

## MIT Open Access Articles

### *From problems to protocols: Towards a negotiation handbook*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Marsa-Maestre, Ivan et al. "From Problems to Protocols: Towards a Negotiation Handbook." *Decision Support Systems* 60 (April 2014): 39–54 © 2013 Elsevier B.V.

**As Published:** <https://doi.org/10.1016/j.dss.2013.05.019>

**Publisher:** Elsevier

**Persistent URL:** <http://hdl.handle.net/1721.1/113089>

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

**Terms of use:** Creative Commons Attribution-NonCommercial-NoDerivs License



# From Problems to Protocols: Towards a Negotiation Handbook

Ivan Marsa-Maestre<sup>a\*</sup>, Mark Klein<sup>b</sup>, Catholijn M. Jonker<sup>c</sup>, and Reyhan Aydoğan<sup>c</sup>

<sup>a</sup> *Computer Engineering Department, University of Alcala, 28803 Alcala de Henares, Madrid, Spain*

<sup>b</sup> *Center for Collective Intelligence, Massachusetts Institute of Technology (MIT), Cambridge, MA 02139, USA*

<sup>c</sup> *Interactive Intelligence Group, Delft University of Technology, Delft, The Netherlands*

## Abstract

Automated negotiation protocols represent a potentially powerful tool for problem solving in decision support systems involving participants with conflicting interests. However, the effectiveness of negotiation approaches depends greatly on the negotiation problem under consideration. Since there is no one negotiation protocol that clearly outperforms all others in all scenarios, we need to be able to decide which protocol is most suited for each particular problem. The goal of our work is to meet this challenge by defining a “negotiation handbook”, that is, a collection of design rules which allow us, given a particular negotiation problem, to choose the most appropriate protocol to address it. This paper describes our progress towards this goal, including a tool for generating a wide range of negotiation scenarios, a set of high-level metrics for characterizing how negotiation scenarios differ, a testbed environment for evaluating protocol performance with different scenarios, and a community repository which allows us to systematically record and analyze protocol performance data.

## 1. Introduction

Automated negotiation represents an important class of decision support system, helping self-interested decision makers rapidly reach high-quality agreements (Robu et al., 2005) (Fenghui et al., 2009) (Sycara, 1993). It has been studied extensively in e-commerce settings (Myerson and Satterthwaite, 1983) (Athey and Segal, 2007) (Lau et al., 2008) but can also be seen more generally as a paradigm for solving coordination and cooperation problems in complex multi-agent systems, e.g. for task allocation, resource sharing, or cooperative planning and design (Smith, 1980) (Hemaissia et al., 2007) (Wu et al., 2009) (Sim and Shi, 2009). The protocol used during the negotiation determines the rules of the negotiation. It determines when parties can make offers, how long offers stand, whether or not there are deadlines, or a fixed set of rounds, and all other “rules of encounter” (Rosenschein and Zlotkin, 1994). A wide range of negotiation protocols have been explored in research to date, ranging from auctions to iterative concession to mediated single- and multiple-text protocols, including both bilateral and multilateral variants (Parsons et al., 2011) (Raiffa, 1982) (Fatima et al., 2009) (Lai et al., 2004) (Kraus, 2006) (Sandholm, 1998) (Jennings et al., 2001).

A critical challenge facing negotiation research is that no protocol is the best choice for every possible negotiation scenario. Protocols that work quickly and well for simple negotiations with a few independent issues, for example, fare poorly when applied to complex negotiations with many interdependent issues (Klein et al., 2003). Some protocols may be fast but produce less-than-optimal results, and others the opposite. The question then becomes, *which* protocols should be used for *which* negotiation problems.

The research community has been unable to answer this key question, to date, because the different efforts all tend to test their protocols using their own idiosyncratic negotiation scenarios, varying widely across dimensions that include the number of agents, size of the contract space, degree of issue interdependency, and so on. It’s been like comparing apples and oranges. This problem is exacerbated by the fact that there currently is no systematic enumeration of what types of negotiation challenges exist, nor do we have standard tools for identifying the type of any given real-world negotiation problem.

Our work aims to address this challenge directly, and enable a much more systematic and cumulative effort by the negotiation research community. To achieve this, we have been creating the infrastructure depicted in Figure 1.

---

\* Corresponding author: ivan.marsa@uah.es




Figure 1. The Negotiation Handbook infrastructure.

This infrastructure consists of the following components:

- a *scenario generator* that creates a range of test scenarios
- scenario *metrics* that can be used to characterize key properties of such scenarios
- a *testbed* for evaluating protocols with these scenarios
- *repositories* that allow researchers to upload test scenarios, in addition to experimental results on how well their protocols performed on these scenarios, for use and analysis by the larger community

Researchers can use standard test scenarios already existing in the repository, or upload their own for use by others. Researchers can also use the testbed to evaluate their protocols, or do their own evaluation and simply upload the results to the experiment repository. The experiment repository can then be data-mined by any interested party to find correlations between scenario properties and protocol performance. As this process continues, the research community can converge upon a *negotiation handbook*, suited for use by negotiation practitioners, that specifies design rules concerning which protocols to use for which kinds of problems.

We describe, in the remainder of this paper, the progress we have made towards achieving these goals. We begin by defining what we mean by a negotiation scenario (section 2), describe the algorithms we are developing to generate test scenarios (section 3), review the metrics we have developed to capture key properties of such scenarios (section 4), and describe the protocol testbed (section 5) and repository elements (section 6) that are used to gather empirical data about the relationships between scenario properties and protocol performance. We conclude with a review of the unique contributions of this work and possible directions for future work.

## 2. What is a Negotiation Scenario?

A negotiation scenario should specify everything we need to know about a negotiation problem in order to pick the right protocol for the job. These elements include:

- The *domain* i.e., the range of contracts being negotiated over, which is defined by the set of issues being negotiated over as well as the valid values for each issue. The domain for a purchase negotiation, for example, would probably include *price* (whose values are positive real numbers) and *delivery date* (whose values are dates). A *contract* specifies a single valid value for every issue in the domain, and an *agreement* is generally the contract that all negotiating parties agree upon. Key properties of a domain include the *types* of the issues (e.g., continuous vs. discrete) as well as the *size* of a domain (i.e., the number of possible contracts).
- The number and characteristics of the *players* involved in the negotiation, including:

- *Agent preferences*, which specify how highly an agent values any particular set of issue values. Key factors here include the *types* of the agents preferences (e.g., ordinal vs. cardinal, fixed vs. mutable), *reservation values* (i.e., the threshold, possibly time-sensitive, for deciding whether or not to accept an agreement), as well as the *preferential dependencies* (i.e., whether the agent's preferences for one issue are influenced by the choices made for other issues). An agent's preferences may also include such factors as retaining goodwill with some other agents because they will likely negotiate again in the future.
- *Agent decision-making style*, such as the agent's *risk attitude* (i.e., its willingness to risk poor outcomes in the pursuit of better ones) (Eckel & Grossman, 2008), *cognitive style* (i.e., whether it bases its decision on economic rationality, emotion, or other factors) (Patton & Balakrishnan, 2010), *cognitive capability* (i.e., whether agents have sufficient computational capability to pick the best negotiation moves in the time available), *truthfulness* (i.e., whether agents choose to report their preferences in a truthful vs. a manipulative way) (Ellingsen et al. 2009), and characteristic *strategies* (e.g., some agents may adopt a Boulware strategy, wherein they concede as slowly as possible towards the other agents' proposals, until the negotiation deadline comes close). An agent's strategy, of course, will be deeply influenced by the protocol being used.
- The *context* of other negotiations. An agent, for example, can use knowledge of previous similar negotiations to predict the other agents' preferences and styles (e.g. it can predict that buyers always want to maximize price), which can in turn inform its own decision-making during a negotiation. Furthermore, there might be other negotiations going on at the same time for related topics, and the progress in the negotiations could mutually influence each other.

These elements define a large space, and many dimensions of this space are as yet incompletely understood. The set of possible agent strategies, for example, is large, and possibly unbounded. Accordingly we have chosen to focus, in our initial efforts, on characterizing the two most basic elements of any negotiation: the *domain* being negotiated over, and the *preferences* of the agents doing the negotiation. Since we have already discussed how to represent a negotiation domain, we focus in the discussion below on how to represent agent preferences.

There are many different ways to represent agent preferences. Perhaps the simplest is the *ordinal graph*, which consists of nodes (each representing one possible agreement) connected by links that indicate that the source contract is preferred, by the agent, to the target contract. Ordinal graphs are fully expressive (they can capture any internally consistent set of preferences), and can handle cases where only partial preference information is available (i.e. where some preference orders are not yet known), but they capture only *ordinal* preferences, providing no information on *how much* one contract is preferred over another. Ordinal graphs also scale poorly; capturing the preferences for a domain with 10 issues, 10 values per issue, would require a graph with 10 billion nodes.

If we wish to capture *cardinal* preferences (i.e. where we can specify how much an agent prefers one contract over another), we move into the realm of *utility functions*. A utility function is defined, for each agent  $j$ , as

$$U^j : D \rightarrow \mathbb{R},$$

It assigns, to each possible agreement in the domain  $D = \prod_{i=1, \dots, n} d_i$ , a real number representing the utility that agreement yields for agent  $j$ . Note that cardinal preferences *subsume* ordinal preferences, since we can always derive the preference order of two contracts by comparing their cardinal utility values.

The simplest form of cardinal utility function, and the one most widely used in the negotiation literature, is the *additive utility function* (Raiffa, 1982) (Brafman and Domshlak, 2009). In such functions, an agent's utility for a contract is calculated as the weighted sum of the agent's utility for each issue  $X_i$ :

$$U(X_1, X_2, \dots, X_n) = \sum_{i=1}^n w_i * U_i(X_i)$$

Additive utility functions provide a compact preference representation, but they are unable to capture the important phenomenon of *preferential dependency*, i.e., where the utility of the choice for one issue is dependent upon the choices made for other issues. Preferential dependencies are ubiquitous in real-life negotiations. When we negotiate over the stereo system we are going to purchase, for example, the utility of any particular issue value choice (e.g. which tuner we select) depends fundamentally on the choices made for other issues (e.g. whether we selected speakers that are compatible with that tuner). Utility functions without preferential dependencies always have a simple *linear* structure, with a single optimum, while utility functions with preferential dependencies can have a much more complex, *nonlinear*, structure, with multiple optima.

There is a range of utility function representations suited for dealing with preferential dependencies. *Auction preference languages* (Boutilier and Hoos, 2001) (Cramton et al., 2006) capture preferential dependencies in the context of allocation negotiations (i.e. where there are only binary issues, each representing whether an agent acquires a particular good or not). The preference language expression:

$$\langle (a \wedge b, 0) \vee (c, 5), 50 \rangle$$

for example, expresses the idea that the agent assigns a value of 50 to acquiring either goods *a* and *b*, or good *c*, and it assigns a premium of 5 to the latter.

*Utility graphs* (Robu et al., 2005) capture allocation preferences as a set of nodes (each representing the issue of whether or not a given good was purchased) in addition to a set of links between these nodes that capture the (positive or negative) complementarities between the goods. In the utility graph in Figure 2, for example, there are four goods as well as three dependency relationships.




Figure 2. An example utility graph, adopted from (Robu et al., 2005).

The utility of a given contract, to an agent, is given by the sum of the values for goods acquired, plus the sum of the values for the complementarity links whose source and target goods were both acquired. In this graph, for example, the contract [I1, I2, I3] (i.e. where the agent acquires the goods I1 I2 and I3) has the utility  $5 + 3 + 4 + -3 + 2 = 11$ . Utility graphs only capture binary preference dependencies (i.e. dependencies between pairs, but not higher-order tuples, of issues).

While allocation negotiations are of course very common, especially in e-commerce contexts, there are many real-world negotiations with N-ary issues that concern more than just whether the agent acquires a given good, or not. When engineers are negotiating over the design of a car, for example, many of the issues (e.g. What kind of engine should we have? What kind of transmission? How many passengers should we make room for?) will have more than two possible values.

*CP-nets* (Boutilier et al., 2004) capture preferential dependencies for N-ary issues using directed graphs in which each node represents the agent's preference for an issue, and each link captures the impact one issue choice has on the preferences for another. In the CP-net example in Figure 3a, for example, there are three issues involved in a vacation negotiation: season, location, and duration. The agent prefers Spring over Summer, Spain over Belgium, and has preferences for duration that depend on the selected season and location:




Figure 3a. An example of a CP-net.

CP-nets can be used with N-ary issues (i.e. where an issue can have more than just two possible values), and can be used when we do not fully know an agent’s preferences, but they capture only ordinal, as opposed to cardinal, utility information. Cardinal utilities can only be *estimated*, in this representation, by applying heuristics such as those described in (Aydogan et al., 2011).

UCP-nets (Boutilier et al., 2001) address this challenge by including utility values, rather than simply preference orderings, with each node in the network. In the UCP-net in Figure 3b, for example, the utility for the agreement [Spring, Spain, 2 weeks] is  $5 + 5 + .1 = 10.1$ :

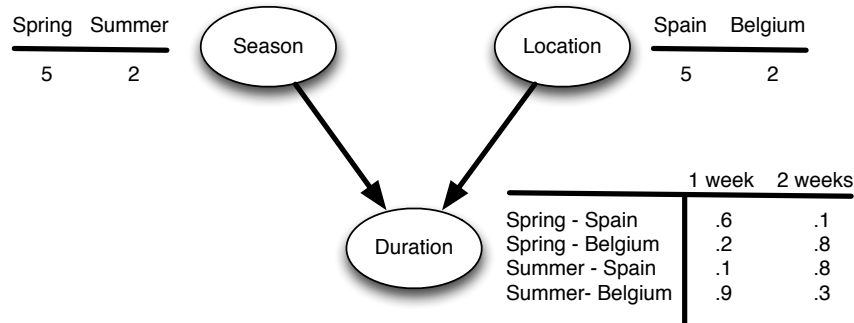


Figure 3b. An example of a UCP-net.

*K-additive utility functions* (Grabisch, 1997) are a generalization of additive utility functions, where the utility of a contract for an agent is the sum of the agent’s value for each issue choice *plus* the values due to the preferential dependencies between two (or more) issues. The following example:

$$U(j) = U(\langle i1 \rangle, j) + U(\langle i2 \rangle, j) + U(\langle i1, i2 \rangle, j)$$

calculates the utility for agent *j* of a contract as the sum of the values of the choices for issues *i1* and *i2*, *plus* the value due to the dependency (synergistic or antagonistic) between those two issues. This approach can be applied to preferential dependencies between any number of issues.

The *weighted hypercube* approach (Ito et al., 2008) represents a utility function as a collection of hypercubes that each assigns a numeric weight to a sub-volume of the domain of possible contracts:

```

utility_function      := (ufun <weighted_hypercube>*)
weighted_hypercube   := (hypercube <weight> + <bound>*)
bound                := (<issue> <min value> <max value>)

```

The utility for a given contract is then calculated as the sum of the weights for the hypercubes that *include* that contract, which means the contract is within all the bounds the hypercube imposes to the different issues. Imagine

we have a scenario with three issues, each of which can have a value between 1 and 20. Imagine, further, that the utility function for an agent consists of the following weighted hypercubes:

| Weight | Bounds  |
|--------|---|
| 10     | (issue1 from 2 to 4)  |
| 12     | (issue2 from 3 to 6) (issue3 from 12 to 17)                       |
| 7      | (issue1 from 1 to 5) (issue2 from 7 to 10) (issue3 from 10 to 12) |

The contract [issue1 = 5, issue2 = 4, issue3 = 11] is included in the second and third hypercubes, so the utility of that contract is  $12 + 7 = 19$ . This representation is capable of representing any imaginable utility function with a discrete domain; at worse, we simply need to create unit-volume hypervolumes for every contract in the domain.

We can summarize these different utility representations using the following table:

|                                  | binary issues<br>(allocation negotiations) | N-ary issues<br>(general negotiations)      |  |
|----------------------------------|--|---|--|
| no preferential dependencies     |  | Additive utility functions                  |  |
| binary preferential dependencies | Utility graphs                             |   |  |
| N-ary preferential dependencies  | Auction preference languages               | <i>Ordinal</i><br>Ordinal graphs<br>CP-nets | <i>Cardinal</i><br>UCP-nets<br>k-additive functions<br>weighted hypercubes |

The most expressive utility representations are the three (in the lower right cell) which support N-ary issues, N-ary dependencies, and cardinal utilities. These three representations are equivalent in terms of expressiveness, since they are theoretically capable of representing any well-defined utility function for discrete issue domains.

We adopted, in our work, an extended version of the weighted hypercube representation, wherein the hypervolume weights can be a constant, linear, *or* sine function of the issue values:

```

scenario := (scenario :name <string>
            :description <string>
            :domain <domain>
            :agents (<agent>*))

domain := (<def>*)

def := (issue <string> real from <number> to <number>) |
       (issue <string> integer from <number> to <number>) |
       (issue <string> one-of <string>*)

agent := (agent :name <string> :ufun ufun)

ufun := ( <hypervolume>* )

hypervolume := (cube (<bound>*) <weight>) |
              (plane (<bound>*) (<weight>*) <bias>) |
              (sine (<bound>*) <height>) |

bound := (limit <string> from <number> to <number>) |
         (limit <string> one-of <symbol>*)

```

This representation can capture any imaginable utility function for both discrete *and* continuous domains, since we know that any waveform can be represented with arbitrary accuracy (using the Fourier transform) as the sum of sine waves (Bracewell, 2000). Figure 4 presents two examples of utility functions (for a two-issue contract domain) that can be generated using a weighted hypervolume approach.




Figure 4: Examples of utility functions expressed using weighted hypervolumes.

---

The utility functions have high regions where many high-weight hypervolumes reside, and lower regions where there are few and/or low-weight hypervolumes.

### 3. Scenario Generation

How can we acquire the scenarios we need to test different negotiation protocols? One obvious approach is to simply collect the scenarios that have been used in previous negotiation research and store them in our scenario repository for general use (we describe how this can be done in section 6: “Repositories”). A key limitation with this approach, however, is that there is a vast number of possible negotiation scenarios, and previous research has necessarily only covered a tiny fraction of this space, predominantly scenarios without preferential dependencies. We need, therefore, a way to cover the scenario space more comprehensively. This is the purpose of the scenario generator. It allow users to generate negotiation scenarios by systematically varying a compact set of high-level parameters, including (1) the *negotiation domain*, including the issues and possible issue values, (2) a *scenario utility histogram* (see below), which defines the distribution of contracts in utility space, and (3) several *structural parameters* that specify the topography (e.g. ruggedness) of each agent’s utility function. Our scenario generator then creates, using a stochastic process, utility functions (sets of weighted hypervolumes) for each agent that satisfy the user’s specifications. We describe this process in the sections below.

#### 3.1. Specifying Scenario Utility Histograms

*Utility histograms* are a generalization of utility diagrams. The latter provide a simple way to visualize the possible agreements between negotiating agents. In such diagrams, there is one axis for every negotiating agent, and one point per possible agreement. The position of a point represents the utility of an agreement for each agent. In the (two-agent) utility diagram in Figure 5a, for example, the y axis represents the utility of an agreement to agent A, and the x axis represents the utility of an agreement to agent B. The Pareto front represents agreements that cannot be improved upon for one agent without sacrificing utility for other agents. The optimality of any proposed agreement can be assessed by measuring its distance from the Pareto front.



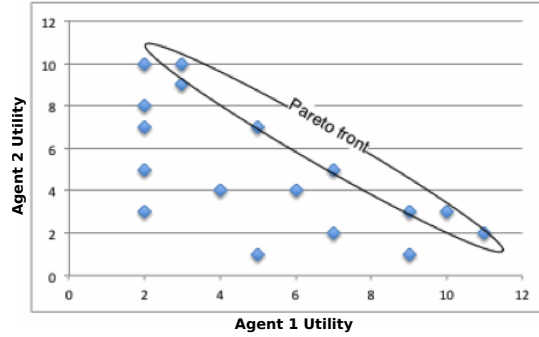


Figure 5a. An example of a scenario utility *diagram*.

The problem of scenario utility diagrams is that they do not capture the multiplicity of agreements. That is, if we have, for instance, one thousand contracts which give the same utility values for the agents, they would be represented in the diagram as a single point. This multiplicity, however, may be highly significant in negotiation. For instance, the scenario utility diagram in Figure 5a may correspond to a negotiation scenario where 80% of the contracts are within the Pareto front, or to an scenario where only 1% of the contracts are, both scenarios being quite different in terms of potential impact on protocol performance. Taking this into account, we generalized scenario utility diagrams into scenario utility *histograms* where each point in the diagram is assigned a number, representing either the number of contracts in the domain which yield the corresponding utility values to the agents or the ratio of contracts relative to the volume of the domain. The former can be used only with discrete domains, while the latter can be used with both discrete and continuous domains. Figure 5b shows a utility histogram for two agents. For a number of agents  $n_a > 2$ , we can represent the utility histogram as a set of  $(n_a + 1)$ -dimensional vectors, where the first  $n_a$  dimensions represent the utility values  $u_i$  for every agent  $i$ , and the last dimension represents the number or volume  $v$  of contracts corresponding to these utility values. We can also use a functional notation  $H(\vec{u}) = v$  to refer to the different points in the histogram.

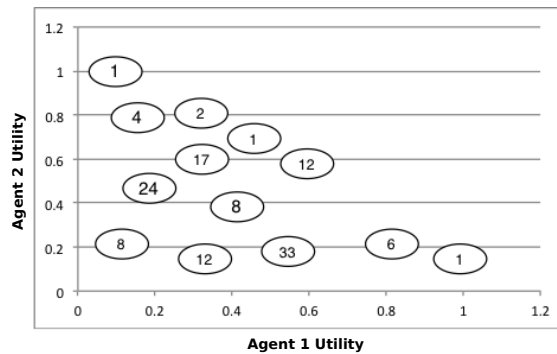


Figure 5b. An example of a scenario utility *histogram*.

Specifying scenarios in terms of their utility histograms has two great advantages. One advantage is that it represents an excellent way, as we shall see, of manipulating important properties of the scenario. We can specify, for example, whether the negotiation is highly competitive (the Pareto front is flat or even concave) or not, what proportion of the contracts are close to the Pareto front (making a good contract easier to find), and so on. The other advantage of this approach is that the Pareto fronts for our scenarios are known up-front, making it easy to get an exact value for the optimality of the agreements negotiated for that scenario. Negotiation researchers typically start by defining utility functions for the individual agents, and then *estimating* the Pareto fronts using an optimization approach such as genetic algorithms. If the contract space is large, and there are preferential dependencies (which as we've noted create multi-optima utility functions) then the Pareto front estimate may be quite inaccurate. Our approach obviates this whole problem by deriving the agent utilities from the Pareto front itself, as opposed to the other way around.

Our scenario generator provides a simple way for users to specify scenario utility histograms, based on a small set of high-level parameters. The user first defines the Pareto front shape as a set of *weighted hyper-spheres*, each which we can define as follows:

$$\sum_{i=1}^{n_a} (u_i^{\alpha_i})^2 = 1$$

where  $n_a$  is the number of negotiating agents,  $u_i$  is the utility obtained by agent  $i$  and  $\alpha_i$  is a weighting parameter which allows to control the shape of the Pareto front. In particular, as seen in Figure 6 (a) to (c), giving the same weights to all agent (that is,  $\alpha_i = \alpha_j \forall i, j$ ) allows to control the concavity or convexity of the Pareto front, and thus the competitiveness of the scenario (see Section 4). Note that we can theoretically model all ranges of competitiveness, including the zero-sum game ( $\alpha_i = 0.5 \forall i$ ).




Figure 6: Pareto fronts generated as 2-dimension weighted hyper-spheres.

In a similar way, as shown in Figure 6 (d), giving different weights to different agents allows us to adjust the symmetry of the Pareto front, thus varying the symmetry of the scenario. In particular, the different weights given to the utilities of the agents may be seen as modeling their bargaining power, as suggested in (Ley, 2006). More complex Pareto fronts may be obtained by aggregating different weighted hyper-spheres, though in these cases the absence of non-dominant points in the resulting front cannot be assured, so residual points need to be removed after generation. Figure 7 shows an example of such a Pareto front. The figures show continuous Pareto Fronts, though our generator uses discrete utility histograms for scenario generation. To discretize the Pareto fronts, the points have been generated uniformly throughout the hyper-sphere, using a generalization of the method in (Lai et al., 2004)




Figure 7: A complex Pareto front generated using two weighted hyper-spheres.

Once we have generated the Pareto fronts, we need to generate the rest of the points in the scenario utility histogram. To do this, we apply different probability distributions to a generalization of the Marsaglia algorithm (Marsaglia, 1972). Figure 8 shows the results of using a uniform and two normal distributions (one centered at the origin and the other at the Pareto front). We can see that varying the distribution used allows us, for instance, to study the effect of different utility or social welfare centroids on the negotiations, or the impact of having a high or low density of points near the Pareto front.




Figure 8: Scenario utility histograms for the same Pareto front, using probability distributions that are (a) uniform, (b) favor points away from the Pareto front, and (c) favor points near the Pareto front.

### 3.2. Specifying Agent Utility Functions

Once the utility histogram has been defined for a scenario, our scenario generator can then define the utility function for the individual agents. Our algorithm does so using *shared hypervolumes*. The idea is to include *similar* hypervolumes (that is, hypervolumes involving the same contracts but with different shapes) in the utility functions of the different agents, adjusting the parameters of the hypervolumes so that they generate appropriate points within the scenario utility histogram. For instance, if we want to generate utility functions for a trivial scenario utility histogram  $H(\bar{u})$  for two agents, where the histogram value is  $v$  for  $\bar{u} = [a, b]$  and 0 otherwise, we could achieve this by generating two utility functions which share a hypercube of volume  $v$ , with weight  $a$  for the first agent and weight  $b$  for the second. Figure 9 shows another simple example for two agents and two integer issues whose values may take values ranging from 1 to 10. In this case, the utility histogram (shown in Figure 9a) includes four points, as follows:

- Two pareto-dominant points:  $H([1,0.6]) = 16$  and  $H([0.5,1]) = 18$ .
- A non-dominant, nonzero utility point:  $H([0,25,0.5]) = 29$ .
- A point representing the rest of the contracts, which yield zero utility for both agents:  $H([0,0]) = 37$ .

The generation algorithm must find a set of hypervolumes which satisfy the volume values imposed by the diagram. In our example, this can be achieved using the three hypervolumes shown in Figure 9b and 9c. Hypervolumes  $C1$  and  $C2$  create an intersection  $C1 \cap C2$  of 3 by 6 (18 contracts), while the union  $C1 \cup C2$  encompasses 63 contracts, leaving 37 contracts unassigned (that is, with zero utility). Finally  $C3$  has 4 by 4 (16) contracts, and the area within  $C1 \cup C2$  which is not inside  $C1 \cap C2$  or inside  $C3$  encompasses a total of 29 contracts. Once the hypervolumes have

been found, the algorithm assigns the appropriate weights to the hypervolumes in the utility function of every agent so that the utility values in the histogram are satisfied. In our example, this is achieved with the weights shown in Figure 9b and Figure 9c. The resulting utility functions are shown together in Figure 9d.

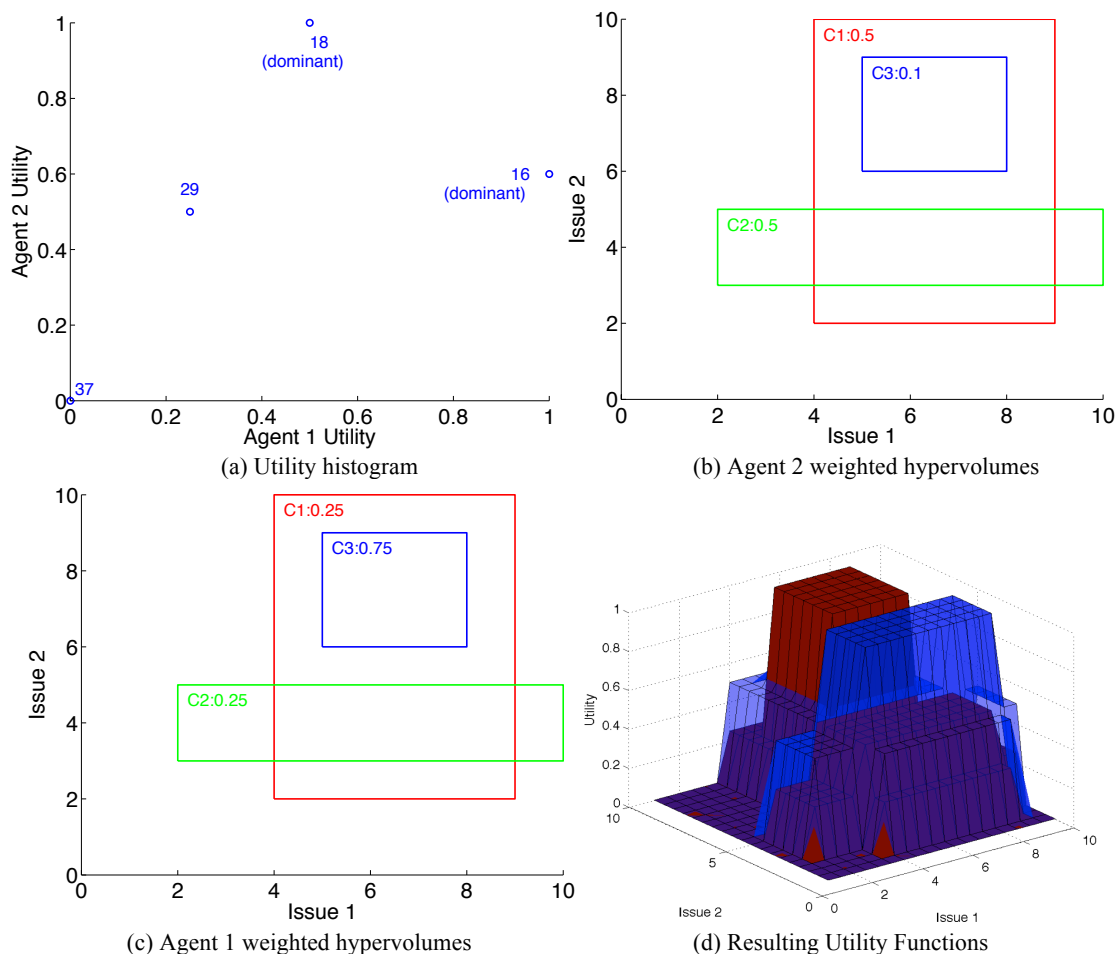


Figure 9: Example of scenario generation from utility histograms.


Of course, as the number of points in the utility diagram increase, the complexity of the generation process also increases, since we have to take into account the effect of the intersection between shared hypervolumes. A first approximation to avoid this problem is to divide the utility space in non-overlapping regions and to assign shared hypervolumes (one for each agent) to each region, with the utility values needed to match the points in the scenario utility diagram. The problem here is that, depending on the shape of the hypervolumes we are using, the resulting scenario utility histogram may differ from the original (e.g. a bell-shaped hypervolume adds more points to the histogram apart from the one corresponding to its peak value). A solution to this problem may be to feed this first approximation of the scenario utility histogram to a nonlinear optimizer which tries to minimize the approximation error. However, the computational cost of running that kind of optimizer for large cardinality contract spaces makes it unfeasible. Other alternatives are to use hypervolumes that do not introduce approximation error (i.e. hypervolumes with a constant weight), or simply accept that there will be some approximation error. As long as the hypervolumes do not overlap, the approximation error will never affect the shape of the Pareto front.

A second downside of the non-overlapping hypervolumes approach is that the utility function representation is not very compact, since it will have at least one hypervolume per point in the utility histogram. In some cases, this can be solved by using compression techniques to convert from non-overlapping hypervolumes to overlapping ones (this works relatively well with hypercubes). However, since the purpose of the generator is not to allow for compact utility representation, but to provide a consistent comparison of negotiation scenarios, we believe succinctness of the


representation is not an issue at this point. Furthermore, the use of shared, non-overlapping hypervolumes has some advantages for the computation of scenario and utility function metrics, so this is the approach we finally adopted. This simplified algorithm works as follows:

1. For each utility histogram point  $H(\bar{u}) = v$ , where  $\bar{u} = \{u_i | i = 1 \dots n_a\}$  and at least one  $u_i$  is nonzero, we partition the associated volume  $v$  into a set of region volumes  $v_j$  such that  $\sum_j v_j = v$ , according to the desired structural properties (see Section 4). For instance, if we want utility functions with low correlation length, the partition will involve small region volumes.
2. We map each region volume  $v_j$  to a particular hypercube region in the contract space which has the appropriate volume and does not overlap any previously allocated region. In this way, we guarantee our hypervolumes do not intersect. Again, the shape of the hypercube may be altered to satisfy desired structural properties. For instance, if we want asymmetric epistatic interactions, the hypercubes should be wider in some dimensions than in others.
3. We include the generated hypercube in the utility function of each agent  $i$ , assigning to it a weight equal to the corresponding utility value associated to that agent in  $\bar{u}$ ,  $u_j$ .
4. The algorithm progresses until all volume has been allocated.


An important property of this scenario generation strategy is that the shape and parameters of the shared hypervolumes may be varied so that additional properties of the generated functions are satisfied. Figure 10 shows an example for two agents and weighted hypercubes, where we have generated two scenarios with identical utility histograms and different correlation lengths. The type of hypervolumes or the aggregation operators used may be adjusted as well, but in these cases we have to take into account the effect that changes may have on the resulting scenario utility histogram.



(a)  $\psi_{0.7} = 5.9$



(b) scenario utility histogram



(c)  $\psi_{0.7} = 1.7$

Figure 10: Utility functions from the same utility histogram but with different correlation lengths.

## 4. Metrics Calculation

In this section we discuss the metrics we have developed to capture the “key properties” of negotiation scenarios, i.e. those that help predict which protocol(s) will work best for which scenarios. Our metrics are divided into three categories: (1) *structural* metrics, which capture properties of the individual agents’ utility functions, (2) *relational* metrics, which capture how the negotiating agents’ utilities relate to each other, and (3) *contextual* metrics, which capture properties of the context in which the negotiations occurs. We describe these metrics in the sections below.

### 4.1. Structural Metrics

Our structural metrics attempt to capture how difficult it is to find high-utility contracts in a given agent’s utility function. Many such metrics have been defined in such fields as artificial intelligence, operations research, and evolutionary computing (Vassilev et al., 2003).

Perhaps the simplest metric is based on the skew of agent’s individual utility histograms. Consider two histograms in Figure 11, giving the number of contracts at each utility value within a given agent’s utility function.




Figure 11. Agent utility histograms with differing skews.

The utility function whose histogram is shown on the left is probably harder to optimize because most of its contracts have a relatively low utility, while the one of the right is probably easier to optimize.

Another metric is *modality variance*, which measures the variance in the heights of the local optima in the utility function. In Figure 12, for example, the modality variance is low in the utility function on the left, and higher in the utility function on the right:






Figure 12. Agent utility functions with different modality variances.

The intuition behind this metric is that, if the modality variance is zero, an agent can simply hill-climb from any random point to find an optimal contract, since every local optimum is also globally optimal. As the modality variance increases, agents are increasingly likely to get stuck in poor local optima.

Our *ruggedness* metric assesses how “bumpy” a utility function is, and therefore how many local optima a search algorithm is likely to encounter when looking for the global optimum. Ruggedness can be assessed using *fitness distance correlation* (FDC) (Tomassini et al., 2005) (Jones and Forrest, 1995), which measures the correlation between the utility of a contract and its Euclidean distance, in the contract space, from the global optimum. If there is a strong correlation between distance and utility, this implies a smooth utility function. If the correlation is weak, this implies a rugged utility function with many “bumps” along the way. Consider the examples in Figure 13:





(a) FDC = -1.0

(b) FDC = -0.61

Figure 13. Agent utility functions with different FDC values.

The top charts in the figure show two different utility functions in a single-issue contract domain, and the bottom charts show utility as a function of a contracts' distance from the globally optimal contract. The left utility function is smooth, and shows a strong correlation between utility and distance (a high FDC). The right utility function is more rugged, and has a correspondingly lower FDC value. The lower the FDC (in absolute value), the more local optima exist and the more difficult it will probably be to find the global optimum. A related metric is the *correlation length* (AKA *correlation distance*) which is defined as the minimum contract space distance which makes the utility correlation between contracts pairs fall below a given threshold (usually 0.5) (Lopez-Carmona et al., 2011) (Marsa-Maestre et al., 2012).

Our *information content* metric takes a different tack, assessing the information content of the utility space based on the intuition that the higher the information content of a utility function, the more difficult it will be to perform a search in that utility space. The extreme cases would be a totally flat utility space (minimum information content, a single sample gives information about all the space) and a totally random utility space (maximum information content). Finding an optimal contract in a random utility space would of course be much harder, because you can't use slope information to guide your search at all. Information content can be assessed by measuring high-frequency components of the utility functions, for instance by using a n-dimensional Fast Fourier Transform (FFT) (Bracewell, 2000). For solution spaces which are too large to allow for comprehensive sampling, we can use *random walk* entropy measurements, where a random walk is performed throughout the utility space, and the entropy of the utilities of the contracts in the walk is measured.

Other metrics assess how effective optimizers are at locating global maximums in the utility functions. One example is the *hill-climbing success ratio*, which measures how often a hill-climbing optimizer can find optimal contracts within an agent's utility function. The *minimum annealing temperature* metric is a variant that assesses the minimum starting temperature needed for simulated annealing-based optimizer (Kirkpatrick et al., 1983) to achieve a given success ratio. The higher the temperature needed, the more rugged and challenging the utility function.

A final metric assesses the degree of preferential interdependency. Utility functions with preferential dependencies create multi-optimum topographies that are typically much harder to optimize than those involving independent issues (Klein et al., 2003). Preferential dependency can be assessed using an information-theoretic measure known as *epistasis* (Ventresca and Ombuki-Berman, 2007), originally developed in the context of genetic biology. It is defined as follows::

$$\epsilon_{ij} = \begin{cases} \frac{I(i;U)+I(j;U)}{I(i,j;U)} - 1 & \text{if } I(i,j;U) \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

where  $I(i;U)$  and  $I(j;U)$  are the amount of information that each issue  $i$  and  $j$  reveals about the value of the utility function  $U$ , and  $I(i,j;U)$  is the amount of information *both* issues reveal about the value of  $U$ . An alternative



technique to assess interdependency relationships between issues is to perform a Walsh analysis of the utility function (Heckendorn and Whitley, 1999).

Structural measures can be less meaningful in *anisotropic* utility functions where there are significant local variations in geography (e.g. where the utility function is smooth in some regions, and rugged in others) (Stadler and Gruner, 1993) (Greenwood and Hu, 1998). In this situation, we can calculate these measures separately for different regions of the space, using local optimum for example to delineate the centroids for the different regions.

## 4.2. Relational Metrics

We have also defined a set of metrics that assess *relationships* between the utility functions of the agents involved in a negotiation. Relational metrics capture important properties that are not obvious from examining any one agent's utility function. We may have, for instance, scenarios where agents have highly rugged utility functions that make it difficult for them to determine their own high utility regions, but where once these regions have been found agreements are fairly straightforward, because high utility regions for the different agents coincide. On the other hand, we may have scenarios where the agents have smooth easy-to-optimize utility functions, but where *mutually acceptable* regions are nevertheless hard to find.

It should be noted that relational metrics will be especially critical in domains with strong preferential dependencies. When there are no preferential dependencies, utility diagrams have simple symmetric structures, with convex Pareto fronts, and differ mainly in whether the agent utility functions are have strongly negative (zero-sum) correlations, strongly positively (win-win) correlations, or something in between (see Figure 14).




Figure 14. Typical utility diagrams for zero-sum, win-win, and uncorrelated scenarios with no preferential dependencies.

In this context, many of our relational metrics will do little to distinguish between negotiation scenarios. This picture changes completely, however, when we consider “nonlinear” utility functions with strong preferential dependencies. As we can see (figure 15), when preferential dependencies are introduced, the utility diagrams can become much more complex and diverse, with non-symmetric structures and concave Pareto fronts, and capturing key dimensions of variation between such scenarios requires a richer set of metrics.

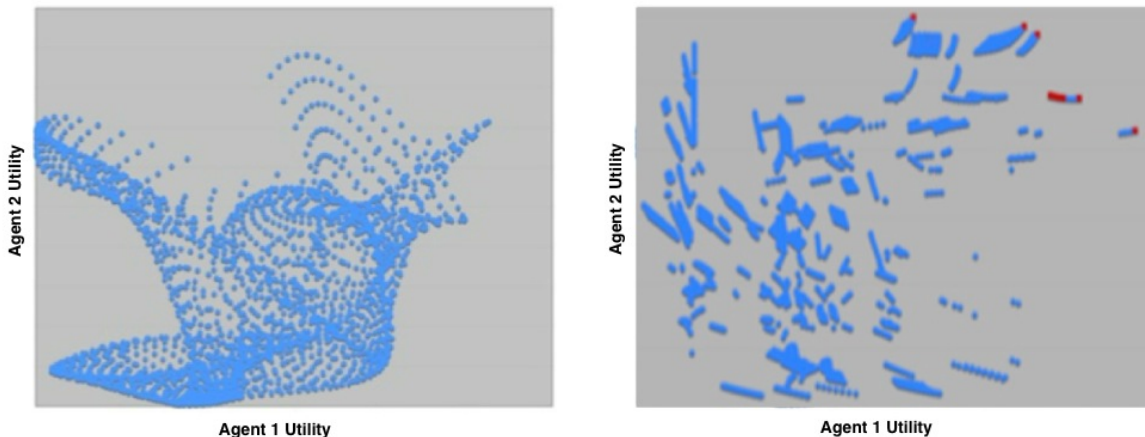


Figure 15. Utility diagrams for scenarios with strong preferential dependencies<sup>1</sup>.

Our *goodness* metrics assesses the proportion of high-utility agreements in the contract domain. This includes *the ratio of high-social-welfare contracts* (which measures what fraction of the contracts have top-quartile social welfare), the *ratio of mutually-acceptable contracts* (which measures what fraction of the contracts lie above the reservation values for the agents) and the *ratio of Pareto-efficient contracts* (which measures what fraction of the contracts are within a distance  $\epsilon$  of the pareto front). A scenario where these ratios are small will probably be more challenging for a negotiation protocol.

Our *symmetry* metric assesses the symmetry of the utility histogram across the equal utility line. Some scenarios may be inherently asymmetric, in the sense that there may be few or no contracts that offer equally high utility values for all the negotiating agents. Scenarios with a low symmetry metric will be more challenging if we want our protocol to find a fair agreement, where fairness can be defined for example as in (Fujita et al., 2010). We compute symmetry by adding up, for every agent, the distances of possible contracts from the equal utility line. The standard deviation of these sums gives an idea of how much an scenario favors some agents against others.

Our *competitiveness* metric captures the extent to which a scenario is win-win, or win-lose, for the agents involved. We calculate this by finding the average fairness for contracts that have high social welfare (and thus are more apt to include the final agreement). Intuitively, a scenario where most such contracts are highly unfair is more competitive than those where most such contracts are fair. A scenario with a highly convex Pareto front (where the highest social welfare contracts are very unfair) is, by this definition, more competitive than one with a highly concave Pareto front (where the highest social welfare contracts are good for all the agents involved).

Our *correlation* metric measures the correlation between the utility functions for different agents. Scenarios with high correlation values will presumably be easier for a protocol seeking win-win agreements.

Our *locality* metric measures the correlation between the distance of contracts in the solution and utility spaces, which is likely to have a strong impact on the suitability of similarity-based protocols.

Our *basin* metrics, finally, checks for the presence and size of *basins of mutual gain* (regions in the contract space where agents may find dominant directions for increasing mutual gain), which favors hill-climbing protocols.

Our relational metrics involve calculating some function of the utility values for all the contracts in the contract space. This can be computationally impractical, of course, if the contract domain is too large. In this case, we can calculate these values for a representative (e.g. random) sample of the contracts in the domain. It is also straightforward to calculate these metrics from the utility histogram for a scenario, so we can get exact metrics values for scenarios created by our scenario generator.

### 4.3. Contextual Metrics

This set of metrics is used to assess the potential impact of a negotiation scenario's *context*. We have focused, in our work so far, on metrics to assess *predictability*, i.e. how well agents can anticipate what will happen in their current negotiation based on the context of previous, similar, negotiations. The notion of negotiation predictability was first introduced in (Hindriks et al., 2007) and later extended in (Hindriks et al., 2011). Predictability is important because it helps negotiating agents make offers that will be appreciated by the opponent, as well as develop strategies that bring the highest utility to the negotiator. For the first we introduce the notion of preference profile predictability and for the latter we introduce bidding strategy predictability.

#### 4.3.1 Preference profile predictability

A preference profile reflects which issues are important to the negotiator, which values per issue are preferred over other values, whether or not there are dependencies between issues, and as a whole provides a (partial) ranking over all possible agreements. Being able to (partly) predict the preference profile of the other negotiators makes it easier to make an offer that the other party can accept, and increases the possibility of reaching a good negotiation outcome more quickly (Zeng and Sycara, 1997) (Faratin et al., 2000). The general question we help to answer is, given a

---

<sup>1</sup> In both cases, the utility functions were generated as the sum of a bell-shaped weighted hypervolumes.

negotiation domain and the negotiators, to what extent is the preference profile of a negotiator predictable? This question can be broken down into several sub-questions (Hindriks et al., 2007) that we can define as follows:

- *Role predictability.* Role predictability refers to the existence of typical negotiation roles. For example, in many domains (e.g., markets) there are sellers versus buyers, in job contract negotiations there are applicants and hiring managers. With a role comes information about the preference profile of the agent playing that role. For example, if the negotiation domain is about selling something, then there is typically an issue price, and it is predictable that sellers prefer higher prices, and buyer prefer lower ones.
- *Issue predictability.* An issue is called predictable if we can predict negotiators' preference profile (or some global properties thereof) for an issue. For example, in a negotiation for a supercomputer, the issue of the number of Teraflops and the issue of the number of Petabytes are predictable: the buyer would prefer more of both. When negotiating with friends about an upcoming holiday, typical issues are location, duration, start date, way of travelling, and price. Of these issues only price is predictable without knowing anything more about the people involved.
- *Issue weight predictability.* Issue weight predictability refers to predicting the importance of an issue in a negotiation, in absolute or relative terms. (Jonker and Robu, 2004) showed that knowing what the most important issue is to the opponent can produce significant improvements in the outcome of the negotiation.
- *Dependency predictability.* This refers to the question of whether it is reasonable to expect that a negotiating agent will have a dependency between some issues. For example, in a negotiation about a car, the price depends on most other issues, such as the age, the motor, size, and which accessories will be put in the car. Although a buyer would always prefer low price, he can adjust his reservation value for the price with respect to those issues. He may be willing to pay more for a new model while he would be reluctant to pay the same amount for an old model. Thus, predicting such interdependencies may help a negotiating agent generate appropriate offers. The same can be said for predicting preferential dependencies.
- *Domain predictability.* This refers to how many predictable issues there are in a given negotiation, and how much we can predict about the importance of those issues. Domain predictability is important in research into bidding strategies (Baarslag et al., 2010) (Hindriks et al., 2011).

Preference profile predictability, or one of the related notions of predictability can be established in various ways. Domains can be studied using conceptual analysis. For a negotiation on computing power, a conceptual analysis will reveal that the number of computing nodes and memory nodes has an effect on the electricity consumption and the cooling required for the nodes (which in turn requires additional electricity). Similarly, conceptual analysis of negotiations on a purchase reveals that role predictability (buyer-seller), issue predictability (price is preferred to be higher by the seller and lower by the buyer), and issue weight predictability (price is typically an important issue). Conceptual analysis can also reveal issue predictability for issues that are in arbitrary domains not predictable. For example, someone buying clothes for a wildlife observation tour will prefer colours that fit in with nature. So understanding the underlying concerns of the negotiation helps. In other cases observing your opponent during the negotiation, or just knowing your opponent better might help. For example, the colours the other negotiator picked for e.g., clothing or home decoration, might make the issue of colour predictable.

Statistical analysis and data mining are techniques that can be applied on a set of preferences of a multitude of people for a comparable issue or domain to extract a preference profile prediction (or part thereof). For example, although colour is in general not predictable, an analysis of the car domain shows that in 2007 (see <http://www.cars.com>), white was the most popular colour in North America, making up 19% of all vehicles. Silver, black, red and grey rounded out the top five.

Various machine learning techniques are used for opponent modelling, for example, to gain more information on the zone of agreement (Zeng and Sycara, 1997) or on the shape of the evaluation functions per issue (Hindriks et al., 2009).

The grounds for drawing a conclusion about preference profile predictability can be economic rationality, surveys over a large enough set of people, or an emergent pattern in a set of preference profiles for that issue or for that domain, or in the subsequent bids made by the opponent.

#### 4.3.2 Strategy predictability

Being able to predict the behavior of the other party enables anticipation on that behavior, see e.g., (Ji and Leung, 2010). The results of the Automated Negotiating Agent Competition (ANAC), see (Baarslag et al., 2010), show which bidding strategies perform best against which other bidding strategy. This knowledge will probably soon be used to create negotiation strategies that will adapt to the current strategy used by the opponent.

An opponent's strategy is called predictable, if on the basis of earlier encounters or a personality/experience profile, the opponents strategy can be characterized in terms of some global style, such as the negotiation styles of (Shell, 2006) (Thompson, 2011), or a known conflict-handling style, as the one shown in Figure 16 (Kilman and Thomas, 1977) (Pruitt, 1986).




Figure 16 – The Thomas-Kilman Conflict Mode Instrument

## 5. Protocol Testbed

The purpose of the protocol testbed is to allow researchers to assess how well their protocols work in different negotiation scenarios. Our testbed is called GENIUS (General Environment for Negotiation with Intelligent multi-purpose Usage Simulation) (Lin et al., 2012) (Baarslag et al., 2012). A user begins by specifying the experiment they want to run, selecting the scenario (domain and utility functions), protocol, and negotiating agents (Figure 17):

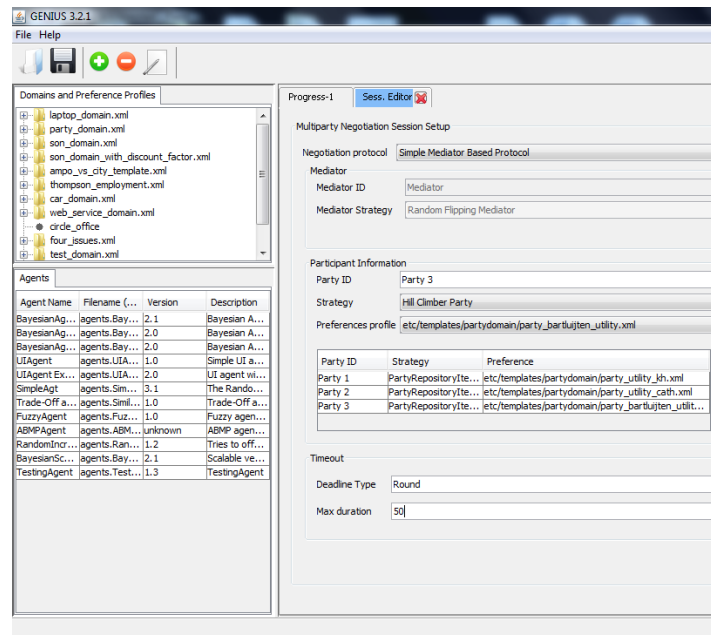


Figure 17: GENIUS: negotiation simulation setup interface

Once the setup is complete, GENIUS can simulate the negotiation, deriving such statistics as the current utility values for each agent, the product of the agent utilities, and so on (figure 18):

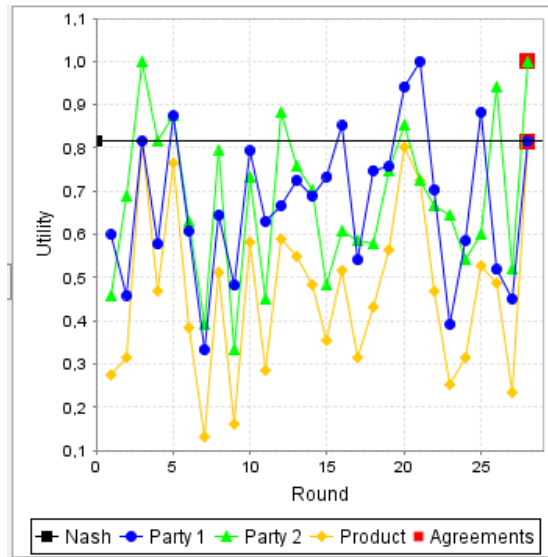


Figure 18.. Overview of a negotiation as it progresses over multiple rounds

The experiment results for different combinations of scenarios and protocols can then be compared with each other. GENIUS has focused, to date, on bilateral alternating offer protocols and linear additive utility functions, and in this form has been used to support the Automated Negotiating Agents Competitions (<http://mmi.tudelft.nl/anac>).

## 6. Repositories

The final component of the Negotiation Handbook are the repositories for scenarios as well as for experimental data concerning the performance of different protocols in these scenarios. At present we have two, complimentary, systems that serve this function. GENIUS, as noted above, includes repositories for scenarios and experiment data, but has focused to date on linear utility functions. The second system, a web-based tool called *Negowiki*, has focused on capturing scenarios and experimental results for *nonlinear* cases. Negowiki supports:

- *Downloading negotiation scenario.* Users can search the Negowiki scenario repository for scenarios with desired properties and download them in an XML format that instantiates the weighted hypervolume formalism described in section 2. The repository currently includes a range of scenarios drawn from the research literature, including (Faratin et al., 2000) (Klein et al., 2003) (Ito et al., 2008) (Lopez-Carmona et al., 2011). Users can search for scenarios based on their properties (figure 19):

| Source  | # Agents   | Domain   |
|---|--|--|
| <input checked="" type="checkbox"/> Ito et al.,2007 | <input type="checkbox"/> 2<br><input type="checkbox"/> 3<br><input type="checkbox"/> 4<br><input type="checkbox"/> 5 | <input type="checkbox"/> 2x[0..9]<br><input type="checkbox"/> 3x[0..9]<br><input type="checkbox"/> 4x[0..9]<br><input type="checkbox"/> 5x[0..9]<br><input type="checkbox"/> 6x[0..9]<br><input type="checkbox"/> 7x[0..9] |

These constraints are satisfied by 116 test scenario(s)

Download Scenarios

Figure 19. The Negowiki scenario selection interface.

Every scenario includes a description of the rationale behind it, as well as the values for the scenario metrics as described in section 4 (figure 20):

**Scenario: Ito et al 2007**  
[EDIT](#) [DOWNLOAD](#)

---

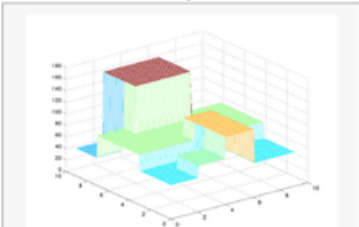
ID E-1M5UXGG-27  
 Agents 10  
 Issues 10  
 Domain 10  
 Description

---

**Category: Ito et al., 2007**

This scenario involves the kind of utility functions used in the paper by Ito et al *Multi-issue Negotiation Protocol for Agents: Exploring Nonlinear Utility Spaces*, published in *IJCAI 2007*<sup>[1]</sup>. The paper used a *weighted\_sum of plain utility functions* called constraints. The utility functions were generated as follows:

- Number of issues: 2 to 10.
- Domain for issue values: integers from 0 to 9.
- Constraint distribution: 5 unary constraints, 5 binary constraints, 5 trinary constraints, etc. (a unary constraint relates to one issue, an binary constraint relates to two issues, and so on).
- Maximum utility value for a constraint :  $100 \times (\text{Number of Issues})$ . Constraints that satisfy many issues thus have, on average, larger weights. This seems reasonable for many domains. In meeting scheduling, for example, higher order



Ito et al., 2007

Weighted sum of random plain utility functions

**Summary**

Weighted sum of random plain utility functions drawn from a uniform distribution

**Key features**

|                         |  |
|-------------------------|--|
| <b>Number of issues</b> | 2,4,6,8,10                                       |
| <b>Domain of issues</b> | $(x_1, x_2)   x_i \in [0, 9] \subset \mathbb{R}$ |

Figure 20. An example of a Negowiki scenario description.

Users can also download function libraries (currently available in Matlab and Common Lisp) that calculate the utility of a given contract for a given scenario, making it straightforward for a research group to use the scenarios in their protocol evaluation experiments.

- *Uploading negotiation scenario for analysis.* Users can upload their scenarios to the Negowiki repository using the scenario formalism described in section 2. When uploaded, the website computes and displays the metrics values for that scenario. Users can archive these scenarios into the repository so they are available for use by other researchers.
- *Uploading negotiation results data for analysis.* If users perform experiments with protocols from the literature or with their own protocols using any of the Negowiki scenarios, they can upload their negotiation results into Negowiki for analysis. When uploaded, the website computes a range of negotiation outcome metrics (such as individual agent utility, social welfare, fairness, distance from the Pareto front) and shows them to the user. The Negowiki team has devoted substantial effort to ensuring that the site calculates accurate outcome metrics for negotiation experiment results. It is important to note that calculating some outcome measures can be very different in nonlinear scenarios. Traditional solution concepts such as the Nash product, for example, can become very misleading in such contexts (Fujita et al., 2010). Users can archive their experimental results in the repository to make them available to the community, and can search this repository for insights into which best protocols are best for solving the negotiation problem at hand, as well as for inspiration on how to design better protocols.

For more information about Negowiki, go to <http://negowiki.mit.edu/>.

## 7. Contributions and Future Work

One of the key challenges in automated negotiation research is that it has been difficult to come up with a systematic picture of which negotiation protocols work best for which problems, due to the diverging scenarios considered by the different research groups. In this paper we have described an infrastructure that aims to fill this gap. This infrastructure provides:

- A tool for *generating negotiation scenarios*, based on the aggregation of hypervolumes to generate utility functions, and on the use of shared hypervolumes and nonlinear regression to generate negotiation scenarios from utility histograms. To our knowledge, despite its obvious advantages, no authors have previously tried to generate negotiation scenarios from utility diagrams.
- A set of *metrics* to measure high-level scenario parameters, taking into account both the structural properties of the agent utility functions, the relationships between the utility functions of the different agents, and the context in which the negotiation occurs.
- A *testbed* that allows researchers to evaluate how well their protocols work on different scenarios.
- *Community repositories* that allow the research community to share scenarios and experiment results and thereby carry on a much more cumulative and systematic exploration of the critical question of which protocols work best for which negotiation problems.

While our approach has been designed to be as general as possible, it does face one inherent limitation. Our scenario representation will not work for domains where it is impractical to capture the entire utility function up front. Imagine, for example, that agents are negotiating over the placement of pieces on a chess board. There is an astronomical number of possible “contracts” (i.e. chess board configurations). It would require an enormous amount of time to calculate the utility values for every possible contract, and an impractically large number of hypervolumes to capture this information. To handle this case, we would need to adopt a scenario representation which represents the *algorithms used to calculate the utility functions* (so contract utilities can be calculated as needed) rather than the utility functions themselves. This is a possible direction for future work.

Our more immediate efforts will involve extending our infrastructure. This will include work on:

- *Scenario generation*: we plan to extend the scenario generator to allow more comprehensive control over other potentially important scenario properties, such as epistasis, fractal statistics in the utility functions (Da Cruz, 1999), and so on.
- *Metrics*: our current set of scenario metrics, while promising, are preliminary and will no doubt have to be refined and extended as we develop a better understanding of which scenario properties most deeply impact the performance of different protocols.
- *Testbed*: currently, the GENIUS testbed environment only supports linear additive utility functions and bilateral protocols. We are extending the GENIUS testbed to support our weighted-hypervolume scenario representation and to enable multilateral negotiations.
- *Repositories*: We will continue to populate the scenario repository with examples from the research literature and from our scenario generator. A critical part of this will be developing a *scenario taxonomy*, i.e. a systematic enumeration of all the important distinct classes of negotiation scenarios, so we can try to assure that the scenario repository has a comprehensive collection of test cases.
- *Experiments*: we will perform comprehensive comparisons of negotiation performance for different protocol and scenario combinations. We are creating a taxonomic representation of known negotiation protocols, based on ideas developed in the Process Handbook project (Klein et al. 2006), which will allow us to systematically explore the space of protocol variants by a recombination process. Since the space of possible scenarios and protocols is very large, so we will need to develop heuristics for focusing our search on scenario-protocol pairs that are apt to be important for real-world challenges.
- *Data mining*: we will explore the experiment data repositories generated by these experiments to develop keener insights into the determinants and patterns of protocol performance.
- *Community development*: we will work to enable a vibrant community process, building on the infrastructure we have developed, wherein researchers share scenarios, protocols, and experiment results, all in the service of creating a Negotiation Handbook that can substantially impact how well humanity is able to solve its complex and critical negotiation challenges.

## 8. Acknowledgments

This work has been supported by: the ITEA-2 project 2008005, "Do-it-Yourself Smart Experiences"; the Spanish Ministry of Education and Science grants TIN2008-06739-C04-04 (project T2C2) and JC-2010-035 (mobility grant); the Dutch Technology Foundation STW, applied science division of NWO and the Technology Program of the Ministry of Economic Affairs; the Pocket Negotiator project with grant number VICI-project 08075; the US National Science Foundation; the New Governance Models for Next Generation Infrastructures project, NGI grant number 04.17.

## 9. References

- Athey, S. & Segal, I. Designing efficient mechanisms for dynamic bilateral trading games. *American Economic Review* 97, 2 (2007), 131–136.
- Aydogan, R., Baarslag, T., Hindriks, K. V., Jonker, C. M., & Yolum, P. (2011). *Heuristic-based Approaches for CP-Nets in Negotiation*. Proceedings of The Fourth International Workshop on Agent-based Complex Automated Negotiations (ACAN 2011), pp. 56 - 59.
- Baarslag, T., Fujita, K., Gerding, E. H., Hindriks, K. V., Ito, T., Jennings, N. R. et al. (2012). Evaluation Methods for Negotiating Agents: Results and Analysis of ANAC 2011. *Artificial Intelligence*, submitted).
- Baarslag, T., Hindriks, K., Jonker, C. M., Kraus, S., & Lin, R. (2010). *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*. Proceedings from The First Automated Negotiating Agents Competition (ANAC 2010).
- Boutilier, C., Bacchus, F., & Brafman, R. I. (2001). *UCP-Networks: A directed graphical representation of conditional utilities*. Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, pp 56–64.
- Boutilier, C., & Hoos, H. H. (2001). *Bidding languages for combinatorial auctions*. Proceedings of the International Joint Conference On Artificial Intelligence, pp. 1211–1217.
- Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., & Poole, D. (2004). CP-nets: a tool for representing and reasoning with conditional ceteris paribus preference statements. *Artificial Intelligence Research*, 21, 135–191.
- Bracewell, R. N. (2000). *The fourier transform & its applications* 3rd Ed. McGraw-Hill Science/Engineering/Math.
- Brafman, R. I., & Domshlak, C. (2009). Preference Handling --- An Introductory Tutorial". *AI Magazine*, 30(1), 58–86.
- Clayton, N., Biddle, R., Tempero, E. (2000). *A study of usability of Web-based software repositories*. Proceedings of the International Conference on Software Methods and Tools (SMT 2000), vol., no., pp.51-58.
- Cramton, P., Shoham, Y., & Steinberg, R. (2006). *Introduction to combinatorial auctions*. MIT Press.
- Da Cruz, W. (1999). Fractons and fractal statistics. *Arxiv preprint hep-th/9905229*.
- Eckel, C. C., Grossman, P. J. (2008), *Forecasting risk attitudes: An experimental study using actual and forecast gamble choices*, *Journal of Economic Behavior & Organization*, Volume 68, Issue 1, October 2008, Pages 1-17.
- Ellingsen, T., Johanneson, M., and Lilja, J. (2009). *Trust and Truth*. *Economic Journal* 119 (534):252–276.
- Faratin, P., Sierra, C., & Jennings, N. R. (2000). Using Similarity Criteria to Make Negotiation Trade-Offs. In (pp. 119–126). Washington, DC, USA: IEEE Computer Society.
- Fatima, S., Wooldridge, M., & Jennings, N. R. (2009). *An analysis of feasible solutions for multi-issue negotiation involving nonlinear utility functions*. Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems, Budapest, Hungary, pp 1041–1048.
- Fenghui, R., Minjie, Z., & Kwang, M. S. (2009). Adaptive conceding strategies for automated trading agents in dynamic, open markets. *Decision Support Systems*, 46(3), 704 - 716.
- Fujita, K., Ito, T., & Klein, M. (2010). A Secure and Fair Protocol that Addresses Weaknesses of the Nash Bargaining Solution in Nonlinear Negotiation. *Group Decision and Negotiation*, 1-19.
- Grabisch, M. (1997). k-order additive discrete fuzzy measures and their representation. *Fuzzy Sets Syst.*, 92(2), 167–189.
- Greenwood, G. W., & Hu, X. (1998). On the use of random walks to estimate correlation in fitness landscapes. *Computational Statistics Data Analysis*, 28, 131–137.
- Heckendorn, R. B., & Whitley, D. (1999). Predicting epistasis from mathematical models. *Evolutionary Computing*, 7, 69–101.



- Hemaissia, M., El, F. S., Amal, Labreuche, C., & Mattioli, J. (2007). *A multilateral multi-issue negotiation protocol*. Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems, New York, NY, USA.
- Hindriks, K., Jonker, C., & Tykhonov, D. (2007). Analysis of Negotiation Dynamics. In M. Klusch, K. Hindriks, M. Papazoglou, & L. Sterling (Eds.), 4676 (pp. 27-35). Delft University of Technology, Man-Machine Interaction, Mekelweg 4, 2628CD Delft The Netherlands: Springer Berlin / Heidelberg.
- Hindriks, K., Jonker, C. M., & Tykhonov, D. (2009). *The Benefits of Opponent Models in Negotiation*. Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, Washington, DC, USA.
- Hindriks, K., Jonker, C. M., & Tykhonov, D. (2011). Let's dans! An analytic framework of negotiation dynamics and strategies. *Web Intelligence and Agent Systems*, 9, 319–335.
- Ito, T., Klein, M., & Hattori, H. (2008). A Multi-Issue Negotiation Protocol among Agents with Nonlinear Utility Functions. *Journal of Multiagent and Grid Systems*, 4(1), 67-83.
- Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Sierra, C., & Wooldridge, M. (2001). Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2), 199-215.
- Ji, S.-j., & Leung, H.-f. (2010). *An Adaptive Prediction-Regret Driven Strategy for Bilateral Bargaining*. Proceedings of the 2010 22nd IEEE International Conference on Tools with Artificial Intelligence, Washington, DC, USA.
- Jones, T., & Forrest, S. (1995). *Fitness distance correlation as a measure of problem difficulty for genetic algorithms*. Proceedings of the 6th international conference on genetic algorithms.
- Jonker, C., & Robu, V. (2004). *Automated Multi-Attribute Negotiation with Efficient Use of Incomplete Preference Information*. Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, Washington, DC, USA New York, New York.
- Kilmann, R. H., & Thomas, K. W. (1977). Developing a Forced-Choice Measure of Conflict-Handling Behavior: The MODE Instrument. *Educational and Psychological Measurement*, 37(2), 309–325.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220, 671-680.
- Klein, M., Faratin, P., Sayama, H., & Bar-Yam, Y. (2003). Negotiating Complex Contracts. *Group Decision and Negotiation*, 12(2), 111 - 125.
- Klein, M., & Petti, C. (2006). A Handbook-Based Methodology for Redesigning Business Processes. *Knowledge and Process Management*, 13(2), 108-119.
- Kraus, S. (2006). Automated Negotiation and Decision Making in Multiagent Environments. In M. Luck, V. Marik, S. Olga, & R. Trappl (Eds.), 2086 (pp. 150-172). Bar-Ilan University Dept. of Mathematics and Computer Science Ramat-Gan 52900 Israel: Springer Berlin / Heidelberg.
- Lai, G., Li, C., Sycara, K., & Giampapa, J. (2004). Literature Review on Multiattribute Negotiations. CMU-RI-TR-04-66).
- Lau, R. Y. K., Li, Y., Song, D., & Kwok, R. C. W. (2008). Knowledge discovery for adaptive negotiation agents in e-marketplaces. *Decision Support Systems*, 45(2), 310 - 323.
- Ley, E. (2006). Statistical inference as a bargaining game. *Economics Letters*, 93(1), 142-149.
- Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K. V., & Jonker, C. M. (2012). Genius: An Integrated Environment for Supporting the Design of Generic Automated Negotiators. *Computational Intelligence*, In press).
- Lopez-Carmona, M. A., Marsa-Maestre, I., Hoz, E. d. l., & Velasco, J. R. (2011). A Region-Based Multi-Issue Negotiation Protocol for Non-Monotonic Utility Spaces. *Computational Intelligence*.
- Marsa-Maestre, I., Lopez-Carmona, M. A., Klein, M., Ito, T., & Katsuhide, F. (2012). Addressing Utility Space Complexity in Negotiations Involving Highly Uncorrelated, Constraint-based Utility Spaces. *Computational Intelligence*, In press.
- Marsaglia, G. (1972). Choosing a Point from the Surface of a Sphere. *Ann. Math. Stat.* 43 (2): 645–646.
- Myerson, R., & Satterthwaite, M. Efficient mechanism for bilateral trading. *Journal of Economic Theory* 28 (1983), 265–281.
- Parsons, S., Rodriguez-Aguilar, J. A., & Klein, M. (2011). Auctions and bidding: A guide for computer scientists. *Communications of the ACM*, 43(2).
- Patton, C., Balakrishnan, P. V. (2010). *The impact of expectation of future negotiation interaction on bargaining processes and outcomes*, *Journal of Business Research*, Volume 63, Issue 8, Pages 809-816
- Pruitt, D. G. (1986). Trends in the scientific study of negotiation and mediation. *Negotiation Journal*, 2, 237-244.

- Qing, F. & Ruhua, H. (2008) Evaluating the usability of discipline repositories. Proceedings of the IEEE International Symposium on IT in Medicine and Education, (ITME 2008), vol., no., pp.385-390, 12-14 Dec. 2008
- Raiffa, H. (1982). *The Art and Science of Negotiation*. Cambridge, USA: Harvard University Press.
- Robu, V., Somefun, D. J. A., & La Poutré, J. A. (2005). *Modeling Complex Multi-Issue Negotiations Using Utility Graphs*. Proceedings from AAMAS'05, Utrecht, The Netherlands.
- Rosenschein, J. S., & Zlotkin, G. (1994). *Rules of encounter : designing conventions for automated negotiation among computers*. Cambridge, Mass.: MIT Press.
- Sandholm, T. W. (1998). Distributed Rational Decision Making. In G. Weiss (Ed.), *Multi-Agent Systems*.
- Shell, G. R. (2006). *Bargaining for advantage: negotiation strategies for reasonable people*. Penguin Books.
- Sim, K. M., and Shi, B. Concurrent negotiation and coordination for grid resource co-allocation. *IEEE Trans. on Systems, Man and Cybernetics, Part B* 40, 2 (2009).
- Smith, R. G. The contract net protocol: high-level communication and control in a distributed problem solver. *IEEE Transactions on Computers* C-29, 12 (August 1980), 1104–1113.
- Stadler, P. F., & Gruner, W. (1993). Anisotropy in Fitness Landscapes. *Journal of Theoretical Biology*, 165(3), 373 - 388.
- Sycara, K. P. (1993). Machine learning for intelligent support of conflict resolution. *Decision Support Systems*, 10(2), 121 - 136.
- Thompson, L. L. (2011). *The Mind and Heart of the Negotiator*. Pearson.
- Tomassini, M., Vanneschi, L., Collard, P., & Clergue, M. (2005). A Study of Fitness Distance Correlation as a Difficulty Measure in Genetic Programming. *Evol. Comput.*, 13(2), 213–239.
- Vassilev, V. K., Fogarty, T. C., & Miller, J. F. (2003). Smoothness, ruggedness and neutrality of fitness landscapes: from theory to application. In *Advances in evolutionary computing: theory and applications* (pp. 3–44). New York, NY, USA: Springer-Verlag New York, Inc.
- Ventresca, M., & Ombuki-Berman, B. (2007). *Epistasis in Multi-Objective Evolutionary Recurrent Neuro-Controllers*. Proceedings of the 2007 IEEE Symposium on Artificial Life (ALIFE '07).
- Wu, M., Weerdt, M., & Poutré, H. (2009). *Efficient Methods for Multi-agent Multi-issue Negotiation: Allocating Resources*. Proceedings of the 12th International Conference on Principles of Practice in Multi-Agent Systems (PRIMA '09), Berlin, Heidelberg Nagoya, Japan.
- Zeng, D., & Sycara, K. (1997). *Benefits of learning in negotiation*. Proceedings of the fourteenth national conference on artificial intelligence and ninth conference on Innovative applications of artificial intelligence (AAAI'97/IAAI'97), Providence, Rhode Island.