Introduction to Cognitive Science:

Linking the World and the Mind

Mark Steedman⋆

`http://www.inf.ed.ac.uk/~steedman`

Not only speech, but all skilled acts seem to involve the same problems of serial ordering, even down to the temporal coordination of muscular contractions in such a movement as reaching and grasping. Analysis of the nervous mechanisms underlying order in the more primitive acts may contribute ultimately to the solution of even the physiology of logic.

Karl Lashley 1951:122

School of Informatics, University of Edinburgh, September 21, 2009

Where Cognitive Science came from

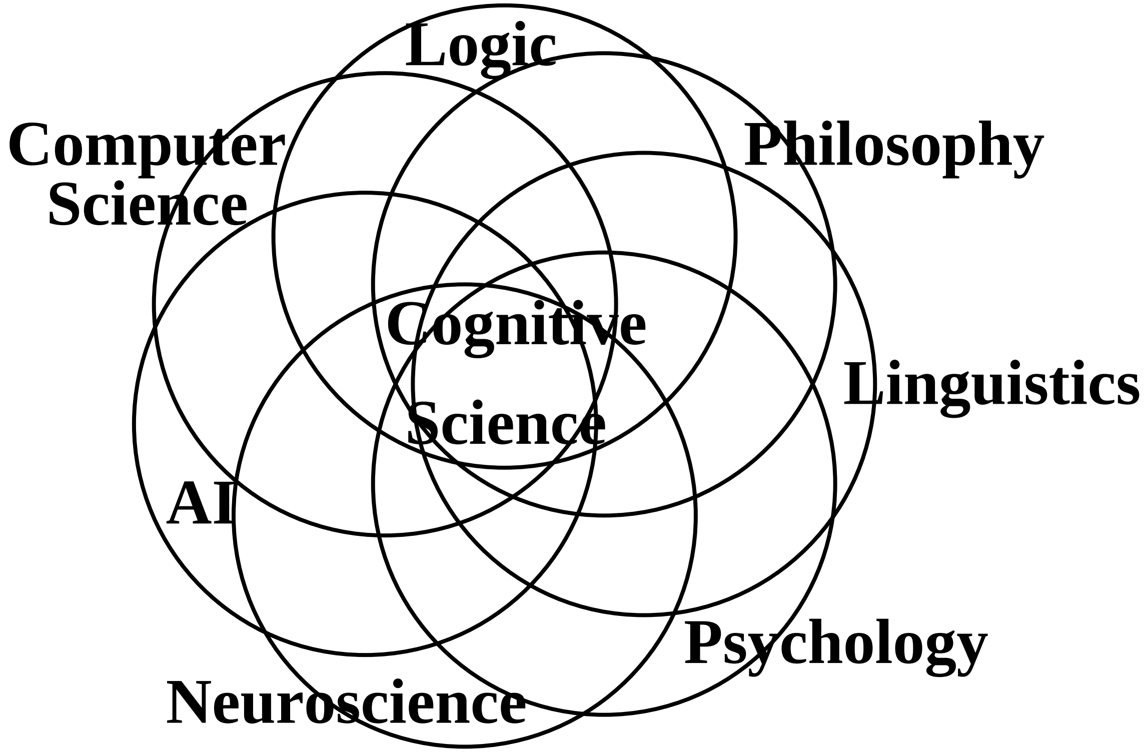

Figure 1: Contributions to Cognitive Science

What Cognitive Science is About

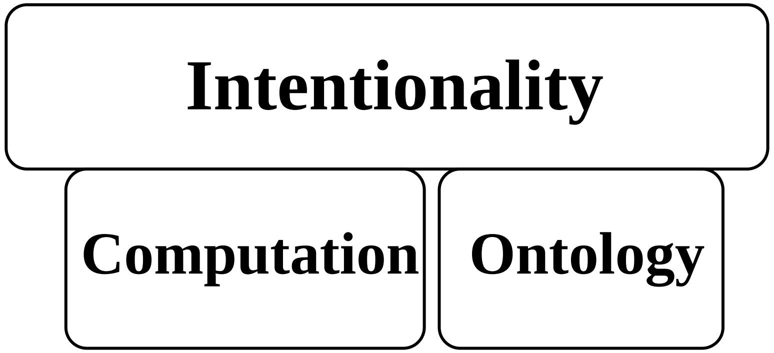

Figure 2: Cognitive Science as the Theory of "Intentionality"

- Intentionality is the property of systems whose actions are contingent on internal representations as well as the external world. Thermostats are not intentional, but bats are, and so are we.
- Computation is manipulation of symbols.
- Ontology is concerned with what those symbols represent.

The Problem

- The problem is that Ontology and Computation are determined by the nature of our being, and interaction with the external world.
- Much of their character has been shaped by evolution over hundreds of millions of years.
- Evolution is capable of coming up with extremely improbable solutions by exhaustive search of vast spaces, consuming resources on a planetary scale. Example: the genetic code (Freeland and Hurst 1998, Morris 2004).
- Cognitive Science tries to work backwards from the detailed structure of behaviour to the hidden and often extremely bizarre ontologies that underlie it, and forwards from neural mechanisms to computations.
- This course focuses on the necessity of purposeful Action in the world as a source of constraints on computation and ontology in intentional systems

Structure of The Course

- I: Mapping the World into the Brain (Stereo Vision).
- II: Representing the World Symbolically (Scene Analysis).
- III: Representing Action in the World Computationally (Planning)
- IV: How Animals and Humans Actually make Plans
- V Neurological and Developmental Substrate of Planning and Language
- VI: How Universal Grammar is Transparent to Prelinguistic Planning.
- VII: How Semantics is Transparent to Prelinguistic Planning
- VIII: How Discourse is Transparent to Prelinguistic Planning
- IX: Human and Computational Natural Language Processing
- X: Human and Computational Language Development
- XI: Envoi: Scope and Limits of Cognitive Science

Introduction to Cognitive Science: Notes

I: Mapping the World into the Brain (Stereo Vision)

- Readings for this section: *Marr (1977); *Marr and Poggio 1976

Two Kinds of Problem in AI: Marr 1977

- “Type 1 theories” decompose Computation into a “Theory” of the Computation and the “Algorithm(s)” that perform the computation. “Type 2 theories” are not decomposable in this way: the system under analysis is in some sense its own simplest description.
- Fourier Analysis is Type 1. Predicting the weather is probably Type 2.
- The linguists’ distinction between Competence and Performance is a distinction between Theory and Algorithm.
- The Type 1/Type 2 distinction is the same as that between “Modular” and “Non-modular” theories.
- A system such as Language use may be Type 1 at some high level (grammar), and Type 2 at some lower level (lexical meaning).
- The involvement of Type 2 systems may place severe limits on our ability to capture intelligence by machine.

Why is Artificial Intelligence so Slow to Arrive?

- Ontologies are embedded, non-modular, and grounded. (Simple example: bat auditory cortex. Complex example: natural language semantics)
- The computation that works with them are distributed, probabilistic and probably rather simple.
- Our knowledge representations (Logics) are not like that.
- Cognitive Science tries to help with this, by asking:
- How do we get from the world to such representations?
- In particular, how do we get from retina-like maps to symbolic representations?
- How do we use the latter to act in the world?
- How does this shape thought and language?

Mapping the World into the Brain

- Maps are ubiquitous in parts of the brain concerned with perception and action.
- For example, patterns of activation in the visual striate cortex at the back of the head are isotopic to the image on the retina.
- The motor and somatosensory cortex are maps of the surface of the body (Figure below from Currie 2000):
- This is not simply an artefact of the fact that they are projections of a planar image: auditory cortex in bats is also a map of the surrounding space.
- There must be a computational advantage to isotopy.

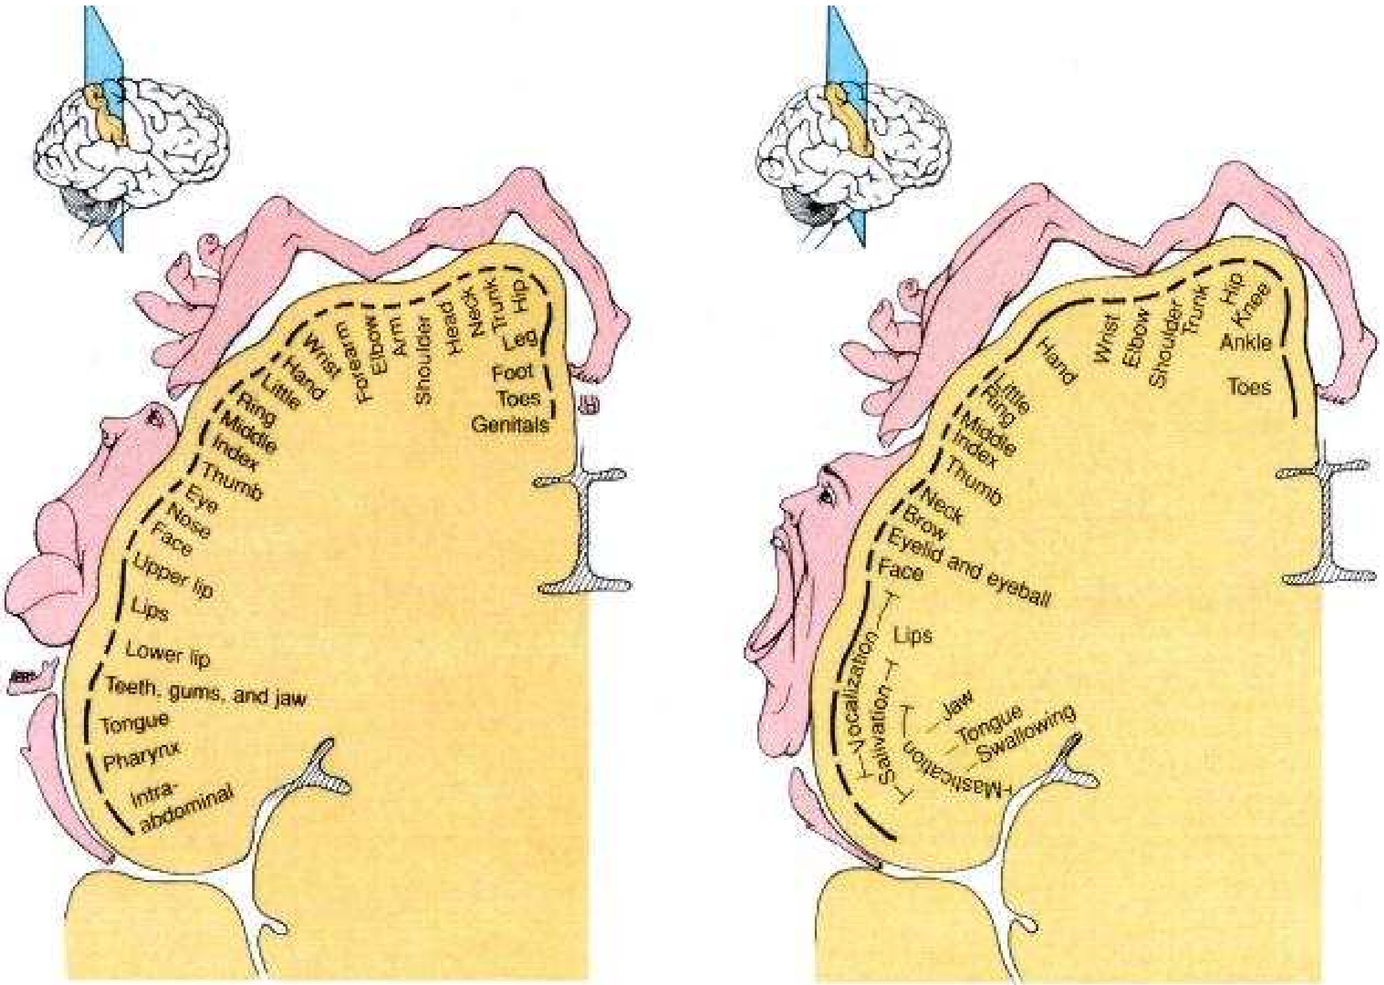

(a) Somatosensory cortex in right cerebral hemisphere

(b) Motor cortex in right cerebral hemisphere

Sensory-Motor Homunculi (from D. Currie 2000).

Marr and Poggio (1976)

- Marr and Poggio (1976) present a theory of stereoscopic binocular vision which shows how three-dimensional maps can be exploited computationally.
- Remarkably, stereoscopy is independent of prior object identification.
- This was shown by Julesz in the early sixties using random dot stereograms, in which a region of random dots is undetectably shifted sideways in one image:

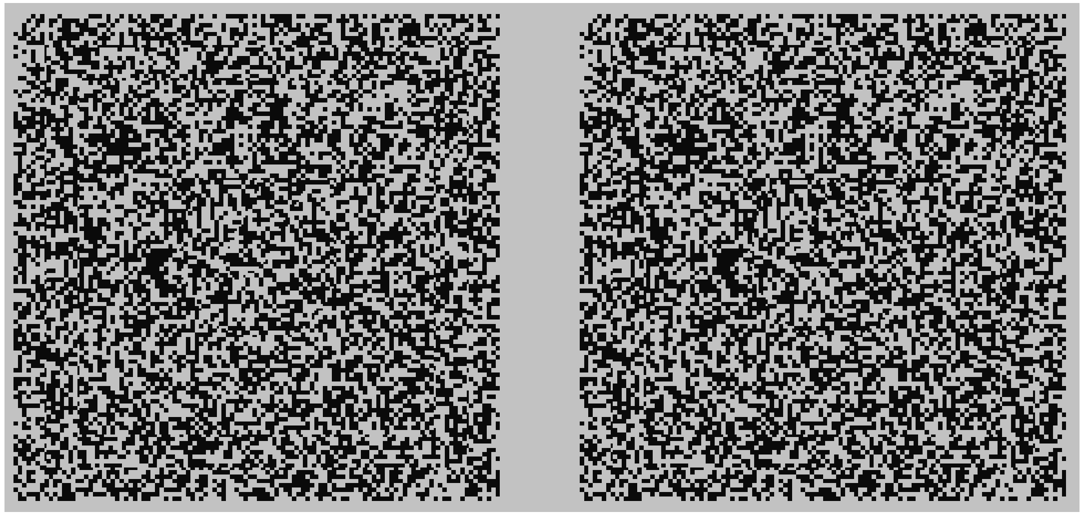

- To view such a stereogram, cross your eyes sufficiently to create a third central combined image and wait for a 3D image to emerge.

Stereo Vision

- Marr and Poggio specified a cellular automaton that computes depth from random dot stereograms
- It exploits two constraints on real-world images:
  - *Uniqueness*: a point along any line of sight from a given eye can only correspond to an object at *one* depth.
  - *Continuity*: adjacent points in the image tend to belong to the same region at the *same* depth.

Marr and Poggio

- The automaton is constructed from a three-dimensional array of computational units corresponding to a discretized version of the three dimensional space in the field of view.
  - *Uniqueness* is modeled by making *inhibitory* connections between units corresponding to points along the same line of sight for each of the two eyes.
  - *Continuity* is modeled by making *exitatory* connections between units corresponding to adjacent points in the same depth plane.
- The cellular automaton itself is complex, with several parameters defining a neighborhood of influence and the relative importance of inhibition and excitation.

Marr and Poggio

- The next slide shows a horizontal slice across a tiny (and unrealistically coarse-grained) fragment of the cellular automaton and a pair of random dot stereograms. (The angle of convergence of the two eyes is greatly exaggerated.)
- units with dots in them represent possible fusions of points in the stereogram.
- There are a lot of them and almost all of them are false fusions.
- The activity of units increases in proportion to the activity of their neighbours in the depth plane, and decreases in proportion to the activity of their neighbours along the line of sight.
- The algorithm makes large sets of adjacent fusion points on the same depth plane grow in activation while isolated fusion points diminish
- After a while, only units corresponding to the position and depth of the depicted objects are active.

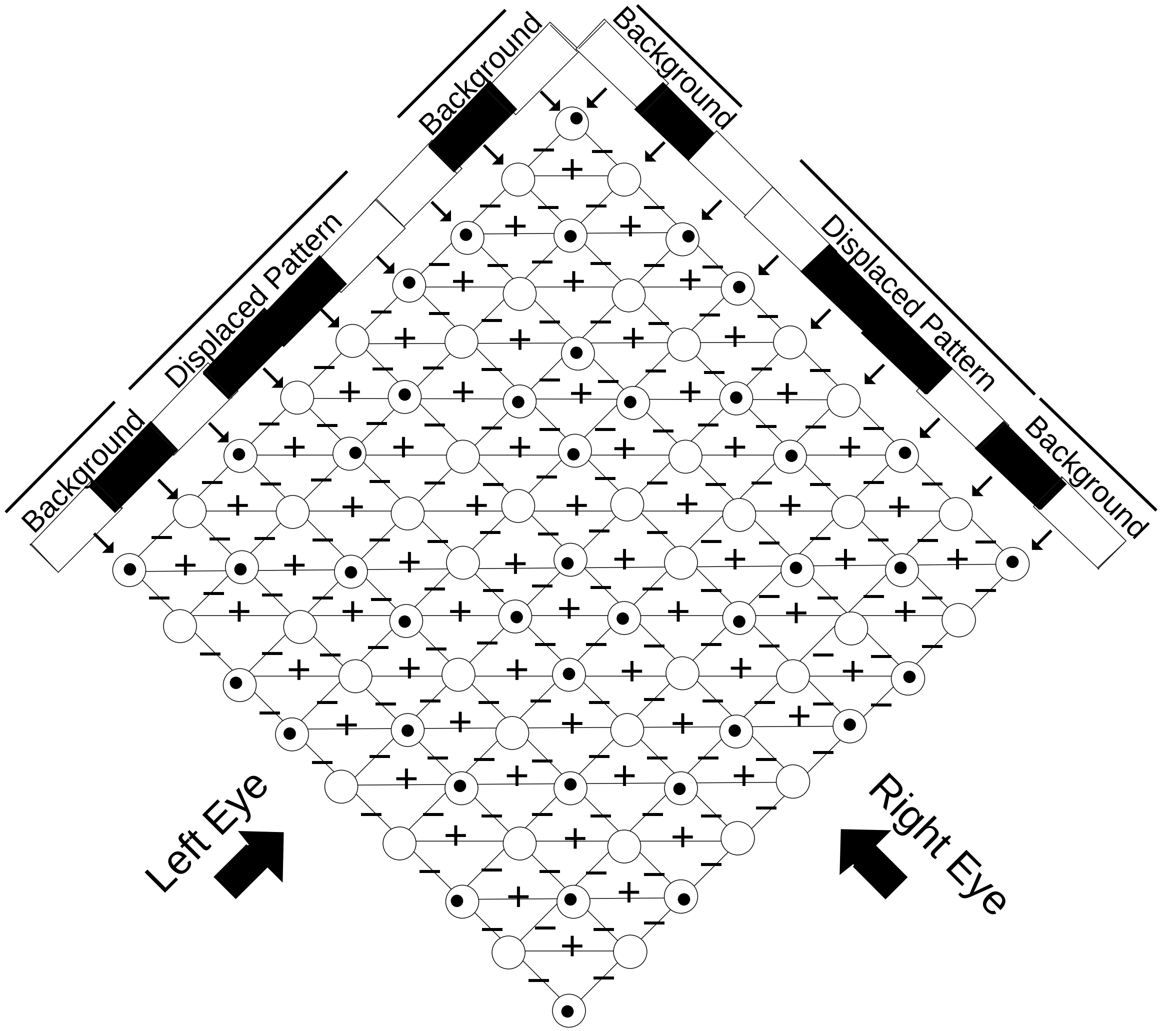
Background
Displaced Pattern
Background
Background
Displaced Pattern
Background
Left Eye
Right Eye

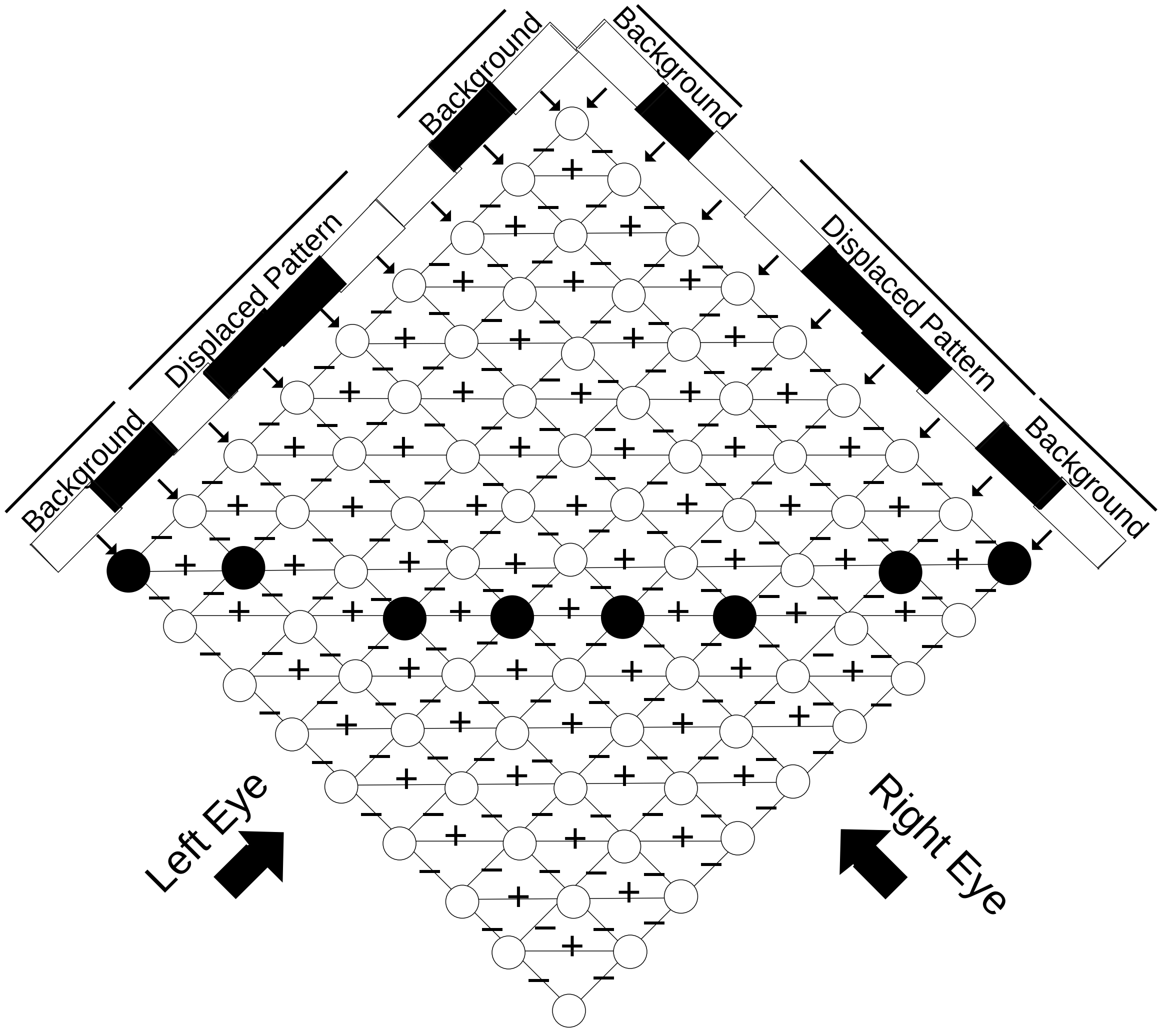
Background
Displaced Pattern
Background
Background
Displaced Pattern
Background
Left Eye
Right Eye

Marr and Poggio

- There is evidence from single cell recording in cat visual cortex of columns of units corresponding to different binocular disparities along a single line of sight.
- Such columns resemble the inhibitory lines in the figures.
- However, the corresponding receptive fields on the cat retina are more complex than single dots in a random dot stereogram. In particular, they are *orientation-sensitive*, corresponding to directional units at the level of the retinal ganglia or higher.

References

Freeland, Stephen and Hurst, Laurence, 1998. "The Genetic Code is One in a Million." *Journal of Molecular Evolutiion* 47:238–248.

Lashley, Karl, 1951. "The Problem of Serial Order in Behavior." In L.A. Jeffress (ed.), *Cerebral Mechanisms in Behavior*, New York: Wiley. 112–136. reprinted in Saporta (1961).

Marr, David, 1977. "Artificial Intelligence: A Personal View." *Artificial Intelligence* 9:37–48.

Marr, David and Poggio, Tomasso, 1976. "Cooperative Computation of Stereo Disparity." *Science* 194:283–287. Reprinted in Vaina 1991.

Morris, Simon Conway, 2004. *Life's Solution: Inevitable Humans in a Lonely Universe*. Cambridge: Cambridge University Press.

Saporta, Sol (ed.), 1961. *Psycholinguistics: A Book of Readings*. New York: Holt Rinehart Winston.

Vaina, Lucia (ed.), 1991. *From Retina to Neocortex: Selected Papers of David Marr*. Boston, MA: Birkhauser.