

© [2009] IEEE. Reprinted, with permission, from [John Debenham, Simeon Simoff, Believable electronic trading environments on the Web, 2009, IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, 2009]. This material is posted here with permission of the IEEE. Such ermission of the IEEE does not in any way imply IEEE endorsement of any of the University of Technology, Sydney's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it

Believable electronic trading environments on the Web

John Debenham

*Centre for Quantum Computation & Intelligent Systems
University of Technology, Sydney
NSW 2007 Australia
debenham@it.uts.edu.au*

Simeon Simoff

*School of Computing & Mathematics
University of Western Sydney
Sydney, Australia
S.Simoff@uws.edu.au*

Abstract—Contemporary Web-based electronic markets reflect the dominating content-based systems approach of Web 2.0. Though useful, these electronic markets are far from being believable trading places. Marketplace is where things and traders have presence, constituting a rich interaction space. The believability of the place depends on the believability of the presence and interactions in it, including the players' behaviour and the narrative scenarios of the marketplace. This paper discusses what constitutes the believability of electronic marketplaces and presents the technologies that support it. Believability of electronic marketplaces can be described through three metaphors: “marketplaces where people are”, “marketplaces that are alive and engaging”, and “market places where information is valuable and useful”. The paper presents the core technologies that enable the perceivable believability of electronic marketplaces. It describes a demonstrable prototype of a Web-based electronic marketplace that integrates these technologies. This is part of a larger project that aims to make informed automated trading an enjoyable reality of Web 3.0.

Keywords-trading environments; virtual spaces; believability;

I. INTRODUCTION

A market is commonly defined as a physical or virtual location, where price is determined and buy and sell orders are matched to create trades according to a set of rules that govern the processing of these orders [1]. Electronic markets have been viewed as information systems “that allow buyers and vendors to exchange information about prices and product offerings [2]. This and similar views have guided the development of “soulless” electronic markets, focussed primarily on enabling standardised or complex transaction processes. Thus automation of electronic markets have been focused on the secure backend transaction processing. A recent review of the area (see Chapter 18 “Electronic Marketplaces and Resource Exchanges” in [3]) provides a broader picture from various perspectives, including agent-based negotiation, brokering, and partnership formation. Still, the operation and the interactions in such Web-based electronic markets reflect the dominating content-based systems approach of Web 2.0. Though useful, these electronic markets are far from being believable trading places.

In this paper we consider electronic (virtual) marketplace to be a regulated space populated by computerised players

that represent a variety of human and software traders, intermediaries, and information and infrastructure providers. Such marketplace is where things and traders have presence, constituting a rich interaction space [4]. The agreed regulations operating in the space structure the interactions between the different contributors. We borrow the metaphor from [5].

Believable electronic marketplaces are perceived as “marketplaces that are alive and engaging, marketplaces where people are. Further we discuss what constitutes the believability of an electronic marketplace.

A. Believability

The notion of believability originated in the creation of human-like computerised agents – virtual humans, which is an active area of research at the crisp of computer science and psychology. It attempts to model the full richness of interactions including natural language communication, gestures, and emotional expression, as well as the cognitive apparatus that underlies these capabilities [6]. Most virtual human research has focused on the cognitive behaviour on the source side of the interaction [7] with a recent shift towards the “recipient” [8].

Earlier research in believability has been heavily influenced by the Carnegie-Mellon set of requirements for believable agents, which is based on research in drama and story telling [9]. These include personality, self-motivation, change, social relationships, and “illusion of life”. Personality infuses everything that a character does - their behaviour, style, “thought”, “emotion”, e.g. their unique ways of doing things. Self-motivation assumes that agents have their own internal drives and desires which they pursue whether or not others are interacting with them, and they demonstrate their motivation. Change implies that characters change with time, in a manner consistent with their personality. Behaviour of agents and interactions between them should be in a manner consistent with their social relationships (in turn, these relationships change as a result of the interaction). “Illusion of life” is used as a label for a collection of features such as: pursuing multiple, simultaneous goals and actions, having elements of broad capabilities (e.g. movement, perception, memory, language), and reacting quickly to stimuli in the

environment. In this sense believable does not necessarily mean realistic.

B. *Believability of electronic markets*

In order to address believability aspects of electronic trading, our work extends Carnegie-Mellon set of features with attributes that relate believability of electronic trading to regulations, processes, spaces, interactions and narrative. We discuss briefly the issues in each of these categories.

Regulations: Norms are part of interactions between trading partners. Collectively they constitute a complex, structured, regulatory system that should be consistent. In a believable trading environment, in addition to compliance with regulations, some times there could be some modifications based on mutual agreements. Background details to the operationalisation of norms in 3D virtual spaces are considered in [10].

Processes: In terms of believability, the structure of the business processes in electronic markets define the narrative of the marketplace. Market players operate in the context of the process structures under the constraints of the regulatory framework.

Spaces: Humans are embodied in space in all their behaviour. They inhabit and operate in it; rely on and use various cues related to space, like pointing and referring to areas of and things in it (for more details see the first two chapters in [11]). This is an essential factor driving the technological conquest for moving us from being on the Internet to gradually being in the Internet space, i.e. towards what is labelled 3D Internet [12]. The evolution relies on several technologies that enable primarily perceptual immersion, including virtual worlds and immersive access to digital content [13]). In terms of believability of electronic environments, the virtual space is essential part of what constitutes an *intelligent environment populated with intelligent artefacts*. Intuitively, to be believable electronic markets should have arrangement of their virtual spaces that are aligned with the business processes in them.

Interactions: Believability of interactions goes beyond traditional focus on modeling the visual presence and communication style (e.g. short vs long utterances, usage of specific vocabulary) of the agents. As a result of their capability to dig out and paste together various pieces of useful information, traders usually are informed to a different extent about various aspects of the deals they are pursuing. When necessary they rely on various relations with other traders that they have established over the time. Their decisions are result of the mix of being rational, informed, impulsive, and the ability to influence others and cope with the influences from others. All these nuances impact the richness of market interactions, hence, must be taken into account when considering the believability of interactions in electronic markets.

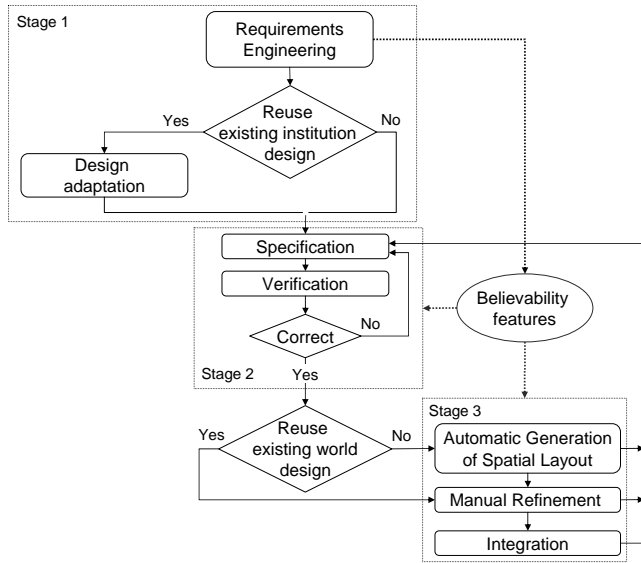
This paper presents the core technologies developed to address these issues. It describes the underlying theoretical and practical solutions of demonstrable prototype of a Web-based electronic marketplace. Section II considers the methodology and mechanisms that enable believable regulated virtual spaces. Section III presents the machinery that enables market players to act believably in the uncertain, diverse and very dynamic environment of Web-based electronic marketplaces. As the machinery, based on the information-based agency [14], relies on information, Section IV discusses the mechanisms of information requesting, mining and delivery as process, without the details of the mining agents, which were presented in our earlier works.

II. BELIEVABLE REGULATED VIRTUAL SPACES

In order to address regulatory requirements that operate in real world research in multiagent systems adapted social science theories and concepts like norms [15]. Normative multiagent systems relate agent theory and the social sciences such as sociology, philosophy, economics, and legal science. The electronic institutions (EI) methodology and technology for normative multi-agent systems (MAS), developed by IIIA CSIC in Barcelona [16], elegantly formalises and implements the institutional approach for MAS. The work on virtual institutions (VI) [10] developed the institutional approach further for inhabited 2D and 3D spatial environments, including virtual worlds. The VI concept and development methodology [10] extends the EI approach, enabling the implementation of institutional commitments that ensure rich and reliable interaction between embodied entities - avatars, whether they are driven by autonomous agents or humans. Central to the implementation of institutional behavioural norms in the EI methodology is the notion of performative structure which formalises processes in terms of scenes, agent roles and communication language. A set of (business) processes are modeled as a discrete collection of interlinked ordered scenes. The involvement of participants (agents) in these processes is modeled through a set of roles, where roles are related to the scenes by the set of participation rights (constraints) for each role in respective scenes, including the subset of the language that can be used in each scene. The later defines the set of permitted dialogues [16]. The VI concept [10] enables the institutionalisation of a virtual world with respect to a performative structure in terms of (i) the spatial layout of the virtual world that reflects the performative structure of the (business) processes; (ii) the objects and avatars and their behaviour within the institution with respect to their roles; and (iii) the rich interaction based on natural language and embodiment of humans and software agents in the institutionalised world.

We extend the VI methodology [10] to embed believability components. Figure 1 presents the extended methodology with the steps grouped into three stages. Stage 1

Figure 1. Extending the VI steps described in [10] for designing believable electronic markets

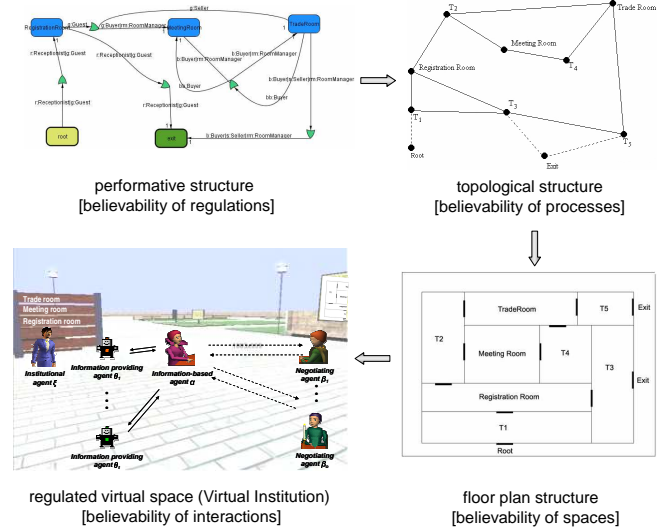


considers the initial requirements engineering of the institutional environment, based on the type of the market in consideration. This stage can generate specific requirements towards the believability of the environment. For instance, if the electronic market is a type of supermarket, then a believability requirement may be the identical arrangement of the spatial layout of the shelves and the presentation of the goods in the same order as in the respective physical shop, emulating what customers are used to. In a property market, the software agent that acts on behalf of a property trader may need to look like "being well informed" and to have the ability to deal with the new information that supersedes the existing information. The set of believability requirements usually translates into elements of the specification of the performative structure, feeds into Stage 2 and propagates further into the layout of the virtual institution in Stage 3. Requirements that are mostly related to the visual presence, like the style of movement (e.g. trajectories, gestures) as well as believability through graphical appearance are considered directly at Stage 3 and may require refinement of the performative structure.

Figure 2 presents a high level view of how the VI technology creates a regulated virtual space on the Web, including the respective believability components as discussed in Section I-B. The topology of the space is extracted from the performative structure and then converted into an initial set of spatial envelopes optimally packed in a bounded institutional space.

The technological architecture of virtual institutions that supports believable electronic markets is shown in Figure 3. The Normative Institutional Layer, specified by the steps in

Figure 2. Enabling the dimensions of believability in electronic markets during the generation of respective virtual institution



"Stage 1" and Stage 2 in 1) takes care for the functioning of the institution as a normative multi-agent system and relies on the EIDE platform [16]. An institutional governor agent is associated with each 'player' G_1, \dots, G_n (whether G_i is a human, a software agent, or another institution). Together with the regulatory mechanism and the execution state of the virtual institution it ensures that the player operates according to the regulatory protocols in each step of the business process. The Intercommunication layer and the Translation Layer enable the causal connection between the institutional infrastructure and the 3D Institutional layer, transforming the actions in the 3D Institutional Layer into messages in the language of the Normative Institution Layer and vice versa. The role of the Translation Layer is to process interactive 3D content, compliant with the X3D standard [17] and translate it to different virtual world platforms (currently - SecondLife [18]).

III. BELIEVABLE MARKET PLAYERS

A. Acting Believably

It is one thing for an agent to appear to be believable, or to move in a believable way, but another for it to interact to *interact* in a believable way. For an agent's utterances to be believable it must act in a way that demonstrates that it understands:

- the significance of each of its utterances *to the observer*
- that on-going interaction are seen as relationships with the other agents that carry implied *social obligations* to act appropriately
- what it should *not* do

A formal model is described that addressed these issues that is based on the observation that an agents beliefs and

Figure 3. The extended architecture of virtual institution technology

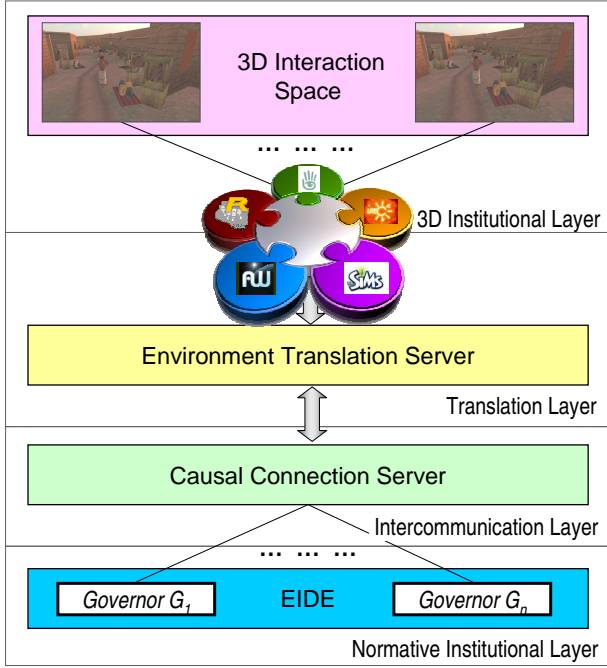
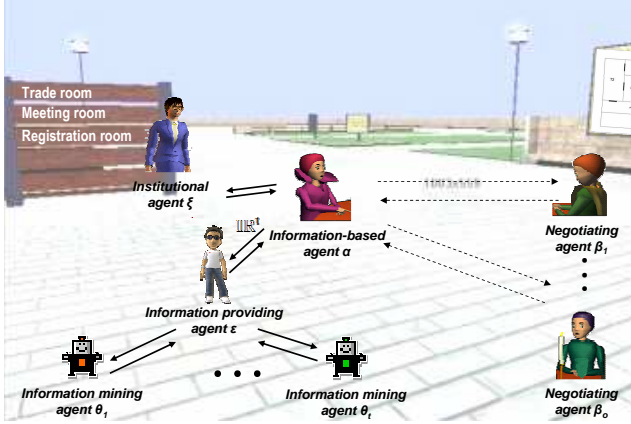


Figure 4. Main types of players and their interaction paths



understanding are necessarily uncertain.

1) *Foundation: Uncertain Beliefs:* As illustrated in Figure 4, this discussion is from the point of view of an information-based agent α in a multiagent system where α interacts with negotiating agents, β_i , information mining agents, θ_j via information providing agent ϵ , and an *institutional agent*, ξ , that represents the institution where we assume that all interactions happen [19]: $\{\alpha, \beta_1, \dots, \beta_o, \xi, \epsilon, \theta_1, \dots, \theta_t\}$.

A pair of agents interact by passing messages to each other. We assume that they share a common ontology and that their interactions are organised into dialogues, where a *dialogue* is a finite sequence of inter-related utterances. A *commitment* is a consequence of an utterance by an agent

that contains a promise that the world will be in some state in the future. A *contract* is a pair of commitments exchanged between a pair¹ of agents. The set of all dialogues between two agents up to the present is their *relationship*.

We represent commitments using conditional probabilities, $\mathbb{P}(\varphi'|\varphi)$. If φ is a commitment that is expected to be fulfilled to some degree, and φ' the corresponding subsequent enactment then $\mathbb{P}(\varphi'|\varphi)$ is the probability that φ' will be observed given that φ was expected. For example, if φ is a commitment signed by β then the conditional probability, $\mathbb{P}(\varphi'|\varphi)$, is an estimate of α 's expectation of β 's execution of that commitment, and the uncertainty in β 's execution of his commitments is the entropy $\mathbb{H}(\varphi'|\varphi)$.

In a multiagent system it is natural to measure the uncertainty of a random variable in terms of the cost, in some sense, of communicating the true value of it from one agent to another. One such sense is the lower bound on the number of binary questions that are always guaranteed to discover the true value of a random variable, X ; this is given by the *entropy*, $\mathbb{H}(X) = \sum_i -p_i \log p_i$, where the p_i are values of the probability mass function for X , [20].

Agent α 's world model, \mathcal{M}^t , at time t is a set of random variables, $\mathcal{M}^t = \{X_i, \dots, X_n\}$ each representing an aspect of the world that α is interested in. In the absence of new information the integrity of \mathcal{M}^t should decrease in time. α may have background knowledge concerning the expected probability mass function for a variable X as $t \rightarrow \infty$. Such background knowledge is represented as a *decay limit distribution*. Given a probability mass function, $\mathbb{P}(X_i)$, for variable X_i , and a decay limit distribution $\mathbb{D}(X_i)$:

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

, where d_i is the *decay function* for the X_i satisfying the property that $\lim_{t \rightarrow \infty} \mathbb{P}^t(X_i) = \mathbb{D}(X_i)$. Either the decay function or the decay limit distribution could also be a function of time: d_i^t and $\mathbb{D}^t(X_i)$ [14]. For simplicity we use a linear decay function; $\mathbb{P}(X_i)$ decays by:

$$\mathbb{P}^{t+1}(X_i) = \lambda \times \mathbb{D}(X_i) + (1 - \lambda) \times \mathbb{P}^t(X_i) \quad (1)$$

for some constant λ : $0 < \lambda < 1$, where λ is the *decay rate*.

Suppose that α receives some new information in an utterance u from agent β at time t where u states that something is so with some probability z , and suppose that α attaches an epistemic belief $\mathbb{R}^t(\alpha, \beta, u)$ to u — this probability reflects α 's level of personal *caution*. [14] gives a method for estimating $\mathbb{R}^t(\alpha, \beta, u)$ on the basis of observation. [14] also describes how 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 $J_s^{X_i}(u)$ is a set of linear constraints on the posterior distribution for X_i . Let $\mathbb{P}^t(X_{i,(u)})$ be the distribution with minimum

¹Sets of commitments between more than two agents are beyond the scope of the paper.

relative entropy with respect to the prior $\mathbb{P}^t(X_i)$ subject to the constraints $J_s^{X_i}(u)$. Then let $\mathbb{P}^t(\bar{X}_{i,(u)}) =$

$$\mathbb{R}^t(\alpha, \beta, u) \times \mathbb{P}^t(X_{i,(u)}) + (1 - \mathbb{R}^t(\alpha, \beta, u)) \times \mathbb{P}^t(X_{i,(u)})$$

The posterior $\mathbb{P}^{t+1}(X_i)$ is $\mathbb{P}^t(\bar{X}_{i,(u)})$ as long as $\mathbb{P}^t(\bar{X}_{i,(u)})$ is “more interesting²” than $\mathbb{P}^t(X_i)$, otherwise it is the prior. This process takes account of both the belief $\mathbb{R}^t(\alpha, \beta, u)$, and the probability z that will be represented in $J_s^{X_i}(u)$. The update functions, J , are the bridge between the illocutionary language and the agent’s world model.

In his 1957 paper [21], E.T. Jaynes wrote: “Information theory ... leads to a type of statistical inference which is called the maximum entropy estimate. It is the least biased estimate possible on the given information; i.e., it is maximally noncommittal with regard to missing information.” Entropy-based inference is a form of Bayesian inference that is convenient when the data is sparse [22] and encapsulates common-sense reasoning [23]. It has three difficulties. First, it assumes that what the agent knows is “the sum total of the agent’s knowledge, it is not a summary of the agent’s knowledge, it is all there is” [23] — this assumption is referred to as *Watt’s Assumption* [24]. Second, it may only be applied to a consistent set of beliefs. Third, its knowledge base is expressed in first-order logic, and so probability distributions are built over finite sample spaces. The way in which the sample space is chosen will affect the inferences drawn — this is referred to as *representation dependence* [25]. Despite these difficulties, maximum entropy inference is an elegant formulation of common sense reasoning that is useful when data is sparse.

2) *The significance of utterances*: If β passes an utterance to α , α evaluates this act in two ways. First, it is valued for the strategic significance of the information that it contains, precisely it is measured as the expected increase in utility that α expects to enjoy given that it has the information — this is the *utilitarian measure*. Second, it is valued because the sending agent *was prepared to divulge* the information in the utterance, precisely it is measured as the decrease in uncertainty that the receiving agent has over the sending agent’s private information — this is the *information measure*. All utterances received are qualified by α with a belief probability as described in Section III-A1.

From α ’s point of view, β ’s *private information* is everything that β knows and that α does not know with certainty. Due to the persistent effect of Equation 1, this will include much of what β knows.

An agent may wish to decide which action, $\{a_i\}$, to take where the payoff depends on which state, $\{s_j\}$, the world is in when the action is taken (possibly in the future). The payoff, \vec{v}_i , from taking action a_i is a vector where v_{ij} will be the payoff from taking action a_i and the state of the world is s_j . Let \vec{p} be the probability mass function of a random

variable representing the prior expectation about the state of the world when the action is taken. Then the *expected monetary value* gained by choosing action a_i is $m_i = \vec{p} \cdot \vec{v}_i$.

Armed with this information suppose that the agent applies some decision criterion, c , to decide what to do — perhaps c will choose the action with the greatest expected payoff: $\arg \max_i \vec{p} \cdot \vec{v}_i$. Now suppose that the agent receives information in an utterance u that enables him to refine his expectation of the state of the world when the action is to be taken ($\vec{p}|u$), and that he applies the same criterion c . Then one *utilitarian value* of utterance u to criterion c is the difference between the payoffs of the respective outcomes. For each state of the world s_j let $b_j = \max_i v_{ij}$ i.e. b_j is the ‘best’ action that the agent can take if the state of the world is s_j then the *expected value of perfect information* is

$$\vec{p} \cdot \vec{b} - \max_i \vec{p} \cdot \vec{v}_i$$

this is an upper limit on the total value of all possible utterances with respect to the application of criterion c .

Utilitarian measures of information are expressed in terms of: if you know information x when applying criterion y to determine which action to perform then you will gain utility z over not knowing x [26]. That is, they are defined in the context of some decision making act — they do not place an intrinsic value on information. So if an agent learns x at time t and is unaware of what future decisions he will make that will benefit from knowing x , then he will be unable to value x until he knows what those future decisions are. But, by the time he is aware of all of those decisions it may not be possible to reconstruct with certainty how he and the other agents would have behaved if he had *not* known x at time t . In summary, it is only possible to attach an intrinsic utilitarian value to information when the future decisions that are relevant to it are known.

We have described the value *gained* by acquiring information, we now consider the value *lost* by an agent’s private information becoming public knowledge — that is, known to all agents in the system. Once information becomes public knowledge it has no tradable value until the integrity of the public’s belief of it decays in time.

Utilitarian measures of information may be used when all the relevant future decisions are either known with certainty or a probability distribution expressing their likeliness to occur is known.

α ’s world model, \mathcal{M}^t , is a set of probability distributions. If at time t , α receives an utterance u that may alter this world model (as described in Section III-A1) then the (Shannon) *information* in u with respect to the distributions in \mathcal{M}^t is: $\mathbb{I}(u) = \mathbb{H}(\mathcal{M}^t) - \mathbb{H}(\mathcal{M}^{t+1})$. Let $\mathcal{N}^t \subseteq \mathcal{M}^t$ be α ’s model of agent β . If β sends the utterance u to α then the *information* about β within u is: $\mathbb{H}(\mathcal{N}^t) - \mathbb{H}(\mathcal{N}^{t+1})$. We give structure to the measurement of information using an *illocutionary framework* to categorise utterances, and an *ontology*.

²One simple criterion is: $\mathbb{H}(\mathcal{M}^{t+1}) < \mathbb{H}(\mathcal{M}^t)$.

The illocutionary framework will depend on the nature of the interactions between the agents. The LOGIC framework for argumentative negotiation [27] is based on five categories: Legitimacy of the arguments, Options i.e. deals that are acceptable, Goals i.e. motivation for the negotiation, Independence i.e. outside options, and Commitments that the agent has including its assets. The LOGIC framework contains two models per agent: first α 's model of β 's private information, and second, α 's model of the private information that β has about α . Generally we assume that α has an illocutionary framework \mathcal{F} and a categorising function $v : U \rightarrow \mathcal{P}(\mathcal{F})$ where U is the set of utterances. The power set, $\mathcal{P}(\mathcal{F})$, is required as some utterances belong to multiple categories. For example, in the LOGIC framework the utterance “I will not pay more for Protos³ than the price that John charges” is categorised as both Option and Independence.

We assume an ontology, and \mathcal{O} denotes its concepts that are organised in an *is-a* hierarchy.⁴ δ measures the semantic distance between two concepts c_1 and c_2 , for example [28]:

$$\delta(c_1, c_2) = e^{-\kappa_1 l} \cdot \frac{e^{\kappa_2 h} - e^{-\kappa_2 h}}{e^{\kappa_2 h} + e^{-\kappa_2 h}}$$

where l is the shortest path between the concepts, h is the depth of the deepest concept subsuming both concepts, and κ_1 and κ_2 are parameters scaling the contribution of shortest path length and depth respectively.

3) *Acting to Respect Social Obligations*: In [27] two central concepts are used to describe relationships and dialogues between a pair of agents. These are *intimacy* — degree of closeness, and *balance* — degree of fairness. Both of these concepts are summary measures of relationships and dialogues, and are expressed in the LOGIC framework as 5×2 matrices.

More generally, the intimacy of α 's relationship with β_i , I_i^t , measures the amount that α knows about β_i 's private information and is represented as real numeric values over $\mathcal{G} = \mathcal{F} \times \mathcal{O}$. Suppose α receives utterance u from β_i and that category $f \in v(u)$. For any concept $c \in \mathcal{O}$, define $\Delta(u, c) = \max_{c' \in u} \delta(c', c)$. Denote the value of I_i^t in position (f, c) by $I_{i(f,c)}^t$ then:

$$I_{i(f,c)}^t = \rho \times I_{i(f,c)}^{t-1} + (1 - \rho) \times \mathbb{I}(u) \times \Delta(u, c)$$

for any c , where ρ is the discount rate. The *balance* of α 's relationship with β_i , B_i^t , is the element by element numeric difference of I_i^t and α 's estimate of β_i 's intimacy on α .

[29] describes measures of: *trust* (in the execution of contracts), *honour* (validity of argumentation), and *reliability* (of information). The execution of contracts, soundness of argumentation and correctness of information are all

³A fine wine from the ‘Ribera del Duero’ region, Spain.

⁴A simplified way of understanding an utterance u is as a set of concepts in \mathcal{O} , that is $u = \{c_i \mid c_i \in \mathcal{O}\}$.

represented as conditional probabilities $\mathbb{P}(\varphi'|\varphi)$ where φ is an expectation of what may occur, and φ' is the subsequent observation of what does occur.

[14] describes a single computational framework for these three measures that summarise α 's observations of β 's behaviour. One of these summary measures is:

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

where the “1” is an arbitrarily chosen constant being the maximum value that this measure may have, and $\mathbb{P}_I^t(\varphi'|\varphi, e)$ is a distribution of enactments that represent α 's “ideal” in the sense that it is the best that α could reasonably expect to happen in *context* e . If α repeatedly observes φ' then the amount of information that those observations convey about the associated commitments, φ , is the *mutual information*: $\mathbb{I}(\varphi'; \varphi) = \mathbb{H}(\varphi') - \mathbb{H}(\varphi'|\varphi)$, this measures the mutual dependence of the two variables, where $\mathbb{I}(\varphi'; \varphi) = \mathbb{I}(\varphi; \varphi')$.

These summary measures are all abstracted using the ontology; for example, “What is my trust of John for the supply of red wine?”. These measures are also used to summarise the information in some of the categories in the illocutionary framework. For example, if these measures are used to summarise estimates $\mathbb{P}^t(\varphi'|\varphi)$ where φ is a deep motivation of β 's (i.e. a Goal), or a summary of β 's financial situation (i.e. a Commitment) then this contributes to a sense of trust at a deep social level.

4) *Knowing what not to do*: [27] advocates the controlled revelation of information as a way of managing the intensity of relationships. In Section III-A2 we noted that information that becomes public knowledge is worthless, and so respect of confidentiality is vital to maintaining the value of revealed private information. We have not yet described how to measure the extent to which one agent respects the confidentiality of another agent's information — that is, the strength of belief that another agent will respect the confidentiality of my information: both by not passing it on, and by not using it so as to disadvantage me.

Consider the motivating example, α sells a case of Protos to β at cost, and asks β to treat the deal in confidence. Moments later another agent β' asks α to quote on a case of Protos — α might then reasonably increase his belief in the proposition that β had spoken to β' . Suppose further that α quotes β' a fair market price for the Protos and that β' rejects the offer — α may decide to further increase this belief. Moments later β offers to purchase another case of Protos for the same cost. α may then believe that β may have struck a deal with β' over the possibility of a cheap case of Protos.

Confidentiality is the mirror image of trust, honour and reliability that are all built by an agent “doing the right thing” — respect for confidentiality is built by an agent *not* doing the wrong thing. As human experience shows, validating respect for confidentiality is a tricky business. One

proactive ploy is to start a false rumour (e.g. “My wife is a matador.”) and to observe how it spreads. The following reactive approach builds on the Protos example above.

An agent will know when it passes confidential information to another, and it is reasonable to assume that the significance of the act of passing it on decreases in time. In this simple model we do not attempt to value the information passed as in Section III-A2. We simply note the amount of confidential information passed and observe any indications of a breach of confidence.

If α sends utterance u to β “in confidence”, then u is categorised as f as described in Section III-A2. C_i^t measures the amount of confidential information that α passes to β_i in a similar way to the intimacy measure I_i^t described in Section III-A2:

$$C_{i(f,c)}^t = \rho \times C_{i(f,c)}^{t-1} + (1 - \rho) \times \Delta(u, c)$$

for any c where ρ is the discount rate; if no information is passed at time t then:

$$C_{i(f,c)}^t = \rho \times C_{i(f,c)}^{t-1}$$

C_i^t represents the time-discounted amount of confidential information passed in the various categories.

α constructs a companion framework to C_i^t , L_i^t is an estimate of the amount of information leaked by β_i represented in \mathcal{G} . Having confided u in β_i , α designs update functions J_u^L for the L_i^t . In the absence of evidence imported by the J_u^L functions, each value in L_i^t decays by:

$$L_{i(f,c)}^t = \xi \times L_{i(f,c)}^{t-1}$$

where ξ is in $[0, 1]$ and probably close to 1. The J_u^L functions scan every observable utterance, u' , from each agent β' for evidence of leaking the information u , $J_u^L(u') = \mathbb{P}(\beta' \text{ knows } u \mid u' \text{ is observed})$. As previously:

$$L_{i(f,c)}^t = \xi \times L_{i(f,c)}^{t-1} + (1 - \xi) \times J_u^L(u') \times \Delta(u, c)$$


for any c .

This simple model estimates C_i^t the amount of confidential information passed, and L_i^t the amount of presumed leaked, confidential information represented over \mathcal{G} . As with most things that information-based agents do, the ‘magic’ is in the specification of the J_u^L functions. A more exotic model would estimate “who trusts who more than who with what information” — this is what we have elsewhere referred to as a *trust network*. The feasibility of modelling a trust network depends substantially on how much detail each agent can observe in the interactions between other agents.

IV. BELIEVABLY INFORMED: INFORMATION MINING AND DELIVERY

Being (well) informed affects the decision making of α , hence, it is another key component of believability discussed in this paper. In the current model information u that α

Figure 5. Example of a tree-structured domain of a taxonomic variable



receives from β , includes fact ϕ with probability $z = \mathbb{P}(\phi)$. When operating on the Web, a more sophisticated estimate of the epistemic belief $\mathbb{R}^t(\alpha, \beta, u)$ can accommodate related information from Web sources. Technically, α sends a request for information $\mathbb{I}\mathbb{R}^t = \{\phi, \mathcal{O}^\Phi, \mathcal{L}^t, q\}$, which includes the fact ϕ , the ontology \mathcal{O}^Φ , level of agreement \mathcal{L}^t between collected information about ϕ from different external sources, and a parameter q which tells information providing agent θ_j the representation form of the information. From information mining point of view \mathcal{O}^Φ can be viewed as a constraint on the tree-structured domain of a taxonomic variable Φ . This limits the size of the ontology with which the information providing agent operates. For instance, α can request from ϵ information on the standing of Protos wine, mentioned in Section III-A2, among other Spanish red wines that can be formulated in terms of “identify the top three wines in the region of ‘Ribera del Duero and the wines closest to those wines from another region’”. ϵ decomposes such request into information mining tasks (e.g. “identify top three red wines in a region”, “cluster the closest ranked wines from two regions”) and pass respective requests to $\{\theta_1, \dots, \theta_t\}$. A view of the results obtained by a θ_j information mining agent in our example is shown in figure 5. In order to combine it with similar results delivered by $\{\theta_1, \dots, \theta_t\}$, ϵ uses De Carvalho’s agreement/disagreement indices [30] - if the value is acceptable with respect to α ’s required level of agreement \mathcal{L}^t then the information is passed to α translated in the form specified by q (currently in probabilistic form), where it is used to refine epistemic belief $\mathbb{R}^t(\alpha, \beta, u)$.

V. CONCLUSIONS

The evolution of the electronic market places is an intrinsic part of the evolution of the Web, hence, it will be essential in Web 3.0 technology. When Web 2.0 is centered around humans engagement, interaction and sharing, we view the forthcoming Web 3.0 to be about placing humans “within” an intelligently behaving Web. Consequently, the notion of believability, discussed in this paper, is essential to Web 3.0. Central to this notion is believable agent behaviour, including the smart ways of gaining advantage from being well informed and the ability to utilise relevant information. We also discussed the enabling technology. At the end we would like to sum up believability using Gary Kasparov’s quotation “I sensed an alien intelligence in the program.” after the 1997 defeat of the world chess champion by the computer program Deep Blue II (as quoted in [31]).

REFERENCES

- [1] J. Mäkiö, I. Weber, and C. Weinhardt, "Electronic negotiations — a generic approach with action systems," in *E-Commerce and Web Technologies*, ser. Lecture Notes in Computer Science, K. Bauknecht, M. Bichler, and B. Pröll, Eds., vol. 3182. Springer Verlag, 2004, pp. 135–143.
- [2] J. Bakos, "A strategic analysis of electronic marketplaces," *MIS Quarterly*, vol. 15, no. 3, pp. 295–310, 1991.
- [3] C. Voudouris, G. Owusu, R. Dorne, and D. Lesaint, *Service Chain Management: Technology Innovation for the Service Business*. Berlin Heidelberg: Springer-Verlag, 2008.
- [4] J. Debenham and S. Simoff, "An e-Market Framework for Informed Trading," in *proceedings 15th International World Wide Web Conference, WWW-2006*, L. Carr, D. D. Roure, A. Iyengar, C. Goble, and M. Dahlin, Eds., Edinburgh, Scotland, May 2006.
- [5] V. Smith, "Markets, institutions and experiments," in *Encyclopedia of Cognitive Science*, L. Nadel, Ed. New York, NY: Nature Publishing Group, 2003.
- [6] N. Magnenat-Thalmann, H. Kim, A. Egges, and S. Garchery, "Believability and interaction in virtual worlds," in *Proceedings of the 11th International Multimedia Modelling Conference (MMM'05)*, 2005, pp. 2–9.
- [7] B. Tomlinson, M. L. Yau, and E. Baumer, "Believable agents: Embodied mobile agents," in *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems AAMAS'06*, Hakodate, Japan, 2006, pp. 969–976.
- [8] R. M. Maatman, J. Gratch, and S. Marsella, "Natural behavior of a listening agent," in *Intelligent Virtual Agents*, ser. Lecture Notes in Computer Science, T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, and T. Rist, Eds. Heidelberg: Springer-Verlag, 2005, pp. 25–36.
- [9] A. B. Loyall, "Believable agents: Building interactive personalities," Ph.D. dissertation, Department of Computer Science, Carnegie Mellon University, Pittsburgh, 1997.
- [10] A. Bogdanovych, "Virtual institutions," Ph.D. dissertation, Faculty of IT, University of Technology, Sydney, November 2007.
- [11] J. O'Keefe and L. Nadel, *The Hippocampus as a cognitive map*. Oxford University Press, 1978.
- [12] T. Alpcan, C. Bauckhage, and E. Kotsovinos, "Towards 3d internet: Why, what, and how?" in *Proceedings of the International Conference on Cyberworlds, 2007*, 2007, pp. 95–99.
- [13] G. Favalora, "Volumetric 3d displays and application infrastructure," *IEEE Computer*, pp. 37–44, August 2005.
- [14] C. Sierra and J. Debenham, "Information-based agency," in *Proceedings of Twentieth International Joint Conference on Artificial Intelligence IJCAI-07*, Hyderabad, India, January 2007, pp. 1513–1518.
- [15] G. Boella, L. V. der Torre, and H. Verhagen, "Introduction to normative multiagent systems," *Computational Mathematics and Organisational Theory*, vol. 12, pp. 71–79, 2006.
- [16] EIDE, "Electronic institution development environment: <http://e-institutor.iiia.csic.es/>," 2005.
- [17] Web3D Consortium, "X3d international specification standards: <http://www.web3d.org/>," 2009.
- [18] Linden Lab, "Second life: <http://secondlife.com/>," 2009.
- [19] J. L. Arcos, M. Esteva, P. Noriega, J. A. Rodríguez, and C. Sierra, "Environment engineering for multiagent systems," *Journal on Engineering Applications of Artificial Intelligence*, vol. 18, 2005.
- [20] D. MacKay, *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [21] E. Jaynes, "Information theory and statistical mechanics: Part I," *Physical Review*, vol. 106, pp. 620 – 630, 1957.
- [22] A. Golan, G. Judge, and D. Miller, *Maximum Entropy Econometrics: Robust Estimation with Limited Data*, ser. Financial Economics and Quantitative Analysis. John Wiley and Sons, Inc., 1996.
- [23] J. Paris, "Common sense and maximum entropy," *Synthese*, vol. 117, no. 1, pp. 75 – 93, 1999.
- [24] M. Jaeger, "Representation independence of nonmonotonic inference relations," in *Proceedings of KR'96*. Morgan Kaufmann, 1996, pp. 461–472.
- [25] J. Halpern, *Reasoning about Uncertainty*. MIT Press, 2003.
- [26] D. Lawrence, *The Economic Value of Information*. Springer-Verlag, 1999.
- [27] C. Sierra and J. Debenham, "The LOGIC Negotiation Model," in *Proceedings Sixth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2007*, Honolulu, Hawai'i, May 2007, pp. 1026–1033.
- [28] Y. Li, Z. A. Bandar, and D. McLean, "An approach for measuring semantic similarity between words using multiple information sources," *IEEE Transactions on Knowledge and Data Engineering*, vol. 15, no. 4, pp. 871 – 882, July / August 2003.
- [29] C. Sierra and J. Debenham, "Trust and honour in information-based agency," in *Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006*, P. Stone and G. Weiss, Eds. Hakodate, Japan: ACM Press, New York, May 2006, pp. 1225 – 1232.
- [30] F. Esposito, D. Malerba, and V. Tamma, "Dissimilarity measures for symbolic objects," in *Analysis of Symbolic Data*, H. H. Bock and E. Diday, Eds. Heidelberg: Springer Verlag, 2000, pp. 165–196.
- [31] N. Cross, *Designerly Ways of Knowing*. Springer Verlag, 2006.