Ontological Issues in Agent Ownership[1]

(Reasoning About Agent Ownership)

Alexander Yip & Jim Cunningham
{ay1, rjc}@doc.ic.ac.uk
Communicating Agents Group, Department of Computing
Imperial College of Science, Technology and Medicine
South Kensington Campus, London SW7 2AZ

Abstract

Ownership is a common concept underlying many social activities – buying, selling, renting and giving etc. For a software agent to operate in an open environment like Agentcities[2], it may be necessary for it to reason about ownership. This paper provides a brief overview of ownership issues for computer agents and an approach towards a coherent ontological framework for reasoning about agent ownership.

Introduction

A software agent is usually considered to be an autonomous computer process entity acting on the user's behalf. However, a software agent is not a legal person under the law, so it does not have the legal capacity to take on responsibility like a human agent and existing business agency law cannot be directly applied to software agents. The common opinion of legal professionals is that an active software agent is to be treated purely as a tool – property owned by an owner. For future software agents to operate legally (as pseudo agents) under the existing legal system, we would like to give the agents the ability to reason about ownership, so that when an agent makes a deal with another software agent it is able to express the ownership of the objects involved, e.g. the buyer may want to verify that the seller really owns the object. In the following sections we give a brief overview of ownership and exercising ownership in a multi-agent environment, present an ontological framework with elements of a concrete ownership ontology, and discuss how it provides a basis for reasoning about ownership.

A brief overview of Ownership

In this paper, we concentrate on the Anglo-American legal system. In such systems, ownership is a relationship between a legal person (individual, group, corporation, or government) and an object. The object of concern may be corporeal, such as furniture, or completely the creature of law, such as a patent, copyright, or annuity. The most basic meaning of ownership is that one's government or society will help to exclude others from the use or enjoyment of one's possession without one's consent, which may be withheld except at a price. [EB]

---

[1] In "LEA 2003: The Law and Electronic Agents", Proceedings of the 2nd LEA Workshop, University of Edinburgh, UK, 24th June 2003, pp. 113-126

[2] We have in mind proposals such as those from EU Agentcities.RTD project IST-2000-28385

One way to model the legal relationship of ownership is to see it as a set of legal states of affairs influenced by a set of legal actions by legal persons. The core states of affairs may be regarded as "ownership" states, for which familiar words like owner, possession, and ownership can be used to describe subsidiary roles or relationships. A statement such as *person A owns object B* is true if and only if *person A owns object B* is part of the state of affairs of the applicable legal jurisdiction, or more simply, "in law". A legal state exists only at the metaphysical level of the legal domain, which is different from the physical world or even the mental state of a person. Legal states of affairs must satisfy appropriate laws. For example, a person who owns an object has the exclusive rights to that object; in a naïve system we assume no-one else can own that object, *etc*. Nevertheless, at any one time more than one party may exert authority over an ownership. For example, a legal court may take control of an object owned by a guilty party and transfer its ownership to the plaintiff as compensation. In this case, the court exerts authority over the ownership of the possession and authorises the transfer of ownership; technically the original owner may continue to own the object until the transfer is completed.

The ownership of an object is itself a legal object, and can be transferred. Familiar actions to transfer ownership are to "buy", or "sell" an object; these two actions are complementary and only differ in whether the role of the performer of the action is beneficiary (gaining ownership – buyer) or originating (seller). Either action, if legal, will modify ownership, the original owner will lose ownership and grant ownership to the person who acquired the object. In this model there are other actions and states that complement the core state and action, for example a "renting" action grants a temporary "ownership" to a special type of owner, a beneficial owner. Whether an action is legal or not depends on the authority of those who participate in the action. Under normal conditions a legal person has the right to buy an object and thus acquire ownership of an object, but the actual circumstances of an event by which ownership is transferred are circumscribed by many issues of legal authority: does the seller own the object, or represent the owner, *etc*. These are discussed by Gelati et al. [Gelati 2002]

Ownership can also be viewed as a social creation – it exists only because there is a common belief of such concept in the society. A member of a society that recognises ownership has typically been brought up to have the belief that certain objects belong to him/her while other objects belong to others. This belief may be reinforced by the belief that a person who does not respect another person's ownership will be punished by the society. It is difficult to see how ownership could be sustained in an anarchist society where no one believed in ownership or had authority to protection it. Searle [Searle 1995] gives a detailed analysis on the concept of property ownership from the sociological point of view. He used the term "common intentionality" to encapsulate this concept of common belief and the individuals' intention to believe in the belief. In an attempt to unify these two points of view, the legal and the social, one can see the concept of ownership actually has two components – a legal state of "ownership" that affects an owner, a person, and an object; and individual beliefs in the existence of such a state.

Exercising Ownership in an Open Multi-agent environment

Most presently deployed agent systems can be considered closed environments, where a single person or a single company implements and owns the whole system – from the agent platform to the agents deployed on it. In this case, one can say that a single authority is exerted on the agent system where a common belief and intention exists to co-operate and to act to the best interest of the owner. However, in the foreseeable future we will see more interactions between agent systems controlled by different authorities, an example is the open agent network created by the Agentcities project [Agentcities 2002]. The project envisages an open network of multi-agent systems, which support business activities that begin in the agent world but with the end result realised and consumed in the human world. For instance the Agentcities "Evening Organiser" recruits resources like restaurant review agents, restaurant booking agents and cinema representatives from the network to create and provide new services (entertainments for an evening) to be used by human users.

Before such open network becomes an integral part of our daily life, it is only proper to address the issues of ownership. One can foresee that the following questions need to be answered:

If an agent is promoting tickets for a concert, how can the Evening Organiser verify that the agent really is representing the owner of the tickets and has the authority to sell the tickets?

The simplest case for the Evening Organiser is that the same authority owns the Ticket agent and the agent platform is a closed environment. An analogy for this situation is Person A owns both bank account X and Y within the same bank B; a transfer between X and Y is automatically endorsed (subject to availability of funds) since they are both under the same authority (A) and environment (B). However, if the owners of the Evening Organiser and the Ticket Agent are different, so that the agents are under different authorities, the agents may need to communicate this fact to allow further authentication and verification to proceed. A more complex situation arises when the ticket owner for a one-off event hires a Ticket Agent to represent him for this one occasion only. In this case, there are three different authorities: the Evening Organiser agent is owned by person A; the tickets are owned by person B and the Ticket Agent is owned by person C who has a hired relationship with ticket owner B.

If a personal agent approaches an Evening Organiser to organise an evening, how can the personal agent express its own ownership to the Evening Organiser and thus convey the authority it is representing?

In this situation the relationship between a personal agent and its owner entails three different legal relationships: Ownership, Mandate and Representation. [Gelati 2002] The owner owns a software agent in a form of an actual possession where the owner can exert physical control (in the form of starting and stopping a software execution) and has the exclusive right to use and dispose of the agent. The owner also acts as the mandator of the agent since a software agent usually takes its owner's will as its own goal. For example, owner A wants to buy a ticket for a concert and sends a personal agent to do so, then "buying ticket" becomes the goal of the personal agent. The

personal agent also holds a representation relationship with the owner as all its actions can be seen as acting on behalf of the owner.

If something goes wrong during the execution of an evening plan e.g. the personal agent got a confirmation from the Ticket Agent but the connection flight from London to Paris is actually double booked, so that the person cannot get to Paris, who can that person sue for compensation?

As an agent is acting as a representative of the owner, this may provide a simpler solution to the situation, as usually the legal capacity for the act to be performed by the representative is required in the principal and not in the representative. Therefore it is possible to treat the situation as simply as: Person A books a ticket from the Airline B, A receives a confirmation from B, but B actually double booked the same seat.

And how can the Evening Organiser know that the personal agent is not a rogue agent performing a denial of service attack?

Only when the relationships between entities are expressed in a coherent fashion can one perform the necessary security check to confirm or deny the claims. The authors have suggested some possible mechanisms that could be used in software agent environment to verify an ownership claim e.g. code watermarking and function hiding [Yip 2002].

To express these relationships in an open environment in a coherent fashion, an ontological framework is required to provide a model for communication.

An Ontological Framework

Projects like Agentcities are attempting to compile and publish concrete "ontologies" for agent communication on specific domains, including, for example, the lexicon associated with restaurants, or with theatre performances, associated with an evening organiser. These concrete ontologies are lexical taxonomies for utilitarian purpose. They are expressed in special file formats like the DAML-OIL variant of XML, so that the types of lexical items can be modelled using terms familiar to programming system designers, but of uncertain connection with the world at large.

Our objective here is more abstract, although we wish to establish the links with such programming ontologies. We think that these links will be established as the meaning of natural language is better explicated in both logical and metaphysical terms. Without this we are left with a semantic gap between our models of the computer system and our understanding of the human world. In the mean time we wish to emphasise that the notations we choose for presenting concepts of concern are intended as intuitive guides which are not fully explicated here and would be subject to revision if they were.

In some areas there is little conflict between the human world and computer models, at least of we allow computer agents to have identity, physical substance and temporal behaviour, all of which have some reality through the detail of a computer systems. Because a computer agent has physical substance and identity, it too can be an object

in the role of a possession in an ownership relation with a legal person, and so too, legal property. But because such an agent is not a legal person, but an autonomously animated object, its ontological status is reminiscent of a physical animal.

The poverty of conventional temporal models of both language and computing makes the behaviour of computer agents difficult to relate to the human world. To illustrate the sort of problems which arises, consider the ontological status of buying a ticket for a cinema. Depending on perspective, this can be considered an action, an event, an occurrence, or a process, but it is not normally considered a state because its dominant meaning is associated with change of state. So in conventional program models such an "action" may be represented as some sort of executed function, or it may become a communicative act in agent models. In each case the temporal perspective is lost and so too the ability to describe an extended purchasing process. Yet the statement *John is buying a ticket* can obviously be a true declaration about a human occurrence in the right circumstances, and so too the more primitive *John buys a ticket*, at least if we imagine a narrator describing an evening as it progresses. The key point is that for the first to be true when uttered, the second must be true of some event whose duration embraces the times when the first statement is uttered.

Our "solution" to these problems will include the presumption that a declarative statement is in general evaluated on intervals from a point based time domain presumed here to be model by the real numbers. For some indication, see [Cunningham 2003]. The stativeness of an utterance like *John has a ticket* is captured by the observation that on an interval when the statement is true, it will also be homogeneously true on all sub-intervals. We expect an occurrence of this state on some interval immediately after the event *John buys a ticket*. This event is also an occurrence true of a specific, although undetermined, temporal interval. We are not concerned if in a simpler computer model the event is reduced to an action point, but it needs to be durative for the progressive activity *John is buying a ticket* to have temporal interval on which such an utterance could be true, and for a human being to perceive the delay involved in the activity of an agent *John*. When perceived as a processing state, a durative event need not be homogeneous, merely a composition of process steps.

Thus equipped with a framework for durative occurrences as events and states, we can seek to import and extend a suitable structural ontology for roles and relations. Such a framework has been provided by Schneider [2001]. A motivation for Schneider's work was that agent communication lacks a top-level ontology to allow domain ontologies to relate to each other. Schneider's ontological framework is an attempt to unify Strawson's theory of individuals [Strawson 1959] with Parson's theory of events and thematic roles in natural language [Parsons 1990]. Hence it can provide a shared semantics for a multi-agent system involving human and software agents interacting with each other, at least partially, via natural language communication. The idea is that by breaking a sentence down to grammatical components associated with thematic roles, the implicit ontological framework of language will provide a semantically consistent basis for reasoning. The ontological framework presented by Schneider is summarised by figure 1.

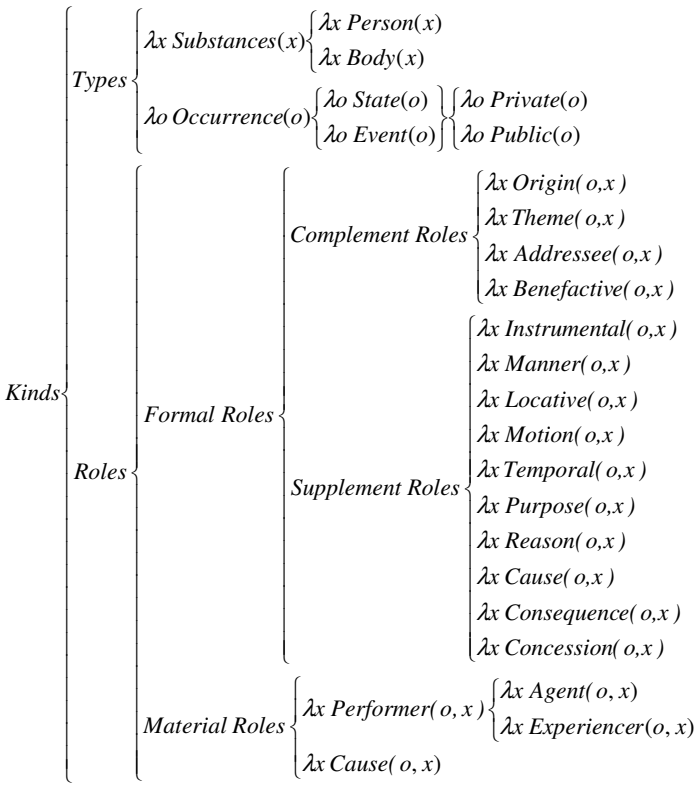

Figure 1: Schneider's Minimal Semantic Ontology Framework

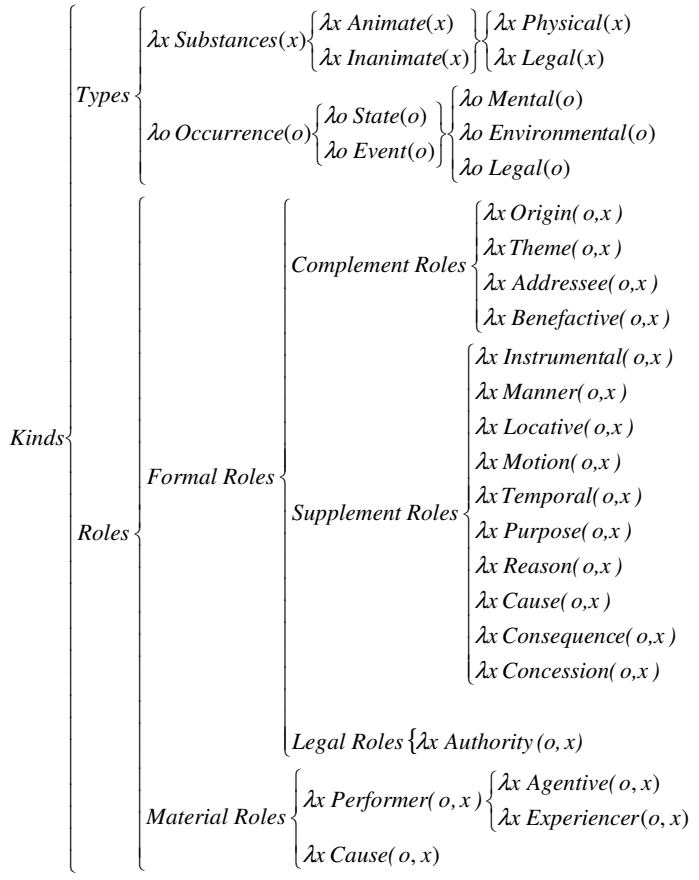

Figure 2: Modified Minimal Semantic Ontology Framework

Some of the notation of figure 1 can only be understood from a perspective of linguistic philosophy. (Lambda notation is used to help explicate the underlying extensions of predicates, whereas contextually superfluous variables and quantifiers for these extensions are suppressed in the intensional interval logics we use for the informal axiomatisation of concrete concepts). Schneider himself found the need to introduce persons as substances capable of agentive roles and distinguish them from objects that are merely experiencers – passive participants in events. But because computer agents are not legal persons we have introduced separate distinctions between the animate and the inanimate. Schneider's terms "private" and "public" are perhaps better read as distinguishing mental from environmental occurrences, although this seems insufficient for our purposes. In figure 2 we have slightly modified the notation to accord with our readings, but moreover, we find that allowing for the legal domain requires a whole new level of sub-categorisation, since as already discussed, legal persons, legal states, and legal roles, exist in a new and distinct metaphysical sense. As a sub-classification, animate and inanimate things can be either legal or physical.

Under this classification, a legal person is a legal animate substance (or *thing*); a dog, or a computer agent is a physical animate thing and a legal right is a legal inanimate thing. So an ownership is also an inanimate legal substance. Our earlier discussion illustrates the distinctions between a legal state, such as a temporal occurrence where a person owns an object, and events which are occurrences that may change the state. But we also recognise that the distinctions between these types of occurrence may depend on a reference interval. In particular, most verbalised events can be perceived on a finer timescale as processing states, in which case the English verbal expression uses the progressive form. A mental state is typically signified by a modal relation of belief, perception, or intent, and the use of such concepts in designing intentional computer agents is well known. A mental event like a decision will alter a human mental state, or a computer agent's internal state. These events are distinguished from environmental events which happen in the world, legal events which affect legal states, and communicative events, which may be considered environmental, but which we speculate are a category worthy of sub-classification. It seems clear that authority in its various manifestations is a legal role which can be distinguished from other formal and material roles, not withstanding recognition that the exercise of such authority is a processing state. A basic semantic classification of lexical elements for a concrete ownership ontology is given in table 1.

Lexical Element	Semantic Category
$\lambda x.person(x)$	(Legal or physical) animate
$\lambda x.jurisdiction(x)$	Legal inanimate
$\lambda x.object(x)$	(Legal or physical) inanimate
$\lambda x.agent(x)$	Physical animate
$\lambda x.money(x)$	Physical inanimate
$\lambda y.ownership(x,y)$	Legal inanimate
$\lambda o.owns(o,x,y)$	Legal state
$\lambda o.transfers(o,x,y,z)$	Event
$\lambda o.buys(o,x,y,z)$	Legal event
$\lambda o.sells(o,x,y,z)$	Legal event
$\lambda o.borrows(o,x,y,z)$	Legal event
$\lambda o.hires(o,x,y,z)$	Legal event
$\lambda o.pays(o,x,y,z)$	Legal event
$\lambda x.authority(o,x)$	Legal role
$\lambda x.represents(o,x,y)$	Legal role
$\lambda x.consideration(o,x)$	Legal role
$\lambda y.delegatee(o,x,y)$	Legal role
$\lambda x.mandator(o,x,y)$	Legal role
$\lambda x.owner(o,x,y)$, $\lambda x.seller$(o,x,y) *etc*.	Derived legal roles

Table 1: A Concrete Ownership Ontology

Reasoning about ownership

Although occurrences like buying and selling have legal significance, their temporal effect is similar to other everyday events which transfer properties between substances. In buying and selling it is the ownership of a legal object which is transferred, but for some consideration. The temporal effect can be logically axiomatised, and ideally explicitly defined with the help of interval operators, for example, using *beforehand*, *afterwards* and *during* to intuitively signify underlying temporal relationships which here we do not define further. There are then many ways of providing formal definitions for mechanised interpretation, just as there are many ways to draft legislation, all can be open to criticism. For simplicity, suppose:

$$buys(x,y,z) \leftrightarrow \ beforehand\ owns(z,y) \wedge \ afterwards\ owns(x,y) \wedge \ during\ pays(x,c,z)$$

Obviously such axiom depends on the definition of *pays*, while a *sells* event has a complementary definition and a more general *transfers* event can be defined. Because the interval of occurrence is implicit but accessible in this style of logic, auxiliary properties can be indicated by adding conjunctions, for instance a property like *jurisdiction*(*j*), helping with scalability of the formal system. (The use of an explicit definition also has certain advantages when dealing with what is usually called the frame problem of qualitative reasoning). We are not seeking here to demonstrate an adequate formalisation, merely remind ourselves that at least the

simpler characteristics of certain classes of everyday legal and commercial activity can be brought within the framework of machine reasoning[3].

The subtle but important differences between buying, borrowing, hiring *etc*. can readily be expressed, albeit naively, if for the purpose we treat ownership as beneficial:

$$borrows(x,y,z) \leftrightarrow beforehand\ owns(z,y) \wedge during\ owns(x,y) \wedge afterwards\ owns(z,y)$$

Similarly, hiring is a borrowing with a consideration:

$$hires(x,y,z) \leftrightarrow beforehand\ owns(z,y) \wedge during\ (owns(x,y) \wedge pays(x,c,z)) \wedge afterwards\ owns(z,y)$$

Perhaps a greater concern for computer agents than the everyday aspects of the law are the adequacy of mechanisms for informing and mandating computer agents. This requires firstly that computer agents can be given coherent computational models of beliefs, goals, obligations, permissions *etc.* both logical and contingent, within in the domain of intended application and its legal framework, but secondly that the range of communicative acts be adequate for delegating responsibility and reporting affairs between human and computer agents. The first of these challenges has been a concern for many agent researchers. It is not a solved problem, and there are no standard solutions, but architectures have been proposed for facets of these challenges, and shown to be feasible, if not ideal e.g. [Barbuceanu]. Whether these architectures can be exploited for adequate reasoning about legal ownership and responsibility is yet to be seen.

We briefly discuss the second issue because it is equally essential, and complementary for reasoning about ownership in its various ramifications for computer agents. Gelati et al., already cited, and earlier work by Jones and Sergot [Jones 1996], use the context of signalling acts to address the formalisation of concepts such as authorisation and delegation in terms of institutional empowerment. Our concern is to ensure that the use of communicative acts in the interface between human and computer agents; and between computer agents, can be of adequate scope and within a coherent ontological framework to achieve similar effect. As pointed out elsewhere [Cunningham 2002], communicative acts can be given a simple temporal axiomatisation in terms only of the effects on the beliefs of agents, so do not depend on sophisticated architectures for realisation unless the semantics of the acts themselves require it. A problem for ownership is that computational realisation of mechanisms associated with authorisation and delegation may indeed be complex.

Our final concern from an ontological perspective is that the position of communicative acts in the ontological framework has not been elaborated. A communicative act could be considered an environmental event, although in the progressive it may also be a processing state. However "Speech Acts" can be considered of sufficient philosophical and practical importance to have their own

[3] This has previously been demonstrated by others, e.g. Kowalski 's British Nationality act publication [Kowalski 1992]

semantic category. Figure 3 is an ontological fragment suggesting an extension to the framework to support communicative acts.

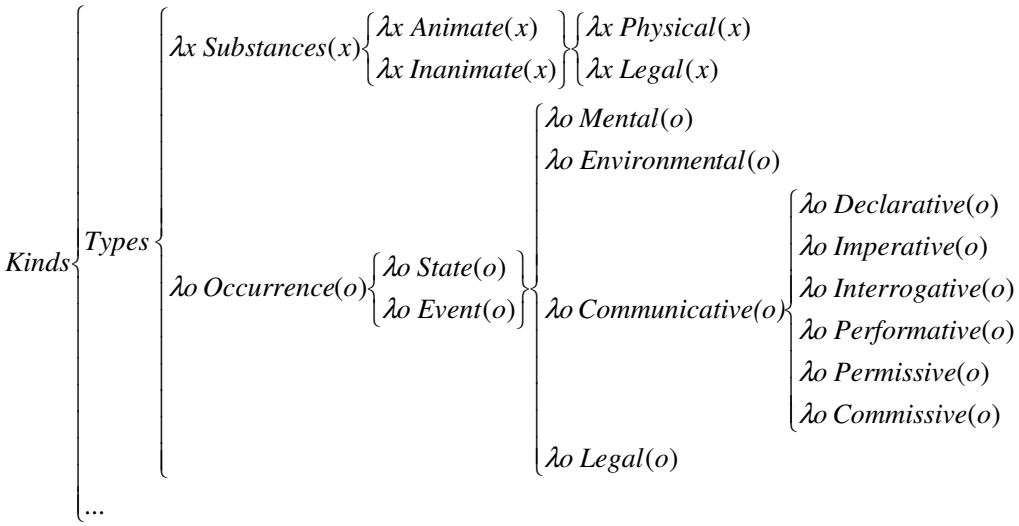

Figure 3: Suggested Extension

Conclusions

Although ownership is legally and sociological complex, key aspects should be sufficiently explicated for adequate realisation by computational agents which may make legal commitments on behalf of humans. For as the range of services that becomes available "electronically" grows, society may need to cope with unforeseen ramifications where ownership and the rights that go with it are challenged. In order to anticipate some of these situations, and have reasoning mechanisms for "intelligent" computer agents; we have sought to clarify the ontological perspective which enables humans to communicate about ownership, and to adjust such a framework for use by electronic agents in an open environment beyond Agentcities.

Acknowledgments

The authors wish to thank Adrian Williams for helpful suggestions on an earlier draft. The research described in this paper is partly supported by the EC project Agentcities.RTD, reference IST-2000-28385. The opinions expressed in this paper are those of the authors and are not necessarily those of the Agentcities.RTD partners.

Bibliography

[Agentcities 2002] EU Agentcities.RTD project IST-2000-28385 http://www.agentcities.net/

[Barbuceanu] Barbuceanu, M., An architecture for agents with obligations, http://citeseer.nj.nec.com/57306.html

[Cohen 1990] Cohen, P.R. and Levesque, H.J., Intention is choice with commitment, *Artificial Intelligence*, 42(2-3)(1990) 213~261

[Cohen 1991] Cohen, P.R. and Levesque, H.J., Teamwork, *Nous 25* (1991) 487-512

[Cunningham 2002] Cunningham, J. 2002, Elements of a temporal semantics for Speech Acts, An unpublished note to FIPA, http://www.doc.ic.ac.uk/~rjc

[Cunningham 2003] Cunningham, J. 2003 Temporal Agents, In *Mistakes of Reason: in honour of John Woods*, Eds. Peacock,K. & Irvine, A., University of Toronto Press. *To Appear*

[EB] "Ownership", *The New Encyclopaedia Britannica (15th Ed.)* Encyclopaedia Britannica, Inc., Chicago

[Gelati 2002] Gelati, J., Sartor, G., Jones, A.J.I., Sergot, M., 2002. Normative relations between agents and humans: Representation, mandate and ownership. *EU IST-1999-10298 ALFEBIITE project Deliverable*.

[Jones 1996] Jones, A.J.I., Sergot, M.J., June 1996. A Formal Characterisation of Institutionalised Power. *Journal of the IGPL* 4(3):429--445

[Kowalski 1992] Kowalski, R., 1992, Legislation as Logic Programs, *Logic Programming in Action* (eds. G. Comyn , N. E. Fuchs, M. J. Ratcliffe), Springer- Verlag, pages 203-230 (1992)

[Parsons 1990] Parsons, T., 1990. *Events in the Semantics of English: A Study in Sub-atomic Semantics*. MIT Press, Cambridge/MA.

[Rao 1995] Rao, A.S., Georgeff, M., BID Agents from Theory to Practice, *Technical Note 56, AAII*, April 1995

[Rao 1991] Rao, A.S. and George, M. P., Modeling rational agents within a BDI-architecture, *Proc. KR'91*, pp. 473-484. Morgan Kaufmann, 1991

[Schneider 2001] Schneider, L.N., 2001. Naive Metaphysics: Merging Strawson's Theory of Individuals with Parson's Theory of Thematic Roles as a Basis for Multiagent Semantics. Imperial College London, Department of Computing MSc Thesis

[Searle 1990] Searle, J.R., *Collective Intentions and Actions,* Intentions in Communication, P.Cohen, J. Morgan and M.E. Pollack eds. Cambridge, Mass.: Bradford Books, MIT Press 1990

[Searle 1995] Searle, J.R., 1995. *The Construction of Social Reality*, New York: Free Press.

[Strawson 1959] Strawson, P., 1959. *Individuals: An Essay in Descriptive Metaphysics*. Routledge, London

[Strawson 1992] Strawson, P., 1992. *Analysis and Metaphysics*. An Introduction to Philosophy. Oxford University Press, Oxford.

[Yip 2002] Yip, A., Cunningham, J., 2002 Some Issues on Agent Ownership. *In "The Law of Electronic Agents: Selected Revised Papers", Proceedings of LEA Workshop*, CIRSFID, University of Bologna, July 2002, pp. 13-22