A Recurrent Neural Network for Game Theoretic Decision Making

Sudeep Bhatia (s.bhatia@warwick.ac.uk)
University of Warwick
Coventry, UK

Russell Golman (rgolman@andrew.cmu.edu)
Carnegie Mellon University
Pittsburgh, PA. USA

Abstract

We describe the properties of a connectionist network that is able to make decisions in strategic games. We use the structure of Bidirectional Associative Memory (BAM), a minimal two-layer recurrent neural network with binary activation functions and binary connection weights. We apply BAM to finite-strategy two-player games, and show that network activation in the long run is restricted to the set of rationalizable strategies. The network is not guaranteed to reach a stable activation state, but any pure strategy profile that constitutes a stable state in the network must be a pure strategy Nash equilibrium.

Keywords: Decision making; Constraint satisfaction; Neural networks; Game theory; Nash equilibrium

Introduction

Strategic decision making is an important feature of high-level cognition. From coordinating meeting times to cooperating on research projects or negotiating household chores, interdependent scenarios – in which the choices of an individual are contingent on the choices of others – are ubiquitous in everyday life. Game theory provides a mathematical framework with which rational decision making in interdependent scenarios can be formally represented and analyzed. It is an important area of research in economics, political science, computer science, and philosophy, and it is widely used for studying behavior in social settings, where it describes and predicts the types of cooperation, conflict, and coordination observed in groups of human decision makers (see, e.g., Camerer, 2003).

Despite its importance for the study of human behavior (as well as rational interdependent behavior, more generally) the cognitive basis of game theoretic decision making is still unclear. Solving game theoretic problems can be computationally intractable (Daskalakis et al., 2008). How do humans represent choices, reason through contingencies involving the choices of others, and ultimately make decisions, in strategic settings?

In this paper we present a connectionist model that is able to reason through game theoretic problems in two-player games. Our model is based on Kosko's (1988) Bidirectional Associative Memory (BAM) network, which is a minimal two-layer recurrent neural network with binary activation functions and binary connection weights. We assume that the two layers in our network represent the strategies available to the two players, with connections between these layers encoding best-responses to any given strategy.

Like the Hopfield network, which it generalizes, BAM can make decisions through constraint satisfaction, with final choices corresponding to activated nodes once the network stabilizes. Constraint satisfaction networks (see e.g. McClelland & Rumelhart, 1981) are commonly used to model decision making in non-strategic settings, including causal reasoning, stereotype formation, analogical mapping, legal reasoning, and preferential choice (Bhatia & Chaudhry, 2013; Glöckner & Betsch, 2008; Holyoak & Simon, 1999; Mischel & Shoda, 1995; Simon et al., 2004; Thagard, 1989; see also Van Overwalle, 2007). The implicit assumption in constraint satisfaction networks is that long-term activation is necessarily described by a stable state. We find that our implementation of game theoretic decision making in the BAM network restricts long-term activation to the set of rationalizable strategies, even if the network never reaches a stable activation state. In settings in which network stabilizes with the activation of a particular strategy profile, this profile is guaranteed to be some pure strategy Nash equilibrium. Conversely, every pure strategy Nash equilibrium corresponds to some stable state in the BAM network. These results show that rational choice in strategic games can be described by the same cognitive principles at play in nonstrategic settings, and that models based on constraint satisfaction are able to provide a wide range of valuable insights regarding the cognitive basis of human reasoning, judgment, and decision making.

Game Theoretic Decision Making

In strategic games, two or more players make choices over a set of strategies. Crucially, the strategies chosen by the players collectively determine the outcomes of the game, so that each player's utility depends on the other's choice as well as on their own. We define a finite-strategy two-player game with a set of pure strategies for each player, $S_1 = \{s_{11}, \dots s_{1N}\}$ and $S_2 = \{s_{21}, \dots s_{2M}\}$ respectively, and a pair of payoff functions u_1 and u_2 that give each player's utility for each profile of pure strategies (s_{1i}, s_{2j}) (see, e.g., Hart, 1992). Thus if player 1 selects s_{1i} and player 2 selects s_{2j} the utility for player 1 is $u_1(s_{1i}; s_{2j})$ and the utility for player 2 is $u_2(s_{2j}; s_{1i})$. We use the notation $\boldsymbol{u}_{ij}$ as a shortcut for $\left(u_1(s_{1i}; s_{2j}), u_2(s_{2j}; s_{1i})\right)$.

The most standard solution concept for a strategic game is Nash equilibrium, which relies on common knowledge of rationality and accurate (so-called "rational") expectations. A Nash equilibrium is a strategy profile in which no player can obtain higher utility by unilaterally changing her strategy; each player is already playing a best response to the equilibrium strategy profile. We define the set of best responses for player μ to an opponent's strategy $s_{-\mu}$ as $\text{BR}(s_{-\mu}) = \arg \max u_\mu(s_\mu; s_{-\mu})$. Then a pure strategy Nash equilibrium can be defined as a strategy profile (s_{1i}, s_{2j}) such that $s_{1i} \in \text{BR}(s_{2j})$ and $s_{2j} \in \text{BR}(s_{1i})$.

The concept of Nash equilibrium can be generalized to relax the assumption that players somehow have correct expectations about what others will do. The solution concept of rationalizability (Bernheim, 1984; Pearce, 1984) retains the assumption of common knowledge of rationality, but imposes no additional constraints on behavior. As in a Nash equilibrium, players best respond to the strategy they expect their opponent to select, but in contrast to a Nash equilibrium, this expectation is not necessarily correct. Players must only be able to "rationalize" their strategy choice as a best response to one of the opponent's rationalizable strategies. We define the set of rationalizable strategies for each player as the maximal sets R_1 and R_2 such that any $s_1 \in R_1$ satisfies $s_1 \in \text{BR}(s_{2j})$ for some $s_{2j} \in R_2$ and any $s_2 \in R_2$ satisfies $s_2 \in \text{BR}(s_{1i})$ for some $s_{1i} \in R_1$. Clearly, any Nash equilibrium profile is rationalizable, and if the sets of rationalizable strategies are singletons, then these strategies form a Nash equilibrium.

Bidirectional Associative Memory

Nash equilibrium and rationlizability are two of the most important solution concepts in game theory. Here we will examine how choices corresponding to these solution concepts are generated by human decision makers, modeled with constraint satisfaction neural networks. These types of networks rely on bidirectional (recurrent) connectivity between their component nodes, which is able to generate sophisticated dynamics and subsequently explain a range of human behavior. Although these networks have traditionally been used only to explain behavior in non-strategic settings (Bhatia & Chaudhry, 2013; Glöckner & Betsch, 2008; Holyoak & Simon, 1999; Mischel & Shoda, 1995; Simon et al., 2004; Thagard, 1989), they can be easily be applied to strategic game theoretic decision making. Indeed these networks are particularly suitable for this task, as game theoretic decision making features complex interactions between the choices of different decision makers; interactions that can be captured through bidirectional connectivity.

The Bidirectional Associative Memory Network (BAM) is a particularly powerful (and mathematically tractable) constraint satisfaction neural network (Kosko, 1988). It consists of two layers with binary connections between their respective nodes and binary activation functions for any given node. When the connections between its nodes are symmetric then BAM is guaranteed to reach a stable pattern of activation, regardless of its starting state. This property has been used by scholars to solve a variety of practical tasks involving associative memory and pattern completion (Kosko, 1988; see also Cao, 2003 or Gopalswamy & He, 1994 for additional results) and also model human decision making and the biases that it often involves (Bhatia & Chaudhry, 2013).

In this paper, we model how an individual decision maker reasons through two-player finite strategy games, using the BAM network. Strategies in these games for each of the two players can be represented in each of BAM's two layers. We will assume that the first layer in the BAM network represents strategies available to the decision maker (or *self*). If the decision maker can choose from the set of strategies $S_1 = \{s_{11}, \dots s_{1N}\}$, then the first layer in our network consists of N nodes, with node i representing strategy s_{1i}. The activated nodes in this layer represent the strategies that the decision maker considers playing in the game.

Correspondingly we will assume that the second layer in the BAM network represents strategies available to player 2 (or *other*). If player 2 can choose from the set of strategies $S_2 = \{s_{21}, \dots s_{2M}\}$, then the second layer in our network consists of M nodes, with node j representing strategy s_{2j}. The activated nodes in this layer represent the strategies that the decision maker thinks *other* might play in the game.

As mentioned earlier, node activation in BAM is binary, with each node being on or off. We will assume that every node has the same binary activation function, with activation triggered by strictly positive input. For any node k (in either layer of the network) with input I_k, the activation function f_k is specified by the following equation:

$$f_k(I_k) = \begin{cases} 1 \; if \; I_k > 0; \\ 0 \; if \; I_k \leq 0. \end{cases}$$

In a slight abuse of notation, we denote by f the activation function for either layer of the network, with components f_k for every node k.

Connections between nodes are also binary, with each node in the first layer either connected or not connected to each node in the second layer, and each node in the second layer either connected or not connected to each node in the first layer. There are no connections between two nodes in one layer. The network structure can thus be described by the matrices $\mathbf{W^{12}}$ and $\mathbf{W^{21}}$ which represent connections from layer 1 to layer 2, and from layer 2 to layer 1 respectively. We will assume the pattern of connections in our network captures best-responses. Particularly, if $s_{2j} \in \text{BR}(s_{1i})$ then we assume the connection from node i in the first layer to node j in the second layer is $W_{ij}^{12} = 1$. If $s_{1i} \notin \text{BR}(s_{2j})$ then we assume $W_{ij}^{12} = 0$. We assume a similar pattern of connectivity from the second layer to the first, so that $W_{ji}^{21} = 1$ if $s_{1i} \in \text{BR}(s_{2j})$ and $W_{ji}^{21} = 0$ otherwise. Figure 1 provides an illustration of the proposed network.
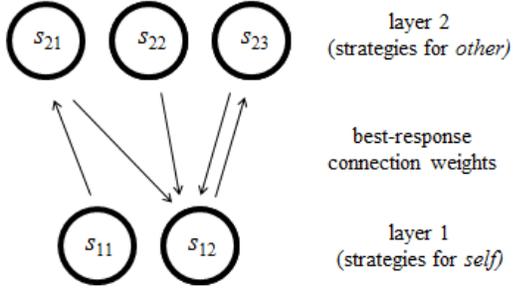
Figure 1: Example of a BAM network encoding a game with two strategies for *self* and three strategies for *other*.

We write the activation of any node i in the first layer, at time t, as $A^1_i(t)$, and any node j in the second layer, at time t, as $A^2_j(t)$. With the connectivity specified above, vectors $\mathbf{A^1}(t)$ and $\mathbf{A^2}(t)$ together represent network activation at time t. We can describe their dynamics in the following equations:

$$\mathbf{A^1}(t) = f(\mathbf{A^2}(t\text{-}1)\cdot\mathbf{W^{21}})$$
$$\mathbf{A^2}(t) = f(\mathbf{A^1}(t)\cdot\mathbf{W^{12}})$$

As formalized in the above equation, node updating in our network is sequential, with layer 1 updating before layer 2. This does not affect the network's behavior, except at the starting point $t = 0$. Here, the above assumption implies that the network begins processing the decision when some nodes in layer 1 are activated exogenously (intuitively, *self* begins the decision process by first considering his strategies). The choice of the starting point activation in our network can affect subsequent node activation and sometimes will determine selection among multiple stable states.

Like related recurrent neural network models, the BAM network can make decisions through constraint satisfaction, that is, by settling into a stable activation state. A stable activation state in the network is a state from which endogenous deviations are not possible. Activation states $\mathbf{A^{1*}}$ in layer 1 and $\mathbf{A^{2*}}$ in layer 2 are stable if $\mathbf{A^{1*}}(t) = \mathbf{A^{1*}}(t+1)$ and $\mathbf{A^{2*}}(t) = \mathbf{A^2}(t+1)$. We assume that a decision maker chooses one of the strategies that are activated in layer 1 and expects *other* to choose one of the strategies activated in layer 2, in the network's stable state.

Note that the connections assumed in this paper are not necessarily symmetric, as strategy s_{1i} can be a best response to strategy s_{2j} without s_{2j} being a best response to strategy s_{1i}. This means that the network is not always guaranteed to stabilize. If the network does not stabilize, then it enters a pattern of oscillating activation in which a certain subset of nodes are activated and deactivated consecutively. We assume that the nodes that are activated (but then deactivated) as part of this oscillating pattern correspond to the set of strategies from which decision makers make their final choice. Nodes that are not activated as part of this oscillating pattern correspond to strategies that are ignored by the decision maker.

Results

First, suppose the network reaches a stable state corresponding to a pure strategy profile, that is, a stable state of activation in which only one node is activated in each layer of the network. Our first result characterizes that stable activation state as corresponding to a Nash equilibrium.

Proposition 1. Suppose that $\lim_{t\to\infty}\left(\mathbf{A^1}(t), \mathbf{A^2}(t)\right) = (\mathbf{A^{1*}}, \mathbf{A^{2*}})$ with unique nodes i and j in each layer for which $\mathbf{A^{1*}}_i > 0$ and $\mathbf{A^{2*}}_j > 0$. Then (s_{1i}, s_{2j}) is a Nash equilibrium.

Proposition 1 will turn out to be a special case of Theorem 1 below, so we omit proof here.

Proposition 1 tells us that the neural network will find a Nash equilibrium if it is able to converge on a single strategy profile. The process through which the network finds this Nash equilibrium is constraint satisfaction. The example of the *traveler's dilemma* in the section below illustrates how the network converges on a pure strategy Nash equilibrium.

The game of *rock-paper-scissors* discussed in the section below illustrates that the network may not converge to a stable state. Even if the network does not reach a stable state, however, we can characterize the nodes which may experience recurrent activation. In the long run, activation is restricted to the set of rationalizable strategies, R_1 and R_2 respectively. Let $\bar{A}_1(t)$ and $\bar{A}_2(t)$ respectively denote the sets of strategies that are activated at time t, i.e., $s_{1i} \in \bar{A}_1(t)$ if $A^1_i(t) > 0$.

Theorem 1. There exists τ such that for any $t > \tau$ we have $\bar{A}_1(t) \subseteq R_1$ and $\bar{A}_2(t) \subseteq R_2$.

Proof. We show that if $s_{1i} \notin R_1$, then for large enough t, $A^1_i(t) = 0$. (The argument for player 2's strategies is analogous.)

If $s_{1i} \notin R_1$, then any chain of best responses can include s_{1i} at most once. Strategy s_{1i} is activated at time t (i.e., $A^1_i(t) > 0$) if and only if there exists $s_{2j} \notin R_2$ such that $s_{1i} \in \text{BR}(s_{2j})$ and $A^2_j(t-1) > 0$. The players only have a finite number ($N+M$) of strategies, so for $t > N+M$, there are no more strategies available to seed a chain of best responses, so $A^1_i(t) = 0$. $\square$

Theorem 1 tells us that we should expect strategically sophisticated individuals only to play rationalizable strategies. The network may never converge to a state with stable activation, so we may not be able to identify a single strategy that will necessarily be chosen, but we can make testable predictions about what will *not* be chosen.

We can also recognize that the lack of a point prediction creates space for contextual factors to matter. An individual's eventual decision may depend on which strategy she considers first, which could in turn depend on the salience of different options, how the options are framed, or how the decision maker's attention is anchored. The *coordination game* discussed in the section below illustrates how the starting point determines which of the multiple Nash equilibria is eventually selected by the network.

It is straightforward to observe that any pure strategy Nash equilibrium would constitute a stable state in our network. The strategies in the Nash equilibrium would, due to the nature of the connection weights, reinforce each other and, once activated, sustain their activation.

Theorem 2. If (s_{1i}, s_{2j}) is a Nash equilibrium, then there exists a stable state $(\mathbf{A^{1*}}, \mathbf{A^{2*}})$ with unique nodes i and j in each layer for which $\mathbf{A^{1*}}_i > 0$ and $\mathbf{A^{2*}}_j > 0$.

Proof. This follows from our assumption that strategies are connected to their best responses. □

Illustrations

In this section we apply the BAM network to three representative games. The games vary in the number of Nash equilibria and in the size of the set of rationalizable strategies. These examples demonstrate that activation in the BAM network in different contexts either may converge to a unique pure strategy Nash equilibrium from any initial state, may fail to converge at all as it oscillates through multiple rationalizable strategies, or may converge to one of many stable profiles depending on the initial state.

Traveler's Dillema

The traveler's dilemma is a generalization of the famous prisoner's dilemma, conceived in order to demonstrate unraveling in a strategic game (Basu, 1994). In the original parable, two travelers have lost identical antiques and must request compensation between $2 and $100. The airline (which is responsible for the lost luggage) will accept the lower claim as valid and pay that amount to both players, and, to deter lying, will penalize the higher claimant with a $2 fee and will reward the lower claimant with $2 bonus. We represent the game with the strategy sets $S_1 = S_2 = \{2, 3, \dots, 100\}$, where x_{1i} and x_{2j} correspond to the dollar amounts associated with strategies s_{1i} and s_{2j}, and with the following utilities:

$$\boldsymbol{u}_{ij} = \begin{cases} (x_{2j} - 2, x_{2j} + 2) & if \ x_{1i} > x_{2j}; \\ (x_{1i}, x_{2j}) & if \ x_{1i} = x_{2j}; \\ (x_{1i} + 2, x_{1i} - 2) & if \ x_{1i} < x_{2j}. \end{cases}$$

The airline's scheme, of course, does not actually reward honesty; it rewards undercutting the other traveler. The best response is always to claim exactly $1 less than the other traveler does (if it is feasible to do so). As it turns out, the only rationalizable strategy for either player is to claim $2, and the unique Nash equilibrium has both players claim $2.

The network connectivity implied by the traveler's dilemma is illustrated in Figure 2. Given enough time this network is guaranteed to stabilize with the activation of the node corresponding to $2 in layer 1, the node corresponding to $2 in layer 2, and the deactivation of all the other nodes (corresponding to higher claims). If, for example, the deliberation process begins with *self* considering claiming $100, i.e., node $100 being activated in layer 1, then node $99 will become activated in layer 2, and in turn node $98 will become activated in layer 1, and so on, until only the nodes corresponding to $2 in each layer are activated. The unique Nash equilibrium corresponds to the only stable state of activation here, because it consists of the only rationalizable strategy for each player. Intuitively, when a decision maker is given enough time to reason, our model predicts that he will choose to claim $2 and will expect the other player to do so as well.
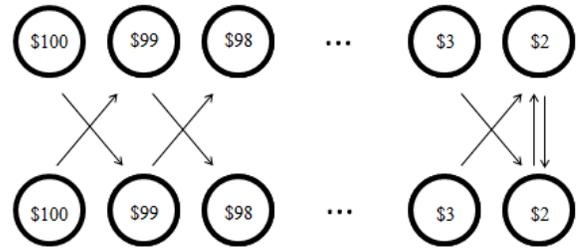

Figure 2: Example of a BAM network encoding traveler's dilemma game.

Rock-Paper-Scissors

The classic game of rock-paper-scissors is the simplest symmetric, zero-sum game with non-transitive winning strategies. Each player has three pure strategies: rock, paper, or scissor. The loop is that rock "defeats" scissors, scissors "defeats" paper, and paper "defeats" rock. If both players play the same strategy, then the game is a tie. We can represent the rock-paper-scissors game with the utilities $\boldsymbol{u}_{ij}$ described in table 1.

Table 1*: Utilities in rock-paper-scissors*

			Other	
		Rock	Paper	Scissors
Self	Rock	(0,0)	(-1,1)	(1,-1)
	Paper	(1,-1)	(0,0)	(-1,1)
	Scissors	(-1,1)	(1,-1)	(0,0)

The game of rock-paper-scissors has no pure strategy Nash equilibrium, but every strategy is rationalizable. Every strategy is a best response to some other strategy, but no strategy is a best response to itself.

The network connectivity implied by the rock-paper-scissors game is illustrated in Figure 3. With a single node initially activated, the network will never stabilize, regardless of which node is the starting point. Instead, the network will display an oscillating pattern of activation, in which each node in each layer activates consecutively. If, for example, the network begins with the activation of the rock node in layer 1, then the node corresponding to paper (the best response to rock) will activate in layer 2. In turn, the node corresponding to scissors will then activate in layer 1. This leads to activating rock in layer 2, and so on. Intuitively, our model predicts that a decision maker will cycle through all three possible strategies as she reasons about the game, and any of these strategies may eventually be chosen.
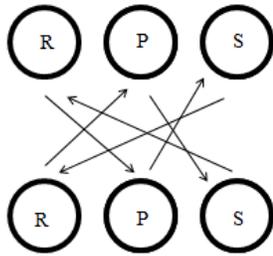

Figure 3: Example of a BAM network encoding rock-paper-scissors game.

Coordination Game

A coordination game captures situations in which the players' primary incentives are to behave similarly, as for example in the case that two friends would like to get together at a meeting place and each has to choose where to go. If they both arrive at the same location, then they each obtain a high reward (e.g., they get to enjoy each other's company). If they arrive at different locations, then they each obtain a low reward or a punishment (e.g., a solitary evening). We can represent a coordination game with $S_1 = S_2$ and a utility function that has the following property:

$$\boldsymbol{u}_{ij} = \begin{cases} (v_{1i}, v_{2j}) \ if \ s_{1i} = s_{2j}; \\ (0,0) \quad if \ s_{1i} \neq s_{2j}. \end{cases}$$

with $v_{1i} > 0$ and $v_{2j} > 0$ for all i and j. In a coordination game, *self* is always incentivized to play the strategy that he expects *other* to play and vice versa, regardless of the specific strategy involved. Thus, there are $N = M$ pure strategy Nash equilibria in the game, with each Nash equilibrium corresponding to an outcome in which *self* and *other* choose the same strategy. As in the rock-paper-scissors game, every strategy is rationalizable. However, whereas the network never stabilizes for the rock-paper-scissors game, it immediately stabilizes for a coordination game.

The network connectivity implied by a coordination game is illustrated in Figure 4. The network is guaranteed to stabilize with the activation of the same strategy nodes in layers 1 and 2. The precise strategies activated in the stable state depend on the starting point in the deliberation process, so that if deliberation begins with the activation of node i in layer 1 (representing strategy s_{1i} for *self*), then the network will stabilize with the activation of node i in layer 1 and the activation of the corresponding node $j=i$ in layer 2. Any pair of nodes corresponding to a Nash equilibrium creates a mutually reinforcing pattern of activation, a stable state in the network. Intuitively, our model predicts that a decision maker will choose to play the strategy that he first begins thinking about and will expect *other* to play this strategy as well.
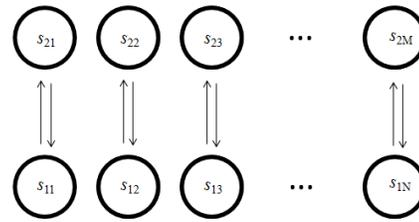

Figure 4: Example of a BAM network encoding coordination game.

Discussion

Constraint satisfaction is a key feature of high-level cognition, and models of constraint satisfaction --often formalized using recurrent neural networks-- are frequently used to study human reasoning, judgment, and decision making. In this paper we extend this research to game theoretic decision making. In particular, we adapt the Bidirectional Associative Memory (BAM) model (Kosko, 1988), a minimal two-layer recurrent neural network, to make decisions in finite strategy two-player games.

BAM is well-suited for this task. Choices in these game theoretic settings are interdependent, with the payoff generated by choosing any one strategy being a function of the choice made by the other decision maker. The recurrent connectivity in BAM can be used to model this type interdependence. We assume that the two layers in the BAM represent strategies available to the *self* and strategies available to the *other*, and connections between these two layers capture the best-responses to the various strategies. With this structure, we show that activation in the BAM network in the long run can only be sustained for rationalizable strategies. Decision making with the BAM network can thus achieve full strategic sophistication (i.e., rational strategic choice), although it dispenses with the assumption of perfect foresight (i.e., rational expectations).

Of course, real people do not exhibit so much strategic sophistication. In the aforementioned traveler's dilemma, for example, people often make higher claims, which cannot be supported by any rationalizable strategy (Capra et al.,

1999). Actual behavior is often better modeled by level-k thinking, which captures bounded rationality (Nagel, 1995). The level-k model assumes that players can reason through *k* steps of best response analysis, where the value of *k* may be heterogeneous in a population. Our proposed BAM network could capture this sort of level-k thinking if we did not associate decisions with asymptotic patterns of activation, but rather terminated the network dynamics at a finite time horizon and associated decisions with activation when this horizon was reached.

Although game theory is most commonly used to describe economic phenomena, such as price setting in markets and bidding in auctions, it also has a large number of applications for the study of cognition and behavior. For example, games such as the traveler's dilemma can help us understand some of the difficulties involved in maintaining cooperation in social settings (Axelrod & Hamilton, 1981). Likewise rock-paper-scissors can be used to represent evolutionary predator-prey dynamics (Nowak & Sigmund, 2004), and the coordination game provides a perspective for understanding the evolution of language, as different individuals have to agree on the meaning of words (Demichelis & Weibull, 2008).

Our paper leaves open the question of strategy learning. We have assumed the network is able to encode best-responses, but how do decision makers learn these best-responses? One solution to this problem may involve a simple form of supervised learning. If, after playing each game, the decision maker is able to infer both the best-response to the strategy that was played by *other*, and *other's* best-response to the strategy that was played by *self*¸ then a variant of the perceptron learning rule (that is restricted to binary connections) could over time allow the network to learn the pattern of best-response connectivity assumed in our BAM model. This could also be accomplished by reinforcement learning, as the individual considers various possible responses to each strategy. In the long run, reinforcement learning would ensure that the response with the highest payoff would be selected for, and the connection to the best response would be the strongest of all connections. Future work should examine these conjectures formally and should compare the predictions of BAM given different forms of learning with real choice observed in experimental settings.

References

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489), 1390-1396.

Basu, K. (1994). The traveler's dilemma: Paradoxes of rationality in game theory. *American Economic Review*, 84(2), 391-395.

Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica,* 1007-1028.

Bhatia, S. & Chaudhry, S. (2013). The dynamics of anchoring in bidirectional associative memory networks. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 1899-1904).

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Cao, J. (2003). Global asymptotic stability of delayed bi-directional associative memory neural networks. *Applied Mathematics and Computation*, 142(2), 333-339.

Capra, C., Holt, C., Goeree, J., & Gomez, R. (1999). Anomalous behavior in a traveler's dilemma? *American Economic Review*, 89(3), 678-690.

Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1), 195-259.

Demichelis, S & Weibull, J. W. (2008). Language, Meaning, and Games: A Model of Communication, Coordination, and Evolution. *American Economic Review*, 98(4), 1292-1311.

Glöckner, A., & Betsch, T. (2008). Modeling option and strategy choices with connectionist networks: Towards an integrative model of automatic and deliberate decision making. *Judgment and Decision Making*, 3(3), 215-228.

Gopalsamy, K., & He, X. Z. (1994). Delay-independent stability in bidirectional associative memory networks. *Neural Networks, IEEE Transactions on*, 5(6), 998-1002.

Hart, S. (1992). Games in extensive and strategic forms. *Handbook of Game Theory with Economic Applications*, 1, 19-40.

Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, 128(1), 3.

Kosko, B. (1988). Bidirectional associative memories. S*ystems, Man and Cybernetics, IEEE Transactions on*, 18(1), 49-60.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: An account of basic findings. *Psychological Review*, 88(5), 375.

Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review*, 102(2), 246.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5), 1313-1326.

Nowak, M. A., & Sigmund, K. (2004). Evolutionary dynamics of biological games. *Science*, 303(5659), 793-799.

Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 1029-1050.

Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint satisfaction. *Psychological Science*, 15(5), 331-336.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, 12(3), 435-502.

Van Overwalle, F. (2007). *Social connectionism: A reader and handbook for simulations*. Psychology Press.