Computational Psycholiguistics
Lecture 1: Introduction

Matthew W. Crocker

Marty Mayberry

What is it?

- Using computational techniques to better understand and model how people produce and comprehend language
  - Competence: How do utterances relate to underlying meaning?
  - Performance: How do people establish this relationship during on-line language processing?
- Computational psycholinguistics seeks cognitively plausible theories about about both mental rules and representations, and about cognitive processes
- Computational psycholinguistics seeks to realize such theories as implemented, predictive models of human knowledge and behaviour

Areas of Psycholinguistics

- Speech perception and articulation
- The mental lexicon: how is it represented?
- Lexical access and lexical choice
- Sentence processing: syntactic, semantic, discourse, dialogue
- Situated language processing: interaction of language with task/context
- Embodied language processing:
  - intertwining of language with other cognitive and perceptual systems
  - modeling of these accounts through cognitive robotics

Computational *Psycho*linguistics

- "To understand and model the processes that underlie the human capacity to understand language"
  - How does the human language processor work?
  - How is it realized in the brain?
  - How can we model it computationally?
  - Where does it come from?
  - How does language interact with other cognitive systems and the environment?

Different from NLP?

- Early NLP (e.g. Winograd, 1983) clearly viewed itself as building models of human understanding
  - Proposals were heavily informed by intuitions about how people understand, and linguistic theories about mental representations
- Modern NLP has shifted emphasis:
  - Application: do limited tasks accurately and robustly, often without real understanding (e.g. spam filters, IR, document clustering, summarization)
  - Deep NLU: Emphasis is on representations, coverage and efficiency. Little concern with cognitive plausibility

Human language processing

- People are highly accurate in understanding language
- People process language rapidly, in real-time
- People understand and produce language incrementally
- People rapidly adjust to context, and are robust
- People achieve this despite limitations on processing resources
- People do make some interesting errors, and exhibit breakdown in certain situations ...

Human language processing

- People are highly **accurate** in understanding language
- People process language **rapidly**, in **real-time**
- People understand and produce language **incrementally**
- People rapidly adjust to **context**, and are **robust**
- People achieve this despite **limitations** on processing resources
- People do make some interesting **errors**, and exhibit **breakdown** in certain situations ...

So what ...

- Speech streams include no discrete boundaries to indicate where one word ends and another begins.
- We understand stammering non-fluent politicians and non-native speakers. Incomplete sentences are no problem for us.
- We deal with ambiguity all the time without breaking down. Computer parsers often maintain thousands of possible interpretations.
- We have a vocabulary of about 60,000 words. We access somewhere between 2-4 words/second (error rates around 2/1000 words)
- We understand speech even faster than we can produce it. We are so fast, we can even finish each others sentences.

Sentence processing

- Sentence processing is the means by which the words of an utterance are combined to yield and interpretation
  - All people do it well
  - It is a difficult task: complexity and ambiguity
  - Not simple 'retrieval', like lexical access
  - Compositional: interpretation must be built, rapidly, even for novel word/ structure input

Simple example

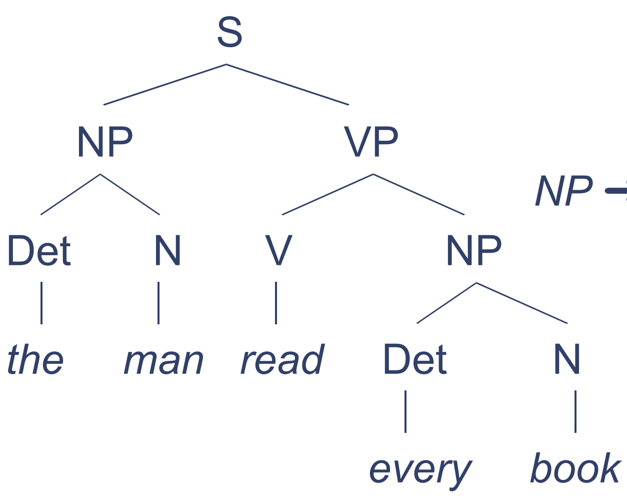

S → NP VP

NP → Det N *VP → V NP*

NP → Det N

Theories of Linguistic Knowledge

- Theories of Syntax
  - **Representations**: Trees, feature structures, dependencies
  - **Structure building**: PS-rules, transformations, unification, composition, tree substitution
  - **Constraints on representations**: Case marking, theta-Criterion, c-command, binding principles, head-foot principle
- Competence Hypothesis
  - The mechanisms of language comprehension directly utilize the rules and representations of the linguistic theory

The Competence Hypothesis

- Knowledge: Competence hypothesis
  - Need to recover the meaning of sentences/utterances
  - Assumptions about (levels of) representations
    - Linguistic theory is isomorphic to human linguistic knowledge
    - Comprehension and production share same knowledge
- Weak competence: people recover *representations* that are isomorphic to those of linguistic theories
- Strong competence: people *directly use* the grammatical knowledge & principles of linguistic theories

The Modularity Issue

- Is language distinct from other cognitive processes?
  - e.g. vision, smell, reasoning ...
- Do distinct modules exist *within* the language processor?
  - e.g. word segmentation, lexical access, syntax ...
- What is a module anyway!?

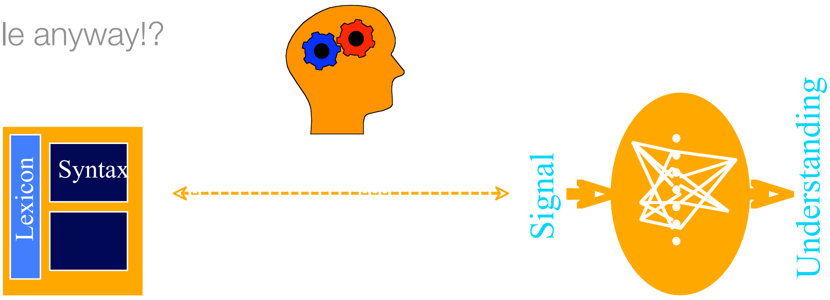

Architectures and Mechanisms

- What does “distinct” mean:
  - Representational autonomy: e.g. phonological versus syntax representations
    - Possibly interactive processes
  - Procedural autonomy: e.g. lexical access versus syntax
    - Possibly shared representations
- How is the language module organized/interact with other systems?
  - Does architecture affect possible mechanisms?
  - Theoretical, computational and empirical arguments concerning modularity?

Modularity and Computation

- The brain is the natural computer, par excellence:
  - Perception occurs in real time, and is highly strategic
- Traditional views on human perception: *Cognitivist* and *Behaviourist*
  - Inferential, unencapsulated: cognitive penetration of perceptual processes
  - Non-inferential, encapsulated: perception reduces to conditioned reflexes
- Fodor: inferential but encapsulated
  - Perception is performed by: *“informationally encapsulated systems which may carry out complex computations”*

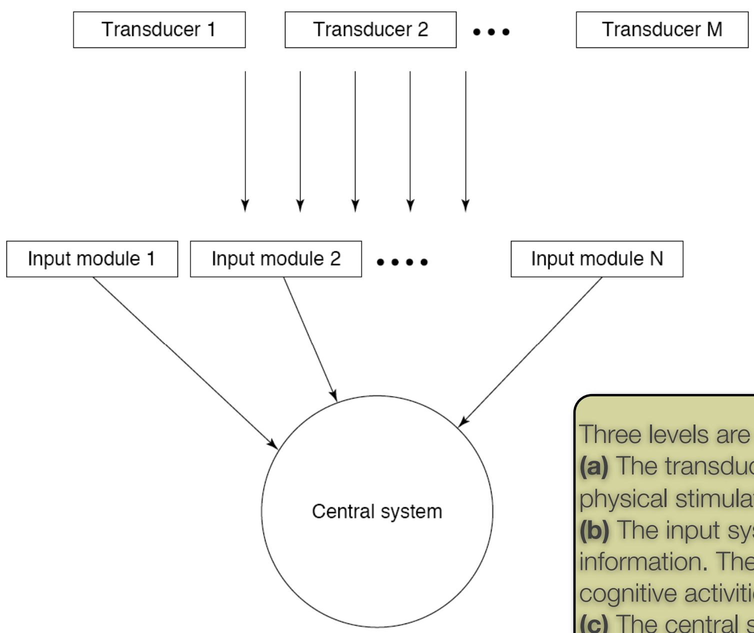

Modules are:
- domain specific
- innately specified
- informationally encapsulated
- fast
- hardwired (neurally specific)
- autonomous
- not assembled

Three levels are distinguished:
(a) The transducers, whose function is to convert physical stimulation into neural signals.
(b) The input systems, interpret transduced information. They are responsible for basic cognitive activities and are modular.
(c) The central system, is responsible for more complex cognitive activities such as analogical reasoning, and is not modular.

Reference
a Fodor, J.A. (1983) *The Modularity of Mind*, MIT Press/Bradford Books

Proving Modularity

- The best proof of Modularity would be evidence for a “Double Dissociation”:
- #1 Damaged linguistic abilities, but intact general cognition
- #2 Damaged cognitive abilities, but intact language

#1 Broca’s aphasia
• normal IQ
• language comprehension is relatively unimpaired
• language production is non-fluent, few words, short sentences, few function words, no intonation

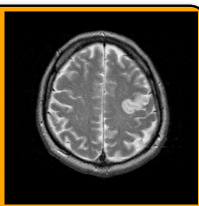

#1 Specific Language Impairment
• normal IQ and hearing
• language is meaningful, appropriate
• problem with grammatical morphemes

#2 Williams Syndrome
(Genetic defect in .001% births)
• low IQ, overly social, poor spatial reasoning
• good language ability, nearly age appropriate

#2 Senile Dementia
• poor memory and diminished general cognitive function
• language production and comprehension remain intact

Strong competence & modularity

- Fodor’s proposals emphasis language as a module, distinct from other perceptual cognitive abilities
- Linguistic theories suggest that language itself may consist of sub-levels: phonology, morphology, syntax, semantics ...
  - Each with different rules and representations
  - Do these correspond to distinct processes?
  - Are these processes modules?
  - Which of Fodors characteristics might they have/not have?

A Modular Architecture

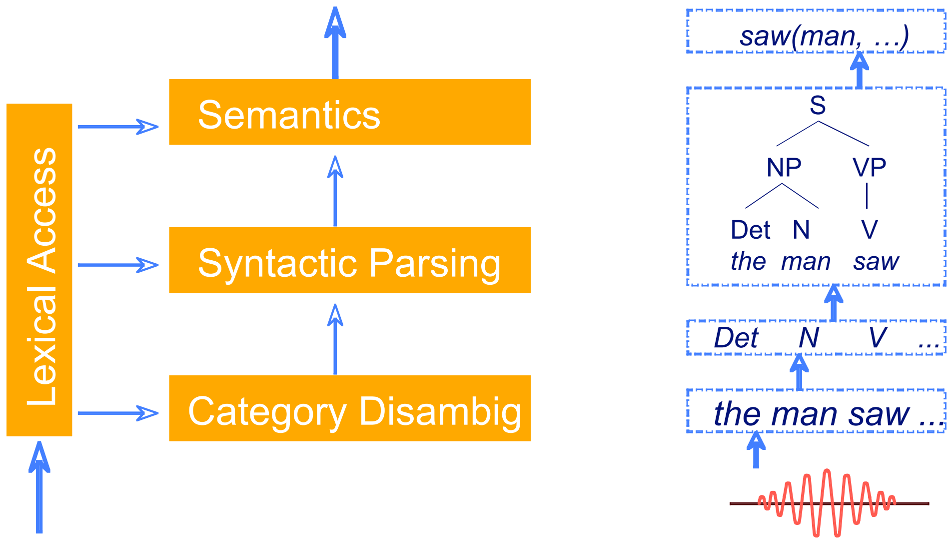

Support for Linguistic Modularity

- Modular lexical access versus syntax: Forster
  - all possible word meanings temporarily available
  - no immediate influence of syntactic context
- Modular syntax versus semantics: Frazier
  - initial attachment ambiguities resolved by purely structural preferences
  - no immediate effect of semantics or context
- Dissociation in language impairment at different levels
  - lexical, syntactic, semantic; production versus comprehension

Attachment Preferences

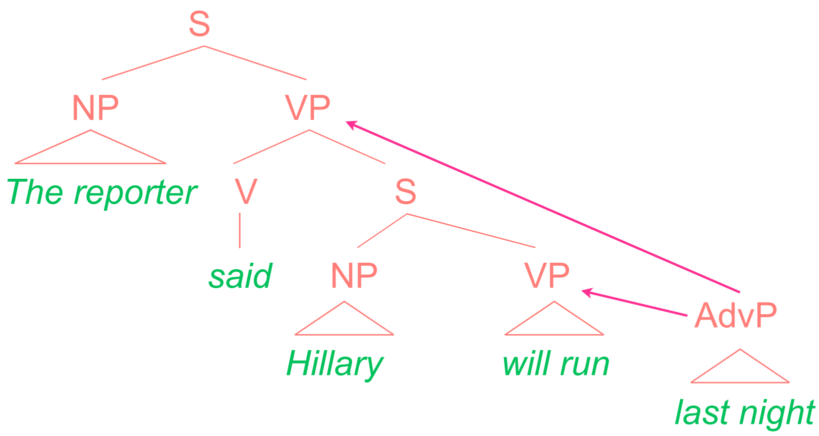

Against linguistic modularity

- Empirical evidence from on-line methods
  - later evidence for “immediate” (very early) interaction effects of animacy, frequency, plausibility, discourse context …
    - *The woman/patient sent the flowers was pleased*
- Appropriate computational frameworks:
  - symbolic constraint-satisfaction systems
  - connectionist systems & competitive activation models
- Homogenous/Integrative Linguistic Theory: HPSG
  - multiple levels of representation within a unified formalism

Human Language Processing

- We understand language incrementally, word-by-word
  - How do people construct interpretations
- We must resolve local and global ambiguity
  - How do people decide upon a particular interpretation
- Decisions are sometimes wrong!
  - What information is used to identify we made a mistake
  - How do we search for an alternative
- Answers can reveal important details about the underlying mechanisms

Experimental Methods

- Reading times
- Neuroscientific methods
- Situated spoken sentence comprehension

The Problem

- How do people recover the meaning of an utterance in real-time?

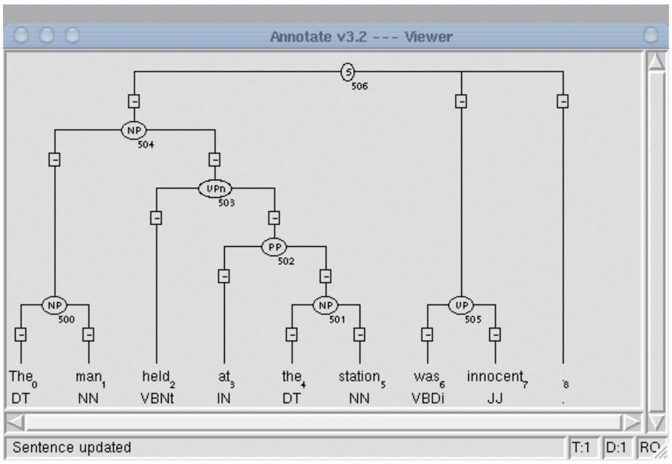

“The man held at the station was innocent”

Crocker & Brants, *Journal of Psycholinguistic Research*, 2000.

Reading time studies

- We can use controlled experiments of reading times to investigate local ambiguity resolution
- (a) The man held at the station was innocent (LA)
- (b) The man who was held at the station was innocent (UA)
- Compare the reading times of (b) where there is no ambiguity, with (a) to see if and when the ambiguity causes reading difficulty.
  - Need a “linking hypothesis” from theory to measures
  - Can then manipulate other linguistic factors to determine their influence on on RTs in a controlled manner

Reading Methods

- Whole sentence reading times:

The man held at the station was innocent

- Self-paced reading, central presentation:

- Self-paced reading, moving window:

The man held at the station was innocent

Eye-tracking: Difference Measures

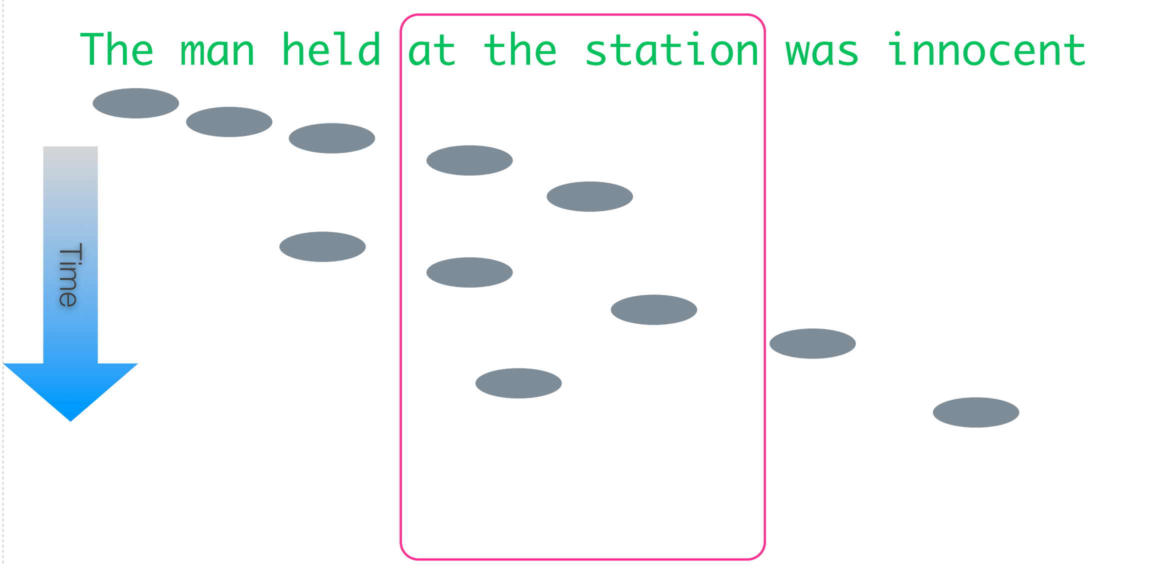

Eye-tracking: First Fixation

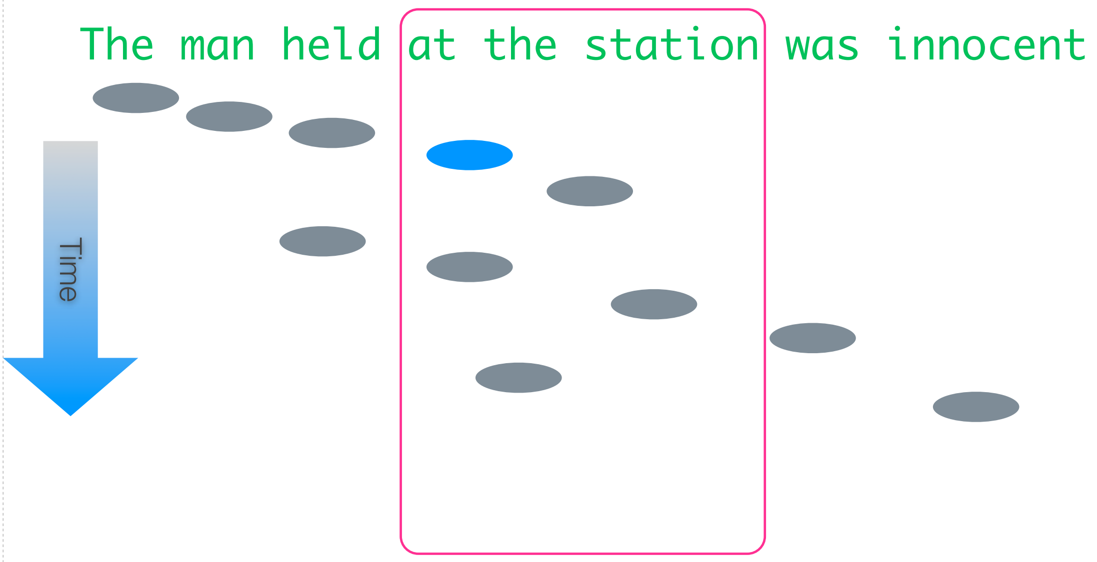

Eye-tracking: First Pass

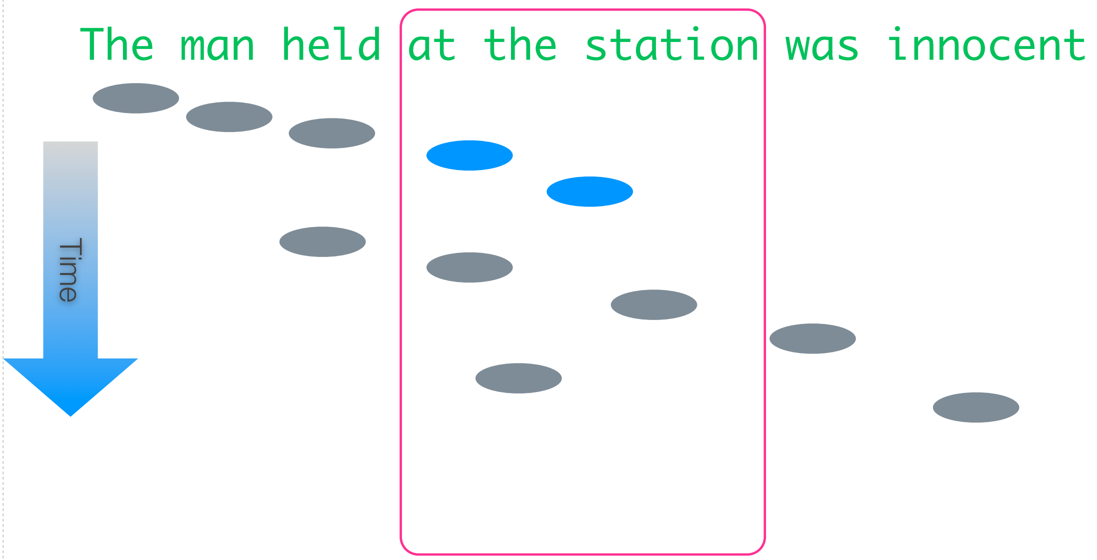

Eye-tracking: Total time

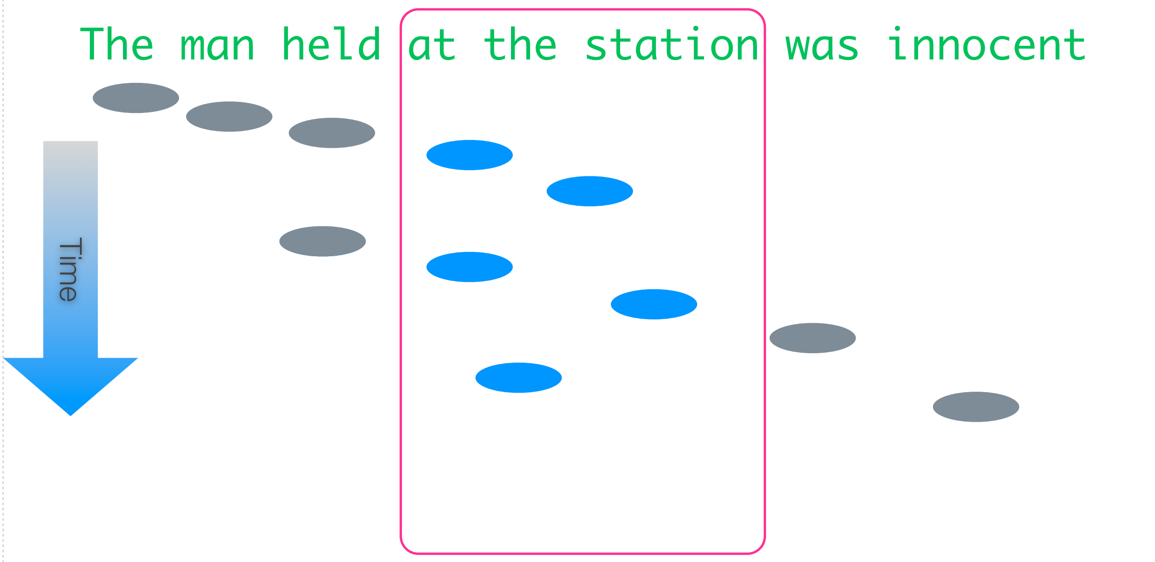

Eye-tracking: Regression Path

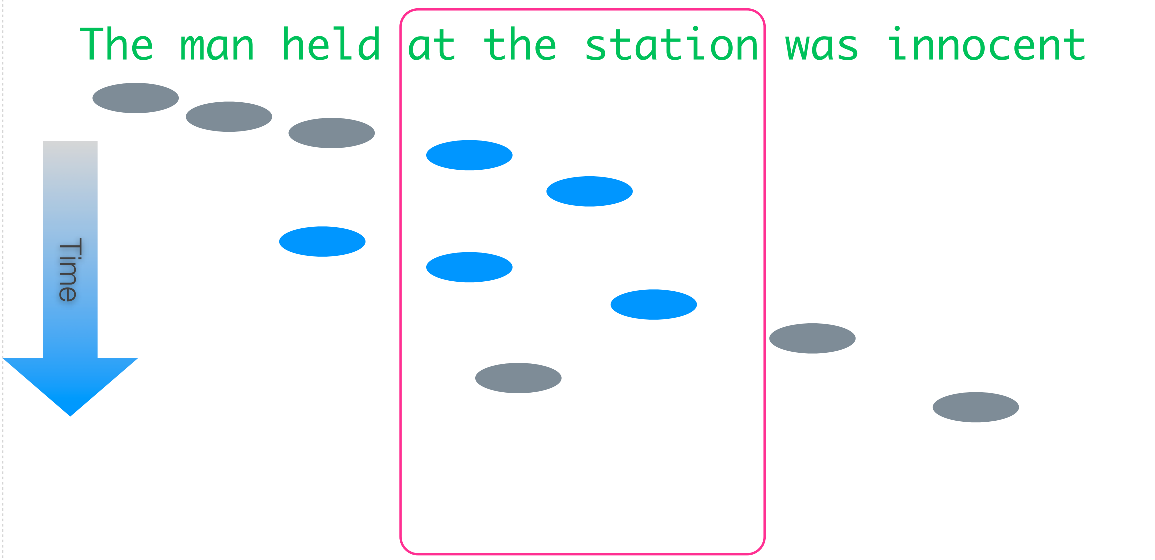

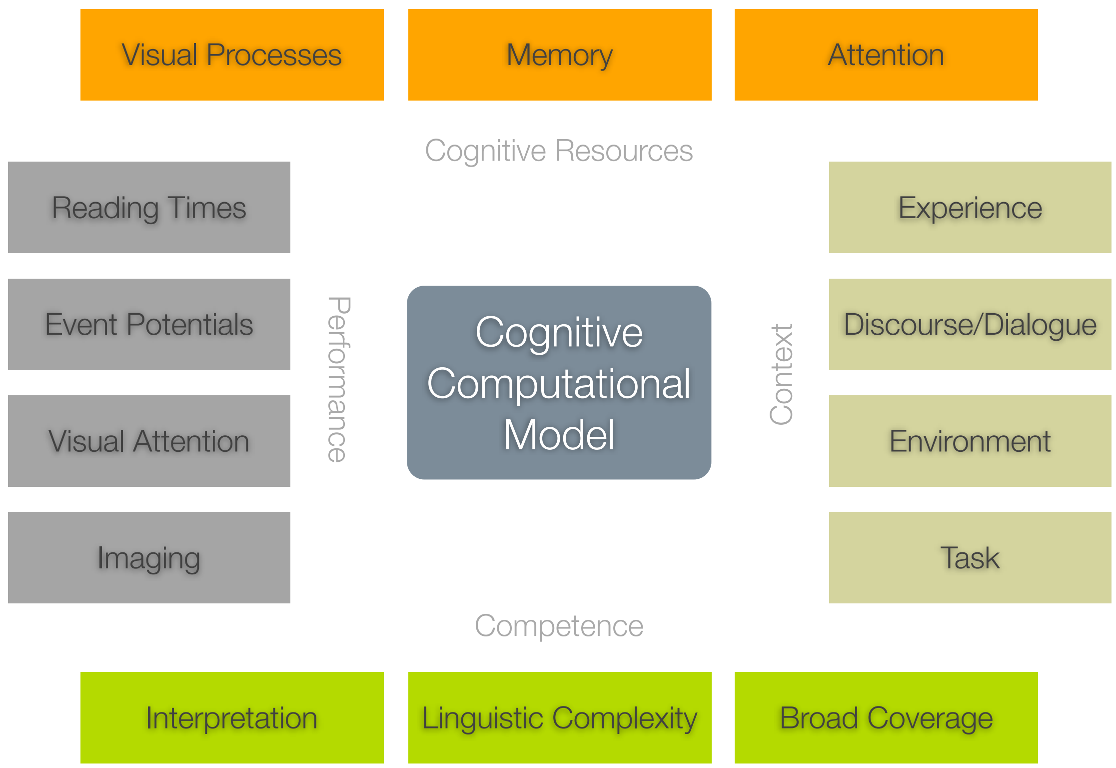

Roadmap

- Theories of sentence processing:
  - modularity, parsing strategies, information sources, reanalysis
- Symbolic parsing models:
  - incremental parsing, ambiguity resolution, memory load, probabilistic models
- Probabilistic parsing models:
  - Symbolic parsers augmented with probabilities, derived from experience
- Connectionist models
  - Distributed models of language learning and language processing

Date	Lecture	Tutorial
	Monday	Wednesday
29.10.07	Lecture 1: Introduction, methods	Tut 0 (Cogent: Experimental Model)
05.11.07	Lecture 2: Syntactic parsing	Tut 1 (Parsing in Cogent)
12.11.07	Lecture 3: Theories of human parsing	Tut 2 (Backtracking/Reanalysis)
19.11.07	Lecture 4: Reanalysis	Tut 3 (LC-Parsing, Mem.load)
26.11.07	Lecture 5: Probabilistic Models I	Tut 4 (Cogent Projects)
03.12.07	Lecture 6: Probabilistic Models II	Tut 5 (Cogent Projects)
10.12.07	Lecture 7: Interactive Models	Lecture 8: Intro to Connectionism
17.12.07	Lecture 9: Learning in Neural Nets	Tut 6 (Group work)
07.01.08	Tut 7 (Group work)	Tut 8 (Proj. Demos)
14.01.08	Lecture 10: Pattern Association	Tut 9 (Using Tlearn)
21.01.08	Lecture 11: Morphology & Phonology	Tut 10 (Reading aloud)
28.01.08	Lecture 12: Simple Recurrent Networks	Tut 11 (English past-tense)
04.02.08	Lecture 13: More on SRNs	Tut 12 (SRNs I)
11.02.08	Lecture 14: Advanced architectures	Tut 13 (SRNs II)
18.02.08	Exam	

Matthew Crocker	crocker@coli.uni-sb.de
Marshall Mayberry	martym@coli.uni-sb,de
Garance Paris	gparis@coli.uni-sb.de

Course details

- Weekly lectures (Monday 2-4pm) and tutorials (Wednesday 2-4pm)
  - Participation and completion of all tutorials is expected!
- Assessment: Final Exam (100%), date: Monday 18 February, 2008
  - All tutorial assignments must be successfully completed to sit the exam
- **Course materials (overheads and most readings) will be made available on the course homepage (linked from general course page)**
- Contact: please e-mail first!
  - crocker@coli.uni-sb.de or martym@coli.uni-sb.de or gparis@coli.uni-sb.de