Using Focal Point Learning to Improve Tactic Coordination in Human-Machine Interactions

Inon Zuckerman[1] and **Sarit Kraus**[1] and **Jeffrey S. Rosenschein**[2]
[1]Computer Science Department, Bar-Ilan University, Ramat Gan, Israel
{zukermi,sarit}@cs.biu.ac.il
[2]School of Engineering and Computer Science, The Hebrew University, Jerusalem, Israel
jeff@cs.huji.ac.il

Abstract

We consider an automated agent that needs to coordinate with a human partner when communication between them is not possible or is undesirable (*tactic coordination games*). Specifically, we examine situations where an agent and human attempt to coordinate their choices among several alternatives with equivalent utilities. We use machine learning algorithms to help the agent predict human choices in these tactic coordination domains.

Learning to classify general human choices, however, is very difficult. Nevertheless, humans *are* often able to coordinate with one another in communication-free games, by using *focal points*, "prominent" solutions to coordination problems.

We integrate focal points into the machine learning process, by transforming raw domain data into a new hypothesis space. This results in classifiers with an improved classification rate and shorter training time. Integration of focal points into learning algorithms also results in agents that are more robust to changes in the environment.

1 Introduction

Agents often need to coordinate their actions in a coherent manner. Sometimes, achieving coherent behavior is the result of explicit communication and negotiation. However, communication is not always possible, for reasons as varied as high communication costs, the need to avoid detection, damaged communication devices, or language incompatibility.

Schelling [1963] called coordination-without-communication scenarios *tactic coordination games*, and named these games' "prominent solutions" *focal points*. A classic example is the solution most people choose when asked to divide $100 into two piles, of any sizes; they should attempt only to match the expected choice of some other, unseen player. Usually, people create two piles of $50 each, and that is what Schelling dubbed a focal point.

The use of focal points to solve coordination games has previously been studied, both theoretically and experimentally, from two different perspectives. In the first, there were assumed to be two human coordinators, and research explored a formal theoretical framework for focal points [Janssen, 1998]. In the second approach, there were assumed to be two automated agents, and both agents ran the same focal point finder algorithm [Kraus *et al.*, 2000]. However, the use of focal points when an automated agent needs to coordinate with an arbitrary human partner has yet to be explored, and it raises new challenges.

The main motivation for such research comes from the increasing interest in task teams that contain both humans and automated agents. Human-Agent collaboration can take the form of robots working with human partners. Another scenario is the development of user interfaces that diverge from a master-slave relationship with the user, adopting a collaborative, task-sharing approach in which the computer explicitly considers its user's plans and goals, and is thus able to coordinate various tasks [Grosz, 2004].

One important type of natural human-machine interaction is the *anticipation of movement*, without the need for prior explicit coordination. This movement can be physical, such as the movement of a robotic arm that is assisting a human in a construction task (e.g., a machine helping a human weld pipes). As humans naturally anticipate their partners' choices in certain situations, we would like automated agents to also act naturally in their interactions with humans. Coordinated anticipation can also take place in virtual environments, including online games, where humans and automated agents can inhabit shared worlds and carry out shared activities.

There are several constraints implicit in the above scenarios. First, the human partner with whom our automated agent is trying to coordinate may not always be known ahead of time, and we want coordination strategies suitable for novel partners. Second, the environment itself is not fully specified ahead of time, and may be configured somewhat randomly (although the overall domain *is* known, i.e., the domain elements are a given, but not their specific arrangement). Third, there is no option to "hard-wire" arbitrary coordination rules into all participants, since we are not dealing with coordination between two centrally-designed agents.

We specifically consider environments in which a human and automated agent aspire to communication-free coordination, and the utilities associated with coordinated choices are equal. Clearly, if utilities for various choices differed, the agent and human could employ game theoretic forms of analysis, which might specify certain strategies.

Our integration of machine learning algorithms and focal point techniques (which we call *Focal Point Learning* [FPL]) is done via a semi-automatic data preprocessing technique. This preprocessing transforms the raw domain data into a new data set that creates a new hypothesis space, consisting solely of general focal point attributes. We demonstrate that using FPL results in classifiers (a mapping from a coordination problem to the choice selected by an arbitrary human coordination partner) with a 40% to 80% higher correct classification rate, and a shorter training time, than when using regular classifiers, and a 35% higher rate than when using only classical focal point techniques without applying any learning algorithm. In another series of experiments, we show that applying these techniques can also result in agents that are more robust to changes in the environment.

In Section 2, we describe the motivation and approach of Focal Point Learning. We then describe the experimental setting (Section 3), its definitions (Section 3.1), and the domains that were used in the experiments (Section 3.2). In Section 4 we discuss our results. Related work is considered in Section 5, and we conclude in Section 6.

2 Focal Point Learning

Coordination in agent-human teams can be strengthened by having agents learn how a general human partner will make choices in a given domain. Learning to classify human choices in tactic coordination games, however, is difficult: 1) **No specific function to generalize** — there is no known mathematical function nor behavioral theory that predicts human choices in these games; 2) **Noisy data** — Data collected from humans in tactic coordination games tends to be very noisy due to various social, cultural, and psychological factors, that bias their answers; 3) **Domain complexity** — training requires a large set of examples, which in turn increases required training time. On the other hand, human-human teams are relatively skilled at tactic coordination games. Experiments that examined human tactic coordination strategies [Schelling, 1963; Mehta *et al.*, 1994] showed that people are often able to coordinate their choices on focal points, even when faced with a large set of options.

Several attempts have been made to formalize focal points from a game theoretic, human interaction point of view ([Janssen, 1998] provides a good overview). However, that work does not provide the practical tools necessary for use in automated agents. However, Kraus et al. [2000] identified some domain-independent rules that *could* be used by automated agents to identify focal points (and discussed two approaches for doing so). The following rules are derived from that work, but are adjusted in our presentation (by adding Firstness, and making minor changes in others).

- **Centrality** — give prominence to choices directly in the center of the set of choices, either in the physical environment, or in the values of the choices.
- **Extremeness** — give prominence to choices that are extreme relative to other choices, either in the physical environment, or in the values of the choices.
- **Firstness** — give prominence to choices that physically appear first in the set of choices. It can be either the option closest to the agent, or the first option in a list.
- **Singularity** — give prominence to choices that are unique or distinguishable relative to other choices in the same set. This uniqueness can be, for example, with respect to some physical characteristics of the options, a special arrangement, or a cultural convention.

We employ learning algorithms to help our agent discover coordination strategies. Training samples, gathered from humans playing a tactic coordination game, are used to create an automated agent that performs well when faced with a new human partner in a newly generated environment. However, because of the aforementioned problems, applying machine learning on raw domain data results in classifiers having poor performance. Instead, we use a *Focal Point Learning* approach: we preprocess raw domain data, and place it into a new representation space, based on focal point properties. Given our domain's raw data O_i, we apply a transformation T, such that $N_j = T(O_i)$, where i, j are the number of attributes before and after the transformation, respectively.

The new feature space N_j is created as follows: each $v \in O_i$ is a vector of size i representing a game instance in the domain (world description alongside its possible choices). The transformation T takes each vector v and creates a new vector $u \in N_j$, such that $j = 4\times$[number of choices]. T iterates over the possible choices encoded in v, and for each such choice computes four numerical values signifying the four focal point properties presented above. For example, given a coordination game encoded as a vector v of size 25 that contains 3 choices (c_1, c_2, c_3), the transformation T creates a new vector $u = (c_1^c, c_1^e, c_1^f, c_1^s, c_2^c, c_2^e, c_2^f, c_2^s, c_3^c, c_3^e, c_3^f, c_3^s)$ of size 12 (3 possible choices $\times$ 4 focal point rules), where $c_l^{c/e/f/s}$ denotes the *centrality/extremeness/firstness/singularity* values for choice l. Note that j might be smaller than, equal to, or greater than i, depending on the domain.

The transformation from raw domain data to the new representation in focal point space is done semi-automatically. To transform raw data from some new domain, one needs to provide a domain-specific implementation of the four general focal point rules. However, due to the generic nature of the rules, this task is relatively simple, intuitive, and suggested by the domain itself (we will see such rules in Section 3.2). When those rules are implemented, the agent can itself easily carry out the transformation on all instances in the data set.

3 The Experimental Setting

We designed three domains for experiments in tactic coordination. For each domain, a large set of coordination problems was randomly generated, and the solutions to those problems were collected from human subjects.

We used the resulting data set to train three agents: 1) an agent trained on the original domain data set (a *Domain Data agent*); 2) an agent using focal point rules without any learning procedure (an *FP agent*); and 3) an agent using *Focal Point Learning* (an *FPL agent*). We then compared coordination performance (versus humans) of the three types of agent.

In the second phase of our experiments (which tested robustness to environmental changes), we took the first domain described in Section 3.2, and designed two variations of it;

one variant (*VSD*, a very similar domain) had greater similarity to the original environment than the other variant (*SD*) had. Data from human subjects operating in the two variant settings was collected. We then carried out an analysis of automated coordination performance in the new settings, using the agents that had been trained in the original domain.

3.1 Definitions

Definition 1 (Pure Tactic Coordination Games). Pure Tactic Coordination Games (also called *matching games*) are games in which two non-communicating players get a positive payoff only if both choose the same option. Both players have an identical set of options and the identical incentive to succeed at coordination.

Our experiments involve *pure tactic coordination games*.

Definition 2 (Focality Value). Let R be the set of selection rules used in the coordination domain, $c \in C$ be a possible choice in the domain, $r \in R$ be a specific selection rule, and $v(r, c)$ be its value. Then the *focality value* is defined as:

$$FV(c) = \frac{\sum_{r \in R} v(r, c)}{|R|}.$$

A focality value is a quantity calculated for each possible choice in a given game, and signifies the level of prominence of that choice relative to the domain. The focality value takes into account all of the focal point selection rules used in the coordination domain; their specific implementation is domain dependent (e.g., what constitutes Centrality in a given domain). Since the exact set of selection rules used by human players is unknown, this value represents an approximation based on our characterization of the focal point rule set. In the experiments, our *FP agent* will use this value to determine its classification answer to a given game.

Definition 3 (Focality Ratio). Let C be the set of all possible choices in the coordination domain and $FV(c)$ be the focality value of $c \in C$, *max* be the maximum function and *2nd_max* be the second maximum function. Then the *focality ratio* is defined as:

$$F\_Ratio(C) = \max_{c \in C}(FV(c)) - 2nd\_\max_{c \in C}(FV(c)).$$

A Focality Ratio is a function that gets a set of possible choices and determines the difficulty level of the tactic coordination game. Naturally, a game with few choices that have similar focality values is harder to solve than a game that might have more choices, but with one of the choices much more prominent than the others.

For each of the experimental domains, we built classifiers that predict the choice selected by most human partners. We worked with two widely used machine learning algorithms: a C4.5 decision learning tree [Quinlan, 1993], and a feed forward back-propagation (FFBP) neural network [Werbos, 1974]. Each of these was first trained on the raw domain data set, and then on the new preprocessed focal point data.

The raw data was represented as a multi-valued feature bit vector. Each domain feature was represented by the minimal number of bits needed to represent all its possible values. This simple, low level representation helped standardize the experimental setup with both types of classifiers using exactly the same domain encoding.

The transformation to focal point encoding provides focality values in terms of our low-level focal point rules (Firstness, Singularity, Extremeness, and Centrality) for each of the possible choices. Their values were calculated in a preprocessing stage, prior to the training stage (and by an agent when it needs to output a prediction). It is important to note that following the transformation to the focal point encoding, we deprive the classifier of any explicit domain information during training; it trains only on the focal point information.

Changes in the Environment

To check agent robustness in the face of environment changes, we took the "Pick the Pile" domain (described below), and designed two variations of it, in which one variant is more *similar* to the original environment than the other is:

Definition 4 (Environment Similarity). Similarity between environments is calculated as the Euclidean distance:

$$d_{ij} = \sqrt{\sum_{k=1}^{n} (x_{ik} - x_{jk})^2},$$

where the environment vector x is constructed from the number of goals, number of attributes per goal, number of values per attribute, and the attribute values themselves.

In our experiments, we put original agents (i.e., agents that had been trained in the original Pick the Pile game) in the new environments, and compared their classification performance.

3.2 The Experimental Domains

We now present three experimental domains that were designed to check FPL's performances. The design principles for constructing them were as follows: 1) create *tactic coordination games* (equal utility values for all possible choices); 2) minimize implicit biases that might occur due to psychological, cultural, and social factors; 3) consider a range of tactic coordination problems, to check the performance of FPL learning in different domains.

Pick the Pile Game

We designed a simple and intuitive tactic coordination game that represents a simplified version of a domain where an agent and a human partner need to agree on a possible meeting place. The game is played on a 5 by 5 square grid board. Each square of the grid can be empty, or can contain either a pile of money or the game agents (all agents are situated in the same starting grid; see Figure 1). Each square in the game board is colored white, yellow, or red. The players were instructed to pick the one pile of money from the three identical piles, that most other players, playing exactly the same game, would pick. The players were told that the agents can make horizontal and vertical moves.

In a simple binary encoding of this domain, for encoding 25 squares with 9 possible values (4 bits) per square, we used 100 neurons for the input layer. Training such a massive network required a large training set, and we built the game as a web application to be played online. The web site was publicized in various mailing lists. To maintain the generality

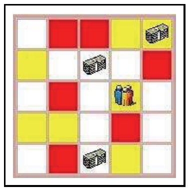

Figure 1: Pick the Pile Game board sample

(1)	Male	25	1.75	Lawyer
(2)	Female	45	1.85	Lawyer
(3)	Male	25	1.78	Doctor
(4)	Male	25	1.78	Engineer
(5)	Female	31	1.75	Lawyer

Figure 2: Candidate Selection Game sample

of the data, each game sequence was limited to 12 game instances; thus, the data set of approximately 3000 games was generated by 250 different human players from around the world. Each instance of the game was randomly generated.

The transformation to the focal point space was done in the following way: the only distinguishable choice attribute is the color, thus the Singularity of each pile was calculated according to the number of squares having the same color. Naturally, a pile of money sitting on a red square in a board having only 4 red squares, would have a higher degree of singularity than a pile of money sitting on a white square, if there were 17 white squares on that board. The Firstness property was calculated as the Manhattan distance between the agent's square and each pile of money. Centrality was calculated as an exact bisection symmetry, thus giving a positive value to a pile that lies directly between two other piles either horizontally, vertically, or diagonally. The Extremeness property was intuitively irrelevant in this domain, so we gave it a uniform constant value.

Candidate Selection Game

Players were given a list of five candidates in an election for some unknown position. The candidates were described using the following properties and their possible values:

1. sex$\in${Male, Female}
2. age$\in${25, 31, 35, 42, 45}
3. height (in meters)$\in${1.71, 1.75, 1.78, 1.81, 1.85}
4. profession$\in${Doctor, Lawyer, Businessman, Engineer, Professor}

Each list was composed of five randomly generated candidates. The (pen and paper) experiments were carried out when subjects (82 first-year university students) were seated in a classroom, and were told that their coordination partners were randomly selected from experiments that took place in other classes, i.e., their partner's identity is completely unknown. For a candidate to be elected, it needs to get these two votes (the player's and its partner's); thus, both sides need to choose the same candidate. To create the necessary motivation for successful coordination,[1] we announced a monetary reward for success. Figure 2 shows a sample question.

The binary encoding for this domain is a set of 65 input neurons in the input layer that encodes 5 candidates, each with 13 possible property values. The focal point transformation had the following intuitive implementation: the Singularity of a candidate was calculated according to the relative uniqueness of each of its values (i.e., a sole female candidate in a set of males will have a high singularity value). The Extremeness property gave high values to properties that exhibit extreme values in some characteristics of the candidate (for example, a candidate who is the oldest or youngest among the set of candidates would get a higher Extremeness value than a candidate who is not). The Firstness and Centrality properties simply gave a constant positive value to the first and third candidates on the list.

Shape Matching Game

Players were given a random set of geometric shapes, and had to mark their selected shape in order to achieve successful coordination with an unknown partner (presented with the same set). The shapes were presented in a single row. The game was randomly generated in terms of the number of shapes (that ranged from 3 to 7) and the shapes themselves (which could be a circle, rectangle, or triangle). Questionnaires containing ten game instances were distributed to students (78 students overall). As before, monetary prizes were guaranteed to students with the highest coordination scores. Figure 3 shows a sample question in the domain.

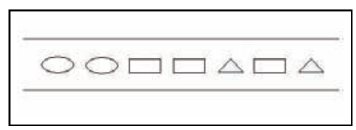

Figure 3: Shape Matching Game sample

This domain is the easiest among our games to represent as a simple binary encoding, because each goal has only a single property, its type. In any game instance, each shape can be a circle, rectangle, triangle, or "non-existing", in the case where the randomized number of shapes is lower than 7. The focal point transformation was implemented as follows: the Singularity of a choice was determined by the number of choices with the same shape (for example, in a game where all shapes are circles and only a single shape is a triangle, the triangular shape will have a high singularity value). The Extremeness property gave higher focality values to the first and last choices in the set. Those values became higher as the number of shapes increased (the extremeness in a game with 3 shapes was lower than the extremeness in a game with 6 shapes). Centrality gave additional focality value to the middle choice, when the number of shapes was odd. In an even number of shapes, no centrality value was given. Firstness gave a small bias to the first shape on the list.

This domain is an example of a transformation in which $j > i$; the transformation actually *increases* the search space.

[1]It is a well-known phenomenon that coordination in these games deteriorates without sufficient motivation.

4 Results and Discussion

4.1 Prediction Performance

For each of the above domains, we compared the correct classification performance of both C4.5 learning trees and FFBP neural network classifiers. As stated above, the comparison was between a domain data agent (trained on the raw domain encoding), a focal point (FP) agent (an untrained agent that used only focal point rules for prediction), and a focal point learning (FPL) agent. "Correct classification" means that the agent made the same choice as that of the human who played the same game.[2]

We optimized our classifiers' performance by varying the network architecture and learning parameters, until attaining best results. We used a learning rate of 0.3, momentum rate of 0.2, 1 hidden layer, random initial weights, and no biases of any sort. Before each training procedure, the data set was randomly divided into a test and a training set (a standard 33%–66% division). Each instance of those sets contained the game description (either the binary or focal point encoding) and the human answer to it. All algorithms were run in the *WEKA* data mining software, which is a collection of machine learning algorithms for data mining tasks. The classification results using the neural network and the decision tree algorithms were very close (maximum difference of 3%). Figure 4 compares the correct classification percentage for the agents' classification techniques, in each of the three experimental domains. Each entry in the graph is a result averaged over five runs of each learning algorithm (neural network and C4.5 tree), and the average of those two algorithms.

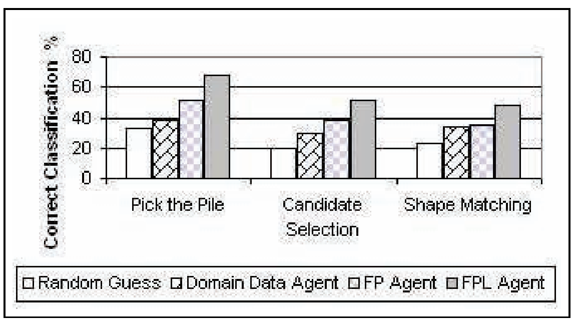

Figure 4: Average Correct Classification Percentage

Examining the results, we see a significant improvement when using the focal point learning approach to train classifiers. In all three domains, the domain data agent is not able to generalize sufficiently, thus achieving classification rates that are only about 5%–9% higher than a random guess. Using FPL, the classification rate improved by 40%–80% above the classification performance of the domain data agent.[3] The results also show that even the classical FP agent performs better than the domain data agent. However, when the FP agent faces coordination problems with low *focality ratios*, its performance deteriorates to that of random guesses.

[2] If there were multiple occurrences of a specific game instance, the choice of the majority of humans was considered the solution.

[3] Since even humans do not have 100% success with one another in these games, FPL is correspondingly the more impressive.

Note also that in the first domain, when using FPL instead of regular raw data learning, the marginal increase in performance is higher than the improvement that was achieved in the second domain (an increase of 30% vs. 21%), which is in turn higher than the marginal increase in performance of the third domain (an increase of 21% vs. 14%). From those results, we hypothesize that the difference in the marginal performance increase is because the first domain was the most complex one in terms of the number of objects and their properties. As the domain becomes more complex, there are more possibilities for human subjects to use their own subjective rules (for example, in the Pick the Pile domain, we noticed that few people looked at the different color patterns that were randomly created, as a decision rule for their selected goal). As more rules are used, the data becomes harder to generalize. When an agent is situated in a real, highly complex environment, we can expect that the marginal increase in performance, when using FPL, will be correspondingly large.

An additional advantage of using FPL is the reduction in training time (e.g., in the *Pick the Pile* domain we saw a reduction from 4 hours on the original data to 3 minutes), due to the reduction of input size. Moreover, the learning tree that was created using FPL was smaller, and can be easily converted to a rule-based system as part of the agent's design.

4.2 Robustness to Environmental Changes

Follow-on experiments were designed to check the robustness of agent-human tactic interaction in a changing environment.

We created two different versions of the *Pick the Pile* game, which had different similarity values relative to the original version. In the first variant (*VSD*), we added a fourth possible value to the set of values of the color attribute (four colors instead of three). In the second variant (*SD*), in addition to the first change, we also changed the grid structure to a 6 by 4 grid (instead of the original 5 by 5). Moreover, in both variants, we changed all four color values from actual colors to various black and white texture mappings.

Additional experiments were conducted in order to collect human answers to the two new variants of the game. The agents that had been trained on the original environment (using the neural network algorithm), were now asked to coordinate with an arbitrary human partner in the new environments. Figure 5 summarizes performance comparison of the agents in each of the new environment variants.

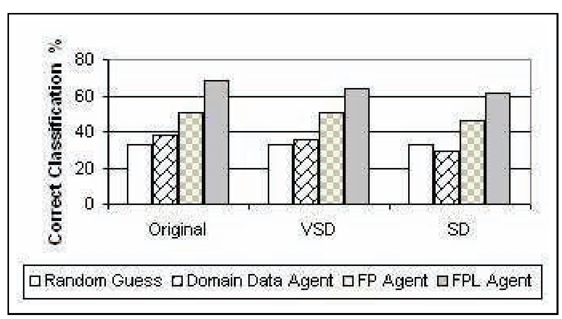

Figure 5: Classification Percent in Changing Environments

The prediction results on the first variant (VSD) show that all three agents managed to somehow cope with the new, *very similar domain*, and suffered only a small decrease in performance. However, when looking at the results of the *similar domain* (SD), we see that the domain data agent's performance decreased all the way to its classification performance's lower bound, that of random guessing. At the same time, our FPL agent did suffer a mild decrease in performance (around 5%), but still managed to keep a considerably high performance level of around 62%. We can also notice that the classical FP agent copes with the environmental changes better than the domain data agent, with performance level of around 45%; however, it is still low when compared to the FPL agent's performance level.

5 Related Work

The first use of focal points in artificial intelligence [Kraus *et al.*, 2000] employed focal point techniques to coordinate between two computerized agents in communication-impoverished situations. The authors modeled the process of finding focal points from domain-independent criteria using two approaches: decision theory, and step-logic. They also defined a *robot rendezvous* domain, and used simulations to show that their focal point-based algorithm, tailored specifically for that domain, effectively solved the problem. Focal points have also been widely studied from the point of view of a human who is coordinating with another human partner. [Janssen, 1998] provides a good overview of the formal models that attempt to describe the focal point phenomenon.

For teams of automated agents, a variety of methods have been developed to achieve coordination, and cooperation without communication, for specific, well-defined tasks. [Gervasi and Prencipe, 2004] provides an example solution to the *flocking problem*, involving simple robot teams that cannot communicate with one another. A comparison of experiments with and without communication was done in [Arkin, 1992], to check robot performance in the *retrieval task*. In a recent paper, [Schermerhorn and Scheutz, 2006] used predefined social laws for achieving coordination without communication in the multiagent territory exploration task.

6 Conclusions

We have presented the *Focal Point Learning* (FPL) approach to building automated agents that play tactic coordination games with general human partners in changeable environments. The technique makes use of learning algorithms to train agents to coordinate with general human partners in specific domains; focal points are integrated into the learning process through the use of focal point selection rules. Training data is preprocessed and transferred into a new representation space, where each vector contains quantified focal point values, and these are used to train the agent.

We created three experimental domains, and collected data to be used for training agents to predict human coordination choices in those domains. Results showed that when trained solely on the domain-encoded data, the classifiers resulted in a close-to-random correct classification percentage, while the FPL agents managed to achieve a significantly higher correct classification rate.

In the next step we created two variants of one of the domains, collected human data for them, and then checked the coordination performance, in these variants, of each of the agents that had been trained in the original domain. Here again, the FPL agents outperformed the others, and demonstrated robustness in a changing environment.

When building agents to coordinate with unfamiliar human partners, without communication, machine learning classifiers have a difficult task generalizing data to predict human choices. Focal point learning can improve performance and robustness to environmental changes. In future work we will explore FPL's performance on more complex problems. The first stage will likely be problems in which the agent has some uncertainty regarding the environment that the human sees.

7 Acknowledgment

This work was partially supported by grants #1211/04 and #898/05 from the Israel Science Foundation.

References

[Arkin, 1992] R. C. Arkin. Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3):351–364, 1992.

[Gervasi and Prencipe, 2004] V. Gervasi and G. Prencipe. Coordination without communication: The case of the flocking problem. *Discrete Applied Mathematics*, 144(3):324–344, December 2004.

[Grosz, 2004] Barbara J. Grosz. Beyond mice and menus. In *Proceedings of the American Philosophical Society*, 2004.

[Janssen, 1998] M. C. W. Janssen. Focal points. In P. Newman, editor, *The New Palgrave of Economics and the Law*, pages 150–155. MacMillan, 1998.

[Kraus *et al.*, 2000] Sarit Kraus, Jeffrey S. Rosenschein, and Maeir Fenster. Exploiting focal points among alternative solutions: Two approaches. *Annals of Mathematics and Artificial Intelligence*, 28(1–4):187–258, October 2000.

[Mehta *et al.*, 1994] J. Mehta, C. Starmer, and R. Sugden. The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, 84(3):658–673, 1994.

[Quinlan, 1993] J. Ross Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, 1993.

[Schelling, 1963] T. Schelling. *The Strategy of Conflict*. Oxford University Press, New York, 1963.

[Schermerhorn and Scheutz, 2006] Paul Schermerhorn and Matthias Scheutz. Social coordination without communication in multi-agent territory exploration tasks. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 654–661, Hakodate, Japan, May 2006.

[Werbos, 1974] Paul Werbos. *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD thesis, Committee on Applied Mathematics, Harvard University, Cambridge, MA, November 1974.