Putting a Face on Algorithms: Personas for Modeling Artificial Intelligence

Amela Karahasanović, Asbjørn Følstad and Patrick Schittekat

SINTEF Digital, Forskningsveien 1, 0373 Oslo, Norway

Amela@sintef.no, Asbjorn.Folstad@sintef.no, Patrick.Schittekat@sintef.no

Abstract. We propose a new type of personas, artificial intelligence (AI) personas, as a tool for designing systems consisting of both human and AI agents. Personas are commonly used in design practices for modelling users. We argue that the personification of AI agents can help multidisciplinary teams in understanding and designing systems that include AI agents. We propose a process for creating AI personas and the properties they should include, and report on our first experience using them. The case we selected for our exploration of AI personas was the design of a highly automated decision support tool for air traffic control. Our first results indicate that AI personas helped designers to empathise with algorithms and enabled better communication within a team of designers and AI and domain experts. We call for a research agenda on AI personas and discussions on potential benefits and pitfalls of this approach.

Keywords: Personas, Interaction with AI, AI Agents.

1 Introduction

Human-automation systems increasingly includes and relies on artificial intelligence (AI) agents that range from robots, autonomous vehicles and chatbots to decision-support tools. AI agents are not just tools executing well-defined tasks but 'team players in joint human-agent activity' [1, 2]. They are designed to improve our lives and reduce our workload. However, the complexity that comes with them might result in humans having a poor understanding of how the automation is performed and of the real-world situation that the automation is helping to control; this is called human-out-of-the loop syndrome [3]. Further, AI agents often employ black-box algorithms that can affect our behaviour without our knowledge or explicit consent [4].

The design of systems containing AI is receiving increasing attention from the Human Computer Interaction community [5-8]. AI has been recognized as a new design material [8] that requires appropriate design methods and tools. Human and AI agents both show forms of agency and behaviour in the sense that they are capable of making a difference [9]. Although design, as a multidisciplinary activity that aims to clarify the purpose and meaning of the world around us [10], should model the behaviour of human and AI agents, this is not an easy task. The design challenges of human-AI interaction include challenges in understanding AI capabilities and in collaborating with AI engineers throughout the design process [11]. Designers should also model the partnership of the human and the artificial [6] and be able to empathise with algorithms to better understand their nature [12]. We follow this line of thinking and propose the use of AI personas in the design of human-automation systems.

1.1 Personas

Personas were first proposed by Cooper [13] and are commonly used in design and development practices. A persona is an archetypical character representing users' behaviour, goals, needs and frustrations, and it should help designers and developers empathise with users. Although the usefulness of the persona approach has been debated, it is considered a standard part of the designer toolbox [14]. Personas vary in the extent to which they rely on user data [14]. As Norman explained, personas are communication aids and tools that add an empathetic focus to design [15]. They can be based on qualitative and quantitative data user data, such as interviews, surveys and clickstreams [16, 17], on designers' experience and intuition, as with Norman's ad-hoc personas [15]; on fiction [18]; or on extreme characters [19]. However, one should be aware of potential biases that might be introduced by automatically generated personas [20] as well as the effects of cognitive styles [21] and experience [20] on the perception of personas. A comprehensive survey of different approaches to personas can be found in [22].

A commonality of all of these approaches is that they describe humans. As Turner and Turner [23] put it, they 'encapsulate users as rounded human beings' and 'put a face on the users'. However, there may be substantial potential in applying a persona approach to the design of AI agents. In particular, as such agents increasingly are partners in team work whom which humans relate to in ways resembling how they relate to their human colleagues or partners. There are emerging initiatives for this within the chatbot and conversational agent research communities, where chatbots are made into beings with a personality often at an early stage of the design process e.g., [24] and where researchers are already addressing the effect of such personalities on user experience e.g., [25]. We seek to contribute this thread of research by exploring the potential of applying personas approach to the design and development of AI agents. More specifically we aim to explore:

- How can personas capture the properties of AI agents?
- How can personas support collaboration in interdisciplinary teams?

2 Method

This paper explores AI personas as a means for modelling the behaviour of AI agents and supporting collaboration in multidisciplinary teams. A participatory action research approach [26] was applied. The design team consisted of a small group of researchers working in mathematical optimization, artificial intelligence and interaction design. The participants had between two and seven years of experience in development of Air Traffic Management (ATM) systems and evaluating them with controllers. One of the participants had attended a course in air traffic control dedicated to developers and researchers. All discussions during workshops and produced artefacts were recorded and then analysed using thematic analysis.

The first step in developing these artefacts was to search the literature for relevant concepts and practices. We found two recurring concepts on the properties of AI robots and algorithms: personality and attraction. The second step was to identify a case. We selected the design of a highly automated decision support tool for air traffic control. The third step was developing the AI personas. To do this, we conducted two workshops; the first one was to discuss the format of the AI personas, and the second one was to develop the personas for our case.

In their survey of socially interactive robots, i.e. robots for which social interaction plays a key role, Fong, Nourbakhsh and Dautenhahn [27] identified the following human social characteristics: expressing and/or perceiving emotions, communicating with high-level dialogue, learning/organising models of other agents, establishing/maintaining social relationships, using natural cues (gaze, gestures), exhibiting a distinctive personality and character and learning/developing social competencies. Although designing engaging personalities for chatbots based on their backstories has been proposed, it has been pointed out that building trust requires that chatbots be upfront about their machine status [28].

Research on the personalities of automated agents includes the work of Mennicken et al. [29] on the perceived personality traits of smart homes. Moreover, Ahrndt, Aria, Fähndrich and Albayrak [30] discussed how automated agents could be bestowed with personality traits to improve predictability in interactions with humans during planning tasks. Culley and Madhavan [31], however, raised the concern that anthropomorphic agents may strengthen human trust due to increased emotional appeal but decrease user sensitivity to actual performance because users would make trust-based judgements.

In personal interactions with computers, humans not only recognise their personalities but also apply the personality-based social rules of similarity attraction and complementary attraction [32]. An exploratory study investigating links between human and attributed robot personalities found that participants' evaluations of their own personality traits correlated with their evaluations of robots' personality traits [33]. Lee, Peng, Jin and Yan [34] found that participants could recognise a robot's personality based on its behaviour and that they preferred interacting with robots that had a personality that was complementary to their own. Research by de Visser et al. [35] suggested that the degree of anthropomorphism in machine agents may affect how users experience them and the level of trust they bestow upon them.

In the present study, we adapted the process for constructing personas described by Cooper, Reimann and Cronin [36] to develop AI personas as follows. First, we identified a preliminary list of behavioural variables. We extended the lists of persona-related variables suggested in the literature, such as the one given by Cooper et al. [36], to include additional properties that we found in other studies. The literature on interaction with robots/social robots was the closest to our needs. Second, we selected relevant behavioural variables (properties) after discussing the preliminary list. During this discussion, we also identified potential redundancies and checked for completeness. Third, we expanded the descriptions of the variables and behaviours. Finally, we designated the persona types.

The first step was done by the first author of this paper, whereas the other three steps were done by a small team of experts (three designers and one mathematician) during two half-day workshops. Knowledge regarding algorithms and their agency in this process was revealed during the discussion among team members. All team members were domain experts, meaning that they had a good understanding of Air Traffic Management, the job of air-traffic controllers, and the tasks, procedures and tools they currently use. The mathematician in the team had expert knowledge on existing and future optimisation algorithms, including their potential and limitations, and was acting as an “advocate” for the algorithms.

3 AI Persona Prototype

An AI agent is anything artificial that is capable of acting based on information it perceives and its own experiences. These agents may have different levels of freedom to choose between different actions, i.e. different levels of autonomy. An AI agent can have a body, such as a vacuum cleaner, or be a program installed on a computer, such as a web browser search engine. AI agents interact with their environment, including other human and AI agents. The core of an AI agent is the algorithm or algorithms that enable it to act. An AI agent can perceive the environment through sensors, such as cameras or GPS signals, and a keyboard, and it can act via actuators, such as robot arms, speakers or screens.

The case we selected for our study of AI personas was the design of a highly automated decision support tool for air traffic control. Air traffic controllers’ tasks include directing aircrafts on the ground and through controlled airspace, organising and expediting the flow of traffic and preventing collisions. The agents in this system are called planning agents [37]. They consider anticipated future situations caused by their own actions to decide the best course of action.

In the ATM domain, decisions are time- and safety-critical. Further, a decision made by one controller (human or AI agent) impacts the performance of the entire system. Although our findings have shown that AI agents that support the generation, selection and implementation of decision alternatives in this domain can improve performance by 20% to 50% [38], introducing them in the working environment of controllers is a non-trivial task. AI agents have to adapt their behaviour to humans’ needs, preferences and current situations. Furthermore, understanding the rationale behind decision alternatives generated by the AI agent and ensuring that decision implementation is manageable for humans is a prerequisite of trusting and accepting the tool [39, 40]. The algorithms behind such planning AI agents are not always based on human decision-makers’ reasoning but on mathematically proven properties and procedures which are difficult to explain to non-mathematicians. So, how should we describe such AI agents in a way that alleviate design process and leads to the development of AI agents that air traffic controllers would like to work with?

Cooper et al. [36] proposed focusing on the following behavioural variables: activities, attitudes, aptitudes, motivations and skills. They also suggested that variables related to job roles should be listed separately for applications in a work context. As AI agents interact with humans, the human social characteristics identified by Fong et al. [27] such as expressing/perceiving emotions and communicating with high-level dialogue, are also relevant.

The properties of AI personas that we discussed during the first workshop are shown in the table below. The left column lists the properties, and the right one gives a brief description. During the first workshop, the team discussed whether the variables were relevant and their meaning in our case, and they tried to understand the controller and AI agent points of view. Descriptions of properties and behaviours were expanded upon with examples. The quotes given in the rest of the paper are from the audio-recorded workshop sessions. During the process, we also identified some redundancies and checked if anything was missing.

The appearance of the AI agent was the least-discussed property. We all agreed that the AI agent would somehow be embodied in the controller's working position, that it should be discrete and non-intrusive and that it should support a combination of interaction forms, including speech, touch and visual presentation.

Table 1. Properties of AI Personas

Property	Description
Appearance	The agent's look/size/layout/interaction forms, such as a pop-up window at the controller's working station or a small robot siting on the desk. Specification of how it interacts with the controller, such as through speech or visual communication.
Type of communication	The high-level or simple dialogue used with the controller, such as brief instructions or longer sentences. Whether the agent will take over communication with the pilot to reduce the workload of the controller.
Social relationship and trust	Ability to establish a social relationship with the controller and build trust over time.
Controller's state	Recognising the state of the controller
Personality	Individual personal characteristics of the agent, such as extraversion (outgoing/reserved) and agreeableness (friendly/challenging) as well as the interpersonal behaviour of the agent, such as dominance/submissiveness and affiliation (warmth/hostility).

Social competence	Ability of the agent to learn and develop social competence
Algorithm's job-related properties /limitations	Ability of the agent to perform its job
Adaptiveness	Ability of the agent to receive input and adapt
Transparency	Ability of the agent to explain its actions, including the reasons for and impact of those actions
Role	The agent's role such as being supervisor or a colleague.

During discussion regarding the type of communication, a question was raised on whether we should accept that the communication method would be as it is today—simple instructive dialogue—or whether the AI agent would reduce the workload and thus make space for more natural, complex dialogues.

The development of a social relationship and trust were considered things that would be built over time. The most important element for trust would be showing users that previous decisions proposed by the AI agent were good. An AI agent could, for example, show how many of its proposed decisions were accepted by controllers or how its proposals led to a certain level of reduction in CO2 emissions. Small talk could be used to increase the closeness between the agent and the controller. For example, after briefly discussing weather, it might be easier for the agent to be bossy and tell the controller to follow his recommendations.

We also discussed the possibility of the agent using other senses to recognise the state of the controller. For example, the agent could recognise that the controller is stressed, sick or unable to perform the job based on gaze, gestures, eye tracking or other unobtrusive physiological measurements.

The following personal characteristics were identified as the most relevant ones for our case: extraversion, agreeableness, dominance/submissiveness and affiliation. The discussion of these characteristics brought up several interesting questions. For instance, would human personality types be enough to describe AI agents? Would we need additional characteristics? Would agents with different personalities propose different solutions? Would a risk-taking AI agent be more likely to violate separation rules when sequencing planes on the runway if it is sure that it is still safe? Would an impatient AI agent would try to send out all of the planes from its sector as soon as possible. It was also suggested that the AI agent should adapt to the personality of the controller. When we discussed the ability to learn/develop social competence, we agreed that the AI agent should propose the solution that is most appropriate to the controller's personality. For example, the agent should not push a riskier solution on a controller who likes to be on the safe side or talk too much to an introverted controller. The agent should also recognise and adapt to the controller's patterns of behaviour. For example, if a controller is under more stress than usual, the AI agent should be able to recognise this and adapt not only its communication style but also the decisions it suggests.

We agreed that the properties and limitations of the AI agent’s algorithms were very important in our case. For example, we needed to determine how quickly the agent can provide a solution, whether it provides a broader view of the airport than the controller’s view and the balance between the quality of the proposed solution and responsiveness.

In a way, agent transparency was the most difficult property to discuss. Whereas everybody agreed that the algorithms should be transparent, the mathematician judged the designers’ expectations related to transparency as difficult or impossible to implement. He explained that the optimisation algorithms explore an extremely large number of possible solutions until the optimal one is found and that this cannot be presented to the controller in a meaningful way. The alternative that was discussed was presenting the impact of the solutions proposed by the agent and the controller based on certain key performance indicators, such as the time an airplane spends on the runway (taxi time) or the levels of CO_2 emission. The agent can also try to briefly explain the reasons for the proposed solutions.

Regarding the possible roles for the AI agent, having the AI agent act as a coach, supervisor or colleague was mentioned. All proposed properties were found to be relevant. Although some of them were closely related, such as type of communication and social skills, including similar properties led to a richer discussion among the team members.

The next step in our process was designating the persona types. We first presented the results of the first workshop. Then, each team member worked on his or her own. After that, the team members presented their personas in plenum, and the others commented on them and suggested additions. Figure 1 depicts an example of an AI persona we developed.

Self-aware ATM robot	
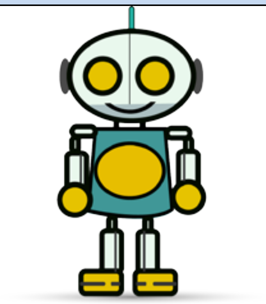	This agent is not human and has no gender. Its voice is male 50% of the time and female 50% of the time. It communicates via speech and text. If everything is normal, it always starts with the same message, perhaps with some variations in tone. If there is some form of negotiation, it will use high-level dialogue, including commands/instructions, and it should support small talk. More high-level talk will be used in the feedback session. The robot explains its decisions. It is not arrogant or patronising; its tone engenders trust. It accepts criticism. It is not passive aggressive. It responds to insults, but it is friendly. It shows some empathy to controllers, but it is not servile or too humble. It communicates sadness if the controller does not follow its advice. It does not use humour very often as it is difficult to echo the culture of the country and of the workplace. If the controller is stressed and does not ask for help, the agent is quiet and says nothing.

Fig. 1. Self-aware ATM robot

4 Experience with Collaboration

During both workshops, the level of interaction and engagement among the team members was very high. They were constantly finishing each other's sentences, interrupting each other and giving their own examples. Independent of their background, all team members empathized with the (human) end user and the AI agent.
They often drew on their own experiences with AI agents when explaining how the agent should not behave. A story from one team member was often followed by similar stories from other team members. For instance, one member noted, 'If I often travel to one place for business meetings, it [the algorithm thinks that it] means that I want to travel there all the time and offers me amazing holidays there'. Another member responded by recalling, 'The last time I was in the cottage [hundreds of kilometres from work], Google explained to me how long it would take to bike to the office'.

When discussing communication with the agent, the mathematician tried to consider the situation from the perspective of different potential end users: 'I have no idea, maybe some people want it different. Perhaps they want less information when they are stressed. What you definitely don't want when you are under stress is contradictory information'.

When we started discussing the properties of the algorithm, the mathematician described his own algorithm and very quickly identified himself with the algorithm, referring to it using 'we' and talking about the algorithm as if it were a human being. He definitely 'put a face on the algorithm', as seen below:

> If we [the algorithms] are keeping track of what a controller can do, we can update the objectives of the model, and you would not suggest some solutions if you know that the controller has only two minutes or something and cannot do it anyway.

He also spoke for the algorithm when talking about the previous generation of decision support algorithms that definitely had no natural language interface: 'Before it was like, "Hi, just give me the input", and the algorithm thinks about it, and gives the output. It expects that it will be followed'.

When considering the controller's situation, the HCI experts often referred to the algorithm with the pronoun 'you': 'If it is a black box, I don't care how you do it, but I see that you [the algorithm] remember that I usually travel'.

5 Discussion and Future Work

This paper introduces a novel method for the design of human-automation systems—the use of AI personas—and reports our first experience with this endeavour. To fully exploit the increasing capabilities of AI, humans and AI agents should learn how to collaborate more effectively. We started our journey towards designing improved human-automation systems by investigating whether and how the personification of AI agents may help multidisciplinary teams. Our initial assumption was that AI personas would facilitate empathy with algorithms, which would, in turn, advance the understanding of human-automation systems. The process that we proposed for the construction of AI personas involves the following steps: (i) identify preliminary behavioural variables, (ii) select relevant behavioural variables, (iii) expand the descriptions of variables and behaviours and (iv) designate types of AI personas.

Our process differs from the one put forward by Cooper et al. [36] in that it is not based on user data. It does, however, require the participation of experts with knowledge of the algorithms that are used or may be used for the development of AI agents. These experts can explain an algorithm's point of view, and experts with domain knowledge can explain the points of view held by different stakeholders. We found that our proposed process paves the way for empathising with algorithms and examining partnerships between humans and AI, moving design practice in the direction envisaged by several scholars [5, 6, 12]. Additionally, the process advances empathising with (human) end users beyond the level achieved with standard personas that model user behaviour. Stepping back and appraising what is happening from an algorithm's point of view requires consideration of the finer details of human-automation collaboration and communication and the enrichment of empathy with end users. Finally, the process recommended in our work helps achieve a holistic view of human-automation systems.

Björndal, Rissanen and Murphy [41] supplemented the process of constructing personas through steps that are useful in the context of industrial robots: globalisation, validation by end users, prioritisation of personas, creation of a common vocabulary, identification of critical business scenarios and identification of critical safety situations. Although we did not include these steps in our process, some were incorporated into our discussions. Specifically, we discussed cultural differences amongst various countries and ATM centres, e.g. hierarchical versus flat organisational structures; formal versus less formal organisational cultures; the need for engagement in different behaviours in critical safety situations, e.g. no negotiation during critical situations; and the importance of a common vocabulary.

Another issue that emerged during the process was the need to clarify whether AI personas should consist of one agent or a mashup of agents, as discussed in Chang, Lim and Stolterman [42]. The advantage of having several personas in our context is that the users, i.e. air traffic controllers, can select a personality in accordance with their preferences and on the basis of similarity or complementary attraction, as described above. The argument for having combined personas in our context is that the controllers would experience the decision-support system as a holistic system that would modify its behaviour based on the situation. In critical situations, for instance, the system would reduce communication to avoid additional stress or assume control over certain functions.

Our other assumption was that AI personas and the associated process would improve communication within a multidisciplinary team. This is exactly what happened; working on AI personas helped bridge gaps amongst team members with different backgrounds and expertise. The workshop participants quickly adapted their highly specialised language to one another and created common ground for understanding the human-automation systems that they were designing.

Our approach goes beyond simply understanding users. It allows us to ‘put a face on’ algorithms. We believe that this is increasingly important as we shift from AI agents that mimic human decision-making to new algorithms that are incomprehensible to non-experts.

This paper presents ongoing research. It points to future research directions: i) investigating potential benefits and limitations of AI personas throughout the entire design process, including evaluating developed solutions, and ii) investigating designers’ experiences with AI personas more comprehensively as well as effects of personas on collaboration and creativity within multidisciplinary teams. We call for a research agenda focused on AI personas. What are the possible pitfalls of using AI personas as an interaction design method? How should this method be adapted to AI as a new design material? How can it capture AI's distinguished nature?

Acknowledgements. The project was funded by the NextGenDST project.

References

1. Klien, G.,Woods, D.D., Bradshaw, J.M., Hoffman, R.R., Feltovich, P.J.: Ten challenges for making automation a "team player" in joint human-agent activity. IEEE Intelligent Systems 19, 91 - 95 (2004)
2. Grudin, J.: From Tool to Partner: The Evolution of Human-Computer Interaction. Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18), vol. Paper C15, pp. 1-3. Association for Computing Machinery (2018)
3. Endsley, M.R., Jones, D.G.: Designing for Situation Awareness: An Approach to User-Centered Design, Second Edition. CRC Press (2011)
4. Bond, R.M., Fariss, C.J., Jones, J.J., Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, James H. Fowler.: A 61-million-person experiment in social influence and political mobilization. Nature 489, 295–298 (2012)
5. Shneiderman, B., Plaisant, C., Cohen, M., Jacobs, S., Elmqvist, N., Diakopoulos, N.: Grand challenges for HCI researchers. interactions 23, 24-25 (2016)
6. Farooq, U., Grudin, J.: Human-Computer Integration. interactions 32, 26-32 (2016)
7. Abdul, A., Vermeulen, J., Wang, D., Lim, B.Y., Kankanhalli, M.: Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, pp. Paper 582. Association for Computing Machinery, Montreal QC, Canada (2018)
8. Holmquist, L.E.: Intelligence on tap: artificial intelligence as a new design material. interactions 24, 28–33 (2017)
9. Rose, J., Jones, M.: The Double Dance of Agency: A Socio-Theoretic Account of How Machines and Humans Interact. Systems, Signs & Actions, An International Journal on Communication, Information Technology and Work 1 (2005), 19-37 (2005)
10. Giacomin, J.: What is Human Centred Design? The Design Journal 17, 606-623 (2014)
11. Yang, Q., Steinfeld, A., Rosé, C., Zimmerman, J.: Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1-13. Association for Computing Machinery (2020)

12. Gajendar, U.: Empathizing with the smart and invisible—algorithms! interactions 23, 24-25 (2016)
13. Cooper, A.: The Inmates Are Running the Asylum: Why High Tech Products Drive Us Crazy and How to Restore the Sanity. Sams Publishing, USA (2004)
14. Bødker, S., Klokmose, C.N.: Preparing students for (inter)-action with activity theory. International Journal of Design 6, 99-112 (2012)
15. Pruitt, J., Adlin, T.: The Persona Lifecycle. Morgan Kaufmann, San Francisco, CA (2006)
16. Zhang, X., Brown, H-B., Shankar, A.: Data-driven Personas: Constructing Archetypal Users with Clickstreams and User Telemetry. Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, pp. 5350-5359. ACM, Santa Clara, California, USA (2016)
17. McGinn, J., Kotamraju, N.: Data-driven persona development. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08), pp. 1521-1524 Association for Computing Machinery, (2008)
18. Blythe, M., Wright, P.C.: Pastiche scenarios: Fiction as a resource for user centred design. Interacting with Computers 18, 1139–1164 (2006)
19. Djajadiningrat, J.P., Gaver, W.W., Fres, J.W.: Interaction relabelling and extreme characters: methods for exploring aesthetic interactions. Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques, pp. 66-71. ACM, New York City, New York, USA (2000)
20. Salminen, J., Jung, S-G., Jansen, B.J.: Detecting Demographic Bias in Automatically Generated Personas. In: In Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19). pp. 1-6. Association for Computing Machinery, New York, NY, USA (2019)
21. Pröbster, M.M., Haque, M.E., Hermann, J.: Cognitive styles and personas: designing for users who are different from me. In: Proceedings of the 29th Australian Conference on Computer-Human Interaction (OZCHI '17), pp. 452–456. Association for Computing Machinery, New York, NY, USA (2017)
22. Mardsen, N., Haag, M.: Stereotypes and Politics: Reflections on Personas. In: CHI 2016, pp. 4017-4031. ACM, (2016)
23. Turner, P., Turner, S.: Is stereotyping inevitable when designing personas. Design Studies 32, 30-44 (2011)
24. Shevat, A.: Designing bots: Creating conversational experiences. O'Reilly Media, Inc (2017)
25. Go, E., Sundar, S.S.: Humanizing chatbots: The effects of visual, identity and conversational cues on humanness perceptions. Computers in Human Behavior 97, 304-316 (2019)
26. McIntyre, A.: Participatory Action Research SAGE Publications. Kindle Edition. (2008)
27. Fong, T., Nourbakhsh, I., Dautenhahn, K.: A survey of socially interactive robots. Robotics and Autonomous Systems 42, 143–166 (2003)
28. Mone, G.: The Edge of the Uncanny. Communications of the ACM 59, 17-19 (2016)
29. Mennicken, S., Zihler, O., Juldaschewa, F., Molnar, V., Aggeler, D., Huang, E.M.: It's like living with a friendly stranger": perceptions of personality traits in a smart home. 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16), pp. 120-131. Association for Computing Machinery (2016)
30. Ahrndt, S., Aria, A., Fähndrich, J., Albayrak, S.: Ants in the OCEAN: Modulating agents with personality for planning with humans In: Bulling, N. (ed.) Multi-Agent Systems. EUMAS 2014. Lecture Notes in Computer Science, vol 8953., pp. 3-18. Springer, Cham (2015)

31. Culley, K.E., Madhavan, P.: A note of caution regarding anthropomorphism in HCI agents. Comp. Hum. Behav. 29, 577-579 (2013)
32. Isbister, K., Nass, C.: Consistency of personality in interactive characters. Int. J. Hum.-Comput. Stud. 53, 251-267 (2000)
33. Woods, W., Dautenhahn, K., Kaouri, C., te Boekhorst, R., Koay, K.L., Walters, M.L.: Are robots like people?: Relationships between participant and robot personality traits in human–robot interaction studies. Interaction Studies 8, 281-305 (2007)
34. Lee, K.M., Peng, W., Jin, S.-A., Yan, C.: Can Robots Manifest Personality?: An Empirical Test of Personality Recognition, Social Responses, and Social Presence in Human–Robot Interaction. Journal of Communication 56, 754-772 (2006)
35. de Visser, E.J., Monfort, S.S., McKendrick, R., Smith, M.A., McKnight, P.E., Krueger, F., Parasuraman, R.: Almost human: Anthropomorphism increases trust resilience in cognitive agents. Journal of Experimental Psychology: Applied, 22(3), 331. 22, 331-349 (2016)
36. Cooper, A., Reimann, R., Cronin, D.: About Face 3: The Essentials of Interaction Design. Wiley Publiching, Inc., Indianopolis, USA (2007)
37. de Weerdt, M., Clement, B.: Introduction to planning in multiagent systems. Multiagent Grid Syst. 5, 345-355 (2009)
38. Kjenstad, D., Mannino, C., Schittekat, P., Smedsrud, M.: Integrated surface and departure management at airports by optimization. In: Modeling, Simulation and Applied Optimization (ICMSAO), 2013 5th International Conference, Hammamet, 2013, pp. 1-5: IEEE Xplore Digital Library (2013)
39. Karahasanovic, A., Nordlander, T., Schittekat, P.: Optimization-Based Training in ATM. In: Schmorrow. D.D., Fidopiastis, C.M. (eds.) Foundations of Augmented Cognition, 9th International Conference, AC 2015, Held as Part of HCI International 2015, Los Angeles, CA, USA, August 2-7, 2015, Springer. ISBN 978-3-319-20815-2. (2015)
40. Nordlander, T., Karahasanovic, A., Schittekat, P.: Increasing trust in optimization based ATM systems through training. Lecture Notes in Management Science. ISSN 2008-0050. 7, s 41- 44 (2015)
41. Björndal, P., Rissanen, M.J., Murphy, S.: Lessons Learned from Using Personas and Scenarios for Requirements Specification of Next-Generation Industrial Robots. In: Marcus, A. (ed.) Design, User Experience, and Usability. Theory, Methods, Tools and Practice: First International Conference, DUXU 2011, Held as Part of HCI International 2011, Orlando, FL, USA, July 9-14, 2011, Proceedings, Part I, pp. 378-387. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
42. Chang, Y., Lim, Y., Stolterman, E.: Personas: From Theory to Practices. Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges (NordiCHI '08). Association for Computing Machinery, Lund, Sweden (2008)