Affect and Interaction in Agent-Based Systems and Social Media: Guest Editors' Introduction

CHLOÉ CLAVEL, Institut-Mines-Telecom, TelecomParisTech, LTCI-CNRS, France

ROSSANA DAMIANO and VIVIANA PATTI, Dipartimento di Informatica, University of Turin, Italy

PAOLO ROSSO, PRHLT Research Center, Universitat Politècnica de València, Spain

CCS Concepts: • **Human-centered computing** → **Social media**; • **Computing methodologies** → **Intelligent agents**; **Natural language processing**;

Additional Key Words and Phrases: Affective computing, socio-affective behavior, agent interaction

ACM Reference Format:

Chloé Clavel, Rossana Damiano, Viviana Patti, and Paolo Rosso. 2017. Affect and interaction in agent-based systems and social media: Guest editors' introduction. *ACM Trans. Internet Technol.* 17, 1, Article 1 (March 2017), 8 pages.

DOI: <http://dx.doi.org/10.1145/3018980>

1. INTRODUCTION

Today's Internet is evolving toward an open society of humans and computational entities, where intelligent agent systems increasingly support the interaction between users and computational components. In this scenario, affect plays a key role, with functions that span from creating and maintaining interpersonal relations, to establishing cooperation and trust. Artificial systems—which can act as both actors and facilitators of these interactions—are more and more requested to integrate affective components in order to achieve truly realistic social behaviors and foster the creation of bonds with the users, with timely reactions to their affective input and appropriate expressions of affect. Achieving this integration requires to understand and reproduce the role of affect in human expressive capabilities, and to account for affect-related phenomena (e.g., sentiment, emotions, and mood) that engage social abilities, such as empathy, and expressive means, such as irony. The expectation for complex phenomena such as empathy and irony is that effective approaches require an interdisciplinary approach and, above all, the integration of representational models and data-oriented processing techniques. The aim of this special section is to bring together leading research on computational models of affect-related phenomena in interactions occurring either in social media or agent-based systems, by attaining cross-fertilization between two relevant perspectives: on the one hand, research on agent architectures and cognitive

The work of Viviana Patti was partially carried out at the Universitat Politècnica de València within the framework of a fellowship of the University of Turin co-funded by Fondazione CRT (World Wide Style Program 2). The work of Paolo Rosso has been partially funded by the SomEMBED TIN2015-71147-C2-1-P MINECO research project and by the Generalitat Valenciana under the grant ALMAMATER (PrometeoII/2014/030).

Authors' addresses: C. Clavel, Institut-Mines-Telecom, TelecomParisTech, LTCI-CNRS, France; email: chloe.clavel@telecom-paristech.fr; R. Damiano and V. Patti, Dipartimento di Informatica University of Turin, Corso Svizzera 185, 10149, Turin, Italy; emails: {rossana, patti}@di.unito.it; P. Rosso, Pattern Recognition and Human Language Technology, DSIC, Building 1F - Campus de Vera Universitat Politècnica de València Camino de Vera, s/n - 46022 Valencia, Spain; email: proso@dsic.upv.es.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

2017 Copyright is held by the owner/author(s).

1533-5399/2017/03-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/3018980>

models, mainly concerned with the integration of affective states into agents and open to the creation of virtual and embodied agents; on the other hand, research on techniques for sentiment analysis and opinion mining, mainly focused on the processing of affective information in social media, typically (but not only) expressed through text. The integration of models and methods between the two perspectives can open the way to the development of a new generation of social and interactive applications that leverage the affective dimension to promote improved, spontaneous technology-mediated interactions, including human–computer and human–human interactions, on small and large scale.

Our goal here is to provide an overview of the open research challenges for the community of researchers interested in analyzing and modeling the interplay between affect and interaction in agent-based systems and social media, as well as an introduction to the special section. The rest of the article is structured as follows. Section 2 discusses a set of research challenges relevant in the context of this special issue, Section 3 briefly introduces the articles included in this ACM TOIT special section. Section 4 concludes the article.

2. RESEARCH CHALLENGES

Research in sentiment and affect spans across different research communities, characterized by different concerns and domains of interest. Focusing on the *interaction* issue, it is possible to identify a set of research topics that are relevant for the scope of this special section. In the field of Intelligent Agents, they include, beyond the integration of emotion models into agent architectures, the role of emotions in the agent's decision making needed to obtain emotionally believable reactions, and the relationships of the generated emotions with the expressive behavior of agents needed to deliver affective reactions. At the interactional level, in fact, the design of an emotionally-aware system requires interoperable models of emotions between the sub-systems for the detection of others' affect (in both interpersonal and social contexts) and the expression of the system's affective state. At the social level, models of the reciprocal relations of emotions and social functions (e.g., normative behavior and cooperation) are needed to manage high-level phenomena like trust and social cohesion. On the other hand, when we focus on spontaneous interactions that can be observed on social media—where users freely express what is on their mind at any moment in time, at any location, and about any subject—relevant research topics range from the analysis of affect-related phenomena, such as sentiment, mood, emotions, or personality traits (i.e., sentiment analysis, opinion mining, emotion and personality detection), attracting an increasing interest among researchers in Computational Linguistics and Natural Language Processing, to the analysis of figurative language devices and creative means of expressions, such as irony and humor [Reyes et al. 2012], which impact the affective dimension and have an important role in interaction with humans [Binsted et al. 2006].

Having this background in mind, in the following we will highlight some research challenges emerging in the area of the special issue.

2.1. Merging Communities

Given the leading thread depicted above, one key point is understanding the best ways to fruitfully merge—in terms of methods, resources, and results—the agent community, focused on the integration of emotions into agent architectures and models, and the sentiment analysis and opinion mining community, aimed at the processing of affective information in user generated texts. On this perspective, the general question to answer is “Which are the mutual benefits and the challenges emerging from this merging?”

Tracks to answer this general question can be found in the more specific challenges discussed in the following.

2.2. Models of Affect, Emotion, and Sentiment

Emotion, opinion, sentiment, mood, affect, attitude, interpersonal stance, personality traits, judgment, appreciation, argumentation. . . The variety of communities working on computational models of affect-related phenomena triggers the use of different terminologies to define a same phenomenon or the use of a same terminology to define different phenomena [Clavel and Callejas 2016]. Opinion mining and sentiment analysis community tends to use opinion, sentiment, and affect, but existing works do not always provide an in-depth definition of the related phenomena. Agent community relies on psychological models to model phenomena rather centered on emotions and affects and, more recently, on social interactions. The two communities (agent community and sentiment analysis community) have delivered a set of valuable linguistic resources (emotion-annotated or sentiment corpora) and multimodal datasets which are consistent with some affect model [Bosco et al. 2013; Nissim and Patti 2016]. They have to be connected in order to tackle this crucial challenge of the definition of the relevant affect-related phenomena. On the one hand, the challenge for agent community is to go beyond emotions and to consider other phenomena, such as attitude or interpersonal stances [Scherer 2005]. On the other hand, the challenge for the opinion mining community is to go beyond the classical positive/negative distinction and to go deeper in the definition and the modeling of the studied phenomena [Munezero et al. 2014], according to the corpora and the applications—i.e., the detection of the affective information of interest in a given application domain.

In parallel with the study of the expressions of affect delivered by the sentiment analysis community, in the agent community, the availability of agent architectures and 3D graphics solutions has prompted the implementation of embodied emotional agents for contexts that range from storytelling to conversation (see, e.g., Hartholt et al. [2013]). In order to leverage the models delivered by the research in emotional systems, agent architectures need to integrate different models of emotions (and different expressive means) into open and highly customizable frameworks, bridging the boundaries between appraisal-based models (mostly oriented toward deliberation; see Gratch and Marsella [2004]) and the expression of affect, rooted in the evolutionary foundations of communication [Pelachaud 2009]. Integration can also be the key to the development of agents and systems that rely on the full range of affect-related phenomena—including personality, stance, and the like—to implement high-level social behaviors, such as empathy and social presence [Paiva 2011].

2.3. Detecting and Using Social Sentiment

The availability of huge amounts of user generated data from social media has been proven to be a great opportunity for advancements in sentiment analysis and opinion mining in the last few years [Cambria et al. 2013; Pozzi et al. 2016]. Thus, social media is emerging as a precious testbed for research in emotions and sentiment. One still open issue is if it also brings opportunities for the agent community interested in modeling the affective component of agent behavior. In many contexts of human-agent interactions, as, for instance, the creative ones related to expressive media [Rosso et al. 2016], which naturally include the interaction with a social audience, research efforts have mostly focused on the generative aspects, while the reception by the users/audience has only occasionally been tackled, with few exceptions [Langlet and Clavel 2016]. However, the study of reception can raise interesting issues, some related to the communicative level (reception as an immediate feedback or as asynchronous response,

like tags and social comments; see, e.g., Bertola and Patti [2016]), and some related to the audience value system (e.g., reputation and trust). Are sentiment methodologies and technologies now mature enough to allow for a tighter integration with emotions in modeling feedback in context like the ones of expressive media? Some inspiration can be taken by embodied conversational agents, where the same models (e.g., facial expressions, gaze) are shared between generation and analysis, leading to effective adaptive agents.

2.4. Moral Emotions and Ethical Issues

Last, but not least, the ethical implications of the implementation of affect in artificial systems deserve further investigations. Besides the issues raised by the recently established science of captology [Fogg 1999], in the last decade ethical questions have been posed about the moral aspects implicit in the paradigm of autonomous agents. For example, according to Guerini and Stock [2005], persuasive agents should be able to reason about their own actions from an ethical perspective (see also Anderson and Anderson [2011]). A tighter integration of models for generation and reception in affect-aware interactive and social systems may result in ethical implications: the possibility of analyzing sentiment and emotions in the behavior of users (be it text or other types of communicative behavior) opens the way to the possibility of manipulating the affective state of the users, e.g., through the accurate and reliable planning of the expressive resources. This possibility requires the community to raise its awareness of ethical issues connected with the manipulation of the affective states of the users, especially in domains where users may be more fragile, such as elder people or children. Finally, the relevance of ethical issues in emotionally-aware systems requires the development of computational models of moral emotions, which constitute a bridge between the individual and the society in the regulation of moral behaviors [Tangney et al. 2007].

3. IN THIS ISSUE

In this section, we briefly introduce each of the articles that were selected for this special issue. These contributions are a varied representation of the challenges and solutions that are emerging in this research context. Interestingly, contributions in the intelligent agents' field converged on the following topics: development of frameworks where affect-related phenomena contribute to implement socially believable virtual agents [Yuksel et al. 2016; Mancini et al. 2016], and embedding of cognitive models of emotions in BDI (belief-desire-intention) agent architectures [Espinosa et al. 2016], with a special focus on moral emotions in Dastani and Pankov [2016]. On the other hand, social media, such as Twitter and Facebook, emerged again as a very interesting testbed for in-depth analysis both of affect [Meo and Sulis 2016] and of complex social abilities and expressive means, such as empathy and irony [Otterbacher et al. 2016; Hernández Fariás et al. 2016]. Furthermore, social networks and online communities provided a very interesting context for investigating new perspectives on reputation and trust [De Meo et al. 2016; Binmad and Li 2016].

3.1. Brains or Beauty: How to Engender Trust in User-Agent Interactions

In Yuksel et al. [2016], the authors investigate factors which can engender user trust in agents by focusing, in particular, on reliability and attractiveness. They show that it is not enough for agents to be simply reliable; they need to be attractive, as well, in order to encourage user interaction. Subjective user ratings of agent trust and perceived accuracy suggest that attractiveness may be even more important than reliability.

3.2. Implementing and Evaluating a Laughing Virtual Character

The article by Mancini et al. [2016] addresses the complex relation that links empathetic behavior with the manifestation of laughter in interpersonal relations, exploring the valence of laughter for the establishment of a positive mood in the user. The authors propose a novel experimental setting for the evaluation of the laughing behavior, where the effect of the laughter is investigated under different varied conditions, which include synchronization with the user, and its effect is assessed on the user's perception of the shared experience (here, music). The experimental results confirm the relevance of the emotional aspects for the communicative behavior of virtual agents in an interpersonal context and provide operational guidelines for the implementation of laughter in embodied characters.

3.3. Other-Condemning Moral Emotions: Anger, Contempt, and Disgust

The article by Dastani and Pankov [2016] proposes a semi-formal model of the so-called CAD triad (Contempt, Anger, and Disgust) that specifies the conditions that motivate each emotion type and the goals they trigger when established. Geared to an operational account of moral emotions for BDI agents, this model accommodates into a unifying framework both the process of emotional appraisal and the generation of the coping behavior that follows from them. Being intrinsically of social nature, these emotions are relevant for the type of interpersonal and public exchanges occurring in social media, as exemplified in the article through a social media scenario.

3.4. Toward Formal Modeling of Affective Agents in a BDI Architecture

In Espinosa et al. [2016], the authors describe an extension of Jason, the BDI agent architecture, called GenIA3 (General-purpose Intelligent Affective Agent Architecture). The proposed architecture assumes a cognitive model of emotions, but is not committed to a specific theory, so different emotion models can be implemented through it, thanks to the extension of the AgentSpeak reasoning cycle. By accommodating different motivational and behavioral models into its architecture, GenIA3 provides a tool for implementing emotional agents with a fine-grained customization, suitable for different scenarios and interoperable with the applications of the Jason architecture.

3.5. Show Me You Care: Trait Empathy, Linguistic Style, and Mimicry on Facebook

In Otterbacher et al. [2016], the studied social media is Facebook and its written human-human conversations. The affect-related phenomenon is here investigated from the perspective of empathic communication with the study of participants' cognitive, emotional, and behavioral empathy. The article proposes a statistical study of the relationship between users' empathy profile (obtained by questionnaires concerning users' cognitive and emotional empathy) and the specific linguistic features that are used when engaging friends (behavioral empathy). The chosen linguistic features focus on the modeling of linguistic mimicry (the adoption of another's language patterns). The article identifies four empathy profiles and highlights their correlation with the various writing styles. The statistical analysis also shows the presence of the linguistic mimicry process between friends in the Facebook conversation.

3.6. Processing Affect in Social Media: A Comparison of Methods to Distinguish Emotions in Tweets

In Meo and Sulis [2016], the specific linguistic features contained in the data of the Twitter microblog are studied and integrated in different emotion classification systems. The psychological emotion model that is here used to define the emotion categories is the Plutchik one [Plutchik 1980]. The focus of the article is put on the text

content representation of the tweets. It compares three different methods of text content representation as an input of supervised emotion classification using classical machine learning algorithms. The article provides an in-depth analysis and discussion of the results obtained by the different processing of the linguistic content of the tweets.

3.7. Using Centrality Measures to Predict Helpfulness-based Reputation in Trust Networks

The work by De Meo et al. [2016] focuses on how to fruitfully combine information coming from different perspectives on user reputation in online social networks. In such contexts, user reputation scores are computed according to two orthogonal perspectives: helpfulness-based reputation and centrality-based reputation. According to the first perspective, the most reputable users post the most helpful reviews; according to the second one, the most reputable users occupy the most central positions in the network encoding trust links. In this study, different datasets extracted from CIAO, Epinions, and Wikipedia, and five centrality measures have been used to calculate centrality-based reputation scores. Authors show that such scores can be used to predict helpfulness-based reputation ones: because user reviews are sparse, they can leverage trust relationships to spot those users producing the most helpful reviews for the whole community.

3.8. Improving the Efficiency of an Online Marketplace by Incorporating Forgiveness Mechanism

The work by Binmad and Li [2016] investigates the impact of affect-related phenomena on reputation and trust, with a particular focus on online marketplace communities, where reputation plays an important role, improving trust among community members. In such business-related context, reputation works as a decision-making tool for interpreting the behavior of the partners so that often success of any online business depends on the trust the business agents share with each other. However, untrustworthy agents do not have any place in online marketplaces and are forced to leave the market, even if they may potentially cooperate. In this study, authors propose a new strategy based on forgiveness mechanism for untrustworthy agents to recover their reputation. A number of experiments based on NetLogo simulation are performed to validate the applicability of the proposed mechanism. Experimental results show that the online marketplace incorporating forgiveness mechanism can be used with the existing reputation systems and improve the efficiency of online marketplaces.

3.9. Irony Detection in Twitter

In Hernández Farías et al. [2016], the authors focus on the role of affect-related phenomena in ironic tweets. Irony has been proven to be pervasive in short texts created via social media, posing a challenge to sentiment analysis systems [Hernández Farías and Rosso 2016]. The affective content involved in ironic texts is multi-faceted, encompassing aspects which can be related to the emotional state of the ironist and of the recipient, or to the evaluative meaning of the ironic utterance, i.e., to the expression of a negative or (rarely) positive opinion toward a target. In this work, authors propose a novel model for irony detection, specifically focused on investigating the role of affective information in a comprehensive manner, by exploring the use of affective features based on a wide range of lexical resources available for English. Evaluation across different datasets, where samples of ironic utterances were collected from various domains by using different collection and selection criteria, allowed to test the robustness of the approach. Overall, results suggest that affective information helps in distinguishing among ironic and non-ironic tweets. A comparative evaluation of the performances of the proposed model is presented, showing that the approach outperforms the state-of-the-art, in almost all cases.

4. CONCLUSION

To conclude, the selection of presented contributions reflect the spirit of the issue, oriented toward the cross-contamination among the communities that investigate the complex manifestations of sentiment, affect, and emotions in interactive agent and social systems. The topics addressed in the issue, are well distributed over the challenges listed in Section 2 and reflect the research goals of the different communities involved in the study of affect, ranging from social networks to embodied interaction, with partially overlapping methodologies and models. We express our hope that the mapping provided by this issue can become a useful compass for research in this field and a positive stimulus to fruitful interchange for the future.

ACKNOWLEDGMENTS

We would like to thank the Editor-in-Chief, Munindar P. Singh, and all the authors for their valuable contributions. In addition, we are deeply grateful to the reviewers that helped with the decision process and contributed with excellent reviews to make this issue possible.

For editorial reasons, the article by Hernández Farías et al. [2016] has been published in Volume 16, number 3 of ACM Transactions On Internet Technology (TOIT), 2016.

REFERENCES

- Michael Anderson and Susan Leigh Anderson. 2011. *Machine Ethics*. Cambridge University Press.
- Federico Bertola and Viviana Patti. 2016. Ontology-based affective models to organize artworks in the social semantic web. *Inf. Process. Manage.* 52, 1 (2016), 139–162.
- Ruchdee Binmad and Mingchu Li. 2016. Improving the efficiency of an online marketplace by incorporating forgiveness mechanism. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–20.
- Kim Binsted, Benjamin Bergen, Seana Coulson, Anton Nijholt, Oliviero Stock, Carlo Strapparava, Graeme Ritchie, Ruli Manurung, Helen Pain, Annalu Waller, and Dave O'Mara. 2006. Computational humor. *IEEE Intell. Syst.* 21, 2 (2006), 59–69.
- Cristina Bosco, Viviana Patti, and Andrea Bolioli. 2013. Developing corpora for sentiment analysis: The case of irony and senti-TUT. *IEEE Intell. Syst.* 28, 2 (2013), 55–63.
- Erik Cambria, Björn W. Schuller, Yunqing Xia, and Catherine Havasi. 2013. New avenues in opinion mining and sentiment analysis. *IEEE Intell. Syst.* 28, 2 (2013), 15–21. <http://dx.doi.org/10.1109/MIS.2013.30>
- Chloé Clavel and Zoraida Callejas. 2016. Sentiment analysis: From opinion mining to human-agent interaction. *IEEE Transactions on Affective Computing* 7, 1 (2016), 74–93.
- Mehdi Dastani and Alexander Pankov. 2016. Other-condemning moral emotions: Anger, contempt and disgust. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–24.
- Pasquale De Meo, Katarzyna Musial-Gabrys, Domenico Rosaci, Giuseppe Sarne, and Lora Aroyo. 2016. Using centrality measures to predict helpfulness-based reputation in trust networks. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–26.
- Bexy Alfonso Espinosa, Emilio Vivancos, and Vicente Botti. 2016. Towards formal modeling of affective agents in a BDI architecture. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–23.
- Brian J. Fogg. 1999. Persuasive technologies. *Commun. ACM* 42, 5 (1999), 27–29.
- Jonathan Gratch and Stacy Marsella. 2004. A domain-independent framework for modeling emotion. *Cognit. Syst. Res.* 5, 4 (2004), 269–306.
- Marco Guerini and Oliviero Stock. 2005. Toward ethical persuasive agents. In *Proceedings of the IJCAI Workshop on Computational Models of Natural Argument*. Edinburgh, Scotland.
- Arno Hartholt, David Traum, Stacy C. Marsella, Ari Shapiro, Giota Stratou, Anton Leuski, Louis-Philippe Morency, and Jonathan Gratch. 2013. All together now. In *Proceedings of the International Workshop on Intelligent Virtual Agents*. Springer, 368–381.
- Delia Irazú Hernández Farías, Viviana Patti, and Paolo Rosso. 2016. Irony detection in Twitter: The role of affective content. *ACM Trans. Internet Technol.* 16, 3, Article 19 (2016), 24 pages.
- Delia Irazú Hernández Farías and Paolo Rosso. 2016. Irony, sarcasm, and sentiment analysis. In *Sentiment Analysis in Social Networks*, Federico Pozzi, Elisabetta Fersini, Bing Liu, and Enza Messina (Eds.). Elsevier, Chapter 7, 113–128.
- Caroline Langlet and Chloé Clavel. 2016. Grounding the detection of the user's likes and dislikes on the topic structure of human-agent interactions. *Knowl. Based Syst.* 106 (2016), 116–124.

- Maurizio Mancini, Beatrice Biancardi, Florian Pecune, Giovanna Varni, Yu Ding, Catherine Pelachaud, Gualtiero Volpe, and Antonio Camurri. 2016. Implementing and evaluating a laughing virtual character. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–22.
- Rosa Meo and Emilio Sulis. 2016. Processing affect in social media. A comparison of methods to distinguish emotions in tweets. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–25.
- Myriam D. Munezero, Calkin Suero Montero, Erkki Sutinen, and John Pajunen. 2014. Are they different? Affect, feeling, emotion, sentiment, and opinion detection in text. *IEEE Trans Affective Comput.* (2014).
- Malvina Nissim and Viviana Patti. 2016. Semantic aspects in sentiment analysis. In *Sentiment Analysis in Social Networks*, Federico Pozzi, Elisabetta Fersini, Bing Liu, and Enza Messina (Eds.). Elsevier, 31–48.
- Jahna Otterbacher, Chee Siang Ang, Marina Litvak, and David Atkins. 2016. Show me you care: Trait empathy, linguistic style and mimicry on Facebook. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–22.
- Ana Paiva. 2011. Empathy in social agents. *Int. J. Virtual Reality* 10, 1 (2011).
- Catherine Pelachaud. 2009. Modelling multimodal expression of emotion in a virtual agent. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, 1535 (2009), 3539–3548.
- Robert Plutchik. 1980. A general psychoevolutionary theory of emotion. In *Emotion: Theory, Research, and Experience: Vol. 1. Theories of Emotion*, R. Plutchik and H. Kellerman (Eds.). Academic Press, New York, 3–33.
- Federico Alberto Pozzi, Elisabetta Fersini, Enza Messina, and Bing Liu. 2016. *Sentiment Analysis in Social Networks, 1st Edition*. Morgan Kaufmann, Elsevier.
- Antonio Reyes, Paolo Rosso, and Davide Buscaldi. 2012. From humor recognition to irony detection: The figurative language of social media. *Data Knowl. Eng.* 74 (2012), 1–12.
- Paolo Rosso, Cristina Bosco, Rossana Damiano, Viviana Patti, and Erik Cambria. 2016. Emotion and sentiment in social and expressive media: Introduction to the special issue. *Inf. Process. Manage.* 52, 1 (2016), 1–4.
- Klaus R. Scherer. 2005. What are emotions? And how can they be measured? *Soc. Sci. Inf.* 44, 4 (2005), 695–729.
- June Price Tangney, Jeff Stuewig, and Debra J. Mashek. 2007. Moral emotions and moral behavior. *Ann. Rev. Psychol.* 58 (2007), 345.
- Beste Filiz Yuksel, Penny Collisson, and Mary Czerwinski. 2016. Brains or beauty: How to engender trust in user-agent interactions. *ACM Trans. Internet Technol.* 17, 1 (2016), 1–20.

Received October 2016; revised November 2016; accepted November 2016