Aspectos sociales de los MAS:

Reputación y Credibilidad

Carles Sierra

IIIA CSIC

In collaboration with John Debenham
University of Technology, Sydney

Talk plan

- Information-based agency
- Information-based trust
- An inspiring Example for reputation
- Forming individual opinions
- Forming group opinions
- SNA modulating group opinions
- Conclusions and future work

Sharing experiences and opinions:

The route to Trust and reputation

$$
\begin{aligned}
u &::= \text{inform}(agent, agent, content, time) \\
content &::= \text{opinion}(agent, agent, [term,](eval)) \mid \\
& \quad \text{experience}(agent, agent, term, term) \\
term &::= \varphi | \phi | \ldots (*expression\ from\ ontology\ O*) \\
eval &::= e = p \mid e = p, eval \\
e &::= good \mid bad \mid \ldots (*qualitative\ term*) \\
p &::= a\ point\ in\ [0, 1] \\
time &::= a\ point\ in\ time \\
agent &::= \alpha \mid \beta \mid \ldots (*agent\ identifiers*)
\end{aligned}
$$

Examples

$$\text{inform}(John, me, \text{opinion}(John, Carles, wrapping(package), (ghastly = 0.7)), t)$$

$$\text{inform}(John, me, \text{opinion}(Carles, John, suggesting(wine(Margaret\ River)), (excellent = 0.9)), t)$$

$$\text{inform}(John, me, \text{experience}(John, Carles, package(date(Monday)), package(date(Friday)), t)$$

$$\text{inform}(John, me, \text{experience}(John, Carles, fly(elephant), \neg fly(elephant)), t)$$

Information-based agency

Agent α receives all messages expressed in $\mathcal{C}$ in an in-box $\mathcal{X}$ where they are time-stamped and sourced-stamped.

A message μ from agent β (or θ or ξ) is then moved from $\mathcal{X}$ to a *percept repository* $\mathcal{Y}^t$ where it is appended with a subjective belief function $\mathbb{R}^t(\alpha, \beta, \mu)$ that normally decays with time. α acts in response to a message that expresses a *need*.

A need may be exogenous such as a need to trade profitably or may be triggered by another agent offering to trade, or endogenous such as α deciding that it owns more wine than it requires.

Each plan contains constructors for a *world model* $\mathcal{M}^t$ that consists of probability distributions, (X_i), in first-order probabilistic logic $\mathcal{L}$. $\mathcal{M}^t$ is then maintained from percepts received using *update functions* that transform percepts into constraints on $\mathcal{M}^t$

Integrity Decay

α may have background knowledge concerning the expected integrity of a percept as $t \to \infty$ — the *decay limit distribution*.

Given a distribution, $\mathbb{P}(X_i)$, and a decay limit distribution $\mathbb{D}(X_i)$, $\mathbb{P}(X_i)$ decays by:

$$\mathbb{P}^{t+1}(X_i) = \Delta_i(\mathbb{D}(X_i), \mathbb{P}^t(X_i))$$

where Δ_i is the *decay function* for the X_i satisfying the property that $\lim_{t\to\infty} \mathbb{P}^t(X_i) = \mathbb{D}(X_i)$. For example, Δ_i could be linear: $\mathbb{P}^{t+1}(X_i) = (1 - \nu_i) \times \mathbb{D}(X_i) + \nu_i \times \mathbb{P}^t(X_i)$, where $\nu_i < 1$ is the decay rate for the i'th distribution.

Either the decay function or the decay limit distribution could also be a function of time: Δ_i^t and $\mathbb{D}^t(X_i)$.

Reactive Reasoning

This procedure updates $\mathcal{M}^t$ for all percepts expressed in $\mathcal{C}$.

Suppose that α receives a message μ from agent β at time t. Suppose that this message states that something is so with probability z, and suppose that α attaches an epistemic belief $\mathbb{R}^t(\alpha, \beta, \mu)$ to μ — this probability reflects α's level of personal *caution*. Each of α's active plans, s, contains constructors for a set of distributions $\{X_i\} \in \mathcal{M}^t$ together with associated *update functions*, $J_s(\cdot)$, such that $J_s^{X_i}(\mu)$ is a set of linear constraints on the posterior distribution for X_i.

Denote the prior distribution $\mathbb{P}^t(X_i)$ by $\vec{p}$, and let $\vec{p}_{(\mu)}$ be the distribution with minimum relative entropy with respect to $\vec{p}$: $\vec{p}_{(\mu)} = \arg\min_{\vec{r}} \sum_j r_j \log \frac{r_j}{p_j}$ that satisfies the constraints $J_s^{X_i}(\mu)$.

Reactive Reasoning /contd

Then let $\vec{q}_{(\mu)}$ be the distribution:

$$\vec{q}_{(\mu)} = \mathbb{R}^t(\alpha, \beta, \mu) \times \vec{p}_{(\mu)} + (1 - \mathbb{R}^t(\alpha, \beta, \mu)) \times \vec{p}$$

and then let:

$$X_{i(\mu)} = \begin{cases} \vec{q}_{(\mu)} & \text{if } \vec{q}_{(\mu)} \text{ is more interesting than } \vec{p} \\ \vec{p} & \text{otherwise} \end{cases}$$

A general measure of whether $\vec{q}_{(\mu)}$ is more interesting than $\vec{p}$ is: $\mathbb{K}(\vec{q}_{(\mu)} \| \mathbb{D}(X_i)) > \mathbb{K}(\vec{p} \| \mathbb{D}(X_i))$, where $\mathbb{K}(\vec{x} \| \vec{y}) = \sum_j x_j \ln \frac{x_j}{y_j}$ is the Kullback-Leibler distance between two probability distributions $\vec{x}$ and $\vec{y}$.

Reactive Reasoning /contd /contd

Finally merging the above we obtain the method for updating a distribution X_i on receipt of a message μ:

$$\mathbb{P}^{t+1}(X_i) = \Delta_i(\mathbb{D}(X_i), \mathbb{P}^t(X_{i(\mu)}))$$

This procedure deals with

- integrity decay
- two probabilities:
  - the probability z in the percept μ that will appear in the constraints
  - the belief $\mathbb{R}^t(\alpha, \beta, \mu)$ that α attached to μ.

An Example

In a simple multi-issue contract negotiation α may estimate $\mathbb{P}^t(\mathrm{acc}(\beta, \alpha, \delta))$, the probability that β would accept δ, by observing β's responses.

Using shorthand notation, if β sends the message $\mathrm{Offer}(\delta_1)$ then α may derive the constraint: $J^{\mathrm{acc}(\beta,\alpha,\delta)}(\mathrm{Offer}(\delta_1)) = \{\mathbb{P}^t(\mathrm{acc}(\beta, \alpha, \delta_1)) = 1\}$, and if this is a counter offer to a former offer of α's, δ_0, then: $J^{\mathrm{acc}(\beta,\alpha,\delta)}(\mathrm{Offer}(\delta_1)) = \{\mathbb{P}^t(\mathrm{acc}(\beta, \alpha, \delta_0)) = 0\}$.

In the not-atypical special case of multi-issue bargaining where the agents' preferences over the individual issues *only* are known and are complementary to each other's, maximum entropy reasoning can be applied to estimate the probability that any multi-issue δ will be acceptable to β by enumerating the possible worlds that represent β's "limit of acceptability".

Empirical estimate of $\mathbb{R}^t(\alpha, \beta, \mu)$

Suppose that μ is received from agent β at time u and is verified by ξ as μ' at some later time t. Denote the prior $\mathbb{P}^u(X_i)$ by $\vec{p}$. Let $\vec{p}_{(\mu)}$ be the posterior minimum relative entropy distribution subject to the constraints $J_s^{X_i}(\mu)$, and let $\vec{p}_{(\mu')}$ be that distribution subject to $J_s^{X_i}(\mu')$.

The *observed reliability* for μ and distribution X_i:

$$\mathbb{R}^t_{X_i}(\alpha, \beta, \mu)|\mu' = \arg\min_k \mathbb{K}(k \cdot \vec{p}_{(\mu)} + (1-k) \cdot \vec{p} \parallel \vec{p}_{(\mu')})$$

If $\mathbf{X}(\mu)$ is the set of distributions that μ affects, then the *observed reliability* of β on the basis of the verification of μ with μ' is:

$$\mathbb{R}^t(\alpha, \beta, \mu)|\mu' = \frac{1}{|\mathbf{X}(\mu)|} \sum_i \mathbb{R}^t_{X_i}(\alpha, \beta, \mu)|\mu'$$

Commitment, Enactment, and Semantics

Denote $\mathbb{P}^t(\text{Observe}(\varphi')|\text{Commit}(\varphi))$ simply as $\mathbb{P}^t(\varphi'|\varphi) \in \mathcal{M}^t$ Set of possible enactments be $\Phi = \{\varphi_1, \varphi_2, \ldots, \varphi_m\}$ with prior distribution $\vec{p} = \mathbb{P}^t(\varphi'|\varphi)$. We estimate the posterior $\vec{p}_{(\mu)}$.

First, if $\mu = (\varphi_k, \varphi)$ is observed estimate the posterior $\vec{p}_{(\varphi_k)} = (p_{(\varphi_k)j})_{j=1}^m$ satisfying the single constraint: $J^{(\varphi'|\varphi)}(\varphi_k) = \{p_{(\varphi_k)k} = d\}$.

Second, we consider the effect that the enactment ϕ' of another commitment ϕ, also by agent β, has on $\vec{p}$. Given the observation $\mu = (\phi', \phi)$, define the vector $\vec{t}$ by

$$t_i = \mathbb{P}^t(\varphi_i|\varphi) + (1- \mid \text{Sim}(\phi', \phi) - \text{Sim}(\varphi_i, \varphi) \mid) \cdot \text{Sim}(\varphi', \phi)$$

for $i = 1, \ldots, m$. $\vec{t}$ is not a probability distribution. The posterior $\vec{p}_{(\phi', \phi)}$ is defined to be the normalisation of $\vec{t}$.

Trust based on Ideal enactments

A distribution of enactments that represent α's "ideal". This distribution will be a function of β, α's history with β, anything else that α believes about β, and general environmental information including time — denote all of this by e, then we have $\mathbb{P}^t_I(\varphi'|\varphi, e)$. For example, if α considers that it is unacceptable for the enactment φ' to be less preferred than the commitment φ then $\mathbb{P}^t_I(\varphi'|\varphi, e)$ will only be non-zero for those φ' that α prefers to φ. Trust is the relative entropy between this ideal distribution, $\mathbb{P}^t_I(\varphi'|\varphi, e)$, and the distribution of expected enactments, $\mathbb{P}^t(\varphi'|\varphi)$. That is:

$$T(\alpha, \beta, \varphi) = 1 - \sum_{\varphi'} \mathbb{P}^t_I(\varphi'|\varphi, e) \log \frac{\mathbb{P}^t_I(\varphi'|\varphi, e)}{\mathbb{P}^t(\varphi'|\varphi)}$$

where the "1" is an arbitrarily chosen constant being the maximum value that trust may have.

Trust based on Preferred enactments

Given a predicate $\text{Prefer}(c_1, c_2, e)$ meaning that α prefers c_1 to c_2 in environment e. An evaluation of $\mathbb{P}^t(\text{Prefer}(c_1, c_2, e))$ may be defined using $\text{Sim}(\cdot)$ and the evaluation function $\vec{w}(\cdot)$ — but we do not detail it here. Then if $\varphi \leq o$:

$$T(\alpha, \beta, \varphi) = \sum_{\varphi'} \mathbb{P}^t(\text{Prefer}(\varphi', \varphi, o))\mathbb{P}^t(\varphi' \mid \varphi)$$

and:

$$T(\alpha,\beta, o) = \frac{\sum_{\varphi:\varphi \leq o} \mathbb{P}^t_\beta(\varphi) \left[\sum_{\varphi'} \mathbb{P}^t(\text{Prefer}(\varphi', \varphi, o))\mathbb{P}^t(\varphi' \mid \varphi)\right]}{\sum_{\varphi:\varphi \leq o} \mathbb{P}^t_\beta(\varphi)}$$

Trust based on Certainty in enactment

The idea here is that α will trust β more if variations, φ', from expectation, φ, are not random. The trust that an agent α has on agent β with respect to the enactment of a commitment φ is:

$$T(\alpha, \beta, \varphi) = 1 + \frac{1}{B^*} \cdot \sum_{\varphi' \in \Phi_+(\varphi, o, \kappa)} \mathbb{P}^t_+(\varphi'|\varphi) \log \mathbb{P}^t_+(\varphi'|\varphi)$$

where $\mathbb{P}^t_+(\varphi'|\varphi)$ is the normalisation of $\mathbb{P}^t(\varphi'|\varphi)$ for $\varphi' \in \Phi_+(\varphi, o, \kappa)$,

$$B^* = \begin{cases} 1 & \text{if } |\Phi_+(\varphi, o, \kappa)| = 1 \\ \log |\Phi_+(\varphi, o, \kappa)| & \text{otherwise} \end{cases}$$

Ontology

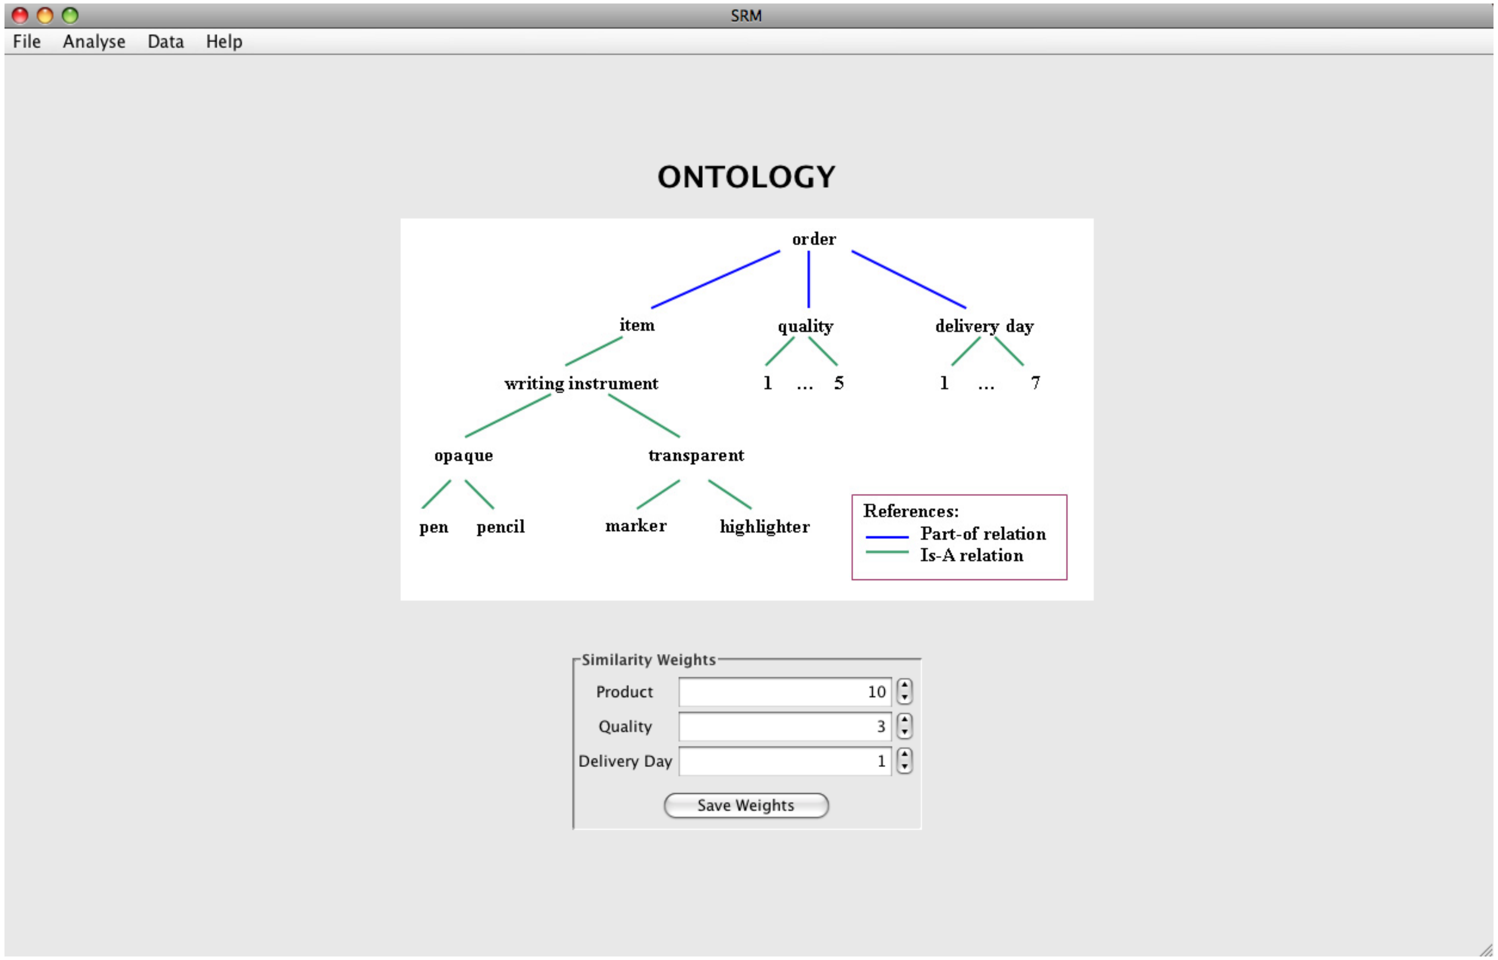

SRM
File
Analyse
Data
Help
ONTOLOGY
order
item
quality
delivery day
writing instrument
1 ... 5
1 ... 7
opaque
transparent
pen
pencil
marker
highlighter
References:
Part-of relation
Is-A relation
Similarity Weights
Product 10
Quality 3
Delivery Day 1
Save Weights

Contracts

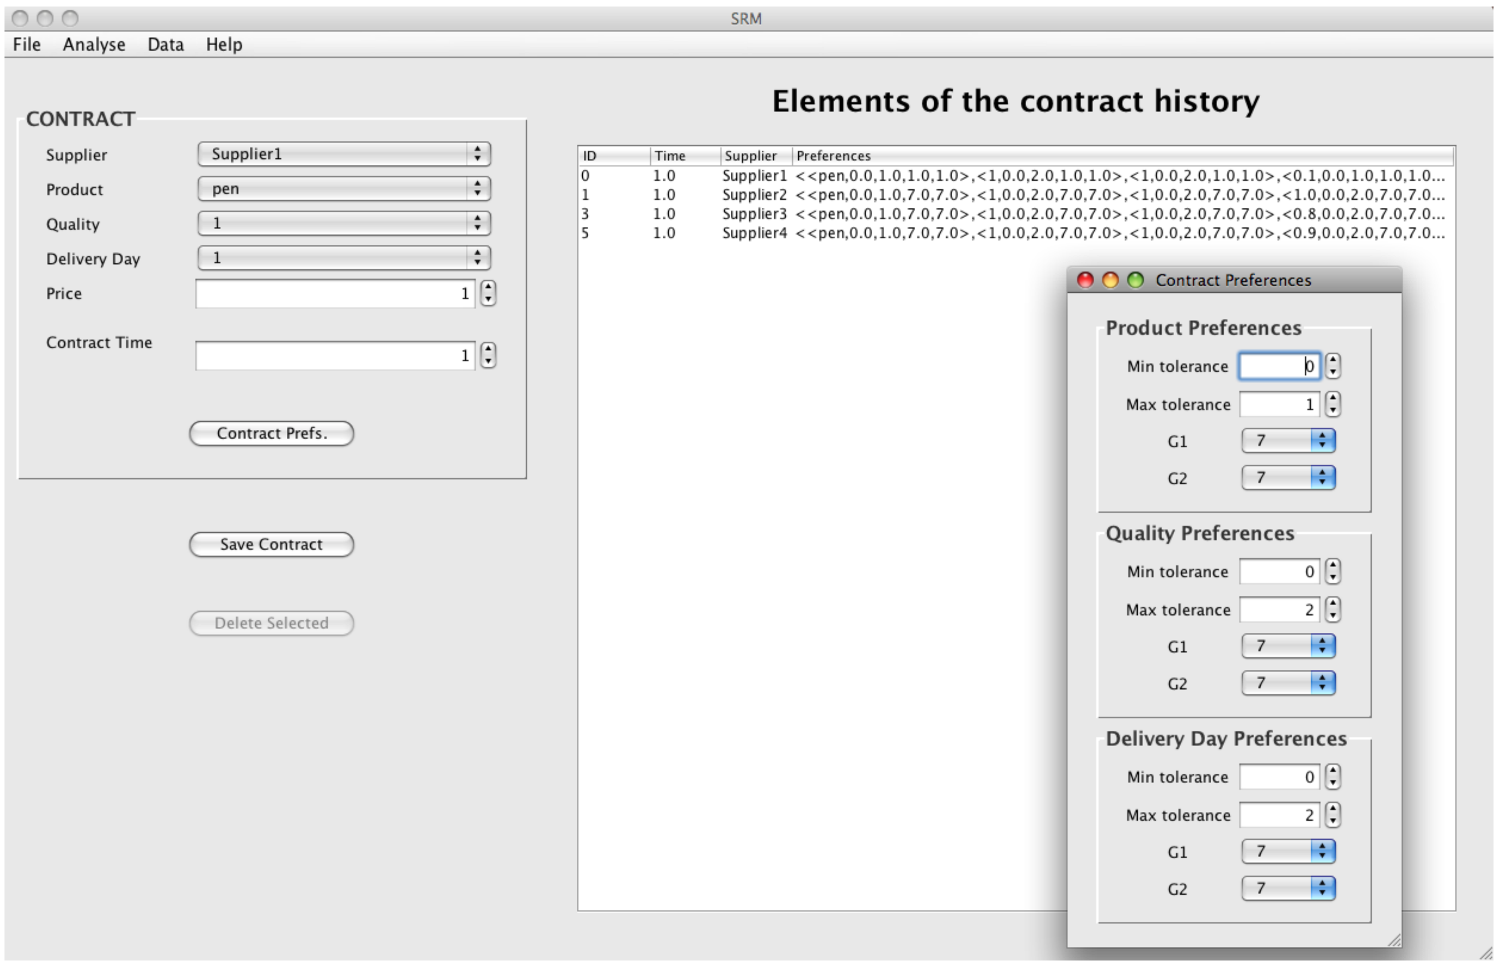
SRM
File Analyse Data Help
CONTRACT
Supplier
Supplier1
Product
pen
Quality
1
Delivery Day
1
Price
1
Contract Time
1
Contract Prefs.
Save Contract
Delete Selected
Elements of the contract history
ID Time Supplier Preferences
0 1.0 Supplier1 <<pen,0.0,1.0,1.0,1.0>,<1,0.0,2.0,1.0,1.0>,<1,0.0,2.0,1.0,1.0>,<0.1,0.0,1.0,1.0,1.0...
1 1.0 Supplier2 <<pen,0.0,1.0,7.0,7.0>,<1,0.0,2.0,7.0,7.0>,<1,0.0,2.0,7.0,7.0>,<1.0,0.0,2.0,7.0,7.0...
3 1.0 Supplier3 <<pen,0.0,1.0,7.0,7.0>,<1,0.0,2.0,7.0,7.0>,<1,0.0,2.0,7.0,7.0>,<0.8,0.0,2.0,7.0,7.0...
5 1.0 Supplier4 <<pen,0.0,1.0,7.0,7.0>,<1,0.0,2.0,7.0,7.0>,<1,0.0,2.0,7.0,7.0>,<0.9,0.0,2.0,7.0,7.0...
Contract Preferences
Product Preferences
Min tolerance 0
Max tolerance 1
G1 7
G2 7
Quality Preferences
Min tolerance 0
Max tolerance 2
G1 7
G2 7
Delivery Day Preferences
Min tolerance 0
Max tolerance 2
G1 7
G2 7

Experiences

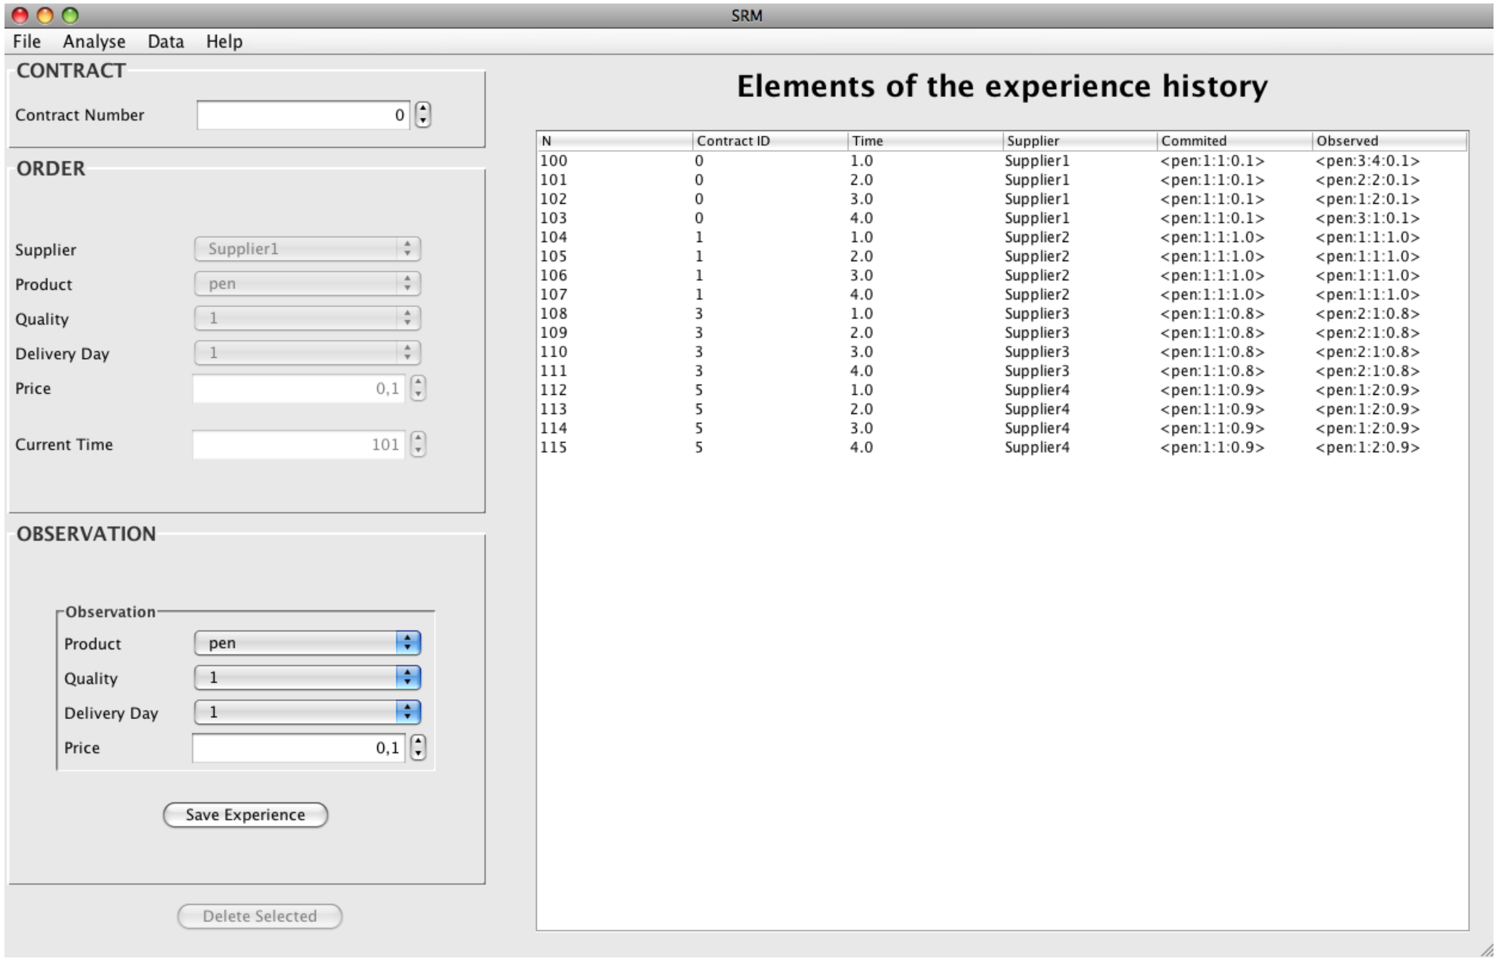

Probabilities

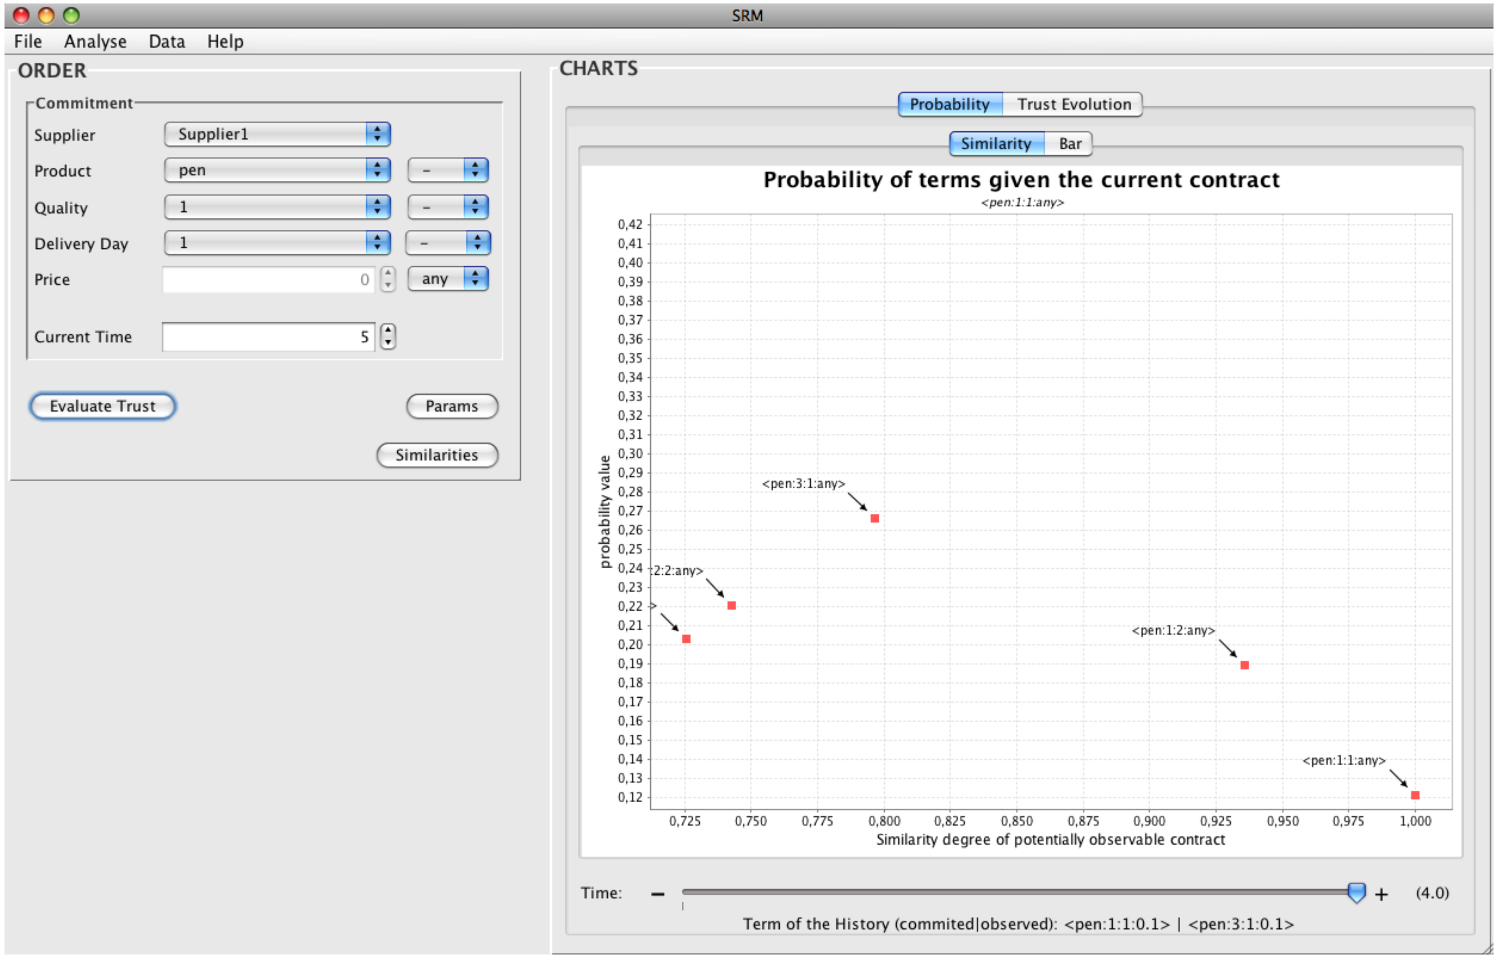

Probabilities

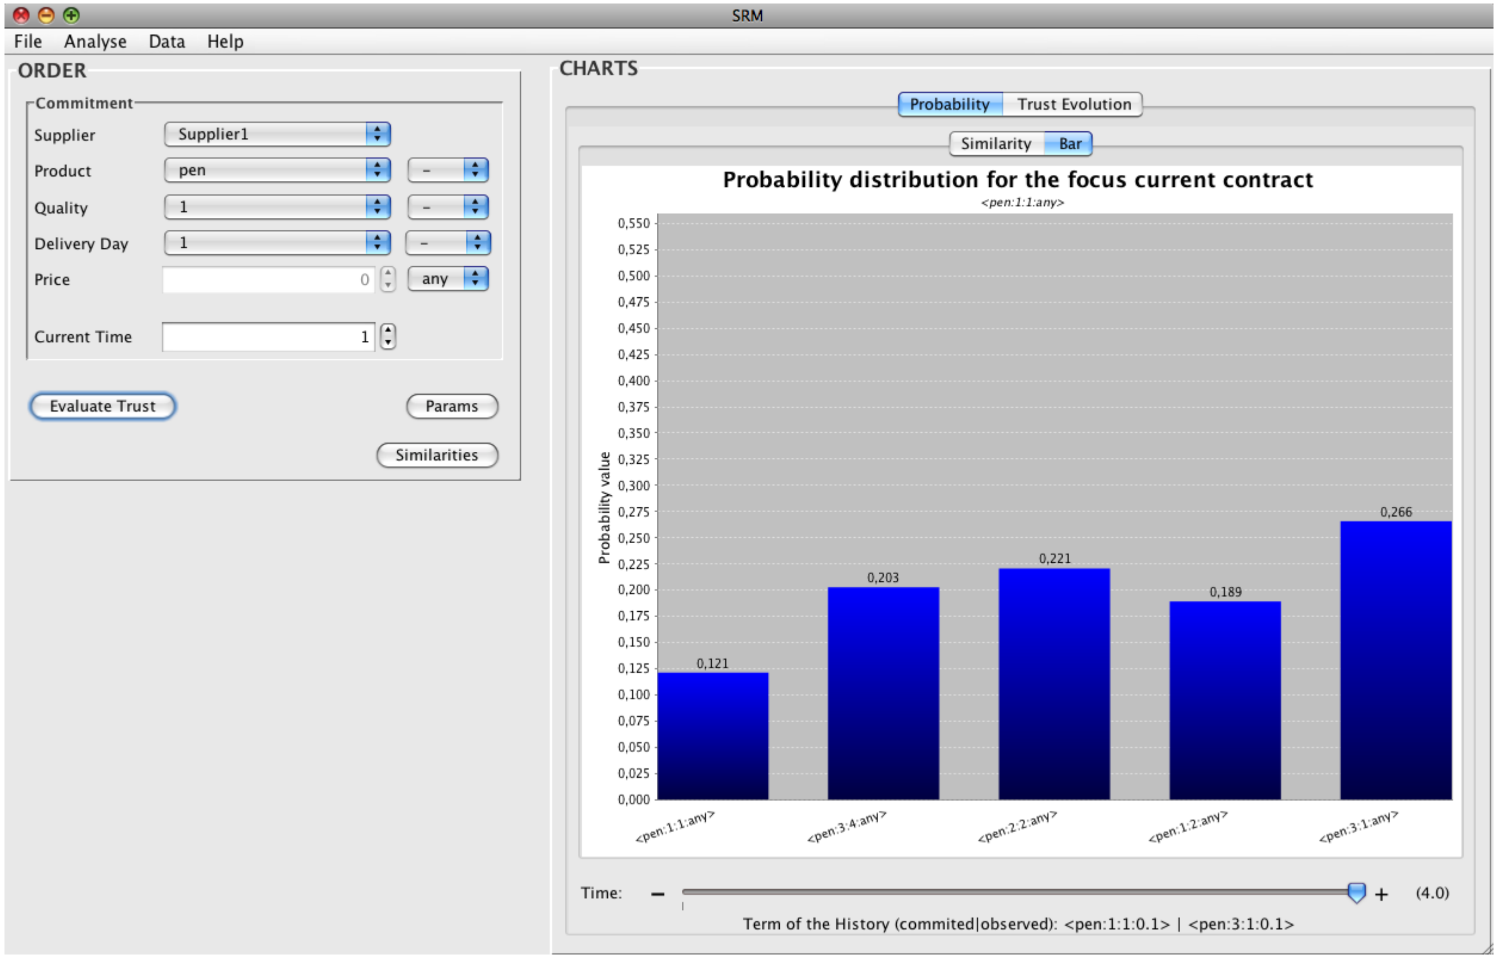

Trust

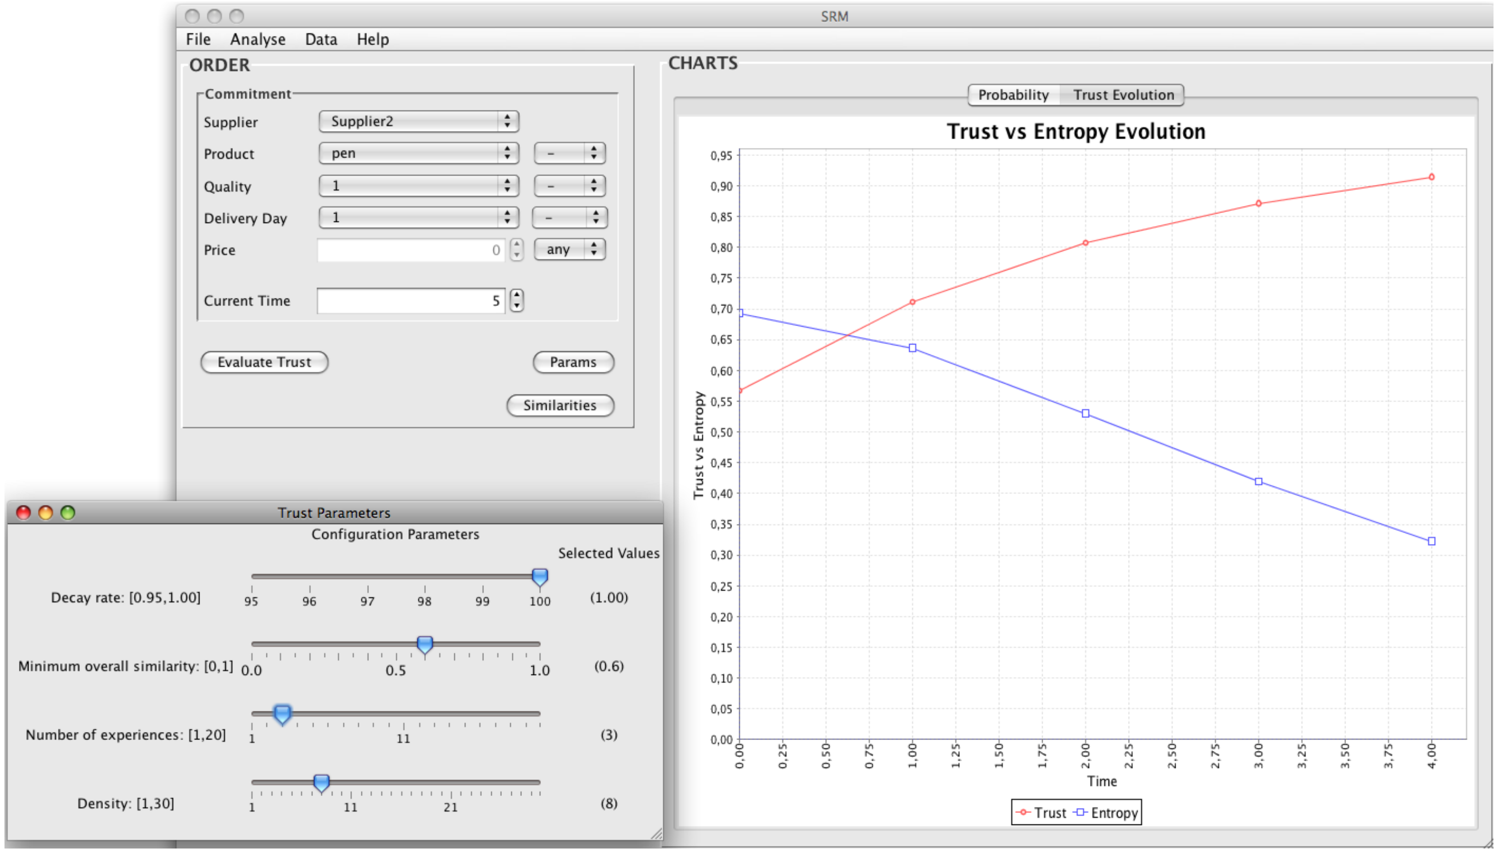

SRM
File Analyse Data Help
ORDER
Commitment
Supplier Supplier2
Product pen
Quality 1
Delivery Day 1
Price 0 any
Current Time 5
Evaluate Trust
Params
Similarities
CHARTS
Probability Trust Evolution
Trust vs Entropy Evolution
Trust vs Entropy
Time
Trust Entropy
Trust Parameters
Configuration Parameters
Selected Values
Decay rate: [0.95,1.00] (1.00)
Minimum overall similarity: [0,1] (0.6)
Number of experiences: [1,20] (3)
Density: [1,30] (8)

Suppliers

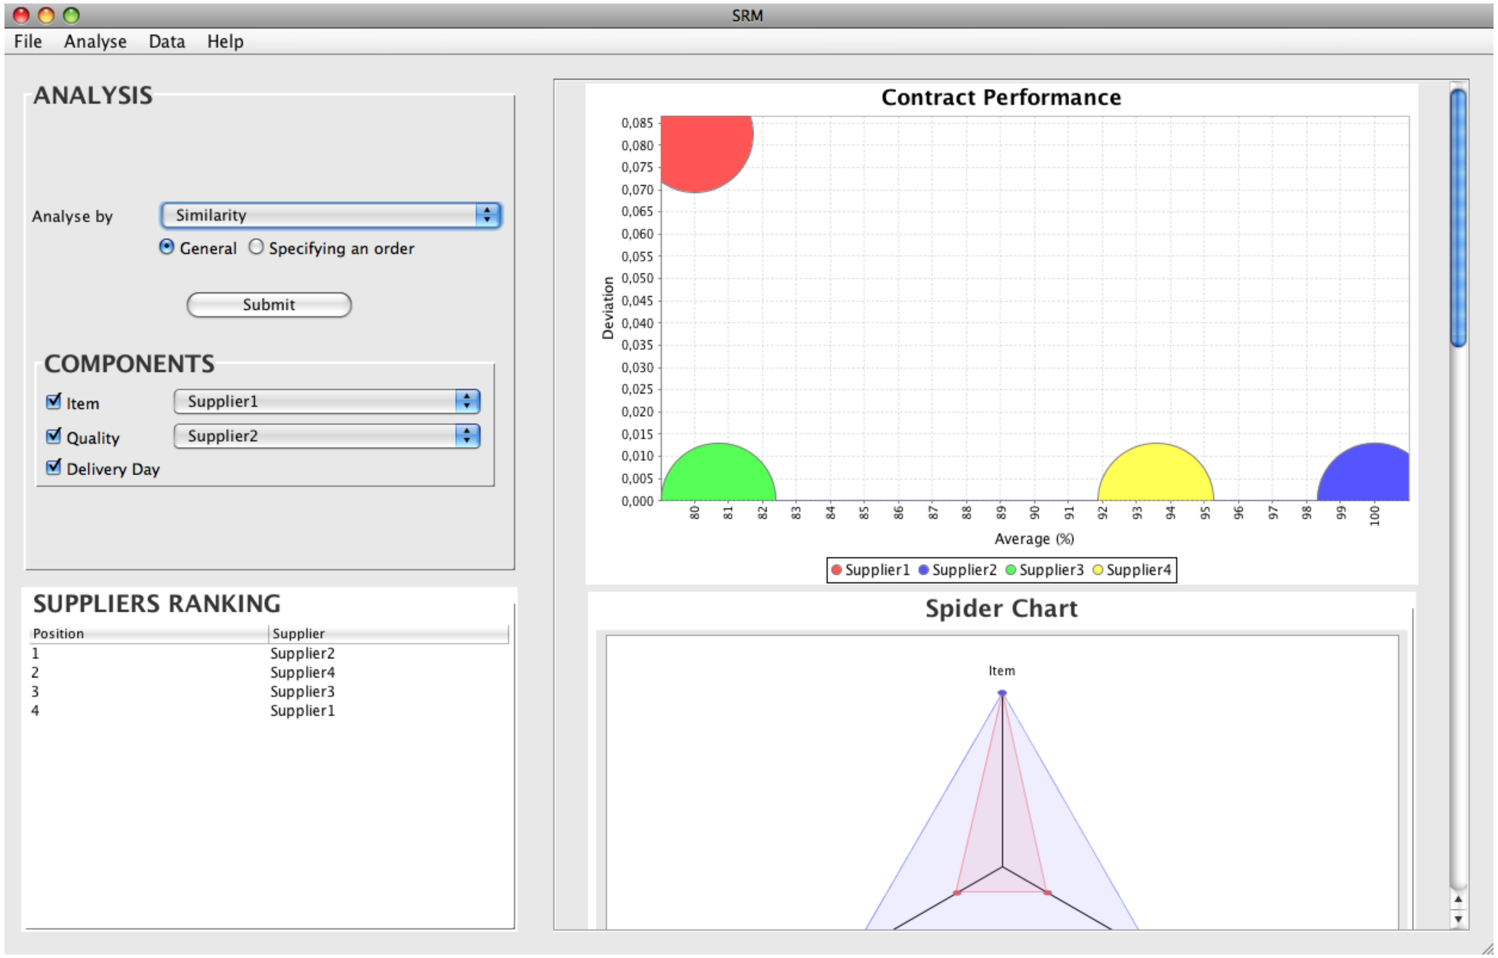

SRM
File Analyse Data Help
ANALYSIS
Analyse by
Similarity
General
Specifying an order
Submit
COMPONENTS
Item
Supplier1
Quality
Supplier2
Delivery Day
SUPPLIERS RANKING
Position Supplier
1 Supplier2
2 Supplier4
3 Supplier3
4 Supplier1
Contract Performance
Deviation
Average (%)
Supplier1 Supplier2 Supplier3 Supplier4
Spider Chart
Item

Suppliers Analysis Charts

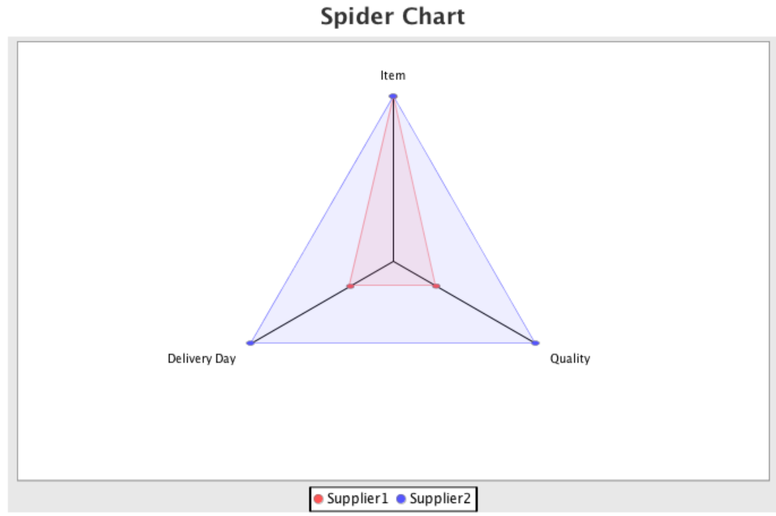

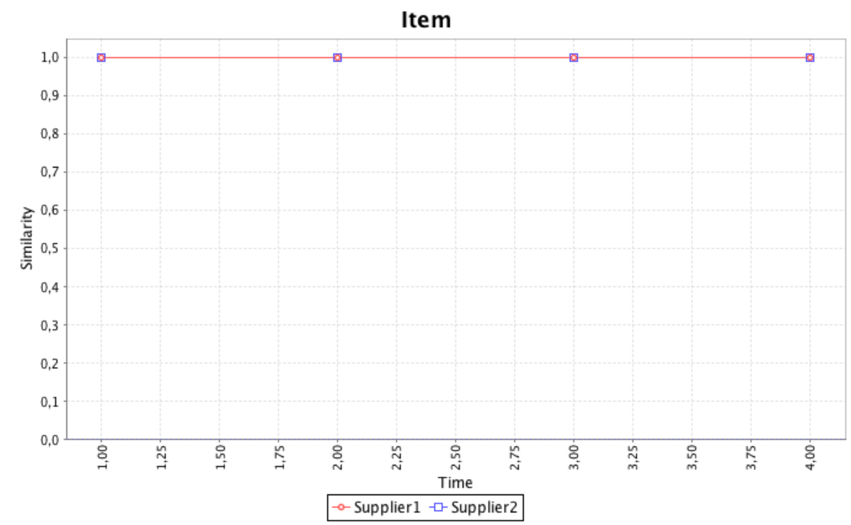

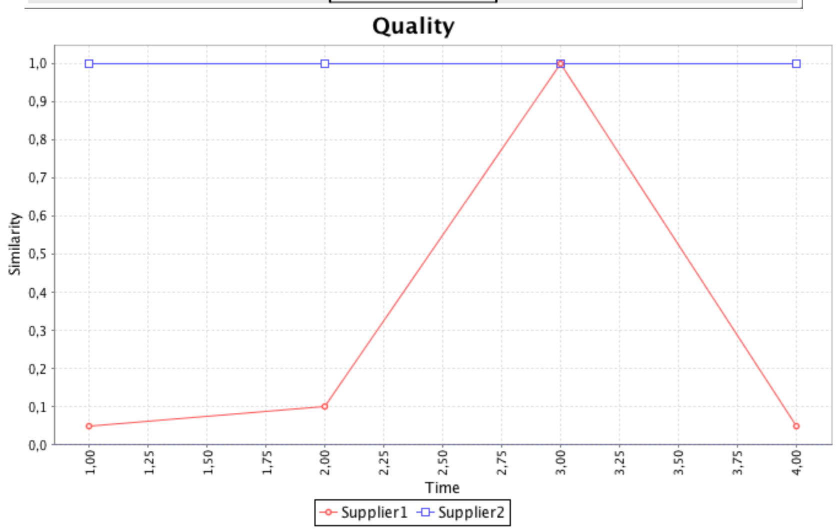

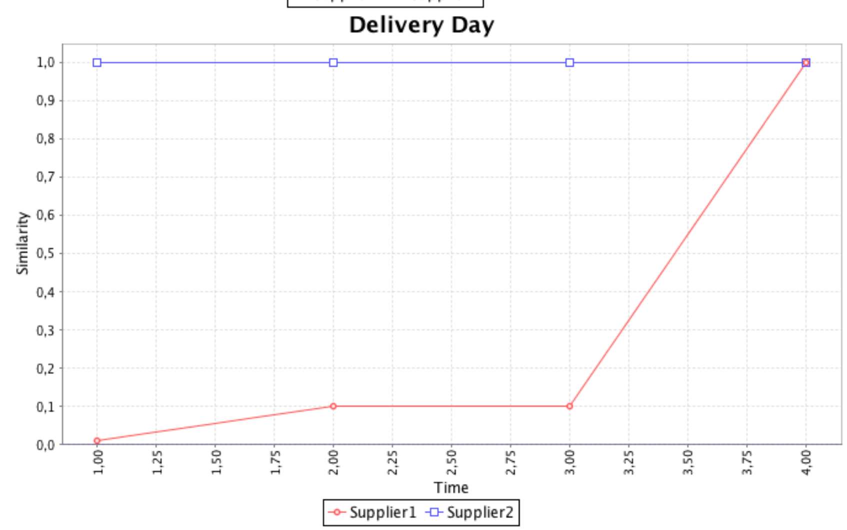

Critical Order

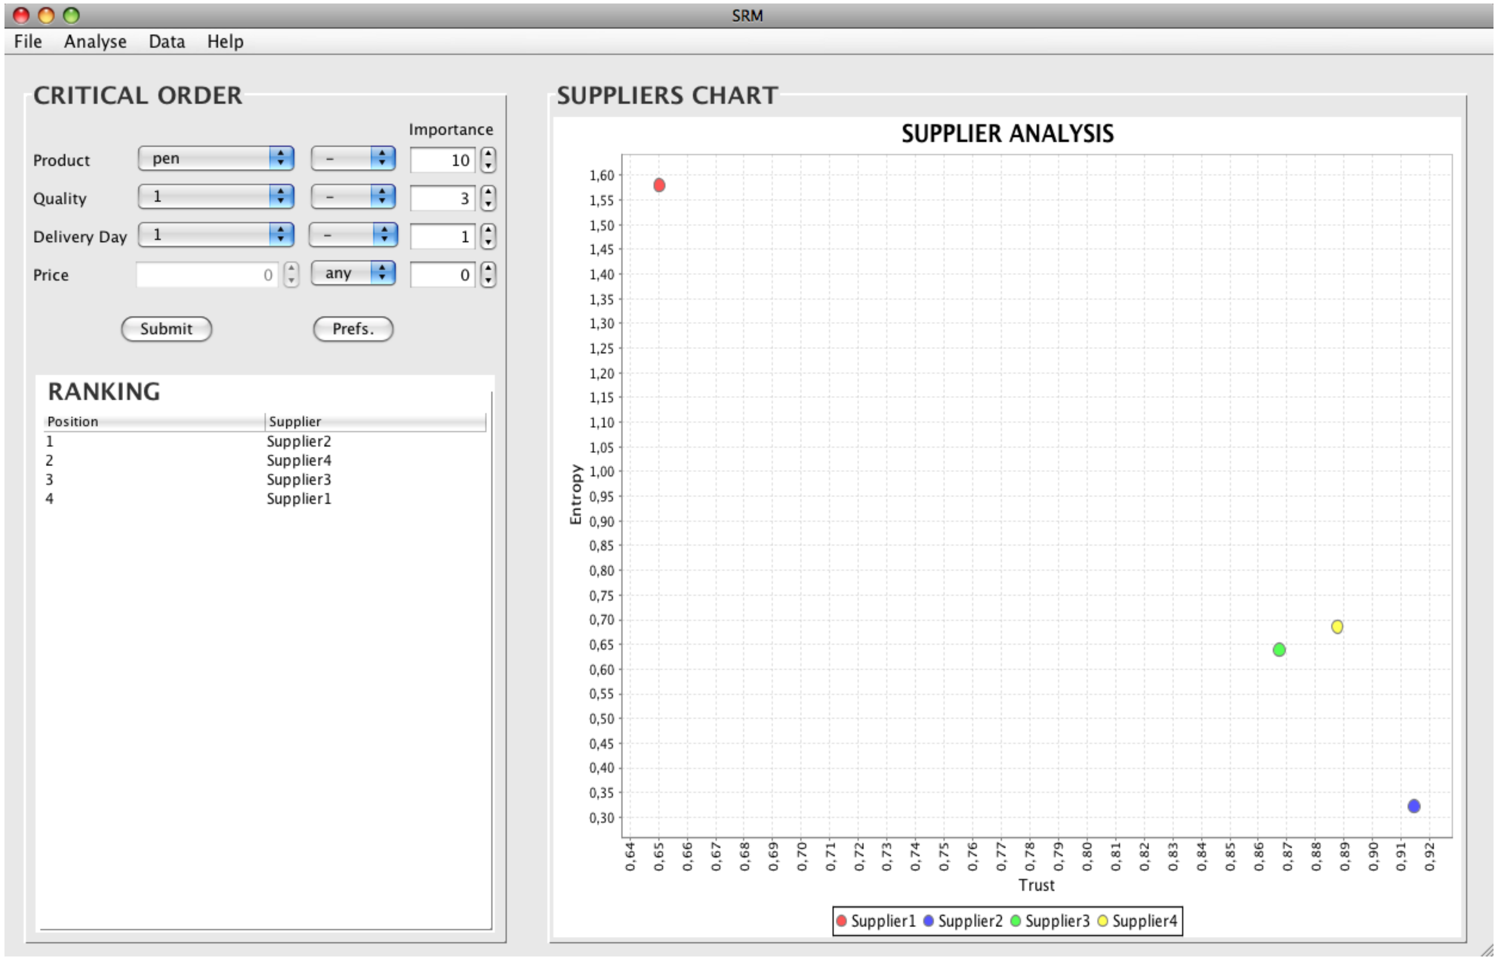

Cost Minimization

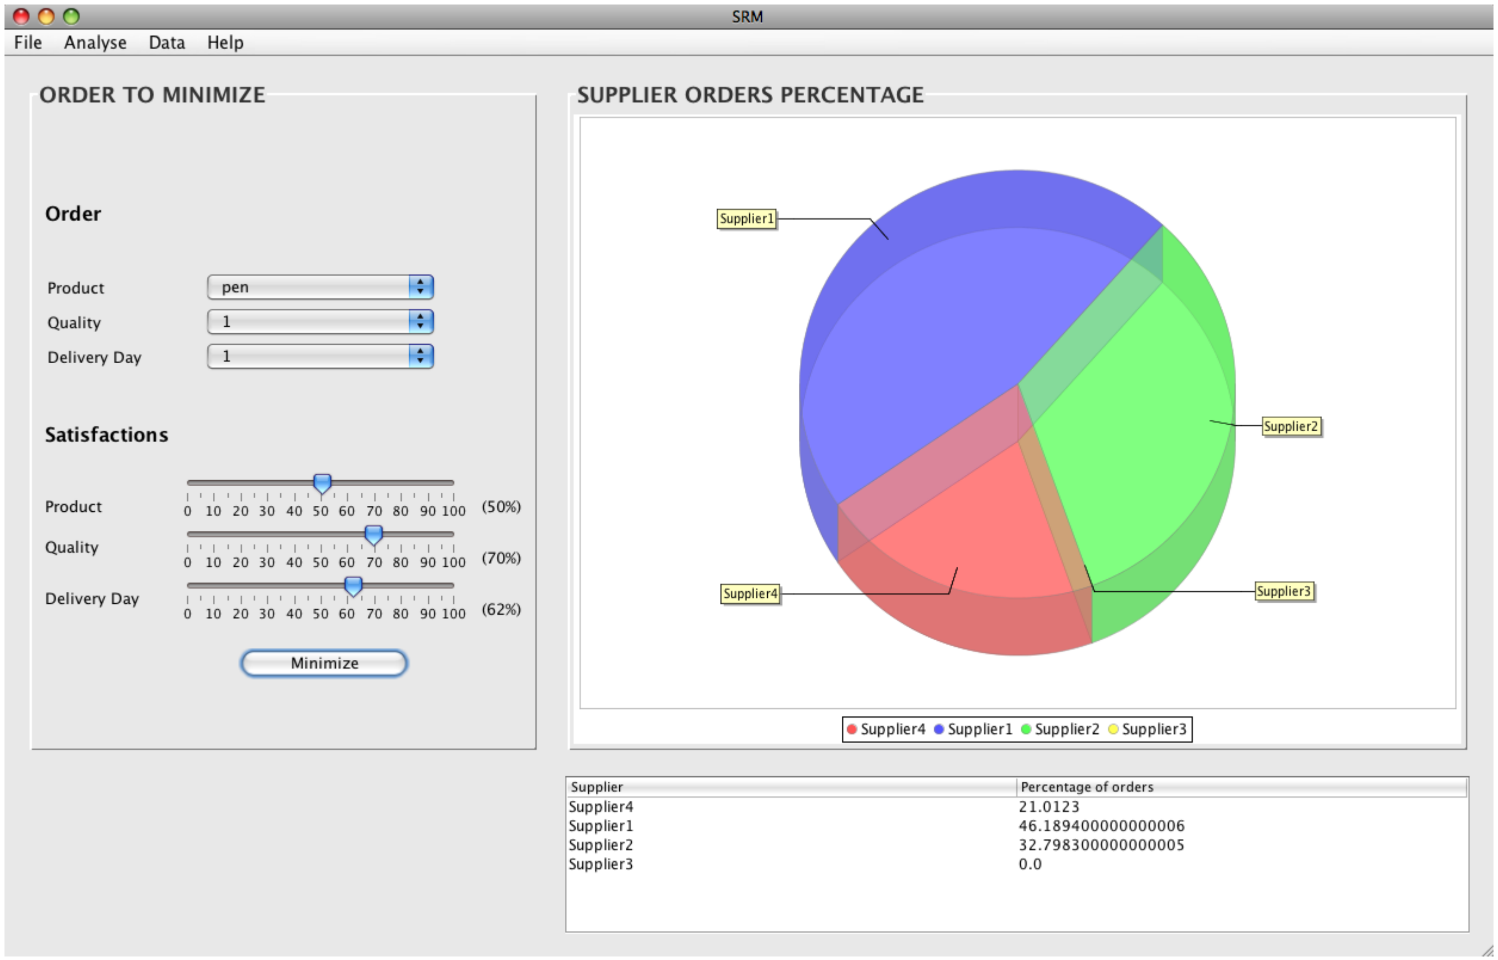

A Reputation example: Liquid publishing (I)

The IJCAI conference uses a reputation-enhanced Conference Management system. The following happens:

- Luigi Bamboozel, the famous professor of sociology at the University of Torino submits a paper "The Logic of Social Networks" to IJCAI. Mike is the senior PC member in charge of the paper, the reviewers are: Cristiano, Stephen and Paul.

- At the review stage they rate the paper as:
  - Cristiano —Accept on the grounds that it opens up a new area
  - Stephen —Reject on the grounds that the typography is poor and ambiguous (Microsoft Word again) and so is the grammar and spelling
  - Paul —Strong Accept on the grounds that he was most impressed by Bamboozels latest book

Liquid publishing (II)

Then follows the discussion phase (directed by Mike) and after it the scores are:

- Cristiano —Weak Accept
- Stephen —Weak Reject
- Paul —Weak Accept

and the paper is accepted as a poster.

Liquid publishing (III)

Our model should help in determining the reputation of:

- Bamboozel as an author and as a MAS author
- The ideas in the paper
- Cristiano as a reviewer and as a scholar
- Stephen as a reviewer and as a scholar
- Paul as a reviewer as as a scholar
- Mike as a senior PC member
- The IJCAI Conference

Reputation

- Reputation = Group opinion on someone or something.
- Opinions use a shared evaluation space E.
- An opinion is an agents evaluation of a particular aspect of a thing in context. A representation will contain: the thing, its aspect, its context, and a probability distribution on E representing the evaluation of the thing.
- Difference in opinion = distance between distributions.
- We will model reputation in the Information-based agency framework.
- Fundamental for e.g. recommender systems, trading sites,

Opinions

- Verifiable opinions. "tomorrow's maximum temperature will be over 30°". We can establish the relationship between opinion and fact, $\mathbb{P}^t(\text{Observe}(\varphi')|\text{Commit}(\varphi))$ simply as $\mathbb{P}^t(\varphi'|\varphi) \in \mathcal{M}^t$. This allows to compute $\mathbb{R}^t(\alpha, \beta, \mu)$.

- Unverifiable opinions. "the Earth will exist in 100,000 years time". If an opinion can not be verified then one way in which it may be evaluated is to compare it with the corresponding individual opinions of a group of agents. See next.

Forming group opinions

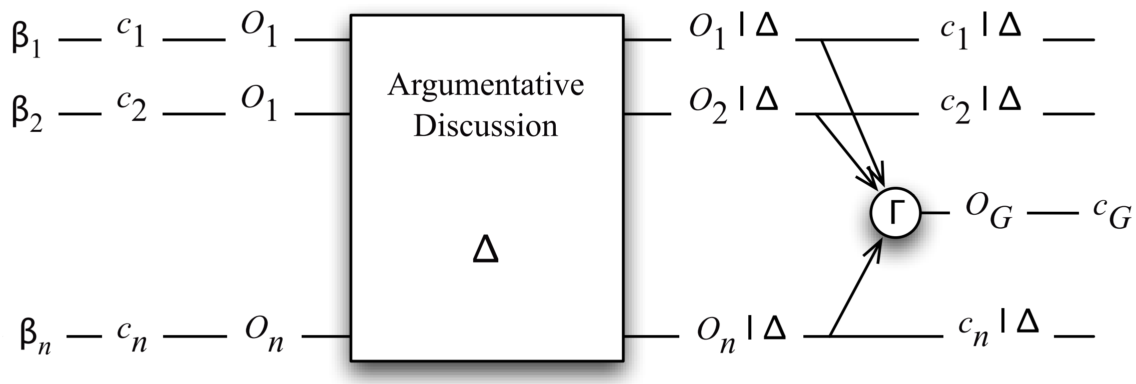

We propose next:

- some distances between opinions
- three methods to aggregate opinions

Distances between distributions: Kullback

Given two probability distributions P and Q over a random variable $X = \{x_1, \ldots, x_n\}$ the Kullback-Leibler divergence is:

$$DIST(P, Q) = D_{KL}(P||Q) = \sum_i P(i) \log\left(\frac{P(i)}{Q(i)}\right)$$

Distances between distributions: EMD

Given two probability distributions P and Q over a random variable $X = \{x_1, \ldots, x_n\}$ and a distance matrix $D = [d_{ij}]$ between any two values $x_i, x_j \in X$, we want the flow $F = [f_{ij}]$ with f_{ij} the flow from x_i to x_j that minimises the overall cost

$$DIST(P,Q) = EMD(P,Q) = WORK(P,Q,F) = \sum_{i=1}^{n}\sum_{j=1}^{n} f_{ij}d_{ij}$$

$$\begin{array}{llll} f_{ij} \geq 0 & 1 \leq i,j \leq m & \sum_{j=1}^{n} f_{ij} & \leq p_i 1 \leq i \leq n \\ \sum_{i=1}^{n} f_{ij} \leq q_j & 1 \leq j \leq n & \sum_{i=1}^{n}\sum_{j=1}^{n} f_{ij} & = 1 \end{array}$$

NOTE: If the space is not ordered, the distance can be the probability of that change during dialogues.

Method 1: Joint Distribution Method

Given two opinions: $\mathbb{P}(X = x)$ and $\mathbb{P}(Y = y)$ where opinion X is correct c% of the time and Y d%. Construct the joint distribution $W = (X, Y, Z)$ and impose the constraints:

$$\left(\sum_i \mathbb{P}(W = w_i) \mid X = x_j\right) = \mathbb{P}(X = x_j)$$

$$\left(\sum_i \mathbb{P}(W = w_i) \mid Y = y_j\right) = \mathbb{P}(Y = y_j)$$

$$\left(\sum_i \mathbb{P}(W = w_i) \mid X = Z\right) = c$$

$$\left(\sum_i \mathbb{P}(W = w_i) \mid Y = Z\right) = d$$

then the *joint opinion* is $\mathbb{P}(Z = z)$.

Method 2: Yager Method.

Given a prior distribution $\mathbb{P}(W = x_j)$, a pair of opinions, $\mathbb{P}(X_i = x_j)$ $i = 1, 2$, with their respective certainties c_i, assuming that the agents are independent, let $w_{i,j} = c_i \times \mathbb{P}(X_i = x_j)$, $i = 1, 2$, and let

$$v_j = \frac{\prod_i w_{i,j}}{\prod_i w_{i,j} + \prod_i (1 - w_{i,j})}$$

then the posterior is:

$$\mathbb{P}(Z = x_j) = v_j + \left(1 - \sum_k v_k\right) \times \mathbb{P}(W = x_j)$$

Method 3: Maximum Uninorm

$P(\alpha, d)$ as the probability that an opinion expressed by α is at distance d of the true (or of α's group) distribution. This can be learned from past cases.

Given a group G, a set of independent opinions $\{O_i\}_{i \in G}$ the Maximum Uninorm group opinion, O_G is the distribution that maximizes the uninorm value of the probabilities of the agents being at the corresponding distance: $P(i, DIST(O_i, O_G))$. That is,

$$O_G = \max_Q \Upsilon(\{P(i, DIST(O_i, Q))\}_{i \in G})$$

where

$$\Upsilon(p_1, p_2, \ldots, p_n) = \frac{p_1 p_2 \ldots p_n}{p_1 p_2 \ldots p_n + (1 - p_1)(1 - p_2) \ldots (1 - p_n)}$$

Examples

Prior 1	0.1000	0.5000	0.2000	0.1000	0.1000	Strength = 0.9	$P = \langle 0.9, 0.05, 0.03, 0.01, 0.01 \rangle$
Prior 2	0.0500	0.8000	0.0500	0.0500	0.0500	Strength = 0.7	$P = \langle 0.7, 0.2, 0.05, 0.03, 0.02 \rangle$
Joint	0.0919	0.5590	0.1653	0.0919	0.0919		
Yager	0.0978	0.6044	0.1022	0.0978	0.0978		
MaxUni	0.0700	0.7000	0.1700	0.0700	0.0700	Strength = 0.95	

Prior 1	0.1000	0.6000	0.1000	0.1000	0.1000	Strength = 0.8	$P = \langle 0.8, 0.1, 0.04, 0.01, 0.01 \rangle$
Prior 2	0.0500	0.8000	0.0500	0.0500	0.0500	Strength = 0.9	$P = \langle 0.9, 0.06, 0.03, 0.01, 0.01 \rangle$
Joint	0.0683	0.7266	0.0683	0.0683	0.0683		
Yager	0.0601	0.7596	0.0601	0.0601	0.0601		
MaxUni	0.08	0.63	0.08	0.08	0.08	Strength = 0.97	

Prior 1	0.0500	0.8000	0.0500	0.0500	0.0500	Strength = 0.8	$P = \langle 0.8, 0.1, 0.04, 0.01, 0.01 \rangle$
Prior 2	0.0500	0.8000	0.0500	0.0500	0.0500	Strength = 0.9	$P = \langle 0.9, 0.06, 0.03, 0.01, 0.01 \rangle$
Joint	0.0573	0.7707	0.0573	0.0573	0.0573		
Yager	0.0363	0.8548	0.0363	0.0363	0.0363		
MaxUni	0.05	0.8	0.05	0.05	0.05	Strength = 0.97	

Reputation labels

Inexorable. If agent β_i is such that: $\mathrm{Dist}(O_i, O_i|\Delta) \ll \mathrm{Dist}(O_i, O_j|\Delta), \forall j \neq i$ consistently holds then β_i is *inexorable*.

Predetermination. If: $\mathrm{Dist}(O_i, R_G) \ll \mathrm{Dist}(O_j, R_G), \forall j \neq i$ consistently, then β_i is a good '*predeterminer*'. Such an agent will have a high c_i value.

Persuasiveness. If β_i is such that: $\mathrm{Dist}(O_i, O_j|\Delta) \ll \mathrm{Dist}(O_j, O_j|\Delta), \forall j \neq i$ consistently then β_i is *persuasive*.

Compliance. If β_i is such that: $O_i|\Delta \approx \arg\min_X \sum_{j \neq i} \mathrm{Dist}(O_j|\Delta, X)$, then β_i is *compliant*.

Dogmatic. If β_i is such that: $O_i = O_i|\Delta$ consistently then β_i is *dogmatic*. A dogmatic agent is highly inexorable.

Adherence. If β_i is such that $O_i|\Delta = O_j$ where $j = \arg\max_{k, k \neq i} c_k$ consistently then β_i is *adherent* (in this round adherent to agent β_j).

Social Network Measures

Given a matrix $R(n,n)$ that represents in $r_{ij} \in [0,1]$ the intensity of the relation R from i to j we define:

- Normalised Degree Centrality. $C_d(i) = \frac{\sum_{j=1}^{n} r_{ij}}{n-1}$
- Normalised Closeness Centrality. $C_c(i) = \frac{n-1}{\sum_{j=1}^{n} d(i,j)}$ where $d(i,j)$ is the minimum distance between i and j in the graph
- Normalised Betweenness Centrality. $C_b(i) = \frac{2}{(n-1)(n-2)} \cdot \sum_{j,k \neq i, j \neq k} \frac{s_{jk}(i)}{s_{jk}}$ where $s_{jk}(i)$ is the number of shortest paths between j and k including i, and s_{jk} is the total number of shortest paths between j and k.
- Prestige Degree. $P(i) = \frac{\sum_{j=1}^{n} r_{ji}}{n-1}$

Information based social measures

There are three relevant information-based measures among agents:

- Information $= avg(\Delta H)$ The average increase/decrease in entropy of the distributions of an agent due to information received. How well informed and informative an agent is.
- Persuasion $= avg(\Delta O)$ The average change in opinion due to dialogues with an agent. How persuasive is an agent.
- Closeness $= avg(dist)$ The average distance in opinions between both agents. The smaller value, the closer the way both agents see things.

More relationships. LiquidPub case

A LiquidPub Network is defined as $LPN = \langle V, R, \sigma \rangle$

- $V = V_\alpha \cup V_p \cup V_k$ is the set of nodes, union of individuals, publications, and keywords.
- $R = \{\mathit{Authorship}, \mathit{Citation}, \mathit{Version}, \mathit{Part}, \mathit{Review}, \mathit{Area}, \mathit{College}, \mathit{Affiliation}\}$ is a set of relationships on the nodes of the network.
  - $\mathit{Authorship} \subseteq V_\alpha \times V_p$. $(i, p) \in \mathit{Authorship} = i$ is author of p.
  - $\mathit{Citation} \subseteq V_p \times V_p$. $(p, p') \in \mathit{Citation} = p$ cites p'.
  - $\mathit{Version} \subseteq V_p \times V_p$. $(p, p') \in \mathit{Version} = p'$ is an improved version of p.
  - $\mathit{Part} \subseteq V_p \times V_p$. $(p, p') \in \mathit{Part} = p$ is part of p'.
  - $\mathit{Review} \subseteq V_\alpha \times V_p$. $(i, p) \in \mathit{Review} = i$ is a reviewer of p.
  - $\mathit{Area} \subseteq V_p \times V_k$. $(p, k) \in \mathit{Area} = p$ is about keyword k.
  - $\mathit{College} \subseteq V_\alpha \times V_\alpha$. $(i, j) \in \mathit{College} = i$ is a colleague of j.
  - $\mathit{Affiliation} \subseteq V_\alpha \times V_\alpha$. $(i, j) \in \mathit{Affiliation} = i$ and j belong to the same organisation.
- $\sigma = \{\sigma_r\}_{r \in R}$ is a labeling function.

Measures

It is possible to define what is the relative certainty (expertise) of the opinion of agent i on a paper on topic X, $c_i(X)$. For instance, an individual is expert in an area (keyword) if it is author of highly cited papers on the topic, has reviewed prestigious papers on the area, and has a central role in the college.

$$c_i(X) = f\left(\sum_{\substack{(i,p)\in Authorship,\\(p,X)\in Area}} P^{Citation}(p), \sum_{\substack{(i,p)\in Review,\\(p,X)\in Area}} P^{Citation}(p), C_b College(i)\right)$$

Currently used measures (and perhaps alternatives to the previous one) are easy to compute, e.g. the h index is simply:

$$h(i) = \arg\max_k \left|\{p \mid (i,p) \in Authorship, P^{Citation}(p) \geq k\}\right| \geq k$$

Group Opinion: SNA modulating dependence

The *persuasion* relationship should modify (flattens, increases the entropy) $P(\alpha, d)$. Makes the agents more uncertain of their opinion in the particular case.

The *information* and *closeness* relationships could be used as a heuristic on the possible dependency among opinions. But, how to factor it in is unclear.

Overall, the SNA analysis should help in determining the *reliability* of an opinion.

When opinions are fully dependent the $\max$ operator over the reliability is to be used. When they are fully independent the Yager operator is the adequate one. SNA dependency measures may determine a point between both extremes.

Future work

- Study the role of SNA in the equations.
- (ongoing) Implement the reputation model.
- Test with real data. Conference data in project LiquidPub.
- Apply it to Supplier Relationship Management.