Focal Points and Hierarchical representations

A cognitive perspective on focal points in a spatial environment

Final project in Amirim program
submitted to Prof. Jeff Rosenschein

By: Reshef Meir
I.D.: 040097503
Date: 23.2.2007

Introduction

People are often able to coordinate their actions in day-to-day tasks without any communication before or during the task. Certain ways of action just seem somehow more "obvious", or salient, to all the involved participants. Those strategies are called Focal points. In this work I will try to explain how focal points emerge in situations involving way finding in a geographical environment. That is, how do agents who are trying to meet each other go to the same place, in face of the many available options?

In the first chapter, I will review the recent progress that was made in defining focal points, explaining where they come from, and how people may use them to coordinate. Specifically I will show why Focal points exceed the field of rational game theory. I will review some of the work done so far in the field, which try to explain focal points both in game-theoretic and in decision theory context. As these attempts improve our understanding of the choice procedure, they still come short of explaining the actual reason that some solution is a salient one. My assumption is that in order to understand these reasons, we must explore the domain specific cognitive representation a person use when facing a coordination problem

In the second chapter I briefly review cognitive and psychological literature, which describes how spatial environment is represented in the human mind. I will concentrate in evidence for the existence of a graph-like representation, where features of the environment are remembered relatively to other features which are "more important" in some ways.

In the third chapter I will put forward my claim that focal points in spatial settings can be explained in light of the cognitive findings presented in chapter two. Specifically, that features playing important role in navigation, may play a similar role in coordination, as focal points.

In the forth chapter, a simulation is described where automated agents, which implement the hierarchical representation scheme, are trying to coordinate with other automated and human agents. I present a new testbed designed to simulate a spatial environment where agents can be trained and tested. Formal metrics to measure focality and performance of coordinating agents are defined as well.

Finally, I will conclude with some future directions, and present some ways in which the theory can be generalized to other problems and fields.

Background

In this chapter, I will present the terms Focal Points and Pure coordination games, and give examples. I will show how such games are formally defined in game-theoretic framework, and why such a framework is not enough when trying to solve them.
Finally, I will review some work in the field regarding both explaining the nature of focal points, and how such points may be found.

What are focal points?

The term "focal point", also referred to as "Schelling point", was introduced by the Nobel Prize winning American economist Thomas Schelling in his book *The Strategy of Conflict* (1960). It describes a place, a way of action, or choice, which people tend to use more often, when they are trying to coordinate with other people with whom they have no communication.
The best explanation can perhaps be given by an example:
Suppose you discussed to meet with your friend at 20:00 in the center of Jerusalem, but the battery of your cell phone is out, so you and your friend cannont talk. Where will you go?
9 of 10 people (who live in Jerusalem) answered they would go to the "Mashbir"[1].
Certainly one does not get any immediate profit from going there rather than to any other location, but people still converge toward this specific point.
Other examples:

1. People asked to coordinate on a positive number often chose "1" (also "7","100" or "a million")[2].
2. In the New-York version of the meeting place question, most people said they would go to Grand Central Station.
3. When also required to coordinate on meeting time, almost all chose 12:00am.
4. Most people selected the top-left square in a 4X4 array of identical squares.
5. Most people choose "Heads", when trying to coordinate on the outcome of a coin toss.

These qualitative findings and others were later measured and reproduced by [Mehta et al. '94b].

[1] An informal survey among students in HUJI
[2] These games and others are presented with their results in (Schelling '60), pp. 56-57

Motivation

The situations described to demonstrate the presence of focal points are often not very useful in real life. Still, if focal points have such a dramatic effect on decision making in such situations, they may also play important role in real life situations where other factors are involved. In Game theory, focal points may bring into games considerations that are not always rational, as I will show in the next section. Therefore, their understanding may be important for game-theoretic analysis of real life situations also where communication is possible and even free. Schelling concentrates in the role focal points may have in negotiations and bargaining.

Another field that may benefit from focal points is Multi Agent Systems. In such systems agents are often required to coordinate with each other, and communication may be expensive or impossible. Understanding the way people find focal points, can reduce the amount of communication between agents, especially when some of them are human.

Focal points in Game theory

In this section I give a brief, formal definition of a game. I will use some of the notations in the next sections.

A game in its strategic form is defined[3] by <P,S,M>, Where:

- $P = \{1,2,..,n\}$ is the set of players.
- $S = \{ S_1, S_2,.., S_n\}$. S_i is the set of strategies allowed by player i.
- M is a $|S_1|x|S_2|x..x|S_n|$ matrix called the Payoff matrix. For each i, denote $k_i = |S_i|$.

  Each cell of M contains n real values, one for each player.

Each player i simultaneously choose a strategy $j_i \in \{1,2,..,k_i\}$

Each player get the i'th value in the cell $M_{j_1,j_2,...,j_n}$

A mixed strategy of player i is a distribution over Si. $(p_1^i, p_2^i,.., p_{k_i}^i)$

The payoff of the i'th player in a game with mixed strategies is:

$$\sum_{j_1=1}^{k_1}\sum_{j_2=1}^{k_2}\cdots\sum_{j_n=1}^{k_n}\prod_{i=1}^{n} p_{j_i}^{i} M_{j_1,j_2,...,j_n}$$

I will only be interested in games with final sets of players and strategies.

[3] Based on http://en.wikipedia.org/wiki/Normal_form_game

It is possible to formalize the games Schelling described in the following way:
Each player has k possible strategies, and all players get the same value, which is “1” iff all of them chose the same strategy, and “0” otherwise (e.g. fig.1). These games are known as pure **coordination games** [Mehta et al. `94b][4].

1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

Figure 1. the payoff matrix of a pure coordination game.

Fig. 1 can also serve as a description of a specific game, in which players should choose matching row and column.

Elimination of game-theoretic explanations to focal points

we saw that formalizing the game is rather easy. As in any game, an important question that arises is “how can a player get a good score in this game?”.
I will show that if we ignore all context data which is not part of the strategic form of the game, the expected payoff of a player cannot exceed 1/k.

Computing the value of the game

Can a player **ensure** herself a good score? Unfortunately no. With pure strategies, the value of this game is 0, since player 1 has no strategy which guaranties something higher than zero, and same goes for all other players. (is there a place in the city which guaranties a meeting?)
If mixed strategies are allowed, the value is 1/k, where k is the number of strategies. This is possible by using the strategy (1/k,1/k,..1/k). Actually, this strategy guaranties that all the players will get exactly 1/k.
The proof is rather trivial (for n=2):

Let (p1,..,pk) be the strategy of player 2.

Player 1 use (1/k,..1/k).

[4] [Sudgen '95] distinct between Pure coordination games and matching games. Such a distinction is not important for my purposes.

The payoff is $\sum_{j=1}^{k}(1/k * p_j) = \frac{1}{k}\sum_{j=1}^{k} p_j = \frac{1}{k}\cdot 1 = 1/k$

Nash equilibrium

In many games special points can be identified which ensure something better then the value, provided that the other player also identifies this point and use the correct strategy. These points are known as Nash equilibria, or Nash points[5].

Player 1 \ 2	A	B	C
A	1 \ 3	6 \ 6	1 \ 1
B	1 \ 3	1 \ 1	10 \ 1
C	20 \ 1	2 \ 1	1 \ 10

Figure 2. A game with a pure Nash equilibrium

In the game presented in Fig. 2, player 1 can guarantee only 2 (using C) and player 2 can guarantee only 3 (using A). Still, there is one pure Nash point in this game (highlighted). By recognizing this fact and take the relevant strategies (A\B) both players can get 6.

Are there pure Nash Equilibria in pure coordination games?

It is easy to see that the main diagonal of the game matrix (Fig. 1) is consisted of Nash points. But the problem is that there is **no unique** Nash point. Rather, the number of pure Nash equilibria is as the number of possible strategies. If players are not allowed to communicate, they have no way to coordinate by merely analyzing the Nash points of the game.

Repeated games

In a repeated game, agents may learn a way to coordinate on a specific equilibrium (Crawford & Haller '90), but this does not solve the problem of a single game. Since human agents seem to be quite successful even in games that do not repeat, any solution to the iterated problem will not explain how focal points are recognized by people.

[5] The Nash theorem guarantees that there is *always* an equilibrium point when mixed strategies are allowed, but for simplicity I will only refer to pure Nash points in this section. The fact that there may be more Nash points only stress the problem I present.

Correlated equilibria

Given a game, it can be extended by an external random event, which the players observe fully or partially. The players can then base their strategy on what they observed.

Informally, A correlated equilibrium is[6]:

- A distribution on the possible outcomes of the random event.
- A strategy for each player, for each possible outcome.

Such that for each outcome no player can do any better than to follow her predefined strategy, assuming all other players will do the same. In situations where the players can agree on such an event in advance, it can indeed solve the problem of coordination.

I will use the "junction game"[7] as an example:

1 \ 2	*Go*	*Halt*
Halt	1	0
Go	-1	1

Figure 3. the junction game. The highlighted cells are both Nash points.

This is a pure coordination game since there is no conflict between the players. It contains 2 pure Nash points, as can be seen in Fig. 3. however , this fact alone does not help the players to choose a strategy.

Now, we extend this game by adding a traffic light (or a random traffic light), with two outcomes {Blue, Yellow}. (we will note the outcome as the color player 1 sees).

A possible correlated equilibrium is the following:

- The strategy of player 1 is Go iff Blue.
- The strategy of player 2 is Halt iff (player one sees) Blue.

The traffic light actually chooses at random between the two lighted entries of the matrix, then each player can do no better than to follow what her strategy implies, assuming the other player will act in the same way.

It is important to emphasize that the correlated equilibrium is possible due to strategies that are **not part of the original game**. Thus, the players can avoid collision (or time wasting), only in the new game that includes the traffic light as well.

[6] A formal definition and analysis of correlated equilibrium games can be found in (Aumann '74).

[7] The traffic light example is based on a similar example by Prof. Aumann in a lecture held in HUJI as part of the Nobel Prize series. I purposely changed the colors to emphasize that without prior agreement, it is not clear which color means "go".

In a situation where such an external signal is not possible, the strategies for correlated equilibrium cannot be used. We face the same problem even if such signals are available, but the agents have no way to communicate before the game and **agree on which signal they will use**, and how. In the example of the traffic light, the agents agreed on the meaning of the signal "Blue". Another correlated equilibrium would be achieved by assigning the same meaning to "Yellow" instead. In order to use any of these two options, the agents have to agree on the meaning of the signal, or the problem will be regressed to choose among correlated equilibria instead of Nash.

Decision theory and salience

In order to explain the ability of people to find focal points, researchers focused on the data which is not part of its strategic form. The main work related **labels,** or **attributes** of strategies.

	A	***B***	***C***	***D***
A	1	0	0	0
B	0	1	0	0
C	0	0	1	0
D	0	0	0	1

Figure 4. a pure coordination game with labels.

This course of explanation asserts that each label has some inherent "salience". This salience has no objective measure, and different strategies may seem salient to different players. Nevertheless, the nature of the decision process is responsible to the fact that in some domains people do attribute salience to the same strategies, especially when these people come from the same cultural background.

This assertion gives rise to at least three explanations of the focal point phenomenon:

1) People always pick the solution that they perceive as most salient. Since large parts of the population attributes salience in a similar way, strategies that are salient to many people will become focal points. This explanation is known as **primary salience**.
2) When asked to coordinate, People always pick the solution they believe has primary salience to most of the people. This is known as **secondary salience**,

and can be extended to third-order salience (choose the solution that has secondary salience to most people), and so on.

3) When asked to coordinate, People pick the solution they believe salient to other people that are *trying to coordinate*. This explanation is known as **Schelling salience**.

Assuming the strategy "A" is the most salient, then the fact that the entry A\A is the focal point in fig. 4 follows from all three explanations.

Schelling salience and the Principle of Coordination

One line of explanation (Sudgen '95) is the logical "principle of coordination". We assume that there exists some irrational process of labeling, on which the player has no direct control and that this process is a common knowledge. Then, players can rationally exploit asymmetries in the given labeling to identify a unique Pareto-efficient Nash equilibrium[8], if there is one. Schelling salience can be explained by such a two-stage process. In Fig. 5, the labeling process will necessarily assign one label (say "dotted") to three of the disks, and another (say "wavy") to the other two. if the player sees the disks in random order, the only allowed strategies are: a) choose a wavy disk ; b) choose a dotted disk ; c) choose any disk. Among those, (a) is a unique Pareto-efficient equilibrium, since it maximizes the probability of coordination.

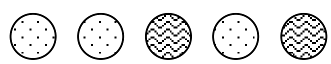

Figure 5. Choose a Disk game. The players see the disks in random order, and must choose the same disk.

(Janssen '95, Bacharach '93) extend the principle of coordination to situations where there are multiple labels and there is also a chance that some players will not perceive some of them.

(Janssen '97) notes, that in a 2X2 game (e.g. Heads or Tails) if there is a salient solution, the other solution is also salient in a way. Thus the principle of coordination can never be used. Since people are quite successful in such games as well, it cannot be the only explanation to focal points. The same criticism may be used against all games in which the k strategies are completely distinct, and hence the labeling cannot yield any asymmetry.

[8] A Pareto-efficient Nash point is a Nash point that is strictly better for at least one agent than any other Nash point, and is not worse for all the others. By definition there can be only one such point, if any.

Heuristic attributes of Focal Points

A different, independent, line of explanation explores the reasons making some options or features more salient than others. Such insights will directly explain primary/secondary salience, but are also needed to understand the labeling process underlying Schelling salience. These reasons are naturally more domain-specific and rely on the actual features of the problem (Mehta '94b).

(Kraus et al. 2000) suggested four heuristic properties used by players that are trying to meet each other in geographical settings, those being Uniqueness, Uniqueness compliment, Centrality, and Extremeness. Each possible solution gets a grade in each property, and these grades are merged in some way to a single grade reflecting the salience, or "focality" of a solution. (Zuckerman et al.'07) used similar heuristics (Uniqueness, Centrality, and Firstness) to improve machine learning of focal points in various domains.

All the explanations I presented here take into account that there is some irrational process, which takes place before the game itself and determine somehow the labels or attributes of the different strategies. Some of them explicitly connect this process to the representation of the game in the minds of the players, and describe some of its characterizations. The question of exploring the nature of this representation remained so far out of scope.

Hierarchical representation of space

In order to understand how people find their way in certain environment, the cognitive representation of such environments should be studied.

In this section I will present findings from cognitive and developmental psychology, showing that spatial concepts such as land marks, paths and regions are represented in the human mind in a schematic **hierarchical way**. I will also show some of the factors affecting the creation of such hierarchy, and ways people may use such a hierarchical representation to find their way.

Empirical findings

A common analogy to the representation of space is a Map. In a map all the items have some fixed coordinates. By visualizing a map one can tell the distance between items, angles, sizes of regions and further describe in a more or less accurate way spatial properties of the depicted environment. In contrast, Experiments where people were asked to answer spatial questions regarding both familiar and newly learned environments show different findings. These findings are not related to the way people perceive their environment (e.g. look around or at a map), but to the way they represent, or remember, what they saw.

- Inconsistency: Distances between landmarks in the same region seem smaller than between land marks in different regions.
- Asymmetry: distances towards some land marks seem shorter than the distances from them (Tversky '93).
- Perspective: land marks which are distant from the person seem closer to each other.
- Partial description: sometimes only the spatial relation between land marks is remembered, with no notion of distance. Oblique land marks may get "rectified" or forgotten altogether (Olson & Bialystok '83, Tversky '93). Distinctive features of landmarks are remembered more often than other features (Hart & Berzok'82, Pick & Rieser '82). Landmarks that have some meaning to the subject (i.e. pose some benefit or danger) are much more likely to be remembered (Hart & Berzok '82).
- Contradictions: different descriptions of the spatial relation between two non-adjacent land marks are often contradictive (Tversky '93).
- Sexual differences: as women tend to remember better the features of landmarks (e.g. color), men remember better their spatial properties (Newcomb '82).

- Age differences: young kinds remember only basic spatial relations, and often cannot understand the spatial relation between distant areas.
- Environment structure: environments of a hierarchical nature support and facilitates way finding (Freksa ‘99). Way finding is also supported by irregularities in the environment.

Cognitive models

An acceptable claim is that the cognitive representation of an environment should be suitable for real way finding and orientation tasks, rather than for an accurate description of space. The presented models show how useful data on the environment may be recovered in a simple way. This simplicity comes on the account of accuracy.

(Tversky ’93) explains the findings in the previous section by ruling out the Cognitive map hypothesis. Instead she suggests two mechanisms of representation, which exist in parallel and can be used in simple environments.

- In **route** description land marks are represented relatively to the viewer. Such description is gained mainly during the actual movement a person performs in some environment.
- In **survey** description, landmarks are represented relatively to each other, or to regions, and the relations used are schematic rather than accurate.

(Hart & Berzok ’82) make a similar distinction between representations, and describe the connection between them. A person uses some survey description of the entire environment. Route descriptions of specific paths or sub-regions are used as more detailed "shortcuts" between land marks. Usually such that are visited more often than others. In addition, it is suggested that these description come in different accuracy level, ranging from simple sequential order to exact path description in route, and from basic spatial relations to more detailed ones in survey.

(Olson & Bialystok ‘83) offer a theory of spatial representation. Among possible representations, they describe the relational one as both flexible enough to incorporate knowledge from different sources (unlike maps), and rigid enough to allow a proper distinction between landmarks and situations (unlike verbal description). They describe complex spatial relations as constructed from simpler ones, thus explaining some of the findings about over simplified representation.

Factors affecting the hierarchy

I tried to mine from the presented findings and models some key features of the way people store and use spatial information. I will later rely on these features in the implementation of navigating agents.

As I emphasized, the most important feature is that land marks are represented relatively to each other, and in a hierarchical way.

Many factors affect whether a certain landmark or region will be represented, which spatial and other properties of it will be represented, and what will be its location in the hierarchy.

A "high", or prominent landmark is usually one that is very distinctive from others, and its distinctive features will be remembered. Other characteristics of prominent landmarks are being more meaningful to the subject, more familiar or more reliable.

It seems to me that two basic factors imply the rest: the potential benefit of a land mark (by itself or by its role in finding other landmarks), and how much resources are required to represent it.

Navigation and way finding

When a person is required to navigate her way to a certain landmark, she may use a specific route description of the way (if she has one). This is probably the mechanism used when people travel a path that is very familiar to them.

When one is required to find some target, to which she has no such ready-made route description, she may use the hierarchical survey description. she may either move physically along some path in the hierarchical representation, or reason about the way to her target using the hierarchy, but without actually moving through it.

Anyway she will use more prominent landmarks and regions, in order to find the less prominent ones. The inaccuracy in the representation does not have a grave effect on navigation, due to the redundancy in the description of routs and landmarks, and the fact that features of the environment can be used to find others.

To summarize, people hold (at least) two representations of their environment. One of these representations is a hierarchical graph-like structure, where most landmarks are represented relatively to others that are more prominent. I will later refer to prominent landmarks as **anchors**.

The research question

The example in Fig. 1 shows that in order to find the focal points of a game one cannot rely strictly on mathematics. Nevertheless, the success rate of people playing this game implies that there is something else. Something, which allows people to solve the problem presented in front of them. If we assume that no paranormal activity takes place, and that people do not cheat somehow, then the answer must be in the **representation** of the problem. Indeed, the different Heuristics used to analyze Focal points take advantage at either the mathematical form of the game, or its representation. For example, when Fig. 1 is used as the representation of the game, people often choose the **first row.**

Focal points in a geographical setting

The general field of Focal Points is very large, and the types of problems may differ from each other in many ways. Since the representation of different problems can also be very different, finding a general method for analyzing Focal points based on representation does not seem feasible to me.

I have noticed that in some domains it is much easier to find focal points than in others. Most of the examples in Schelling original work and in other places (Mehta et al. '94a,b) use either alternatives including integers, or some spatial setting, usually geometry or geography.

In this work I looked at the sub-domain of spatial navigation. The different strategies are locations in the environment where a player can go, and the payoff is relative to the chance that a meeting will occur. I tried to analyze the sort of Focal Points we should expect, considering the knowledge we have on spatial representation in the human mind.

Hypothesis

The observed effect of "focal points" is caused by the spatial representation mechanism mentioned in the previous section, and can be explained in its frame:

- In the lack of any other information or reason, people will be strongly biased towards their "anchors".
- People tend to have similar hierarchies and therefore similar "anchors", those land marks end up as the Focal points of the task.

I do not claim that the mentioned mechanisms are the only ones used for navigation. In complex environments people can and do use other mechanisms as well (this entails though, that in environments where other mechanisms are dominant, I will expect to witness strong focal points)

Method

It is notable that my main claim is cognitive in nature, rather than computer-theoretic. Such claims cannot be "proved" directly. Nevertheless, they can be supported in many ways: cognitive experiments, neurological findings, philosophical arguments, computer simulation, and more. Aside from some intuitive arguments, the project intent was to support the mentioned claims by using a simulation.

First, a testbed was created, simulating an environment in which agents can perceive their surroundings, move and perform tasks.

The tasks are simple in nature: the agent is presented with a set of spatial locations, and is required to pick one of them and **go there**. Each agent is encouraged to pick the location which all other agents are likely to pick (i.e. to find the focal point).

Agents can be controlled either by a human (through GUI), or by a computer program.

The testbed and simulation process are covered in detail later on.

Applications

Since my claim was a cognitive one, the main gain of this work was to get a better understanding of the focal point phenomenon. Specifically, eliminating the game-theoretic approach to focal points, and instead treating focal points from a cognitive point of view.

These insights on focal points might be used in real world applications where agents share some mutual goal, but communication is impossible or expensive.

There are two primary settings:

- Coordination of automated agents and human agents.
- Coordination of automated agents with each other.

In the first setting we can use our newly gained insight on focal points, along with existing psychological/cognitive models, to develop effective algorithms and heuristics for coordination of human and robots. Such an attempt was made by (Zuckerman et al. '07). The idea of developing computational cognitive models in order to improve coordination of men and machine is overviewed by (Schultz & Trafton ‘05).

The second setting is a special case of a more general problem known as “Decentralized decision making with uncertainty”, with strict limitations on the allowed communication. Traditional ways to solve such problems use AI techniques such as DEC-POMDP[9]. Such techniques provide ways for solving general coordination problems (meeting on a grid being just a special case), but require the use of very heavy time and space resources.

There are two other downsides, which are common to DEC-POMDP and other decentralized algorithms for coordination such as in (Schermerhorn & Scheutz '06): the first is that agents are to be **designed together in advance**. The second is that features of the environment are usually ignored or treated as obstacles, rather than exploited.

By using features of the environment, agents using focal points can increase the chance for coordination following only some simple guidelines. This idea was successfully implemented by (Kraus et al. 2000).

Hierarchical models have been explored and used for some time now for representation of space, exploration and navigation tasks by robots (Egenhofer et al. '96, Kuipers & Byun '88, Kuipers 2000). Many of these models had been inspired by psychological findings such as those appearing in chapter two. The focal points framework allows us to benefit from this research on single agents, in order to design efficient multi-agent systems that will require less communication, both among themselves and with their human operators.

General arguments that support the hypothesis

As I have shown earlier, hierarchical structures have significant roles in the way a person represents her environment. Such structures are efficient and flexible, and instead of accuracy they allow a quick action that does not require much deliberation.

[9]Decentralized Partially Observed Markov Decision Process (or DEC-POMDP) is a probabilistic model of the world, which consists of states and transition probabilities. A utility is assigned for each state or transition. Each agent has a different belief on the state (or distribution over states) of the world, and tries to maximize utility. (Goldman & Zilberstein '04) show how agents can use a strategy that will reach high expected utility. A pure coordination game can be viewed as a special case, where utility is 1 in all the states indicating agents are in the same place.

When a person has more information on a certain environment, she might use it to find her way in a quicker way. When such information is missing, unreliable, or cannot be used, she must “stick to what she knows”. Usually this means using the higher, more prominent landmarks in the hierarchy, to visualize some path to the target, even if this path is not the shortest one. People often start moving towards their anchors even if they did not visualize yet the whole path to the target.

Trying to coordinate without communication is a private case of such situations where crucial information is missing. Specifically, the person does not know where to go. She may solve this problem by going to the most prominent places that she knows best and will be the first to think of. Moreover, since the context has a crucial effect on representations, a different hierarchy (with different prominent landmarks) may be “attached” to different people whom she is trying to meet in different situations. Thus, Home would be the obvious place to meet her boyfriend, as the cafeteria would be the natural choice to meet a friend while in school.

Most heuristic properties of Focal points mentioned by (Kraus et al. 2000) and (Zuckerman et al.'07) can be explained using the properties of prominent landmarks, mainly since they are easy to represent:

Uniqueness: a unique object is very distinctive from others. Such an object would be easy to represent, since the person only needs to encode one or two features of it, where as many features are required to distinct between similar objects.

Extremeness: an object with some extreme property (e.g. highest) can often be represented by this property alone, with no need for quantities information at all. These object are referred to by (Olson & Bialystoc '83) as “canonical objects” (pp.235-236). Other objects may be represented relatively to such an object.

Centrality: again, representing a central object only requires one or two predicates, and no quantitative information. Representing a central object also reduces the cost of remembering many other objects, which can now be represented as “near” the central one.

Firstness: it makes sense that later objects will be represented relatively to those who appeared earlier, especially in an unfamiliar environment.

The experiment

Preliminary work

In order to perform an experiment that will test my main hypotheses, few steps had to be done:

- Formally define the metrics what will be used to evaluate the agents. I.e. their ability to find focal points.
- Building an artificial environment. In such an environment agents can play and their performance may be tested.
- Evaluating the "Focality" of different land marks in the environment.

The following sections describe the formal definitions I used, the artificial environment, or testbed, which was created for the project, and the preliminary experiment in which focal points were evaluated using human agents.

Definitions

In this section I define some basic metrics that I will use later in the analysis section. These metrics will allow me to grade different agents based on their performance. The metrics I present here are not specific to the problem in question, and can be used in any domain to measure focal points and biased behavior.

Definition 1: a *Task* is a tuple (k,**A**) where:

- k is the number of Goals available in the task.

  k defines a 2-player pure coordination game, where the pure strategies of each agent are the alternative goals of the task: $\{1,2,..,k\}$. The payoff matrix of this game is the standard matrix of pure coordination games. i.e a diagonal matrix of 0's and 1's (as in fig. 1).
- **A** is the base agent, defined by his strategy.

  Since a strategy is a distribution on $\{1,..,k\}$, I denote $(p_1^{\mathbf{A}}, p_2^{\mathbf{A}}, ..., p_k^{\mathbf{A}})$ as the strategy of **A**.

Task defines a single player game, whose payoff is the payoff in the two-player game, played against **A.**

Definition 2: the *Focality* of a goal $j \in \{1,2,...,k\}$ in the Task (k,**A**), is the probability that **A** will choose j. Denote:

$F(j,\mathbf{A}) := p_j^{\mathbf{A}}$

Definition 3: An agent A is *Biased* if its strategy is not the equal distribution (1/k,1/k,..,1/k). If p_j^A>1/k then agent I is *biased towards* goal j. if p_j^A<1/k then agent A is *biased away* from goal j.

Definition 4: a Task (k,**A**) is *biased* if **A** is biased.

Definition 5: Let T be an agent playing a Task. The *Coordination index* of T is defined as its probability to coordinate with player **A**. let $(q_1^T, q_2^T, ..., q_k^T)$ be the strategy of T, thus:

$$Coord(T,\mathbf{A}) := \sum_{j=1}^{k} q_j^T p_j^A$$

This definition is not good enough to be used as a grade, since the coordination index ignores the fact that players may also coordinate by chance, especially when k is low. I would like to grade the agent according to the probability of *coordinating intentionally*.

Let *p(I)* be this probability, and assuming that meeting by chance has a probability of 1/k, then it holds true that:

$$Coord(T,\mathbf{A}) = p(I) + (1 - p(I))\tfrac{1}{k}$$

Definition 6: the *Grade* of Agent T in the Task (k,**A**), is *p(I).* Therefore:

$$Gr(T,\mathbf{A}) := p(I) = \frac{k * Coord(T,\mathbf{A}) - 1}{k - 1}$$

The *Grade* has the following desired properties (for k>1):

Result 6a*: Gr(T,**A**) = 0 iff T is not biased or the Task is not biased. The Grade of T is positive iff T is biased towards the same goals as **A**. Thus, the expected Grade of a player T who chooses a goal at random will be 0.

Result 6b: Gr(T,**A**) is an increasing function of k and Coord(T,**A**).

Result 6c: $Gr(T,\mathbf{A}) \xrightarrow[k\to\infty]{} Coord(T,\mathbf{A})$

Result 6d: $Gr(T,\mathbf{A}) \le 1$, with equality iff T coordinates perfectly with **A**.

Result 6e: Gr(T,**A**) = Gr(**A,**T)

* This only holds if T,**A** are indeed distribution vectors. If the sum of **A** or T is less than one, grade may be negative. Such a case may happen if some agents fail to complete a task.

Definition 7: the *Focality of a Task*, is the probability that two agents playing like **A** (i.e. both are playing according to the mixed distribution $p^{\mathbf{A}}$) will meet. Denote:

$$F(\mathbf{A}) := Gr(\mathbf{A}, \mathbf{A})$$

Thus all the results of Grade also apply to Focality:

Result 7a: F(**A**) > 0 iff **A** is biased. Otherwise F(**A**)=0.

Result 7b: F(**A**) is an increasing function of k and Coord(**A**,**A**)[10].

Result 7c: $F(\mathbf{A}) \xrightarrow[k\to\infty]{} Coord(\mathbf{A}, \mathbf{A})$

Result 7d: $F(\mathbf{A}) \leq 1$, with equality iff **A** is a pure strategy (i.e. the agent **A** always chooses the same goal).

I will now extend the notion of grade to a set of agents.

Claim 1: there is a complete equivalence between a single agent (playing a mixed strategy), and a set of agents (playing pure or mixed strategies). i.e. for every agent T there is a set of agents using the same strategy, and for every set of agents **T** = {T1,..Tn}, there is a single agent that simulates the strategy of the set **T**.

Proof: the first direction of the equivalence is trivial – simply define {T} be the set of agents. For the other direction define a new agent T', playing according to the following distribution:

$$q_j^{T'} = \tfrac{1}{n}\sum_{i=1}^{n} q_j^{Ti}$$

It is left to show that $(q_1^{T'},...,q_k^{T'})$ is indeed a distribution:

$$\forall i,j \;\; q_j^{Ti} \geq 0 \;, thus \; q_j^{T'} \geq 0$$

$$\sum_{j=1}^{k} q_j^{T'} = \sum_{j=1}^{k} \tfrac{1}{n}\sum_{i=1}^{n} q_j^{Ti} = \sum_{i=1}^{n} \tfrac{1}{n}\sum_{j=1}^{k} p_j^{Ti} = \sum_{i=1}^{n} \tfrac{1}{n} = 1$$

From claim 1 all the definitions 1-7 and their results apply in the same way to sets of agents, i.e. when **A,** and T are sets of agents.

Thinking on **A** as a set of agents, F(**A**) can also be seen as their average grade on the Task (k,**A**).

[10] Coord(**A,A**) is very similar to the coordination index of a task defined in (Mehta '94a p. 6), except my index treats only the relative part of agents who chose some goal, and not their actual number. The difference is very small when n is big enough. Note that Mehta's index do not possess the same desirable properties as F(**A**), especially when k is low.

proof:

$$\frac{1}{n}\sum_{i=1}^{n}Gr(a_i,\mathbf{A})=\frac{1}{n}\sum_{i=1}^{n}\frac{k*Coord(a_i,\mathbf{A})-1}{k-1}=\frac{1}{n}\sum_{i=1}^{n}\frac{k*\sum_{j=1}^{k}p_j^i p_j^{\mathbf{A}}-1}{k-1}=\frac{k*\sum_{j=1}^{k}\frac{1}{n}\sum_{i=1}^{n}p_j^i p_j^{\mathbf{A}}-1}{k-1}=\frac{k*\sum_{j=1}^{k}p_j^{\mathbf{A}} p_j^{\mathbf{A}}-1}{k-1}=$$

$$=\frac{k*coord(\mathbf{A},\mathbf{A})-1}{k-1}=Gr(\mathbf{A},\mathbf{A})=F(\mathbf{A})$$

The testbed

As most of my work was theoretic, the testbed should have only demonstrated part of the ideas. Due to obvious problems I could not build a general system that will simulate the real world. Even in a synthetic environment, I had to be careful about what would be included.

The environment itself was constructed as a game in which the agent is represented by a small Mars rover. The rover perceives its immediate environment, and is required to perform tasks of navigation to certain landmarks. Some of these tasks are defined as Coordination tasks, in which there is more than one possible goal, and the agent has to choose one of them and go there. Although agents may start a given scenario from different locations, in my setting all the agents started from the same location[11]. Agents do not see other agents – when the agent reaches his chosen landmark, this landmark is registered as its choice. All the choices of the agents are logged for later analysis.

The environment I eventually used was consisted of 24 Tasks, divided to 4 scenarios. Scenarios contained about 10-25 landmarks, whereas in each task the agents were required to choose from 3-8 options. A more detailed description of the testbed appears in appendix I. Examples of scenarios appear in appendix II.

Evaluation of focal points

In order to test automated agents, it was necessary to rank the focality of the different goals. That is the Focality of each goal j. Since I wanted to compare my agents with human behavior, a group of HUJI students played the game to determine the Focality of each goal, in each task.

---

[11] I now question the fact whether that was the right thing to do. Although starting from the same location raises the chance of coordination, it may do so by other means than the hierarchical structure of the survey representation. Thus my main argument may me missed.

Subjects

76 students played the game and recorded their logs. Each student also added his/her sex and army background for later analysis.

Method

A version of the game was placed on CS computer system, allowing any CS user to play the game. A mail was sent to all undergraduate and graduate students of CS, inviting them to participate. An instruction page was attached to the mail (see appendix II).

Each subject was required to complete all the scenarios by performing all the tasks in each scenario. The first scenario was a tutorial and was not recorded.

The subjects were encouraged to complete the coordination tasks by going to the goal they believe most other players would go. A prize was offered to the 3 players with the most "average" behavior. The details of the grading mechanism were not explained. The log of each player was automatically saved in my private folders. The experiment took place between the dates 7-26/11/2006.

Evaluation

I used the definitions from the previous section, and the group of students as the base **A.** the Focality of each landmark could now be calculated as F(**A**). also, any new agent T which played the game, human or automated, could be evaluated using Gr(T,**A**). This grade is higher as T succeeds in finding the Focal points of the task (k,**A**), i.e. in choosing the same landmarks as most of the students in **A**.

It was stressed before that common cultural background is important for coordination. Therefore, all the students were divided to subgroups according to their sex and army background. Those are two factors I suspected might have a strong effect on both navigation and coordination. In tasks were no clear focal point emerges, each subgroup can be used as base (**A**) to test whether the agent T replicates the behavior of this subgroup.

Expected results in the evaluation phase

This work involves many factors, all of which introduce noise and uncertainty to the results: the use of statistics, representations and manipulation of data records, limitations of interface, and above all the fact that we deal with human behavior.

Still, I think that on the qualitative level I might be able to show some interesting things. (from now on, the term behavior refers to the points an agent selects in coordination tasks)

First, I presume that focal points are, indeed, a part of the human behavior and we hope to witness them clearly in the results. This is a pre-requisition to any other result or claim in this work.

If I presume correctly, many Tasks will have high Focality, i.e. they will contain one landmark whose focality substantially exceeds all others'.

Possible effect 1: clear focal points may rise in part of the tasks (scenarios) and be less obvious in others.

Possible effect 2: Instead of one dominant pattern of behavior, we may witness two or more clear patterns, dividing the participants to a similar number of subgroups according to cultural background (e.g. male vs. females).

I expect to get a clear focal point (say, F(**A**)>0.3) in at least 50% of the tasks when **A** is the entire group of 76 students.

I expect this rate to be a bit higher inside every subgroup.

Results

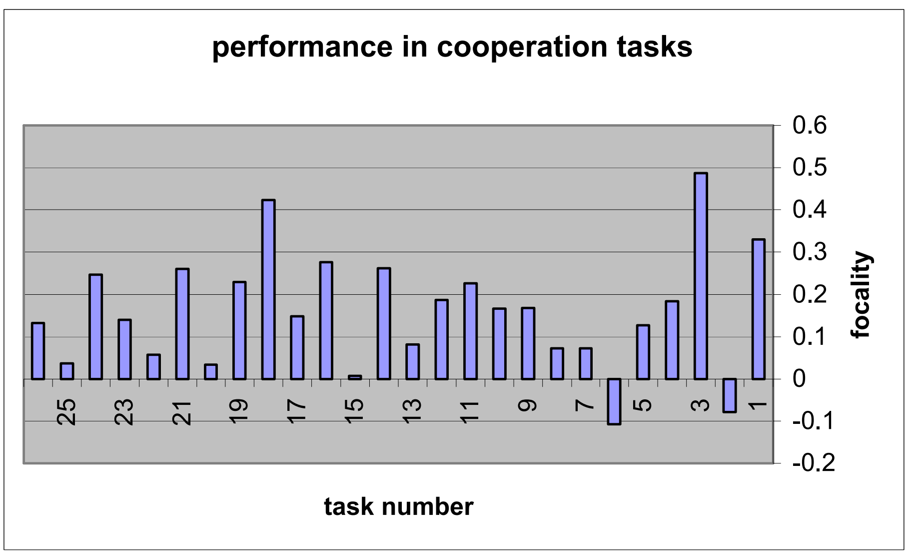

Figure 6. the performance (i.e. the Focality as defined earlier) of the entire group of human players. There where 26 cooperation tasks. Negative grade are due to tasks which some subjects did not complete.

As one can see in Fig. 6, most of the tasks did not have high focality. Only 3 tasks got a grade over 0.3, and the average grade was 0.16. These results indicate that there was indeed some bias in most tasks (i.e. people did not choose at random), but that bias was rather small[12]. Thinking that the results might be improved when the grade of subgroups is measured, I divided the subjects to the predetermined groups, and calculated the focality again. The only division that left subgroups of reasonable size (i.e. over 30) was according to army background. I also grouped the cooperation tasks according to their 4 types (see appendix I).

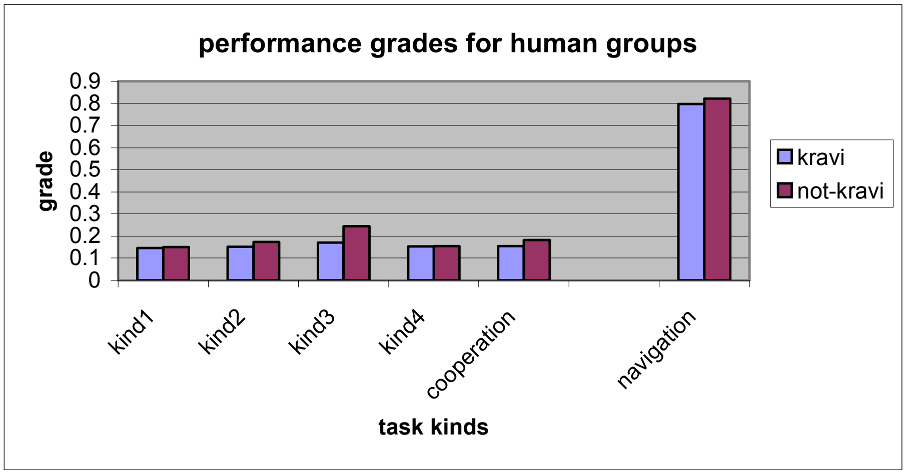

Figure 7. The performance of each subgroup in each type of mission is displayed. The "cooperation" column is the average of the 4 kinds.

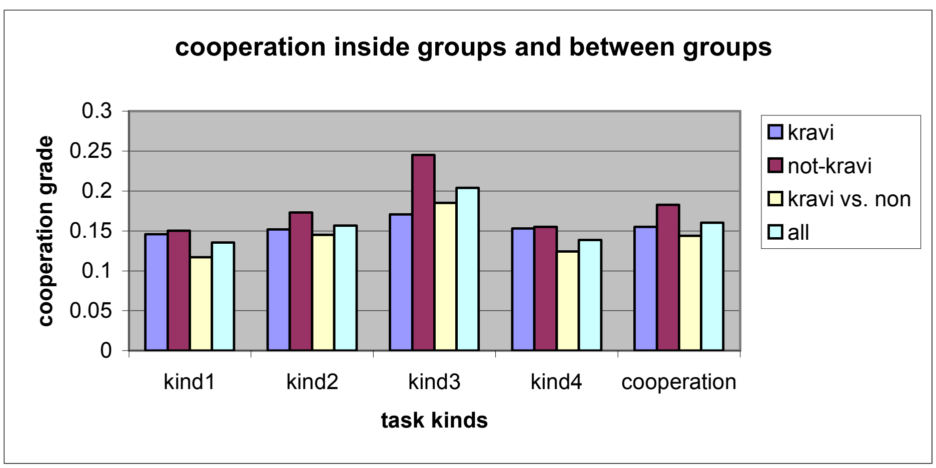

Figure 8. Cooperation inside each subgroup was slightly better than the cooperation between the two subgroups.

In Fig. 7 and Fig. 8 we can see that the same picture is kept in each subgroup as well. The slight improvement in focality when testing subgroups is expected, and is too

[12] For comparison, I calculated the Focality grade of the subjects in questions 11-20 in (Mehta et al. '94a), most of which were in a geometric setting. The grade for all questions was in the range 0.14 – 0.52, with average 0.31.

small to indicate a real difference in the behavior of groups. For that reason I will use only the unified group of all subjects as the base behavior.

Conclusions from the preliminary experiment

The relatively low grade that subjects achieved in most tasks can come from a few main sources. First, the amount of considerations that could be applied in selecting a goal was quite large, and mixed the problem of finding a goal (navigation), and deciding which goal to find (coordination). This difficulty however is an intrinsic part of such geographic settings, and the idea that coordination is affected by other factors through representation is exactly what I was trying to show. Second, the game might have been too complicated, with too many rules and instructions. If subjects understood the game differently, their behavior may also be different. Third, subjects may have used different mechanisms of way finding. I described two such mechanisms in previous chapters (route and survey), but the environment I planned cannot isolate one of them or distinct between them.

All the above problems result from the fact that the experiment was too complicated and had too many parameters.

Testing the hypothesis

The implemented Agents

There are 4 main types of automated agents:

- Type Random. Agents with no model of the environment, which act in a greedy way. That is, wandering around the environment, and completing a Task if any of its' goals are in sight. These agents demonstrate the random behavior to which other algorithms may be compared.
- Type MUA (Map Using Agent). Agents with a perfect model of the environment (i.e. a map), which choose coordination goals randomly, and use the map in order to find them.
- Type HUA (Heuristic Using Agent). Similar to MUA, only using heuristics to coordinate. Each type of map item has a fixed prominence, predetermined according to its size, uniqueness, and meaning.
- The rest of the agents create a simplified representation of the environment as they see parts of it, and use this representation both to navigate and to coordinate. I used several representations and algorithms:
    - TypeD2. Builds a relative mapping of the objects in sees. The information is stored in a graph, where each edge stores the accurate distance and direction between 2 nodes (items).

      The graph is used to recognize previously items, and the agent always represents itself relatively to the set of the last items it saw (the reference set). On seeing an unfamiliar item, the agent adds it to the graph, plus edges to all the items in the reference set. Typically this set is has 1-4 items, and the degree of nodes ranges in 1-6.

      When a Navigation/Coordination task arrives, if there is a goal in sight, the agent walks to it. Otherwise, the agent uses the graph to find the shortest path to a goal of the task, and then follows this path (this may not be the closest goal).

      There is no hierarchy between items in the representation.

      Regions are not represented at all.

- TypeD3. Acts quite similar to TypeD2, with a few differences. (a) each item has a fixed prominence, based on its size and meaning. This is the same prominence used by HUA. (b) representation is limited to a tree-like structure, where each new item A is only connected to the most prominent item B1 in the reference set. Typically, this results in a tree where prominent nodes also have higher degree.

I will refer to HUA and MUA as the "map using agents" and to TypeD2, TypeD3 as the "relative representation agents". The Random agent will also be referred to as the "simple" agent.

Originally the experiment was supposed to have two phases. In the first phase the way finding behavior of human players in some scenarios was to be recorded. The empirical findings would have been used to enhance and calibrate the parameters of the cognitive models inspired by the theoretical findings in chapter two.

In the second phase, the built agents were to be tested against human players, in different scenarios of course.

Due to time constraints, the first phase was cancelled.

Testing

Each type of agent "played" in each scenario 100 times, and its performance was logged. Thus, each log contained a distribution on the goals of each task. This distribution was compared to itself to get the focality of each type. It was also compared to the distribution of the human base group **A**, and a grade was calculated for each type of agent on each task.

Apart from the four scenarios I designed specifically for this purpose, I created 3 more random scenarios, in which the automated agents tried to coordinate when starting from a **random** location. In these scenarios only the focality could be calculated, and there was no comparison with human players.

Expected results

Navigation

I expect the relative representation agents to be better than the simple agent and not much worse than the map using agents, which use much more resources.

Meeting with other automated agents

If deterministic agents will start the same task from the same place, they will always meet of course. In order to test such abilities agents must start from different locations. I expect that HUA will do well on Tasks where there are distinctive objects (e.g. unique in some way). The relative representation agents may be better in other Tasks.

Meeting with human agents

to evaluate the agents ability to find the "real" focal points, i.e. those which human agents use, I shall use Gr(T,**A**) where T can use nondeterministic behavior by running multiple tests and merge the logs, and **A** represents the set of students the agent is trying to coordinate with.

First of all, I would not expect the agents to succeed in doing any better than the average human. That means I expect $Gr(T,\mathbf{A}) \leq F(\mathbf{A})$.

If all my assumptions and claims are correct, the grade of my agents will be close to the focality of the group. I.e, the agent "thinks" like the group.

I expect that the relative representation agents will beat the simple agent and also the MUA. HUA may be better in some tasks.

The best results will be if one of the relative representation agents will always get a score close to the Focality of the group.

Results

Navigation

As seen in Fig. 9, all agent types got high grade in navigation tasks, with a slight edge to the map using agents (MUA & HUA). Since even the simple agent, which did not use any model of the environments succeeded in over 80% of the tasks, I cannot really compare the benefit of the different models to navigation.

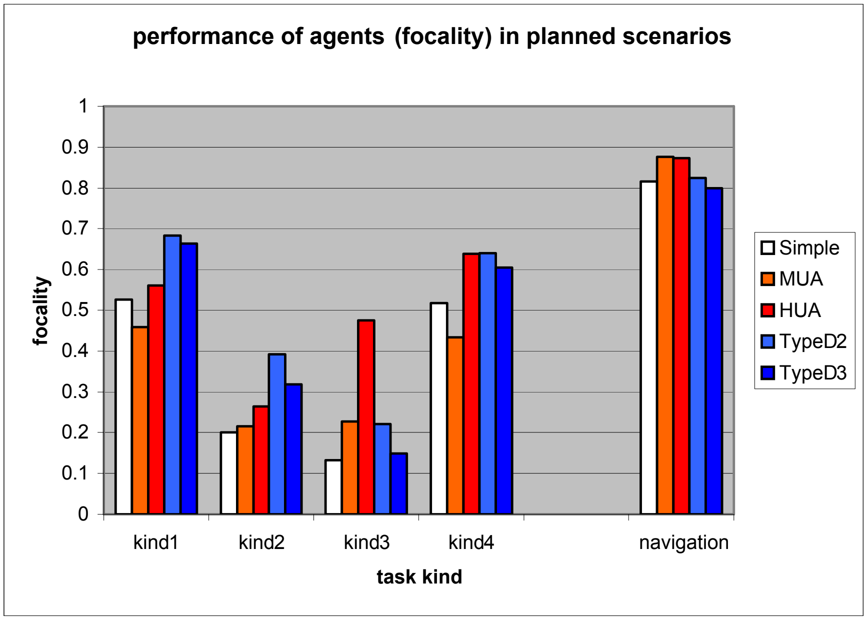

Figure 9. here we can see the performance of each type of agent in the different tasks. The right-most columns demonstrate the percentage of completed navigation tasks.
The other columns show how well each type of agent coordinated with agents of the same type.

Meeting with other automated agents

Here the results vary between task kinds. While the relative representation agents succeeded more in most planned tasks, they showed poor performance in coordinating on an item in a specific region (kind3 tasks). This is not surprising though, if we remember that regions were not represented at all by these agents.

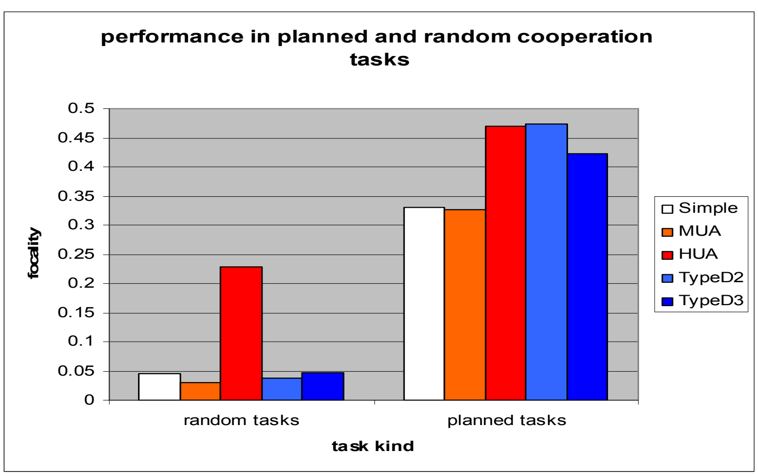

Figure 10. the left columns show the performance of agents in random scenarios, starting from random locations. The right columns are the "coop-all" columns from fig.9.

When tested on random tasks, the performance of the relative representation agents drops significantly, and only HUA shown any ability to coordinate. This shows that HUA was the only agent who could use global information effectively.

Meeting with human agents

In first glance, it appears that all agents performed poorly in all task kinds, due to the low average grades on figure 11. The results seem even worse since the relative representation agents performed even worse than the other agents in most tasks.

But averaging over task kinds may hide a more complicated picture.

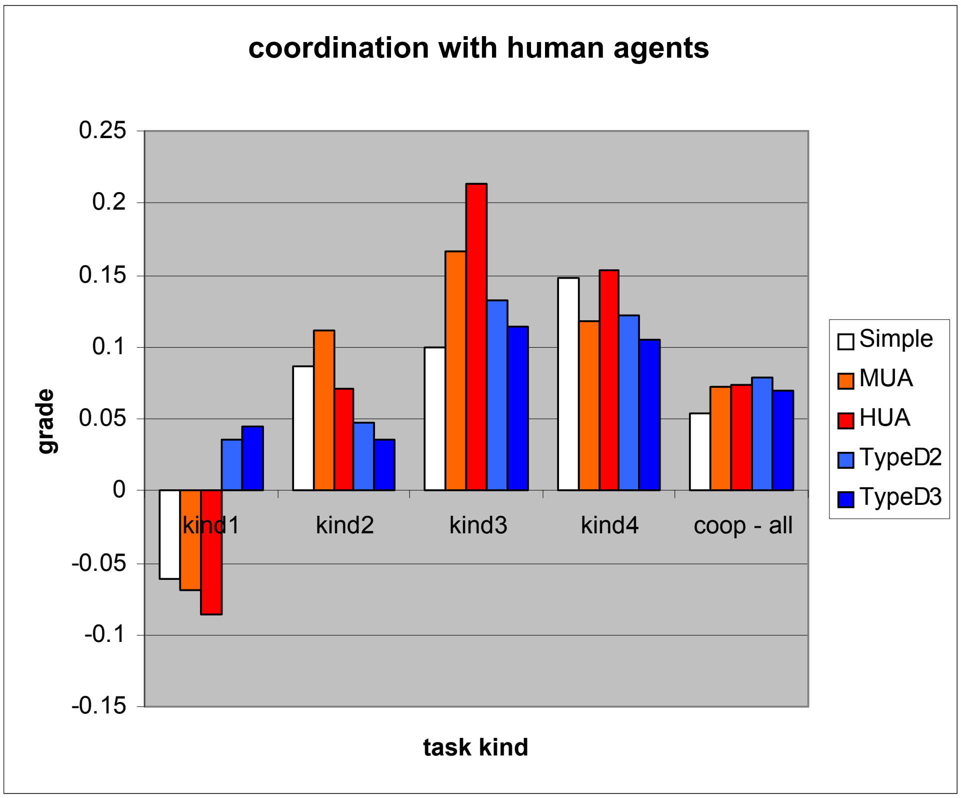

Figure 11. coordination grade between the automated agents, and the base group of human agents.

In Figure 12 we can see that the relative representation agents showed reasonable performance in most tasks. HUA, on the other hand, performed excellent on some tasks, and very poorly on others.

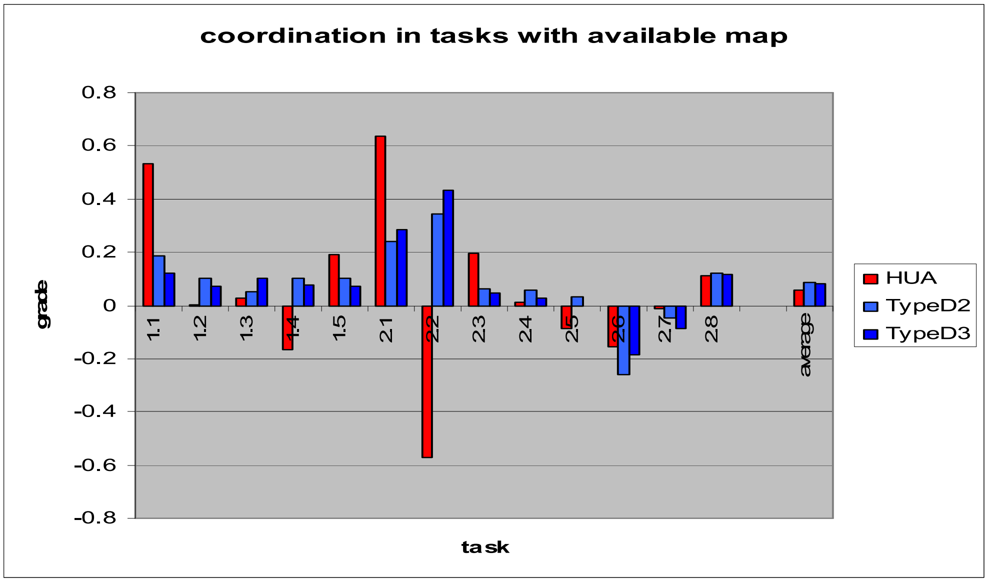

Figure 12. coordination with human agents on each task in the first two scenarios, where the map was available to the map using agents.
Here the relative representation agents are only compared to HUA, which out-performed MUA and the simple agent in most tasks.

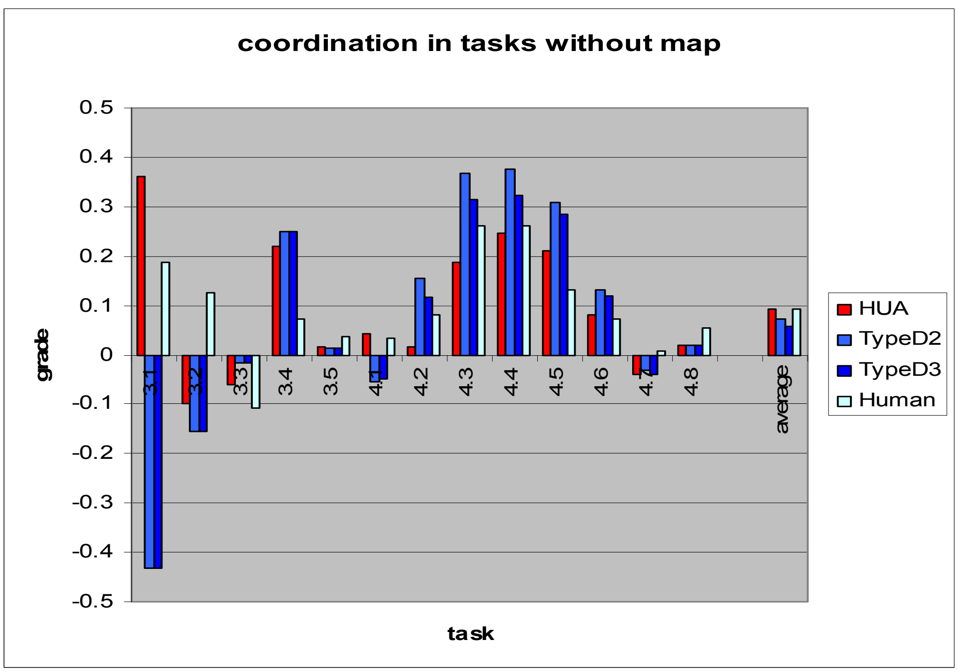

Figure 13. Coordination with human agents on each task in the last two scenarios.
The performance (focality) of human agents in these tasks is also shown.
Since the map was not available, HUA, MUA and simple agents had the same behavior.

In the scenarios with no available map (fig. 13), all agents performed very poorly (<0.05) on 6 tasks. In 5 of these tasks the human agents themselves showed poor performance (3.3, 3.5, 4.1, 4.7, 4.8).
In all but one of the remaining tasks, the relative representation agents outperformed all other agents.

Discussion

Taking into account the strategies taken by the different agents, most of the results make sense, and fit my previous expectations.

The HUA showed the best ability to coordinate with agents of their kind, both in planned and random scenarios. As I mentioned in the background, (Kraus et. Al 2000) have already demonstrated coordination between agents, using heuristics similar to mine, and others.

The agents that used relative representations of their environment demonstrated better ability to coordinate with human agents, although this ability varied between tasks, and was not significant.

Specific pitfalls that affected agents

- The relative representation agents constructed an internal representation that was strongly biased by their starting point. This bias helped them to coordinate in the planned scenarios, where they all started in the same place (and so did the human agents). In the random tasks however, those biases made them go to different places and fail to coordinate.
- Entire parts of the environment, such as regions, were not at all represented or handled by some agents. The effect of this over-simplifying was less than I feared though.
- Due to the lack of training phase during the implementation of agents, the experiment may have failed to show the full power of models and heuristics.

I claimed in the second chapter that the mental representation of the environment consists of two interrelated descriptions: the global **survey** description, and the viewer-dependent **route** description.

The heuristics used by HUA completely ignored spatial organization of map items, thus all route information is lost, but some aspects of the survey description are captured by the used heuristics.

TypeD2 encoded the spatial relations between items, but with no global structure, and thus all survey information is lost.

This may explain why some tasks were addressed well (i.e like a human) by HUA, while others where better handled by TypeD2.

I tried to incorporate the best from the two worlds in TypeD3, which used both relative descriptions and a hierarchy based on items' attributes, but the very simple tree representation and the fact that it could not modify existing parts of it, left it still worse than TypeD2.

General critique

The main problem, as I noted earlier regarding the preliminary experiment, is that I used two models (survey and route), but did not have a robust way of combining them. Nor even I could tell which one will dominate the behavior of human agents in different situations. Thus, trying to model such unpredictable behavior is problematic from two reasons. First, it is harder, since I had to make many extra assumptions that are not part of the original models. Second, I can explain the results in a more flexible way by choosing post-hoc the model used in each task. This may seem as an upside, but a model that can explain a wide range of results is harder to refute, and therefore suffers from lack of parsimony.

Summary

In this work I intended to cover the work done so far in the field of focal points, and show that while existing techniques may under some circumstances give us some way of predicting focal points, they cannot account for **why** some points have a higher focality than others. I tried to come with a new way for understanding spatial focal points from the cognitive point of view. I suggested a simple model for spatial representation based on findings from cognitive and developmental psychology. This model consists of the following principles:

- Spatial landmarks in the environment are represented relatively to other landmarks.
- Some landmarks are more prominent than others, due to considerations of representation difficulty and utility.
- The spatial relations are encoded in a simplified way, using a small set of predicates.
- Landmarks may be represented with redundancy.

These principles induce a hierarchy over the different landmarks. As the model is quite abstract and contains many parameters, this hierarchy is not strict and may vary between people.

Focal points arise in situations where this variability is low, and people do represent the environment using similar hierarchy.

To demonstrate the ability of this model to reproduce focal points, I built two agents that implement aspects of the model and tested their performance in coordinating with human agents, as well as with other agents of their kind.

Other agents that acted at random, or used pure heuristics were used to compare the results. While the agents built according to my model showed poor results in coordinating among themselves, they out performed the other agents in tasks involving coordination with human agents. I see this as reinforcement to my claim that the understanding of cognitive representation is crucial to the understanding of focal points. I attribute the fact that the results are inconclusive and suffer from low significance to the following problems:

a. The experiment was too complex. A good experiment in cognitive psychology usually tries to isolate one or two parameters, while keeping everything else constant. My environment had too many parameters, and also other aspects beyond my control. Since the results of the human agents were of low quality, any other result based on it would necessarily be insignificant.

b. my model was too general, and did not go into important details of the representation. Specifically I did not handle the relation between route and survey description. I went over the implications of this vagueness in the last chapter.

c. most of the effort in building the agents was spent on the technical parts of handing input and maintain the internal state. Thus the implementation of some important aspects of the model was postponed to a later stage (yet to have happened…).

To improve the results and testability of a similar project, it is required to focus on a much simpler environment and tasks. Alternatively, much more effort and experimentation are needed in order to fit the theoretic model to the test environment.

Lessons learned

The first lesson regards the tradeoff between developing an idea and testing it.

I think that the notion of having a good and thorough testbed is important and even crucial when trying new ideas, but time scale should also be a consideration. On the bottom line, I spent almost the entire time on building the testing mechanism, and too little time in implementing the model and developing the agents.

A second thing is the amount of effort and consideration required when conducting an experiment with human participants. I don't see how I could complete this project without such an experiment, but I will keep this in mind in choosing projects from now on.

But most importantly, I enjoyed very much working on the project, and I believe I learned a lot – on cognitive models, programming, focal points, GIS, and everything between them. It was a good opportunity to engage in a multi-disciplinary research, and I hope to continue working on similar projects in the future.

Future directions

The work presented in this project only tried to demonstrate the qualitative link between the way people represent their environment and the way focal points are found in geographical settings. It remains to further explore this connection in order to be able to apply the knowledge acquired in cognitive psychology on **real** day-to-day situations, and use it to forecast the focal points in those situations. As situations are more complex, other descriptions of the environment affect way finding, and therefore affect coordination and focal points. Those should be explored as well. It could be interesting to take existing complex agents that already use hierarchical representation and reasoning, and see if adding some simple coordination rules will enable them to find focal points in real world situations.

Another direction is trying to generalize the hypothesis about hierarchical representation to other fields than spatial domains. It is important to note though, that the presence of a hierarchical structure is not sufficient to the appearance of focal points. For focal points to occur, people must use similar (more or less) hierarchies as well. I hypothesize that a good rule of thumb to determine the location of an item in the relevant hierarchy would incorporate the difficulty raised by representing the specific item, versus the expected utility gain.

I suspect both hierarchical structures and similarity may be found in semantics and language, for example.

Bibliography

Aumann, R. J. (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality." *Econometrica,* Vol. 55, No. 1, pp. 1-18.

Bacharach, M. (1993) "Variable Universe Games." in Binmore K., Kirman A. and Tani P. (eds.), *Frontiers of Game Theory*, MIT Press, Cambridge

Crawford, V.P. and Haller, H. (1990) "Learning How to Cooperate: Optimal Play in Repeated Coordination Games." *Econometrica*, Vol. 58, No. 3, pp. 571-595.

Egenhofer, M. J., Sharma, J. and Papadias, D. (1996) "Hierarchical reasoning about direction relations." *Proceedings of the 4th ACM Workshop on GIS.*

Freksa, C. (1999) "Spatial Aspects of Task-Specific Wayfinding Maps – A Representation-Theoretic Perspective." In Gero, J. and Tversky, B. (eds.), *Visual and Spatial Reasoning in Design*, University of Sidney, Key Centre of Design Computing and Cognition, pp. 15-32.

Goldman, C. and Zilberstein, S. (2004) "Decentralized Control of Cooperative Systems: Categorization and Complexity Analysis." *Journal of Artificial Intelligence Research* (JAIR-04), Vol. 22, pp. 143-174.

Hart, R. and Berzok, M. (1982) "Children's Strategies for mapping the Geographic-Scale Environment." In Poetgal, M. (ed.), *Spatial Abilities*, Academic Press, New York, pp. 147-168.

Janssen, M. (1998) "Focal points." In Newman P. (ed.), *New Palgrave Dictionary of Economics and the Law* (vol. II), MacMillan, London, pp. 150-155.

Kuipers, B. J. (2000) "The Spatial Semantic Hierarchy." *Artificial Intelligence* Vol. 119, pp. 191–233.

Kuipers, B. J. and Byun, Y.T. (1988) "A robust qualitative method for spatial learning in unknown environments." *Proceedings of the National Conference on Artificial Intelligence* (AAAI-88), Los Altos, CA: Morgan Kaufman.

Kraus, S., Rosenschein, J. S. and Fenster, M. (2000) "Exploiting Focal Points Among Alternative Solutions: Two Approaches." *Annals of Mathematics and Artificial Intelligence*, Vol. 28, No. 1-4, pp. 187-258.

Mehta, J., Starmer, C and Sugden, R. (1994a) "The Nature of Salience: An Experimental Investigation of Pure Coordination Games." *The American Economic Review*, Vol. 84, No. 3, pp. 658-673.

Mehta, J., Starmer, C and Sugden, R. (1994b) "Focal points in pure coordination games: an experimental investigation." *Theory and decision,* Vol. 6, Issue 22, pp. 163-185, Kluwer.

Newcombe, N. (1982) "Sex-Related differences in spatial Ability: problems and gaps in Current approaches*."* In Poetgal, M. (ed.), *Spatial Abilities*, Academic Press, New York, pp. 223-243.

Olson, D.R. and Bialystok, E. (1983) *Spatial Cognition*. Hillsdale, New Jersey.

Pick, H.L. and Rieser, J.J. (1982) "Children's Cognitive Mapping." In Poetgal, M. (ed.) *Spatial Abilities*, Academic Press, New York, pp. 107-127.

Tversky, B. (1993) "Cognitive Maps, Cognitive collages, and Spatial Mental Models." *Lecture Notes in Computer Science*, Vol. 716 (Spatial information Theory), pp.14-24.

Schelling, T. (1960) *The strategy of conflict.* Oxford university press.

Schermerhorn, P. and Scheutz, M. (2006) "Social Coordination without Communication in Multi-Agent Territory Exploration Tasks." *Proceedings of Autonomous Agents and Multi-Agent Systems* (AAMAS).

Schultz, A. C. and Trafton, G. J. (2005) "Towards collaboration with robots in shared space: spatial perspective and frames of reference**."** *Interactions*, Vol. 12, Issue 2, pp. 22-24.

Sudgen, R. (1995) "A theory of Focal Points." *The economic Journal*, Vol. 105, No. 430, pp. 533-550.

Zuckerman, I., Kraus, S., and Rosenschein, J. S. (2007) "Using Focal Point Learning to Improve Tactic Coordination in Human-Machine Interactions." *The Twentieth International Joint Conference on Artificial Intelligence* (IJCAI 2007), Hyderabad, India, pp. 1563-1568.

Appendix

Appendix I.* Details *of the testbed

After choosing the sub domain of geographical navigation, I had to decide what kind of environment to implement. Although it is tempting to create an environment that will simulate real situations, it may pose difficult problems.

Two (negative) examples:

1) a **3D environment** would have taken a huge effort in implementation.

2) **Natural language messages** are simple to support, but will require agents to be able to interpret them, and the agents' designers to put a lot of effort into something of minor relevance to project.

Therefore, i chose a constrained set of properties of a spatial environment, which could be implemented with reasonable effort. Then, i picked a subset of those properties which agents should be able to deal with, and implemented them. At a later stage this subset may be extended, as agents get more sophisticated (i.e. in possible future work). The environment is built in a way that allows such extensions.

In selecting properties, I tried so pick those that are most relevant to spatial navigation according to the literature (e.g distinguishable landmarks, the ability to physically move in the environment). If my theory is correct, then these are the very same properties that support focal point.

Originally planned environment

With unlimited time and resources, the environment would have contained all these features:

- a 2D, look-from-above playground, where agents can move simultaneously and perform coordination and cooperation tasks.
- A large set of "ground types" and artifacts, which agents can perceive and perhaps manipulate. (e.g. move objects)
- Ground types and artifact have free semantics (e.g. quicksand, bonus artifact)

  That is not known in advance, and might even change. (requires agents to have learning abilities)
- Landmarks of 0D (e.g. trees) ,1D (e.g. trails) and 2D (e.g. swamps or lawns).
- Barriers of 0D,1D and 2D.
- Unlimited size, cyclic environment.

- N agents, possibly of different kinds.
- A variety of ways in which an agent can be exposed to the environment:
    - using a map
    - "see" the immediate surroundings
    - Get GPS like coordinates
    - Get a verbal description
    - More…
- Perceived information can be partial / wrong / contradictive.
- A multi-directional messages system:
    - a synthetic language of predicates
    - a cost for sending a message
    - unreliability (noise)
    - more…
- The environment has many controllable parameters, to test agents in a variety of situations.

The implemented environment

- a 2D top, agent centered view.
- Agents can move in any direction.
- Movement is performed in synchronic steps. (i.e. in turns)
- a set 5 of ground types and ~10 item types.
- Each map contains no more than 5 item types.
- The semantics of all items is pre-defined and known.
- 0D and 2D land marks are ground types and artifacts, or combinations of them
- Maps can have any size. A typical can be crossed in 30-60 "steps", and contains about 10-40 landmarks.
- Maps are cyclic. (walking all the way to one direction will result in returning to the same place)
- There is no concurrency or networking.
- There are 3 game modes:

- ONLINE – 2 agents (typically a human and a computer) play simultaneously (in turns). Their behavior is compared to see how many times the succeeded in cooperation.
- COLLECT – 1 human agent is playing. logs of his behavior are recorded for later analysis (more on logs later)
- TEST – 1 agent is playing. His behavior is compared to a log which represents the typical behavior of a group of other agents (say, female human) in the **same scenario**. The behavior in the log defines the **focal point**, and so the new agent is tested against these focal points.

- The application allows 2 more options for analysis:
  - Comparing logs from previous runs.
  - Unifying a set of logs, to represent the behavior of a group in a single log.
- Data available to the agent:
  - All the rules of the game and semantics of items is available via written instructions and a graphical legend.

	Availability	Human agent (player)	Computer agent
Game rules	Full	Via written instructions	Via agent creator
Semantics of items	Full	Via graphical legend	Via agent creator
Immediate surroundings	Full	Via GUI, as described	A "perception buffer" is passed
Map, when available	Location of all items	Via GUI, in a small window	The map is passed
State of the agent	Score, energy, location	Via GUI	All data elements are passed
Tasks	On getting a new task	A description of the task in English, via GUI	A formal description is passed
Other agents	(was not used in the project)	Via GUI	Via the "percept buffer"

- In some scenarios, the map is not available.
- All available data is fully credible.
- Controllable env. Parameters:
    - Map size
    - Quantity and location of map items
    - Map availability
- Agents can receive messages (events/tasks) from the "mothership".

  agents have no transmitters.
- There is a small, predefined set of messages that can be communicated, which can be defined using few parameters. Example:

  [C,1,40,(2;4;6;5)], stands for "cooperation task no. 1, meet the other agent at one of the locations (2,4,6,5), you have 40 seconds".
- a run of an agent is a finite discreet sequence of actions. Each run is recorded

logs

by the driver and saved in 2 log files:

- Raw data (full_log) contains the entire sequence of actions performed by the agent. It can be used to replay the run in the GUI and observe the agent's behavior.
- Higher level data (task_log) contains only the result of each task the agent performed. It is used for analysis of focal points and the grading of agents.

All logs are written on the disk in a human readable format.

The map

- Every location on the map is represented by 2 floats – (X ,Y) coordinates.
- There are 0D landmarks like CRATERs and HOLEs.

  Each type of landmark has a predefined semantics. e.g. if the agents touches a TOWER it gets Energy recharge. A HOLE can kill it.

  A legend is available at any time, for human players.

  - There are 2D (rectangle) landmarks like ENERGY_FIELD, also with predefined semantics.
  - There are "drop points" (DP), all marked with a unique name.

Each agent sees different names. That means agent1 can see DP “alpha” where agent2 sees DP “Charlie”. If both go to point “alpha”, they may get to different locations.

- A map is saved on the disk as a file in a simple human readable format. An example can be found in appendix II.

The task set

0. (default task) wander around collecting crystals.

Navigation tasks:

1. go to a specific DP (e.g go to DP “Charlie”)
2. go to a 0D landmark of a specific type. (e.g find a BIG_CREATER)
3. go to the center of a 2D landmark of a specific type.
4. go to a specific corner of a 2D landmark of a specific type.

Cooperation tasks:

1. meet the other agent in a drop point, from a limited subset. (e.g. meet the other agent at one of the locations {b,d,f,g})
2. meet the other agent at a 0D landmark of a specific type.
3. meet the other agent at a 2D landmark of a specific type
4. meet the other agent (anywhere)

scripts

Scripts, or scenarios are constructed from a map, a set of game parameters, and a sequence of Tasks. A script is saved on the disk as a file.

Scripts are written manually. An example can be found in appendix II.

Logs are comparable iff they belong to runs made on the same script.

A screenshot of the game

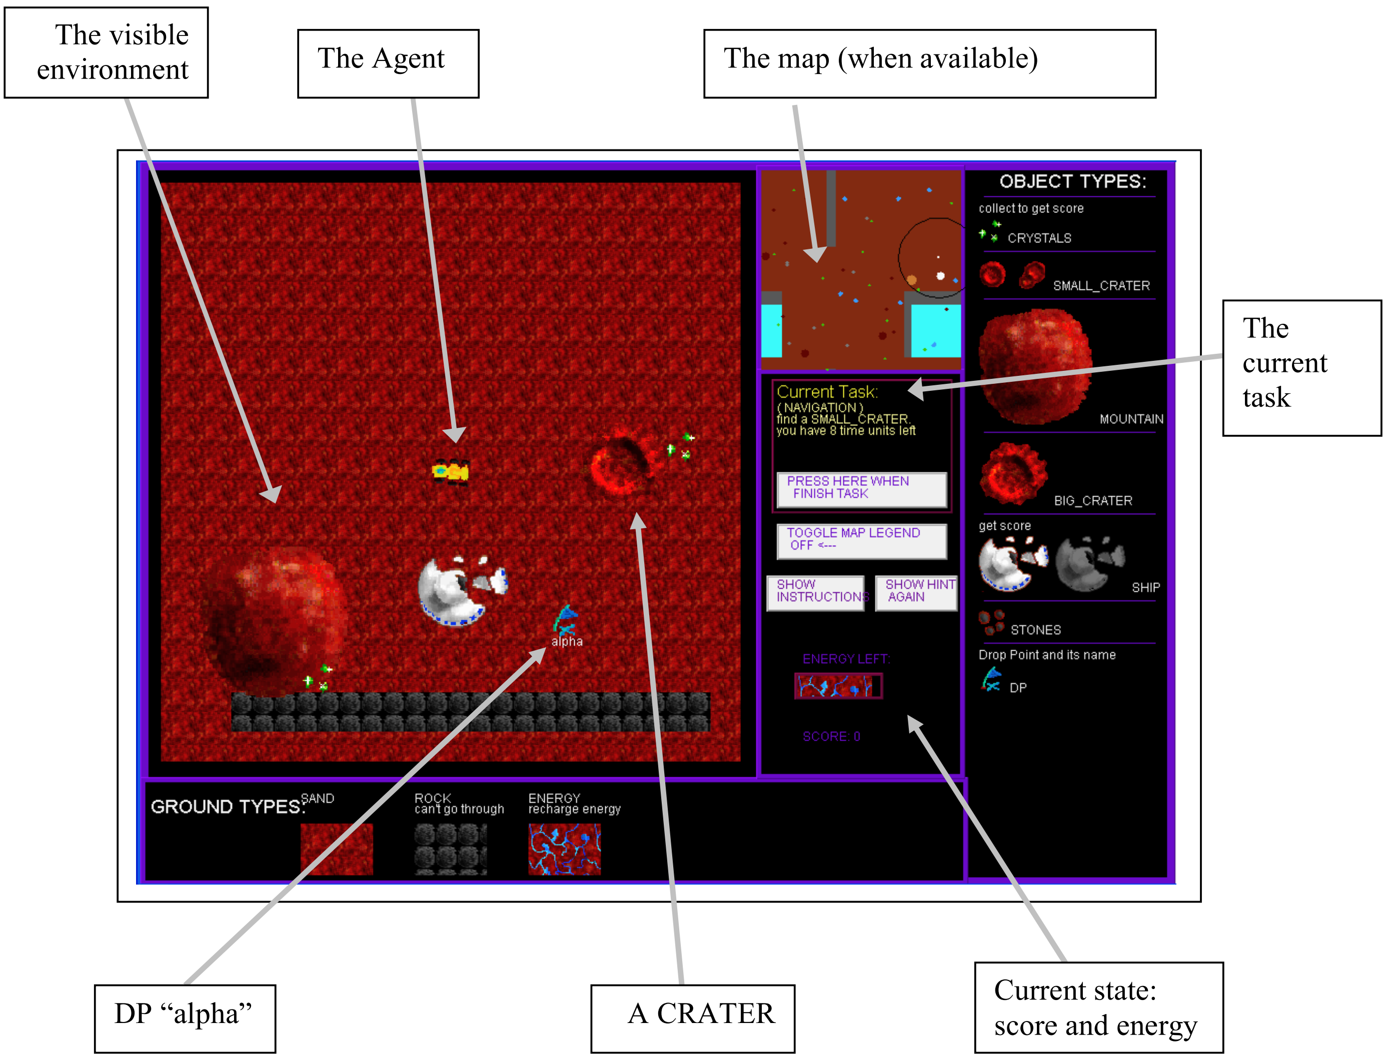

Appendix II. examples of maps and scenarios

Map of the first scenario

```
100 100        // size of the map
10 2           // number of landmarks and regions

STONE_PI       // a landmark and its coordinates
;80;25
STONES
;93;25
STONES
;80;40
STONES
;88;48
STONES
;70;12
BIG_CRATER
;25;80
SMALL_CRATER
;37;80
SMALL_CRATER
;25;70
SMALL_CRATER
;25;58
SMALL_CRATER
;15;58

DARK_SAND       // a region and its cordinates (corners)
;45;93;5;50
LIGHT_SAND
;99;49;65;5
```

Script of the first scenario

```
mission1_e.map                   // the relevant map filename
9                                // number of tasks
this is a very simple mission. take your time, and make correct
decisions.                       // description
0 0
MAP_CENTERED
SAVE
400 400
 20 TASK [N,2,20,SMALL_CRATER]  // navigation task, issued on time 20
 45 TASK [N,2,20,STONES]
 55 TASK [N,2,30,STONE_PI]
 65 TASK [C,3,0,LIGHT_SAND]   // coordination task, issued on time 65
 75 TASK [C,2,0,STONES]
 85 TASK [N,3,30,DARK_SAND,1]
 95 TASK [C,2,0,SMALL_CRATER]
 105 TASK [C,3,0,DARK_SAND]
 115 TASK [C,4,0]
```

Appendix III. The mail that was sent to CS students

----

this mail was sent with the permission of Prof. Jeff Rosenschein

and Prof. Michael Ben-Or

----

to all students hello,

i am a graduate student conducting a small experiment as part of my research project. participating takes 10-15 minutes in which you will be playing a game. prizes of 100nis will be given to the best 3 players.

you can play the game in your free time on any CS computer.

if you are interested, please read the attached instructions (in Hebrew)

thanks,

Reshef Meir.

reshef24@cs.huji.ac.il

(the attached instructions are on the next page)

הוראות להשתתפות בניסוי
הניסוי בנוי כמשחק מחשב מחולק לשלבים. הוא נמשך כ10-15 דקות.
מטרת הניסוי היא לבדוק את היכולת של השחקן לפעול בצורה דומה לשחקנים האחרים (שאותם הוא אינו רואה).

השחקן שולט בעזרת העכבר ברכב-חלל קטן שמנווט על פני כוכב לא מוכר.
במהלך המשחק תדרשו לבצע משימות שונות הקשורות להתמצאות וניווט בסביבה.
המשימות מתחלקות לשני סוגים:

משימות **ניווט** (NAVIGATION) – במשימה זו תידרשו להגיע לנקודה ספציפית בסביבה על פי התיאור או השם שלה. במשימות אלה הגבלת זמן, ובמידה שתצליחו תקבלו ניקוד.

משימות **תיאום/שיתוף פעולה** (COOPERATION) – במשימה זו תוכלו לבחור בין כמה מקומות אפשריים. מטרתכם להגיע אל הנקודה ה"טבעית" ביותר, כלומר לאותה נקודה בה לדעתכם יבחרו שאר השחקנים, עפ"י מיקומה.
(שמות המקומות הם אקראיים ומשתנים בין משחק למשחק)
לא תוכלו לראות שחקנים נוספים, ולכן לא תוכלו לדעת האם הצלחתם, עד לניתוח התוצאות שיתבצע בשלב מאוחר יותר.
במשימות אלה לרוב אין הגבלת זמן.

בתחילת המשחק תישאלו שאלות רקע, המסייעות בארגון המידע.
השלב הראשון אינו נרשם, והוא למטרת לימוד המשחק בלבד (tutorial).

הפעלת המשחק
היכנסו לספריה
~reshef24/mars_game/
הפעילו את הקובץ play
אם מופיע מסך ריק – הקליקו עליו עם העכבר.

איך מקבלים את הפרס?
בסיום המשחק, אתם אמורים לקבל mail לתא של cs, המציין כי תוצאותיכם נרשמו. בעוד כחודש (עד 31.11.06) כל מי ששיחק יקבל mail נוסף ובו הדירוג שלו בין כלל השחקנים.
שלושת השחקנים עם הדירוג הגבוה ביותר יצרו איתי קשר ע"מ לקבל את הפרס (100 ₪ במזומן).
- כיצד הדירוג נקבע?
לאחר איסוף כל התוצאות, מחושבת ההתנהגות ה"ממוצעת" במשימות התיאום. התוצאות של כל שחקן מושוות להתנהגות זו, וכל שחקן מקבל ציון עפ"י קירבתו לממוצע. *אין התייחסות לניקוד המושג במשימות הניווט, אלא למשימות התיאום בלבד.*

ניתן לשחק פעם אחת בלבד.

לדיווח על באגים / שאלות / הבהרות - reshef24@cs.huji.ac.il

תודה מראש,
רשף מאיר