

AI-based Digital Assistants: Opportunities, Threats, and Research Perspectives

Alexander Maedche, Karlsruhe Institute of Technology, Germany

Christine Legner, University of Lausanne, Switzerland

Alexander Benlian, Technische Universität Darmstadt, Germany

Benedikt Berger, Ludwig-Maximilians-Universität München, Germany

Henner Gimpel, University of Augsburg and Fraunhofer FIT, Germany

Thomas Hess, Ludwig-Maximilians-Universität München, Germany

Oliver Hinz, Goethe University Frankfurt, Germany

Stefan Morana, Karlsruhe Institute of Technology, Germany

Matthias Söllner, University of Kassel, Germany and University of St. Gallen, Switzerland

The final version of this paper will be published as:

Maedche, A., Legner, C., Benlian, A. et al. (2019).

AI-based Digital Assistants: Opportunities, Threats, and Research Perspectives.

Business and Information Systems Engineering, Issue 4, 2019

<https://doi.org/10.1007/s12599-019-00600-8>

AI-based Digital Assistants: Opportunities, Threats, and Research Perspectives

1. Introduction

Artificial intelligence (AI) is becoming omnipresent; it permeates our work and private lives in many areas. A key area of application is AI-based digital assistants, which are now becoming available in large numbers and a wide variety of usage scenarios. Research into AI-based digital assistants has a long history, dating back to Joseph Weizenbaum's well-known ELIZA in 1966. In parallel, global technology companies such as Microsoft, IBM, Google, and Amazon have been working intensively for decades on advancing AI-based digital assistants and have recently made them suitable for the mass market. Empowered by recent advances in AI, these assistants are becoming part of our daily lives. We are observing the ever-growing usage of various digital assistants, for instance, voice-based assistants such as Amazon Alexa, or text-based assistants (chatbots), such as those embedded in Facebook Messenger. It is foreseen that AI-based digital assistants will become a key element in the future of work. Today's enterprise communication platforms such as Slack or Microsoft Teams already provide many different bot types to augment work, and Gartner (2019) predicts that by 2021, one-quarter of all digital workers will use a virtual employee assistant daily.

AI-based digital assistants provide significant opportunities, but also might become a threat. On the one hand, they are expected to take over routine tasks from humans and to free up time and resources for more demanding tasks. For instance, IBM (2017) argues that chatbots can help to reduce customer service costs by 30%. On the other hand, a recently announced advanced AI-based digital assistant by Google named Duplex (Google AI Blog, 2018) has led to a debate about potential misuses for deception and fraud, owing to its human likeness. More generally, while the pervasiveness of AI-based digital assistants increases, most people ignore

their underlying architecture and algorithms (Frey & Osborne, 2017), resulting in serious concerns and user aversion regarding their uses (Dietvorst, Simmons, & Massey, 2015, 2018).

From a conceptual perspective, AI-based digital assistants – like every IS – can be understood from two different yet complementary perspectives (see Figure 1): First and broadly speaking, AI-based digital assistants represent a socio-technical system that relies on the interplays of three key elements (Goodhue & Thompson, 1995; Heinrich, Heinzl, & Riedl, 2011): the individual *user*, who seeks to achieve certain goals; the *tasks* the user needs to accomplish so as to achieve their goals, and the *technology*, as the computer system (i.e. software, hardware, and data) an individual may interact with to carry out tasks. Second, the AI-based digital assistant is an applications class (i.e. a combination of software components and data structures) and can be characterized by its input, output, and processing. Digital assistants generally have a specific extent of interactivity and intelligence in order to help users to perform tasks (Maedche, Morana, Schacht, Werth, & Krumeich, 2016). By using AI technologies, such as natural language processing, machine translation, speech recognition, machine learning, or knowledge representation (Russel & Norvig, 2010), AI-based digital assistants augment human task performance with higher extents of interactivity and intelligence than previous generations of digital assistants or traditional software applications. Most contemporary AI-based digital assistants rely on some form of conversational user interface, such as speech-based or text-based conversational agents, both for receiving input from and delivering output to users using natural language processing. Advanced AI-based digital assistants may also apply computer vision to recognize visual inputs. Further, AI-based digital assistants have the capability to represent and process domain knowledge as well as to learn and generate new knowledge from collected data by applying machine learning algorithms.




Figure 1. Two Conceptualization of AI-based Digital Assistants

Based on this generic conceptualization, Figure 2 depicts a simple example of an interaction with an AI-based digital assistant in a smart home scenario. A child seated in the living room can interact with a speech-based conversational agent (e.g. in Amazon's Alexa) to switch on the overhead light without touching the wall light switch. The AI-based digital assistant takes the recorded audio data as input, recognizes the speech, and tries to understand the spoken language. As part of the dialogue management, the expected action is decided (e.g. the light is turned on via the connected smart home sensors). Finally, a confirming response is generated and delivered to the child via text-to-speech synthesis.




Figure 2. Example of an AI-based Digital Assistant in a Smart Home

With further technological advances in AI, more sophisticated scenarios will be realized in the future, with both positive as well as negative consequences for humans. Thus, human-AI interaction ranges from substitution (AI replaces humans), to augmentation (humans and AI augment one another), to assemblage (AI and humans are dynamically brought together to function as an integrated unit) (Dellermann, Ebel, Söllner, & Leimeister, 2019; Rai, Constantinides, & Sarker, 2019). As a socio-technical discipline, the BISE community is challenged to provide scientifically grounded and practice-relevant answers to the question how the interplays between users, tasks, and technologies in AI-based digital assistants should be shaped so as to achieve a good tradeoff between positive and negative consequences.

This discussion section follows a panel at the International Conference on Wirtschaftsinformatik in March 2019 in Siegen (WI 2019) and presents different perspectives on AI-based digital assistants. It sheds light on (1) application areas, opportunities, and threats as well as (2) the BISE community's roles in the field of AI-based digital assistants. The different authors' contributions emphasize that BISE, as a socio-technical discipline, must address the designs and the behaviors of AI-based digital assistants as well as their interconnections. They have identified multiple research opportunities to deliver descriptive and prescriptive knowledge, thereby actively shaping future interactions between users and AI-based digital assistants. We trust that these inputs will lead BISE researchers to take active roles and to contribute an IS perspective to the academic and the political discourse about AI-based digital assistants.

Alexander Maedche

Karlsruhe Institute of Technology (KIT)

Christine Legner

University of Lausanne

2. Modes of Human-Machine Collaboration and Opportunities for Future BISE Research

The question how AI-based digital assistants can be fruitfully applied in our daily lives, in firms, and in society is a matter of human-machine collaboration. It is not about how machines can trump people or how people can tame machines. Instead, the biggest application potential for AI-based digital assistants is mutually beneficial cooperation – in symbiosis. Both people and machines have relative strengths. While machines are ideal for conducting repeatable, highly structured tasks, in collecting, storing, and processing huge amounts of data, and in predicting the future in fairly stable environments, persons can handle abstract problems and can deal with fragmented information much more efficiently. Further, persons are much better at putting information into a bigger context and at drawing on intuition, empathy, and ethics to underpin their decisions. When persons and machines combine and meaningfully complement their relative strengths, AI-based decisions can lead to most beneficial outcomes. But how strongly should machines and persons combine their strengths in decision-making processes?

The answer must be: It depends. It depends on every concrete application context and on the potential harm at stake. I visualize a continuum between two poles, *humans decide on their own* and *autonomous decision-making*. The most interesting collaboration modes lie between these two poles, with varying intensities of AI inclusion in decision-making processes, such as *assisted decision-making*, *verified decision-making*, and *delegated decision-making*. For instance, actuaries in insurance companies face cases with different standardization levels to calculate insurance premiums. Standard cases with risk models that require similar, highly structured data sources (e.g. car insurance) could be supported in a collaboration mode in which an actuary would only verify the suggestions made by AI-based digital assistants. However, for complex and unique cases (e.g. the insurance of large production plants) in which highly fragmented data from various sources must be combined and actuaries' experience and

oversight is crucial, the most adequate collaboration mode is *assisted decision-making*, if not *humans decide on their own*. Keeping human decision-makers closely involved in and in the loop of decision-making processes maximizes the opportunities to apply AI-based digital assistants, and minimizes AI's potential downsides (e.g. the 'black box' problem).

In my view, BISE can make several original contributions to research into AI-based digital assistants. First, the BISE community is well-equipped, with a long tradition of combining design science with behavioral research. Especially in recent years, both research approaches have coalesced to allow full-cycle or multi-cycle research journeys and programs (e.g. Chatman & Flynn, 2005; Sturm & Sunyaev, 2019). These journeys and programs uniquely enable researchers to anticipate both the beneficial and harmful implications of technologies and to systematically incorporate ethical, social, and psychological theories into the design of AI-based digital assistants. Second, the BISE discipline has a long interdisciplinary tradition – BISE researchers have always worked on socio-technical phenomena at interfaces with other disciplines, such as computer science, psychology, or management. Seeking to understand and integrate theories, concepts and empirical findings from different fields, BISE scholars are well versed in looking at research phenomena from different scholarly perspectives and at aligning hard and soft sciences. BISE should be well positioned to assume the role of a linchpin in interdisciplinary research endeavors to address AI-based research problems that cannot be addressed by single disciplines alone. Third, while the BISE community may not be able to design and develop superior AI-based information systems compared to global digital giants such as Google or IBM, BISE scholars have several key capabilities, such as rigorous scientific methods as well as balanced and neutral perspectives, which can serve as correctives and counter-balances to global digital companies' interest-driven activities. In my view, in the future, an increasingly important task of the BISE community will be to ask uncomfortable questions concerning AI's power and to point out where we must draw the line.

In both the BISE and IS disciplines, we are still at an early research stage into AI generally and AI-based digital assistants in particular. According to Gregor (2006)'s seminal paper on the nature of theory in IS (Gregor, 2006), we are in a stage in which we are predominantly developing taxonomies and describing attributes of AI-based digital assistants (*theory for analyzing and description*). We are only slowly advancing to higher-level theorizing (*theory of explanation or prediction*). There are multiple opportunities for future research about how design artifacts can be developed and used based on sound theories to explain and predict outcomes at different levels of analysis (e.g. individuals, teams, firms, society). To illustrate potential avenues for future research, human-computer interaction research or digital nudging research can benefit greatly from investigating the infusion of digital design artifacts with AI features and capabilities at human-computer interfaces. For instance, it would be interesting to advance our understanding on how anthropomorphism – defined as the attribution of human-like (physical or non-physical) features, behaviors, emotions, characteristics, and attributes to a non-human agent or an inanimate object – affects the interaction quality between users and anthropomorphic smart devices at work and at home (Benlian, Klumpe, & Hinz, 2019; Pfeuffer, Benlian, Gimpel, & Hinz, 2019). As another example, the business model and digital transformation research (Riedl, Benlian, Hess, Stelzer, & Sikora, 2017; Veit et al., 2014) may look into the novel and innovative ways in which AI-based logics are influencing the core elements of business models and how they may shape companies' IT function or digital transformation strategies (Haffke, Kalgovas, & Benlian, 2017; Hess, Benlian, Matt, & Wiesböck, 2016).

In sum, like all technologies before it, AI is not an inevitable fate. AI-based digital assistants can be carefully and mindfully shaped. The BISE community should contribute its fair share in this regard.

Alexander Benlian

3. AI-based Systems' Explanations: Still a Topic for IS Research

Advancements in artificial intelligence (AI) research have resulted in technological capabilities that open additional potential for automation, specifically of cognitive tasks. AI-based systems assist users in an increasing number of contexts and are supposed to bring about profound changes in the ways we work (vom Brocke et al., 2018). Possible applications in business include handling customer complaints, allocating advertising budgets, and optimizing warehouse logistics. AI-based systems' recent performance gains in these and other tasks are based on deep learning methods that employ artificial neural networks. However, the high performance of deep learning methods comes at the cost of high model complexity and low interpretability (Lipton, 2018), that is, these systems constitute 'black boxes' for their users. Thus, the application of deep learning methods has important implications for individuals, organizations, and society, while offering promising starting points for future IS research.

Over its history, AI research has been characterized by a wide variety of different goals and methods. However, two fundamental approaches to realize AI-based systems have emerged (Russel & Norvig, 2010). In the symbolic reasoning approach, developers encode and store knowledge in a knowledge base to solve tasks, drawing rule-based inferences from that knowledge. Knowledge-based systems can provide justifications for their solutions or recommendations, but codifying knowledge requires substantial effort, and inference rules become increasingly complex with the breadth of the knowledge domain, which restricts these systems' capabilities. In machine learning, on the other hand, systems perform tasks using statistical models. Developers optimize (i.e. train) these models by extracting patterns from data of solutions to similar past problems or by letting the system gain experience from feedback over time.

Artificial neural networks are a model class that has been proven useful to solving a wide variety of tasks. However, artificial neural networks' complexity usually prohibits determining why a system based on such models have reached a specific solution. The lack of interpretability exacerbates the application of such systems in several contexts. For instance, in 2017, Amazon decided to abandon a recruiting recommendation tool after finding out that it discriminated against women (Dastin, 2018). The developers had trained the system using data from more than 10 years of incoming past applications and hiring decisions, in which males dominated. This bias went unnoticed during the development of the system, because it did not derive and justify its recommendations based on reasoning, but applied a complex statistical model that developers and users could not interpret.

Building interpretable AI-based systems is a widely addressed topic in AI research (Lipton, 2018) and poses important questions to the IS discipline. One approach to alleviate a lack of system interpretability is to build explanation facilities for AI-based systems. Explanations' roles have been a topic in IS research into earlier AI-based systems, such as expert and recommender systems (Rzepka & Berger, 2018). This research has shown that providing explanations enhances trust in AI-based systems' recommendations and improves users' performance (Wang and Benbasat 2007). However, to date, the implications of a lack of model interpretability and thus explanations owing to the technological differences between AI-based systems in the past and today are less clear.

Thus, IS researchers may investigate whether users recognize and how they perceive a lack of system interpretability. Such perceptions may have negative consequences for users' trust in these systems and, thus, lower usage intentions. Companies that employ AI-based systems must understand how their reliability can be judged and ensured. This may have implications for these systems' admissibility in various business contexts. Organizations must also determine who accepts responsibility for mistakes following the use of AI-based systems. The new

contexts and tasks in which AI-based systems can assist users today offer opportunities to further investigate whether and how explanation facilities can provide effective remedies for these issues, including the questions which explanation type(s) users require and how these explanations should be designed. Given the existing conceptual foundations in this area, the IS community is well positioned to address these questions and thus to make important contributions towards responsible applications of AI-based systems.

Benedikt Berger and Thomas Hess

Ludwig-Maximilians-Universität München

4. AI-based Interactive Assistance Systems and BISE Research, and Why We Must Swiftly Adopt this Topic

The rise of AI-based digital assistants has opened a wide research area for IS scholars. It is a technology with an explicit interface to persons and could therefore provide a fruitful avenue for human-computer interaction (HCI) research. Researchers from this area must think about effective designs and may build on research from the area of robotics and anthropomorphism (for a discussion, see Pfeuffer et al., 2019). Also, behavioral insights in this area are likely to be new and may be publishable in our field's top journals (Benlian et al., 2019).

Further, such interactive assistance systems are likely to impact on individual economic behaviors as well as at the aggregate level (Robertson et al., 2015). Assistance systems and ecosystems around such systems are likely to be of economic importance, and IS researchers who work at the intersection of IS and economics need to help us to understand this technology's economic impacts.

IS researchers who work at the intersection with organization science may be interested in understanding how this new technology type may shape enterprises of the future (Meyer von

Wolff, Hobert, & Schumann, 2019). In this regard, there are many open questions concerning the responsibility of and accountability for decision-making when non-human agents get involved.

Design science researchers must figure out how to build effective, performant systems that address the needs of prospective users. Our discipline could benefit from working together and then presenting holistic approaches and insights that incorporate our engineering expertise and our expertise in behavioral sciences. It is likely that this will be a race against time or against other disciplines such as marketing or economics (see e.g. the new section ‘Frontiers in Marketing Science’ in *Marketing Science* that welcomes manuscripts that focus on “generating early insights about novel business practices” and promises fast turnarounds) and computer science, which is making more and more important contributions in the area of HCI.

We will find such assistants in a plethora of application areas, including smart homes (Benlian et al., 2019), smart cars (Mihale-Wilson, Zibuschka, & Hinz, 2019), robo-advisory (Adam, Toutaoui, Pfeuffer, & Hinz, 2019; Jung, Dorner, Glaser, & Morana, 2018), customer services (Gnewuch, Morana, & Maedche, 2017), in electronic commerce (Qiu & Benbasat, 2009), in healthcare (Laranjo et al., 2018), or as pedagogical agents (Fryer, Ainley, Thompson, Gibson, & Sherlock, 2017).

The implementation of AI-based interactive assistance systems will likely shape the future of many IT-based ecosystems, and companies and economists may be interested in the risks of winner-takes-all-markets or oligopolies and their own market structure.

Besides the economic impacts, these systems will also impact on humans. Especially children may get used to these systems and may relate to them as family members. These systems may also make us lazier, and we may lose the ability to solve some tasks ourselves (finding a route without route guidance assistants, calculating without pocket calculators, driving a car, etc. pp).

It is unlikely that humans are able to solve the special cases, when the assistants fail to deliver; for instance, taking over a self-driving car on a road covered by ice, once humans lack prior experience because assistants usually do the routine tasks.

Finally, many of these systems rely on big data, and new privacy issues are arising. Policymakers need experts (who may come from our discipline) to adequately address these new challenges. Overall, the area of AI-based interactive assistance systems, with all its challenges and opportunities, is perfectly suited for IS research.

Oliver Hinz

Goethe University Frankfurt, Germany

5. AI-based Digital Assistants May Be Valuable Companions – if We Get the Values Right

The proliferation of AI-based digital assistants is worth supporting – it requires us to carve out the opportunities and to constrain the dark side. Among the multiple opportunities of AI-based digital assistants to improve our lives and organizations, two stand out: at the level of the individual user, this is the path toward unbounded rationality. At the team level, this is advancing collective intelligence.

The rational homo oeconomicus, with unlimited cognitive capacity to make optimal decisions, is a standard neoclassical economic perspective. While it allows for elegant mathematical treatment of human behavior, it lacks real world fidelity. A more descriptively valid perspective assumes bounded rationality: Individuals maximize their benefits under cognitive constraints. Individuals are assumed to make tradeoffs between effort and the marginal utility of further information processing. We use heuristics and are prone to biases. Digital assistants with computer-implemented cognitive skills allow us to relax our cognitive boundaries. Especially

in situations of information overload, they can help us to filter, sort, navigate, and process information. This can bring us closer to the ideal of unboundedly rational individual decision-making. Some examples: Many cars have emergency brake assistants operating at an information processing speed that exceeds human capabilities. E-commerce recommender systems help us to navigate an overflow of offerings. Automated e-mail filtering distinguishes spam from clutter and from high-priority e-mails so as to focus our attention. Workplace data discovery tools help us to navigate the sheer unmanageable trove of data provided by colleagues. The move to unbounded rationality is a continuance of work at the intersection of HCI and AI. In light of the increasing prevalence of anthropomorphic IS, it is important to consider not only the computer-implemented cognitive features but also the visual, auditory, emotional, and behavioral features of such assistants (Pfeuffer et al., 2019).

I will now focus on the level of groups and teams, and the benefits offered by AI-based digital assistants. Collective intelligence is desirable for groups and teams (Malone in Gimpel, 2015). Teams nowadays primarily use digital technologies purely as communication media. In the near future, AI-based team assistants could act as facilitators or team members. As facilitators, they could support team processes, for instance, by highlighting areas of agreement, conflict, and progress, upholding team norms and uncovering team dynamics. As team members, they could work ‘at eye level’ with human participants in hybrid human AI assemblages (Rai et al., 2019). This will be a substantial expansion of the work at the intersection of computer-supported collaborative work, collective intelligence, and AI.

By supporting a trend towards individually unbounded rationality and collective intelligence, AI-based assistants contribute to a brighter society. This requires us to redefine the division of labor between humans and machines. It requires substantial technological development to build better tools that support human users and will require human adaptation to harness these new tools’ power. As any transition, this one also has a dark side, one of adverse risks and side-

effects (Gimpel & Schmied, 2019). The major challenges include assuring that AI-based digital assistants support moral principles such as doing good, doing no harm, being transparent, maintaining human autonomy, and being non-discriminatory (AI HLEG – European Commission’s High-Level Expert Group on Artificial Intelligence, 2018). We need *ethics-by-design* as a non-functional requirement for AI-based systems.

I will now address a specific ethical challenge: maintaining human autonomy. Human autonomy builds on freedom of will. Ambient, persuasive AI-based assistants have the possibility to boost human autonomy by supporting deliberate and intentional actions. However, such systems’ (technical) autonomy may interfere with human agency. Digital nudging (Weinmann, Schneider, & Brocke, 2016) may be benevolent paternalism, but the borderline to manipulation is not clear-cut and is sometimes crossed. Persuasive systems can be supportive and engaging, but may lead to addiction. Automated decisions (e.g. IoT devices ordering products) may be convenient, but deprive us of control. Delegating tasks to digital assistants may free up cognitive and physiological resources, but is accompanied by increasing incompetence to perform the tasks oneself. Here, we need a societal discourse about what is desirable or acceptable, and we need methods to engineer systems that comply with these principles.

Beyond maintaining human autonomy, being non-discriminatory is an important moral principle. Yet, historical and current individual, business, and political practices don’t always adhere to this principle. Machine learning offers great potential to power AI-based assistants. However, when the training data are biased, the machine learning algorithms may pick up these biases and may then perpetuate them; this leads to a continuance of discriminatory decisions. Even if this is against the intentions of the developers and users, it may happen without their knowledge. Examples of discriminatory algorithms are Amazon’s presumably sexist

recruitment support system (no longer operational)¹, Northpointe's presumably racist recidivism scores used in the U.S. criminal justice system (currently operational)², or Google's presumably racist image tagging service (fixed in this respect)³. We need more and better tools to debug and audit systems based on machine learning, and we need higher awareness among developers and users of machine learning-based systems.

As researchers working on business and information systems engineering, we should seek to support a bright digital future. We have the theoretical and methodological background to contribute descriptive and prescriptive knowledge about AI-based assistants as well as our individual and collective interactions with such systems. Being versed in interdisciplinary cooperation is an asset to lever here. Our interdisciplinary discourses with colleagues from management, economics, and computer science remain key. In light of the opportunities outlined above, an intensification of the collaboration with (cognitive and social) psychologists is a fruitful avenue. In light of the challenges outlined above, our discipline will benefit from strengthening discourses with scholars from philosophy, especially ethics. Filling the void between ethics, management, economics, and computer science in both behavioral and design science is a challenge, but it is a fun challenge for which we as a community are well equipped.

Henner Gimpel

University of Augsburg, FIM Research Center and Project Group Business and Information Systems of Fraunhofer FIT, Germany

¹ <https://www.theguardian.com/technology/2018/oct/10/amazon-hiring-ai-gender-bias-recruiting-engine>. Accessed: 2019-03-26. Archived by WebCite® at <http://www.webcitation.org/77ACKVmSo>

² <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>. Accessed: 2019-03-26. Archived by WebCite® at <http://www.webcitation.org/77ACdF4bj>

³ <https://edition.cnn.com/2015/07/02/tech/google-image-recognition-gorillas-tag/index.html>. Accessed: 2019-03-26. Archived by WebCite® at <http://www.webcitation.org/77ACoiDyn>

6. Designing Cooperation between Humans and AI-based Digital Assistants

Currently, there is an intensive discourse on the different ways humans are interacting with AI-based technologies (e.g. Rai et al., 2019) and how the performance of a certain task should be divided between these two entities. It is important to involve humans to an appropriate level in the task performance, depending on the task characteristics and the context (Davenport & Kirby, 2016). This is particularly important for AI-based digital assistants, since the task performance here is a cooperative effort. Thus, an important challenge for future research is to investigate how to distribute the task performance between these two entities at an appropriate level in order to achieve desired outcomes. When investigating this distribution, we must consider not only performance-related outcomes (e.g. effectiveness, efficiency), but also individuals' cognitive states (e.g. mental effort, situation awareness) as well as individuals' attitudes to and perceptions of the AI-based digital assistant in question (e.g. trust, usage intentions). The generic conceptualization proposed in this article (see Figure 1) can be a starting point for further research into the multifaceted, contextualized interplays between humans, tasks, AI-based technologies, and the various resulting outcomes.

The BISE community is especially qualified to address this research opportunity from both the behavioral as well as the design research perspective. I will now outline three potential avenues for future research into cooperation between humans and AI-based digital assistants.

First, there is need to investigate from a conceptual perspective the interplays between humans and tasks when using AI-based digital assistants. Specifically, in my view, there is a need to classify the different task types and their characteristics carried out with AI-based digital assistants as well as the resulting outcomes. There are instrumental outcomes, such as effectiveness and efficiency, as well as humanistic outcomes, such as the mental effort during task performance. Further, the task performance's context should be considered in such

conceptualization, because it will influence the outcomes of cooperative task performance. This first research endeavor should result in an agreed-upon conceptual framework that describes cooperation between humans and AI-based digital assistants.

Second, based on this framework, the BISE community can investigate the design of AI-based digital assistants, focusing on the different conceptual dimensions. For instance, there is a need to investigate design variants of AI-based digital assistants for different task types. Depending on the task type, different cooperation forms between humans and AI-based digital assistants should be instantiated in order to achieve specified outcomes. For instance, simple tasks, such as creating an appointment from an e-mail request, could very well be handled by the digital assistant, with only little involvement by a person. However, for more complex tasks or decisions, it may be necessary that the human and the digital assistant jointly perform a specific task, or that the human takes over the primary task performance and the AI-based digital assistant is only supportive. Depending on the extent of human involvement in task performance, different designs of the AI-based digital assistant may be appropriate. In addition to the task type, the context should also be considered when designing AI-based digital assistants. For instance, the interaction with a digital assistant in a private life context differs to an organizational context. While it is very convenient to interact with a voice-based digital assistant (such as Amazon Alexa or Google Home) in spoken language at home, it may be confusing or even disturbing to use such an interaction mode in an open-plan office. Thus, providing generic and context-specific design knowledge for AI-based digital assistants is an interesting future research opportunity in the BISE community.

Further, human characteristics must be considered in the design of AI-based digital assistants. Individual characteristics such as expertise with the technology and personality have important roles in how humans interact with AI-based digital assistants to perform tasks. Moreover, persons' attitudes and perceptions of the resulting interactions must also be investigated. For

instance, we must understand in which conditions humans establish trust in AI-based digital assistants.

In sum, cooperation between humans and AI-based digital assistants provides multiple research opportunities for the BISE community to contribute both descriptive and prescriptive knowledge on this interesting and current phenomenon.

Stefan Morana

Karlsruhe Institute of Technology (KIT), Germany

7. AI-based Digital Assistants as Tutors and BISE Researchers as Integrators of Interdisciplinary Insights and Creators of Design Knowledge

Increasing student-teacher ratios are a challenge for many schools and universities worldwide and an opportunity for the members of the BISE community in their roles as researchers and teachers. For instance, according to the Organisation for Economic Co-operation and Development (OECD), the number of students at universities in Germany rose by 29 percentage points from 2004 to 2014 (OECD, 2016), while public spending for education decreased by 1 percentage point. This development leads to situations most of us have already experienced: larger class sizes, especially in curricula that cover current topics such as digitalization. For instance, while the number of students in University of St. Gallen's Master in Business Innovation increased from 224 in 2014 to 507 in 2018, such an increase in student numbers typically has not resulted in increases in budgets or in teaching resources. Thus, factors such as interaction, feedback, and individualization that have been shown to strongly impact on learning outcomes (Hattie, 2015) suffer, on average.

Within these boundary conditions, AI-based digital assistants have the potential to help counter this development, since they provide novel opportunities to increase the levels of interaction, feedback, and individualization in the learning process, and – therefore – learning outcomes.

Especially concerning abilities, such as problem-solving skills, early work shows that assistants such as Amazon Alexa can take a role comparable to a human tutor, and can help learners to increase their task performance (Winkler, Söllner, Neuweiler, Rossini, & Leimeister, 2019) and their skills levels over time. At the same time, we need to be aware of the limitations of today's AI-based digital assistants. For instance, in our research and teaching, we focus on using these assistants to help students in their learning process, or concerning the structural features of an essay. We don't rely on such systems to assess the quality of student essays' content, since the semantic capabilities of today's systems are not yet well enough developed to understand the content at the level we deem necessary.

Another point – that does not address the uses of AI-based digital assistants in education, but rather their impacts on education in general – arose during the lively discussion with the audience in our panel discussion at WI 2019 in Siegen. Given the assumption that AI-based digital assistants may complete many routine tasks in the future, and that humans focus more on complex tasks, how does this development impact on the education of future employees? Typically, career paths rely on employees gathering experience by completing easier tasks, to prepare for taking over more complex tasks later on. How do we need to redesign our curricula to ensure that employees are ready to complete tasks that are more complex when this period of completing easier tasks disappears? These are important questions that we must consider now, given the speed at which AI-based digital assistants are entering different areas of society.

In my view, BISE researchers are well positioned to shape the future development of AI-based digital assistants and to ensure that their design is in line with goals and values of the society they are used in. Researchers interested in the domain of collaboration may also be interested in the research agenda for machines as teammates, co-authored with 11 colleagues (Seeber et al., 2018). BISE researchers can contribute to the design of such assistants through research in three areas, which will guide the following discussion: understanding human users' needs of

AI-based digital assistants, integrating theoretical and normative insights from multiple disciplines, and codifying design knowledge that helps designers to design desirable AI-based digital assistants.

Understanding human users' needs of AI-based digital assistants. One pillar for fostering the design of AI-based digital assistants that are in line with goals and values of the society they are used in is creating a rich understanding of the needs of the potential users of such assistants. Here, we can build on our experience in research domains such as technology acceptance and task-technology fit, as well as on established approaches such as design thinking. The goal in this research stream should be the creation of nuanced theoretical knowledge on what different user types expect from using AI-based digital assistants in certain contexts.

Integrating theoretical and normative insights from multiple disciplines. The second pillar is enriching knowledge about user needs by adding further theoretical and normative lenses, from the IS discipline as well as from adjacent disciplines that are relevant for the design of AI-based digital assistants that are in line with goals and values of the society they are used in. Examples are normative and theoretical insights from law, ethics, or sociology. Here, we can build on our experience as mediators between different stakeholders, for instance managers and developers, in order to integrate different perspectives on a specific object of interest. The goal in this research stream should be the identification and resolution of potential conflicts between user needs and demands or insights that stem from different normative or theoretical bases, in order to lay the foundation for the codification of design knowledge that accounts for as many as possible of the interdisciplinary facets necessary to design AI-based digital assistants that are in line with goals and values of the society they are used in.

Codifying design knowledge that helps designers to design desirable AI-based digital assistants. In this research stream, we can leverage our experience in conducting rigorous and

relevant design science research. The focus should be the translation of the created theoretical knowledge into properly codified design knowledge that can help both researchers and practitioners to design proper AI-based digital assistants. Examples of such design knowledge are design principles, requirements, and design patterns, but also methodological approaches that can guide researchers and practitioners throughout the design process.

Matthias Söllner

University of Kassel, Germany & University of St. Gallen, Switzerland

References

- Adam, M., Toutaoui, J., Pfeuffer, N., & Hinz, O. (2019). Investment Decisions with Robo-Advisors: The Role of Anthropomorphism and Personalized Anchors in Recommendations. In *Proceedings of the 27th European Conference on Information Systems (ECIS), Stockholm & Uppsala, Sweden*. https://aisel.aisnet.org/ecis2019_rp/33/
- AI HLEG – European Commission’s High-Level Expert Group on Artificial Intelligence. (2018). *Draft Guideline for Trustworthy AI*. Brussel, Belgium.
- Benlian, A., Klumpe, J., & Hinz, O. (2019). Mitigating the Intrusive Effects of Smart Home Assistants by using Anthropomorphic Design Features: A Multi-Method Investigation. *Information Systems Journal*, (forthcoming). <https://doi.org/10.1111/isj.12243>
- Benyon, D. (2014). *Designing interactive systems: A comprehensive guide to HCI, UX and interaction design*. Pearson Education Limited.
- Chatman, J. A., & Flynn, F. J. (2005). Full-Cycle Micro-Organizational Behavior Research. *Organization Science*, 16(4), 434–447. <https://doi.org/10.1287/orsc.1050.0136>

- Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. Retrieved March 14, 2019, from <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>
- Davenport, T. H., & Kirby, J. (2016). Just How Smart Are Smart Machines? *Sloan Management Review*, 57(3), 20–25.
- Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid Intelligence. *Business & Information Systems Engineering*, (forthcoming). <https://doi.org/10.1007/s12599-019-00595-2>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them. *Management Science*, 64(3), 1155–1170. <https://doi.org/10.1287/mnsc.2016.2643>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Fryer, L. K., Ainley, M., Thompson, A., Gibson, A., & Sherlock, Z. (2017). Stimulating and sustaining interest in a language course: An experimental comparison of Chatbot and Human task partners. *Computers in Human Behavior*, 75, 461–468. <https://doi.org/10.1016/j.chb.2017.05.045>
- Gartner. (2019). Gartner Predicts 25 Percent of Digital Workers Will Use Virtual Employee

- Assistants Daily by 2021. Retrieved March 18, 2019, from <https://www.gartner.com/en/newsroom/press-releases/2019-01-09-gartner-predicts-25-percent-of-digital-workers-will-u>
- Gimpel, H. (2015). Interview with Thomas W. Malone on “Collective Intelligence, Climate Change, and the Future of Work.” *Business & Information Systems Engineering*, 57(4), 275–278. <https://doi.org/10.1007/s12599-015-0382-4>
- Gimpel, H., & Schmied, F. (2019). Risks and Side Effects of Digitalization: A Multi-Level Taxonomy of the Adverse Effects of Using Digital Technologies and Media. In *Proceedings of the 27th European Conference on Information Systems (ECIS 2019), Stockholm & Uppsala, Sweden*. https://aisel.aisnet.org/ecis2019_rp/145/.
- Gnewuch, U., Morana, S., & Maedche, A. (2017). Towards Designing Cooperative and Social Conversational Agents for Customer Service. In *Proceedings of the International Conference on Information Systems (ICIS 2017), Seoul, Korea*. <https://aisel.aisnet.org/icis2017/HCI/Presentations/1/>
- Goodhue, D. L., & Thompson, R. L. (1995). Task-Technology Fit and Individual Performance. *MIS Quarterly*, 19(2), 213. <https://doi.org/10.2307/249689>
- Google AI Blog. (2018). Google Duplex: An AI System for Accomplishing Real-World Tasks Over the Phone. Retrieved March 18, 2019, from <https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html>
- Gregor, S. (2006). The nature of theory in Information Systems. *MIS Quarterly*, 30(3), 611–642.
- Haffke, I., Kalgovas, B., & Benlian, A. (2017). Options for Transforming the IT Function Using Bimodal IT. *MIS Quarterly Executive*, 16(2), 101–120.

- Hattie, J. (2015). The applicability of Visible Learning to higher education. *Scholarship of Teaching and Learning in Psychology*, 1(1), 79–91. <https://doi.org/10.1037/stl0000021>
- Heinrich, L., Heinzl, A., & Riedl, R. (2011). *Wirtschaftsinformatik: Einführung und Grundlegung*. Springer.
- Hess, T., Benlian, A., Matt, C., & Wiesböck, F. (2016). Options for Formulating a Digital Transformation Strategy. *MIS Quarterly Executive*, 15(2), 123–139. <https://doi.org/10.1108/10878571211209314>
- Jung, D., Dorner, V., Glaser, F., & Morana, S. (2018). Robo-Advisory: Digitalization and Automation of Financial Advisory. *Business and Information Systems Engineering*, 60(1). <https://doi.org/10.1007/s12599-018-0521-9>
- Laranjo, L., Dunn, A. G., Tong, H. L., Kocaballi, A. B., Chen, J., Bashir, R., ... Coiera, E. (2018). Conversational agents in healthcare: a systematic review. *Journal of the American Medical Informatics Association*, 25(9), 1248–1258. <https://doi.org/10.1093/jamia/ocy072>
- Lipton, Z. C. (2018). The mythos of model interpretability. *Communications of the ACM*, 61(10), 36–43. <https://doi.org/10.1145/3233231>
- Maedche, A., Morana, S., Schacht, S., Werth, D., & Krumeich, J. (2016). Advanced user assistance systems. *Business and Information Systems Engineering*, 58(5). <https://doi.org/10.1007/s12599-016-0444-2>
- Meyer von Wolff, R., Hobert, S., & Schumann, M. (2019). How May I Help You? – State of the Art and Open Research Questions for Chatbots at the Digital Workplace. In *Proceedings of the 52nd Hawaii International Conference on System Sciences (HICSS 2019)*, pp. 95–104.
- Mihale-Wilson, C., Zibuschka, J., & Hinz, O. (2019). User Preferences and Willingness to Pay

- for In-Vehicle Assistance. *Electronic Markets*, 29(1), pp. 37–53. <https://doi.org/10.1007/s12525-019-00330-5>
- OECD. (2016). *Education at a Glance 2016: OECD Indicators*. Paris, France.
- Pfeuffer, N., Benlian, A., Gimpel, H., & Hinz, O. (2019). Anthropomorphic Information Systems. *Business & Information Systems Engineering*, (forthcoming). <https://doi.org/10.1007/s12599-019-00599-y>
- Qiu, L., & Benbasat, I. (2009). Evaluating Anthropomorphic Product Recommendation Agents: A Social Relationship Perspective to Designing Information Systems. *Journal of Management Information Systems*, 25(4), 145–182. <https://doi.org/10.2753/MIS0742-1222250405>
- Rai, A., Constantinides, P., & Sarker, S. (2019). Next-Generation Digital Platforms: Toward Human–AI Hybrids. *MIS Quarterly*, 43(1), iii–ix.
- Reddy, T. (2017). How chatbots can help reduce customer service costs by 30%. Retrieved March 22, 2019, from <https://www.ibm.com/blogs/watson/2017/10/how-chatbots-reduce-customer-service-costs-by-30-percent>
- Riedl, R., Benlian, A., Hess, T., Stelzer, D., & Sikora, H. (2017). On the Relationship Between Information Management and Digitalization. *Business & Information Systems Engineering*, 59(6), 475–482. <https://doi.org/10.1007/s12599-017-0498-9>
- Robertson, S., Solomon, R., Riedl, M., Gillespie, T. W., Chociemski, T., Master, V., & Mohan, A. (2015). The Visual Design and Implementation of an Embodied Conversational Agent in a Shared Decision-Making Context (eCoach). In P. Zaphiris & A. Ioannou (Eds.), *Learning and Collaboration Technologies* (pp. 427–437). Springer, Cham. https://doi.org/10.1007/978-3-319-20609-7_40

- Russel, S., & Norvig, P. (2010). *Artificial Intelligence: A Modern Approach*. Upper Saddle River, New Jersey: Pearson.
- Rzepka, C. C., & Berger, B. (2018). User Interaction with AI-enabled Systems: A Systematic Review of IS Research. In *39th International Conference on Information Systems, San Francisco*. In *Proceedings of the 39th International Conference on Information Systems (ICIS 2018), San Francisco*. <https://aisel.aisnet.org/icis2018/general/Presentations/7/>
- Seeber, I., Bittner, E., Briggs, R. O., de Vreede, G.-J., de Vreede, T., Druckenmiller, D., ... Söllner, M. (2018). Machines as Teammates: A Collaboration Research Agenda. <https://doi.org/10.24251/HICSS.2018.055>
- Sturm, B., & Sunyaev, A. (2019). Design Principles for Systematic Search Systems: A Holistic Synthesis of a Rigorous Multi-cycle Design Science Research Journey. *Business & Information Systems Engineering*, 61(1), 91–111. <https://doi.org/10.1007/s12599-018-0569-6>
- Veit, D., Clemons, E., Benlian, A., Buxmann, P., Hess, T., Kundisch, D., ... Spann, M. (2014). Geschäftsmodelle. *Wirtschaftsinformatik*, 56(1), 55–64. <https://doi.org/10.1007/s11576-013-0400-4>
- vom Brocke, J., Maaß, W., Buxmann, P., Maedche, A., Leimeister, J. M., & Pecht, G. (2018). Future Work and Enterprise Systems. *Business & Information Systems Engineering*, 60(4), 357–366. <https://doi.org/10.1007/s12599-018-0544-2>
- Weinmann, M., Schneider, C., & Brocke, J. vom. (2016). Digital Nudging. *Business & Information Systems Engineering*, 58(6), 433–436. <https://doi.org/10.1007/s12599-016-0453-1>
- Winkler, R., Söllner, M., Neuweiler, M. L., Rossini, F. C., & Leimeister, J. M. (2019). Alexa,

Can You Help Us Solve This Problem? How Conversations With Smart Personal Assistant Tutors Increase Task Group Outcomes. In *Proceedings of the 2019 CHI Conference Late-Breaking Work on Human Factors in Computing Systems (CHI LBW '19)*.