



GRASP agents: social first, intelligent later

Gert Jan Hofstede¹

Received: 24 August 2016 / Accepted: 18 November 2017 / Published online: 26 December 2017
© The Author(s) 2017. This article is an open access publication

Abstract

This paper urges that if we wish to give social intelligence to our agents, it pays to look at how we acquired our social intelligence ourselves. We are born with drives and motives that are innate and deeply social. Next, as children we are socialized to acquire norms and values and to understand rituals large and small. These social elements are the core of our being. We capture them in the acronym GRASP: Groups, Rituals, Affiliation, Status, Power. As a consequence, economic rationality or logical reasoning do not suffice when it comes to social intelligence. Basic features of our social behaviour, of the kind that one sees early in the lives of children, need to be prominent. These include fear, love, and aggression. They extend to the combination of these drives with basic social clues from the environment such as big and small, good and bad, as well as culture-specific specializations of these. This will make agents respond differentially to inferred attributes such as gender, age, group membership. This level of universality in social intelligence should receive our full attention. The general insights gained can then be re-used in myriad implementations to specific modelling issues.

Keywords Virtual agents · Volition · Drives · Social intelligence · Status-power theory · Culture

1 Introduction

“Imagine a superhero and a police agent facing a villain. There is also somebody else: a girl has been taken hostage by the villain, and although the villain thinks she is just a girl, she is indeed the sidekick of the superhero, who knows her moves”.

This story, in more elaborate form, is the example used by Felli et al. (2014) that they employ to discuss their sophisticated model of social intelligence, in particular two modes of reasoning: projection (reasoning as others) and stereotyping (reasoning about others). In this paper we shall argue that in real people these kinds of reasoning do take place, but not in a void. People live in a social landscape. Their thinking is underpinned by a system of social perceptions and drives. We therefore take a step back and consider perception of the social world rather than reasoning. We thus hope to contribute to the issues raised in recent discussions (Dignum et al. 2014). We hope to provide some key elements

for what pioneer Cristiano Castelfranchi (2006) termed an ‘architecture of mind’, limiting ourselves to the social.

Suppose you had to explain this story to 6-year-old children that came from a place where the social role of police agents and girls was very different from what they are in Australia (where these authors are based). How would you explain? You would probably start by telling them who was big or small, strong or weak, good or bad, and perhaps who was a boy or a girl. But would that answer all their questions?

These questions might very well be about ‘why’ and ‘with whom’. They might ask whether you knew the superhero, why the superhero was so strong, who was the boss, who was friends or family with whom, whether the girl was the child of one of the others, why there was no boy. They would ask about the things that we shall call GRASP: groups, rituals, affiliation, status and power.

Indeed, children from different continents might definitely have widely different pre-conceived ideas. For instance, is a policeman good or bad? Is he ‘one of us’ or ‘one of them’? ‘Policeman’ could be just another word for ‘villain’. Or take the girl: does it make any sense for a girl to be a ‘sidekick’ of a superhero? What would a sidekick do? Is it a sort of wife, or sister, or child?

✉ Gert Jan Hofstede
gertjan.hofstede@wur.nl

¹ Social Sciences, Information Technology Group,
Hollandseweg 1, 6706 Wageningen, The Netherlands

I am introducing and discussing this example here since it is from a recent and up to date article, but I could have used any of a number of rather similar examples to point out that the examples used in AI papers tend to assume social knowledge on the part of their readers that cannot be taken for granted in the real world.

The point of using the children is that AI agents are like children, but even more ignorant, in the sense that they do not know any of the preconceived categories from any civilization. The designers have to teach agents everything about their world. There is also a difference: real children have to learn to understand the full complexity of the world, while agents typically live in a very narrowly finite world. This will have to change though now that agents are starting to be used in socially richer contexts such as health care (Johnson et al. 2014) or cross-cultural training (Degens et al. 2014).

In this essay-style paper I explore how the comparison between children and agents can inform the design of socially intelligent agents. Learning in children is introduced. Then, the big picture is used to search for theory that can allow agents to socialize. There is a discussion that touches on levels of analysis, on emotions, norms and values, on language and embodiment. We revisit the villain–superhero story in terms of the GRASP model primitives from the literature. A brief conclusion ends the article.

In order to arrive at a big picture, the paper takes giant leaps. It does so at the price of superficiality. Any reader wishing to use this paper for designing social agents is left with a lot of work to do. The paper offers a helicopter view, but no road map.

2 The primacy of the social over the intellect

A child is handicapped more by lack of social capacities than by lack of intellect. Suppose that we wish to create more socially intelligent agents. Can we educate our agents as we raise our children? Children are born with certain capabilities, and then raised by their caretakers and the wider society to bring these capacities to fruition. A new-born wants to be fed and held. As it grows up the range of needs becomes wider, but the drive to have them satisfied remains. The first thing that a baby learns is to express intentions and to respond to the intentions of others, e.g. by smiling, crying, babbling or becoming agitated. In this way a baby can show whom it loves and whom it fears. It also develops empathy. These innate capacities precede the development of basic motor skills as well as language skills.

Then, ‘it takes a village to raise a child’. Every day, while attempting to satisfy its drives, the child is subjected to endless interactions and also performs millions of experiments by behaving in a certain way and receiving feedback. This is how it learns which behaviours are rewarded and which are

punished, which distinctions are meaningful (such as big/small, good/bad, boy/girl, clean/dirty) and which categories or groups of people should be treated differently from others.

It is thus that our children acquire a mental architecture of social life that underpins and precedes their reasoning in any particular situation. This includes the things captured by the acronym GRASP: group identity, rituals, affiliation, status, and power. The question now becomes whether we can use this in developing socially intelligent agents. Can we give agents a coherent basic architecture of social life? Ideally, such social intelligence could be re-used for all kinds of applications, notably virtual agents and robots. Naturally, such applications would require supplementing social intelligence with other capacities, such as sensory skills, language skills and motor skills.

The child-agent metaphor could be pursued in various ways. One could try to mimic the years-long inductive learning process. This invites the use of neural networks, mimicking brain circuitry. Here we shall take the simpler stance that we might socialize our agents in one fell swoop by inculcating drives, as well as a full-fledged architecture of the social world, in their minds. Concepts and theories used should be so general that we could build our models of social intelligence on them without needing to assume all kinds of preconceptions. They should build on a limited number of primitives such as status, power, fear, aggression, relationships, and love. Selecting which primitives to use and designing a meta-model containing them is a crucial task that the theories should help us with.

Can we do this without recourse to neuroscience? We believe we can, just as Darwin came up with his theory without having any knowledge about genes. One can model a phenomenon without knowing the proximate mechanisms that it operates through.

3 Basic theory for social intelligence

So many disciplines exist in the social sciences, and so many theories in each of them, that it might at first seem a hopeless task to select theories that can help us model socially intelligent agents in the generalist sense defined above.

We propose the following criteria by which to select theories:

- Distinguish levels of analysis that have shown in practice to be relevant to social life and select theories at the most important levels;
- Select parsimonious theories, so that they do not complicate but simplify the modeller’s life;
- Use theories that have proven their real-world relevance;
- Use theories that can be combined into a coherent meta-model.

Levels of analysis could be the ones targeted by, in order of descending universality: biology, sociology, cross-cultural psychology, social psychology, individual psychology, neuroscience. This paper will limit itself to three of these.

- First of all there is the level of ‘all people’ dealt with by biology and sociology. We are social mammals of a clever, cultural kind. We categorize people in groups, we recognize rituals, we are driven by motives of affiliation, status and power. In short we are ‘GRASP’ creatures and our social lives are GRASP games (‘groups, rituals, affiliation, status, power’), the rules of which are unwritten but deep-seated.
- Second, there is the level of a single group or category of people. This level involves cultural and social identity issues and is quite intricate, because there are so many categories and groups to which an individual can belong. In each of these, as well as between them, we play the GRASP game.
- Third, there is the level of the individual. This involves individual variation in motives, attributes and capabilities.

I believe these three levels to be essential, and a good starting point. The text will introduce theories at each of these levels that can be used for modelling agents that are socially intelligent in a generic way.

The theories chosen here are examples that fit the criteria. Other theories at each level could work equally well, or even fit some purposes better.

3.1 All people: sociology

Sociology can tell us things about the dynamics of social life, generalized across people. They shed light about what motivates people to engage in social interaction.

A theory about social interaction is status–power theory by Theodore Kemper (Kemper 2011). The theory grants people basic drives, allows for individual deliberations, is sparse, and generic across all people. I believe this theory to be eminently suitable for agent world design. It posits that people are driven to love and be loved, respect and be respected. Backed up by years of study as well as by empirical support, Kemper posits that people are driven to both claim status for themselves and confer status upon one another. His concept of status is very generic. Kemper uses the word ‘status’ for all forms of attention, respect, esteem, love, politeness etc. People engage in this mutual status–conferral game in a context of reference groups. They will commit to those groups from which they expect the best conferral of status. The result of this game is that in any group the current status of an individual is a sort of running average of the conferrals made to that individual. Status has

a process side (conferrals and claims happen) as well as a result side (a person assumes a certain status in a group).

Because of its basis in group life, status–power theory fits with social identity theory. The second tenet of the theory is power, or involuntary enforcement. Only when someone receives less status than they feel they deserve will they be tempted to resort to power moves. Status and power moves can take thousands of forms, discussed at length by Kemper. There is no room here to do full justice to the theory, particularly the parts on authority, roles and rituals.

Kemper himself stresses the unifying potential of status–power theory. He describes derivations for theories of emotions, ideas, social identity, and motivation, among others. He also indicates that much of what we do could be seen as a cover-up of underlying status–power motivations, to which there is a coherent system. It is that underlying system that can be the focus of agent-based models.

Though Kemper explicitly denies pretences to biological validity, his theory pleases the population biologist writing this paper. Ethologists of nonhuman primates and other social animals always consider status–power dynamics, whether in observations (De Waal 2007), or in agent-based simulations (Hemelrijk et al. 2005). They may use different terms such as dominance, aggression, submission, grooming; these are all status–power phenomena.

Kemper’s work has been used in ABM before (Gert Jan Hofstede et al. 2015). This simulation reproduces gender-based status differences based on cultural acceptance of power use. Agents have properties directly linked to status–power theory. Their ‘beauty’ determines how much status they attract. Their ‘kindness’ determines how much they confer and how easily they are provoked. If given too little status an agent may pick a fight. Their ‘power’ determines how strong they are in case of a fight. As a result of their interactions the agents have an evolving public social ‘status’. This simulation thus includes the phenomenon of emergence highlighted by Castelfranchi (2006) as crucial.

A related theory, smaller in scope of application but more detailed, that has recently been used in an agent-based model is Affect Control Theory by US sociological social psychologist David Heise (Heise 2013). Heise posits that people trade ‘affective meaning’ during group meetings. Affective meaning effectively consists of a perceived status conferral element (‘evaluation’, ranging from pleasant to unpleasant) and a power element (‘potency’ ranging from powerful to powerless), as well as an ‘activity’ element that is less easily interpreted in Kemperian terms. Heise used transcripts of US jury meetings in which the utterings had been classified according to ‘Interaction process analysis’, which distinguishes classes such as ‘shows solidarity’, ‘disagrees’, ‘shows antagonism’. With his agent-based model, Heise was able to convincingly reproduce the distribution of utterings by participants, including gender patterns.

3.2 Groups: cross-cultural differences and social identity

3.2.1 Culture

Culture at the deep level is about the social issues that societies contend with and have solved each in their own way (Hofstede et al. 2010). This concerns ways of dividing social goods that people are driven to strive for: leadership, status, power, love and belonging. It operationalizes the ways in which agents, or people, live in different worlds. Here ‘worlds’ can refer to societies, countries, or other larger groups that exist from birth. Down to the most basic psychological phenomena, such as the relative importance of the basic drives and distinctions, behaviours are culture-specific (Smith et al. 2006). The earlier in life one becomes a member of a group or category, the deeper its effect on culture. People from different parts of the world do not live for the same things; culture resides in their hearts.

Models of culture that allow implementing agent worlds with different cultures have to be comparative. There are a number of these, with different empirical bases and different track records when it comes to describing or predicting society-level phenomena. The theory that has proved most useful so far is the one by Hofstede, including additions made by others in recent years (Hofstede et al. 2010). It has stood up to many replications by many people, shows continued validity over the years across many domains of life (Kirkman et al. 2006) and is extensively used both in research and in practical disciplines such as management and organizational behaviour. The concept of dimensions of culture is amenable to formal models.

3.2.2 Social identity

There is overwhelming evidence (Brown 2000) that people, from wherever in the world they might be, are prone to dividing the social world into ‘we’ and ‘they’, and that this is a process that happens at many scales of time and intensity. We have a family, a people, various groups of very different persistence and closeness. We tend to prefer those groups that will provide us with social status.

It is surprisingly easy to call social identities into existence. At the most ephemeral level we can don yellow and red shirts and immediately change our behaviour accordingly. People understand group boundaries and commit to roles in groups. Infants tend to consider everyone as a group member, but are quick in seizing cues to the contrary from their caretakers. Such capacities for perceiving us/them distinctions would much benefit virtual agents.

Social identity has been used for agent design (Prada and Paiva 2005) but the search for good models is by no means finished (Dignum et al. 2014; Gert Jan; Hofstede

et al. 2015). The implicit default in the agent world so far is that there is just one group; social identity is usually simply disregarded. In a simulation that includes agents with norms this becomes problematic. Norms vary by group. Adherence to, or violation of, norms are typically used by people as social signals, not just about the validity of these norms, but also about group membership (Ferreira et al. 2013).

3.2.3 Social identity, culture, and change

Social identity can sometimes coincide with culture, but this is not necessarily so. People can choose their social identity, change it, and use it strategically. All this does not hold for culture. So in agents, culture would consist of parameters shared between agents of a culture but not visible to others. Agents from different cultures would differ in the parameters of their basic architecture of social life. Social identity would also be shared between members of a group, but contrary to culture, it would normally be visible to other agents.

Culture creates many problems in social interaction precisely because it is unconscious and invisible, so that cross-cultural interactions are rife with misunderstandings. Differences in social identity create inter-group dynamics. The conjunction of these two phenomena is a very good arena for model development in virtual agents.

Culture is part of human nature. It self-organizes, emerging from myriad everyday interactions, and in the real world, a deep societal culture cannot be consciously created or changed. Our cultures are the unwritten rules according to which we self-organize our societies (Gert Jan Hofstede 2015). Most of the change that we observe from hype to hype, or from year to year, happens as pattern caused by cultural rules and values. Some of that rapid change in practices is captured in unspoken norms that are commodities around which group membership issues are played out, for instance what to wear to be considered important in a certain group, or how to greet various categories of people to show appropriate respect. These changes in norms affect practices all the time, in the form of fashions or hypes. They leave the underlying values comparatively unaffected. For instance, the level of respect to grant to certain categories of people (parents, strangers, ...) when one meets them, changes much less than the ways in which to convey such respect (Hofstede et al. 2010).

Such changes in norms function as signifiers of social identity. They are waves rippling the sea of culture, the deep levels of which are unaffected. Social identities can change on a scale of days or years, without necessarily altering the cultural value system of the society in which they occur. In consequence, agents with cross-cultural skills should distinguish social identity from culture.

3.2.4 Rituals

The behaviours of people are not haphazard. They hang together in sequences that have symbolic meaning. The word rituals captures this. Rituals mediate the social identities of people. They can range from major rituals that change relations between groups or membership of groups, such as signing a peace treaty or getting married, to minute rituals such as mutual nodding between individuals who meet on the street (Rothenbuhler 1998). Rituals can be useful for social agents precisely because a ritual as a whole carries a social meaning in terms of status conferrals or power moves.

Rituals also come with authority vested in roles defined by a ritual, and thus they structure the status–power world. They confer power without the need for confrontation. Pointers about how to model rituals are provided in an earlier paper (Gert Jan Hofstede et al. 2011).

3.3 Individual differences and cognition

The ideas put forward so far offer only the roughest way to tell agents apart, based on variables related to status and power. This is often good enough. Depending on a simulation's aim, designers need not be concerned about the precise differences between agents at individual level, as long as the individuals have basic human drives, plus some heterogeneity.

Likewise, the ideas in this paper are pre-cognitive. All cognitive processes are flattened until only their status–power intention and effect remains.

The reason for making this choice is that we believe the social precedes the cognitive; so in order to have intelligent agents that are not social idiots we need to model the social first. We can then build cognitive capacities and individual differences on top of that social layer.

4 Creating socially intelligent agents

4.1 Embodiment

How much change would there be in the need for social intelligence depending on whether it was meant for a robotic application, a realistic 3-D virtual world, or a simple 2-D grid?

Embodiment changes many things. One has to link the physical world to the social. If the superhero wants to kill the villain, he needs a heavy, solid object, such as a stick—is one to be found in the environment? In a non-embodied world there would not need to be a stick, nor a dexterous hit that kills. That would not change the intention though, and the superhero's skill could be modelled by a probability of succeeding in killing the villain.

My position is that the agents' social intelligence should be applicable to any agent world, regardless whether it is embodied or not. Embodiment requires a lot of design decisions. These concern individual differences, social identities, status–power motives and cultures. All of these must be bound to the physical reality, which adds enormously to design complexity and required computational power but does not in any way limit the social intelligence needed.

Embodiment adds many things but removes none. The social can exist in a drastically simplified physical environment.

4.2 Basic distinctions and motives

What are essential differences that social agents need to be able to perceive? This depends on the aim they are created for. Here we lay down a basic set of distinctions. The deepest of these are universal, and operationalize status and power in the minds of agents: agents can 'be good', i.e. be deserving of status, in the eyes of those groups that they affiliate with, or they can be 'bad', the opposite. Agents have a 'kindness', which is a tendency to confer status, whatever the circumstances. Agents have a 'status' in the eyes of every group that perceives them, depending on the history of claims and conferrals the agent has with that group. This status determines the capacity of the agent to make others freely comply with its wishes. Depending on the case it could have other attributes that determine how much power it can wield, such as strength, skill or power, or a role in a ritual vested with authority. With power it can coerce others whether they like it or not.

Then there is a category of agent attributes that systematically vary in average value across cultures (Table 1). These are derived from the dimensions of culture found by Hofstede and others in large-scale cross-cultural comparisons. This is admittedly crude, and many applications will require subtler differences. It is a good place to begin though.

In order for a model to make social transactions happen between agents, they need a motivational system based on a limited set of innate drives. I have not found any theory that could be an improvement on Kemper's status–power theory in this respect. According to status–power theory, the agents' basic motive is to be 'good'. They will try to do whatever they believe will earn them status conferrals in the groups with which they affiliate. The theory is also open-ended enough to allow, even require, combination with other theories.

4.3 Groups and individuals

Social Identity Theory can be combined with the reference group element of Kemper, since his reference groups are social identities in the mind of an observer. This level needs

Table 1 Dimensions of culture, their extremes, and status-deservingness

Dimension	Extreme	Status-deserving
Individualism	Collectivistic	Be in-group
	Individualistic	Be independent
Power distance	Large power distance	Old age, ascribed status
	Small power distance	–
Masculinity	Masculine	Be strong/big
	Feminine	Be weak/small
Uncertainty avoidance	Uncertainty avoiding	Be clean/safe
	Uncertainty tolerant	Be relaxed
Long-term orientation	Long-term oriented	Be useful
	Short-term oriented	Reciprocate status conferrals
Indulgence	Indulgent	Follow impulses
	Restrained	Hide emotions

to be present almost in any case. In the superhero story, for example, there is a clear *we/they* distinction.

In order for an agent world to be generalizable beyond a single society, the level of culture needs to be present (Table 1). Culture can be modelled as variations in status–power logic of perceptions, interpretation and action selection. For example, a ‘masculine’ culture according to Hofstede (e.g. Anglo cultures, China) is one in which the use of power to obtain what one wants is socially accepted and showing strength is a way to show status-worthiness. A ‘feminine’ culture (e.g. Netherlands, Scandinavia) is the opposite.

The level of the individual could be simplified to include only basic drives as per Maslow/Kemper, as well as distributions of variation so that agents are heterogeneous. Actual personalities are not often needed. McClelland’s motives or McCrae’s drives could be used to provide them.

Proof-of-concept implementations of agents that operate according to Kemper’s ideas in a cross-cultural context are described in recent work (Degens et al. 2012; Mascarenhas et al. 2013a, b, 2015). A strong feature of these agents is that depending on their culture they can reach different judgement regarding the appropriateness of behaviours.

4.4 Emotions

Emotions are an early warning system. They alert people about the dynamics of status-power issues such as opportunities (hope, anticipation) or threats (fear). They are also linked to the fulfilment of status conferral needs (love) and to status withdrawals (anger, hatred). For an agent to have emotions it requires distinguishing its ideal state from the actual state of its status–power world. The agent also needs to know whether a path from actual to ideal exists, could open, or could disappear.

For simple GRASP agents the most basic emotions would suffice. These include fear of others with high status

if they are in a ‘they’ social identity, love of others if they are deserving, tendency to reciprocate status conferrals, and anger at others that have used power against a ‘we’ group.

Emotions are also means of communicating status–power issues. More advanced agents would need explicit emotions that can be mutually perceived and used for theory of mind reasoning.

4.5 Norms and values

Norms and values are essential concepts for social agents (Dechesne et al. 2013; Dignum 1999; Dignum and Dignum 2015).

The term ‘values’ has been used at various levels of analysis. Many authors would classify values as attributes of individuals, whereas others see them as shared attributes of the members of a culture; Hofstede (Hofstede et al. 2010) calls values the unwritten rules of the game of a society shared between its members. Values direct our drives. As we have seen, a drive for social status will lead agents to behave very differently depending on the values that prevail in the groups to which they affiliate.

To social reality modellers, values will give a mapping from possible agent actions to their social correlates. In terms of deservingness of status, killing a villain is ‘good’ or ‘bad’ for instance. This mapping from specific to its relational meaning in a social intelligence meta-model could be implemented as a ‘counts as’ function (Searle 1995). So ‘killing a villain’ is one way to achieve a certain relational result; in this case a way to achieve more social status among the Good Guys. However, under a different value system, or perhaps under the same value system but if one were a member of the villain’s group, killing that villain could be a reprehensible crime that causes status loss.

For the term ‘norm’, all agree that a norm is shared by people. In real life norms are routinely used to mediate membership of social identity groups; one is urged to ‘behave’,

or to ‘be a Good Girl’. Many modellers use norms as a prescription for behaviour, others argue that one can deviate from a norm, and some note that there can be different kinds of norms: ideal-types of behaviour, versus actual behaviour. For instance, the ideal might be not to kill, but killing the odd villain does happen. Actually this can perhaps be better explained in terms of group identity: the norm for not killing only holds between members of one’s group.

The difference between values as discussed here at society level, and norms at the level of social identity, is that the latter are more changeable. Adherence to, deviation from and creation of norms are used as commodities to mediate group membership. They have the effect of preserving value systems. As Blaise Pascal formulated it: “Plus ça change, plus c’est la même chose.”. An analogous saying in English is “we watch the ripple, and take the lake for granted”. Norms are ripples, while values are the lake. The more norms change, the more they perpetuate the value system that underlies their dynamics. Recent developments in self-organising value-based agents are promising in this light (Dignum and Dignum 2015).

A rich model of social identity will require implementing norms as well as values. The former will be linked to specific agent behaviour at the level of social identity through counts-as logic (McBreen et al. 2011). The latter will be linked to the cultural level.

4.6 Language

A recurring controversy in AI is how much of human behaviour one can understand without language. A full article could be devoted to this. Briefly, the positions are as follows. On one side are those who maintain that language makes humans different in kind from animals; this holds e.g. for Michael Tomasello. On the other side are those who stress the similarities across creatures, finding only differences in degree between humans, bonobos, great whales, etcetera. This includes for instance Frans de Waal.

Historically, the strength of the first position has been on the decline for centuries, since times when not even all human beings were granted humanity (Fernández-Armesto 2012). In fact there could be a social identity issue at stake: are we prepared to consider ourselves as similar to apes, for instance? Each categorical barrier between humans and other creatures used up till today has been discarded on closer scrutiny. This includes language use. All kinds of creatures have been found to be remarkably good at conveying social intention, without using full-fledged languages of the kind I’m writing this article in.

In conclusion, it seems wise to adopt the position that, as far as social intelligence goes, language is a means of expression, not a source. This is in effect the position taken in the agent-based model discussed earlier (Heise 2013).

Therefore we can expect to make progress in socially intelligent agents even if they do not have natural language.

5 The example revisited

If we modelled the villain-superhero story using these basic theories, what could we get?

Of course this is a highly subjective matter, if only because anybody, including the present author, reads with a subjective lens influenced by individual experience, personality, group affiliation, social identity, and culture. For instance one academic Dutch reader of the manuscript commented that she suspected the superhero might be a mafia boss, hence a Bad Guy, and the sidekick a gullible chick, hence good but stupid—so she introduced another distinction, between smart and stupid. The attempt is to be as dispassionate as possible about it, and try to stay true to the intention of the original authors.

Let us first put it in simple terms. The crucial social knowledge for our four protagonists is “we are in a fight between two camps with no chance of peace”, “Who is on my side” and “How strong is everyone”. How does this translate to theories required?

- Individual: the characters are all archetypes, so individual personality is not needed.
- Social identity: there are two enemy group identities: the Good Guys, and the Bad Guys. Part good, part bad has no place. The villain is the only Bad Guy. There are no family relations. The superhero and the girl are friends. The sexual domain is implicit. They are probably sexual partners too, and the villain might aspire to possessing the girl, but this is kept outside the story.
- Sociology: when it comes to power, the superhero is the most powerful, followed by the villain, girl and police officer, perhaps but not certainly in that order. When it comes to status claims, the villain claims superhero status (maximally high) while the girl only claims ordinary status for a young girl (implicitly sexually attractive). Talking about status conferral tendencies, the superhero and girl are friends and will always protect one another from the power of others.
- Culture: at the level of culture, power dynamics are more important in this story than status dynamics. It is taken for granted that the Good Guys will support one another through thick and thin. Opponents are ready to use power against one another, in a violent manner. None respect others’ status to the point of obedience, i.e. to the point of refraining from power use voluntarily. In Hofstede’s terms we thus implicitly have a masculine culture.

Table 2 summarizes the beliefs and knowledge in GRASP terms.

Note that like Felli et al.'s model (Felli et al. 2014), it does not yet say anything about the more proximate aspects of modelling: personalities, capacities, the physical world; it could quite conceivably be paired with such a system though.

Felli et al.'s article contains some elements that could be developed in the direction of this paper's concepts. Their stereotypical reasoning, notably, could be elaborated in that direction. The fact that the villain holds the girl prisoner is not in the GRASP primitives until we create a ritual for it, and we would need to add 'hold prisoner' to the agent world, as a ritual. Another ritual can be freeing someone held. Holding an enemy and freeing a friend need to be status-worthy.

Now, our modelled world is in place. The superhero, driven to do status-worthy things, would discover a strong, most unworthy member of 'them' holding a friend from 'us'. He would develop a goal to free the girl. From here on we enter the realm of the theory of mind, 'thinking as others' and 'thinking about others', that is the main emphasis of Felli et al.'s article, and beyond our present scope.

6 Conclusions

This paper argues that socially intelligent agents need to be social first and foremost. These need to understand basic 'why' and 'with whom' elements of social behaviour, similarly to the way children learn about the social world. We summarize these elements in the acronym GRASP: Groups, Rituals, Affiliation, Status, Power. When these basics are modelled in satisfactory ways, we have generic social agents. They know, for instance, how to be polite and how to return a favour. They can recognize social primitives such as kindness, strength, group membership. Apart from that they are still naïve about the physical and instrumental world. These generic agent models can be used by others to build more cognitively sophisticated, situationally instantiated, 'what' and 'how' logics. Without such a generic level, the field will be left with a body of disparate applications.

Table 2 GRASP expression of the superhero story

	Superhero	Villain	Police	Girl
Groups	Us	Them	Us	Us
Rituals	Frees girl	Holds girl		
Affiliation	Likes girl			Likes superhero
Status	++	--	o	++
Power	++	+	o	o

The scale: ++/+/o/-/-- runs from maximal to minimal

This is a way of saying that the field of modelling social intelligence would profit from adopting some elements from the ontogenetics of social intelligence itself. Be social first, intelligent next. In terms of level of analysis, start with the big picture, not the details. These two elements may prove fertile directions for development of our field. The paper gives an example of how a GRASP architecture can underpin a theory of mind framework.

Acknowledgements The author is grateful to NIAS, the Netherlands Institute for Advanced Studies in the humanities and social sciences, for offering him a fellowship with Frank Dignum and Rui Prada. The manuscript has much benefited from the comments of Katy Hofstede, Sjoukje Osinga, and anonymous reviewers.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Brown R (2000) Group processes: dynamics within and between groups, 2nd edn. Blackwell, Malden
- Castelfranchi C (2006) Cognitive architecture and contents for social structures and interactions. *Cognit Multi Agent Interact* 355
- De Waal F (2007) Chimpanzee politics: power and sex among apes. JHU Press, Baltimore
- Dechesne F, DiTosto G, Dignum V, Dignum F (2013) No smoking here: values, norms and culture in multi-agent systems. *Artif Intell Law* 21(1):97–107
- Degens N, Hofstede GJ, McBreen J, Beulens AJM, Mascarenhas S, Ferreira N, Paiva A, Dignum F (2012) When agents meet: empathy, moral circle, ritual, and culture. In: Paper presented at the workshop on emotional and empathic agents (EEA) at the 11th international conference on autonomous agents and multiagent systems (AAMAS 2012)
- Degens N, Hofstede GJ, McBreen J, Beulens AJM, Mascarenhas S, Ferreira N, Paiva A, Dignum F (2014) Creating a world for socio-cultural agents. Emotion modeling: towards pragmatic computational models of affective processes, vol LNAI 8750. Springer, Berlin, pp 27–43
- Dignum F (1999) Autonomous agents with norms. *Artif Intell Law* 7:69–79
- Dignum V, Dignum F (2015) Value-sensitive design of self-organisation. In: Paper presented at the self-adaptive and self-organizing systems (SASO), 2015 IEEE 9th International Conference
- Dignum F, Hofstede GJ, Prada R (2014) Let's get social! From autistic to social agents. In: Paper presented at the 13th International conference on autonomous agents and multi-agent systems (AAMAS), Paris
- Felli P, Miller T, Muise C, Pearce A, Sonenberg A (2014). Artificial social reasoning: computational mechanisms for reasoning about others. In: Paper presented at the The International Conference on Social Robotics
- Fernández-Armesto F (2012) So you think you're human? Oxford University Press, Oxford
- Ferreira N, Mascarenhas S, Paiva A, Di Tosto G, Dignum F, McBreen J, Degens N, Hofstede GJ, Andrighetto G, Conte R (2013) An

- agent model for the appraisal of normative events based in in-group and out-group relations. In: Paper presented at the twenty-seventh AAAI conference on artificial intelligence
- Heise DR (2013) Modeling interactions in small groups. *Soc Psychol Q* 76(1):52–72. <https://doi.org/10.1177/0190272512467654>
- Hemelrijk C, Wantia J, Gygas L (2005) The construction of dominance order: comparing performance of five methods using an individual-based model. *Behaviour* 142(8):1037–1058
- Hofstede GJ (2015) Culture's causes. the next challenge. *Cross Cult Strateg Manag J*. <https://doi.org/10.1108/CCM-03-2015-0040>
- Hofstede G, Hofstede GJ, Minkov M (2010) *Cultures and organizations, software of the mind*, 3rd edn. McGraw Hill, New York
- Hofstede GJ, Mascarenhas S, Paiva A (2011) Modelling rituals for *Homo biologicus*. In: Paper presented at the ESSA 2011 (7th Conference of the European Social Simulation Association), Montpellier. <http://www.gertjanhofstede.com/research/>
- Hofstede GJ, Dignum F, Prada R, Student J, Vanh e L (2015) Gender differences: the role of nature, nurture, social identity and self-organization. In: Grimaldo F, Norling E (eds) *Multi-agent-based systems workshop (MABS) 2014*, LNAI 9002. Springer, Berlin, pp 72–87
- Johnson M, Bradshaw JM, Feltovich PJ, Jonker CM, Riemsdijk MBv, Sierhuis M (2014) Coactive design: designing support for interdependence in joint activity. *J Hum Robot Interact* 3(1):43–69. <https://doi.org/10.5898/JHRI.3.1>
- Kemper TD (2011) *Status, power and ritual interaction: a relational reading of Durkheim, Goffman and Collins*. Ashgate, Farnham
- Kirkman BR, Lowe KB, Gibson CB (2006) A quarter century of culture's consequences. a review of empirical research incorporating Hofstede's cultural values framework. *J Int Bus Stud* 37:285–320
- Mascarenhas S, Prada R, Paiva A, Degens N, Hofstede GJ (2013a) Can I ask you a favour?: a relational model of socio-cultural behaviour. In: Paper presented at the 2013 international conference on autonomous agents and multi-agent systems
- Mascarenhas S, Prada R, Paiva A, Hofstede GJ (2013b) Social importance dynamics: a model for culturally-adaptive agents. In: Aylett R, Krenn B, Pelachaud C, Shimodaira H (eds) *Intelligent virtual agents*, vol 8108. Springer, Berlin, pp 325–338
- Mascarenhas S, Degens N, Paiva A, Prada R, Hofstede G, Beulens A, Aylett R (2015) Modeling culture in intelligent virtual agents. *Auton Agent Multi Agent Syst*. <https://doi.org/10.1007/s10458-015-9312-6>
- McBreen J, Di Tosto G, Dignum F, Hofstede GJ (2011) Linking norms and culture. In: Paper presented at the second international conference on culture and computing, Kyoto
- Prada R, Paiva A (2005) Believable groups of synthetic characters. In: Paper presented at the AAMAS, fourth international joint conference on autonomous agents and multi-agent systems
- Rothenbuhler EW (1998) *Ritual communication, from everyday conversation to mediated ceremony*. Sage, Thousand Oaks
- Searle JR (1995) *The construction of social reality*. Penguin, London
- Smith PB, Bond MH, Kagitcibasi C (2006) *Understanding social psychology across cultures*. Sage, Thousand Oaks