l		ɭ	ʎ	ʟ				
Lateral flap				ɺ		𝼈						

Where symbols appear in pairs, the one to the right represents a modally voiced consonant, except for murmured ɦ.
Shaded areas denote articulations judged to be impossible. Light grey letters are unofficial extensions of the IPA.

CONSONANTS (NON-PULMONIC)

Anterior click releases (require posterior stops)	Voiced implosives	Ejectives
ʘ Bilabial fricated	ɓ Bilabial	ʼ *Examples:*
ǀ Laminal alveolar fricated ("dental")	ɗ Dental or alveolar	pʼ Bilabial
ǃ Apical (post)alveolar abrupt ("retroflex")	ʄ Palatal	tʼ Dental or alveolar
ǂ Laminal postalveolar abrupt ("palatal")	ɠ Velar	kʼ Velar
ǁ Lateral alveolar fricated ("lateral")	ʛ Uvular	sʼ Alveolar fricative

CONSONANTS (CO-ARTICULATED)

ʍ Voiceless labialized velar approximant
w Voiced labialized velar approximant
ɥ Voiced labialized palatal approximant
ɕ Voiceless palatalized postalveolar (alveolo-palatal) fricative
ʑ Voiced palatalized postalveolar (alveolo-palatal) fricative
ɧ Simultaneous x and ʃ (disputed)
k͡p t͡s Affricates and double articulations may be joined by a tie bar

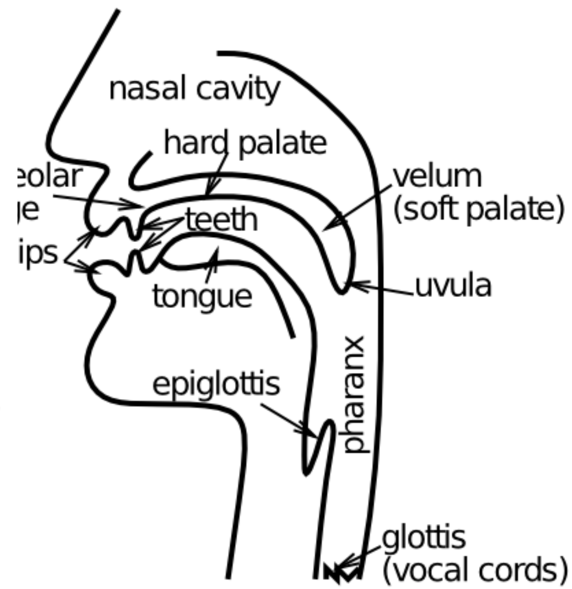

(slide < O'Reilly)

Biology of Language

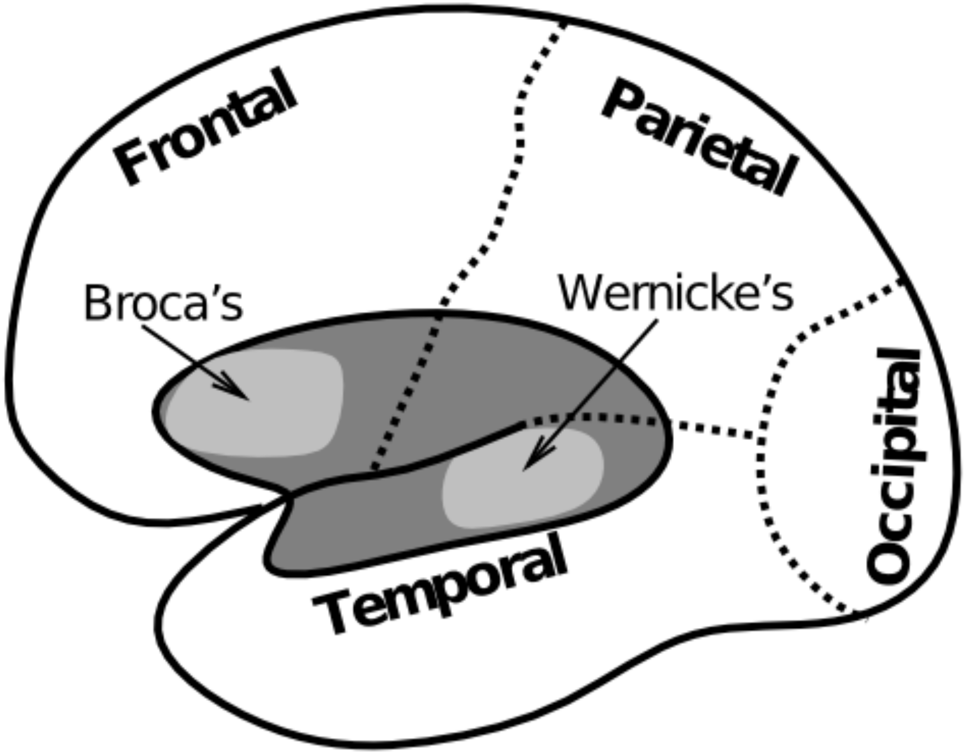

- Broca's = speech output, syntax, grammar (surface production):
  - — active maintenance of context to perform syntactic processing
- Wernicke's = semantic comprehension + output (deep):
  - — interconnected overlapping distributed info about semantics

Aphasias

- *Aphasia* = an acquired language disorder
  - disorder of higher-order processing, not articulation
- Wernicke's aphasia (*receptive aphasia*)
  - damage to Brodmann's area 22 (temporal lobe: semantics)
  - can't understand language, but can speak correctly
  - but not necessarily meaningfully
  - may eventually talk nonstop nonsense (because no self monitoring?)
- Broca's aphasia (*expressive aphasia*)
  - damage to Brodmann's areas 44 & 45 (frontal lobe: motor sequencing)
  - comprehension intact, but difficulty speaking correctly: "dog … walk"
  - may eventually become mute
  - deficits in understanding syntactically complex sentences (due to FC role in temporally extended behavior?)
- LH lesions have analogous effects on signing and on reading and writing

Wernicke and Broca in Right Hemisphere?

- LH is dominant in language, so what do the corresponding areas do in RH?
- They are involved in language *prosody* (rhythm and intonation)
  - — musical aspects
  - — important in conveying meaning
- Lesion in RH Wernicke ⇒ inability to understand intonation
- Lesion in RH Broca ⇒ speaking in flat tones

Reading and Writing

- Written language is an invention, not innate
- LH involved in written language as in spoken
- LH parietal areas involved in written language
- Damage to LH Brodmann's areas 39 & 40 (parietal) ⇒ acquired illiteracy
  - — i.e., an acquired inability to understand written language

B. Reading and Dyslexia in the Triangle Model

Reading: The "Triangle Model"

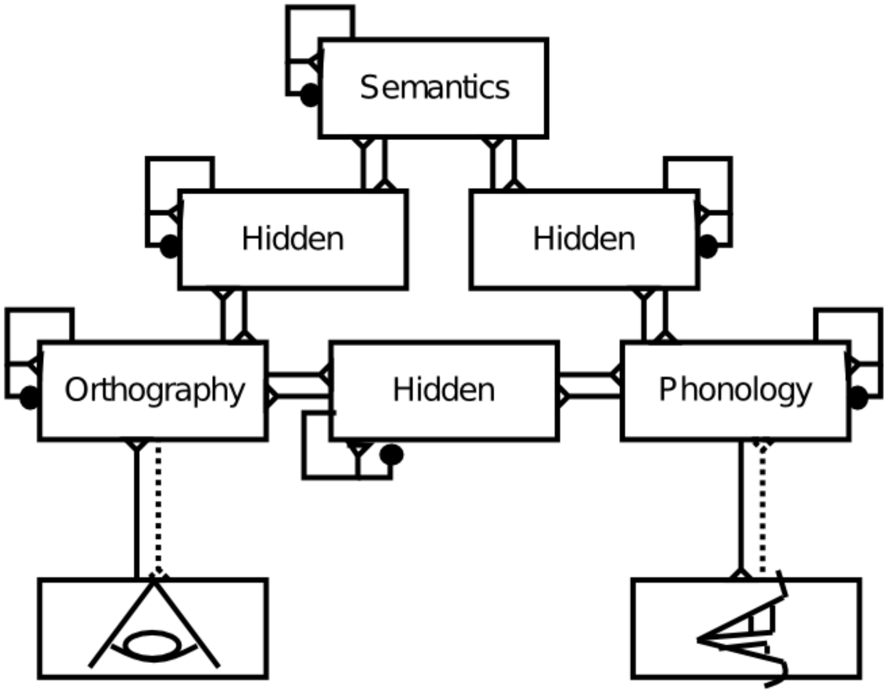

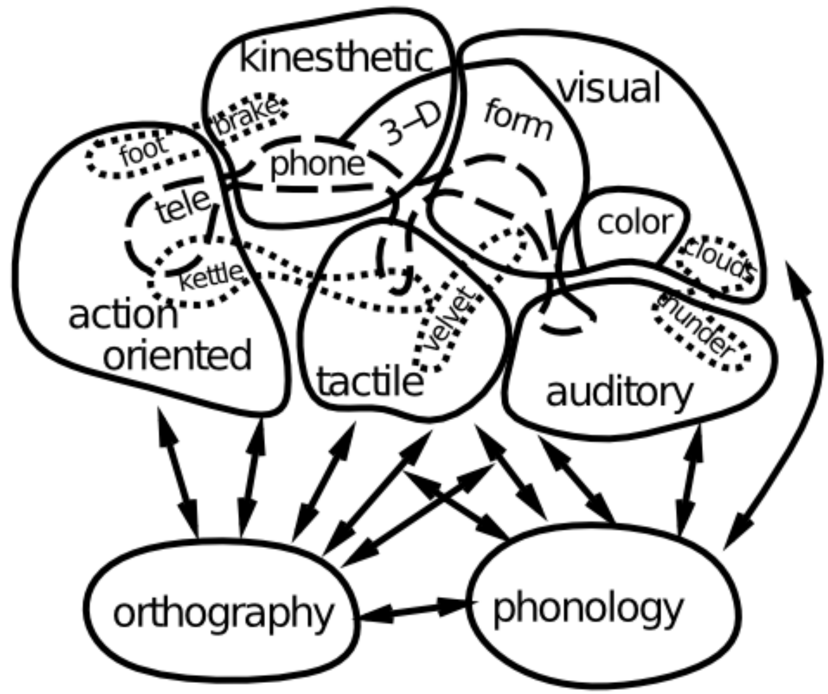

Note: There is no single "lexicon," no "word units"

Acquired Dyslexias

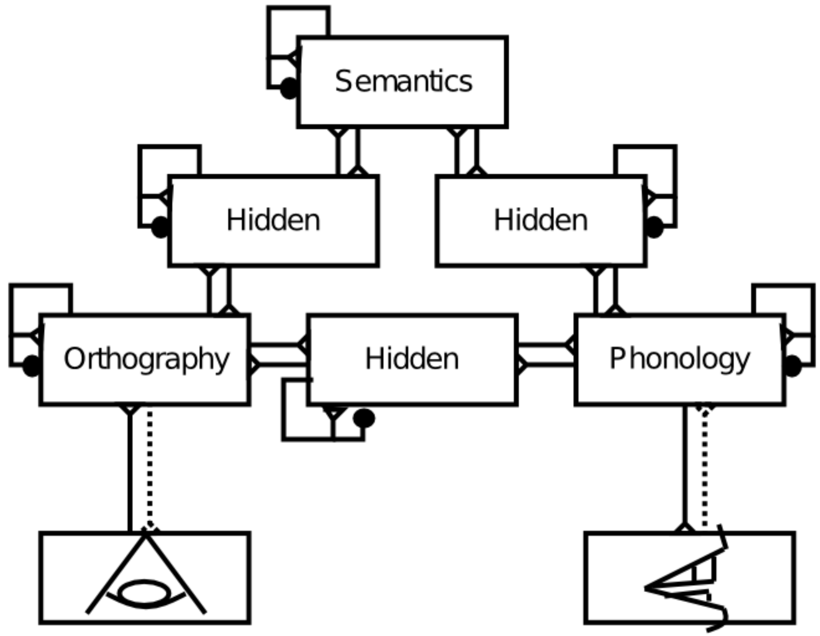

- **Phonological**: nonword (*nust) errors
- **Deep**: phonological + semantic errors (“dog” → “cat”) + visual errors (“dog” → “dot”)
- **Surface**: exception (“yacht”) errors + visual errors + impaired semantic access

(slide < O’Reilly)

Hypotheses

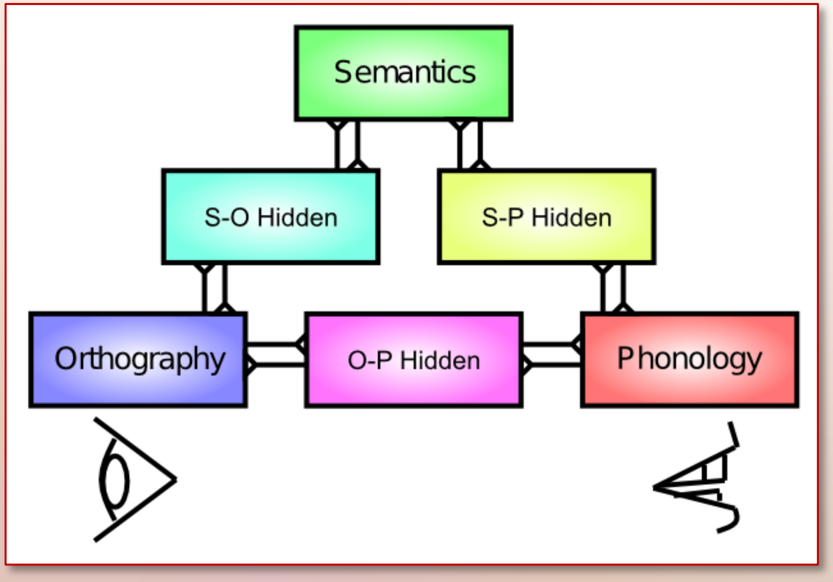

- Phonological dyslexia
  - difficulty reading nonwords (*nust or *mave)
  - explained by damage to direct orthography → phonology pathway
- Deep dyslexia
  - may make semantic substitutions (“orchestra” → “symphony”)
  - may make visual errors (“dog” → “dot”)
  - significant damage to direct pathway
- Surface dyslexia
  - intact nonword reading, but access to semantics impaired
  - impaired pronunciation of exception words (“yacht”)
  - explained by damage to indirect (semantic) pathway

The dyslex Model

- Trained on all pathways (ortho ⇔ phono etc.) for 40 4-letter monosyllabic words (e.g., “flag,” “star”)
- Concrete & abstract words use different pools of semantic units
- Abstract words activate fewer semantic units than concrete words

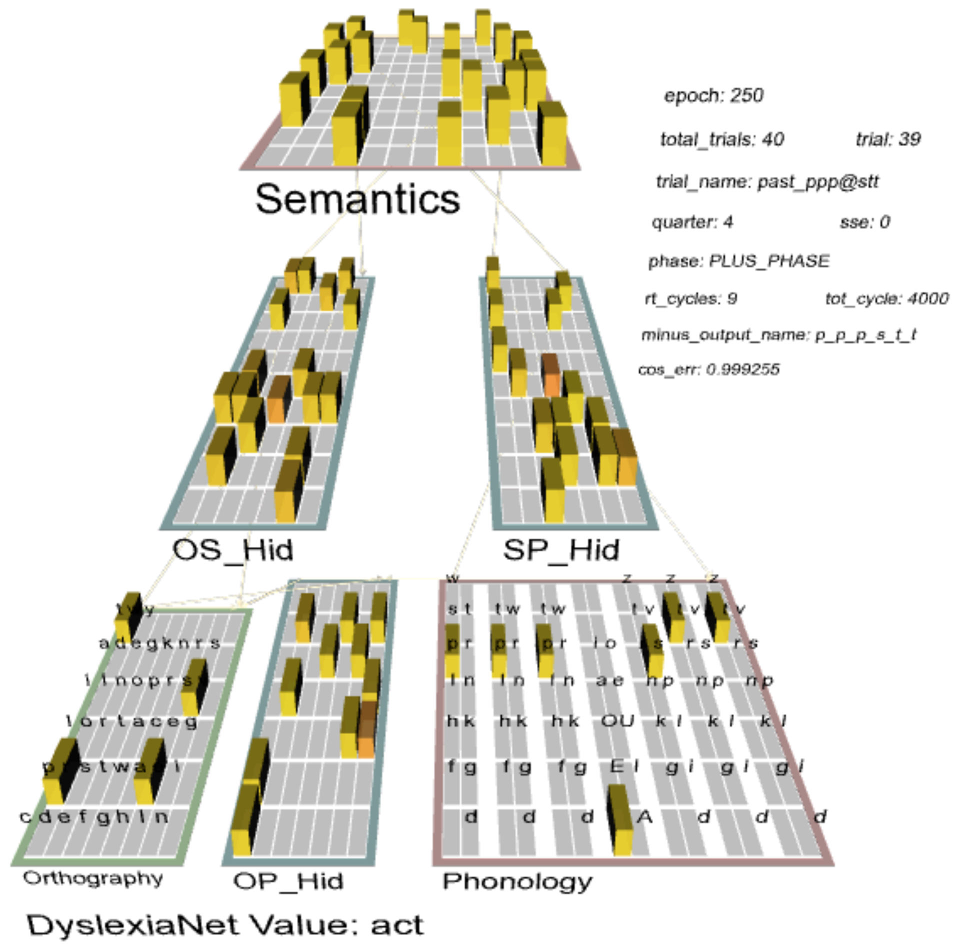

Word Corpus

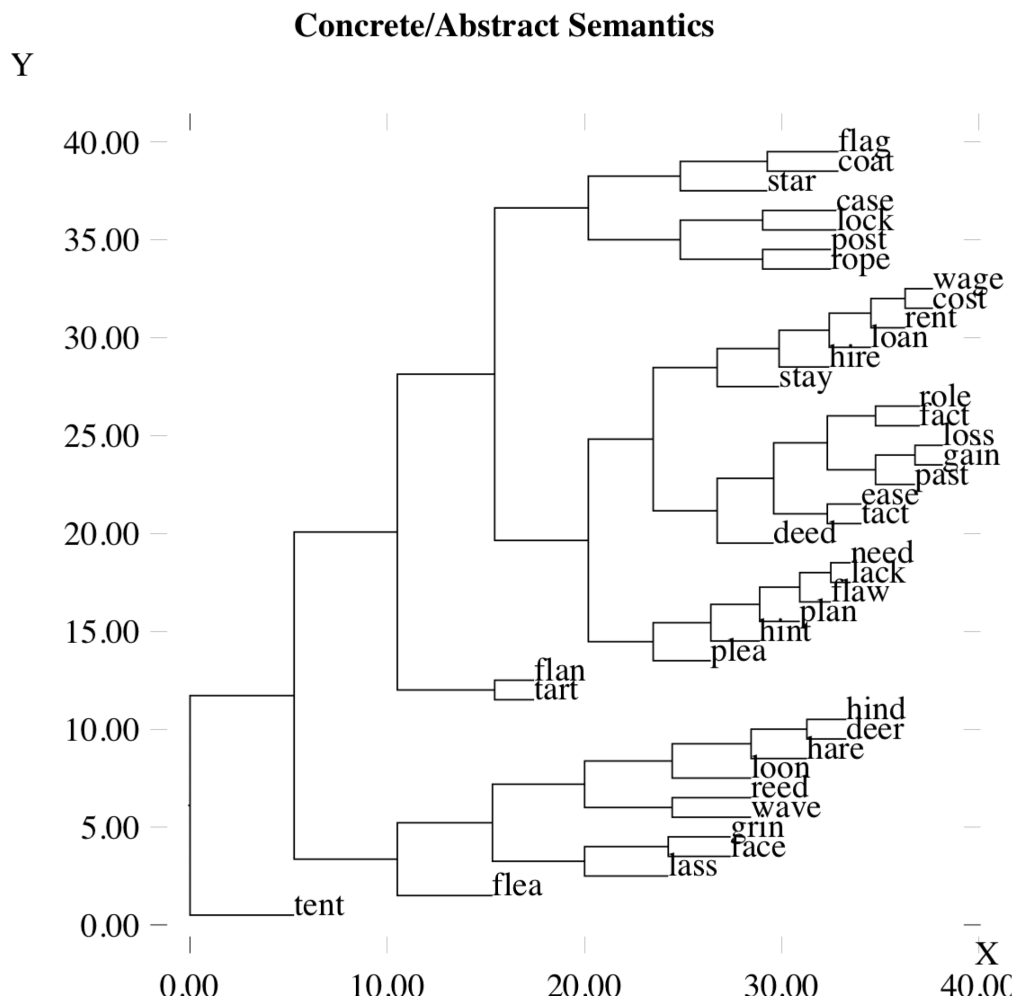

- Cluster plot of semantic similarity for words used in model
- Words that are semantically close are sometimes confused for each other in simulated deep dyslexia

Partial Direct Pathway Lesions (with or without Semantic Pathway)

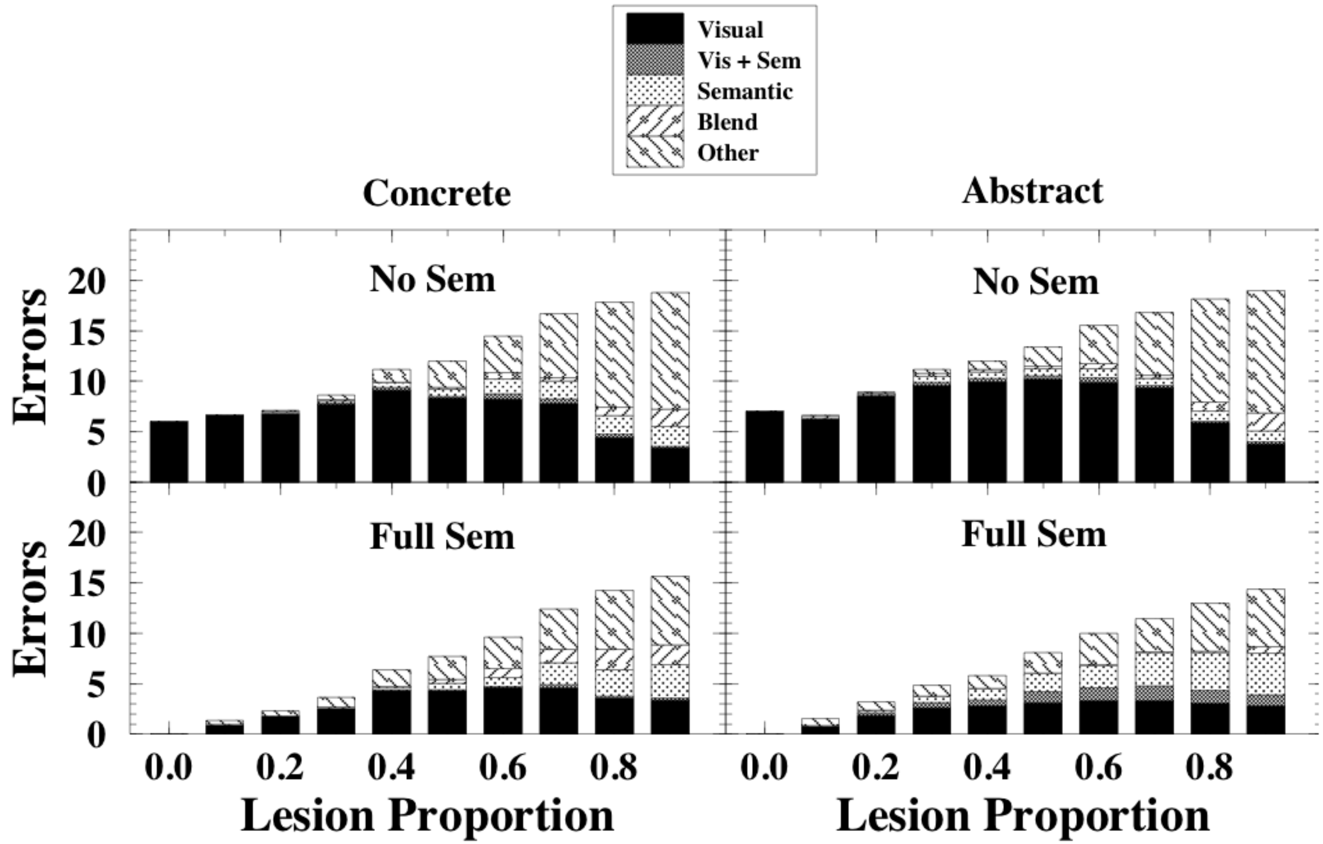

- minor direct damage: just visual errors
- more damage: semantic errors ⇒ deep dyslexia even with full semantics
- cf. phonological and deep dyslexia

(slide < Frank)

Partial Semantic Pathway Lesions

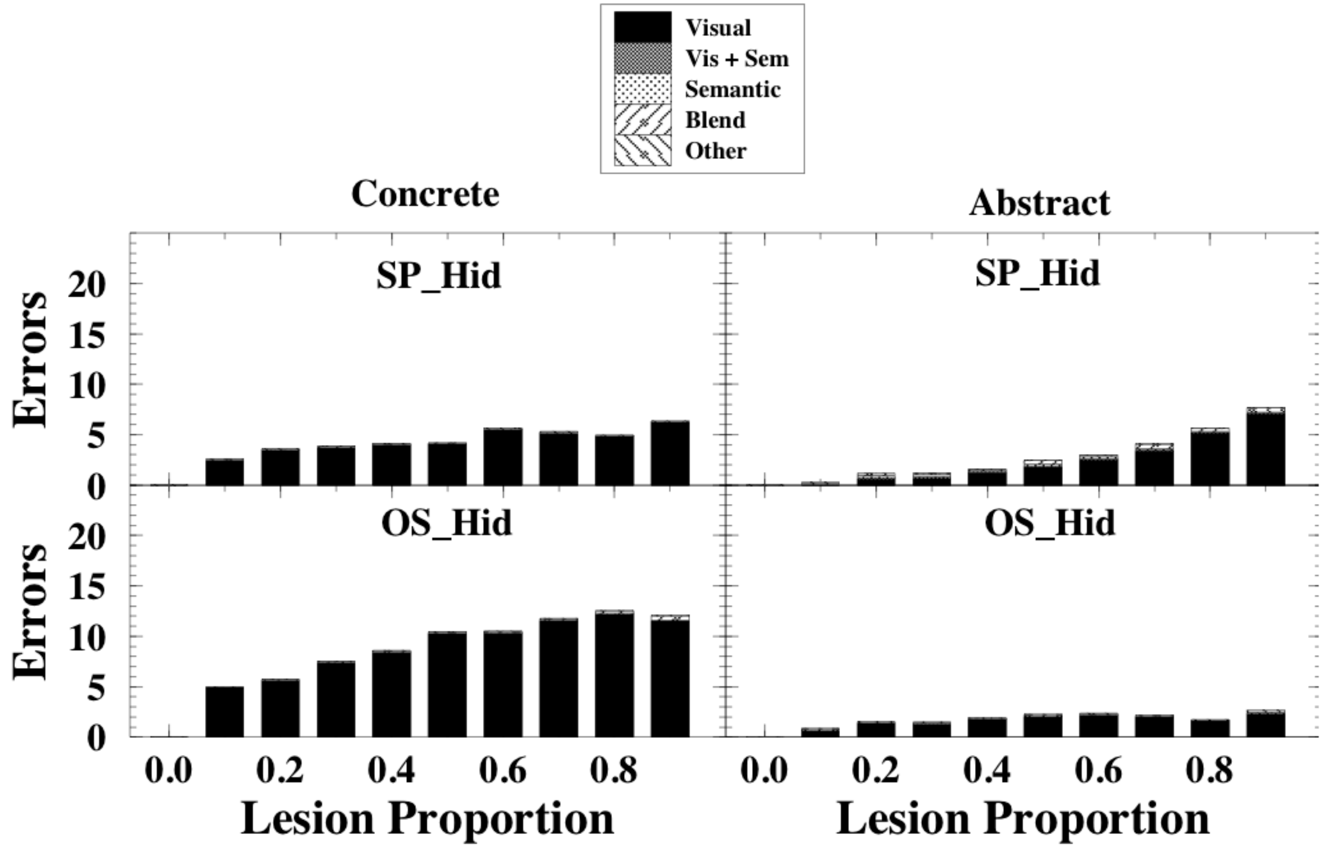

- visual errors with semantic pathway lesions; no semantic errors!
- more for concrete than abstract
- cf. surface dyslexia

(slide < Frank)

Partial Semantic Pathway Lesions with Complete Direct Pathway Lesions

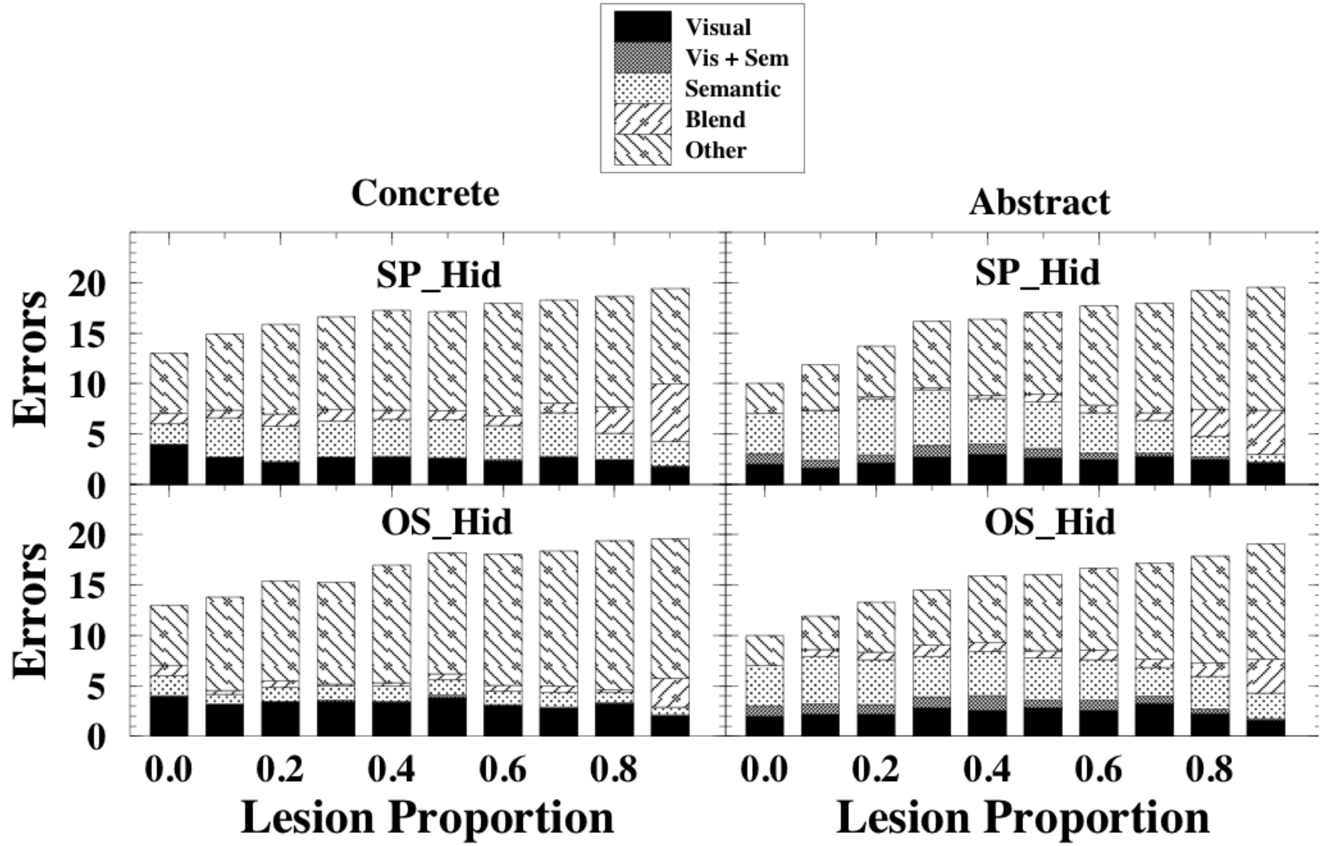

- multiple errors types
- more abstract semantic errors than concrete
- cf. deep dyslexia

(slide < Frank)

Abstract vs. Concrete Summary

- Semantic pathway lesions hurt concrete words more than abstract words
- Concrete words are more strongly represented (more units active) than abstract words in the semantic pathway
- Learning is a function of activation, so the semantic pathway learns more about concrete words
- The more semantic pathway learns about concrete words the less direct pathway learns
- The less the direct pathway learns, the less it is able to support performance on its own
- With full direct pathway lesions, the model makes more semantic errors for abstract than concrete
- Abstract words have less distinctive semantic reps than concrete words
- The model is more likely to fall into wrong semantic attractor for abstract words

(slide < Frank)

emergent Demonstration dyslex.proj

C. Spelling to Sound Mapping in Word Reading

Modeling the direct pathway

Regularities & Exceptions

- Regularities in pronunciation are often partial, context dependent: *bint
- “i” in mint, hint, stint,… (regular) vs. pint (exception)
- but also: mind, find, hind,… (regular) mine, fine, dine,… (regular)
- Pronunciation depends on context
- Exceptions are extremes of context dependence
- Need a range of context dependency for regulars and exceptions.

Gradual Invariance Transformation

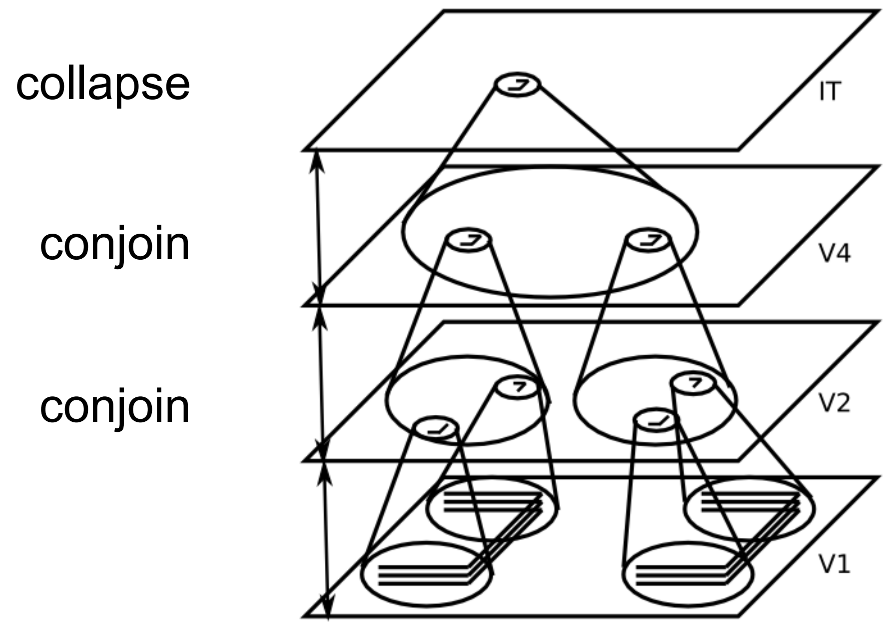

Increasing receptive field size gives two options:

- Encode location-sensitive conjunctions of features
  - e.g., a horizontal line below and to the left of a vertical line
- Collapse over location or size information

Illustration of Spatial Invariance

I cnduo't bvleiee taht I culod aulaclty uesdtannrd waht I was rdnaieg. Unisg the icndeblire pweor of the hmuan mnid, aocdcrnig to rseecrah at Cmabrigde Uinervtisy, it dseno't mttaer in waht oderr the lterets in a wrod are, the olny irpoamtnt tihng is taht the frsit and lsat ltteer be in the rhgit pclae. The rset can be a taotl mses and you can sitll raed it whoutit a pboerlm. Tihs is bucseae the huamn mnid deos not raed ervey ltteer by istlef, but the wrod as a wlohe. Aaznmig, huh? Yaeh and I awlyas tghhuot slelinpg was ipmorantt! See if yuor fdreins can raed tihs too.

Word Reading as Object Recognition

- Word reading modeled as spatially invariant object recognition
- V1-like: Words show up in different locations in input
- V2/V4-like: extracts more complex combinations of letters, also developing more invariant representations
  - integrates individual letters or multi-letter features over multiple different locations
- IT-like: fully spatially invariant representation of the word

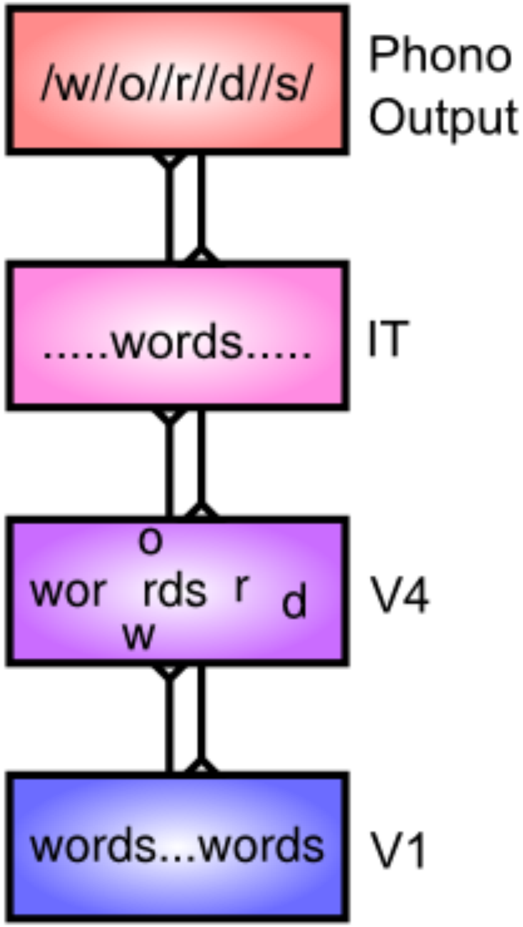

Reading Model (Direct Pathway)

- Detailed model of the "direct" reading pathway (ortho → phono)
- Trained to pronounce large set of regular & exception words
  - nearly 3000 sampled according to frequency
- Generalization testing: nonwords (e.g., *nust)

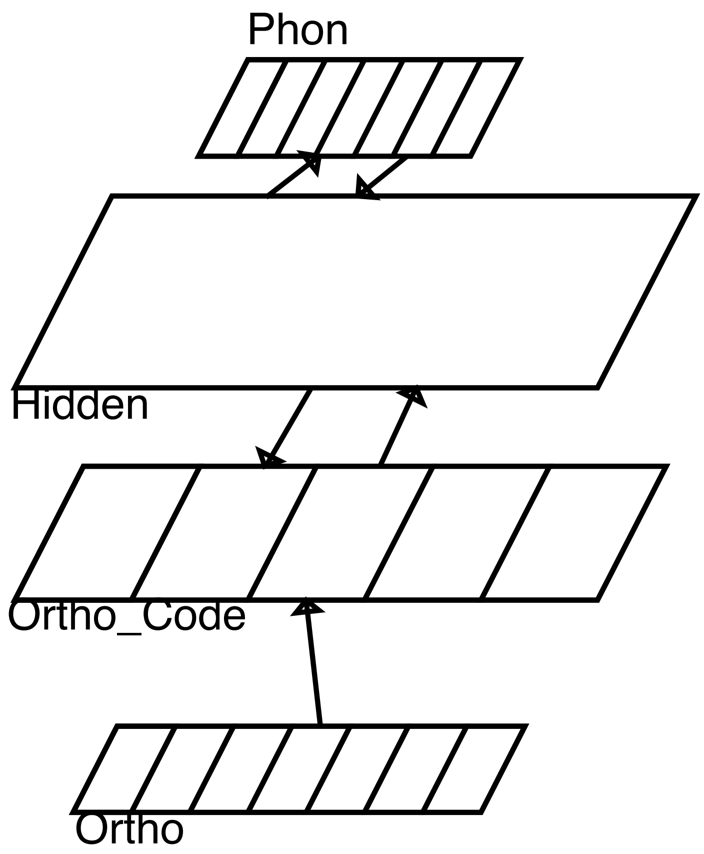

(slide based on Frank)

Phonological Representations

- Same 7 slot vowel-centered representations as before:
  - face = fffAsss
  - grin = grrinnn
  - star = sttarrr
  - post = pppOstt
- Except instead of using a localist representation of each phoneme, we use a distributed representation
- This allows us to represent the fact that phonemes vary in their similarity to one another

ss Network Structure

- Ortho: 7 slots with 27 letter units
  - words shift to right for spatial invariance
- OrthoCode: receptive fields are 3 Ortho slots
  - learns groups of 1–3 letters
- Hidden: spatial invariance + increased complexity
  - encodes context-sensitive regularities
- Phon: 7 slots centered on vowel

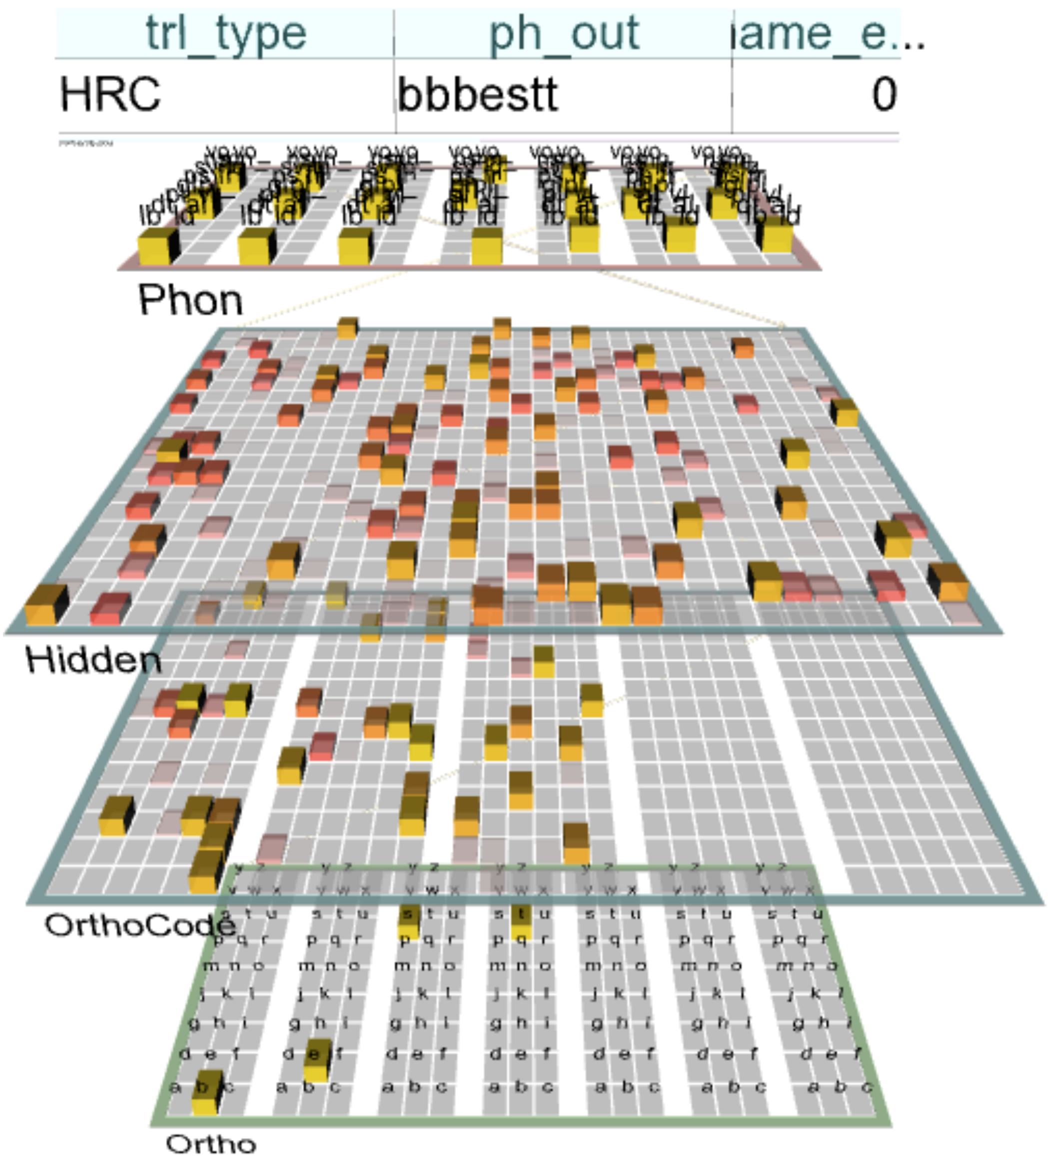

Test Sets of Nonwords

- **Glushko regulars**: nonwords constructed to match strong regularities, for example *nust, which is completely regular (e.g., must, bust, trust, etc.).
- **Glushko exceptions**: nonwords that have similar English exceptions and conflicting regularities, such as *bint (could be like mint, but also could be like pint). We score these items either according to the predominant regularity, or also including close exceptional cases (alt OK in the table).
- **McCann & Besner ctrls**: these are pseudo-homophones and matched controls, that sound like actual words, but are spelled in a novel way, for example *choyce (pronounced like choice), and the matched control is *phoyce.
- **Taraban & McClelland**: has frequency matched regular and exception nonwords, for example *poes (like high frequency words goes or does), and *mose, like lower frequency pose or lose.

Nonword Performance

Regularity tests (Glushko): bint → /bint/

Pseudo-homophones (McCann & Besner):
phoyce → /fYs/, choyce → /CYs/

Matched regularity/exception cases (Taraban):
High freq: poes → /pOz/, goes → /gOz/, does → /d^z/
Low freq: mose → /pOs/, poes → /pOz/, lose → /lUz/

Nonword Set	Model	PMSP	People
Glushko regulars	95.3	97.7	93.8
Glushko exceptions raw	79.0	72.1	78.3
Glushko exceptions alt OK	97.6	100.0	95.9
McCann & Besner ctrls	85.9	85.0	88.6
McCann & Besner homoph	92.3	n/a	94.3
Taraban & McClelland	97.9	n/a	100.0[1]

(slide < O'Reilly)

Reading Summary

- One network can learn both regular pronunciations and exceptions, *and* it can generalize properly to nonwords
- Network learns a good mix of context-dependent and context-invariant representations on its own

emergent Demonstration: ss.proj

D. Latent Semantics in Word Co-occurrence

And the Emergence of Semantic Representations through the Generalized Hebbian Algorithm

How Do Words Come to Mean Anything?

- What gives words their meaning?
- Where does this meaning come from?

Distributed Semantics

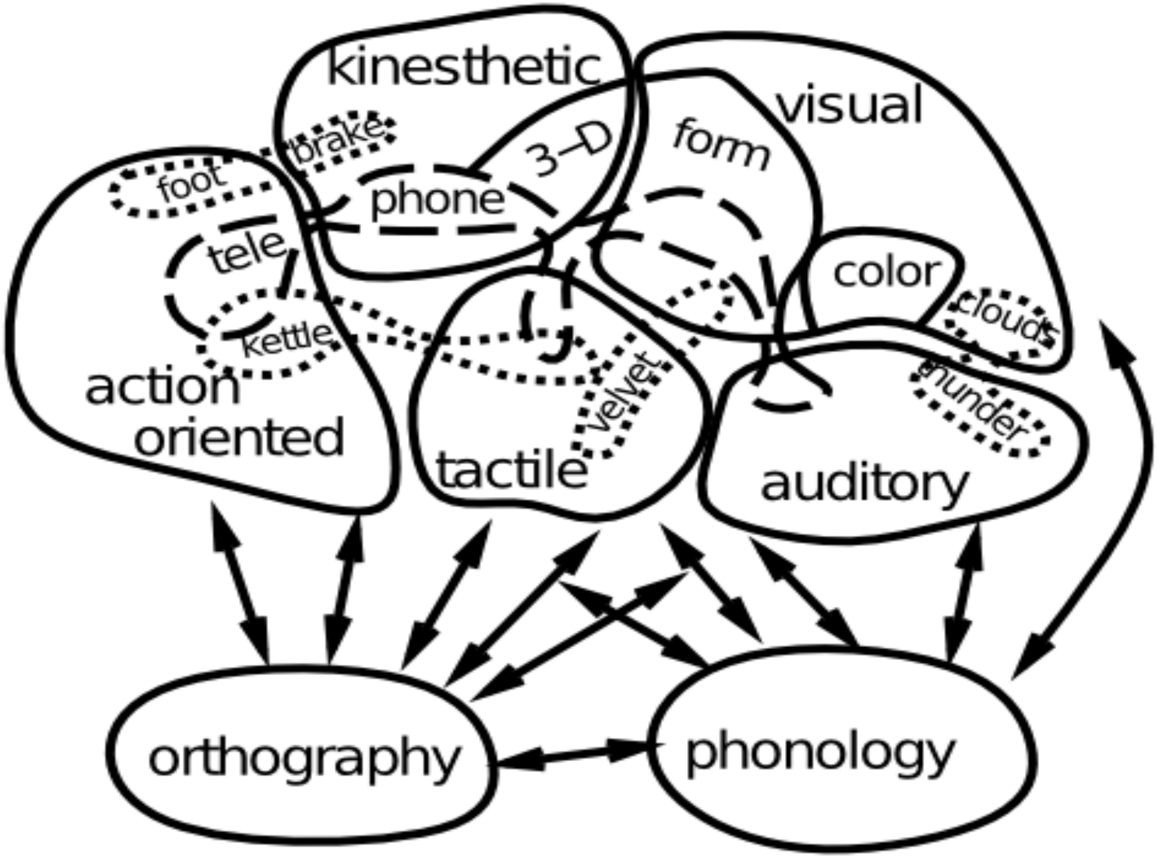

Semantics is distributed across specialized processing areas

Distributed Semantics in the Brain

- Semantics:
  - — distributed in domain-specific areas
  - — fundamentally <u>embodied</u> and <u>grounded</u>
- Anterior pole of temporal lobe
  - — perhaps a central hub for coordinating semantics in distributed areas
  - — perhaps especially important for abstract words

Correlational Semantics

- Hebbian learning encodes structure of word co-occurrence.
- Same idea as:
  - — V1 receptive field learning: learn the strong correlations
  - — Similar to Latent Semantic Analysis (LSA)

Latent Semantic Analysis

- Let X be $N \times d$ matrix such that X_{ij} is the number of occurrences of term *i* in document *j*
- By *singular value decomposition* (SVD): $\mathrm{X} = \mathrm{U}\Sigma\mathrm{V}^{\mathrm{T}}$, where U and V are orthogonal matrices and Σ is $N \times d$ with just diagonal elements $\sigma_1, \sigma_2, \dots, \sigma_n$, where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n$ are the *singular values*

The σ_i^2 are the eigenvalues of XX^{T} and $\mathrm{X}^{\mathrm{T}}\mathrm{X}$

- Compute rank *k* approximation of X by keeping the largest *k* singular values:
$$\mathrm{X}_k = \mathrm{U}_k \Sigma_k \mathrm{V}_k^{\mathrm{T}}$$
(that is, factor X through a lower-dimensional semantic space: a bottleneck to extract meaningful dependencies: truncated SVD)
- The rows $\mathbf{t}_i^{\mathrm{T}}$ of U_k are *k*-dimensional representations of the terms (and the columns of $\mathrm{V}_k^{\mathrm{T}}$ are *k*-dimensional representations of the documents)
- Terms *p* and *q* can be compared by the inner product of $\Sigma_k \mathbf{t}_p$ and $\Sigma_k \mathbf{t}_q$ (cosine similarity) = $\mathbf{t}_p^{\mathrm{T}} \Sigma_k^2 \mathbf{t}_q = \mathbf{t}_p^{\mathrm{T}} \Lambda_k \mathbf{t}_q$
- Reduced representation: $\hat{\mathbf{t}} = \Sigma_k^{-1} \mathrm{U}_k^{\mathrm{T}} \mathbf{t}$

Computation by Generalized Hebbian Algorithm

- Recall from supplementary slides to ch. 4 that normalized Hebbian learning extracts the first principal component (Oja's Rule)
- The (left) singular vectors (columns of U) are the eigenvectors of XX^{T} (term co-occurrences) with corresponding eigenvalues $\sigma_1^2, \sigma_2^2, \ldots, \sigma_k^2$
- We can extract these with a series of neural nets $y_i = \mathbf{w}_i^{\mathrm{T}} \mathbf{x}_i$
- Let $\mathbf{x}_1 = \mathbf{x}$ and $\Delta \mathbf{w}_1 = \epsilon[y_1 \mathbf{x}_1 - y_1 \mathbf{w}_1 y_1]$ (Hebbian with normalization)
- We find weight vector $\mathbf{w}_1 \rightarrow \mathrm{U}_{\cdot 1}$ (1^{st} eigenvector of $\mathcal{E}\{\mathbf{x}\mathbf{x}^{\mathrm{T}}\}$) and output variance $\langle y_1^2 \rangle \rightarrow \sigma_1^2$ (1^{st} eigenvalue)
- This procedure can be repeated to extract the other principal components by a variant of the Gram-Schmidt algorithm

Computation by GHA (continued)

- To compute the 2nd PC, subtract the first PC from input and pass it to 2nd neural net
- Let $\mathbf{x}_2 = \mathbf{x} - \sigma_1^2 \mathrm{U}_{\cdot 1} \approx \mathbf{x}_1 - \mathbf{w}_1 y_1$ (remainder of input to train output 2)
- Therefore,

$$
\begin{aligned}
\Delta\mathbf{w}_2 &= \epsilon[y_2\mathbf{x}_2 - y_2\mathbf{w}_2 y_2] \\
&= \epsilon[y_2(\mathbf{x}_1 - \mathbf{w}_1 y_1) - y_2\mathbf{w}_2 y_2] \\
&= \epsilon[y_2\mathbf{x} - y_2(\mathbf{w}_1 y_1 + \mathbf{w}_2 y_2)]
\end{aligned}
$$

- In general, let $\mathbf{x}_i = \mathbf{x}_{i-1} - \mathbf{w}_{i-1} y_{i-1}$ and

$$
\begin{aligned}
\Delta\mathbf{w}_i &= \boldsymbol{\epsilon}[y_i\mathbf{x}_i - y_i\mathbf{w}_i y_i] \\
&= \boldsymbol{\epsilon}\left[y_i\mathbf{x} - y_i \sum_{\boldsymbol{k \le i}} \mathbf{w}_k y_k\right]
\end{aligned}
$$

- In this way, weight vectors converge to the eigenvectors, which are the rows of U and yield semantic representations (i.e., incremental LSA)

emergent Demonstration: sem.proj

E. Syntax and Semantics in a Sentence Gestalt

How to understand language without parsing

Sentences and Syntax

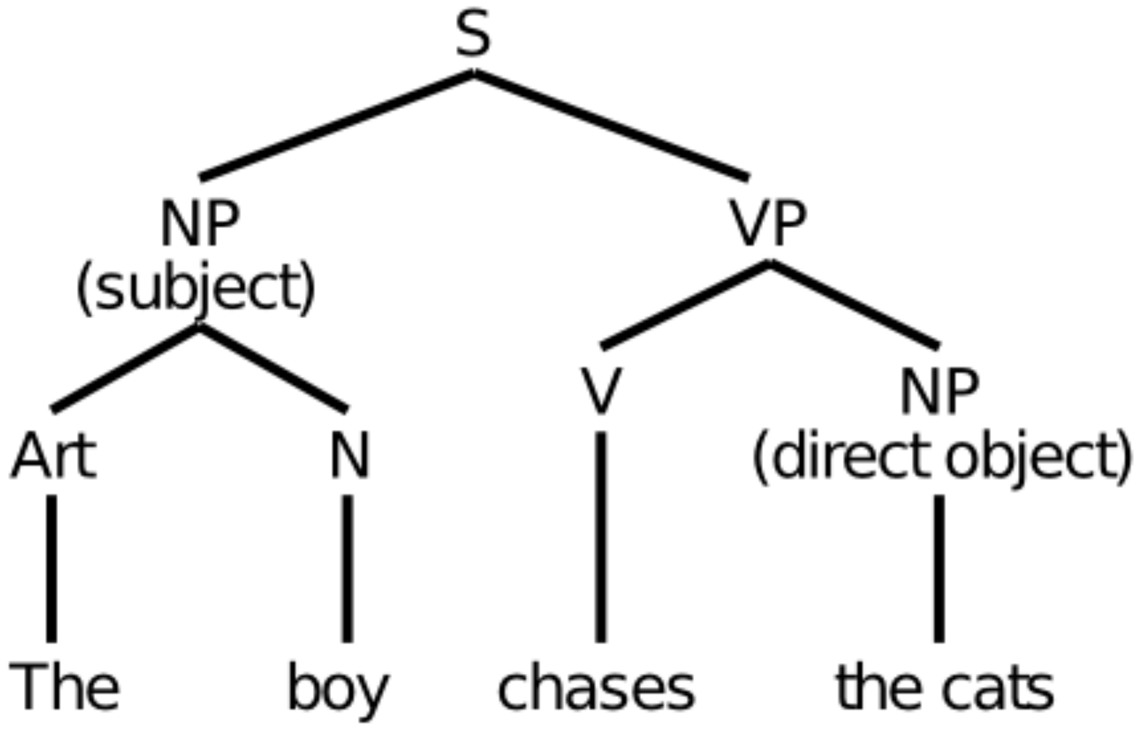

Is this how it really works?

Traditional Language Analysis (used in compilers)

- Lexical analysis: breaks input stream into lexical units (e.g., words)
- Syntactic analysis: parses lexical stream, produces a syntactic structure (parse tree)
- Semantic analysis: "decorates" parse tree with semantic information (meanings)
- Pragmatics: Interpret semantics (action / effect)

Semantics Affects Syntactic Interpretation

- Time flies like an arrow.
- Fruit flies like a banana.

- The slippers were found by the nosy dog.
- The slippers were found by the sleeping dog.

- Syntax depends on semantics very deeply

(slide < O'Reilly)

The “Gestalt” Alternative

- Just try to get the gist of what the sentence is saying:
- G. W. Bush:
- “Families is where our nation finds hope, where wings take dream.”

- Does this really work? How?

Sentence Comprehension

- We want to build an internal model of the situation
- e.g., “The teacher drank Pepsi in the classroom”
  - Who/what is the agent? teacher
  - What is the patient (object)? Pepsi
  - What did the agent do? drink
  - Where? classroom (and so on…)
- Goal: Teach a model to understand sentences
- Present one word at a time
- Want the model to be able to answer questions
  - demonstrates rudimentary understanding
- e.g., Who is the agent? (has to be able to do this even if agent not currently in input)

(slide < Frank)

SG Toy World

- People: *busdriver* (adult male), *teacher* (adult female), *schoolgirl*, *pitcher* (boy). *adult, child, someone* also used
- Actions: *eat, drink, stir, spread, kiss, give, hit, throw, drive, rise*
- Objects: *Spot* (the dog), *steak, soup, ice cream, crackers, jelly, iced tea, kool aid, spoon, knife, finger, rose, bat* (animal), *bat* (baseball), *ball* (sphere), *ball* (party), *bus, pitcher, fur*
- Locations: *kitchen, living room, shed, park*
- Semantic roles: *agent, action, patient, instrument, co-agent, co-patient, location, adverb, recipient*

(slide based on O'Reilly)

Training

- Sentences generated randomly from semantic and syntactic grammars
- Present words & their roles, one at a time; after each word/role pair, quiz the net on what it has seen up to that point
- The busdriver stirred Kool-Aid
- Present “busdriver” + agent
  - Who is the agent? busdriver
- Present “stirred” + action
  - What is the action? stirred
  - Who is the agent? busdriver
- Present “Kool-Aid” + patient
  - What is the patient? Kool-Aid
  - Who is the agent? busdriver
  - What is the action? stirred

(slide based on Frank)

SG Example/Probe Sentences

- **Active semantic**: *The schoolgirl stirred the kool-aid with a spoon.* (kool-aid can only be the patient, not the agent of this sentence)
- **Active syntactic**: *The busdriver gave the rose to the teacher.* (teacher could be either patient or agent — word order syntax determines it).
- **Passive semantic**: *The jelly was spread by the busdriver with the knife.* (jelly can't be agent, so must be patient)
- **Passive syntactic**: *The teacher was kissed by the busdriver.* vs. *The busdriver kissed the teacher.* (either teacher or busdriver could be agent, syntax alone determines which it is).
- **Word ambiguity**: *The busdriver threw the ball in the park. The teacher threw the ball in the living room.* ("ball" is ambiguous, but semantically, busdriver throws balls in park, while teacher throws balls in living room)

(slide < O'Reilly)

SG Example/Probe Sentences

- **Concept instantiation**: *The teacher kissed someone.* (male). (teacher always kisses a male — has model picked up on this?)
- **Role elaboration**: *The schoolgirl ate crackers.* (with finger); *The schoolgirl ate.* (soup) (these are predominant cases)
- **Online update**: *The child ate soup with daintiness.* vs. *The pitcher ate soup with daintiness.* (schoolgirl usually eats soup, so ambiguous *child* is resolved as schoolgirl in first case after seeing soup, but specific input of *pitcher* in second case prevents this updating).
- **Conflict**: *The adult drank iced-tea in the kitchen.* (living-room) (iced-tea is always had in the living room).

Sentence Gestalt Model

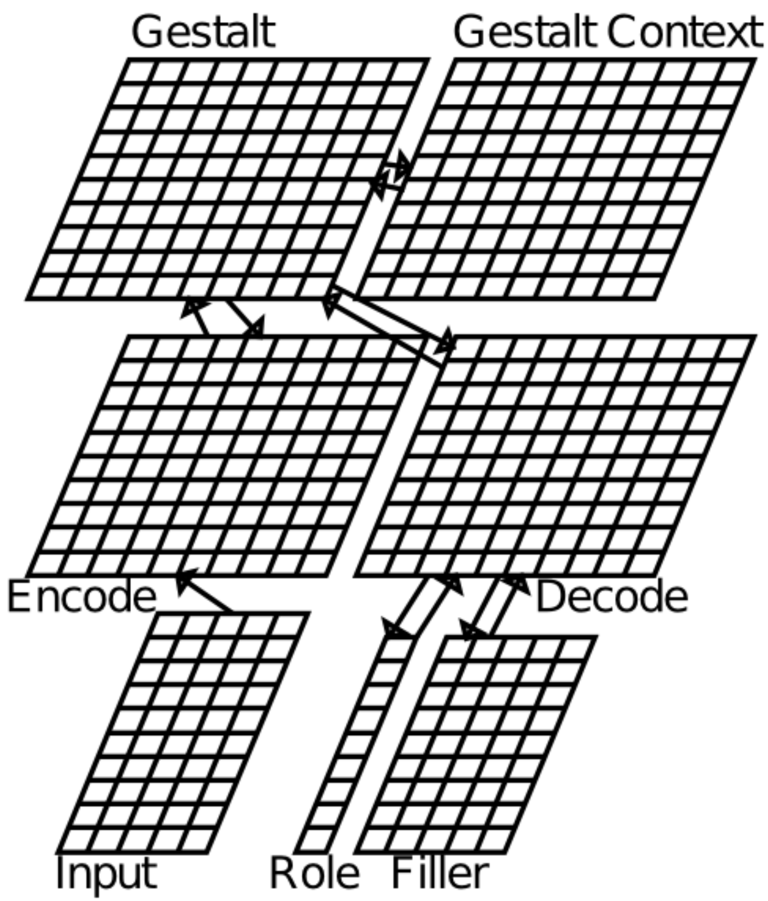

- Single word Input (localist single-unit representations of words)
- Gestalt layer: where distributed representation of sentence meaning develops
- Context: memory for prior words and meaning interpretations of the sentence is encoded
  - copy of Gestalt layer activation state from previous word.
  - simple recurrent network (SRN)
- Role input unit is activated, then network is trained to activate appropriate response in Filler output layer

Gestalt Representations

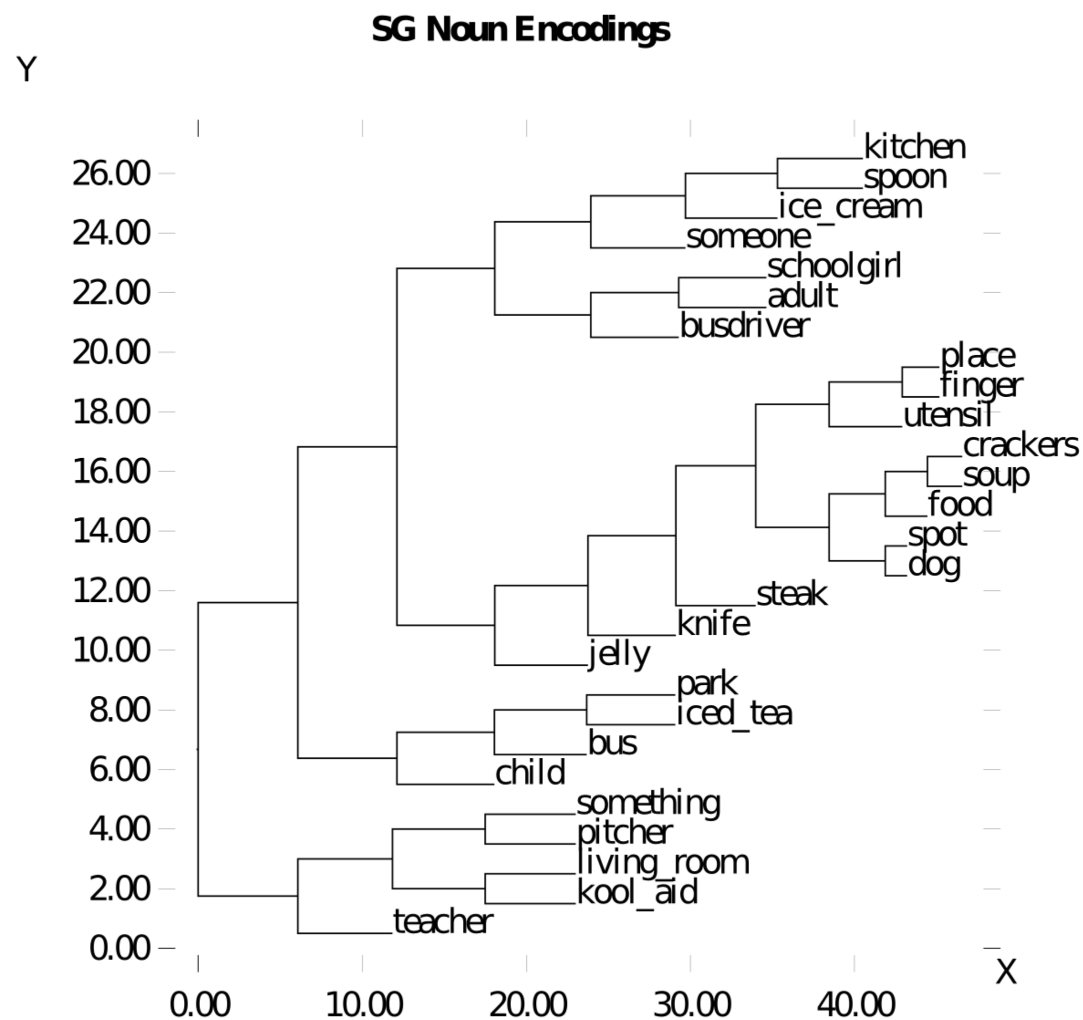

(slide < O'Reilly)

Gestalt Representations

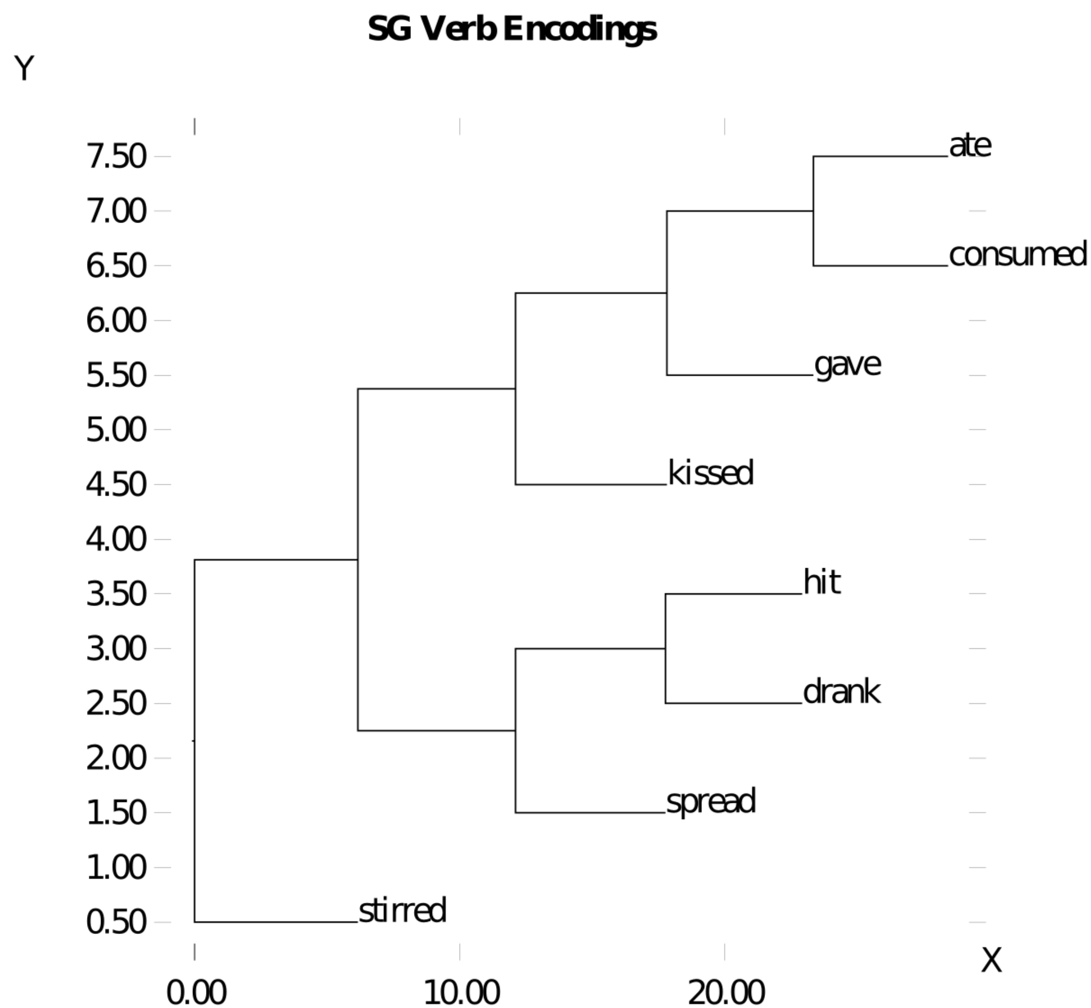

(slide < O'Reilly)

Gestalt Representations

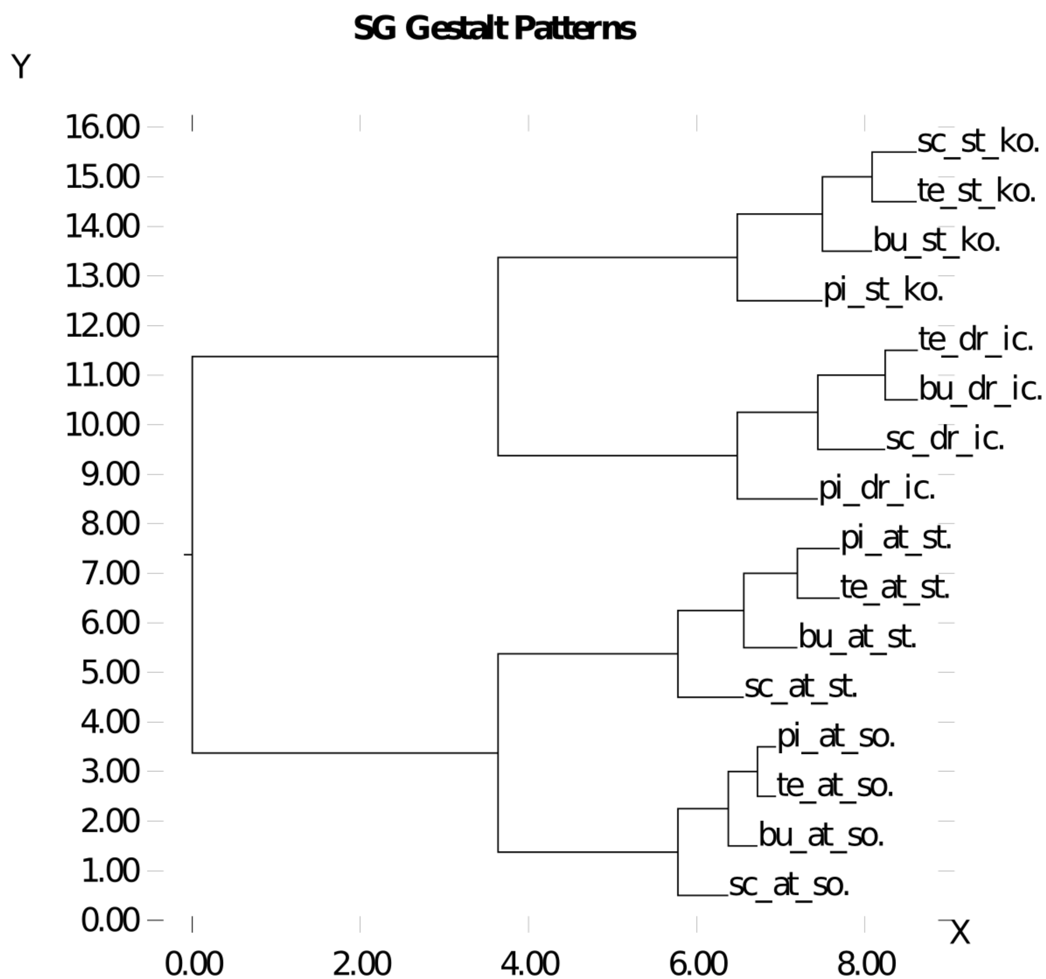

sc = schoolgirl

st = stirred

ko = kool-aid

te = teacher

bu= busdriver

pi = pitcher

dr = drank

ic = iced-tea

at = ate

so = soup

st = steak

Problems with Statistical Approach

- The model makes mistakes for infrequent and/or irregular sentences
- Example: busdriver ate soup; responds with steak as patient
- Explanation: Net saw busdriver eating steak 7× more than soup
- Statistical model overrides reality…
- People suffer from similar biases

Rohde (2002) Model

- Uses structured semantic representations in form of slot-filler propositions
- Includes roles: *agent, experiencer, goal, instrument, patient, source, theme, beneficiary, companion, location, author, possession, subtype, property, if, because, while*, and *although*
- Departs from notion of unstructured gestalt representation of semantic meaning
- Injects externally more of what model should be developing on own
- Challenge: develop a more naturalistic way of training corresponding semantics

Embodied Semantics

- Most of our semantic representations are *grounded* and *embodied*
- The semantic representations used in language understanding are related to those used for:
  - understanding the world
  - controlling our own actions
  - understanding other's intentions and actions
- I.e., we know what it is to be an agent, to be a patient, to be in a place, to eat, to drink, to hit, etc. etc.
- Our linguistic abilities are very much tied to our "form of life"
- To what extent can (even humanoid) robots understand human language?

An Inverted View

- <u>Pragmatics</u> emerges through statistical regularities in our interactions with the world, including other humans
- <u>Semantics</u> (meaning) emerges through statistical regularities in pragmatic interactions (including communication)
- <u>Syntax</u> emerges through statistical regularities in the communication of semantics

emergent Demonstration: sg.proj