

Don't Lose Sight of the Forest

Why the Big Picture of Social Intelligence is Essential

Emma Norling

School of Computing, Mathematics and Digital Technology
Manchester Metropolitan University
United Kingdom
norling@acm.org

ABSTRACT

This paper argues that “social intelligence,” a hot topic in this community in recent years, is still not properly understood, and that this has important implications for multi-agent systems. The term shows up repeatedly in the intelligent agents and distributed artificial intelligence literature over many years, but in most cases the authors refer to just limited sub-components of social intelligence, such as norms, or trust. It is argued here that these (and other) facets of social intelligence are inextricably linked, and that we should only attempt to model them in isolation if we know what the implications of ignoring the other facets will be on model outcomes. We do not necessarily need to *implement* a “big picture” of social intelligence for every agent system, but having a conceptual model of this big picture will help us to understand the implications of what we *do* choose to model. This paper sketches a possible conceptual model of social intelligence, and begins to untangle the implications for different facets of social intelligence that are commonly scrutinised by the multi-agent systems community. The community is challenged to go further in developing this model and understanding of how its various facets make up the whole.

Keywords

social intelligence; human behaviour modelling

1. INTRODUCTION

Computational social intelligence is a topic that has received considerable attention in recent years (e.g. [6, 10]), but the concept of “social intelligence” has had attention in the intelligent agents community for at least thirty years [7]. However when you look at the usage of the term, the way in which it is used often focuses on just a narrow sub-view of what social intelligence really is, such as social structure [2], social situation [4], norms [1, 5], and trust [3, 12]. While each of these things is no doubt an aspect of social intelligence, none on its own fully captures the entirety of the problem. Thus while a superficial scan of the literature seems to indicate that we as a community have been studying this problem for decades, in this paper it is argued that we still do not

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.

Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

understand the big picture of social intelligence; we cannot see the forest because of all the trees. In fact we could go one step further than that, saying that we are not just looking at the trees, but individual aspects of a tightly coupled ecosystem, and that we should not try to understand any one aspect without considering its relationships with all other elements of the forest.

It is argued here that we should start with a conceptual model of social intelligence, giving us understanding of how the various aspects influence each other, before we attempt to model these aspects in isolation. Social intelligence is not necessarily relevant to some theoretical areas of multi-agent system, or even some problem domains, but when any aspect of it is, it is worth considering if other aspects might also be. It is probably not necessary – in fact it would be probably be overkill – to implement *every* aspect of social intelligence within an agent framework, but understanding the interplay between them should make the implications of ignoring particular aspects of social intelligence more clear. For example, what is the implication of considering social situation while ignoring social structure, or the implication of considering norms without taking trust into account? The answers to these questions have obvious implications for systems involving or modelling human behaviour (e.g. socio-technical systems, or social simulation). There may also be subtle lessons that could be learned for purely technical systems (e.g. agent-mediated e-commerce) that integrate aspects of social intelligence, in that they may suggest future developments for such systems.

This paper starts with a definition of social intelligence from the social psychology literature, sketching a conceptual model based on this definition. The difficulty is that the social psychologist’s perspective is that of the forest, the big picture of social intelligence, without much of the detail. The next stage is to begin to understand how various aspects of social intelligence – those mentioned above and others – relate to each other within this framework. Some steps towards this are presented here, but this is only a beginning, and the wider community is encouraged to contribute to this endeavour, and to consider the implications for all types of multi-agent systems.

2. SOCIAL INTELLIGENCE

The term “social intelligence” was coined by Thorndike in 1920 [16], and perhaps the earliest clear definition was given by Vernon, who described social intelligence as

“ability to get along with people in general, social technique or ease in society, knowledge of

social matters, susceptibility to stimuli from other members of a group, as well as insight into the temporary moods or underlying personality traits of strangers.” [17]

This definition breaks down into five components:

1. **Ability to get along with people in general.**

This refers to people’s general ability to “rub along” with others; to co-exist in relative harmony despite different beliefs and goals.

2. **Social technique or ease in society.**

This is the ease with which an individual interacts with others. It is closely linked to the individual’s self-image, and also understanding of relationships between themselves and others.

3. **Knowledge of social matters.**

Societies have expected behaviours and standards, and social intelligence implies a knowledge that these exist, and an ability to work with them.

4. **Stimuli from other members of a group.**

A socially intelligent person should be able to recognise and respond to cues from other societal members, particularly those related to the expected behaviours and standards that were previously mentioned, but also those related to the individual needs of fellow group members.

5. **Insights about strangers.**

Underlying this point is an assumption that the individual recognising other individuals being in some way “like self.” However here it is recognised that while at one level others are similar, they also have differences based on temporary moods and personality traits, and a socially intelligent individual will recognise at least some of these differences.

Amongst these five points, we start to see aspects of social intelligence that are often highlighted in the multi-agent systems literature. For example, points three and four relate closely to norms and normative behaviour, while point one has a strong relationship with work on negotiation. But the point is that all five *together* combine to give social intelligence, and if we examine them in isolation, we may miss important influences on behaviour. There are also other things that are commonly regarded as aspects of social intelligence, such as trust, which are not even mentioned. Does this mean that these aspects are not really part of social intelligence? Or is this definition in fact not sufficient? The assumption made here is that in fact these aspects of social intelligence are hidden in the detail of the definition, and that we should be able to identify them as we develop a conceptual model.

3. DEVELOPING A CONCEPTUAL MODEL

Based on the assumption that we will be able to identify particular aspects of social intelligence as we develop the conceptual model, Figure 1 begins to sketch just what such a model might look like.

Three components have been added to the agent itself: self, normative framework and theory of mind(s) (ToM(s)).

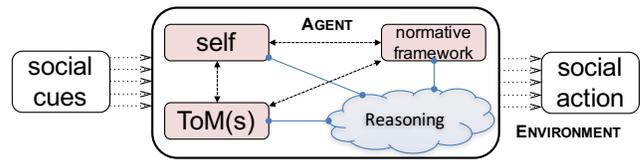


Figure 1: A Sketch of the Conceptual Model

These three components all influence each other, and most importantly, they influence the agent’s reasoning. This reasoning could be one of many existing reasoning frameworks.

“Self” refers to the agent’s self-image and positioning (relationships) within society. This will moderate the way that an agent acts based upon its relationships and the image that it wishes to project.

The understanding of norms and normative behaviour is an intrinsic element of social intelligence, and this is what is required in the “normative framework” component. There has been extensive work in the multi-agent systems community on norms and normative behaviour that could be adapted for use here.

“ToM(s)” refers to the agent’s mental models of the other members of society. Social intelligence requires recognising other societal members as in some sense like-minded individuals, but also allowing for individual differences. The detail of the ToM at this stage is blurred, but it is important to remember that at this point it is a *conceptual* model that is being developed, for the purpose of better understanding social intelligence. Any actual *implementation* of a model of social intelligence may well contain far fewer details for pragmatic reasons. That said, what is needed here? First, it is worth noting that although every member of society can be recognised as an individual, the level of detail in the mental model for each individual will vary tremendously. Those “close” to the agent will typically be represented by quite detailed models, while those distant will have very simple models. These ToMs should moderate the agent’s behaviour in many ways. With knowledge of others, the agent may seek help to further its goals. It may also choose to support the pursuit of others’ goals, or even thwart the goals of others. These choices will be strongly influenced by the agent’s “self.”

Also in this diagram, outside the agent, are social cues and social action. These are elements of the environment in the diagram, but they are still linked to the agent itself, because even if social cues *occur* in the diagram, they may not be recognised by the agent. Similarly, although the agent’s social actions occur in the environment, the reasoning that generates those actions must be within the agent.

It is freely acknowledged that none of these elements are new ideas in the agents research community: examples of normative frameworks have already been discussed, and ideas regarding *self* and *ToM* are rich in both the affective agents research community and the cognitive modelling research community (see for example [9, 13, 18, 19]). The unique contribution is in using these elements to create a framework that can relate the various aspects of social intelligence.

4. UNTANGLING SOCIAL INTELLIGENCE

We now have a sketch of a conceptual model that can capture the given definition of social intelligence. There is

considerably more work to be done in order to fully develop this conceptual model, but let us return to the purpose of this model. The aim of developing this conceptual model is to have a better understanding of social intelligence, and in particular, how the various facets of social intelligence interact with and depend upon each other. This in turn will help us to understand the implications of ignoring some aspects of social intelligence when we model others. We can see already, even in this conceptual model we have three interdependent components, all influencing the agent's reasoning. If the assumption that other aspects of social intelligence will be identified within this framework is correct, it is difficult to imagine how they too will not be tightly coupled.

So what are the aspects of social intelligence that we might want to identify? A non-exhaustive list includes: norms, trust, reputation, negotiation, argumentation, responsibility, and obligations. These are all things which have been discussed in relation to social intelligence, and so should be able to be explained with reference to a conceptual model of social intelligence. A normative framework has already been included within the framework, so let us take the next item: trust.

4.1 Trust as an example

Can we see a way of untangling trust from the model? At this stage there is not sufficient detail in the conceptual model to completely achieve this. However even though we do not have the detail at this stage, it is worth considering whether it would be *possible* to understand trust in terms of this model of social intelligence. There are a myriad of existing models of trust, as evidenced in Yu et al.'s recent survey [20].

For the sake of the thought experiment, let us consider just one model, that of Singh [14]. Singh's model is based on the idea of trust as dependence, with a series of postulates about trust that are connected to an agent framework. In other words, trust is expressed in terms of connections between actions, a trust that one state will lead to some other state (where the first state is typically brought about by an action of the trustor, and the second state is typically – but not necessarily – brought about by an action of the trustee).

Conceptually, what is needed here are:

- Recognition of other entities.
- Recognition of the desired states. It is not completely clear in the described model whether it is necessary that the trusted agent brought about the desired state. (It could for example be brought about by the action of an independent party.) Nevertheless, regardless of the intention in the original work, the perceived contribution of the trusted agent to the outcome should probably contribute to trust.
- Commitment. This is the flip-side to trust: If agent *X* trusts agent *Y* to achieve a state given some pre-condition state, agent *Y* has a commitment to agent *X* that it will (try to) achieve that state given the pre-condition; if agent *Y* fails in that commitment, it will betray the trust that *X* has in it.
- Mutual progress. Whereas it was not clear in the model description whether an agent's contribution to a desired state influenced trust in the general case, for this case it is explicit. When agents trust each other

to achieve a common goal, they expect to see their trustee working towards that goal.

- Teamwork. Although teamwork is explicitly discussed in presentation of the model, teamwork itself is embedded in a subset of the trust postulates. Thus to implement this model of trust it is not necessary to explicitly account for teamwork; support for (at least this view of) teamwork should arise if the other aspects of the trust model are accounted for.
- Scepticism. One agent's trust of another agent is lowered if that trusted agent fails to achieve the outcome state in an activated connection. (For example, if agent *X* trusts agent *Y* to send goods upon payment, and agent *X* has made payment but agent *Y* has not sent goods, agent *X*'s trust in *Y* is lowered.)
- Faith. This is essentially the opposite of scepticism: one agent's trust of another agent is increased each time the state needed to be achieved by trusted agent in an activated connection is achieved. (So in a buyer-seller relationship, the buyer's trust in the seller is achieved each time goods are successfully shipped.)

The first two points are already explicit in the conceptual framework: other entities are present as ToM models, and also in the relationships in the “self” component. Recognition of social cues will capture the recognition of working towards desired states. Commitment is not explicit in the conceptual framework, but may actually be part of the underlying reasoning framework. Many reasoning frameworks (for example the BDI framework [8]) already include commitment as an intrinsic element, but in many cases this is a “light-weight” version of commitment: not a commitment to another agent, but a commitment to achieving a goal. To fully capture this idea of commitment to another, the relationships with other agents would again come into play, such as discussed by Singh [15]. It is not necessary to explicitly account for teamwork with this model of trust, as discussed previously, so the remaining two points to be addressed are faith and scepticism. Both of these relate to the degree of trust an agent has in another, or in other words, it is part of the relationship with the other agent. This can even take into account the fact that you might trust an individual in some contexts but not others, such as the example in the original paper: “you can trust your friend to take your side in a dispute, but not against his employer.” [14]

We can see from this thought experiment that it should be possible to capture *this* model of trust within the conceptual framework, with the addition of some detail (regarding a measure of trust and its manipulation, plus commitment to others) to the relationship structures. As it happens, this model is neatly constrained within the “self” component of the conceptual model, other than the recognition of social cues. This does not mean however that it is neatly isolated from other aspects of social intelligence, as these too may use and manipulate the same elements. Even if we just look at the next item in the list, reputation, it quickly becomes obvious that these are closely related. (Indeed, trust and reputation are often tightly coupled, see [11] for a survey.)

4.2 The Bigger Picture

The point here is that although we were able to demonstrate that a particular model of trust *could* be implemented

within the conceptual framework, it still doesn't capture how trust relates to other aspects of social intelligence. In order to this, we need a categorical list of what these aspects are, and how each and every one of them fits into the conceptual framework. One of the challenges here is that many aspects of social intelligence – trust being the perfect example – have not just one but many potential models that could be considered. Given this variation, can we be sure that we will correctly identify the links between different aspects of social intelligence if we map a particular computational model of an aspect to the conceptual framework? It would be better perhaps to again take a step back, identifying the important aspects of social simulation from the multi-agent systems literature, but turning to psychology and sociology for explanations of their relationships to social intelligence. Once this bigger picture has been developed, we will have a yardstick against which we can consider particular computational models of aspects, such as the one discussed above.

5. THE CHALLENGES

This paper has presented a problem – that we are losing track of the big picture of social intelligence as we focus on specific aspects of it – and started to sketch a solution to this problem. This big picture of social intelligence is important because the many different aspects of it that we study are in fact interrelated, and if we do not understand these relationships, we cannot be sure of the implications of studying a single aspect (or set of aspects) in isolation. Understanding the relationships will give us a better understanding of the scope of any of these models, and also insights into how we could grow and enhance them.

There are a series of challenges presented in this paper in which the wider multi-agent systems research community is invited to take part:

1. Development of the high level conceptual model of social intelligence. Only a rough sketch has been provided here, with many details to be completed.
2. Creation of a comprehensive list of the aspects of social intelligence in which we as a community are interested.
3. Development of a relational map of each of these aspects of social intelligence within the conceptual model, particularly highlighting their relationships with other aspects of social intelligence.

Each of these three challenges is non-trivial, but the resultant conceptualisation will be invaluable to the community.

REFERENCES

- [1] M. Boman and H. Verhagen. Social intelligence as norm adaptation. In B. Edmonds and K. Dautenhahn, editors, *Socially Situated Intelligence: a workshop held at SAB'98*, pages 17–24, Zürich, 1998.
- [2] C. Castelfranchi. Modelling social action for AI agents. *Artificial Intelligence*, 103(1–2):157 – 182, 1998.
- [3] C. Castelfranchi and Y.-H. Tan. The role of trust and deception in virtual societies. *International Journal of Electronic Commerce*, 6(3):55–70, 2002.
- [4] R. Conte. Social intelligence among autonomous agents. *Computational and Mathematical Organization Theory*, 5(3):203–228, 1999.
- [5] R. Conte. Agent-based modeling for understanding social intelligence. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7189–7190, 2002.
- [6] F. Dignum, R. Prada, and G. J. Hofstede. From autistic to social agents. In *The 2014 international conference on Autonomous agents and multi-agent systems*, pages 1161–1164. IFAAMAS, 2014.
- [7] L. Gasser. 1985 DAI workshop report. *AI Magazine*, 8(2):91–98, 1987.
- [8] M. Georgeff, B. Pell, M. Pollack, M. Tambe, and M. Wooldridge. The belief-desire-intention model of agency. In J. P. Müller, M. P. Singh, and A. S. Rao, editors, *Intelligent Agents V, Proceedings of 5th ATAL Workshop*, volume 1555 of *Springer LNAI*, pages 1–10. Springer, 1999.
- [9] J. Gratch. Understanding the mind by simulating the body: virtual humans as a tool for cognitive science research. In S. Chipman, editor, *Oxford Handbook of Cognitive Science*. Oxford University Press, Oxford, UK, in press.
- [10] G. A. Kaminka. Curing robot autism: a challenge. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 801–804. IFAAMAS, 2013.
- [11] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.
- [12] M. Schillo, P. Funk, and M. Rovatsos. Using trust for detecting deceitful agents in artificial societies. *Applied Artificial Intelligence*, 14(8):825–848, 2000.
- [13] M. Shayganfar, C. Rich, and C. Sidner. An overview of affective motivational collaboration theory. In *Proceedings of AAAI Workshop on Symbiotic Cognitive Systems*, 2016.
- [14] M. P. Singh. Trust as dependence: a logical approach. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 863–870. IFAAMAS, 2011.
- [15] M. P. Singh. Commitments in multiagent systems: Some history, some confusions, some controversies, some prospects. In F. Paglieri, L. Tummolini, R. Falcone, and M. Miceli, editors, *The Goals of Cognition. Essays in Honor of Cristiano Castelfranchi*, pages 601–626. College Publications, London, 2012.
- [16] E. L. Thorndike. Intelligence and its uses. *Harper's Magazine*, 140:227–235, 1920.
- [17] P. Vernon. Some characteristics of the good judge of personality. *Journal of Social Psychology*, 4:42–57, 1933.
- [18] H. de Weerd, R. Verbrugge, and B. Verheij. Negotiating with other minds: the role of recursive theory of mind in negotiation with incomplete information. *Autonomous Agents and Multi-Agent Systems*, pages 1–38, 2015.
- [19] A. B. Youssef, M. Chollet, H. Jones, N. Sabouret, C. Pelachaud, and M. Ochs. Towards a socially adaptive virtual agent. In *Intelligent Virtual Agents*, pages 3–16. Springer International Publishing, 2015.
- [20] H. Yu, Z. Shen, C. Leung, C. Miao, and V. R. Lesser. A survey of multi-agent trust management systems. *Access, IEEE*, 1:35–50, 2013.