Artificial Agents - Personhood in Law and Philosophy

Samir Chopra[1] and **Laurence White**[2]

Abstract. Thinking about how the law might decide whether to extend legal personhood to artificial agents provides a valuable testbed for philosophical theories of mind. Further, philosophical and legal theorising about personhood for artificial agents can be mutually informing. We investigate two case studies, drawing on legal discussions of the status of artificial agents. The first looks at the doctrinal difficulties presented by the contracts entered into by artificial agents. We conclude that it is not necessary or desirable to postulate artificial agents as legal persons in order to account for such contracts. The second looks at the potential for according sophisticated artificial agents with legal personality with attendant constitutional protections similar to those accorded to humans. We investigate the validity of attributes that have been suggested as pointers of personhood, and conclude that they will take their place within a broader matrix of pragmatic, philosophical and extra-legal concepts.

1 Introduction

A recurring philosophical debate concerns when or whether to ascribe personhood to artificial agents. Typically, contributions to this debate involve drawing up a list of necessary and sufficient conditions, which must be met by an artificial agent in order to be classified as a genuine cognizer on par with human beings. Several issues are - unavoidably - conflated in this debate: the ascription of intentionality or conscious phenomenal experience, the possibility of the exercise of free will by - and autonomy for - artificial agents and so on. Unsurprisingly there is disagreement about what such a list of conditions should look like. One view on this situation is that philosophical theorizing about the cognitive status of artificial agents should draw inspiration from legal theorizing - which carries a strong pragmatic flavour - about the status of these agents in our society. Some of these legal arguments would support classifying agents as intelligent beings on par with human beings. Others would not. Conversely, legal arguments could draw upon philosophical arguments in arguing for the ascription of elevated cognitive status to artificial agents. A virtuous circle of complementary theorizing is possible. Theorizing in this area is of crucial importance to *designers* of artificial agents: will the agents designed be autonomous enough to deserve and warrant legal rights?

A legal person is an entity that is the subject of legal rights and obligations. Typically, a legal person has the capacity to sue and be sued, and to hold property, in its own name, although some kinds of entity - notably corporations, children and the mentally incapacitated - may need to act through agents to exercise their legal capacities. Not all legal persons have the same rights and obligations; some rights (e.g., marriage) depend on age. Other rights (e.g., voting) and obligations (such as the liability to be imprisoned) are typically restricted to humans.

Being human is not a necessary condition of being accorded legal personality, an obvious example being the modern business corporation. English admiralty law treats a ship as a legal person capable of being sued in its own right; other legal systems have recognized temples, dead persons, spirits and idols as legal persons[12, 1]. The law has also not, historically, considered being human a sufficient condition of being recognized as a legal person. In Roman law, the paterfamilias or free head of the family was the subject of legal rights and obligations on behalf of his household; his wife and children were only indirectly the subject of legal rights and his slaves were not legal persons at all. In English law before the middle of the 19th century, the married woman was not, for most purposes, accorded separate legal personality from that of her husband. In United States law slaves were considered non-persons. Currently, human foetuses in many jurisdictions are not considered legal persons; all other living human beings are generally accorded legal personality.

Legal scholars have considered the problem of personhood for artificial agents in a number of contexts. This debate has often taken shape as a possible solution to the legal doctrinal problem of accounting for the formation of electronic contracts. In the first case study, we sketch out this doctrinal problem, and some possible solutions, only one of which involves according legal personality to artificial agents. On the basis that it is conceptually possible that the legal system will accord legal personality with civil rights to artificial agents, the second case study examines the factors that the legal system could be expected to take into account in coming to this decision. In what follows, we confine ourselves to doctrines of Anglo-American law. While civil law (based on Roman and Napoleonic law) will differ in details, most of the concepts discussed below will have their analogues in these legal systems[15, 11]. We use 'agent' to denote complex computational systems that could individually - or as part of a multi-agent system - represent human persons or corporations. We use 'operator' rather than 'user' to denote the entity on whose behalf the agent operates, as 'user' can be confused with the person interacting with the agent (e.g., a shopper interacting with a shopping website agent). Lastly, we use 'principal' interchangeably with 'operator', by analogy with the principal of a human agent.

2 The contracting problem

Artificial agents, and the contracts they make, are ubiquitous. Every time we interact with a shopping website we interact with a relatively autonomous interface that queries the operator's database, uses our input to populate the operator's database, and confirms the terms of the transaction. The operator does not exercise direct control over the

[1] Department of Computer and Information Science, Brooklyn College of the City University of New York, Brooklyn, NY 11205. Email: schopra@sci.brooklyn.cuny.edu

[2] Financial Services Authority, 25 The North Colonnade, London, E3 5NZ. Email: laurence.white@fsa.gov.uk. The views expressed do not purport to represent the official position of the FSA.

agent's 'choices', at least until the operator has a chance to confirm or reject the transaction entered into. Right now, websites such as the ebay auction website offer users agent-like functionality, by optionally bidding incrementally up to a specified maximum, allowing the user to 'set and forget' the bidding process. Eventually, agents such as shopbots and pricebots capable of collecting information and engaging in transactions with very limited operator input are envisaged.

Various legal doctrinal difficulties are associated with contracts made by artificial agents[1, 7, 8]:

1. In relation to the requirement that there be two parties involved in contract-making, artificial agents are not considered by current law to be legal persons; therefore, only the buyer and seller can be the relevant parties to the contract.
2. There are therefore difficulties in finding an agreement about terms between the parties where one party is unaware of the terms of the particular contract entered into by its artificial agent.
3. In relation to the requirement that there be an intention to form legal relations between the parties to the contract, a similar issue arises here as with agreement: if the agent's principal is not aware of the particular contract being concluded, how can the required intention be attributed?

2.1 Possible solutions to the contracting problem

There is some disagreement as to the effectiveness of potential solutions to the contracting problem as we move up the sophistication scale of artificial agents. The first three are 'tweaks' - involving minor changes of law, or pointing out that existing law perhaps with minor modifications or relaxations can accommodate the problem. The fourth is more radical and involves treating artificial agents the same as human agents, but without legal personality; the fifth, the most radical, proposes according artificial agents legal personhood.

Most current approaches to the problem of electronic contracting adopt the first potential solution - by treating artificial agents as *mere tools* of their operators, or as mere means of communication. All actions of artificial agents are attributed to the agent's operator, whether or not they are intended, predicted, or mistaken. On this basis, contracts entered into through an artificial agent will always bind the operator - a stricter liability principle than that which applies to human agents and their principals. The 'artificial agent as tool' approach predominates in proposed legislative attempts to deal with electronic contracting[7, 8]. It is not clear whether legal change will be required in order to entrench this approach, since it currently appears so common-sense. However, as the autonomy of agents increases, it will be less realistic to approach agents as mere tools of their operators and as mere means of communication. The limitations of the approach, in cases where it would seem unjust and economically inefficient to burden the principal with losses caused by erratic behaviour of the agent, have not become sufficiently evident because of the limited autonomy displayed by existing artificial agents. There may be a natural limit to the 'means of communication' doctrine too - e.g., where the person receiving the communication no longer can reasonably rely on the communication as having emanated from the purported principal.

The second potential solution, which addresses difficulties (2) and (3), is to deploy the *unilateral offer doctrine* of contract law. Contracts can be formed by a party's unilateral offer addressed to the whole world, together with acceptance - in the form of conduct stipulated in the offer - by the other party. Competitions and terms and conditions of entry to premises are among the most common examples involving unilateral contracts. In a simple sell-side agent example, the user's interaction with the website, in legal terms, can be equated with an interaction with a vending machine, where the contractual terms of particular contracts are not determined by the agent.

Many will see in the unilateral offer doctrine a theory that could justify many electronic contracts. The offer being made to the world is to be bound by contracts made through the artificial agent. That offer is accepted by interaction by a user with the artificial agent. What are the limits of such a doctrine? According to [7], the analysis breaks down when we are dealing with agents that are able to determine contractual terms autonomously - for then the seller cannot be said to have intended the terms of each particular contract. However, we suggest that the law might imply into the offer made to all the world a reasonableness requirement i.e., the limits of the doctrine would be reached when it is unreasonable for the user to believe that the agent's principal would assent to the terms of the contract. This will only be the case where the agent is acting erratically or unpredictably as opposed to merely autonomously. This is a distinction that is relevant to the next solution too.

The third potential solution is to deploy the *objective theory of contractual intention*, which is dominant in United States law. Here, a contract is an obligation attached by the force of law to certain acts of the parties, usually words, which accompany and represent a known intent. The party's assent is not necessary to make a contract; the manifestation of intention to agree, judged according to a standard of reasonableness, is sufficient, and the real but unexpressed state of the first party's mind is irrelevant. It is enough that the other party had reason to believe that the first party intended to agree[10]. [7] suggests that the objective theory cannot be relied on for assistance when an offer can be said to be initiated by the electronic device "autonomously i.e., in a manner unknown or unpredicted by the party employing [i.e., operating] the electronic device". This is because a party employing such a device cannot be thought to assent to contracts of which he is unaware, and which he cannot predict.

We do not share these doubts about the potential utility of the 'objective theory'. As above, we distinguish autonomy from unpredictability. Almost by definition, autonomous action takes place without the knowledge (at least contemporaneously and of the specific transaction) of the principal. But just as a good employee can, while exercising autonomy in decision-making, stay within well-defined boundaries and act in predictable ways, so can an artificial agent. We suggest that the 'objective' theory of intent might be relied on as an alternative underpinning to many contracts reached autonomously through artificial agents. However, [7] is correct to suggest that there would be limits to the applicability of the objective intent theory. We suggest that those limits would be reached not when agents behave autonomously but when agents behave erratically or unpredictably. Under those circumstances, it becomes difficult to argue that a reasonable person would conclude that, merely by reason of operating the agent, the operator should be taken to assent to the terms agreed by the agent. This limit is similar to the limit we proposed on the applicability of the unilateral offer doctrine.

It is not clear however, that the 'unilateral offer' doctrine or the 'objective theory' doctrine should be used as an attribution rule at all, whereby the actions of an artificial agent can be attributed to its operator. In current law relevant to human agents, the doctrinal work is done by the *law of agency*, which is intimately connected with the notion of the agent's authority to act. Within the scope of the agent's actual authority, legal acts done by it on behalf of its principal - such as entering a contract or giving or receiving a notice - become, in law, the acts of the principal. The doctrine extends to cases of

apparent authority where the agent has no actual authority, but where the principal permits third parties to *believe* that he has authority.

The fourth potential solution, therefore, involves taking the artificial agent metaphor seriously, and treating artificial agents literally as the *legal agents* of their operators. Under this solution, within the scope of the agent's authority, contracts entered into by agents would become contracts of their operators. The agent's authority could be readily understood as that field of contracts which the agent had instructions/permissions - and the means - to conclude. A number of objections to the possibility of treating artificial agents as true agents in the legal sense have been discussed [7, 8]. Refutations to the most important of these follow:

1. Artificial agents necessarily lack legal power to give consent because they are not persons. But the example of Roman slaves - who were not considered legal persons but who did have capacity to enter contracts on behalf of their masters - disproves this objection.
2. Artificial agents lack the intellectual capacity or ability to exchange promises. But many artificial agents (such as the interfaces operated by shopping websites) display all the 'intelligence' or 'intentionality' required of them - no less so and much more reliably than the human telephone clerks that they replace, and which in many cases are practically indistinguishable from recorded voices. We believe the objective theory of contract should be deployed in this context On this approach, an ability to perform the requisite actions - such as emailing the buyer with the correct details in a notification and updating the seller's databases correctly - would be sufficient to qualify a seller's artificial agent as having the right 'intentional' or 'mental' states. We do not see how else courts could adjudicate on these matters other than by adopting an 'intentional stance' towards agents: it would be the best way to make sense of their behavior[3].
3. In some legal systems agency requires a contract between the agent and the principal, and as artificial agents are not persons, they cannot enter contracts in their own name. However, in Anglo-American law a contract between principal and agent is not necessary; all that is necessary is that the principal is willing for the agent to bind him as regards third parties. In legal systems where a contract is necessary, it will be necessary to investigate personhood for artificial agents in order to deal with contracts made by artificial agents, or move to the Anglo-American model.
4. Another objection relates to the mental capacity of artificial agents. While in Anglo-American law an agent need not have the contractual capacity of an adult legal person - for example, children who cannot contract for themselves can contract on behalf of adults - nevertheless, an agent must be of sound mind in order for the agency to begin or to continue. This means that the agent must understand the nature of the act being performed. This doctrine would need to be adapted, to the case of artificial agents without human-like intelligence, before a true agency treatment of artificial agency could be undertaken.

Scholars have postulated a fifth potential solution, where the legal system whould treat *artificial agents as legal persons* in order to solve the doctrinal difficulties cited[1, 7, 8, 15, 11]. But as pointed out above, a contracting agent need not necessarily be treated as a legal person to be effectual. The advantage of treating artificial agents as legal persons in the contracting context is that it would provide a complete parallel with the situation of human agents, in terms of the ability of the innocent third party to sue the agent for breach of authority. There are however a number of objections to the idea of according artificial agents with legal personality as a way of dealing with the contracting problem:

1. It is unnecessary, as other solutions are adequate.
2. The warranty of authority is not a particularly important or much-relied upon doctrine of law, so the inability to sue the errant agent (in the agent-as-slave scenario) is not in practice a significant loss for the innocent user.
3. The personal identity conditions for artificial agents are not well-understood. For example, a multi-agent system in the form of 'swarmware' - consisting of multiple copies of the same program in communication - might alternately be seen as one entity and a group of entities. [6, 15] have suggested a so-called 'Turing register' where agents - and their principals - would be registered and recognised by the legal system, much as companies are registered today. [7] points out that the costs of the Turing registry might outweigh its benefits.
4. In civil law countries, the concept of a legal person is intimately bound up with the concept of patrimony[11], i.e. the assets under the control of the person which might be used to satisfy a judgement against the person. It may seem unclear where such assets can be derived from - although one possibility would clearly be ordinary gainful employment, on behalf of users or operators.

2.2 Reflections on the possible solutions

Which solution one chooses depends on whether one believes that the risk of unpredictable and erratic activity of artificial agents should fall on their operators, or on those interacting with them. This question can be seen as a difficult issue of efficient risk allocation in the economics of law, which is outside the scope of this paper. The five approaches sketched above have different results:

1. Under the 'agent as mere tool' solution, the principal would be liable for erratic behaviour - unless a reasonableness requirement were imported into it - but the principal could in many cases recover his loss from the designer of the agent, with whom he might have a contractual relationship, or under product liability laws.
2. Under both the 'unilateral offer' solution and the 'objective intent' solution, the principal would generally not be liable for erratic behaviour, if it were not reasonable for the user to believe that the principal would have assented to that behaviour, had it been aware of it.
3. Under the agent as slave solution, all behaviour outside the scope of the agent's authority - actual or apparent - could be disclaimed by the principal, and the user would have no right of action against the agent, which would not be a a person. The user might however have rights under product liability laws against the designer of the agent.
4. Under the agent as person solution, the principal would not be liable for the agent's behaviour outside the scope of its authority, but the user would be able to sue the agent for breach of its authority. The agent might conceivably in turn be able to sue its own designer under product liability laws, for instance.

We suspect that the 'agent as mere tool' doctrine will survive for some time on grounds of convenience and justice. As agents become more sophisticated, a slave analysis, and then an analysis involving personhood, might be embraced. The contractual problem is not enough, on its own, to motivate a personhood analysis. This is the closest so far to a real-world problem of according personhood to artificial agents. It is intended to show the extent to which 'system-level' concerns will continue to dominate for the foreseeable future.

Factors such as whether it is necessary to introduce personhood in order to explain all relevant phenomena, efficient risk allocation and whether alternative explanations gel better with existing theory, will count for more than the qualities or capacities of artificial agents in this debate.

3 'Constitutional personhood'

The legal literature about artificial agents, we have argued, provides a testing-ground for the pragmatic value of philosophical positions about the nature and possibility of artificial intelligence and of moral personhood. A strong strain in that literature takes the possibility of personhood for artificial agents seriously in the context of seeking to address the contracting problem. Assuming continuing development in the intelligence and autonomy of artificial agents, we believe that questions will arise whether the legal system should consider such agents as legal persons in the context of the constitutional rights accorded by the law to legal persons generally.

In what follows, we concentrate on the factors and arguments relevant to legal personhood, but some will also be relevant to the separate question whether, given legal personhood, *constitutional protections* should be accorded to artificial agents. One salient example of the latter[14] is Article 13 of the United States constitution, which states 'neither slavery nor involuntary servitude, except as a punishment for crime whereof the party shall have been duly convicted, shall exist within the United States, or any place subject to their jurisdiction.' Because each constitutional protection has differing public policy bases, it is hard to generalise as to what factors will be relevant to the question whether artificial agents should be accorded them. So we restrict our attention to the preliminary question of whether agents should be accorded legal personality.

We do not think that the issue of legal personhood for artificial agents will come ready-formed into the courts. Rather, a system for granting legal personhood may need to be set out by the legislature, perhaps through a registration system as in the case of companies.

As demonstrated in the contracting problem case study, unless personality for artificial agents is needed doctrinally, there is unlikely to be overwhelming internal pressure from the legal system to accord it to them. However, where legal convenience is increased by granting legal personhood to artificial agents, this will obviously be a relevant consideration, but it should not be overstated. The convenience of the modern business corporation's legal entity status is significantly greater than that of the unincorporated partnership, but many large businesses continue to survive as partnerships. In the modern law-making context, cost-benefit analysis is of increasing importance. Few laws are proposed today without an argument that the costs imposed by the measure outweigh the benefits. Such considerations would play an important role in the decision whether to accord artificial agents with legal personality.

One argument against legal personhood for agents is their limited susceptibility to punishment. But the modern corporation is accorded legal personality; although it cannot be imprisoned, it can be subject to financial penalties. More crucially, artificial agents can be imagined that have a moral sense, in the sense of responding to the threat of punishment by modifying their behaviour, goals and objectives in appropriate ways. Once artificial agents have a moral sense, the law has something to act upon directly. However, infants, who have little or no moral sense, and few legal responsibilities, are nevertheless accorded legal personality by modern law; intelligent animals such as dogs, which 'know' what their masters do and do not wish them to do, are punished for disobedience - but it is not thought for that reason that they should be accorded legal personality. Here, recognition of 'species resemblance' and similarities in potential dispositions between children and adult humans seems to drive the ascription of legal personality.

If it starts to make sense to discuss the behaviour of artificial agents as morally good or bad, then the legal system might well 'sit up and take notice'. An important element of 'moral personhood' is the ability to distinguish right from wrong - the sense, when doing something wrong, that that is what is being done. An artificial agent with a set of internalised 'commands', but which nevertheless was autonomous enough to disobey those commands, might well qualify as a 'moral person'. If so, this would be an argument in favour of legal personhood. Under the 'moral personhood' banner is one factor - susceptibility to deterrence - which we have treated above. Another consideration in this debate is whether we should ever fully delegate responsibility to artificial agents - see [9] for arguments against such delegation.

Legal personality is an important step towards being accorded full constitutional ('human') rights, as it is only when an artificial agent could qualify as a legal person that constitutional protections come into play. The result is that arguments relevant to whether artificial agents should be accorded constitutional protections such as those against slavery will be relevant to the decision whether to accord artificial agents with legal personality, even though not all constitutional protections are accorded to every legal person. The arguments leading up to the abolition of slavery in the United States in 1865 were many and varied, but all rested on the fundamental equality of all human beings. The similarity between humans was held more important than the difference between races. The more artificial agents come to resemble humans, the more our legal system will be drawn to accord them legal personality. It should not be forgotten that it took a long civil war, and not just intellectual or emotional arguments, to abolish slavery in the United States.

In response to the objection that artificial agents are not humans, [14] observes that corporations have constitutional rights and concludes that "our concept of a person may change in a way that creates a cleavage between human and person. Our current linguistic practice will not be binding in the imagined future".

Some philosophical arguments against personhood for artificial agents might find favour in a legislative or courtroom setting. [14] examines these and concludes that:

- 'missing-something' arguments, which try to show that artificial agents cannot be legal persons because they lack some essence of personhood, whether it be souls, consciousness, intentionality, feelings, or free wills, either would not survive contrary intuitions and actual experiences of jurors (or legislators), or depend on values or assumptions not shared in a modern pluralistic society;
- the Lockean argument that artificial agents ought to be property of their creators, cannot apply to children - and would not to hypothetical humans built by scientists using DNA, so should not therefore be thought to apply to artificial agents.

[14] suggests that if the behavioural evidence and knowledge of underlying processes both pointed to *actual rather than simulated features of human mentality*, this would be a good reason to believe that artificial agents did possess these features. This accords a role to a similarity in underlying cognitive processes between humans and artificial agents - and therefore to cognitive science in theorising and/or discovering what those might be. It is not clear however whether a legal system would deny an agent personality on the basis simply of its internal architecture as opposed to whether it engaged in the right

kinds of behaviour, because its behaviour is what will regulate its social interactions. Note too, that the distinction made above assumes that we know what the actual features of mentality are - which brings us back to the starting point of this discussion.

What attributes of particular agents should the legal system take into consideration in determining whether to grant personhood to particular agents? The practical capacity to perform cognitive tasks will presumably be of primary importance. However, something like the Turing test is neither a necessary or sufficient condition for this capacity. Children and the mentally incapacitated are accorded legal personality by the legal system, while having limited mental capacity; companies and non-human legal persons such as ships or temples do not exhibit intelligence of their own - although they are associated with human owners or representatives who display those attributes.

An important pragmatic factor is the ability to control money: being able to receive, hold and pay money and other property such as securities, and to remain financially solvent. Without this ability a legal system might be reluctant to impose liabilities on such entities - as the civil law concept of patrimony attests. Prior to being accorded legal personhood, agents would need to be able to do these things on behalf of other people, and agents that were unable to undertake economic transactions without relying on intermediaries would find their case for legal personhood diminished. On the other hand, the most common form of legal person other than humans at present is the modern business corporation. This can only act by its agents - by itself it is helpless. So a technical inability to perform a task is not, ultimately, a bar to being accorded legal personality.

A persistent debate in the philosophy of mind is whether artificial agents can ever be said to be conscious. The legal system has not seen consciousness as a necessary or sufficient condition of legal personality. Historically, many categories of fully conscious humans - such as married women, slaves and children - have been denied legal personhood. Conversely, persons in comas or asleep i.e., humans temporarily lacking consciousness - are not denied legal personality on that basis (although those in a permanent vegetative state often suffer fairly drastic curtailments of their human rights). The many categories of non-human legal persons lack consciousness. We believe that the question of whether artificial agents have phenomenal experience is misguided (see [4] for an amplification of this view). As [2] points out, artificial agents simply do not share the same perceptual world as us - and since intersubjective agreement about phenomenal experience is the only way we have of confirming that other human beings have the same phenomenal experiences as ourselves, why not extend this to artificial agents?

Some artificial agents might be so complex and their reasoning so adaptive, that they effectively would be the best immediate authority in terms of reporting on their own mental states[5, 13]. Such an ability plays a role in our interactions with human beings - for rather than examining a human's neurological structure to determine their reasons for a particular action, all we have to do is ask, and we are assured that in most cases the reports we receive will be reliable and the best source of information. Any legal interaction with such an agent would perforce rely primarily on the agent itself for eliciting reasons for its actions - an important determinant of autonomy.

4 Conclusion

While the legal theory of artificial agents is not fully fleshed out, whatever path is adopted will have significant effect on philosophical theorizing about artificial agents. The non-reliance of computing and mentality on a particular physical substrate has made possible speculation that the cognitive status of agents will be a matter of pragmatic judgement. The high point of such pragmatic deliberation is the legal sphere. As noted above, a crucial determinant in courtroom arguments is historical or legal precedent and pragmatic considerations of our society's best interests. Attributes such as the practical ability to perform cognitive tasks, the ability to control money, and 'legal system wide' considerations such as cost benefit analysis, and moral and extra-legal arguments such as are relevant to slavery, will play their part in legal personhood with constitutional rights implications. Our conclusion from the discussion on agency law is that doctrinal convenience and neatness, as well as just outcomes, will influence personhood decisions for less sophisticated agents. A mix of the pragmatic and the conceptual will play a role in future determinations of the legal status of artificial agents. Much more needs to be said, and interplay between philosophical and legal theorizing will remain intense as artificial agents gain in sophistication - an exciting prospect for computer scientists, legal scholars and philosophers.

A final note on these entities that cause us such perplexity by their presence in our midst. Philosophical discussions on personal identity often take recourse in the Wittgensteinian idea that ascriptions of personal identity are of most importance in a social structure where that concept plays the important legal role of determining responsibility and agency. We ascribe a coherent unity to a rapidly changing - both physically and psychologically - object because very little social interaction would otherwise make sense. Similarly, it is unlikely that in a future society where artificial agents wield tremendous amounts of executive power, that anything would be gained by continuing to deny them personhood. At best it would be a chauvinistic preservation of a special status for biological creatures like us.

REFERENCES

[1] Tom Allan and Robin Widdison, 'Can computers make contracts?', *Harvard Journal of Law and Technology*, **9**, 25–52, (1996).

[2] Rodney Brooks, 'Intelligence without representations', in *Mind Design II*, ed., John Haugeland, MIT Press, (1997).

[3] Daniel Denett, *The Intentional Stance*, Bradford/MIT Press, 1987.

[4] Daniel Dennett, 'Quining qualia', in *Readings in Philosophy and Cognitive Science*, ed., Alvin Goldman, MIT Press, (1993).

[5] Daniel Dennett, 'The case for Rorts', in *Rorty and his critics*, ed., Robert B. Brandom, Blackwell Publishers, (2000).

[6] CEA Karnow, 'Liability for distributed artificial intelligences', *Berkeley Technology Law Journal*, **11**, 147–204, (1996).

[7] Ian R. Kerr, 'Providing for autonomous electronic devices in the uniform electronic commerce act', in *Proceedings of the Uniform Law Conference of Canada*, (1999).

[8] Ian R. Kerr, 'Ensuring the success of contract formation in agent-mediated electronic commerce', *Electronic Commerce Research*, **1**(1/2), 183–202, (2001).

[9] A. Kuflik, 'Computers in control: rational transfer of authority or irresponsible abdication of autonomy', *Journal of Ethics and Information Technology*, **1**(3), 173–184, (1999).

[10] Jean F. Lerouge, 'Uniform computer information transaction act', *John Marshall Journal of Computer and Information Law*, **18**, 403, (1999).

[11] Federica De Miglio, Tessa Onida, Francesco Romano, and Serena Santoro, 'Electronic agents and the law of agency', in *Proceedings of the Workshop on The Law of Electronic Agents (LEA-02)*, (2002).

[12] Jane Nosworthy, 'The Koko dilemma', *Southern Cross University Law Review*, **2**, 1, (1998).

[13] Richard Rorty, 'Incorrigibility as the mark of the mental', *Journal of Philosophy*, **67**, 399–424, (1970).

[14] Lawrence B. Solum, 'Legal personhood for artificial intelligences', *North Carolina Law Review*, **2**, 1231, (1992).

[15] Steffen Wettig and Eberhard Zehendner, 'The electronic agent: A legal personality under German law?', in *Proceedings of the Law and Electronic Agents workshop (LEA'03)*, pp. 97–112, (2003).