

Using Focal Point Learning to Improve Human-Machine Tacit Coordination*

Inon Zuckerman¹, Sarit Kraus¹, Jeffrey S. Rosenschein²

¹Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel
{zukermi, sarit}@cs.biu.ac.il

²The 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 (*tacit 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 tacit coordination domains.

Experiments have shown that 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 point rules into the machine learning process, by transforming raw domain data into a new hypothesis space. We present extensive empirical results from three different tacit coordination domains. The *Focal Point Learning* approach results in classifiers with a 40% to 80% higher correct classification rate, and shorter training time, than when using regular classifiers, and a 35% higher correct classification rate than classical focal point techniques without learning. In addition, the integration of focal points into learning algorithms results in agents that are more robust to changes in the environment. We also present several results describing various biases that might arise in Focal Point based coordination.

*A preliminary version of this article appeared in the Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI 2007).

1 Introduction

One of the central problems in multi-agent systems is the problem of coordination. Agents often differ, as do humans, in their subjective view of the world and in their goals, and need to coordinate their actions in a coherent manner in order to attain mutual benefit. Sometimes, achieving coherent behavior is the result of explicit communication and negotiation [35, 20]. 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.

Several methods have been developed for achieving coordination and cooperation without communication, for teams of automated agents in well-defined tasks. The research presented in [9] provides a solution to the *flocking problem*, in which robots need to follow their leader. The robots, grouped into mobile teams, move in a two-dimensional space and cannot communicate with one another. A comparison of experiments with and without communication in the *retrieval task*, in which agents need to scout and retrieve resources, was presented in [1]. In [30], agents used predefined social laws for achieving coordination without communication in the multiagent *territory exploration task*. All of the above research considers specific methods that are tailored for a single, well-defined task, and for pure autonomous agent teams (i.e., humans do not take part in the interactions).

In experimental research presented by Thomas Schelling [29], it was shown that people are often able to successfully solve coordination-without-communication scenarios (which he named *tacit coordination games*) in an impressive manner, usually with higher coordination rates than that predicted by decision theoretic analysis [4]. It appears that in many of those games there is some sort of “prominent solution” that the players manage to agree upon without communication, and even without knowing the identity of their coordination partner. Those “prominent solutions” were named *focal points* by Schelling (and are also sometimes referred to as *Schelling points*).

A classic example of focal point coordination is the solution most people choose when asked to divide \$100 into two piles, of any size; they should attempt only to match the expected choice of some other, unseen player. More than 75% of the subjects in Schelling’s experiments created two piles of \$50 each; that solution is what Schelling dubbed a focal point. At the same time, using decision theory would result in a random selection among the 101 possible divisions, as the (straightforward) probability distribution is uniform.

Previous coordination-without-communication studies were directed at coordinating a team of automated agents; the main motivation for our research, however, comes from the increasing interest in task teams that contain both humans and automated agents [10]. In such cases, augmenting an automated agent with a mechanism that imitates *focal point* reasoning in humans will allow it to better coordinate with its (human) partner.

Human-Agent collaboration can take the form of physical robots or of software agents that are working on a task with human partners ([10] provides a good survey). For example, the “DEFACTO” system [31] used artificial agents in the fire-fighting domain to train incident commanders. In the area of space exploration, NASA has explored the possibilities of having collaborative agents assisting human astronauts in

various activities [23, 32, 27]. Another scenario is the development of user interfaces that diverge from a limited master-slave relationship with the user, adopting a more collaborative, task-sharing approach in which the computer explicitly considers its user’s plans and goals, and is thus able to coordinate various tasks [11].

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 [27]). As humans naturally anticipate their partners’ choices in certain situations, we would like automated agents to also act naturally in their interactions with humans [13]. Coordinated anticipation can also take place in virtual environments, including online games and military simulations, where humans and automated agents (“synthetic forces” in their terminology) can inhabit shared worlds and carry out shared activities [18].

Regardless of the specific problem at hand, there are several general constraints implicit in the above scenarios:

- 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.
- 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).
- 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, such as Nash equilibria selection (e.g., [12, 15]) which might specify certain strategies. However, game theory does not address the problem of choosing among multiple choices with equivalent utility, all other aspects being equal, in a tacit coordination game.¹

In this paper, we present an approach to augmenting the focal point mechanism in human-agent interactions through the integration of machine learning algorithms and focal point techniques (which we call *Focal Point Learning* [FPL]). The integration 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. The transformation is done according to four general focal point rules: Firstness, Centrality, Extremeness, and Singularity, and their intuitive interpretation in the coordination domain.

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

¹Even the question of how to choose among multiple Nash equilibria is not necessarily straightforward.

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.

We begin by providing background on focal points in Section 2. In Section 3, we describe the novel Focal Point Learning approach. We then describe our experimental setting in Section 4, its definitions, methodology, and the domains that were used in the experiments. Next, in Section 5, we discuss the robustness of our agents to dynamically changing environments. Additional experimental results and insights on the nature of focal points is discussed in Section 6, and we conclude in Section 7.

2 Focal Points

Focal points were introduced by Schelling in [29] as a prominent subset of solutions for *tacit coordination games*, which are coordination games where communication is not possible. In such games (also known as *matching games* in game theory terminology) the players only have to **agree** on a possible solution, regardless of the solution itself. In other words, they receive a reward by selecting the same solution, regardless of the solution. When their solutions differ, both players lose and do not get any reward. A solution is said to be “focal” (also “salient”, or “prominent”) when, despite similarity among many solutions, the players somehow converge to this solution.

2.1 Focal Point Examples

To better understand the notion of focal points, we will now review several coordination tasks that were investigated by Schelling in his original presentation [29].

The classic example, presented above, is a coordination task in which two players need to divide a pile of 100 identical objects (e.g., 100 coins) into two piles. A player’s only concern is that his objects should be divided in the same way as the other player’s objects, regardless of the piles’ sizes. Schelling found that players with a strong incentive for success would divide the pile into two identical piles of 50 objects each. The player’s reasoning process would dictate that, as at the basic level of analysis all choices are equivalent (that would be the expected analysis when applying a straightforward decision theoretic model), the players must apply higher-level reasoning by focusing on some property that would distinguish a particular choice—and at the same time, rely on the other person’s doing likewise. Here, the property that causes the 50–50 choice to be more “prominent” than others can be regarded as a symmetric uniqueness property.

In another example, Schelling asked his subjects to coordinate by naming a positive integer. If both players select the same number, they both get a positive reward (here again, regardless of the number itself), otherwise, they get nothing. His results show that despite there being an infinite number of positive integers, players did manage to converge to a few potential choices, and often coordinate. The most prominent choice in this experiment (which got $\frac{2}{5}$ of the answers) was the number 1. This number has an obvious property that distinguishes it from others, as it is the smallest positive integer.




Figure 1: 3x3 Coordination Grid

At times, physical or geographical ordering of the environment can help focus choices. For instance, the coordination task in Figure 1 is to check one square on a 3x3 grid board; again, the only need is to coordinate with another player regardless of the square itself. Here, most people manage to do better than the $\frac{1}{9}$ predicted by straightforward mathematical analysis. Most people selected the central square as it is considered prominent according to various subjective properties related to symmetry and centrality.

However, a small change in the environment can result in a more challenging coordination task. Looking at Figure 2, we can notice that now there is no prominent solution according to the symmetry and centrality properties that were found in the previous version. However, Schelling's experimental results in this task suggest that subjects converged to the upper-left square as the prominent focal point, and generally speaking most selections were on the squares residing in the upper-left to lower-right diagonal.




Figure 2: 4x4 Coordination Grid

The focal point phenomena can be observed in various coordination domains: finding a meeting place at an airport, where to leave a note for a spouse, voting for the same candidate in an election. However, the underlying idea is that the players are motivated to coordinate and that they are doing so by a kind of higher-order reasoning: reasoning about what the other player would reason about me.

2.2 Related Work

Schelling [29], after the presentation of his experimental results, claimed that when searching for prominent solutions, there are two main components: Logic and Imagination. Logic is some logical explanation for a choice (for example, choosing 1 when asked to pick a positive integer, because it is the smallest positive integer). Imagination

includes the unknown predisposition that makes people tend to choose “Heads” over “Tails” in a simple Heads or Tails coordination game.

2.2.1 Game Theory

The problem of selecting a choice among alternatives with equal utility values is also present in game theory. There, interactions are often represented as normal form games, in which a matrix is used to represent the players’ strategies, and each player’s payoffs are specified for the combined actions played inside the matrix. This type of representation allows us to find dominating strategies and different sorts of equilibrium points. One example is the Nash equilibrium [26]: two strategies S_1 and S_2 are said to be in Nash equilibrium if, assuming that one agent is using strategy S_1 , the best the other agent can do is to use S_2 . When the coordination game has a single equilibrium point one might argue that it should be selected, but there are games where there are multiple equilibria. In Figure 3 there are two equilibria: one for strategies ac , and the other for strategies bd .

2x2 game		Player 2	
		Action c	Action d
Player 1	Action a	(2, 1)	(-1, -1)
	Action b	(-1, -1)	(1, 2)

Figure 3: Normal Form 2x2 Game with Two Equilibrium Points

Game theory provides various solutions to cases where the payoff matrix is asymmetric [15, 12]; other solution concepts deal with the evolution of equilibria in games played repeatedly within a specific population [19, 36, 37], or which iteratively converge on a pattern of coordinated play [6]. However, none of the above solutions from game theory addresses the problem of solving non-repeated, multiple equilibria, symmetric coordination games without communication, such as tacit coordination games.

		Player 2			
		a_2	b_2	...	n_2
Player 1	a_1	(1, 1)	(0, 0)	...	(0, 0)
	b_1	(0, 0)	(1, 1)	...	(0, 0)
	\vdots
	n_1	(0, 0)	(0, 0)	(0, 0)	(1, 1)

Figure 4: n -Action, Two-Player Tacit Coordination Game

A tacit coordination game for two players can be presented as the following normal form matrix (see Figure 4). In this example, each player has a set of n possible actions, labeled from a to n (with the player number as a subscript). In this game, the players get a payoff only if they are able to agree on an action, regardless of the action itself. The divide-100-objects focal point example (Section 2.1) can be seen as a similar matrix where $n = 101$, for the number of possible strategies. Yet again, while game theory

does not provide a solution for such cases, human beings are often able to do better than might be predicted by decision theory.

2.2.2 Labeling Theories

Game theory lacked a formal model that would explain the findings in Schelling's experiments. Gauthier [8] was the first who addressed this topic. He introduced the notion of Saliency when rational players are engaged in coordinated interaction. Players seek salience in order to increase the equilibrium utility value (i.e., distinguishing one of the choices) using additional knowledge about the situation by forming expectations about what I expect you to expect from me. This additional knowledge is the player's own description of the world (which is not handled in classical game theory's analysis), according to the way it is conceived by him. The player, when making a choice, follows a *principle of coordination*, which ensures that the most distinguished equilibria are selected.

Following Gauthier, Bacharach [2] introduced the notion of *availability*, which is the probability that the player will conceive of certain aspects of the coordination problem. For example, given a coordination problem where one has to select one brick from a set of eight bricks, it is easy for the player to notice the color of the bricks (thus, this dimension will have high availability), but it might be hard to notice that a single brick is made of a different material (this dimension will have lower availability). Bacharach argued that the number of possible choices (or equilibria) is given according to the number of aspects the player grasps in the coordination problem. Bacharach and Bernasconi [3] presented a variable frame theory (VFT), where they refer to features as *frames* which subjectively describe the *game* to each player. Frames are sets of variables that are used to conceptualize the game.

Janssen [17] continued building on Bacharach's model, generalized it for general classes of dimensions, and showed that players in all cases would receive a higher payoff by following his selection principle rather than neglecting the label information.

Sugden [33] presented a different theory by showing how labels can influence decisions in pure coordination games. He argued that the labeling of choices is beyond the conscious control of the player, and is influenced by psychological and cultural factors. He showed that his *collective rationality principle* may or may not create coordination, depending on the labeling procedure correlation.

The economic theories presented above all give appropriate retroactive justification to answers for coordination games, but are not highly descriptive of human behavior, nor are they applicable to our mixed human-agent scenario, for several reasons. First, they do not give any consideration to social conventions and cultural factors [22, 36, 37].² Second, their analysis does not quantify the notion of availability, nor give the labeling conceived of by players. Moreover, these theories have no explanatory power; for example, in the Heads/Tails question, most people choose Heads. This could be explained by saying that Heads has a conventional priority over Tails, but the same explanation could have been used if the majority had picked Tails.

²We will see later that these are very important in the process of focal point discovery.

2.2.3 Experimental Work

There has been some experimental research that has tried to establish the existence of focal points in coordination games.

Mehta et al. [25], in a series of controlled experiments, managed to verify people's ability to coordinate their answers. In one experiment, subjects were divided into two groups. Group A was instructed to try and coordinate their answers; Group B was instructed to give a response without any incentive to match their partner's choice. The questions were generally "open" coordination questions, without a well-defined set of possible answers, for example, "Choose a year", or "Choose a mountain". Results showed that Group A had a significantly higher success rate in coordinated answers than Group B, which had no incentive for coordination.

In the second experiment, the focus was on assignment games where the purpose was to isolate some focal point selection rules that the authors hypothesized would be used in the coordination problem: (1) closeness, (2) accession, or (3) equality. The game was as follows: there was a board containing two squares and an arbitrary number of circles. The squares were always located at the same position, and the circles were positioned at different places on the board. The task was to assign each circle to a square (by painting the circles red or blue), and to be able to coordinate your assignment with another player. The results showed support for the hypothesis that subjects used the rules mentioned above. The authors claimed, "It seems clear that subjects are drawing on each of the three rules to identify focal points". However, in an analysis that we did of each of the rules, it appears that the accession rules had a very limited impact on the overall results, in comparison to the other rules.

Mehta published another experiment [24], in which she interviewed players, after completing various coordination games, about their behavior in the games. There were some insights that were common to the majority of the interviewed subjects: (a) subjects tried "going into the heads" of their partner, to figure out what he would do; (b) subjects tried using cultural information shared with their partner; (c) the use of rules—most subjects followed some rules in choosing their answer; (d) the use of "fairness" played a significant role.

Another set of experiments, which strengthens the notion of focal point usage, was done by Van Huyck et al. [14], studying games with multiple equilibria; they checked how human subjects make decisions under conditions of strategic uncertainty. Cooper et al. [5] experimented to discover what happens when players play a non-cooperative game with multiple Nash equilibria. His results showed that the outcome would come from the set of Nash equilibria.

In the artificial intelligence literature, Kraus et al. used Focal Point techniques to coordinate between agents in communication-impooverished situations [7, 21]. In [21] they modeled the process of finding focal points from domain-independent criteria using two approaches: decision theory, and step-logic. They devised a focal point algorithm tailored to a *robot rendezvous* coordination game (where two robots have to agree on a single object from a small set of objects with various properties), and showed that their algorithm managed to converge to focal points in a very high percentage of cases. However, though their approach initiated the use of focal points in the agent-agent coordination problem, their results are not very surprising: two agents running the same

algorithm would necessarily converge to the same solution (though the authors’ approach was not hard-wired). We consider that a major advantage that automated agents could gain from using focal points is specifically when working with human beings, who inherently seem to exhibit such reasoning capabilities.

3 Focal Point Learning

To enhance human-agent coordination, we would like the automated agent to have a cognitive mechanism which is similar to the one that exists in human beings. Such a mechanism would allow agents to reason and search for focal points when communication-impooverished situations occur.

Coordination in human-agent teams can be strengthened by having agents learn how a general human partner will make choices in a given domain. It is possible to use various machine learning algorithms to explicitly learn the focal choices in a given game. However, learning to classify the choices of a general human partner in tacit coordination games is difficult for the following reasons:

1. **No specific function to generalize** — there is no mathematical function nor behavioral theory that predicts human choices in these games. Specifically, no function can capture the notion that for some tacit coordination games, different human players can select different choices.
2. **Noisy data** — data collected from humans in tacit coordination games tends to be very noisy due to various social, cultural, and psychological factors that bias their answers. When collecting data, any experience before or during the game can impact the focal points (we will see examples of that phenomenon below).
3. **Domain complexity** — in complex domains, training a classifier not only requires a large set of examples, but in order to generalize an arbitrary human partner, those examples should be taken from different sources in order to remove cultural and psychological biases. This results in a very difficult data collection task.

These difficulties suggest that using classical machine learning methods to build a focal point reasoner, which works similarly to that exhibited by humans, is not an easy task. As we will see in the experimental section below, the main problem is that we want to classify a “general” human coordination partner and not a specific partner (which would be a considerably easier task for classical machine learning algorithms).

As mentioned above, several attempts have been made to formalize focal points from a game theoretic, human interaction point of view ([16] provides a good overview). However, as we said, that research does not provide the practical tools necessary for use in automated agents. In [21], Kraus et al. identified some domain-independent rules that *could* be used by automated agents to identify focal points. The following rules are derived from that work, but are adjusted and refined in our presentation.³

³Kraus et al. used the following intuitive properties: Uniqueness, Uniqueness complement, centrality, and extremeness.

- **Centrality** — this rule gives prominence to choices directly in the center of the set of choices, either in the physical environment, or in the values of the choices. For example, in the 3x3 grid coordination scenario (Figure 1, above), the center square had the centrality property as it resides directly in the center of the physical environment (both horizontally and vertically).
- **Extremeness** — this rule gives prominence to choices that are extreme relative to other choices, either in the physical environment, or in the values of the choices. For example, when asked to select one number out of a set, the highest or smallest numbers will have the extreme property. In a physical environment, the tallest, smallest, longest, etc., can be named as the extreme choices.
- **Firstness** — this rule gives 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. For example, when asked to select one number out of a set, the number that appears first has the firstness property.
- **Singularity** — this rule gives 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. There are many examples of this rule, from a physical property such as the object’s color or size, to some social norm which singles out one of the options.

We employ learning algorithms to help our agent discover coordination strategies. Training samples, gathered from humans playing a tacit 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 $j = 4 \times [\text{number of choices}]$.⁴ 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 three 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 and the number of rules used.

⁴This can be generalized to a different number of rules by taking $j = [\text{number of rules}] \times [\text{number of choices}]$.

Input: Original coordination task Encoding
Output: Focal Point based coordination task Encoding
 $V[num'Choices \times num'Rules]$;
foreach $c \in choices$ **do**
 foreach $r \in Rules$ **do**
 $V[c \times r] = ComputeRule(r,c)$;
 end
end
return V ;

Algorithm 1: The General Transformation Algorithm

The transformation from raw domain data to the new representation in focal point space is done semi-automatically using Algorithm 1. This linear-time algorithm is a general transformation algorithm to any number of focal point rules. In order to transform raw data from some new domain, one needs to provide a domain-specific implementation of the four general focal point rules. There is currently no automated way to suggest the optimal set of rules for a given problem domain, and we will not claim below that our choices are optimal. 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 4.3). When those rules are implemented, the agent can itself easily carry out the transformation on all instances in the data set.

4 The Experimental Setting

We designed three domains for experiments in tacit 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 types of agents, and compared their coordination performance (versus unknown human partners). The agent types are as follows:

1. *Domain Data agent* — an agent trained on the original domain data set.
2. *Focal Point agent (FP agent)* — an agent using focal point rules without any learning procedure.
3. *Focal Point Learning agent (FPL agent)* — an agent using the *Focal Point Learning* method.

In the second phase of our experiments we tested robustness to environmental changes (Section 5). We took the first domain described in Section 4.3, 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*, a Similar Domain) had. Data from human subjects operating in the two variant settings were 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. In addition to the main results, we

will discuss below several insights into the nature of focal points that came from different stages of the experiments, pre-experiments, and while interviewing the subjects after they participated in the experiments.

4.1 Definitions

Definition 1 (Pure Tacit Coordination Games). Pure Tacit 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 involved pure tacit coordination games. We demonstrate the definitions below with the following example of such a game. Two non-communicating players are faced with the following set of numbers: $\{ 15, 18, 100, 8, 13 \}$. Their instructions are to select a single number of that set, where successful coordination is rewarded with \$50 for each player (regardless of the coordination choice); in the case of unsuccessful coordination, the players receive nothing.

Obviously, in the above example, straightforward decision theory would suggest that all five options have similar probability ($p(c) = 0.2$, where c is a choice), for successful coordination; from a game theoretic point of view, all five choices/actions have arbitrary labels and an arbitrary ordering. However, we have reason to believe that using focal point reasoning, we will be able to coordinate with a human partner with higher probability than the expected 0.2.

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.

Going back to our running example, we can now calculate the focality value for each of the possible choices of the coordination problem (five choices). The first step is to provide an intuitive domain-dependent implementation to the suggested focal point rules:

- Centrality — will increase prominence of the central choice of that set (the number 100);
- Extremeness — will be defined to increase prominence of the two extreme choices: the smallest and largest numbers (numbers 100 and 8);

- Firstness — the intuitive implementation would be to increase the prominence of the first number in the list (number 15);
- Singularity — in this setting can be defined using odd/even division (15 and 13), prime numbers (only 13 is prime), or simply according to the number of digits (according to this property, 8 and 100 are the distinguished choices).

We should choose one or more singular properties which are believed to be intuitive interpretations that will be used by most humans.

We can give different weight to different rules, but here we continue the example using uniform weight for the rules. The following list will specify the *FP* values according to the above rules of all choices:

FP(15) = 2 — using the firstness and singularity property with the parity division.

FP(18) = 0 — none of the above rules are prominent for this choice.

FP(100) = 3 — using the centrality, extremeness and singularity rules according to the number of digits.

FP(8) = 2 — using the extremeness and singularity rules according to the number of digits.

FP(13) = 2 — using singularity according to parity and prime numbers.

According to our specific implementation, a human player can have the following observations: choosing 18 is the least recommended choice as its focality value is the lowest. The most prominent choice would be the number 100, with the highest focal value, 3. Naturally, different interpretations as to the rules, and different weights, would result in different focal values. However, as we see in our experimental section, the most intuitive descriptions using those rules will assist us in focusing the answers, or will at least help us in eliminating some of the options (e.g., in our example 18 can be easily pruned, leaving us with an easier coordination task).

Definition 3 (Focality Difference). 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 difference* is defined as:

$$F_Diff(C) = \max_{c \in C}(FV(c)) - 2nd_max_{c \in C}(FV(c)).$$

A focality difference is a function that gets a set of possible choices and determines the difficulty level of the tacit 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.

In our example, the focality difference is $F_Diff(example) = 3 - 2 = 1$. This difference allows us to compare different tacit coordination games and understand which game is easier to solve using a focal-point-based answer. The higher the focality difference is, the easier it is to coordinate on a focal answer. Moreover, we can see that the focality difference is not a function of the cardinality of the set of possible choices.

4.2 Methodology

For each of the experimental domains presented below, we used the same methodology. First, we collected a large set of samples from different human players. Each such sample was a randomly generated instance of the coordination domain; thus, there were instances that were generated more than once and were “played” by different players.

The next step was to build machine learning 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 [28], and a feed-forward back-propagation (FFBP) neural network [34]. Obviously, the different domains have different numbers of input and output neurons, thus they require using different network architectures.

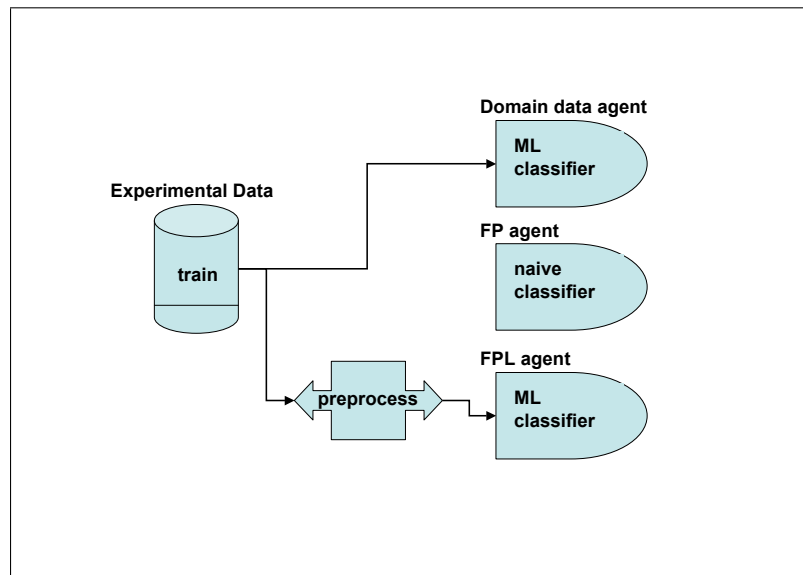


Figure 5: The Training Stage

Each of these was first trained on the raw domain data set, and then on the new preprocessed data based on the focal point rules. Figure 5 describes the training stage that was done for each of the experimental domains. As can be seen, the domain data agent is trained on the raw experimental data, the FP agent does not undergo any training, and the FPL agent is trained on the preprocessed 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 of 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). In

the training session, the algorithms learn the best values and weight for each rule, as the individual impact of each individual rule may vary across domains. 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.

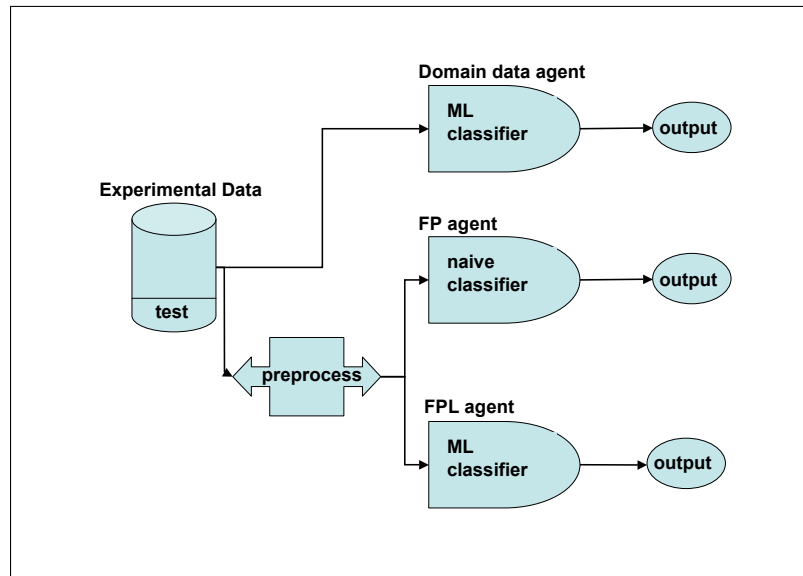


Figure 6: The Testing Stage

Finally, we compared the performance of our three agents in each domain according to their correct classification of the test samples. This process is described in Figure 6, where we can see that the test examples are fed in their original encoding to the domain data agent, while the FP and FPL agents classify the example after it has gone through the preprocessing phase. For a given game instance, each of the agents outputs one of the possible choices, the one that it predicts most people will select.

4.3 The Experimental Domains

We now present three experimental domains that were designed to check FPL's performance. We designed the coordination games with the following design principles in mind:

1. Make the domains *tacit coordination games* (equal utility values for all possible choices).
2. Avoid implicit biases that might occur due to psychological, cultural, and social factors (i.e., remove possible biases).
3. Use a variety of tacit coordination problems to check the performance of focal point learning in different domains.

4.3.1 Pick the Pile Game

We designed a simple and intuitive tacit 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. 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 7). 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.




Figure 7: Pick the Pile Game Board Sample

Data was collected using an Internet website (Figure 8) which allowed players from all over the world to participate in the game, and their answers were recorded. Upon entering the website, each player was requested to read the instructions and was asked to play the game only one time. The instructions specified that each player is paired with an unknown partner and that their score would be given at the end. Each game session was constructed of 10 randomly generated instances of the domain. We enforced the one-game-per-person rule by explicitly requesting the players to play only once, and by recording the player's IP and removing multiple instances of the same IP from our database.⁵ The *call for players* was published in various AI related forums and mailing lists all over the world, and eventually we gathered approximately 3000 game instances from over 275 different users from around the world.

The first step was to build a feed-forward back-propagation network on a simple encoding of the domain. 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. The output layer consisted of 3 neurons (as there are 3 piles from which to choose), where the first neuron represented the first pile in the game, that are ordered horizontally from the top-left corner to the bottom-right corner. The other network parameters, such as number of hidden neurons, learning rate (η), and momentum constant, were set by way of trial and error to achieve the best performance.

⁵Our method was not fully secure, and could be manipulated by requesting a new IP address, or by using other computers. However, due to the nature of the game it is reasonable to assume that most players did not do that.

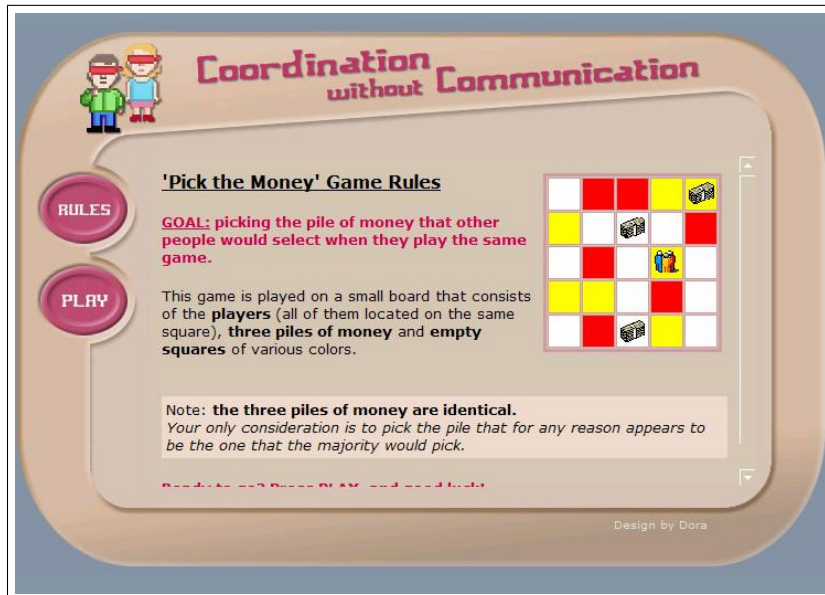


Figure 8: Screenshot from Game Website

The transformation to the focal point space was done in the following way:

1. Centrality — 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.
2. Singularity — 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.
3. Firstness — The Firstness property was calculated as the Manhattan distance between the agent's square and each pile of money.
4. Extremeness — The Extremeness property was intuitively irrelevant in this domain, so we gave it a uniform constant value.

After preprocessing the data according to the above interpretation of the focal point rules, we built two additional agents: the focal point agent (FP) is one that selects the pile with the highest focality value (without employing any learning procedure). The focal point learning agent (FPL) built a new neural network after discretizing the focal point rules to a set of eight possible values {0.125, 0.25, 0.375, 0.5, 0.625, 0.75, 0.875, 1}. Now, the number of input neurons is 36 (4 rules \times 8 discrete values (3 bits) \times

3 possible choices), there are 3 output neurons, and we will train the network on the newly transformed data.⁶

Example: Looking at Figure 7, we can compute the focality value according to the above rules. Before doing so, let us enumerate the piles as follows: pile1 is the upper right pile (on the yellow-colored grid), pile2 is the pile on the 2nd row (counting down from above), and pile3 is the pile on the bottom row. In addition, note that the *ComputeRule* function will be denoted as v for ease of presentation. Now, neither of the piles in that example have any Centrality property (as they are not residing on the same row, column, or diagonal); formally, $v(\text{centrality}, \text{pile1}) = v(\text{centrality}, \text{pile2}) = v(\text{centrality}, \text{pile3}) = 0$. In terms of Firstness, pile1 is 3 grid squares away from the agents, thus $v(\text{firstness}, \text{pile1}) = 0.75$. Pile2 is two grid squares away, thus $v(\text{firstness}, \text{pile2}) = 0.875$, and pile3 is three grid squares away, $v(\text{firstness}, \text{pile3}) = 0.75$. The Singularity property finds piles 2 and 3 on a white grid square; thus, as the board contains 12 white grid squares out of 25 squares overall, their values will be $v(\text{singularity}, \text{pile2}) = v(\text{singularity}, \text{pile3}) = 1 - 0.48 \approx 0.5$. Pile1, residing on a yellow grid square would be valued as $v(\text{singularity}, \text{pile1}) = 1 - 0.28 \approx 0.75$.

Summing their focality values, we have $FV(\text{pile1}) = 0.75$, $FV(\text{pile2}) = 0.6875$, $FV(\text{pile3}) = 0.625$. Here, the agent who is only using a focal point calculation will select pile1, as it has the highest focality value in this game. Moreover, we can now compute the *focality difference* of the game as $0.75 - 0.6875 = 0.0625$. This quite low *focality difference* suggests that this specific instance of the domain is difficult to solve.

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

Each list was composed of five randomly generated candidates. The (pen and paper) experiments were carried out when subjects (a total of 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, we announced a monetary reward for success.⁷ Figure 9 shows a sample question in the domain.

⁶The number of input neurons can be reduced, as the Extremeness property was not used in this interpretation of the rules.

⁷It is a well-known phenomenon that coordination in these games deteriorates without sufficient motivation.

(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 9: Candidate Selection Game Sample

The binary encoding for building the neural network in this domain was a set of 50 input neurons in the input layer that encoded 5 candidates, each encoded with 10 bits (1 bit for gender, and 3 bits for each of the others). The output layer was composed of 5 output neurons, one for each possible candidate.

The focal point transformation had the following intuitive implementation:

1. Centrality — gave a positive constant value to the third candidate in the list (which is located in the center of the selection list).
2. Singularity — 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 increase the singularity value by $1 - 0.2 = 0.8$).
3. Firstness — The Firstness property gave a positive constant value to the first candidate on the list.
4. Extremeness — The Extremeness property gave high values to properties that exhibited 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).

Example: let us now compute the rules according to the game instance presented in Figure 9. We will enumerate the candidates according to their order of appearance from top to bottom. The centrality and firstness properties are intuitive, and result in $v(\text{centrality}, \text{candidate3}) = c^1$ and $v(\text{firstness}, \text{candidate1}) = c^2$, while $c^1, c^2 > 0$. The extremeness property gave prominence to candidate2 as she is the sole oldest, and the tallest, candidate; thus $v(\text{extremeness}, \text{candidate2}) > 0$. The same goes for the singularity property, in which we compute the singularity of the values exhibited in each of the candidate’s properties. For instance, candidate1 is one of the three Males (hence, 0.4), one of the three 25-year-old candidates (again, 0.4), one of the two candidates of 1.75m height (0.6) and one of the three lawyers (0.4). Taking the average, we have $v(\text{singularity}, \text{candidate1}) = 0.45$.

4.3.3 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 seven shapes were presented in a single row and were randomized from the set of 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 10 shows a sample question in the domain.




Figure 10: 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, or triangle. Thus, the question was easier in terms of its simple binary representation, as we needed 14 input neurons and 7 output neurons.

The focal point transformation was implemented as follows:

1. Centrality — Centrality gave additional focality value to the middle choice, and increased the focality value of a shape which had the same shape series on both of its sides. A longer series yielded a higher value.
2. Singularity — 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).
3. Firstness — Firstness gave a small bias to the first shape on the left-hand side of the list.
4. Extremeness — The Extremeness property gave higher focality values to the first and last choices in the set.

In this domain we have an example of a domain transformation in which $j > i$, which means that the transformation actually *increases* the search space. From 14 input neurons in the simple binary representation, we now move to a space of 42 input neurons (6 bits \times 7 choices).

4.4 Results

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 particular human player who played the same game.⁸

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 33.3%–66.6% 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*⁹ 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 11 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.

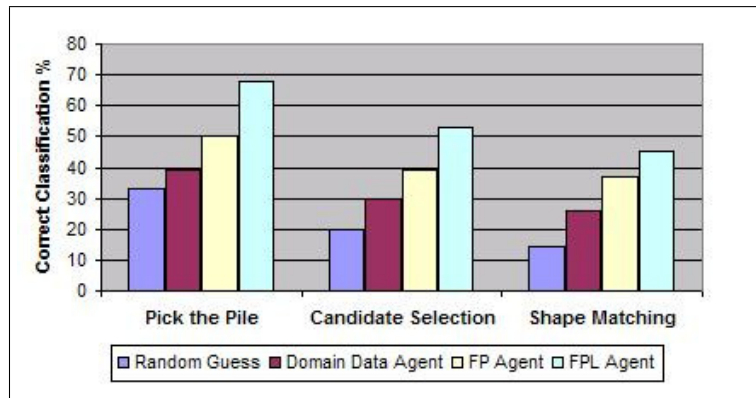


Figure 11: Average Correct Classification Percentage

Examining the results, we see a significant improvement when using the focal point learning approach to train classifiers, rather than the domain data agent ($p < 0.01$ in two-proportion z-tests in all domains). In all three domains, the domain data agent is not able to generalize sufficiently, thus achieving classification rates that are only about 5%–10% higher than a random guess. Using FPL, the classification rate improved by 40%–80% above the classification performance of the domain data agent.¹⁰ The results also show that even the classical FP agent, which does not employ any learning algorithm, performs better than the domain data agent. In an additional analysis that

⁸If there were multiple occurrences of a specific game instance, the choice of the majority of humans was considered the solution.

⁹<http://www.cs.waikato.ac.nz/ml/weka/>

¹⁰Since even humans do not have 100% success with one another in these games, FPL is correspondingly the more impressive.

was done on the FP agent, we saw a tendency in which the FP agent, when facing coordination problems with low *focality difference*, has its performance deteriorate to that of random guesses.

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 28% vs. 22%), which is in turn higher than the marginal increase in performance of the third domain (an increase of 22% vs. 18%). From those results, we hypothesize that the difference in the marginal performance increase is because the first domain was the most complex 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.

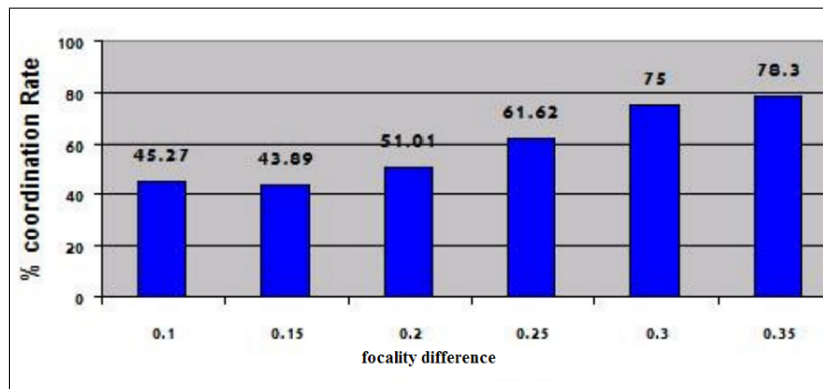


Figure 12: “*Pick The Pile*” Focality Difference Impact

In Figure 12 we checked the impact of the focality difference on the success rate in the “*Pick the Pile*” domain. We computed the focality difference of each of the randomized instances that was played, and divided them into six groups according to their values. We then isolated each group and checked the successful coordination rate on each of the groups. It turns out that in the highest focality instances (> 0.35) the FPL agent managed to achieve a 78.3% successful coordination rate, and in the next group (0.3–0.35) it managed to achieve a 75% successful coordination rate. One can also see that the successful coordination rate increased with the focality difference.

5 Robustness to Environmental Changes

As seen in Section 4, the focal point rules proved to be helpful in handling various tacit coordination domains against an arbitrary human partner. In this section, we examine the robustness of human-agent tacit interaction in changing environments, and analyze the robustness of our agents to changes in the domain environment. Dynamic changes in the environment can occur for many reasons, and having a more robust agent can become a crucial ingredient for succeeding on various missions.

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.

To check agent robustness in the face of environment changes, we took the “Pick the Pile” domain (described in Section 4.3.1), and designed variants; we denote them as the *similar* domain, and the *very similar* domain. To check agent performance, we put the original agents (i.e., the *domain data* and *focal point learning* agents that had been trained on the original Pick the Pile version, and the regular *focal point* agent) in the new environments, and compared their classification performance.




Figure 13: “Pick The Pile” — VSD Example

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 (which we will denote *VSD* for Very Similar Domain), we added a fourth possible value to the set of values of the color attribute (four colors instead of three). In the second variant (which we will denote *SD* for Similar Domain), 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 (see Figure 13 for an example).

Additional experiments were conducted in order to collect human answers to the two new variants of the game (85 first-year computer science and applied mathematics

students took part). 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 14 summarizes performance comparison of the agents in each of the new environment variants.

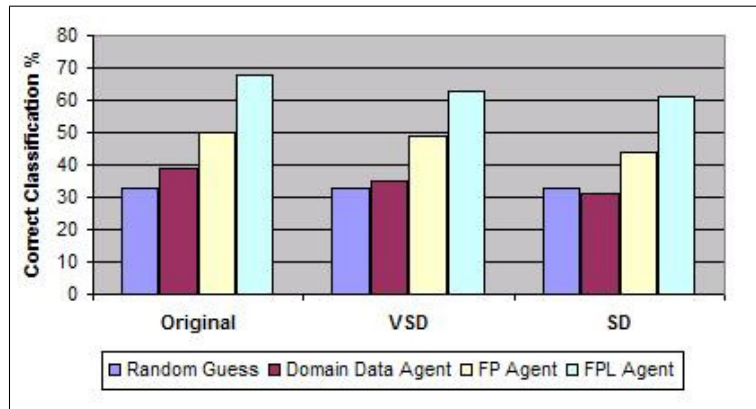


Figure 14: Classification Percentage in Similar Environments

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 reasonably high performance level of around 62% (significantly better than the domain data agent, with $p < 0.01$ in a two-proportion z-test). We can also notice that the classical FP agent copes with the environmental changes better than the domain data agent, with a performance level of around 45%; however, it is still low when compared to the FPL agent’s performance level.

6 Additional Results and Insights

6.1 Generalized Pick the Pile Domain

We conducted an additional set of experiments in order to evaluate the algorithm’s performance on a variation of the Pick the Pile domain that is more generalized and challenging. In this variation, in contrast to the original Pick the Pile domain, the coordination task is not limited to agreeing on one of the three possible piles (i.e., 33% is the expected successful coordination rate according to decision theory), but to agree on any possible grid position. The board was reduced to a 3 by 3 grid, and after removing the pile icons, we were left with 8 possible meeting positions (the grid where the agents are located was not considered a valid meeting place), which resulted in a 12.5% expected successful coordination rate according to standard decision theory. In

addition, instead of the original color attribute, we presented 4 new possible values on each grid square: a picture of a tree, grass, stones, or a lake. This would implicitly cause the players to use more subjective contextual information when needed. Figure 15 presents an instance of the new domain.




Figure 15: Generalized Pick the Pile Example

This experiment was a pen-and-paper experiment with 200 first year students studying computer science, applied mathematics, and engineering. Each subject received a questionnaire with 6 randomly-generated instances. As before, we compared the performance of an agent that was trained on the original domain encoding, and a focal point learning agent that was trained using the same rule implementations as described above for the original domain.

Our results show that this variation is very challenging to both classifiers, as the domain data agent (trained with the original encoding) achieved approximately 16% correct classification, while the focal point learning agent achieved approximately 34% correct classification. While these rates are not as impressive as our results in the original domain, they still provide a significant improvement on the results that the original classifier was able to achieve without using the focal point learning scheme.

6.2 Generalized Shape Matching Domain

In an additional experiment that we conducted on the Shape Matching domain, we generalized the game as follows: instead of having a fixed set of 7 shapes in each instantiation of the domain, we randomized the number of shapes in each row in addition to their type. In the generalized domain, each game instance can include from 3 to 7 shapes (and not always 7 shapes, as in the original version), and each such shape itself was randomized from the set of possible shapes. Figure 16 presents an example of two game instances from a questionnaire, where in the first there are 4 shapes to choose from, while in the second there are 7 shapes to choose from (just as in the original version).

The data was collected in the same way as in the original experiment (each subject played 10 randomized game instances). A comparison was made between a domain data learning agent and a focal point learning agent, so as to evaluate the advantages gained by using the new learning method in a variant on the domain which is more complex than the original one. When the number of randomized shapes was smaller than 7, we used a binary value to represent a “non-existing” shape value for that input neuron. The same focal point rules were used, with the following adjustments to the centrality and extremeness rules, in order to consider the possibility of randomizing the




Figure 16: Generalized Shape Matching Example

number of shapes per row. (1) Centrality — Centrality gave additional focality value to the middle choice, when the number of shapes was odd. With an even number of shapes, no centrality value was given. (2) Extremeness — 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).

Our results show that the domain data agent managed to achieve a correct classification rate of 37%, while the focal point learning agent managed to beat that with a 49% correct classification rate. We also checked the performance of a regular focal point agent (not employing any learning algorithm), and found that it performs as well as the domain data agent, with a 38% correct classification rate.

6.3 Culturally Biased Candidate Selection Domain

One of Schelling’s classic examples can be found in [29], where he showed his subjects a 4 by 4 grid (see Figure 2) and instructed them as follows:

“Put a check mark in one of the sixteen squares. You win if you all succeeded in checking the same square.”

Schelling reported that 24 out of 41 (58%) people selected the upper-left corner (which is of course more than the 6.25% predicted by standard decision theory), and all but 3 of the votes were distributed along the diagonal.

We conducted a similar experiment with a larger set of subjects (81 first-year students) and surprisingly, we had two main focal points: the first was the upper-left corner (with 40 votes), and the second was the upper-right corner (with 32 votes). The upper-right corner did not appear as a possible focal solution to this problem in Schelling’s report, but was found as a strong focal point in our data. Those results seem less surprising when we add the information that the experiment was conducted with students who write from right to left. We hypothesize that this cultural information biased the subjects toward the upper-right corner as an additional focal point, along with the upper-left one, which was found in the original experiment.

In order to further explore the notion of culturally biased tacit coordination games, we took the Candidate Selection game and designed a new, culturally biased, variation of it. In the original version of the game, the five candidates were enumerated using numbers 1 to 5. In the new variant, the five candidates were enumerated using family names that remained constant and were not randomized (all other properties were still randomized). Candidate 1 was named “Brown”, candidate 2 was named “Cohen”,

candidate 3 was named “Jones”, candidate 4 was named “Mitchell”, and candidate 5 was named “White”. The names were selected as to give a cultural bias to candidate 2 as this name is of Jewish origin (and one of the most common family names in Israel, where the experiment took place), while all other candidates have common American family names that we assumed have no special meaning to most of our student subjects.

<u>Name</u>	<u>Gender</u>	<u>Age</u>	<u>Height</u>	<u>Profession</u>
Brown	female	31	1.78	businessman
Cohen	male	29	1.81	doctor
Jones	female	35	1.78	doctor
Mitchel	male	31	1.81	doctor
White	female	35	1.71	businessman

Figure 17: Culturally Biased Example

We used first and second year students to play randomized instances of the game. The results, without even considering the candidate properties, are not surprising: while in the original version (enumerated by numbers), 16% of the students selected candidate 2, in the biased version this number rose to 27%.

Given a tacit coordination domain and such knowledge of a cultural bias (or biases of other forms as we will see below), we can augment the FPL with a cultural bias-based rule that will increase the focality value of the choice (candidate 2 in our example) by some constant amount. This increase will help the learning algorithm to break ties and differentiate towards the biased answer when the focality difference is small. In our experiments, by including such a bias rule in the form of a constant increase of 0.15 to singularity value of candidate 2, we managed to increase the agent’s correct classification rate by 9%.

6.4 Additional Insights

In a series of experiments that were partly conducted to reproduce and verify the results reported in several of Schelling’s classic experiments, and partly as a pre-experimental stage for the domains presented above, we gained additional results that can help in deepening our understanding of the nature of focal points.

6.4.1 Educational Bias

In a pre-experimental stage of the “pick the pile” domain, we conducted pen-and-paper experiments in a small class composed of 14 computer science graduate students. The

pre-experimental stage was meant to check whether the written instructions were comprehensible, as we wanted to avoid answering questions out loud, due to the possibility that such answers might bias the subjects.

The instructions stated that the subjects were to be paired with an unknown partner from that class, and that their answers would be compared. Here again, a monetary prize was given to the couple with the highest number of successful coordinations.

When recording their solutions, we found that there was a tendency to select piles that are on the same column or row, even if they are further away in terms of their Manhattan distance. Those results surprised us, as we hypothesized that the subjects, being computer science graduates, would naturally follow the Manhattan distance measure, as the instructions clearly state that “diagonal moves are illegal”.

After the experiment ended, we interviewed several subjects in order to understand their selection. To our surprise, it appears that those students were working in the field of Robotics, and it was obvious to them that the *refacing* operation is an expensive operation in the robotics domain, and should be avoided. They also added that they used that knowledge when solving the coordination games, as they knew that their coordination partners were to be selected from that group alone.

Those results show again that the more knowledge one has of one’s coordination partner, the more it allows one to focus choices with respect to that information which is common to both partners.

6.4.2 Firstness in “Heads or Tails?” Domain

In several of our paper and pen experiments we decided to explore why, according to Schelling, around 85% of the answers to the simple “Heads or Tails?” question turns out to be “Heads”.¹¹ Why is “Heads” an obvious focal point for a tacit coordination task based on that question? We hypothesized, according to our suggested focal point rules, that the Firstness property has a crucial role in focalizing “Heads”, as it appears first in the presentation of the question.

In order to check our hypothesis, we generated two types of questions: one being the original formulation of the problem “Heads or Tails?”, and the other the reverse presentation, “Tails or Heads?”. In several of our previous experiments, we planted a single question of one of the above forms into the questionnaire. Thus, several of our subjects played the “Heads or Tails?” version, and others played with the “Tails or Heads?” version.

Table 1: “Heads or Tails?” Results

	“Heads” selected	“Tails” selected
“Heads or Tails?”	23	7
“Tails or Heads?”	13	16

Overall, we collected 59 answers, and the results appear in Table 1. The results show that on the “Heads or Tails?” question, we got results similar to those of

¹¹In his experiments, 36 selected Heads and only 6 Tails.

Schelling, as only 7 people selected “*Tails*” out of 30 subjects (i.e., 23% in our experiment compared with Schelling’s 14%). However, in the second variation, the “*Tails or Heads?*” presentation, 55% of the subjects selected “*Tails*”, a greatly increased percentage. It might be hypothesized that although presenting “*Tails*” first made it a far more likely choice, it may not have reached the popularity of “*Heads*”, since “*Heads or Tails?*”, with “*Heads*” first, is still the more standard presentation of the problem in English (and thus the primacy of Heads is still somewhat maintained in many players’ minds).

In any case, it is easy to see that the presentation of the question completely changed the focality of “*Heads*” from a high focality answer in the original presentation, to a lower focality answer in the new presentation. Two lessons can be learned from those results: first, obviously the presentation of the coordination task impacts its focality difference. Second, as the only difference in the new version is the order of the choices, firstness is an important property in that coordination task.

7 Conclusions and Future Work

In this article we have been interested in enhancing agent-human coordination possibilities in situations where communication is not possible. When an agent (or robot) is coordinating with a human partner, having a cognitive mechanism that mimics the reasoning process of that human being will make the agent more robust to changes in the environment, and more adaptable to different domains.

We presented the *Focal Point Learning* (FPL) approach to building automated agents that play tacit coordination games with general human partners in dynamic 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 pre-processed 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. We then created two variants of one of the domains, collected human data for the variants, 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 an altered environment.

As can be seen from our experimental work, even in our simplified domains the suggested FPL technique relies on having a large amount of data. Moreover, this data should be taken from a large enough sample space to remove various biases. It might be that the costs of applying FPL techniques for cases where there is a need to coordinate without communication is not practical for all domains, as acquiring experimental data from humans is not a simple procedure. However, in domains where human lives are at stake, such as search-and-rescue tasks, caring for the elderly, or human-robot

teams, this technique may be very beneficial. In addition, the wide array of electronic data that is available due to the development of the internet (e.g., web searches and electronic transactions), can be used with various data mining techniques to improve user-interface experiences or automated agent development processes, when communication will not be available at runtime.

Future work on Focal Point Learning can be taken in several different directions. First and foremost, more experimental work with human beings is needed to deepen our understanding of the nature of focal points and their properties. Through that effort, we can gain a deeper understanding of the reasoning process that underlies tacit coordination games. Additional experimental work on the psychological component of focal points is most definitely needed.

Another direction is to take a game theoretic approach and look for a robust and formal theory in game theoretic terms which, in contrast to existing focal point theories, will have predictive abilities and not only the capability of retrospective analysis.

From the Focal Point Learning point of view, one obvious next step is to apply the technique to a real-world domain, for instance, a military application where manned and unmanned aircraft or land vehicles are cooperating to carry out some well-defined task. In military applications, due to uncertainty and the dynamic nature of the environment, having a cognitive mechanism could prove to be highly important.

To sum up, when building agents to coordinate with unfamiliar human partners, without communication, machine learning classifiers have difficulty generalizing data to predict human choices. Focal Point Learning can improve performance, and robustness in the face of environmental change.

References

- [1] R. C. Arkin. Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3):351–364, 1992.
- [2] M. Bacharach. Variable universe games. In *Frontiers of Game Theory*, pages 255–275. MIT Press, Cambridge, MA, 1993.
- [3] Michael Bacharach and Michele Bernasconi. The variable frame theory of focal points: An experimental study. *Games and Economic Behavior*, 19(1):1–45, April 1997.
- [4] Herman Chernoff and Lincoln Moses. *Elementary decision theory*. Dover Publications, Inc., New York, NY, USA, 1986.
- [5] Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. Selection criteria in coordination games: Some experimental results. *American Economic Review*, 80(1):218–33, March 1990.
- [6] Vincent P Crawford and Hans Haller. Learning how to cooperate: Optimal play in repeated coordination games. *Econometrica*, 58(3):571–95, May 1990.

- [7] Maier Fenster, Sarit Kraus, and Jeffrey S. Rosenschein. Coordination without communication: experimental validation of focal point techniques. In *Readings in agents*, pages 380–386. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1998.
- [8] D. Gauthier. Coordination. *Dialogue*, 14:195–221, 1975.
- [9] V. Gervasi and G. Prencipe. Coordination without communication: The case of the flocking problem. *Discrete Applied Mathematics*, 144(3):324–344, December 2004.
- [10] Michael A. Goodrich and Alan C. Schultz. Human-robot interaction: A survey. *Foundations and Trends in Human-Computer Interaction*, 1(3):203–275, 2007.
- [11] Barbara J. Grosz. Beyond mice and menus. *Proceedings of the American Philosophical Society*, 149(4):529–543, 2004.
- [12] John C. Harsanyi and Reinhard Selton. *A General Theory of Equilibrium Selection in Games*. MIT press, 1988.
- [13] Robert Hirsh, Jeffrey Graham, Kimberly Tyree, Maarten Sierhuis, and William J. Clancey. Intelligence for human-assistant planetary surface robots. In A. M. Howard and E. W. Tunstel, editors, *Intelligence for Space Robotics*, pages 261–279. TSI Press, 2006.
- [14] John B. Van Huyck, Raymond C. Battalio, and Richard O. Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1):234–48, March 1990.
- [15] John B. Van Huyck, Raymond C. Battalio, and Richard O. Beil. Tactic coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1):234–248, 1990.
- [16] M. C. W. Janssen. Focal points. In P. Newman, editor, *The New Palgrave of Economics and the Law*, pages 150–155. MacMillan, 1998.
- [17] Maarten Janssen. Rationalizing focal points. *Theory and Decision*, 50(2):119–148, March 2001.
- [18] Randall W. Hill Jr., Youngjun Kim, and Jonathan Gratch. Anticipating where to look: Predicting the movements of mobile agents in complex terrain. In *Proceedings of the First International Joint Conference on Autonomous agents and Multiagent Systems (AAMAS’02)*, pages 821–827, New York, NY, USA, 2002. ACM.
- [19] Michihiro Kandori, George J Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56, January 1993.
- [20] Sarit Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, 2001.

- [21] Sarit Kraus, Jeffrey S. Rosenschein, and Maier Fenster. Exploiting focal points among alternative solutions: Two approaches. *Annals of Mathematics and Artificial Intelligence*, 28(1–4):187–258, October 2000.
- [22] D. K. Lewis. *Convention: a Philosophical Study*. Harvard University Press Cambridge, 1969.
- [23] C. Martin, D. Schreckenghost, P. Bonasso, D. Kortenkamp, T. Milam, and C. Thronesbery. An environment for distributed collaboration among humans and software agents. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 1062–1063, New York, NY, USA, 2003. ACM.
- [24] Judith Mehta. Telling tales: Actors’ accounts of their behavior in coordination games. In *Focal Points: Coordination, Complexity and Communication in Strategic Contexts*, Lund, Sweden, 1997.
- [25] Judith Mehta, Chris Starmer, and Robert Sugden. The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, 84(3):658–73, June 1994.
- [26] John Nash. Two-person cooperative games. *Econometrica*, 21(1):128–140, 1953.
- [27] L. Pederson, D. Kortencamp, D. Wettergreen, and I. Nourbakhsh. A survey of space robotics. In *7th International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, 2003.
- [28] J. Ross Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, 1993.
- [29] T. Schelling. *The Strategy of Conflict*. Oxford University Press, New York, 1963.
- [30] 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.
- [31] Nathan Schurr, Janusz Marecki, John P. Lewis, Milind Tambe, and Paul Scerri. The defacto system: Training tool for incident commanders. In *AAAI-05*, pages 1555–1562, Pittsburgh, 2005.
- [32] Maarten Sierhuis, Jeffrey M. Bradshaw, Ro Acquisti, Ron Van Hoof, and Renia Jeffers. Human-agent teamwork and adjustable autonomy in practice. In *7th International Symposium on Artificial Intelligence (i-SAIRAS)*, 2003.
- [33] Robert Sugden. A theory of focal points. *Economic Journal*, 105(430):533–50, May 1995.
- [34] 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.

- [35] Michael J. Wooldridge. *An Introduction to MultiAgent Systems (Second Edition)*. Wiley, 2009.
- [36] H. P. Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.
- [37] H. P. Young. The economics of conventions. *Journal of Economic Perspectives*, 10(2):105–122, 1996.