Scale Invariant Associationism, Liquid State Machines, And General Purpose Ontogenetic Learning In Robotics

Anthony F. Morse

COGS, Department Of Informatics, University Of Sussex
Falmer, Brighton, BN1 9RH, UK
anthonfm@sussex.ac.uk

Abstract

Unsupervised Hebbian learning produces a scale invariant associative phenomena with limited computational power (the linear separation of states). With the inclusion of a Liquid State Machine (LSM), simple associative methods are transformed to provide truly general purpose and computationally powerful ontogenetic learning. This radical expansion of the representational space of a problem enables solutions by computationally primitive means. This not only supercharges unsupervised techniques but also leads to cognitive penetrability and successful psychological explanation. Demonstrated on an embodied robot the approach displays cumulative general-purpose learning and training through human interaction leading to multiple psychological phenomena, sequence learning, navigation, and a strong resilience to catastrophic forgetting.

Introduction

One of the major challenges in building intelligent robots is to provide mechanisms for the learning of new tasks and the solving of new problems. Designers want their robots to be shaped by future owners in unanticipated ways, and through physical interaction rather than reprogramming. For psychology, solutions to such problems may have the potential to unite theories of disparate phenomena through common mechanisms influenced in different ways by different incoming data streams. In neuroscience, increasing evidence suggests that biological brains develop structures that are highly dependent on the incoming data, rather than genetically pre-specified circuitry (Sur et al 2000). In this paper a biologically plausible account of very simple unsupervised plasticity is demonstrated to result in complex cognitive learning far beyond the abilities of normal unsupervised algorithms. Modeling both classical and operant conditioning, the resulting robotic agent adapts well to its environment and remains malleable through learnt social interaction and other new experiences. As a general-purpose approach to learning, the same architecture can be placed in any embodiment and face any task without necessitating reconfiguration of any kind. The subsequent success at solving any particular problem will depend largely on exposure to appropriate sensory motor contingencies over a prolonged developmental period.

Enactive Distributed Associationism

The Enactive Distributed Associative (EDA) approach (Morse 2004) uses a hybrid Hebbian algorithm initially designed for the ongoing construction of symbolic Interactive Activation and Competition (IAC) networks (Morse 2003). The learning rule is implemented such that; iff *a* and *b* commonly co-occur in the context *x*, then the presence of either *a*, or *b*, in the context *x*, will cause expectation of *b*, or *a*, respectively. This is implemented using an Adaptive Resonance Theory network (ART) to identify patterns / contexts in sub-populations of the input stream. These autonomously discovered context patterns are then subject to Hebbian learning from the input layer (see fig 1). Positive connections are forged bi-directionally while negative connections emanate from the context nodes only. This results in an associative algorithm where the consequences of any particular event are free to vary wildly across situations without adversely effecting learning. Rather than apply this algorithm to localist symbolic entities, it is instead applied to the time variant activity vector of a random Liquid State Machine (LSM) cortical microcircuit (Maass et al 2002a,b,c). The LSM is viewed as implementing both the high dimensional space of a Support Vector Machine (SVM) kernel and the temporal variance properties of a fading analogue memory. Thus we can view its activity vector as containing transiently localist micro-symbolic features of the recent input stream. Although LSM networks are typically read with parallel perceptrons explicitly trained to recover specific variables, this is achieved via a linear transformation. Thus, anything at all which is linearly recoverable from the LSM need not be explicitly recovered or trained for as it will have the appropriate associative effects under EDA anyway. From this perspective, the application of ongoing IAC learning can now be applied to a LSM receiving input from any device and in any format (and is thus non-task specific). The resultant architecture simply learns to mimic incoming data streams thus providing prediction. Motivational rewards come into play by selecting specific inputs to strongly inhibit activity in other positively associated neurons when primed. Thus behaviors causing predicted or expected activity in these inhibitory inputs will themselves become inhibited. Attaching these mechanisms to input neurons such as bumper activity, naturally results in the development of appropriate behaviors to minimize these inputs, such as obstacle avoidance.

Trading Spaces & Transient Localism

One of the well-known limitations of Hebbian plasticity is that it can only accurately capture / predict linear relations between the entities it is applied to (Clark & Thornton 1997). Due to this limitation, Hebbian computational power is dependent on the representational space afforded it to find any solution. If for a given problem, appropriate representations are linearly separable (i.e. at least partial physical separation) then Hebbian plasticity will lead to accurate prediction (Grainger & Jacobs 1998), if not then predictive performance will suffer. Providing such limited methods with input filtered through a LSM, we massively expand and warp the representational space afforded. This is in anticipation of providing many more separable micro-features than were present in the original data stream. This approach is very similar to that used in SVM's where linear decision boundaries constructed in a high dimensional space, translate to complex non-linear decision boundaries in the original data space. In addition to implementing the kernel of a SVM without supervision or training, LSM's also provide temporal variance, simplifying problems in sequence learning and time dependant responses (Maass 2002a,b,c). Even though the referent of any given localist interpretation will necessarily vary over time (without necessitating external change), in any particular context (of local LSM activity), the referent remains stable. This provides a transient (it will change) localist (but right now it's stable) interpretation (Morse 2004). Thus because the LSM preserves differences between input streams, and cannot generate any activity independently from an input, a micro-symbolic non-stationary representational account of associative behavior can be formed (see fig 1).

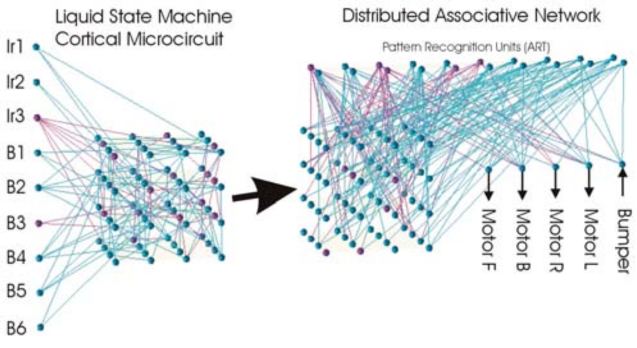

Fig1. *Showing the EDA architecture consisting of a cortical microcircuit (left) and a (context sensitive) distributed associative network (right). Input from sensors perturbs the LSM. ART pattern recognition then creates and activates a context (top right). Finally associative connections grow between these contexts and all non-context nodes (either in the LSM or in the output).*

Modeling basic ontogenetic development, EDA implements a powerful adaptive force capable of capturing highly complex non-linear separations in any data stream. Although EDA is heavily reliant on the LSM involved to provide appropriately separable representations, there is no specialization by design. This ability comes readily from the dynamics of cortical microcircuits. LSM models have previously been utilized in conjunction with parallel perceptrons (also limited to linear separation) and have been shown to fair well against other more complex supervised methods in a number of standard benchmark tests (Maass et al 2003a,b,c).

Psychological Explanation

By analogy to localist symbolic IAC networks, the behavior of EDA architectures can easily be understood. Constructing associative structures from transiently localist micro-symbolic neurons, the pervasive characteristics of associative learning are exhibited at all greater scales (e.g. in associative descriptions of behavior). This simply means that associative rules applied at the sub-symbolic level necessarily result in the same rules accurately describing the macro-symbolic relations (where linear separation is possible). These architectures naturally display context sensitive associations of the kind; if *a* and *b* are typically present together, then the presence of one will cause expectation of the other. Even though this is implemented in neural plasticity, the same functional description equally applies to the high level behaviors the agent produces. The scale invariant phenomena produced by Hebbian associations leads to these rules applying equally well to any aspect of the agent's internal or external behavior. The additional structuring of associations via context (Morse 2003) provides greater sensitivity for relations such as; *a* and *b* together imply *c* (where neither *a* nor *b* individually imply *c*). This additional complexity stabilizes the transient interpretations of activity in the cortical microcircuit by providing a context to each association. Combined with a micro-symbolic account of the interpretation, EDA provides a firm neuroscientific basis from which psychological explanation can begin.

All of the following models implement the same EDA architecture in the same embodiment (see fig 2) and using the same parameter settings. The cortical microcircuit used was randomly generated and consisted of 64 neurons arranged in a 4 by 4 by 4 grid. A pool of 30 ART2 neurons were made available (See Morse 2003 for further details).

Cognitive Modeling

Previous publications on the EDA methodology have demonstrated successful embodiment of a single ontogenetically learning robot displaying multiple psychological phenomena (Morse 2004, 2003). Following Burton's descriptions of IAC networks (in Young and Burton 1999).

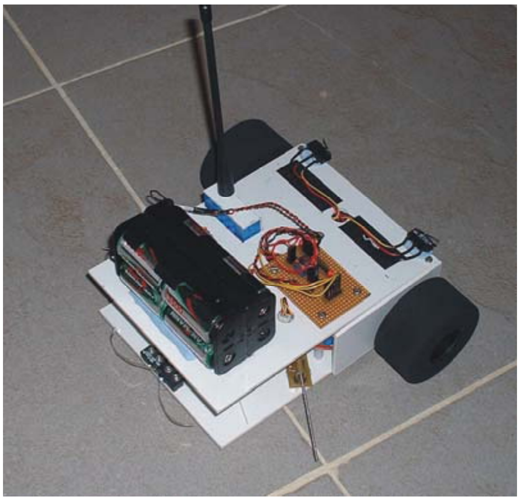

Fig2. *An autonomous robotic embodiment of the EDA algorithm.*

The EDA robot consistently displays - semantic and repetition priming, overt and covert recognition, and can be impaired with various prosopagnosias. In addition to this, the EDA architecture also displays schema learning via classical and operant conditioning, and phobia development. Unlike most cognitive models, these phenomena are not captured separately, require no prior knowledge of embodiment or environment, and require no parameter fixing. Such approaches may ultimately lead to the unification of disparate psychological theories by explanation through common mechanisms influenced differently by different incoming sensory streams. This perspective has much in common with recent theories in enactive psychology (O'Regan, & Noë, 2001), and should be viewed as a cumulative alternative to dichotomous approaches in cognitive modeling.

Semantic Priming is the phenomena whereby the speed of recall of facts associated with a stimulus can be robustly manipulated by recent prior presentation of other similarly associated stimuli. In EDA architectures, as in other spreading activation models (e.g. IAC), associations between symbols, micro-symbols, or behaviors, lead to a transfer of energy between active and associated entities. The target response, having been recently primed by the first stimulus, retains decaying residual activity for a short period of time. If the second stimulus is presented during this time, the subsequent re-priming of the target is speeded up by this residual activity (See Fig3). This phenomena crosses domain but remains short-lived.

Repetition priming describes a far more constrained effect in which repeated stimulation of an association causes temporary strengthening leading to a faster response. This is also related to exposure effects in which strong associations are far more readily recalled than weaker ones.

Overt and Covert Recognition: In addition to priming, figure 3 also shows overt recognition, in that upon presentation of either stimulus, appropriate responses correctly gain in activity. Covert recognition can be induced when low levels of primed activity cause changes in context resulting in behavioral changes. Thus appropriate behaviors can be triggered by very weak changes in activity insufficient for overt recall. Similar behavior can also occur from damage degrading the strength of associations. In specific input regions this would lead to the exhibition of prosopagnosias (Burton et al 1999), and hamper further learning relating to new instantiations of the covertly recognized stimulus.

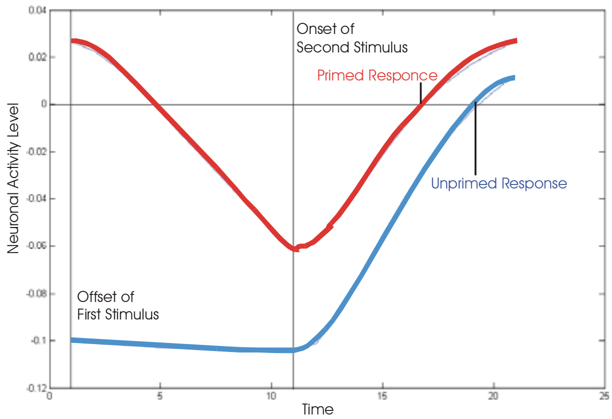

Fig3. *Showing the activity of a motor neuron over time, following sequential presentation of two environmental stimuli associated with that motor action (primed response), and following only the second associated stimuli (unprimed response).*

Classical Conditioning: As embodied EDA models continually learn from their experiences, analogies are possible to both classical conditioning (Pavlov 1927), and operant conditioning (Skinner 1953). In classical conditioning, a natural response to a specific stimuli (such as saliva production in the presence of food) is induced by the presence of a conditioned stimulus (such as a bell ring at every mealtime). The conditioned stimulus initially provokes no specific reaction, however due to the high frequency of co-presentation, the new stimulus (bell) becomes associated with the original stimuli (food) and induces the same response. If strong enough, the future presentation of the conditioned stimulus alone (bell) will provoke the conditioned response (salivation). This process is demonstrated in the EDA learning of obstacle avoidance behavior. The agent initially fails to react to activity from the Infra-Red (IR) sensors (proximity), however as this is typically high during bumper activity, the agent becomes conditioned to produce reactions to collisions before they actually happen. This demonstrates the successful acquisition of obstacle avoidance behavior consistent with classical conditioning.

Operant Conditioning: In operant conditioning, the consequences of a behavior can result in modifications to the production of that behavior. In the embodied robot, this is again demonstrated by changes in the behaviors that lead to obstacle avoidance. Here, the action of the forward motor becomes sensitive to the presence of activity in the IR sensors as a result of previously experienced collisions. Both classical and operant conditioning are demonstrated and explained by the same simple underlying mechanisms producing the same phenomena at slightly different scales. Two different behavior types develop from this model. If the agent experiences early frontal collisions it tends to produce behavior backing into open spaces and spinning (while remaining responsive to new obstacles). Agents initially experiencing rear collisions tend to adopt forward exploratory strategies. This further demonstrates the agents' malleability to environmental experience.

Phobias: Associative theories of learning lead to the possibility of misplaced associations occurring due to coincidences in an agent's experience. Wolpe (1958) describes an account of how conditioning can result in the misplaced association of a stimulus present during a traumatic event, to a fear response. This condition is known as a phobic response, and is modeled in the robotic agent's phobia of narrow corridors. Having associated (through many experiences) IR activity with a collision, the agent incorrectly predicts such a collision when facing a narrow entrance. This is misplaced as the agent could easily pass through without causing a collision. A cognitive behavioral technique for overcoming phobic reactions is to gradually desensitize the patient in a state of relaxation. To model this, the agent is given experience of an additional input during safe situations. The new input associates to inhibit/reduce the prediction of a collision and can therefore be used to 'relax' the agent. With this additional input active, the agent is able to approach the narrow passage more closely, gaining substantially more experience contradicting the misplaced association. Following several such approaches the agent is able to traverse the passage and can subsequently do this without the aid of a 'relaxing' input. This removal of an otherwise robust phobic response provides a clear example of how human interaction can adapt the behavior of the EDA robot. These manipulations also have no adverse affect on any of the agents other learnt behaviors.

ALife Modeling

IAC networks can readily implement environmentally cued sequential learning by associating a sequence of landmarks with each step of a routine. EDA networks however have a far richer and temporally varying representational space. This allows the embodied EDA robot to learn sequences requiring no external change but instead based on the temporally changing internal representational space. Autonomously learning such sequences enables time dependant responses to be generated and further corrected by landmark recognition.

Sequence Learning & Temporal Prediction

The input of the embodied EDA robot was deliberately manipulated to include a specific repeating 6 bit binary sequence, 100 steps in length (including repetition of steps in different orders). Each step of the sequence remained a fixed input for a duration of 1 second, the next step of the sequence was then presented. After 15 minutes (9 complete cycles of the sequence), the additional input was removed and the new input units observed. These units continued displaying the taught sequence of activity (although at lower levels) with gradually decaying performance over time. Subsequent presentation of 5 ordered steps from the sequence, at any point following the initial training, resulted in continuation of the sequence from that specific point with increased accuracy again degrading over time (See fig3). Even though polarity correlation drops in accuracy (as the temporally cued change from one step to the next becomes blurred), taking the average activity over each 1 second period (each step presentation) results in perfect non-decaying polar correlation to the original trained sequence.

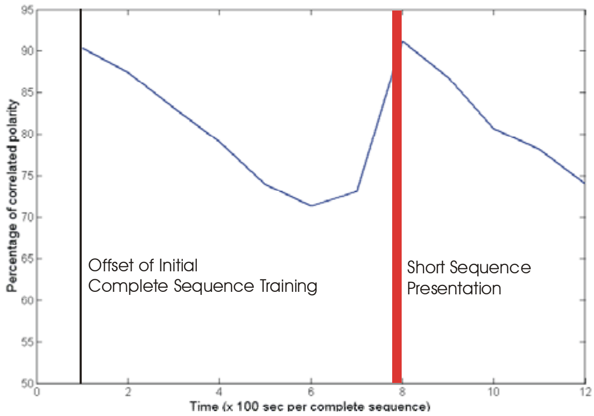

Fig4. *Showing the percentage correlation between the activity generated independently, and the trained sequences over time. External influence is removed at time step 1. A 5 set sequence is then presented at the end of step 7 leading to enhanced performance and a positional shift in the looped sequence. The percentage correlation is the sum of appropriately matched polarities minus the sum of mismatched polarities between target and input neurons, sampled at 1msec intervals and averaged over each sequence step.*

Navigation

The embodied EDA agent was further driven, by external stimulation of its motor output neurons, from open space towards a wall. On reaching the wall, the agent was remotely turned to the right to follow the wall for 5 sec and then turn away again. This procedure was repeated several times until the agent successfully replicated the trained sequence of behavior autonomously. By subsequent manipulation of the environment, the robotic EDA agent was found to move forward until it reached a wall, irrespective of how far that was, before turning right. This is an example of simple landmark recognition, where the change in behavior is cued by environmental phenomena. No such external cue was available for the second turn away from the wall yet the agent robustly began this turn after 5 (+-0.5) seconds of wall following. This demonstrates the successful integration of LSM temporal variance with the previously limited IAC cued sequential learning abilities.

Catastrophic Forgetting

The learning of complex sequences of data raises the issue of catastrophic forgetting. In most forms of neural network modification, as new relationships are learnt, previously acquired relationships are lost. This is due to the algorithms re-tuning weighted connections previously tuned for different responses. The context sensitivity of associations generated using the EDA method means that the re-tuning of weights will only occur (in any strong way) in the specific context in which the weight was originally created. This prevents a great deal of forgetting but can also lead to the prolonging of falsely learned associations, especially if they provoke avoidance behaviors. Useful forgetting can also occur via environmental manipulations such as those detailed in the models of behavioral phobia correction presented earlier.

Discussion

Like all algorithms, there are limitations to the EDA approach. Firstly the cortical microcircuit must be able to separate sufficient relevant features from the data stream. Secondly the agent must have appropriate experience of sensory motor contingencies. Despite these limitations EDA successfully models complex ongoing learning in which neuronal mechanisms provide explanatory accounts of disparate systematic high level behavior and mental phenomena. Fodor's assertion that cognitive systems display systematicity (Fodor & Pylyshyn 1988) is also non-trivially true for the EDA architecture. This is evident in the scale invariance of the associative descriptions given.

The EDA methodology is a development from early localist PDP models (McClelland & Rumelhart 1986) having all the advantages of such localist symbolic architectures without the disadvantages of having to specify a symbolic starting point. Providing no starting point, and using an unsupervised algorithm, the agent has no input whatsoever from the designer and so is truly general purpose. The only factor the designer needs to worry about is how big the EDA network should be. The larger a network is, the more representational power it has, although the evolution of task-specific LSM models can improve performance drastically.

Simple structured associative plasticity, implemented between the neurons of a cortical microcircuit, results in complex associative phenomena in the behavior of an embodied agent. This leads to a formal high level of description explaining multiple known psychological properties. Liquid State Machines further provide a biologically plausible, transiently localist, micro-symbolic basis for the construction of concepts in a language of thought. The nature of this representational space is however far more complex than that of existing static symbolic models and contributes greatly to the computational and cognitive abilities of context sensitive associative mechanisms. As an example of the trading spaces approach, EDA architectures poses only very limited computational power, however as the representational space afforded is so rich, these limited computations far outperform their supposed limited application.

Although adult human cortex is highly structured and modularized, Karmiloff-Smith (2000) suggests that this specialization results from gradual ontogenetic development. Even specialized brain structures normally found only in specific regions of the visual cortex can be induced elsewhere following surgical sensory swap-over (Sur et al. 1989 & Sharma et al. 2000). The cortex's ability to build self-organizing and somewhat modular structures is reflected in the structural development of EDA networks. EDA structures are ontogenetically formed by the incoming data streams. These algorithms each capture and abstract different aspects of known plasticity and processing in the cortex of real biological brains to produce cognitively useful adaptations.

EDA plasticity is highly suited to incorporate inputs from other circuitry such as evolved behavioral networks or Self-Organizing Maps (SOM's). By learning sequences and associating external cues, EDA can mimic the function of most other networks and enable future modification through unsupervised operant conditioning. It is suggested that such associative adaptability can simply be bolted onto any existing system enhancing its existing capabilities. Of specific interest is the potential incorporation of self-organizing maps pre-processing the input stream, and gradually modifying the behavioral repertoire. Should the LSM fail to provide a sufficiently rich representational space, expansion or evolution of the cortical microcircuit may become necessary. Although evolution is a possibility suggested by Maass et al (2002a,b,c) it has not been a necessary stage in any of the experiments conducted so far. Cursory investigations have shown far more robust behaviors and significantly reduced developmental periods can be achieved through the evolution of the LSM. Such developments are however seen as moving away from the non-task specific application of the EDA methodology.

Conclusion

The successful embodiment of unsupervised symbolic learning without design specialization or the imparting of a symbolic basis provides a new approach to cognitive modeling and developmental robotics. EDA is a working demonstration that differences in incoming data can be sufficient to generate multiple disparate and complex psychological phenomena from the same underlying cortical mechanisms as are present in real brains. From a transient localist perspective, the context specific interpretations of activity within the complex cortical microcircuit provides a micro-symbolic basis from which explanation is easily provided. EDA does not implement just any abstract high level formal system, but one that has been demonstrated to produce multiple psychological phenomena in a real robotic embodiment. As a general-purpose model of conditioned learning, EDA is a major step towards the construction of intelligent robots, providing adaptability to new tasks and the solving new problems.

Future Directions

It is intended that the EDA network be tested with a richer embodiment. Various reward based training experiments are planed, including visual attention and auditory command learning. Further developmental psychology experiments are also planed with the intention of demonstrating incremental learning stages. Experiments are also planned to demonstrate the EDA architecture influencing multiple pre-evolved networks to produce useful adaptations in the coordination of bipedal walking behaviors.

References

Clark, A., & Thornton, C., 1997. *Trading spaces: Computation, representation, and the limits of uninformed learning.* Behavioral and Brain Sciences 20 (1): 57-92.

Fodor, J.A., & Pylyshyn, Z.W., 1988. *Connectionism and Cognitive Architecture: A Critical Analysis.* In Macdonald, C., & Macdonald, G., Eds. 1995 Connectionism: Debates on Psychological Explanation, pp 90-163.: Blackwell Press.

Grainger, J., & Jacobs, A.M., (1998) Localist Connectionist Approaches to Human Cognition.: LEA Press

Karmiloff-Smith, A., 2000. *Why Babies' Brains Are Not Swiss Army Knives*. In Rose, H., & Rose, S., (Eds), Alas, poor Darwin, pp 144-156, Jonathan Cape, London.

Maass, W., Natschläger, T., & Markram, H., 2002a. *Real-time computing without stable states: A new framework for neural computation based on perturbations*. Neural Computation, 14(11):2531-2560, 2002.

Maass, W., Natschläger, T., & Markram, H., 2002b. *A model for real-time computation in generic neural microcircuits.*: In the Proc. of NIPS 2002

Maass, W., Natschläger, T., & Markram, H., 2002c. *Computational models for generic cortical microcircuits.* In J. Feng, (ed.), Computational Neuroscience: A Comprehensive Approach.: CRC-Press.

Morse, A.F., 2004. *Psychological ALife: Bridging The Gap Between Mind And Brain; Enactive Distributed Associationism & Transient Localism*. In Cangelosi,A., Bugmann, G., Borisyuk, R., & Bullinaria, J., Eds. Proceeding of the ninth conference on neural computation and psychology.: In Press

Morse, A.F., 2003. *Autonomous Generation Of Burton's IAC Cognitive Models*. In Schmalhofer, Young, & Katz Eds. Proceedings of EuroCogSci03 The European Cognitive Science Conference.: LEA Press.

O'Regan, J.K., & Noë, A., 2001. *A sensorimotor account of vision and visual consciousness*, in Behavioral and Brain Sciences, 2001, 24(5), 939-1011.

Palvov, I., 1927. *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex.*

Sharma, J., Angelucci, A., & Sur, M., 2000. *Induction of visual orientation modules in auditory cortex*. Nature 404, 841-847.

Skinner, B. F. 1953. *Science and human behavior*.: New York. Macmillan.

Sur, M., Garraghty, P.E., & Roe, A.W., 1989 *Experimentally induced visual projections into auditory thalamus and cortex*. Science 242, 1437-1441.

Wolpe, J., 1958. *Psychotherapy by Reciprical Inhibition.* Stanford, CA: Stanford University Press.

Young, A.W., & Burton, A.M., 1999. *Simulating face recognition: Implications for modeling cognition*. Cognitive Neuropsychology*, 16(1),* 1-48