Levels of Analysis and ACT-R

LaLoCo, Fall 2013

Adrian Brasoveanu, Karl DeVries

[based on slides by Sharon Goldwater & Frank Keller]

David Marr: levels of analysis

- Background
- Levels of Analysis

John R. Anderson: ACT-R

- Background
- ACT-R: a cognitive architecture

Readings: Anderson, 1996; Marr, 1977, 1982.

Outline

Two foundational people in cognitive modeling and their ideas:

- **David Marr** (1945–1980), who introduced the idea of different levels of analysis for information processing systems. These levels provide a framework for thinking about cognitive models.
- **John R. Anderson** (1947–), who developed the widely known cognitive architecture, ACT-R. We will compare its high-level commitments with those of Cogent, and we'll return to a different idea of Anderson's, rational analysis, later on.

David Marr (1945-1980)

- Worked in MIT's AI Lab and Department of Psychology
- A founder of Cognitive Neuroscience.
  - First paper (1969) proposed theory of cerebellar function which is still relevant today (Strata, 2009).
- Developed influential computational theory of vision, treating computational results on par with neurobiological findings.
  - *Vision: A computational investigation into the human representation and Processing of Visual Information*, MIT Press, 1982.

Three Levels of Analysis

Marr suggested that the solution to a complex information processing (IP) problem often divides naturally into three parts:

- A characterization of the problem as a particular type of computation based in the physical world — i.e., an abstract formulation of what is being computed and why.
- A choice of algorithms for implementing the computation, including necessary I/O and internal representations — i.e., an abstract formulation of how the computation is carried out.
- A commitment to particular hardware in which the algorithm is implemented and physically realized — i.e., a concrete formulation of how the computation is carried out.

Problems that decompose this way have, in Marr's terms, a Type 1 theory.

Non-cognitive example: Cash Register

Figure: http://www.springdaleark.org/shiloh/image_archive/cash_register.jpg

Non-cognitive example: Cash Register

- **Computational level**: what does it do, and why?
  - Computes sum of inputs using theory of addition;
  - Need correct total of money owed for goods.
- **Algorithmic level**: what is the representation and algorithm?
  - Arabic numerals
  - Add least significant digits first, carry remainder to more significant digit, add, carry, etc.
- **Implementation level**: what is the physical realization?
  - 1880's: Mechanical device using a crank and rotary wheels
  - Later: Electromechanical, electricity replacing manually-operated crank
  - Modern adding machine: Electronic

Cognitive Example: Bats hunting for prey

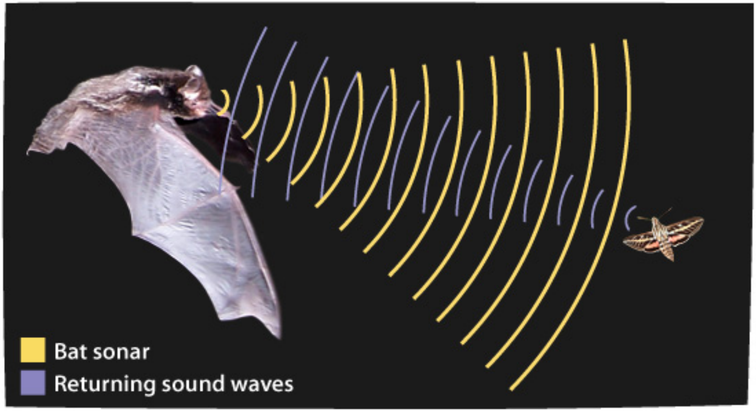

Figure: http://askabiologist.asu.edu/sites/default/files/echolocation.jpg

Cognitive Example: Bats hunting for prey

Bats use echolocation to find food (insects, fruit, nectar).

- **Computational level:** what does it do, and why?
  - Computes distance, motion, and location of objects. (Could be more specific using mathematical equations).
  - Bats hunt at night, so can't rely on vision.
- **Algorithmic level:** what is the representation and algorithm?
  - I/O: delay and Doppler shift between bat calls and returning echo.
  - Algorithm for object recognition?? Active area of research; robots use Kalman filters, artificial neural networks, etc.
- **Implementation level:** what is the physical realization?
  - Bats: neural mechanisms
  - Other: silicon chips; mechanical device??

But...

Biological-based IP problems (such as those posed by cognition) need not have a decomposable (Type 1) theory.

Instead, when a problem is solved by the simultaneous interaction of multiple processes, the interaction may be its own simplest description. Marr refers to this as a Type 2 theory.

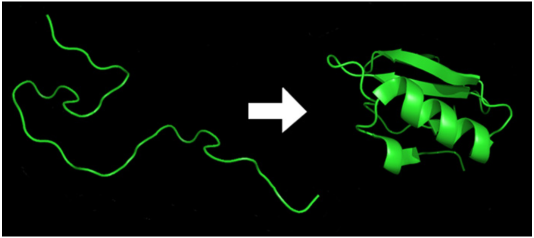

www.math.uwaterloo.ca/AM_Dept/prospective/media/p_fold.jpg

Type 1 and Type 2 Theories

Even an Information Processing problem which has a Type 1 theory may be tied to an IP problem where only a Type 2 suffices.

Example: Language processing may be Type 1 for grammar but Type 2 for word meaning.

Main challenge: determine which problems have Type 1 theories, in part by trying to discover computational-level descriptions of them.

- Marr argues that computational level yields greater insight.
- Some researchers disagree, preferring to work at algorithmic level (either because models are more satisfying, or more practical).

John R Anderson (1947–)

- Professor of Psychology at CMU since 1978
- Early pioneer of work on intelligent tutoring
- Influential work on Cognitive Architectures – *Adaptive Control of Thought* (ACT, ACT*)
- Introduced framework for rational analysis (Anderson, 1990)
- Incorporated into ACT-R (*Adaptive Control of Thought – Rational*), a hybrid Cognitive Architecture in widespread use.

ACT-R Overview

- Unified theory of cognition realized as a production system (a type of cognitive architecture model; similar to Cogent).
- Intended as single model to capture all aspects of cognitive processing.
- originally implemented in LISP; a Python reimplementation available here:

  `https://sites.google.com/site/pythonactr/`
- Example: Learning mathematics involves
  - Reading (both visual processing and language processing)
  - Understanding mathematical expressions
  - Problem solving
  - Reasoning and skill acquisition
- All would be modelled in ACT-R.

Example: Salvucci and Macuga, 2001

Want to predict effects of dialing mobile phone while driving.

- Develop two separate ACT-R models for driving and dialing mobile phone.
- Put them together to predict effects of driving on phone dialling and vice versa.
- Compared four ways of dialing.
- Predicted that only full manual dialing has significant impact on steering abilities.
- Predictions borne out through later experiments.

N.B. Real distraction is talking on mobile phone!

Other domains

Hundreds of papers on ACT-R site (`http://act-r.psy.cmu.edu/`):

- Perception and attention: visual search, eye movements, task switching, driving behavior, situational awareness.
- Learning and memory: list memory, implicit learning, skill acquisition, category learning, arithmetic, learning by exploration and example.
- Problem-solving and decision-making: use and design of artifacts, spatial reasoning, game playing, insight and scientific discovery.
- Language processing: parsing, analogy and metaphor, learning, sentence memory, communication and negotiation.
- Other: cognitive development, individual differences

Fundamental Assumption

Cognition emerges from the interaction between very many small bits of two different types of knowledge (procedural and declarative), stored in corresponding parts of memory.

- Declarative knowledge: facts, things remembered or perceived, goals; represented in ACT-R as chunks (really: feature structures / AVMs)
  - 2+2 = 4
  - Edinburgh is the capital of Scotland.
  - There is a car to my right.
  - I'm trying to get to class.
- Procedural knowledge: processes and skills (represented in ACT-R as production rules)

Procedural Knowledge

- Production rule consists of conditions and actions:

  IF goal is to add two digits d_1 and d_2 in a column and $d_1+d_2=d_3$
  THEN set a subgoal to write d_3 in the column

- In ACT-R, conditions may depend on declarative knowledge, buffer contents, and/or sensory input.
- Actions can change declarative knowledge, goals, or buffer contents, or initiate motor actions.

Modular organization

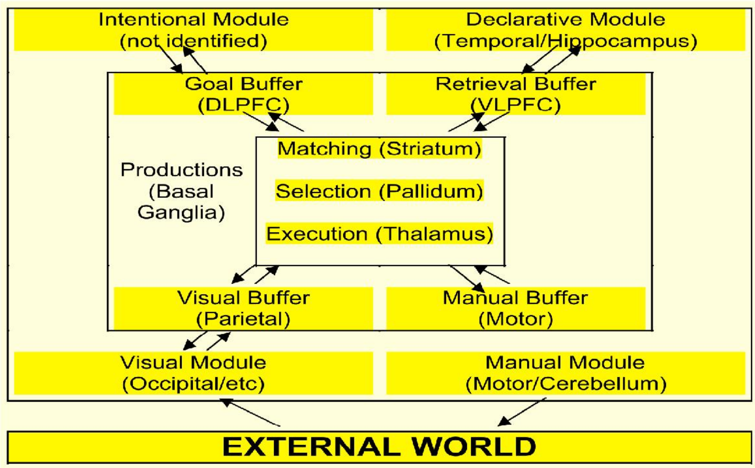

(Anderson et al., 2004)

Modular organization

- Modules: store and process long-term information, which is then deposited in buffers.
  - Goal buffer: tracks state in solving problems.
  - Retrieval buffer: holds information retrieved from long-term declarative memory.
  - Visual buffer: tracks visual objects and their identities.
  - Manual buffer: control and sensation of hands.
- Central production system: executive control and coordination of modules.
  - Not sensitive to activity in modules, only to buffer contents.

Timing and coordination

- Within modules, processing is in parallel.
  - Ex: visual system processes entire visual field at once.
- Overall timing determined by serial processing in central production system. In one critical cycle:
  - Patterns in buffers are recognized and a production fires.
  - Buffers are updated for the next cycle.
- Assumptions:
  - Each buffer may contain only one chunk.
  - Only a single production fires each cycle.
  - Cycle takes about 50 ms (based on experimental data).

Hybrid Architecture

- Behavior determined by interaction between symbolic and sub-symbolic (statistical) systems.
  - Symbolic: production system.
  - Sub-symbolic: massively parallel processes summarized by mathematical equations.
- Each symbol (production/chunk) has sub-symbolic parameters that reflect past use and determine probability of current use.
- Inclusion of sub-symbolic activation levels is a major difference to Cogent.

Example 1: Declarative memory module

- Purpose: retrieve chunks formed previously.
- Each chunk has a sub-symbolic activation level, the sum of
  - Base level activation, reflecting general usefulness in past.
  - Associative activation, reflecting relevance to current context.
- Total activation determines probability of being retrieved and speed of retrieval.

Example 2: Procedural memory

- Many production rules may match at once but only one can fire.
- Each rule has a sub-symbolic utility function combining
  - The probability that the current goal will be achieved if this rule is chosen (based on past experience).
  - The relative cost (time/effort) and benefit of achieving the current goal.
- The rule with the highest utility is executed.

ACT-R Summary

- Complex cognition emerges as the result of (procedural) production rules operating on (declarative) chunks.
- Independent modules encapsulate parallel processing functions, place single chunks in buffers.
- Central production system accesses buffers, detects when rule triggers are satisfied, fires one rule at a time.
- Chunks and rules are symbolic, but sub-symbolic activation levels determine which ones get used.
- Learning involves either acquiring new chunks and productions, or tuning their sub-symbolic parameters.

ACT-R features

- Can predict time-sharing between different tasks.
- Bridges short time-scale processes (retrieval, single productions) with long time-scale processes (e.g., learning to solve algebraic equations), with implications for education.
- Some evidence that modular structure corresponds to different brain regions.

References I

Anderson, John R. (1996). “ACT: A Simple Theory of Complex Cognition”. In: *American Psychologist* 51.4, pp. 355–365.

Anderson, John R. et al. (2004). “An integrated theory of the mind”. In: *Psychological Review* 111.4, pp. 1036–1060.

Marr, David (1977). “Artificial Intelligence: A personal view”. In: *Artificial Intelligence* 9, pp. 37–48.

— (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: W. H. Freeman.

Salvucci, D. D. and K. L. Macuga (2001). “Predicting the Effects of Cell-phone Daialing on Driver Performance”. In: *Procedings of the 4th International Conference on Cognitive Modeling*. Ed. by E. Altmann et al. Mahwah, NJ: Lawrence Erlbaum Associates, pp. 25–30.

Strata, P. (2009). “David Marr’s theory of cerebellar learning: 40 years later”. In: *The Journal of Physiology* 587.23, p. 5519.