

## EVIDENCE OF MODULARITY FROM PRIMATE ERRORS DURING TASK LEARNING

JOANNA J. BRYSON

*University of Bath, Department of Computer Science*

*Bath, BA2 7AY; United Kingdom*

*E-mail: j.j.bryson@bath.ac.uk*

The last two decades have seen a great deal of theorising and speculation about the modular nature of human intelligence, as well as a rise in use of modular architectures in artificial intelligence. Nevertheless, whether such models of natural intelligence are well supported is still an issue of debate. In this paper, I propose that the most important criteria for modularity is specialised representations. I present a modular model of primate learning of the transitive inference task, and propose an extension to this model which would explain task-learning results in other domains. I also briefly relate this work to both neuroscience and established AI learning architectures.

### 1. Introduction — Modularity and its Criteria

The 1980s saw the publication of two books which drew attention to modular theories for the organisation of intelligence. The first, *The Modularity of Mind* (Fodor, 1983) came from psychology and dwelt primarily on specialised modules for perception. Fodor's full theory proposed that a central processor analyses the output of these modules and chooses some action, then further modules generate expressed behavior from these decisions. The second book, *The Society of Mind* (Minsky, 1985) came from artificial intelligence (AI). It described a system where control is divided and negotiated between agents with different goals and skill bases. Like Fodor, Minsky also proposed some cross-module shared resources, such as long-term memory and a module maintenance system.

Now, many leading psychologists and philosophers of mind take as given that natural intelligence has some sort of modular structure (Carruthers, 2003). Similarly, modular architectures have come to dominate the areas of artificial intelligence that deal with problems most similar to humans': autonomous robotics and virtual reality characters (Bryson, 2000). The relevance of the preferences of AI engineers to theories of natural intelli-

gence is this: These engineers face the same problems as nature, such as dealing with uncertainty in perception and action, and searching enormous, combinatorially-explosive spaces of possible actions. Design features which have proved useful for engineering may therefore be treated as reasonable hypotheses for aspects of natural intelligence. It may be that evolution has had the same result of its search for successful control strategies as we have.

Yet the modularity hypothesis in natural intelligence is still controversial (Karmiloff-Smith, 1992; Carruthers, 2003). This is not least because the *definition* of modularity is still not agreed on. Fodor proposed many possible characteristics for identifying a modular system, such as innateness, localisation and domain specificity. Innateness is extremely difficult to determine, however, because of the nature of development (Elman et al., 1996). Development always begins with a uniform (though modular) embryo of a few cells. The shift into a highly specialised adult moves along a continuum. Selecting a single point when all ‘innate’ modules might be expressed might be impossible. Certainly no modules will express themselves without an environment containing oxygen. It’s quite likely that other modules might not be expressed without the sorts of intellectual stimulation that our species typically receives, such as social interaction. Further, there’s no reason to use birth as the defining moment for what attributes have to be present to be considered innate. Given the developmental continuum, other organismic milestones such as puberty might be far more appropriate.

Localisation, on the other hand, is in some sense a given. We know that the brain is divided into organs with different roles and representations. Yet most proposed cognitive modules, e.g. face recognition or language, utilise disjoint sections of brain simultaneously. Similarly, domain specificity has proved difficult. Even V1, the part of the brain identified with ‘early visual processing’ has lately been indicated for performing other spatial tasks, such as reading braille (Kauffman et al., 2002). This is hardly surprising given the nature of the evolutionary process and neural computing — any cell system found producing orderly and useful computations in a particular context is likely to be exploited and encouraged regardless of whether that is the context in which it first proved useful.

Software engineering has also struggled with the correct way to define modularity, where here ‘correct’ means ‘most useful’ — the way most likely to be clear, consistent, maintainable and scalable. Although there is no single universally-accepted strategy for modular decomposition in software engineering, one standard approach is to define a module in terms of its

variable state — that is, what it learns or stores, what it brings from past experience in order to generate behavior in the present. Although many AI systems attempt to store all knowledge in a single representation in order to maintain generality, this ignores the intractability of learning in the most general case Bryson (2001). Modularity provides an opportunity for specialised representations which can greatly increase the probability for any agent (animal or artificial) to successfully acquire the skills and knowledge it needs. This theory is derived from a transition the field of psychology made in understanding learning, from the general-purpose learning of the behaviourists to emphasis on specialised learning and modularity (Gallistel et al., 1991; Roper, 1983).

Differences in representation in animals are not only a good indication of modularity, but should in fact be taken as the primary criteria for modularity. While specialised representations are correlated with many of the traditional criteria for modules (such as innateness, domain specificity and localisation) they can also allow for a more flexible definition of ‘module’ that addresses the difficulties described above. A specialised representation can emerge (be learned) over development as neural systems tune to regularities useful in the agent’s context. This representation may apply to more than one domain, but still not to most. It might be encoded across multiple brain sub-systems that are routinely activated in a particular context through reciprocal connections. It may only exist in certain contexts when activation of its component parts is synchronised (or bound (von der Malsburg, 1995)), while in other contexts its component parts perform other functions independently. The concept of module defined in these terms still explains why some capacities of an individual are present in some contexts and not others. It also still makes the complexity of the overall system more manageable by encapsulating some behavior into a single system.

The remainder of this paper assumes this definition of modularity, which implies that to demonstrate modularity one must only demonstrate multiple distinct representations. We now move to examine whether task learning in primates is modular, beginning with the task of transitive inference.

## 2. Evidence of Modularity in Transitive Inference Learning

Transitive inference (TI) is the process of reasoning whereby one deduces that if, for some quality,  $A > B$  and  $B > C$  then  $A > C$ . This ‘reasoning’ capacity has been highly researched in psychology. Piaget described transitive inference as an example of concrete operational thought (Piaget,

1954). That is, children become capable of doing transitive inference when they become capable of performing mentally the physical manipulations that would determine the correct answer — in this case, visualising and comparing where two items appeared in a sequence. However, Bryant and Trabasso (1971) showed that pre-operational children could perform TI, given enough training. McGonigle and Chalmers (1977) also showed that monkeys (in this case squirrel monkeys *Saimiri sciureus*), although able to perform well on any two items (e.g.  $B, D$ ) drawn from a trained series of pairs  $\{A > B, B > C, C > D, D > E\}$ , if presented with *three* trained items from that set (e.g.  $B, C, E$ ) performed significantly worse. This last result, which also proved true of children (Chalmers and McGonigle, 1984), indicates that the subjects are not drawing their answers from a simple linear representation.

### 2.1. Modelling TI with a Production-Rule Stack

Harris and McGonigle (1994) presented a model that fit the McGonigle and Chalmers (1977) trigram data extremely well, based on production rule stacks. This model matches the performance of the monkeys very well whether they are modelled as a group and as individuals.

The production-rule stack model requires the following assumptions:

- (1) The subject learns a set of rules of the nature “if  $A$  is present, select  $A$ ” or “if  $D$  is present, avoid  $D$ ”.
- (2) The subject learns a prioritising of these rules.

This process results in a *rule stack*, where the first rule is applied if its trigger finds the context appropriate. If not, the second, and so on.

For an example, consider a subject that has learned this stack:

1. ( $A$  present)  $\Rightarrow$  select  $A$
2. ( $E$  present)  $\Rightarrow$  avoid  $E$
3. ( $D$  present)  $\Rightarrow$  avoid  $D$
4. ( $B$  present)  $\Rightarrow$  select  $B$

Here the top item (1) is assumed to have the highest priority. If the subject is presented with a pair  $CD$  it begins working down its rule stack. Rules 1 and 2 do not apply, since neither  $A$  nor  $E$  is present. However, rule 3 indicates the subject should avoid  $D$ , so consequently it selects  $C$ . Priority is critical. For example, for the pair  $DE$ , rules 2 and 3 give different results. However, since rule 2 has higher priority,  $D$  will be selected.

Adding the assumption that in the trigram test cases, an ‘avoid’ rule results in random selection between the two remaining items, Harris and McGonigle (1994) model the conglomerate monkey data so well that there is no significant difference between the model and the data. For example, over all possible trigrams, the stack hypothesis predicts a distribution of 0, 25% and 75% for the lowest, middle and highest items. Binary sampling predicts 3%, 35% and 63%, and logic of course 0%, 0% and 100%. The monkeys showed 1%, 22% and 78%. Further, individual performance of most monkeys were matched to a particular stack.

Without trigram data, there would be no way to discriminate which rule set the monkeys use. However, with trigram data, the stacks are distinguishable because of their errors. For example, a stack like

- 1'. ( $A$  present)  $\Rightarrow$  select  $A$
- 2'. ( $B$  present)  $\Rightarrow$  select  $B$
- 3'. ( $C$  present)  $\Rightarrow$  select  $C$
- 4'. ( $D$  present)  $\Rightarrow$  select  $D$

would always select  $B$  from the trigram  $BCD$  by using rule 2', while the previous stack would select  $B$  50% of the time and  $C$  50% because it would base its decision on rule 3.

## 2.2. A Two-Tiered Modular System for Learning TI

Production-rule models are actually common in cognitive modelling (e.g. Newell, 1990; Anderson, 1993), but are not satisfying, perhaps partially because they don't make a clear connection to brain mechanisms (though see Fincham et al., 2002). In an attempt to better understand how the Harris and McGonigle (1994) model could so accurately represent animal behavior, I built a modular neural learning version of their system. I broke the learning into two problems: the association of a stimulus (an item, e.g.  $B$ ) to an action (select or avoid, as above), and the prioritisation of those perception/action pairs (see Figure 1).

This two-tiered model consists of a list of items seen, each associated with a weight and with a unique 2-item list of possible rules (select or avoid) also associated with weights. When the agent sees a new item, it adds it to its list of stimuli. The biological plausibility of these assumptions is addressed in the discussion.

All of the weights in a single list are normalised (they sum to 1) — new items receive the weight  $1/N$  (where  $N$  is the current number of items). New items in the stimulus list are also associated with a rule-list where both

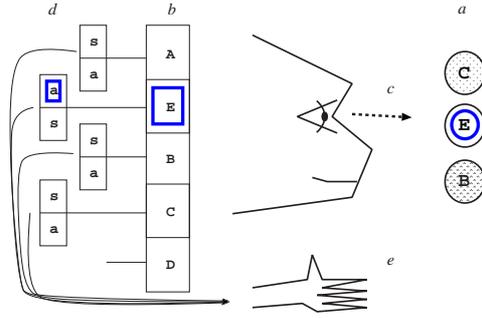


Figure 1. The two-tier model. When the agent observes a set of stimuli ( $a$ ), a weight vector ( $b$ , the first tier) determines which item present is most salient. This attracts visual attention ( $c$ ) and determines which rule vector ( $d$ , the second tier) selects the appropriate action (select or avoid) which controls the monkeys grasp ( $e$ ). The two vectors that were most recently active ( $a$  and one of  $d$ ) are then updated as determined by the result and the learning rule (Equation 1).

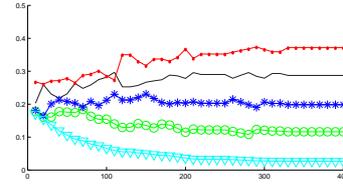


Figure 2. Typical results for single-vector learning (only one rule — select). X-axis: trial number; Y-axis: weights of stimuli vector (sum to one). Free variables are set to Parameters:  $\Xi = .08$ ,  $\delta = .02$ . Key:  $A \bullet$ ,  $B \circ$ ,  $C \ast$ ,  $D \square$ ,  $E \nabla$ .

*select* and *avoid* have weight 0.5. provided. When the agent is given a trial, it examines its list for the rules associated with the stimuli present. For that set of rules, the highest-priority stimulus determines the focus of the agent's attention. Then the highest priority rule in the rule-list associated with that stimuli determines whether the agent selects or avoids the object it has fixated on. In the case of a tie, outcome is arbitrary.

Weights are updated after every trial, where a trial is a presentation of a pair or trigram, a selection by the agent, and a reward. The training rule is a simple step function. The subject learns a weight vector  $\mathbf{v}$ . For the pair  $XY$ , where  $X$  is selected by the subject and  $\mathbf{v}_X$  and  $\mathbf{v}_Y$  are the weights associated with  $X$  and  $Y$  respectively:

$$\begin{aligned} &\text{If } X \text{ is correct and } (\mathbf{v}_X - \mathbf{v}_Y < \Xi), \text{ add } \delta \text{ to } \mathbf{v}_X; \\ &\text{else, if } X \text{ is incorrect, subtract } \delta \text{ from } \mathbf{v}_X. \end{aligned} \quad (1)$$

where  $\Xi$  and  $\delta$  are free parameters.  $\Xi$  is a threshold over which reward is so expected that it no longer prompts learning.  $\delta$  is the amount a weight is changed by a single bout of learning. If a weight-change occurs,  $\mathbf{v}$  is subsequently renormalised.

The results of this model are an even better match to nature than the original Harris and McGonigle (1994) model. First, when the model learns it exactly replicates the production-rule model, so in that sense it is equivalent. But second, like primates (both children and squirrel monkeys) it *doesn't* always learn successfully. Chalmers and McGonigle (1984) document a training regime that is necessary for getting a majority of children to successfully learn the initial pairs ( $A > B$ ,  $B > C \dots$ ). This task is difficult because most items are sometimes rewarded and sometimes not.

Figure 2 shows a typical learning of weights in a more conventional, single-representation model of TI. These results are based on the same learning rule given above, but applied only to a single vector. Here learning occurs very reliably — the only way to disturb it is to constrain the learning resources such that the individual cannot find a stable set of weights, in which case the subject occasionally errs.

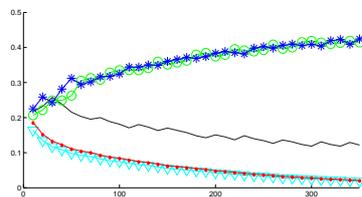


Figure 3. Rule learning with no training regime. Rules learned in descending order of priority: select  $D \circ$ , avoid  $C^*$ , avoid  $B^-$ . The system cannot stabilise because it is far from a complete solution, but it behaves correctly for every training pair except  $CD$ . Parameters:  $\Xi = .08$ ,  $\delta = .02$ .

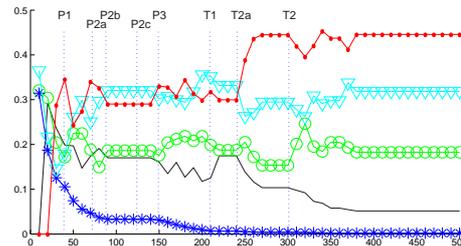


Figure 4. Rule Learning with Phased Training. Labelled lines indicate the *end* of training and testing phases (taken from Chalmers and McGonigle (1984)). This agent arrives at different stable solutions at different points, but they are all correct. The final rules set is: select  $A \bullet$ , avoid  $E \nabla$ , avoid  $D \circ$ . The agent succeeds even with very 'stupid' parameters:  $\Xi = .12$ ,  $\delta = .06$ .

Figures 3 and 4 show outcomes with the two-tier model. These diagrams only show the weights of the first tier, the priorities between rules. The rules themselves converge very quickly and then remain stable over the trial; the captions indicate whether 'select' or 'avoid' dominates for each stimuli. Unless the agents are put through a training regime, only roughly a quarter of the agents successfully learn the task. Most of the rest display the pattern shown in Figure 3, where they have learned what to do in the case of the

endpoints successfully. Interestingly, the dominant tendency is to neglect the endpoints themselves, and focus attention on their neighbours, which are after all the more difficult learning task. Unfortunately for the agents, no correct rule set exists without an endpoint as the highest priority rule. Learning fixates on finding a correct ordering for two adjacent non-endpoint items, but no such ordering exists. In Figure 3, the agent's behavior is identical whether it selects  $D$  or avoids  $C$ : either is right in all training cases except when presented with  $CD$ , in which case either is wrong.

On the other hand, under the sort of training a child receives when being tested on the TI task, the agents learn successfully almost 9 times in 10, and across a broad range of parameter values Bryson and Leong (2004). Thus the two-tier model can not only successfully learn a model which conforms to Harris and McGonigle (1994), but it also *fails* to learn in an appropriate way: occasionally with training, and very frequently without it.

### 3. Modularity and the Two-Tier Model

Despite the quality of its results, the two-tier model is unattractive for several reasons. First, why when a single-vector model of TI learning would be not only simpler but more effective would nature choose something so complex? The answer of course is that TI is almost certainly not being learned on a special-purpose mechanism — if it had, perhaps it would also have applied itself to trigrams. This of course begs the question, what is this learning system, and what is it for?

One possible hint to this answer lies in another set of questions about the model. Why are the two rules 'select' and 'avoid'? Why isn't avoid better informed? Why are there the same number of rules as there are stimuli? My guess is that TI is actually revealing the deeper structure of its learning mechanism, a more general purpose one for task learning. I suspect that for any given problem, the agent attempts to determine a minimum set of important discriminative stimuli and also a minimum set of appropriate actions. When subjects are first taught this task, they are taught to 'select' for a reward. Then they are presented with two stimuli, and penalised for selecting incorrectly. This forces them to learn another behavior, 'avoid', a partially-inhibited form of select. 'Select, but not the thing you are attending to.'

Within this ontology, the animal then has two additional things to learn: in which contexts to apply which behavior, and what the prioritisation is between contexts when more than one applies.

The clearest aspect of the two-tier model is that it splits the problem of task-learning into at least two parts:

- (1) learning to associate actions with environmental cues
- (2) learning to prioritise these cues when more than one association holds.

Neuroscience indicates that long-term rule learning, including categorisation for action selection, seems to be dependent on both the prefrontal cortex (Freedman et al., 2001; Wallis et al., 2001) and the hippocampus. In particular, the entorhinal cortex seems necessary for forming episodic or event memory necessary for such tasks as paired association, while the hippocampal formation proper is necessary for learning more complex relations (Alvarado and Bachevalier, 2000; Baxter and Murray, 2001).

The hippocampus has long been implicated in memory consolidation. Its very sparse, redundant coding system seems ideally suited to recording memories in real time which can then be consolidated — saved to long term storage in an integrated representation — using a slower learning mechanism (McClelland et al., 1995; Louie and Wilson, 2001). Thus the following might be a reasonable neurobiological explanation of the model presented in this paper. The basic association between stimuli and behaviors would be performed by the prefrontal cortex generalising from associations learned by the entorhinal cortex. The hippocampal formation would learn prioritisations between these associations.

#### 4. Conclusion

The two-tier model demonstrates a clearly modular decomposition of human-like intelligence. Associations and priorities are learned separately as different representations and in different parts of the brain. Further refinements on the model, such as explanations of TI-related timing effects, will show further modularity in complex process of learning new tasks.

#### References

- Alvarado, M. C. and Bachevalier, J. (2000). Revisiting the maturation of medial temporal lobe memory functions in primates. *Learning and Memory*, 7(5):244–256.
- Anderson, J. R. (1993). *Rules of the Mind*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Baxter, M. G. and Murray, E. A. (2001). Opposite relationship of hippocampal and rhinal cortex damage to delayed nonmatching-to-sample deficits in monkeys. *Hippocampus*, 11(1):61–71.
- Bryant, P. E. and Trabasso, T. (1971). Transitive inferences and memory in young children. *Nature*, 232:456–458.

- Bryson, J. J. (2000). Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, 12(2):165–190.
- Bryson, J. J. (2001). *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*. PhD thesis, MIT, Department of EECS, Cambridge, MA. AI Technical Report 2001-003.
- Bryson, J. J. and Leong, J. C. S. (2004). Representations underlying transitive choice in humans and other primates. submitted.
- Carruthers, P. (2003). The cognitive functions of language. *Brain and Behavioral Sciences*, 25(6).
- Chalmers, M. and McGonigle, B. O. (1984). Are children any more logical than monkeys on the five term series problem? *Journal of Experimental Child Psychology*, 37:355–377.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., and Plunkett, K. (1996). *Rethinking Innateness. A Connectionist Perspective on Development*. MIT Press, Cambridge, MA.
- Fincham, J. M., Carter, C. S., Veen, V., Stenger, V. A., and Anderson, J. R. (2002). Neural mechanisms of planning: A computational analysis using event-related fMRI. *Proceedings of the National Academy of Sciences, USA*, 99(5).
- Fodor, J. A. (1983). *The Modularity of Mind*. Bradford Books. MIT Press, Cambridge, MA.
- Freedman, D. J., Riesenhuber, M., Poggio, T., and Miller, E. K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science*, 291:312–316.
- Gallistel, C., Brown, A. L., Carey, S., Gelman, R., and Keil, F. C. (1991). Lessons from animal learning for the study of cognitive development. In Carey, S. and Gelman, R., editors, *The Epigenesis of Mind*, pages 3–36. Lawrence Erlbaum, Hillsdale, NJ.
- Harris, M. R. and McGonigle, B. O. (1994). A model of transitive choice. *The Quarterly Journal of Experimental Psychology*, 47B(3):319–348.
- Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Change*. MIT Press, Cambridge, MA.
- Kauffman, T., Theoret, H., and Pascual-Leone, A. (2002). Braille character discrimination in blindfolded human subjects. *Neuroreport*, 13(5):571–574.
- Louie, K. and Wilson, M. A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*, 29(1):145–156.
- McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457.
- McGonigle, B. O. and Chalmers, M. (1977). Are monkeys logical? *Nature*, 267:694–696.
- Minsky, M. (1985). *The Society of Mind*. Simon and Schuster Inc., New York, NY.
- Newell, A. (1990). *Unified Theories of Cognition*. Harvard University Press, Cambridge, Massachusetts.
- Piaget, J. (1954). *The Construction of Reality in the Child*. Basic Books, New York.
- Roper, T. J. (1983). Learning as a biological phenomena. In Halliday, T. R. and Slater, P. J. B., editors, *Genes, Development and Learning*, volume 3 of *Animal Behaviour*, chapter 6, pages 178–212. Blackwell Scientific Publications, Oxford.
- von der Malsburg, C. (1995). Binding in models of perception and brain function. *Current Opinion in Neurobiology*, 5:520–526.
- Wallis, J. D., Anderson, K. C., and Miller, E. K. (2001). Single neurons in the prefrontal cortex encode abstract rules. *Nature*, 411:953–956.