Norm Governed Agents In CSCW

Kecheng Liu and Alan Dix
School of Computing, Staffordshire University
PO Box 334, Stafford, ST18 0DG, UK
K.Liu@soc.staffs.ac.uk, A.J.Dix@soc.staffs.ac.uk

http://www.soc.staffs.ac.uk/~cmtkl/ssig.html

Abstract

Information systems are organisations in which signs are created, processed and consumed for social and economic purposes. In order to fulfil the organisational goals and objectives, members within an organisation must understand their responsibilities and authorities, and must act cooperatively. The key to this organisation and coordination lies in norms which define responsibilities and authorities for each human agent, and establish regularities of behaviour. In the context of computer supported cooperative work (CSCW), where "intelligent" software agents are involved, to understand the norms of behaviour of various human agents becomes critical. Software agents can perform some tasks autonomously on the user's behalf. but this autonomy is the result of delegation of job-functions and responsibilities within that user's authority. Such delegation involves a set of complicated philosophical and legal issues. Following discussion on delineation of various boundaries of responsibility and authorities, this paper addresses norms and normative behaviour of human agents within an organisation. Taxonomies of norms are discussed. A method of norm specification using deontic operators is described, supported by a series of examples of banking and email.

Motivation

Information systems are organisations and are, therefore, social systems. In business operations, signs are created, processed and consumed within organisations for social and economic purposes. In order to fulfil the organisational goals and objectives, "actors" within an organisation must act cooperatively, and the success of an organisation depends on the work of all its members. The coodination and cooperation are achieved by the norms that govern people's behaviour. Advanced information technology lends itself to support and facilitate the articulation of distributed and dispersed work activities, thereby increasing productivity and efficiency dramatically while quality of work is also improved. Such information technology includes computer supported cooperative work (CSCW).

In a large, distributed, network computer-based system, "intelligent" software components are often used. Such software components are labelled "intelligent agents" (Wooldridge & Jennings 1995). This notion may have been potentially confusing with human agents in a social, business situation. Human agents in a business organisation will act according to the social, organisational norms. They can be assigned duties and authorities, and can be held responsible for their actions. Many of these concepts have been adopted by the AI (artificial intelligence) research community. In their definitions of machine "agency", characteristics such as autonomy, social ability and communication with other agents are thought to be fundamental. A software agent is supposed to be capable of expressing beliefs, desires and intentions. However, key questions have to be asked before the agency theory is applied to machines. Do we understand the patterns of human agents well enough? Are we able to capture and represent human beliefs, desires and intentions? How do we express human authorities and responsibilities?

Background

In a collaborative system, people have business responsibilities and commitments, some of which can be delegated to software agents in many roles: proxies, performing actions and interacting with other agents on a user's behalf; mediators, passing on or routing messages; translators, modifying the form of messages; or coordinators, suggesting or determining allowable actions within the collaborative environment. From a semiotic standpoint we can see that most of these roles include aspects of sign processing or transmission. We therefore look to the software agents in a system and examine how they help the human users in their work.

Agency in collaborative work

CSCW systems are composed of human users (human agents), the media through which they communicate and the objects on which they act (Dix, Finlay & Hassell 1994). The objects are important both as the focus of work and as a means of communication *through artefacts* (Dix 1994). However, for this paper we will concentrate on the agents and media. In addition, many CSCW systems include software agents intended to improve the efficiency and quality of work. In fact, the distinction between agents in a software system and the media of communication is itself complex. Even in a physical medium such as the postal system, human agents (the postman, sorting office staff, etc.) act as mediators and routers of letters. The same may happen in a software system and, in addition, parts of the system labelled as agents may perform fairly trivial tasks.

Characteristics of machine agents

An *agent* is normally seen as a person who acts for another, especially one who looks after or represents the business affairs of a person or firm. In organisational semiotics, the term "agent" denotes a person or a group of people in a social system who take responsibilities and perform actions (Stamper 1980; Stamper & Liu 1994). Recently the concept of agency has been adopted in computing science, but used specifically to refer to a software component or an "intelligent" machine as an "agent" The computing technology of "intelligent agent" has proved useful in many application areas, for example the Internet, communications, network management, electronic commerce and database management. It is also highly relevant for CSCW because the number of beneficiaries is greater than in a single user situation, and hence the efficiency gain is much more significant.

Research workers tend to agree on the following characteristics necessary for a machine-based agent (Wooldridge & Jennings 1995): autonomy: it can operate without the direct intervention of humans or others, and has some kind of control over its actions and internal state; social ability: it interacts with other agents (and possibly humans); reactivity: it has the ability to perceive its environment (the physical world, a user, other agents, the Internet, etc.) and responding to changes in the environment; pro-activity: it takes initiatives and behaves pro-actively in order to achieve pre-determined goals.

Researchers in DAI (distributed artificial intelligence) claim that an intelligent software or machine agent has a way to exhibit an intentional stance, such as knowledge, belief, desire, intention, obligation, commitment, choice, and so on (Wooldridge & Jennings 1995). It is believed that these complex notions can be captured and represented to the extent that automatic manipulation and reasoning can be performed on a rational basis. A great deal of effort has been put into representation of "beliefs, desires and intentions" using modal logic, temporal logic and the like (e.g. Rao 1995). Success stories of machine agents can be seen in a number of applications areas. As reported by Smith and Mamdani (1996), they can be used to assist people with routine but important tasks, to undertake complex tasks rapidly, to act as gophers on a human user's behalf, to explore large cognitive search spaces, to find elusive solutions and to deal with data tasks.

Warnings are given regarding issues such as responsibility, authority, liability and ethical questions (Nwana & Ndumu 1996; Eichmann 1994). These issues have not, in general, been put on the immediate research agenda; therefore, little attention has been paid to how to understand and handle the questions of obligation, commitment, responsibilities and authorities. It is argued later in this paper that these questions are important and have immediate relevance.

Norms and normative behaviour

Five types of norms can be identified, each of which governs a certain aspect of human behaviour (see Stamper *et al.* in these proceedings). Perceptual norms deal with how people receive signals from the environment via their senses through media such as light, sound and taste. Cognitive norms enable one to incorporate the beliefs and knowledge of a culture, to interpret what is perceived, and to gain an understanding based on existing knowledge. Evaluative norms help explain why people have certain beliefs, values and objectives. Behavioural norms govern people's behaviour within regular patterns. Finally, denotative norms direct the choices of signs for signifying; such choices are culture-dependent, e.g. the choice of a colour to signify happiness or sadness.

Norms are developed through the practical experiences of people in a culture, and in turn have functions of directing, coordinating and controlling actions within the culture. A research group or a working team may have a sub-culture and therefore may have "local" norms. The norms will provide guidance for members to determine whether certain patterns of behaviour are legal or acceptable within the given context. An individual member in the community, having learned the norms, will be able to use the knowledge to guide his or her actions, though he or she may decide to take either a norm-conforming or a norm-breaking action (Liu, 1993). When the norms of an organisation are learned, it will be possible for one to expect and predict behaviour, and hence to collaborate with others in performing coordinated actions. Once the norms are understood, captured and represented in, for example, the form of deontic logic, this will serve as a basis for programming intelligent agents to perform many regular activities.

Embedded Norms in Electronic Media

Most CSCW systems require some form of sign processing and transmission, often mediated by various levels of electronic media and computer agents. The communication over such media will be governed by norms of behaviour and communication, but with some intermediaries it becomes problematic where these norms are interpreted and understood. It is clear that a human translator working for a senior diplomat would not only use semantic knowledge as part of the language translation, but would also use knowledge of the cultural and diplomatic norms in choosing appropriate phrasing.

Mediation and agency

The most common form of electronic mediation is the telephone. Although there is extensive signal processing (probably digital) and routing of signals, the medium is semantically neutral. The affordances (Gaver 1992) or constraints (Clark 1991) of the medium may modify human communication, but the norms are produced and interpreted by the human users alone.

Email is slightly more complex. In the simplest case an email system consists of mail agents for each user and a transport medium (a computer network such as the Internet). The phrase "mail agent" is used to refer to the actual program the user interacts with; that is, the user

Figure 1: Communicating by telephone – semantically neutral

Figure 2 Communicating by email – agents and routers

interface of the email system. These agents are important as they embody most of the "intelligence" of an email system as a whole.

Low-level email standards

Because email operates across platforms and between proprietary systems, there are extensive standards covering issues such as addressing, routing and message content. These standards can themselves be regarded as a form of low-level norm and have arisen by a variety of processes: international committees such as X500, proprietary standards such as MAPI and standards arising from agreed common usage such as the Internet standards. We shall look at the Internet standards in more detail to see where higher levels of norm are embodied. As we move up through the levels of the system more complex norms are embodied.

Consider an email message such as the one in Figure 3. At the lowest level, email is delivered from machine to machine using the Simple Mail Transfer Protocol (SMTP). The sending machine contacts the recipient or routing machine, tells it who the message is for (A.J.Dix@soc.staffs.ac.uk, counter-staff@bigbank.com, and D.Tracy@fbi.gov), and who it is from (K.Thomas@bigbank.com) and then sends the body of the message including all the headers. This level of protocol does not look inside the message at all. It is like addressing an envelope, putting a return address on the back, putting a letter inside, sticking it down and then posting it. The post office should not look inside the envelope, but simply deliver it. Note especially that the interpretation of headers is left entirely to the mail agents.

The headers themselves and the body are subject to standards governing various common headers with their meanings (From:, To:, Subject:, etc.) and the coding of different kinds of file format and media (MIME). The SMTP does not check that the recipients or sender agree with those in the body, or that the message is conformant with MIME or other standards. Because of this neutrality of the underlying medium, mail agents can use different encoding standards for email content and additional headers (normally prefixed with "X-"). Of course, if email is sent between different kinds of email agent parts of the message may not be properly interpreted, especially file attachments.

The message in Figure 3 has a "To:" field, a "Cc:" (copy to) field and a "Bcc:" (blind copy) field. The email is delivered to the recipients in all these fields. As far as the email agent is concerned, when posting the message "To:" and "Cc:" are treated the same. However, it does recognise the different status of "Bcc:" and before sending the email it removes the "Bcc:" line from the header. That is, the email agent "understands" one aspect of the norms governing email usage. Where then lies the difference

From: K.Thomas@bigbank.com
To: A.J.Dix@soc.staffs.ac.uk
Subject: ac. no. 765 183 0078
Cc: counter-staff@bigbank.com
Bcc: D.Tracy@fbi.gov

Dear Prof. Dix,
The state of your account is causing some concern. In particular, your recent withdrawal of £117,000 has considerably exceeded your overdraft limit of £500. I have therefore asked staff to refuse any further withdrawals and urge you to contact me immediately on your return to the country.
Yours,
Kate Thomas (Senior Account Manager)

Figure 3. Typical email message

between the "To:" field and the "Cc:" field? In fact, the difference lies in the way the recipients interpret them(Dix, Finlay *et al* 1993). If you are named in the 'To:' list, you will feel a greater need to respond than if you are in the "Cc:" list, when you may simply need to take note of the contents. The precise details will depend on the norms of the organisational culture of the sender and recipients (and may be misunderstood because of the neutrality of the medium). However, the important thing is that the interpretation of these norms is performed by the users themselves.

Rules and filters

The knowledge about "Bcc:" is built into the email agent by the designer. However, many email systems allow users to specify filtering, which may alert the user, file messages or even discard them according to simple rules such as:

if field 'Subject:' contains 'special offer'
then move to Trash

These rules embody *individual* norms of behaviour that are executed by the mail agent. However, such rules are limited by the restricted set of fields available. They allow simple decisions based on sender or recipient (mailing list or personal), but more complex decisions require text matching rules(as above) which are notoriously error prone. Also note that participants in an email exchange may not even be aware that their colleague is using an automatic email filter.

Semi-structured message systems such as Group Lens (now known as Oval) offer a more explicit representation of the semantics of a message (Malone 1987). These systems incorporate typed messages. Messages of a particular type will have different fields; for example, a seminar announcement may have a Speaker: and a Title: field. These fields allow more efficient personal filtering of messages based on message types and the fields they contain. However, this is based on the explicit cooperation of colleagues. Different types of message can easily be created and sub-classed from existing types, but the efficiency of filtering depends critically on the common understanding and usage of types and fields. Thus we have a system which can embody, and be tailored to, some of the group norms of particular communicative intra- and inter-organisational sub-groups.

Workflow

Some of the most explicit coding of norms within messaging systems is found in work-flow systems (Kreifelts 1991; Prinz 1996). These embody pre-defined *corporate* norms of behaviour. Users are presented with documents/messages where they must fill in or approve specific parts; depending on their responses the document message will automatically be passed on to the next person who must see it, according to the pre-defined business process. For example, an employee may fill in an electronic expense form. On completion, the form is passed to the line manager for approval, then to the finance department who will assign a payment code and finally to payments who write the cheque. The popularity of business process re-engineering (BPR) has driven the widespread take up of such systems, often implemented using Lotus Notes or similar scriptable messaging systems.

At the extreme of workflow systems is Coordinator (Winograd 1986). This is a typed messaging system where the types, and rules relating to types, are based on speech-act theory (Searle 1969). Generic patterns of "speech acts" have been identified, most well known being the "conversation for action" (CfA). According to Winograd and others these speech acts represent norms of conversation across societies. In Coordinator, the initiator of a conversation must identify what type of conversation is being performed and for each message the users must identify what illocutionary point (encoded in message type) it is intended to convey (a request, a refusal, etc.). The system only allows messages that conform to the speech-act patterns.

Levels of embedded norms

Consider again the move from simple email, through filtering and semi-structured messages, to workflow systems and Coordinator. First this represents a gradual increase in the explicit semantic content encoded within the messages. Second, this explicit semantic knowledge also represents a gradual shift in the communality of norms embodied within the mail agents (or messaging system) as illustrated in Figure 4. In all these cases the rules obeyed by the agents have been explicitly coded by someone. These rules embody different levels of norms, but can we say that the agents interpret or understand these norms, or are they merely obeying orders? Is it that the humans in the system impute understanding of norms to the computational agents? Of course, we may have human agents acting on our behalf (e.g. a stockbroker), whom we also expect to obey our orders. The main difference is that the computational agent obeys fixed deterministic rules, whereas the human agent is expected to act with discretion within parameters and must thus understand as well as be able to obey the norms. However, it is important to be able to specify these norms, otherwise how can we be sure that a human agent is acting within the appropriate boundaries of responsibilities and authorities?

system	embodiment of norms
simple email	naming
email filters	individual norms
semi-structured messages	group norms
workflow	organisational norms
Coordinator	social and linguistic norms

Figure 4: Norms embodied in computer systems

Norms and Responsibility

Still, difficult issues arise in two fold: how to decide the extent to which human responsibilities should be delegated to a software agent, and what are the roles of human agents when machine agents are doing most of the work? Before an attempt is made to explore these issues, a brief review of the most fundamental legal conceptions is necessary.

Authorities and delegation

A legal philosopher, Wesley N. Hohfeld, recognised eight fundamental legal conceptions in two sets (quoted in Allen & Saxon, 1993). In each set, the paired concepts are "opposites" in columns and "correlatives" in rows:

Right Set:	Right	Duty
	No right	Privilege
Power Set:	Power	Liability
	Disability	Immunity

As discussed by Allen and Saxon (1993), all eight concepts are related to the terms "must", "should", "must not" and more. Each term may involve more than one of these legal concepts. For example, "a credit card holder *must* pay any outstanding amount of credit within 25 days of posting without incurring interest", implies that a card holder has the *rights* to pay, and not to pay till the due day, but a liability of paying interest will occur from the due date if the amount is not paid. As pointed out by Allen and Saxon (1993), these terms may also result in multiple interpretations and it is only appropriate to expect a machine to assist the human in interpretation.

A machine agent should be seen as an assistant to the human, who can delegate some of his or her responsibilities. Only when the human user has the right and power can he or she transfer responsibilities to a machine agent. It should be noted here that what is transferred are duties or job functions rather than liability. In the same manner the government of a country can delegate responsibilities and functions to an embassy in a foreign country but cannot relinquish the liability for what the ambassador has done. Of course, this does not stop users from attempting to use software to abrogate responsibility "sorry I can't help, the computer says so"!

To make a machine agent behave in the same way as the owner would wish, behavioural norms of its user have to be specified and stored in the machine agent. Other types of norms, such as perceptual, cognitive and evaluative, can be used to enhance the machine agent's learning ability. Finally, denotative norms will be used to design the interaction between the machine and the owner and other users. However, study of some of these norms may be more difficult because their effects are not as observable as behavioural norms.

Norm specification

Understanding of the norms and patterns of people's behaviour within an organisation is a foundation for designing an effective CSCW system. In business, most rules and regulations fall into the category of behavioural norms. These norms prescribe what people *must*, *may,* and *must not* do, which are equivalent to three deontic operators "is obliged", "is permitted" and "is prohibited". Hence, the following format is considered suitable for specification of behavioural norms:

whenever <condition>
if <state>
then <agent>
is <deontic operator>
to <action>

Adopting this form, a credit card company may state norms governing interest charges as follows:

whenever *an amount of outstanding credit*
if *more than 25 days after posting*
then the *card holder*
is *obliged*
to *pay the interest*.

whenever *an agreement for credit card is signed*
if *within 14 days after commencing*
then the *card holder*
is *permitted*
to *cancel the agreement*.

The first norm says that after 25 days of posting the invoice, if there is still an amount of outstanding credit, the card holder will have to pay the interest. The second norm states that the card holder retains the right of cancellation of the agreement within 14 days of commencing. The next norm says that unless there is a special arrangement made, e.g. with the account manager, the card holder is not allowed to spend more than the credit limit:

whenever *purchasing*
if *no special arrangement is made*
then the *card holder*
is *forbidden*
to *exceed the credit limit*.

The card holders are expected to behave according to the norms stated in the agreements. As understood both by the customers and the credit card company, the company may impose sanctions if a customer fails to observe the norms. With this form of specification of norms, a computer program can be written to execute the norms. As long as the norms are specified, computing technologies such as active databases, object technology and artificial intelligence will have different approaches towards software realisation.

Autonomy and discretion

Consider the email filtering rule we saw earlier. This rule could be restated using the form of the previous section as:

whenever *a mail message arrives*
if *field* 'Subject:' *contains* 'special offer'
then the *mail agent*
is *obliged*
to *move the message to Trash.*

Compare this with the sorts of rules human agents operate under. Clearly the conditions and actions of human rules do not need to be as formally stated as those of a software agent, simply because the human has more intelligence and contextual knowledge. However, perhaps the most significant difference is in the deontic condition. The email filtering rules will always be an obligation. We don't expect the filtering agent to sometimes sort our mail, when it feels like it. No, it must always sort our mail in precisely the way prescribed. Simple email filtering is deterministic.

However, intelligent agents are expected to operate with more autonomy, no longer satisfying deterministic rules, but instead learning from experience and taking initiative. This suggests that the norms embodied in these agents should take the more general deontic form. Yet if an agent is only permitted to perform some action, why should it ever do it, and if it doesn't does it matter?
Human agents operate within this kind of looser deontic context. How do they cope?

Framing norms
Consider the case of a bank manager faced with an unauthorised overdraft. The published condition of the account would state (albeit in plain English!):

> **whenever** *overdraft exceed agreed limit*
> **then** the *bank*
> **is** *permitted*
> **to** *charge a fee.*

This is the rule which the customer has seen when taking out the account and hence legally agreed to. It thus forms a *contractual* norm. The customer cannot complain if the bank exercises its right to charge a fee. In practice, however, the bank manager will exercise discretion. The bank manager will not want to alienate a good customer and hence will often waive the fee.

This sounds as though a valid course of action for the bank manager would be to never charge a fee at all on unauthorised overdrafts. Clearly the customers would be quite happy with this situation, but, of course, the bank would not be! In fact, there will be *corporate* norms within the bank, some explicitly stated in rule books, others implicit. These will constrain the bank manager's discretion to act. One such rule might be:

> **whenever** *overdraft exceed agreed limit*
> **if** *the excess is unduly large*
> **then** the *bank manager*
> **is** *obliged*
> **to** *charge a fee.*

Of course, the bank will neither want the bank manager to exercise discretion in an overly legalistic manner and will have further *corporate* norms which prevent this:

> **whenever** *overdraft exceed agreed limit*
> **if** *the excess is small* **and** *it is a good customer*
> **then** the *bank manager*
> **is** *forbidden*
> **to** *charge a fee.*

The bank manager is the agent of the bank and the corporate norms allow the bank to determine the boundaries of discretionary action. Note that the contractual norm governs the behaviour of the bank with the customer. Of particular importance, the corporate norms should be such that the actions of its agent (the bank manager) ensure that the bank acts within its contractual norms.

Note that the discretion of the bank is limited by contractual norms and the discretion of the bank manager is limited by corporate norms. These are both examples of *framing* norms, which set bounds on the discretionary activity of an autonomous agent. Within these framing norms an agent may have its own individual norms: rules and attitudes which it uses to make individual decisions. Some bank managers are more generous than others!

Policy and meta-norms
However, even within the bounds of framing norms an agent is usually not free to have arbitrary individual norms of behaviour. Consider a bank manager who always waived the fees of family and friends, or one who waived fees of male customers but charged overdrawn women. The former would be seen by the bank as unacceptable favouritism, the latter (in the UK) would be illegal.

In fact, there are normally general policies which govern the sort of individual norms which agents are allowed to operate. In the banking context these polices would include some sense of fairness. At a legal level UK courts use the common law concept of equity to limit the discretion of government officials and the courts themselves operate within the principle of precedent.

Notice that whereas the contractual and corporate norms we have considered operate on individual cases, all these policies operate between cases determining the sort of rule which is acceptable as a behavioural norm. These are *meta-norms*, norms which govern the generation of other norms.

Intelligent agents
Whereas the agents we have considered so far in email systems have been deterministic, there is increasing interest in the use of intelligent agents. These may be pre-programmed, in which case they obey norms laid out by the application developer. However, perhaps the most interesting cases are agents which use machine learning techniques to generate their own rules of behaviour. One of the early examples of this is Eager, a system which watches a user's actions, learns patterns and then when it notices a user beginning a learnt pattern offers to complete the sequence of actions (Cypher 1991). Similar techniques have been used to select interesting news postings, build a personalised newspaper, and to generate database queries from examples of required records (Dix & Patrick 1994).

These learning agents are also being developed for email systems. An agent can watch users' actions as they sort their mail and build automatic filtering rules: effectively capturing the user's individual norms of behaviour. Within an organisation they can even watch the patterns of document passing between individual employees and thus infer the corporate norms of behaviour and generate workflows. Furthermore, pre-programmed deterministic agents have already been proposed as part of the infrastructure of virtual organisations – it will not be long before learning agents are used at this cross-organisational level. This may include a mediation role, adapting corporate norms of one organisation to those of another, rather like translators do between languages.

Such agents will start to take decisions which significantly affect the externally perceived behaviour of the individual or organisation. An email agent which discards a message from the managing director could significantly affect the promotion prospects of an employee! Furthermore, it would be easy for such an agent to generate illegal rules and expose the individual or organisation to litigation (Dix 1992).

Meta-norms for agents

The autonomy and discretion of such agents must clearly be limited in the same way that human agents are. Framing rules can be specified using the deontic rules we have outlined earlier. These can be entered at an individual level or customised for the organisation as a whole (for example, "all messages from the managing director are flagged urgent"). However, perhaps the most interesting area is in the expression of meta-norms.

The learning algorithms can themselves be regarded as one level of meta-norm as they limit the sorts of rules which can be inferred. However, this is a crude and technologically determined limit. More interesting are rules which are explicitly built into the system. However, these are by their very nature far more complex than norms for case-by-case behaviour. It is always hard to step back and externalise one's norms of behaviour, how much more difficult to codify the rules which govern the generation of those norms! At the level of corporate meta-norms this may be possible as professional knowledge engineers can perform the requisite knowledge elicitation and agent programming, but it is unlikely that individual users will be able to customise their own interfaces in this way.

Meta-norms can also be captured within processes. A bank clerk who notices a more efficient way of processing loan applications cannot decide to by-pass the bank manager and approve or deny loans on their own authority. However, it may be acceptable (depending on the organisational culture) for the clerk to approach the manager and propose a rule "we seem to always accept loans that are less than half the customers monthly salary". Similarly, an email agent can infer filtering rules and propose them to the user using the same format as the users do to write their own filtering rules. Notice what is happening here. The software agents operate within *procedural* meta-norms. The agents propose new behavioural norms to the users who decide whether they fit within their own meta-norms of behaviour. The mechanism of norm generation lies with the learning agent, but authority for norm generation belongs solidly with the user.

Conclusions and Future Research

Organisations are information systems where people use signs to conduct business activities. Human agents act in the systems in an organised and coordinated manner. In achieving this organisation and coordination, norms play an essential role in governing people's behaviour. To understand and model the behaviour of members of an organisation becomes essentially the task of understanding and representing norms. There are many types of norms which are rooted in rich social, cultural and linguistic contexts. The norms govern a whole range of human activities from perception, cognition, evaluation to action; and they can be explicit as well as implicit. Therefore, modelling norms can be a difficult task and can only be achieved to a limited extent.

Norms are closely related to legal concepts such as responsibility and authority. *Framing norms* are concerned with delineation of such boundaries, within which human agents have the right to make decisions and to exercise discretion. When machine agents are employed in CSCW, some responsibilities will be delegated to machines. The delegation is only valid within the boundary of users' authority.

In illustrating our discussion, email systems have been used as a driving example. Norms can be embodied into email systems at various levels: individual norms, group norms, corporate norms or social norms. At present these are normally of a deterministic nature. However, the norms within which human agents act are of a richer nature. To capture these we have incorporated deontic operators into explicit rules. Human agents operate <u>within</u> the limits of *framing norms* (including corporate norms and contractual norms). However, their discretion within these limits is <u>governed</u> by meta-norms (equity, fairness, etc.).

Intelligent agent technology has great potential in commercial and industrial applications. But many research issues have to be addressed before it advances much further. Two of them, elicitation and specification of norms, are of immediate relevance to the development of CSCW.

Acknowledgements

Many thanks to Fiona for checking the copy and correcting some of our worst use of English.

References

Layman E. Allen and Charles S. Saxon (1993) A-Hohfeld: A Language for Robust Structural Representation Knowledge in the Legal Domain to Build Interpretation-Assistance Expert Systems. In H.C. Meyer & R.J. Wieringa (eds) Deontic Logic In Computer Science, John Wiley & Sons, Chichester.

H. H. Clark and S. E. Brennan (1991). Grounding in Communication. *Socially Shared Cognition*, Eds. L. B. Resnick, J. Levine and S. D. Behreno.

A. Cypher (1991). Eager: programming repetitive tasks by example. *Proceedings f CHI'91*, New Orleans, ACM Press. pp. 30–39.

A. Dix (1992). Human issues in the use of pattern recognition techniques. *Neural Networks and Pattern Recognition in Human Computer Interaction*, Eds. R. Beale and J. Finlay. Ellis Horwood. pp. 429-451.

A. Dix, J. Finlay, G. Abowd and R. Beale (1993). *Human–Computer Interaction*. Prentice Hall.

A. Dix, J. Finlay and J. Hassell (1994). Environments for cooperating agents: Designing the interface as medium. *CSCW and Artificial Intelligence*, Eds. J. Connolly and E. Edmonds. Springer Verlag, pp23-37.

A. J. Dix (1994). Computer-supported cooperative work — a framework. *Design Issues in CSCW*, Eds. D. Rosenburg and C. Hutchison. Springer Verlag, pp9-26.

A. Dix and A. Patrick (1994). Query By Browsing. *Proceedings of IDS'94: The 2nd International Workshop on User Interfaces to Databases*, Lancaster, UK, Springer Verlag. pp. 236–248.

David Eichmann (1994) Ethical Web Agents. *Electronic Proceedings of the Second World Wide Web Conference,* http://www.ncsa.uiuc. edu/SDG/IT94/Proceedings/Agents.

W. Gaver (1992). The affordances of media spaces for collaboration. *CSCW'92*, Toronto, Canada, ACM Press.

T. Kreifelts, E. Hinrichs, K.-H. Klein, P. Seuffert and G. Woetzel (1991). Experiences with the DOMINO Office Procedure System. *Proceedings of ECSCW'91*, Amsterdam: Kluwer, pp117–130.

C. Krogh (1995) The Right of Agents. *Intelligent Agents II: Agent Theories, Architectures, and Languages,* M. Wooldridge, J.P. Muller & M Tambe (eds), Springer, Berline, pp1-16.

K. Liu (1993) *Semiotics Applied to Information Systems Development*, PhD thesis, University of Twente, Enschede.

T. W. Malone, K. R. Grant, K. Lai, R. Rao and D. Rosenblitt (1987). Semistructured messages are surprisingly useful for computer supported coordination. *ACM Transactions on Office Information Systems,* **5**(2), pp115-131.

H.S. Nwana and D.T. Ndumu (1996) An introduction to agent technology. *BT Technology Journal*, 14(4), pp55-67.

W. Prinz and S. Kolvenbach (1996). Support for Workflows in a Ministerial Environment. *CSCW'96*, Lancaster, ACM Press, pp199–208.

J.R. Searle (1969) *Speech Acts - An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press,.

R. Smith and E.H. Mamdani (1996) Intelligent software technologies. *BT Technology Journal*, 14(4), 24-31.

R. Stamper (1980) LEGOL: Modelling Legal Rules by Computer. *Computer Science and Law*, Ed. B. Niblett, Cambridge University Press.

R Stamper & K. Liu (1994) Organisational Dynamics, Social Norms and Information Systems. *Proceedings of the Twenty Seventh Hawaii International Conference on System Sciences*, Hawaii, pp IV645-654.

Anand S. Rao (1995) Decision Procedures for Propositional Linear-Time Belief-Desire-Intention Logics, in M Wooldridge, J.P. Muller & M. Tambe (eds), Intelligent Agents II, Springer.

T. Winograd and F. Flores (1986). *Understanding computers and cognition: a new foundation for design.* Addison-Wesley, New York.

M. Wooldridge and N. Jennings (1995) Intelligent agents: theory and practice. *The Knowledge Engineering Review,* 10(2), pp115-152.