

Artificial Agents Entering Social Networks

Nikolaos Mavridis

Introduction

Social network sites (SNSs), which have recently become tremendously popular,¹ have so far been exclusively populated by human actors. On the other hand, at least part of the functionality of such networks relies on software agents implementing artificial intelligence techniques—for example, in order to implement recommendation systems for friends or other entities. However, such agents were not playing actor roles within the network. Recently, the monopoly of human actors within SNSs has been broken; disembodied or even physically embodied intelligent software agents are just starting to populate SNSs. A huge range of potentialities exists regarding useful roles for such artificial agents, which might furthermore have varying degrees of autonomy. In this chapter, I will start by introducing a concrete example of such an agent: Sarah the FaceBot, a robotically embodied intelligent artificial agent, which carries out natural language interactions with people, physically present or remote, and which utilizes and publishes social information on Facebook—even having her own automatically updated page. Then, five areas of open questions that have arisen will be presented, as well as an exposition of the potentialities for other artificial agents in SNSs, either in actor or in other roles, which are promising to unleash new possibilities and beneficially transform social networks.

Sarah the FaceBot

Sarah (Mavridis et al., 2009a; Mavridis, Kazmi, & Toulis, 2009c) is a physical mobile robot with Face Recognition, natural language dialogue, as well as navigation/mapping capabilities (Figure 14.1). Furthermore, it is a robot that is also equipped with an interaction database (keeping information about previous interaction episodes) as well as a social database (keeping personal info as well as friendship relations for the people the robot knows). Most importantly, the

robot also has a live connection to the Facebook social network site—and utilizes online information toward crafting dialogues with humans, as well as publishing information on its own Facebook page. Sarah has three kinds of human friends: those that she has met physically, but which are not her friends on Facebook (P), those that she has met physically, and which are also friends with her on Facebook (PF), and those that she has not (yet) met physically, but which are friends with her on Facebook (F). Interaction with human friends can happen over a number of channels: real-time natural language dialogue with physically present friends, real-time dialogue over chat for friends which are currently online on Facebook, as well as offline Facebook messaging and posting.

A transcript of a typical interaction with Sarah follows:

//Greetings/Recognition:

R: Hello! My name is Sarah. Welcome to the interactive robots and media laboratory.

R: Hi! Are you Panos Toulis? (from Face Recognition).

H: Yes (Robot changes its online status to “interacting with Panos Toulis”).



Figure 14.1 Our robot, Sarah the FaceBot, while interacting with a human.

//Query of overall state:

R: Hey, Panos Toulis, are you doing well today?

H: Yes, pretty well.

//General news-based discussion:

R: Let me tell you some interesting news that I've heard.

R: The latest windows release will be available to everyone after a surge in demand crashed the Microsoft website on the 9th of January (item taken from BBC news online).

R: Have you heard about this?

HR: Yes.

R: That's great!

//Robot discusses Facebook status of human:

R: I see you are chasing robots in the desert.

//Robot gives news about a mutual friend, acquired through Facebook:

R: Did you know that our friend Shervin Emami has posted a new photo on Facebook? (information acquired from Facebook minifeed).

H: No.

//Robot sends reminder to human:

R: I am sending you a message about this to check it out.

//Robot remembers past encounter with mutual friend:

R: I saw Shervin yesterday evening and he was doing well. (information acquired from interaction database).

//Robot gives news about a mutual friend, acquired through Facebook:

R: Did you know that our friend Chandan is "waiting for Sarah to have that little sweet conversation with me"? (information acquired through Chandan's Facebook status).

H: Yes, of course.

//Robot connects an online friend with the human:

R: One of our friends, Nikolaos Mavridis, is online—would you like me to send him a message?

H: Yes ("Nikolaos, Panos says hello" is sent).

//End of interaction:

R: Hey, Panos, it was nice talking to you! I have to go now. See you later! (Robot continues wandering).

Thus, currently Sarah is capable of producing dialogue turns for:

D1. Recognition/Greeting.

D2. Querying the *state* of the human.

D3. Relaying customized *general news*.

D4. Relaying Facebook *minifeed-based news* about human or common friends.

D5. Relaying previous *interaction-based memories* about human or common friends.

- D6. Performing a *real-time connection with a third* common friend which is online.
- D7. Saying goodbye.

It is worth noting that all of these turns contribute toward real-time information diffusion within the social net; and, apart from these, Sarah's updated Facebook page contents as well as messages also diffuse information, but in a non-real-time manner.²

Sarah was originally created in order to test an interesting hypothesis in the field of HRI (Human–Computer Interaction), which was formulated in Mavridis et al. (2009a): “Can reference to *shared memories* and *shared friends* in human–robot dialogue create more meaningful and sustainable relationships?”

Motivation for positing this question was provided by disappointing early results on long-term human–robot interaction experiments, as exemplified by Mitsunaga et al. (2006)—although robots seem to be exciting and interesting to humans at first, upon multiple encounters quite quickly humans lose interest. Thus, the following chain of argument led to the postulated hypothesis:

Let us examine random human encounters, without explicit purpose of interaction—say, short chat with a colleague or friend. What is their content? First, there seems to be continuity in these dialogic episodes, connecting the current with the previous encounters; a common, *shared past* is being created, and reference to it is often made in the dialogue. Second, this common past is not exclusive to the two partners conversing at the moment; it actually extends to their circle of mutual acquaintances—and thus news and memories regarding *shared friends* are often being mentioned. Thus, let us try to create a conversational robot that can refer to shared memories and shared friends in its dialogues; and examine whether this will lead to better long-term human–robot relationships.

Upon closer examination, and in AI terminology, in a sense Sarah is a form of a chatterbot; and there exists a long line of such systems in the literature, starting with the classic ELIZA (Weizenbaum, 1966). But there are a number of important differences between FaceBots and classic chatterbots; not only is Sarah physically embodied, but most importantly her dialogues are driven by a rich context of previous interactions as well as social information, acquired physically or online, and which is dynamic and conversational-partner specific.

Two further comments are worth making: first, regarding “shared” entities; and second, regarding implicit teleology. The primary hypothesis that FaceBots were created for, is concerned with two postulated “shared” entities and their effect on human–robot relationships: shared past and shared friends.

Both of these belong to a wider set of shared entities that might prove to be important: shared interests, shared goals—actually often quite correlated with shared past and shared friends, at least in certain contexts/for certain subsets. All of these shared entities can be hypothetically unified under the “intersection” $I(A(t),B(t))$ of the two actors (human and robot in our case), at a given time instant t —a time-varying concept. It might well be that the creation, maintenance, and synergistic co-evolution³ of such an intersection turns out to be a crucial factor toward long-term human–robot relationships.

Before proceeding to five areas of open questions that have arisen from this project, a short note on teleology: the casual conversations that Sarah is attempting to replicate seem not have an explicit purpose from the conversational partner’s point of view. However, their teleology is probably better localized not at the personal or the dialogic-partners level—but at the social network level. The establishment of an adequate intersection enabling understanding and co-reference, the flow of local-context relevant information, and the resulting bonding might well be three main components—ultimately tied to collective social capital.⁴

Five Areas of Open Questions

Apart from the original motivation behind the creation of Sarah the FaceBot, this line of research opened up a number of interesting avenues as well as questions related to artificial agents and social networks:

Q1. *Interaction patterns of agent*: What will be the interaction patterns of such agents with physically present or remote humans? For example, what will be the frequency, duration, and content of such interactions?

In practice, for artificial agents within social networks, this would amount to logging and analyzing the different types of interaction events that will occur—synchronous or asynchronous, mutually visible or unidirectionally visible: viewing a profile or photo, sending a message, chatting, adding a friend, etc. For agents that also have a physical embodiment, such as Sarah the FaceBot, proxemics, gaze, and other such external measurements might also be utilized.

Q2. *Friendship graph of agent*: What will be the form and temporal dynamics of the friendship graph of such agents? (a snapshot of Sarah’s graph can be found in Figure 14.2). What will the connectivity patterns, tie strengths, as well as the individual social capital (Coleman, 1988) be?⁵

One might expect significant differences with human actors in this respect;⁶ for example, the sustainable social circle size of technologically unassisted humans is constrained by cognitive limitations—which seem to be somewhat relaxed in the case of artificial agents. On the other hand, one should also note that there also exist important limitations of the current state of agents as

Given that human actors do not live in a symbolic/language-only world, and they populate SNSs with multimedia content, it is important for artificial agents to be able to handle and/or contribute such content. On the other hand, again given the different domains and activities on which the current state of agents is more capable as compared to humans, and vice versa, this also creates an opportunity for overall benefit.

Q5. *Social engineering potential of such agents for SNSs*: How will such agents be designed/positioned in order to affect connectivity patterns, diffusion patterns, social capital, and other such important parameters at will? How will one exploit the different capabilities of artificial agents for such a purpose?⁸

From a practical point of view, this is the most important question—and we will return to some aspects of this in the last section of this chapter.

Currently, some very early answers to aspects of Q1 and Q2 for the case of Sarah have been reported in Mavridis et al. (2009c), together with an extensive discussion of the synergies between SNSs, interactive robotics, and face recognition. Furthermore, the use of live photos in conjunction with online photos toward better face recognition, as well as algorithms utilizing social context toward better and/or faster recognition through such agents, is discussed and algorithms are given in Mavridis, Kazmi, Toulis, & Ben-AbdelKader (2009b). Also, simple algorithms for empirically estimating the social graph given only photos containing co-occurring faces are presented.

Of course, this is just a very early stage regarding the questions and avenues listed above—and much more work remains to be done in order to reach a more mature stage. Also, one can pose the above questions (Q1–Q5) not only in their predictive form (“What will be?”), but also in their potential form (“What could be?”), their normative form (“What should/would one want to be?”), and their engineering form (“How should we act in order to reach . . .?”). Thus, we can, for example, ask not only: how will social capital change with the introduction of artificial agents? But also: how could it change? As well as: how would one want it to change? And also: what action plan should be followed so that the introduction of artificial agents within social networks changes social capital toward the desired direction?

The Physical vs. Online and Symbolic vs. Sensory Realms

Expanding upon Q4, another interesting observation regarding embodied artificial agents in actor roles arises: such artificial actors, as human actors do, belong to an actual social network, a subset of which is re-represented within Facebook. Also, as mentioned before, they have three categories of friends: physical only (P), physical who are on Facebook (PF), and Facebook only (F).

Their perceived identity thus depends on different primary sources for each of the three categories of friends (physical presentation vs. online); and the effect of differences and misalignments across these can thus be studied.

Yet one more observation is concerned with the relationship of the linguistic/symbolic with the sensory realms for such agents. Both realms are accessible physically as well as online; although different projections/selections of the two realms exist in the two. For example, consider photos; these belong to the sensory realm—and the robot has access to snapshots from its own camera (physically), as well as to Facebook-posted photos (online). For example, consider the friendship relationship between two individuals; say, George and Jack. This linguistic/symbolic information might be available through the online friendship graph on Facebook, or might be acquired by direct/indirect questioning, through the robot's dialogue system. On the other hand, this linguistic/symbolic piece of information is not uncorrelated to the sensory realm; as a simple statistical analysis can show (see Mavridis et al., 2009b), we expect that “The face of X appears in photos together with the face of Y” (a sensory-realm relation) is a strong predictor for “X is a friend of Y” (a linguistic/symbolic-realm relation). In essence, this is yet one more instance of symbol grounding (Harnad, 1990)—which is normally performed by human actors, and which in this case could potentially be transferred over to the artificial actors (Mavridis, 2007). Thus, a quartet of vertices arises: sensory/physical (capturing a photo through the robot's camera), linguistic-symbolic/physical (hearing that X is a friend of Y), linguistic-symbolic/online (reading that X is a friend of Y from Facebook), sensory/online (seeing a photo on Facebook), and the bidirectional connections among these vertices are to be resolved by the actors involved.

Now, having seen a brief introduction to FaceBots as an example of a robotically embodied artificial agent in an actor role within the Facebook SNS, let us move on toward a wider perspective: a basic taxonomy and an exposition of the potentialities for other artificial agents in SNSs (either as an actor or in other roles) will be presented, followed by a discussion of their possible effects toward beneficially transforming human social networks.

The Space of Potentialities for Artificial Agents

The space of potentialities for artificial agents within social networks is quite vast, and a number of basic degrees of freedom/dimensions (D) will be introduced here.

D1. One first obvious choice is concerned with the *Appearance* of the Agent to the human actors of the network; one possibility for the agent is to have an active Actor role within the SNS, with a profile, a friendship network, and

interactions—such as the case of Sarah—and either for it to be declared as an artificial entity or to posit itself as a human actor. Another is for it not to appear as a human actor, but as a distinct entity (for example, an installable Facebook application) or as part of the architecture of the SNS itself (as is the case of the friend recommendation system of Facebook). Yet another, quite interesting, possibility is for its existence to be unknown to the human actors; where the agent can be acting by effectively modulating what might appear as random events; for example, the order of presentation of items within a list, pushing forward and thus emphasizing some items in order to increase their availability in the human's mind.

D2. One other degree of freedom is concerned with the *Physicality* of the agent. One can have, for example, a physically embodied agent; a virtual character with a cartoon-like body; or a totally disembodied entity. Of course, this degree does not only cover form, but also movement and body dynamics of the agent.

D3. Yet one other interesting dimension is *Autonomy*; the artificial agent might be completely autonomous, or exhibiting adjustable autonomy through human assistance at specific times or in certain levels of abstraction. Such a configuration sometimes combines the best of both worlds (artificial and human), and enables successful application of agents to areas where their current state of the art would not allow them to be applied alone. Some recent examples of adjustable and sliding autonomy in the agents and robotics literature are Schurr, Marecki, Tambe, Lewis, and Kasinadhuni (2005) and Sellner, Heger, Hiatt, Simmons, and Singh (2006)—and analogous guiding principles can be followed in creating effective man-machine hybrid agents participating in SNSs.

D4. In the case of an agent in an actor role, another important dimension is that of the apparent perceived *Identity* of the agent; the profile information, linguistic style, dialogue system, posted pictures, friendship circles, as well as interaction behaviors of the agent, all contribute to this. As noted, the agent is performing his or her identity in two stages: the physical and the online stage. Simple software tools for crafting artificial actor identities have not yet appeared; although one would envision that with appropriate machine learning techniques, information mined from the profiles, dialogues, and the other traces of the actor's performed identity would enable the creation of congruent identities for artificial actors, parametrized by a set of simple user choices. For example, one could envision the possibility of learning simplistic mappings from regional-socio-economic background (part of profile information) to linguistic style (mined from dialogues), for a limited dialogic range, and vice versa, and thus using these mappings in order to minimize authoring time when crafting the identities of new artificial actors.

D5. Finally, and quite importantly, there is the question of the overall *Purpose* of the agent. This will be considered in more detail in the next section.

Possible Purposes for Artificial Agents

Let us start with an observation: moving on from actor-role to non-actor-role agents, one of the crucial differences is concerned with their scope of visibility; usually, an actor-role agent can only have direct access to the resources opened to him or her via the adjusted security settings of the other agents that have chosen to connect on the network. In contrast, an overt non-actor agent, for example a Facebook application, often gets wider access to all data of the actors that have installed it; and even more so, an overt- or covert-non-actor agent that is part of the SNS itself, for example the friend recommendation system of Facebook, can have omniscient access to all actors within the SNS as well as their interactions. After this comment regarding the difference in scope of visibility between actor- and non-actor-role agents, let us move back to some possible choices for the purpose of artificial agents within social networks.

The purpose of the example agent presented above, Sarah the FaceBot robot, is to create sustainable relationships with humans—which could be translated into a metric containing components related to frequency and duration of interaction over a longer period, human satisfaction, as well as number of friends, for example. Another possible purpose for actor-role agents is teaching/education, specialist assistance, as well as multiple forms of persuasion (Fogg, 2002).

Also, artificial agents in actor roles can be quite beneficial for setting up experiments in order to test scientific hypothesis related to social networks—for example, questions regarding diffusion—as they are, in a sense, limited but perfectly reliable puppets. As long as their divergence from human behavior is not detrimental for the purpose of the experiment, they can be used to create predictable responses and gather measurements within the social network: for example, when studying diffusion, agents can act as pre-programmed filters or targeted redistribution nodes; or when acquiring friendship request acceptance prediction models, agents can be set up with the desired apparent identities and initial messaging response patterns, and gather results regarding the acceptance of their requests by various actors. The interchange between human actor and artificial actor for social network research is quite parallel to human/robot interchange when bi-directionally informing Human–Robot Interaction (HRI) by Human–Human Interaction studies and vice versa, for example (Mutlu et al., 2009); and as long as the nature of the experiment can benefit from the “limited but perfectly reliable puppet” constraint.

Another possible purpose for actor-role agents is to intervene within the information flow of the network—toward a number of potential goals: re-spreading news, monitoring for possible mutations, even counter-spreading information, or creating parallel flows and adjusting existing two-step flow of communication nets and influencers (Katz & Lazarsfeld, 1955). Another possible goal is the active acquisition of information: actor-role agents could potentially activate their own connections on demand, in order to seek, ask for, and relay back missing information.

One further possible purpose is restructuring the connectivity of the network, through suitable overt or covert recommendations; this might take place towards a variety of goals, for example related to useful matchmaking of actors toward personal or professional goals, which could be beneficial to the network or a sub-network as a whole—perhaps in terms of social capital. For example, an agent might try to actively detect and manipulate structural holes. Due to the benefits of a possible wider scope of visibility and non-interactivity in this case, non-actor agents are more suitable for this purpose.

Another primary role for non-actor agents is supervising/policing the network in order to detect possible criminal or otherwise harmful/illegal activity. Currently, there exist, for example, automated- or human-assisted picture censorship services within SNSs; but there exist many more areas that could potentially benefit from the appropriate form of supervision, given of course appropriate privacy and freedom concerns.

Finally, let us close this brief exposition of some possible purposes for agents within social networks with a relevant comment: when arbitrating visibility/action scope across a number of agents, often hierarchical structures are quite beneficial, sometimes augmented with hierarchy-breaking patches. A recent example of a hierarchical multi-agent cognitive architecture is, for example, EM-1 (Singh, 2006), where the idea of higher-order agents having access to the internals of lower-order agents and acting as “mental critics” is central.⁹ One could thus envision similar hierarchies of visibility and action scope within hybrid multi-human/artificial agent systems operating on SNSs.

Conclusion

In this chapter we have discussed the entry of Artificial Agents, in embodied or disembodied forms, within human social networks. We started by introducing a concrete example of such an agent: Sarah the FaceBot, a robotically embodied intelligent artificial agent, which carries out natural language interactions with people, physically present or remote, and which utilizes and publishes social information on Facebook—and which publishes on her own automatically updated page. Then, there was a brief presentation of five areas of open

questions that have arisen, a short discussion on relevant aspects of the quartet created by the physical/online and symbolic/sensory realms, and an exposition of the potentialities and purposes for such agents; either in actor or in other roles. In conclusion, artificial agents, which are currently increasingly populating social networks, are promising to significantly change these networked publics in a beneficial manner, and unleash numerous new possibilities.

Notes

1. Before the introduction and wider spread of SNSs, the primary means of online self-presentation were homepages, which while changeable, were not dynamic (Papacharissi, 2002).
2. Currently, and mainly due to speech recognition constraints, Sarah is mainly diffusing information acquired through online news, Facebook minifeed and status, and interactions; but there is not much direct acquisition of information from the human, except from a basic state query and “did you know x” queries. This is an active direction for extensions.
3. This co-evolution often indirectly relies on input from personal evolution and interaction with other entities inside or outside the shared circle of friends; such interactions might lead to the growth of the personal non-shared component of each actor, which in turn leads to novel input for co-shaping the intersection.
4. For an interesting and somewhat complementary evolutionary view, including a theory postulating the transformation of primate grooming into gossip, see Dunbar (1996).
5. For a concise introduction to the basic social network analysis (SNA) terms used here, one could look at the opening chapters of Marlow (2005).
6. Ultimately, after a number of layers, reducing to some of the differences between atoms and bits, in sense of Negrepointe (1995), or at least to the differences between biological atoms and the current state of agents comprised by bits.
7. For example, the well-established power law distributions arising from the model of Barabasi (2002), depend on preferential attachment processes—which, for the sake of experimentation at least, artificial agents might not chose to follow—and linear growth of the net.
8. For example, the much larger interaction memory as well as social info storage of such agents, or the possibility of having distributed embodiments spanning large geographical distances, are two basic differences.
9. Such models are arguably quite reminiscent to implementations of the structures of a platonic republic, at least in some respects.

References

- Barabasi, A. L. (2002). *Linked: The new science of networks*. Cambridge, MA: Perseus Publishing.
- Coleman, J. (1988). Social capital in the creation of human capital. *American Journal of Sociology*, 94 (Issue supplement: Organizations and institutions: Sociological and economic approaches to the analysis of social structure), S95–S120.

- Dunbar, R. I. M. (1996). *Grooming, gossip, and the evolution of language*. Cambridge, MA: Harvard University Press.
- Fogg, B. J. (2002). *Persuasive technology: Using computers to change what we think and do*. San Francisco, CA: Morgan Kaufmann.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42: 335–346.
- Katz, E. & Lazarsfeld, P. F. (1955). *Personal influence: The part played by people in the flow of mass communications*. Glencoe, IL: Free Press.
- Marlow, C. A. (2005). *The structural determinants of media contagion*. PhD Thesis, Media Arts and Sciences, Massachusetts Institute of Technology, Cambridge, MA.
- Mavridis, N. (2007). *Grounded situation models for situated conversational assistants*. PhD Thesis, Media Arts and Sciences, Massachusetts Institute of Technology, Cambridge, MA.
- Mavridis, N., Datta, C., Emami, S., Tanoto, A., Ben-AbdelKader, C. & Rabie, T. F. (2009a). Facebots: Social robots utilizing and publishing social information in Facebook. *Proceedings of the IEEE Human–Robot Interaction Conference (HRI 2009)*.
- Mavridis, N., Kazmi, W., Toulis, P., & Ben-AbdelKader, C. (2009b). On the synergies between online social networking, face recognition, and interactive robotics. *Proceedings of the Computational Aspects of Social Networking Conference (CaSoN 2009)*.
- Mavridis, N., Kazmi, W., & Toulis, P. (2009c). Friends with faces: How social networks can enhance face recognition and vice versa. In A. Abraham, A. Hassanien, & V. Snasel (Eds.), *Computational social networks analysis: Trends, tools and research advances*. Berlin: Springer Verlag.
- Mitsunaga, N., Miyashita, T., Ishiguro, H., Kogure, K., & Hagita, N. (2006). Robovie-IV: A communication robot interacting with people daily in an office. *Proceedings of IEEE IROS 2006*, 5066–5072.
- Mutlu, B., Yamaoka, F., Kanda, T., Ishiguro, H., & Hagita, N. (2009). Nonverbal leakage in robots: Communication of intentions through seemingly unintentional behavior. *Proceedings of the 4th ACM/IEEE Conference on Human–Robot Interaction (HRI 2009)*.
- Negroponce, N. (1995). *Being digital*. New York, NY: Vintage Books.
- Papacharissi, Z. (2002). The self online: The utility of personal home pages. *Journal of Broadcasting & Electronic Media*, 46, 346–368.
- Putnam, R. D. (1993). The prosperous community: Social capital and public life. *American Prospect*, 13, 35–42.
- Schurr, N., Marecki, J., Tambe, M., Lewis, J. P., & Kasinadhuni, N. (2005). The future of disaster response: Humans working with multiagent teams using DEFACTO. *American Association Artificial Intelligence (AAAI) Spring Symposium on AI Technologies for Homeland Security 2005*.
- Sellner, B., Heger, F. W., Hiatt, L. M., Simmons, R., & Singh, S. (2006). Coordinated multiagent teams and sliding autonomy for large-scale assembly. *Proceedings of the IEEE*, 94, 7, 1425–1444.
- Singh, P. (2006). *EM-ONE: An architecture for reflective commonsense thinking*. PhD Thesis, Media Arts and Sciences, Massachusetts Institute of Technology, Cambridge, MA.
- Weizenbaum, J. (1966). ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9, 1, 36–45.