

Grand Challenges for Ambient Intelligence and Implications for Design Contexts and Smart Societies

Norbert Streitz^a, Dimitris Charitos^b, Maurits Kaptein^c, and Marc Böhlen^d,

^a*Smart Future Initiative, Frankfurt am Main, Germany*

^b*Department of Communication and Media Studies, School of Economics and Political Sciences, National and Kapodistrian University of Athens, Greece*

^c*Statistics and Research Methods, Tilburg University, Tilburg, The Netherlands and Jheronimus Academy of Data Science, Den Bosch, The Netherlands.*

^d*Department of Art, Computational Media, University at Buffalo, Buffalo NY, 14260 USA*

Abstract. This paper highlights selected grand challenges that concern especially the social and the design dimensions of research and development in Ambient Intelligence (AmI) and Smart Environments (SmE). Due to the increasing deployment and usage of ‘smart’ technologies determining a wide range of everyday life activities, there is an urgent need to reconsider their societal implications and how to address these implications with appropriate design methods. The paper presents four perspectives on the subject grounded in different approaches. First, introducing and reflecting on the implications of the ‘smart-everything’ paradigm, the resulting design trade-offs and their application to smart cities. Second, discussing the potential of non-verbal communication for informing the design of spatial interfaces for AmI design practices. Third, reflecting on the role of new data categories such as ‘future data’ and the role of uncertainty and their implications for the next generation of AmI environments. Finally, debating the merits and shortfalls of the world’s largest professional engineering community effort to craft a global standards body on ethically aligned design for autonomous and intelligent systems. The paper benefits from taking different perspectives on common issues, provides commonalities and relationships between them and provides anchor points for important challenges in the field of ambient intelligence.

Keywords: Ambient Intelligence, Artificial Intelligence, Smart Environments, Smart-Everything, Humane and Sociable AmI, Citizen-Centered Design, Smart City, Self-aware City, Hybrid City, Design Trade-offs, Human Control, Human in the Loop, Non-Verbal Communication, Spatial Communication Interfaces, Data Science, Future Data, Multi-armed Bandit Problem, Uncertainty, Privacy by Design, General Data Protection Regulations, GDPR, Autonomous Intelligent Systems, General Artificial Intelligence, Machine Learning, Opaque AI, Algorithmic Transparency, Traceability of Algorithms, Ethically Aligned Design, Governance of Technology.

1. Introduction

This position paper is part of the thematic issue on the occasion of the 10th anniversary of the Journal of Ambient Intelligence and Smart Environments (JAISE). The objective is to highlight selected grand challenges that concern especially the social and the design dimension of research and development in ambient intelligence (AmI) and smart environments. Social and design contexts have changed during the last 10 years, some anticipated, but not addressed in depth,

some are new arrivals deserving timely attention. Due to the increasing deployment and usage of ‘smart’ technologies determining now a wide range of everyday life activities, there is an increasing need to (re)consider the societal implications and imagine how to address them with appropriate design methods.

With reference to the title of the JAISE journal, it is useful to distinguish between the two parts of the journal’s name due to different connotations. Unfortunately, the term ‘ambient intelligence (AmI)’ created in the late 1990s lost traction in the last years despite

very good work in research and development, a trend partly due to the fact that the term and the field of artificial intelligence (AI) has gained more attention in the public, even though much of what is marketed as AI now are applications of supervised learning with all its problems not to be discussed here. On the other hand, so-called ‘smart’ technologies experience widespread implementations and deployments so that the term ‘smart’ became a ubiquitous buzzword (smart objects, smart environments, smart technologies, smart data, smart phones, smart rooms, smart homes, smart cities, smart airports, smart nations, ...) with no clear definition anymore. One can observe that AmI-conferences are in competition with events that carry different labels as, e.g., Internet-of-Things (IoT), ubiquitous and pervasive computing, intelligent or smart environments, etc. Smart environments can be found now at many levels and increasing in scale: from smart artefacts to smart rooms and smart buildings all the way to smart cities addressing a wide range of activities in urban environments. One can also observe an extension of application areas: from office work and learning activities (as more traditional areas) via services for daily routines organizing your life, health and well-being to manufacturing and production enabled by the Industrial Internet also called Industry 4.0 as well as smart farming and agriculture, to marketing and sales. There is no doubt that this immense proliferation has severe implications for society, especially since many, if not most of these developments are very much technology-driven. Thus, there is a responsibility to analyze and diagnose the situation, to provide frameworks for facilitating this condition and to propose and recommend human-centered design approaches that address the pressing issues. Therefore, it is time for the scientific community to pause, take stock of the situation and to propose methods and design guidelines for remedying the deficits and problems of many current technology-driven developments.

This paper provides different perspectives by four authors on the subject. Streitz speaks to the ‘smart-everything’ paradigm, the resulting design trade-offs for privacy and human control and their application to smart cities. Charitos discusses the potential of non-verbal and spatial communication interfaces for AmI design practices. Kaptein shows how new data categories such as ‘future data’ and the role of uncertainty need to be considered for next generation AmI design. Böhlen discusses the current attempt by the world’s largest professional engineering community to craft

global standards for ethically aligned design in artificial intelligence. While each of the contributions offer a distinct perspective, the paper establishes various relationships between them and provides anchor points for important challenges in the field of ambient intelligence. In this paper we will focus on:

- Redefining the ‘smart-everything’ paradigm by moving beyond ‘smart-only’ approaches and addressing inherent design trade-offs between smartness and privacy as well as human control vs. automation.
- Designing AmI experiences as spatial communication interfaces by acknowledging the significance of physical space and social interaction as important design contexts.
- Reflecting on the role of data for AmI environments and applications by acknowledging and incorporating recent advances in data science.
- Enforcing ethical and privacy considerations in the wake of increased collection, processing and exploitation of large amounts of data, especially personal data captured in smart environments.
- Shifting from an R&D exclusive focus in ethical design to practical interventions in ethical design; considering the political dimensions of data management in smart environments; thinking today about the next generation of General Artificial Intelligence with superhuman abilities.
- Applying the considerations to the application domains of ‘smart’ cities and ‘smart’ societies as well as stating claims and recommendations with general relevance for the field of ambient intelligence and smart environments.

While we start out by describing various challenges in individual sections, there are strong correlations and interactions between them, forming a coherent and comprehensive picture of the AmI-related challenges society is confronted with. To make the interactions and dependencies concrete and transparent, we provide an example of the application of our predictions and recommendations in the domain of future urban environments. They are currently known under the label of ‘smart’ cities, but we show that it is necessary to move beyond ‘smart-only’ cities towards humane, sociable, and cooperative hybrid cities reconciling people and technology. In the final section on conclusions and outlook, we describe also claims with general rele-

vance, not only tied to smart cities, but also to society. All of this requires addressing the issues and challenges we describe in the following sections.

2. Redefining the ‘Smart-Everything’ Paradigm

In his seminal paper in *Scientific American*, Mark Weiser [88] described in 1991 the idea of ‘ubiquitous computing’ as the blueprint for the ‘Computer of the 21st Century’. This was followed by several developments of linking the internet to real-world objects by establishing device-to-device data communication, finally resulting in the notion of an Internet of Things (IoT) in the late 1990s and early 2000s. For current overviews on the history and technology developments of IoT, see Chin et al. [19] and Gomez et al. [34].

While these constituted rather technology-driven developments, Ambient Intelligence (AmI) was proposed around the same time to contrast these developments and paying more attention to user-centered design, social interfaces and the notion of a context-aware and adaptive ambient environment. There is no space and intention to provide a historical or comprehensive account here (see, e.g., [2]). Around 10 years after AmI entered the scene and a scientific community was established with various conferences, the *JAISE* journal was founded and published its first issue in 2009. It started with a prominent article by Aarts and de Ruyter [1] providing new research perspectives on AmI, addressing again the contrast between a system perspective and a human-needs oriented AmI vision. The authors also argued to emphasize more the social, emphatic and conscious dimension of interaction in AmI environments. With similar intentions, Streitz and Privat [82] also took stock of the AmI status in 2009 addressing the relationship between IoT, Artificial Intelligence (AI) and AmI and proposed seven contrasting pairs describing design options and their role for the advancement of the AmI vision. A more recent account of the relationship between AI and AmI is provided by Gams et al. [29].

Now, again 10 years later and thus 20 years after putting AmI on the map, it is again time to evaluate and rethink the situation. How can we advance the original AmI vision in the current context of ubiquitous smart technologies that are not anymore research prototypes, but commercial products in everyday use? What are the new constraints and how can the AmI vision play a role?

As described in the introduction, we are confronted with a situation which Streitz characterized by

a ubiquitous diffusion of the ‘*smart-everything*’ paradigm [78, 79]. It is based on the observation that everything must be ‘smart’: specific devices, software, platforms and services. It results from the combination of the Internet of Things (IoT) and Artificial Intelligence (AI), where especially the latter is increasingly in the public focus and promoted to a large extent. Unfortunately, the notion of ‘Ambient Intelligence’ is not so prominent anymore in a significant way, although its approach has a lot to offer. This is accompanied by a loss of many design imperatives being core to the AmI vision and an uptrend of technology-driven approaches, which we consider to be more than questionable. We follow here the British architect Cedric Price who expressed his concerns about technology-driven approaches in the remarkable provocation “*Technology is the answer, but what was the question?*” [61].

The term ‘smart’ is not a problem, but the way it is interpreted and propagated needs critical reflection and alternative perspectives, especially when combined with increasing automation and autonomous systems. For example, we must look at the underlying rationale of ethical considerations and their implications in more detail (see section 6). The extent of collecting, processing and exploiting data, often without consent of the people as their proper owners resulting in privacy infringements (see subsection 3.2) needs a reevaluation of how data are provided and used.

The alternative is provided by an approach that moves beyond ‘smart-only’ environments towards humane and sociable AmI environments. It is rooted in the initial AmI vision and requires redefining the ‘smart-everything’ paradigm [79]. We think that it is time again to promote this humane and social perspective and adapt it to the new constellations. This alternative view is based on design trade-offs described in the next section.

3. Design Trade-offs

We argue that a human-/people-/citizen-centered design approach is needed for going *beyond* ‘smart-only’ technology-driven ubiquitous instrumentations and installations. The approach is characterized by design goals like “*keeping the human in the loop and in control*” and the proposal that “*smart spaces make people smarter*” [79, 83]. There are several problem sets consisting of general concerns about artificial intelligence and algorithmic automation as well as privacy issues. According to Streitz [79], there are at least

two trade-offs (and their combination) to be considered:

- Keeping the human in the loop and in control, thus empowering humans vs. automation or even autonomous importunate behavior of smart environments.
- Ensuring privacy by being in control of making decisions over the use of personal data vs. intrusion of often unwanted, unsupervised and importunate data collection methods as a prerequisite of providing smartness, for example, in terms of smart services.

3.1. Human Control vs. Automation

The first design trade-off concerns the current shift towards more or even complete automation of previously (partially) human operator-controlled activities. Smart devices and underlying algorithms are gaining ground in controlling processes, services and devices as well as the interaction between devices and humans. Humans are increasingly removed from being the operator, supervisor or at least being in charge and thus from being in control. The problems caused by the ‘smart-everything’ paradigm can be categorized in three problem sets: A) Inability and error-prone behavior. B) Rigidity, and C) Missing transparency and traceability. Since a more elaborate description is provided in [79], we mention here only a few examples.

Error-prone behavior or inability of AI or other algorithmic approaches can be observed in many areas despite manifold promises. A major problem is the unresolved dependency of supervised machine learning on having appropriate, unbiased, and sufficient training data of high quality. As a consequence, differences in training data and algorithmic constraints result in very different results/predictions, although they are supposed to provide the same “right” answer [40]. The high expectations towards autonomous driving are disappointed by failures, for example, in recognizing speed-limit signs or being fooled by so-called ‘scam stickers’ [25]. Why are autonomous cars driving too fast although they are supposed to make traffic safer as in the recent deadly accident caused by an Uber car [52]? Moreover, during the phase of having only level 2 to 4 capabilities (which will be the standard for a long time), according to the SAE [67] classification¹, wrong detection information might result in unjustified legal consequences for the human drivers because

they will be still liable for damages. A practical example: the car system sends (incorrect) messages about discrepancies (driving at a speed limit correctly obeyed by the driver which is higher than a lower speed limit wrongly identified by the car) to the police or insurance company. This might result in a fine or increase in insurance premium, although based on wrong information identified and sent by the vehicle.

Rigid behavior is another problem. Users and customers experience it when confronted with fully automated call centers or on-line shops without humans involved. It needs only small deviations from the standard routine or process and the system cannot handle the requests. The problem is that customers are going to lose control and be completely at the mercy of companies and their algorithms with no recourse. Hotel booking systems repeat recommendations for hotels in cities which are not relevant anymore. Customers of on-line shops are confronted with the same category of items just bought, although one does not need multiple items in this category at the same time.

While the first two problem sets might be remedied (partially) by progress in the field, missing transparency and traceability and incomprehensible decisions are and will stay with us as an essential problem. Assuming further “progress”, AI-based behavior will increasingly become non-transparent and incomprehensible to observers. Being untraceable implies also that there are no reproducible outcomes and a lack of liability. People are already now confronted with the problems and lack of transparency, as demonstrated in the financial domain with high frequency trading or decisions on creditworthiness. When nobody can trace the underlying argumentation or mechanisms we have a really serious problem. A more detailed discussion of these issues and additional references can be found in [79]. There are some attempts to address these problems [4, 5]. The General Data Protection Regulation (GDPR) [23, 24] of the European Union (EU), effective since May 2018, highlights also some of these issues. For example, the GDPR requires to provide users with an explanation or the rationale of a decision made by the underlying algorithms and provides citizens with the right to opt-out or to make different decisions.

Furthermore, it seems obvious that these challenges are intricately related with ethical issues. So, it is no coincidence that the world’s largest professional engineering community IEEE engaged in the definition of a standards body on “Ethically Aligned Design” (EAD). The efforts of the IEEE Global Initiative on

¹ According to the SAE [67], progress towards autonomous driving is categorized by levels from 0 to 5, where “0” is fully manual

with no automation and “5” full automation (no human driver needed for supervision).

Ethics of Autonomous and Intelligent Systems (A/IS) (<https://ethicsinaction.ieee.org>) and its recommendations for EAD are discussed in detail in section 6.

The problems and their implications for the welfare as well as a fair and democratic treatment of people and society at large require changing the design approach. We argue for a ‘people-empowering smartness’ instead of a ‘system-centric, importunate and automated smartness’ [79, 83]. People should not only be in control, but they should “own the loop”. People should be empowered by smart capabilities and thus be able to make more informed and mature decisions. Thus, the design goal is that “smart spaces make people smarter”. This people-oriented, empowering smartness approach is in line with the AmI vision. Obviously, there is a caveat to being in control at all levels of the process and making all the decisions because people have limitations on the amount of data they can process. But they should be in control of the trade-off between the degree of automated (pre)processing and aggregation of data vs. the degree of human intervention and decision making. The degree of system automation must be configurable by the user/ citizen. But to be put into the position at all of making this trade-off requires that the options and the type of balance must be anticipated as a design objective and carefully prepared by the system designers in the first place.

3.2. Privacy vs. Smartness

It is no surprise that there is a tricky trade-off between maintaining privacy and providing smartness. A smart system can potentially be smarter with more data about the person requesting a smart service [77]. The trade-off decision should be under the control of the respective persons. They should be able to make the decision which data are provided and for which purpose and for how long these data should be accessible and available to the system afterwards. The challenge is now to find the right balance which – again – requires transparency about which data are really necessary for providing the smart service. This is also in line with the requirement of the EU-GDPR [23] being in effect since May 2018.

The problem was and still is (despite GDPR), that people are not asked for their permissions to collect and process their personal data on a specific basis and providing the rationale for doing so. The technology aspects of data collection and processing in AmI environments are discussed in more detail in section 5, when we describe how data science can inform the design of AmI environments.

Currently, people do not have the choice to decide and make the trade-off decision between smartness and privacy themselves but are confronted with serious privacy infringements [76,77].

Solving the problem is not trivial, especially when considering public spaces in urban environments, where different AmI environments with different properties and permissions might overlap due to undefined boundaries between them. This is partly also caused by the unobtrusive design of embedding sensors and actuators in the real/ hybrid environment, part of the original AmI vision and instantiated by the Disappearing Computer approach [73, 74, 80]. People are usually not aware of being monitored and which data are collected about them and their context. How can they take control, if they are not provided with the transparency and the options to opt-in or to opt-out to a specific data collection effort?

Of course, we cannot ignore the fact that people provide data also on a voluntary basis to all kind of institutions and service providers. Unfortunately, many of them are not aware of the implications when their personal data are stored on servers in a foreign country where only very weak or no legislation exists to protect their rights and their data. This shows again that much more information and education is necessary about the implications. People should have the freedom making their own decisions, but these should be informed decisions and people should be in control over their data. To remedy the difficult situation, a ‘privacy by design’, respectively ‘privacy by default’ approach is proposed [78, 79].

Finally, it should be mentioned that in real life situations the two trade-off conditions described before are confounded, because they coexist in parallel. Therefore, one must address combinations of trade-offs which makes things even more complex. This is elaborated in more detail in [79]. Nevertheless, we argue strongly that these trade-offs should be made accessible as part of a people-/citizen-centered design approach of current and future AmI environments.

4. Informing the Design and Evaluation of Ambient Intelligence Experiences

Once we have identified conceptual as well as pragmatic constraints and design trade-offs for the development of AmI environments, we are confronted with the multidimensional process of designing AmI expe-

riences. One of the main aspects involved in experiencing an AmI environment is the communicational aspect of interaction between the user and the mediated environment, which will be discussed in this section, ultimately aiming at identifying certain precedents for informing the design of these experiences. Underlying this discussion is the hypothesis that when users experience an AmI environment, they enter a bilateral communicational process with the distributed spatialized computational entity. For this purpose, when we design the communicational aspect of the interaction process, we may learn from interpersonal human-human communication and we may model the human-AmI communication accordingly, while taking into account the specific requirements and limitations of each entity's potential for communication.

The social aspect of interaction amongst users inhabiting a computationally mediated environment has been extensively researched [38, 39, 70]. This section will mainly focus on the communicational process between a single user and an AmI environment.

4.1. The Design of AmI Experiences as Spatial Communication Interfaces

“An ambience is defined as an atmosphere, or a surrounding influence: a tint.”

Brian Eno "Music for Airports"

In his liner notes for the “Music for Airports” album², Brian Eno defined the term “ambient music” as a form of “environmental music suited to a wide variety of moods and atmospheres”. He related the objective of creating such a musical composition to the creation of an environmental situation which induces certain emotional responses to the listeners experiencing this situation. The term “ambience” is an important aspect of the concept of “ambient intelligence”, which stresses the environmental character of these systems and the experiences they evoke.

The concept of ambience in AmI [1] implies that computation becomes non-obtrusively integrated into everyday objects and spaces. In a talk on *Everyware* in 2009, Greenfield³ asserted that in the case of an (appropriately designed) AmI experience, information processing colonizes the environment of everyday life

and the design of the experience “dissolves in behavior”. Computation is also perceived as dissolving in behavior as well as into the physical environmental context. Consequently, the physical environment becomes a medium for supporting interaction between the user and the AmI functionality, hence the interface becomes spatialized.

Mark [47] has acknowledged the significance of physical space as a characteristic of pervasive computing. Most conventional computer interfaces take neither physical space nor the presence or identity of human beings into account. However, as computation gradually becomes part of everyday physical space, the spatial context within which interaction between humans and computation takes place radically changes from a fairly static single-user, location-independent world to a dynamic multi-user situated environment. The physical location of the interface to the computation now becomes relevant. Networked pervasive computation is embedded in the environment as in the Disappearing Computer approach [80,81] and communicates multimodal content which dynamically changes as a result of user interaction. Thus, computation is moved from the center of our attention to the periphery, the area just outside focal attention [48] and added to static spatial elements, forming a coherent whole that offers an enhanced environmental experience [65]. The spatial context within which interaction takes place comprises of both computation and all physical environmental stimuli that may be involved in the process of interaction. This context is a spatial interface.

The term spatial interface⁴ characterizes human computer interfaces that utilize space as a context for supporting navigating within and interacting with information. Since humans use spatial organizing principles in their daily lives, they are used to and skillful in navigating space and communicating easily within space. It is therefore often appropriate to employ a spatial distribution of information as a means for organizing interaction with information and certain applications in a functional, well-structured and meaningful manner. Therefore, a physical environment enriched with pervasive and ubiquitous computation, an ambient intelligence, may be considered as a type of spatial

² Liner notes of Brian Eno's "Music for Airports, the initial American release of musical recording in CD format: Ambient 1", PVC 7908 (AMB 001), 1978.

³ Greenfield refers to Fukasawa's [26] concept of “Design dissolving in behavior” as an approach for conceptualizing AmI:

https://www.youtube.com/watch?v=_PKNbueOF5U&list=PL240CD0E5E91A9BA4

⁴ A spatial interface could be manifested to the user as immaterial (i.e. virtual environments), material (i.e. physical computing) or hybrid (i.e. augmented reality)

interface, a hybrid (physical and digital) spatial configuration, where computation expands in the physical space.

Humans utilize various modalities during direct human-human communication. The implementation of multiple modalities in HCI results in interfaces with reduced cognitive load [66]. Furthermore, the use of other senses besides vision may accelerate user adaptation [46]. Any environmental experience is multimodal. Apart from visual sensory input, all other sensory inputs like the perception of auditory, olfactory, thermal, and tactile input, and the sense of proprioception all contribute to the establishment of a sense of space [32]. This approach is also in agreement with Hall's conception of personal and social space [36]. Therefore, we suggest that the development and use of multimodal interfaces results in spatial interfaces affording an enriched and more complete spatial experience.

When considering the relation between the user, the computation, and the environment within which interaction takes place in the case of an AmI experience, we could suggest that pervasive and ubiquitous computation and linked media communicate information to the user in various modalities. This information escapes the representational context of the limiting two-dimensional space of a screen and is projected onto and manifested via the activity of the technological artifacts located in the physical environment. The activity of these artifacts shapes the context of human beings in a rather implicit manner which is not only attributed to human-machine communication but may also be attributed to machine-to-machine or environment-machine communication⁵.

Communication systems embody and integrate the functions of a communication interface, a series of transmission channels and an organizational infrastructure. Biocca and Delaney [11] define a communication interface as the interaction of physical media, codes, and information with the user's sensorimotor and perceptual systems. As suggested earlier, an important characteristic of the particular interfaces that this section deals with is their environmental character. Following Biocca and Delaney's definition, the spatial interface to an AmI environment could be considered a communication interface that engages the human sensorimotor channels into a vivid communication ex-

perience and that also affords an environmental experience [17]. Accordingly, we may use the term *spatial communication interface* to characterize the type of interface experienced by AmI users. Designing such a communication interface implies the design of the way in which interaction occurs among physical media, codes, and information on the one hand and the user's sensorimotor and perceptual systems on the other hand, as well as the appropriate environmental context, media displays, representations and other actuators, which function as a framework wherein this interaction occurs.

However, as earlier suggested, in AmI environments, computation is perceived as dissolving in behavior as well as into the physical environmental context. The computer "disappears" as a "visible" distinctive device, either physically due to being integrated in the environment or mentally from our perception [73, 74], thus providing the basis for establishing a calm technology as it was envisioned by Mark Weiser [88] and realized in multiple projects of the Disappearing Computer Initiative [80]. A main challenge is that "users" are often not fully aware of the interaction options provided in an AmI environment. A related implication is that they receive no feedback about wrong or inadequate user input or even system failures.

This new constellation requires a rethinking of the notion of "affordances" [31, 32, 54, 55, 56] in this new type of environments. Affordances were introduced by Gibson ([32] p.36) as the relationship, the set of possible actions, between an object of an environment and a living organism that may act upon this object ([56] p. 123). Norman [55] has appropriated and extended this concept to the world of design. He has stressed the significance of designing affordances which, when perceived, may inform the user of which actions can be performed on an object and how they may be performed by the user. When this design objective is successfully achieved, the designed artifact may communicate its purpose and functionality to the user.

Norman also suggests that media have special properties which may enhance and constrain their usage. A communication medium may not be a physical object, but it still has affordances ([56], pp.123-124). We may communicate the affordances of these media by appropriately designing the form of the media objects integrated within the AmI environment and the way in

⁵ Animals as living creatures are a part of the environment. Relevant experiments and artistic interventions involving animal-machine interaction were presented and discussed by Böhlen [13], Böhlen and Rinker [14] and Charitos and Theona [18].

which this form may be transformed over time via interaction with the user. When this design goes wrong, we may have a lack of information (hidden affordances) or wrong information (false affordances) [91] being communicated from the media object and perceived by potential users. Gaver [30] also stresses the fact that the perception of these affordances is partly determined by the observer's culture, social setting, experience and intentions.

In rethinking the new constellation in Aml environments, where users are often not provided with direct clues for interacting with the embedded, invisible computational devices, one must extend the notion of affordances. Streitz et al. [81] proposed the notion of "inherited affordances" for coping with such challenges in integrated smart environments, based on the design experiences with their interactive "Roomware[®]" environments [86].

4.2. On intelligence and non-verbal communication with/in the ambient intelligence

The concept of *intelligence* in Aml implies that the computational aspect of the environment supports some form of intelligent interaction. Intelligent behavior [1] involves four system elements: context awareness, personalization, adaptivity and anticipatory behavior, in which the Aml environment can extrapolate behavioral characteristics and generate pro-active responses. Additionally, this system intelligence must be compliant with societal conventions.

In order to achieve the above, the Aml experience has to somehow initiate and maintain bidirectional communication of meaning with the user. This could either be:

- *explicit communication*, via some kind of display (i.e. screen, framed surface or speaker, distributed in the environment) presenting verbal elements (text, static or moving images, sounds, symbols) and/or various types of representations (comprising abstract, iconic or symbolic content, communicated via visual, auditory or multimodal stimuli), on which the user usually focuses their attention.
- *implicit communication* through non-verbal elements which are presented at the periphery of the user's attention and perception.

This categorization of implicit and explicit communication elements is adopted by van de Ven et al. [87]. Mark [47] has also acknowledged the fact that pervasive computation is implicit.

Schmidt [69] has discussed "Implicit Human-Computer Interaction" (iHCI), where the user offers implicit input and receives implicit output. Implicit input refers to actions and behaviors of the user, which are not considered primarily as interaction-initiating, but are perceived as such by the system. Implicit output, similarly, refers to output, which occurs as a result of the reception and processing of implicit input. Implicit output is seamlessly integrated with the environment and supports the user's task. Essentially, the system detects subtle communicational cues inherent in the behavior of a human through the use of appropriate devices. After processing these data, the system reaches some conclusions about the user's state and the task to be accomplished and may subtly act on the environment towards increasing the possibility of the user successfully completing the task.

Kaptein et al. [42] used explicit measures of users' tendencies to comply with distinct persuasive strategies as well as implicit, behavioral measures of user traits for implementing persuasion profiling, as a method for personalizing the persuasive messages used by a system to influence its users.

From a communicational perspective, a person's experience of reality is altered by an additional layer of mediation that is placed between the user and the environment. This layer may have an impact on the users' conception of the computer and their behavior within such an enhanced environment [65]. Reeves & Nass [62] propose their theory of media equation, according to which humans interact with media technologies as if they were human. Computers are viewed as a social medium [23] and even as potential interlocutors or "social actors" [53]; humans tend to attribute to them abilities and traits they do not have (e.g. intelligence) and are willing to interact with them in the same way as they do with other humans. We could then suggest that the user may partly perceive the Aml experience as a process of communication with an artificial, human-like entity, where implicit communication may be mostly prevalent.

Rizopoulos [64] analyses the potential relation between non-verbal communication and spatial interaction in the context of spatial interfaces and suggests that: from a communicational perspective, iHCI is based on the perception of communication signals produced and transmitted by the human without her intention and which "reveal" her internal state. Implicit communication is largely embodied, since the body is closer to the unconscious and is more difficult to consciously control [91].

Non-verbal communication is often the way for providing input when the implicit HCI paradigm is

adopted [69]. This form of communication entails the information which is communicated through the user's perceptual channels, in a non-verbal manner. There are elements of non-verbal communication (prosodic) which relate to the verbal message [8]. Other elements of non-verbal communication are independent of the message: i.e. paralinguistic signals [8], which refer to the manner in which the message is communicated (i.e. tone, style and intensity of voice, speech). Non-verbal communication relates to the embodied aspect of communication and consists of three main categories: a) tacesics (the study of bodily touch between humans), b) proxemics (the interpersonal distances which are kept for negotiating our personal space and territories) [36] and kinesics (the analysis of the bodily movements and of the meanings related to them) [64]. Argyle [8] also explains that non-verbal elements of communication have the following functions: a) they express emotions, b) they communicate interpersonal attitudes, c) they accompany and support speech, d) they support self-presentation and e) they play a prominent role in rituals of social behavior. Indeed, some non-verbal signals stand for emotions, attitudes or experiences which are not easily expressible in words.

Recapitulating the main arguments presented in this section, we could suggest that a part of the communication of information between the user and an AmI experience is implicit. Users may perceive this experience as a process of communication with an artificial, human-like entity, but since this communication is partly implicit, it may escape the user's attention and although it may "reveal" her internal state, it may be based on the perception of communication signals produced and transmitted by the user without their intention. It should be clarified here that a part of the communication of information between the user and an AmI experience may also be explicit and this serves the functionalist objectives of interacting with the AmI experience to achieve an application task.

When implicit communication is adopted in an AmI experience, non-verbal communication is often the way for providing input for both the user and the system. We could then inform the process of designing implicit communication elements in an AmI experience by learning from the manner in which non-verbal communication signals are exchanged in social interaction amongst humans. It is necessary however to go through a systematic design research process of abstracting the ways in which non-verbal signals are communicated in human social interaction and adapting these ways to the specific characteristics of the

output devices through which this implicit communication signals will be presented in an AmI context. Devices which were or can be used for this purpose are: ambient displays [89], ambient light smart artefacts as, e.g., the Hello.Wall [60, 83], multisensory output devices, motors and other kinetic effectors, other material artifacts (possibly utilizing smart materials) the formal characteristics of which may be transformed via interaction and/or via the transformation of environmental parameters.

Of course, explicit, linguistic or representational elements may also be communicated within an AmI environment. McCullough [49] discusses various ways in which information may be embedded onto the elements of an environment or communicated to users via appropriate multimodal media displays: epigraphs, adhesive electronics (creating links between the digital and the physical context), cultural tagging, frames, screens, urban screens, etc.

As we conclude this section, we should also consider that "every course of action depends in essential ways upon its material and social circumstances" [84]. Humans often act on impulse and adapt to these circumstances, achieving intelligent action. Contexts in communication are not preset; rather, they are co-constructed by the participants. Communication should not be viewed as the process of information exchange, but as the process of the exchange of meanings and interpretations of the situations the actors are involved in [63].

5. The Role of Data Science for Informing Ambient Intelligence

As we saw already in our discussion of design trade-offs with respect to privacy and smartness (section 3), data play a key role for realizing AmI environments, where "massively distributed devices operate collectively while embedded in the environment using information and intelligence that is hidden in the interconnection network" [1]. Thus, it is necessary to take a closer look at the constraints and requirements of data collection, processing, analysis, exploitation and evaluation. In this section, Kaptein further explores topics from the field of *data science* that can contribute to the future development of AmI. Some of the topics discussed here historically originated in neighboring fields such as computer science, machine learning, artificial intelligence, and statistics (data structuring, uncertainty quantification, etc.), but the recent focus

on data science has highlighted (or sometimes reignited) our interest for these topics and has occasionally provided a different viewpoint. In this section, lessons are drawn that are important for AmI.

To realize AmI environments, we need a) *data* that describe the current state of the world to the devices that operate therein, b) *data processing*, either through explicit human-coded rules or more implicit, machine learned, relations, and c) *estimates of the outcomes* of the actions that AmI environments might take. Much of the application-oriented research work in AmI often does not explicitly focus on developing and evaluating methods for *data collection, processing, or estimation*. Rather, its focus is on the development and evaluation of novel applications, and on the user involvement and social responsibilities of such applications [1, 42]. In much of this work, machine intelligence is taken as a given; a useful assumption that has allowed the field to effectively study and reason about future emerging technologies and to involve users in the design of AmI applications even before these could be technically realized. However, it is worthwhile to explicitly evaluate the impact that research themes in data science have on our understanding of data collection, processing, and estimation as these will affect AmI environments. In this section, we pay specific attention to recent developments in the health domain: an application area that has received attention from both AmI and data science researchers. Specifically, in the health domain the use of data to make intelligent and user-centric decisions for the benefits of individuals is of large importance. Novel advances in data science are now shaping the ways in which data can be used to effectively personalize health-care decisions—where we take a broad view on health ranging from care to cure and hence including eHealth applications and health education programs. These advances provide meaningful directions for future AmI research.

The following themes have (re-)emerged in the study of data science and have potential impact on AmI research (and, potentially, scientific research as a whole):

- We need to *structure and organize* our data: Although more and more data is available, and the AmI vision gives rise to extremely large datasets, it remains a problem to effectively organize, combine, and disclose data such that it can effectively be utilized.
- We need to *embrace uncertainty*: given that we only have access to limited data, we will never be fully certain of our conclusions. While the

AmI community has largely relied on the existence of fixed rules for intelligent reasoning, current data science methods actively embrace the uncertainty that is inherent in data-driven decisions.

- We need to make *decisions sequentially*: Whatever actions we—or our technologies—take based on data will produce new data. This new data provides feedback regarding the utility of our decisions and the accuracy of our predictions. We need to actively close this feedback loop.
- We need to actively consider the *future value* of our collected data: The way we collect data will affect its future value: collecting data with insufficient descriptions of the state of the world and the data generating process often renders even extremely large datasets practically useless. The ways in which AmI devices interact with their environment should, at least partly, be driven by the future utility of these actions and the resulting data.
- We need to understand the *mechanisms* that generated our data: The AmI vision has always heavily relied on the existence of data. However, it is becoming more and more clear that a failure to understand the mechanisms that generated our data can lead to erroneous or biased reasoning in the future.
- We need to make *transparent* what drives our data-driven decisions: Finally, data-driven decisions affect the everyday lives of people, whether they are embedded in AmI technologies or not. Consistent with the AmI vision users should be able to understand how and why a certain decision was made.

5.1. Data Structure and Organization

It has been stated before that 80% of data science is effectively data cleaning. Despite the large potential value of all the data that is currently collected, it still proves hard to tie data together, transform it into usable formats, and disclose the data without infringing privacy or violating ethical norms. Health data provides a prime example: while the randomized controlled trial (RCT) is regarded as the pinnacle of evidence in the medical sciences [33], one would be inclined to believe that the estimates of effectiveness of different treatments that originate from RCTs can be further refined by looking at their effects “in the field”. Theoretically such data are readily available; hospitals store the treatment and outcome combination for each

disease for each patient, and often insurance companies will have direct records of the costs efficiency of treatments. However, combining data from hospitals, let alone merging the health outcomes with health care costs, has proven notoriously difficult. Our failure to easily combine data, to disclose data with privacy and security guarantees, and to analyze data resulting from multiple source currently limits the value of data.

AmI could make large contributions to this existing problem: as AmI devices continuously collect data in the field [6], AmI researchers and practitioner could actively contribute to creating standard for data sharing and merging. They could be on the fore-front of developing methods to deal with missing data, and—continuing its focus on contextual factors—could create standards that allow us to not only collect data of the primary processes in play (e.g., disease treatments and outcomes), but also the context in which the process played out. As of now however most AmI prototypes collect diverse types of data without a focus on standards for data sharing and data portability.

5.2. Embracing uncertainty

Even if we manage to share the data originating from diverse sources, we will still need to change our fundamental view of the world and the evidence data brings to the table. AmI applications often take intelligence—in the form of elaborate decision rules that depend on the current context and user—as a given when developing and evaluating applications. In reality, however, such deterministic rules, when derived from finite data, will always contain inferential errors. And, if we believe that the user and the context matter—a position that is strongly held in healthcare with its recent focus on personalized medicine [33]—the data that informs the actions of AmI technologies will inherently be very limited: we will never know with full certainty what the best action is for the current user in the current situation.

This has several consequences: first of all, we should actively model this uncertainty; despite a contemporary focus on point estimation in most of the data science literature effective methods for uncertainty quantification have been developed over the last decades and should not be ignored [59]. Second, if our uncertainty is too large to make a decision, we should inform the user, or perhaps actively illicit user input.

5.3. Learning sequentially

Not only should we embrace uncertainty and make it transparent to users—refraining from making decisions when the uncertainty is too large—we should also actively consider how our actions reduce future uncertainty. AmI technologies will always be imbedded in their environment and they can learn from actively interacting with their environment.

If an AmI technology is trying to make the best choice for a given user in a given context based on limited information, it is abstractly solving a decision problem called the contextual multi-armed bandit problem (MAB) [51]. The MAB problem is easily motivated in a health context: given two different treatments (or *actions*) $a = 1$ and $a = 2$, and given some outcome (or *reward*, r) dependent on these actions, we aim to maximize our outcome over sequential choices of our actions $\sum_{t=1}^T r_t$. If initially we do not know the outcomes of the different actions, we will need to explore the options and learn about the outcomes. However, as our knowledge of the outcomes give the actions grows, we would like to exploit and choose the action we believe serves the current patient best [50].

AmI technologies that actively learn from interacting with their environment will need to solve this *exploration-exploitation trade-off*: since for one specific user, in a specific context, no deterministic rule can be available, we need to balance making choices that inform our future choices, with utilizing the knowledge we currently have. Deterministic rules are asymptotically suboptimal to address this problem. Here embracing uncertainty is key: if we can properly quantify the uncertainty of the outcomes given our actions we can actively explore uncertain outcomes. A large literature that examines effective strategies, or *policies* to balance exploration and exploitation has emerged [see, e.g., 3], and AmI researcher should embrace this sequential learning view on the word: AmI technologies should actively seek information to inform their future decisions.

5.4. Considering the future value of data

Once we approach making intelligent decisions based on data as a sequential problem in which uncertainty is abound, we quickly encounter the following question: can we use the data that we collected using one specific policy to choose our actions to evaluate what would have happened if we had used another policy? This question is particularly relevant in a health-

care setting: does the data originating from a randomized clinical trial in which the patient population is assumed to be homogeneous (e.g., each unit has the same probability of receiving a treatment) allow us to evaluate an alternative policy that selects treatments based on user characteristics [44]?

Emerging answers to this question have direct implications for AmI technologies: it turns out that using data generated by a specific policy to evaluate “what if” questions is possible as long as the probabilities of receiving a treatment conditional on the user and context characteristics are known. This so-called propensity score can subsequently be used to counterbalance the effect of the policy that generated the data and allows us to obtain unbiased estimates of the performance of alternative policies [10]. Interestingly, such counterbalancing is impossible if probabilities are 0 or 1; this again highlights the importance of embracing uncertainty [90]. Policies that consist of deterministic decision rules generate data that is effectively useless for the evaluation of alternative strategies. AmI technologies, when interacting with their environment, should at the very least store the probabilities of the actions they took at each point in time to generate data that is valuable for re-use. Obviously, any data collection, storage and processing by AmI technologies must comply with the GDPR [23] introduced and discussed before in Section 3.1.

5.5. Understanding mechanisms that generate data

A very specific version of the “what if” question abound: we often have access to observational data—thus data that originated from a policy unknown to us—and we want to use it to evaluate alternative policies. For example, we might have data considering the outcomes of two treatments for a specific disease as administered in a hospital, and we want to know what the outcomes will be for future patients if we select one of the treatments.

This question is in general not solvable: there is no guarantee that the observational data originating “in the field” allow one to properly estimate the causal effect of the treatment. For example, a naive comparison of the survival rates for breast-cancer patients receiving chemotherapy or not based on observational data in the Netherlands would lead one to conclude that chemotherapy negatively affects survival rates. However, this conclusion is fully confounded by the severity of the tumor: only women with a severe tumor receive chemotherapy. If we had known the propensity scores in this case, we would have concluded that—

given the fact that these were 1 for those with severe tumors and 0 for those with mild tumors—the observational data was useless to evaluate another scheme of administering treatments.

Luckily, recent advances in our study of causal inference based on observation data have greatly improved our understanding of this problem. Effective methods to estimate propensity scores [27] and to uncover causal structures [57] now exists. AmI technologies that use existing data in their reasoning should actively incorporate these methods to prevent erroneous decisions. As AmI (and AI) applications are becoming more and more prominent in highly impactful areas such as healthcare (see also Section 7 of Gams et al. [29]) proper understanding of the (causal) mechanisms that generate our data is of increasing societal importance.

5.6. Make transparent decisions based on data

A final topic that has recently emerged in data science that should resonate with AmI researchers is the topic of transparency and fairness. As more and more decisions that affect individuals are made based on data, there is a growing need for methods that a) allow individual to understand why a decision was made, and b) control the feasibility of a decision in terms of fairness and avoid possible discrimination. Both areas are active areas of study: on the one hand, researchers are working actively on making decisions of black-box machine learning models transparent to their users [21, 45]. These methods should be incorporated into AmI technologies that autonomously make decisions that affect end-users. This work indirectly highlights a benefit of explicit, rule-based, processing methods: rule-based methods are often easy to understand for users. On the other hand, there is also a growing community of researchers that focusses on developing algorithmic and technical solutions to ensure fair and discrimination aware data science [e.g., 35]: these methods should also be embraced by AmI researchers.

As we have seen, contemporary developments in data science are likely to affect the AmI research field. Most notably of these is the recognition of the uncertainty contained in our data, and the sequential and interactive nature of data collection and decision-making. Whatever we do based on data is likely to generate new data, and this should affect our decisions. This view does not merely influence AmI: developments in data science and neighboring disciplines are currently challenging the use of RCT’s for the collection of

knowledge in the health sciences [44] and have transformed decision making in online marketing (for examples, see [3]). As our views regarding the utility and value of data are constantly changing, the AmI vision, in which interaction with complex environments based on continuously collected data is key, should embrace these changes. This is especially important since our abilities of monitoring the environment and context (see, e.g., Prati et al. [58] for an overview of the state-of-art) are rapidly increasing and hence our ability of making meaningful decisions based on sensor information is likely to improve strongly in years to come.

6. Ethically Aligned Design — Can AmI learn from mistakes?

In this section, Böhlen discusses a development of significance to AmI that is taking place largely outside of AmI, namely the crafting of guidelines for ethical design in autonomous, intelligent systems by IEEE, the world's largest professional engineering community.

Perhaps it is no coincidence that IEEE launches this initiative at the same time as global technology companies begin in earnest to question the motto of “moving fast and breaking things” in favor of a more measured approach to growth. Indeed, recent large-scale deployment of artificial intelligence into everyday products and services, several of which are discussed above, has made the governance of artificial intelligence a new priority. In the wake of these events, ethics of artificial intelligence are under (re)-evaluation.

The problem of ethics in autonomous, intelligent systems (A/IS) is significant to AmI because these new global standards will create new expectations towards AmI.

6.1. Managing ethics of autonomous systems

The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems⁶ (or Ethically Aligned Design - EAD) is an attempt to map out the territory of ethics in artificial intelligence, and to offer actionable recommendations to its constituents of engineers and software designers.

No doubt, there is a real need for, and public interest in, governing artificial intelligence [28]. From autonomous vehicles to virtual assistants, Autonomous and Intelligent Systems (A/IS) are impinging on every aspect of life along a multitude of vectors. It is no secret that engineers and designers who actually build artificial intelligences should consider ethics as a formal part of system development. The real problem is how to go about it in practice.⁷

EAD is both an initiative and a document under active construction and versioning⁸. EAD documents are a collaborative effort to which many people contribute. The stated goals of the initiative are twofold, first, to advance (and moderate) a public discussion on ethics in A/IS, and second, to create a standards and associated certification program that will enable and guide artificial intelligence research and development in practice. The current second version of the EAD document includes a comprehensive overview of relevant references from multiple perspectives on existing and new questions on ethics within artificial intelligence. The document covers the domains of affective computing, mixed reality, well-being, personal data, methodologies, safety, superintelligence, autonomous weapons and economics.

Ethical design is not news to the ambient intelligence community. Ethical concerns are inscribed into the conception of ambient intelligence research ab initio. For over a dozen years [12, 22, 41, 85], ambient intelligence research has addressed ethics of IT systems operating in everyday life. However, AmI has not been able to move ethical design from research contexts into ethical design in the wild, at scale, and has even been forced to see good intentions altered beyond recognition. The intelligent home compromised by its own data harvesting appliances is but one prominent case in point.

What the current EAD initiative offers beyond ambient intelligence's experimental contributions are implementation-oriented guidelines for considering ethics specifically of advanced A/IS for an entrepreneurial global context. With A/IS systems now operating in the wild at scale, this shift in scope is significant. The next subsection describes the EAD initiative in more detail.

⁶ <https://ethicsinaction.ieee.org/>

⁷ See this blog for clues to how a well-known technology company is struggling with defining ethical behavior in AI development: <https://blog.google/topics/ai/ai-principles/>

⁸ At the time of this writing the second version of the EAD document has been released and a third version is under development.

6.2. *From inspiration to recommendation*

The general principles the EAD initiative subscribes to are inspiring: human beneficence as a superset of human rights, the prioritization of benefits to humanity and the natural environment, and the mitigation of risks and negative impacts. With these principles, EAD proceeds across multiple sections to elaborate on the significance of AI ethics in the areas mentioned above.

The initiative clearly struggles to reconcile its ambitions with the vast territory it maps out. The approach the initiative takes is informed by its ultimate goal, namely making artificial intelligence and its application as A/IS manageable. With this ulterior motive in mind, each section of the document contains - as the examples below illustrate - practical recommendations that serve as the glue between the topic overviews and proposed implementations. The goal of the following paragraphs is to understand the logic of the initiative's argumentation through a critical reading of select sections of the current version of the document.

6.2.1 *Examples*

The section *Embedding Values into Autonomous Intelligent Systems* describes the problems designers encounter when attempting to create systems responsive to particular norms and values. The text stresses the importance of identifying norms and the circumstances in which they occur prior to implementing A/IS that operate within those norms. The evaluation of A/IS should, so the authors, continue from design to deployment and include procedures to resolve conflicting evaluation results (EADv2, p.50). After all, one person's helpful robot assistant might be another person's intrusive robot spy.

The section on *Methodologies to Guide Ethical Research and Design* calls for sustained interdisciplinary collaborations and the need to incentivize technical staff to voice ethical concerns throughout the product lifecycle (EADv2, p.62). The section on *Personal Data and Individual Access Control* deals with the challenge of organizing various dimensions of user data. It is well-established that the combination of personal data, technical metadata and inferences gleaned from data analytics create a high value digital footprint with high spatial and temporal granularity. EAD suggests granular-level consent at the time and point data is used (EADv2, p.106) across all data transactions as one

way to counter data misuse. Unfortunately, the promising concept of granular-level consent is not elaborated on in detail.

The section on *Policy* suggests that technology leaders and policy makers should work together to create A/IS systems, using internationally recognized human rights standards, non-discrimination and inclusiveness to assess the impact of an A/IS on individuals (EADv2, p.185). A possible framework towards such a collaborative effort is identified in various forms of exchanges between technologists and policy makers, including for example fellowships in which technologists spend time in political offices or policy makers join organizations at the intersection of engineering and advocacy (EADv2, p.186). Certainly, such exchanges are a good step. If only the document could elaborate on how the outcomes of such interactions would in practice flow into the crafting of ethically aligned design in A/IS.

Finally, the section on *Mixed Reality* is concerned with the various ways in which virtualization impacts personal identity, social interactions, privacy and mental health. The EAD authors foresee the potential for a new kind of social reclusiveness and a detachment from common reality to the point where avatars might redefine death (EADv2, p.222). The call for input from domain experts outside of artificial intelligence - such as mental health professionals - is repeated and that is good; but making good use of such expertise is left as an exercise for the reader.

6.3. *Recipes are not good enough*

Applied ethics generally concerns itself with concepts of good and bad conduct. A/IS, as engineering in general, considers the increase of efficiency as desirable in its own right. But efficient solutions need not be ethically sound solutions. Certainly the history of warfare offers copious examples to the point.

With the preferential positioning of efficiency and solutions-oriented methodologies, EAD risks skewing the interpretation of applied ethics in A/IS from a moral to a requirements management problem.

For example, the section on *Transparency* rightly points out that "transparency is important because it provides a simple way (for stakeholders) to understand what the system is doing and why" (EADv2, p.30). The corresponding recommendation then stresses the need to "develop new standards that describe measurable, testable levels of transparency". As an example of this testable transparency, the text mentions a care robot with a *why-did-you-do-that-button* one can activate to

have a robot explain an action it just performed. While this button might make getting a response from a robot easier, it certainly does not guarantee that the response is helpful. Imagine the robot indifferently stating it “did what is was programmed to do” when asked for an explanation. What is missing in this recommendation are the deeper dimensions of a transparent explanation such as context, and the ability to question the result delivered by the robot. EAD offers only a formal version of transparency, ‘transparency lite’, adequate maybe to satisfy legal requirements, but not even close to bona fide transparent action.

6.4. Take the long road

There is some help from other sources. The idea of algorithm impact auditing, part of ongoing efforts from media scholars [68], legal scholars [20, 71], and institutes [5] is a case in point. Algorithm impact auditing seeks to make algorithms accountable. Auditing includes the concept of disputability, allowing the public not only to see what an algorithm is doing but to dispute its outcome. Auditing also implicitly considers the effects of code in the real world, including pathologies of scaling. Evaluating A/IS in the lab on small sample data is not the same thing as running A/IS in the messy world on data from millions of people⁹. Side effects are much more likely to occur in complex environments, and much harder to counter with optimization approaches. Expanding the reach of algorithm control to the level of accountability [16] is important, legally and politically, as enforceable action is only available from large-scale structures charged with upholding the interests of the public. In this regard, the European Union’s General Data Protection Regulation [23] and its formulation of enforceable, individual rights is an important attempt to apply policy level intervention at least on personal data.

A/IS are complex socio-technical systems composing computers, sensors, data, databases, multi-author algorithms with various levels of autonomy running continuously in remote locations. Then: time constraints, patches not applied, deadlines looming, people under stress in the workplace, etc., etc. These (and many other) intertwined factors contribute to how an A/IS behaves – and fails – in the real world.

Indeed, the term failure hardly catches the many dimensions along which outcomes can deviate from expectation. Even the space of technical failures is vast. Learning systems can go bad simply because of improperly formulated goals, fragile or miss-specified objective functions [7] which an algorithm might try (and succeed) to optimize. Making A/IS safe, let alone ethically aligned is a complex undertaking into which research is far from complete.

To be clear: The issue is not only that ethically aligned algorithms in A/IS are not ready for mass production, but that the scope of the challenge itself has not been adequately established. Ethically aligned interventions require more than a purportedly ethically aligned algorithm. The recent “success” of a state of the art school bus routing algorithm designed to increase equity for Boston public school students¹⁰ is a case in point. The algorithm met its goals of reconfiguring bus start times and cutting transportation costs, but created a political disaster as affluent sections of the city lost established advantages and protested against the algorithmically proposed changes.

Instead of suggesting broad recommendations at this point in time, EAD could call for and support open research into ethics of A/IS in the wild. For example, EAD could suggest and coordinate experiments by which to test and evaluate ethically aligned design concepts, much like it has been proposed for A/IS safety [7]. To these experiments, monitored and evaluated by multi-disciplinary teams, one could add continuously updated field reports from A/IS failures in the real world, operating at scale. Together these sources could constitute a compendium on safety, failures and ethically aligned experiments in A/IS. Importantly, such a compendium should be publicly available. Bringing known problems within A/IS from closed board rooms into public view is one way to increase trust in A/IS, a central concern of the EAD initiative.

At least in the autonomous vehicle industry recent accidents [52] have increased the pressure for safer solutions, making autonomous vehicles a good candidate for the approach outlined above. Safety, as opposed to ethics, is directly linked to economic imperatives. Liability lawsuits in response to lax safety provisions may just be the most effective entry point into making harm-adverse A/IS a reality. Then, with safer systems

⁹ Robustness to distributional shift considers part of this problem. See Amodei [7], p.16ff.

¹⁰ David Scharfenberg. Computers can solve your problems. But you might not like the answer. What happened when Boston Public Schools tried for equity with an algorithm. The Boston Globe. September 21, 2018.

under development, the most effective tested safe algorithms could be used as basis for ethically sensitized algorithms. This would give the engineering community a robust departure point from which to consider the hard non-engineering elements of A/IS ethically aligned design, namely policies and politics, business and global culture.

A cautious iterative approach is not a Luddite retreat. Rather it suggests carefully building a path towards ethical design in AI/S while one can still afford to make mistakes. After all, the currently deployed A/IS are proficient mostly in specific domains. But the coming realm of Artificial General Intelligence, “an intellect that is much smarter than the best human brains in practically every field” [15] might be much less forgiving. Casting Artificial General Intelligence as an extension of A/IS, and not an unrelated alien creature, is a good move on the part of EAD; it allows one to take control of the development of Artificial General Intelligence. In principle, at least.

One candidate recommendation offered for countering malicious Artificial General Intelligence is the safe-by-design (EADv2, p.79) approach. As above, details on precisely how safe-by-design systems might operate in critical situations are missing. For example, how would safe-by-design prevent an armed autonomous drone from optimizing a reward function of minimizing public disturbance by simply picking the most effective action, killing protesters, even though the drone was never explicitly programmed to do so? Likewise, the statement “teams working on developing Artificial General Intelligence should be prepared to put significantly more effort into AI safety research as capabilities grow” (EADv2, p.77) offers little help and even less solace. Instead of debating the merits and drawbacks of Arkin’s ethical governor [9], for example, the recommendations prefer uplifting messages, to wit: “Adopt the stance that superintelligence should be developed only for the benefit of all of humanity” (EADv2, p.82).

As the historian Yuval Harari reminded his audience at the World Economic Forum 2018, it took societies millennia to learn how to organize something as simple as the ownership of land [37] through an evolving set of concepts on contracts, fences and city walls etc. How can one expect to robustly organize the ownership of endless global data streams, let alone the super intelligence that will process and learn from them so quickly? It is too early to craft recipes. A long view is required. No one can afford to be guided by naïve hopes; no one can afford not to learn from past mistakes. Ambient intelligence was early to the game of ethically aligned design but did not succeed in

bringing the concept to industrial scale. The intelligent home is in danger of been compromised by its own data harvesting appliances; the smart city a victim of greedy data collection marketing business models whose motto Bruce Sterling aptly described as “information about you wants to be free to us” [72].

The time is ripe for ethically aligned design; done carefully without the shortcuts the EAD initiative proposes. Whichever rules of AI management will be agreed upon next should be understood as provisional. It is important to anticipate mistakes and remain adaptive; more adaptive yet than the new superintelligences under construction.

7. Conclusions and Outlook

In the preceding sections, we described a selection of challenges and their implications for design contexts and implementations of Aml environments and finally also society. Although they were described in individual sections, there are strong correlations and interactions between them, forming a comprehensive picture of the challenges society is confronted with. To make these interactions and dependencies concrete, we provide first an example of the application of our predictions and recommendations in the domain of future urban environments and then broaden the scope in our claims for future developments.

7.1. Beyond ‘smart-only’ cities and societies

While the analyses and recommendations have general applicability, it is useful to apply them to the domain of current and future urban Aml-environments. Currently, one can observe an increasing hype indicated by the label ‘smart cities’. Sterling [72] even asks for “Stop saying smart cities”. As shown in several examples in the preceding sections, there is a need to move beyond ‘smart-only’ cities by putting a different set of requirements and design goals in the first place. One could use a rephrasing of smart: “smart, but only if cooperative and humane”. In accordance with the design trade-offs mentioned before, the overall goal of designing and realizing future or refurbishing existing cities should be to build humane, sociable and cooperative hybrid cities reconciling people and technology by providing a balance between human control and automation as well as privacy and smartness [75, 76, 79]. This implies that we need to foster and enable the following actions and requirements for designing

and building AmI applications in the context of smart urban environments:

- Establishing a calm technology providing ambient intelligence that supports and respects individual and social life by “keeping the human in the loop and in control”. This includes a transparent dealing with data and a clear knowledge of the limitations of the processing methods used.
- Respecting the rights of citizens, especially in terms of privacy and security. Therefore, personal data should – as much as possible – only be collected based on consent by providing choices and control of the process, including models of temporary provision and access and/or obligations to delete data later. The GDPR regulations issued by the European Union provide a good basis. But we are also aware that the introduction (and perpetual updating) of such a legal framework is a process that evolves at a slower pace than the implementation and embedment of AmI systems in our everyday environments.
- Educating citizens about data acquisition and management. This may enhance their awareness and consequently aid them towards making more conscious decisions about how to manage their own data, in everyday life situations. This can only be a positive move towards protecting citizens from fundamental civil rights’ violations, by states and/or private parties, and ultimately emancipating them with regards to using AmI environments, as techno-social systems mediating everyday life. In sociopolitical terms, ownership of citizen’s data means power. The material implications of this are becoming visible, but there seem to be no simple answers to this issue.
- Viewing the mediated city and its citizens as mutual cooperation partners, where a city is ‘smart’ in the sense of being ‘self-aware’ and ‘cooperative’ towards its citizens by supporting them in their activities. This requires mutual trust and respect for the motives and vested interests of all stakeholders involved.
- Acknowledging the capabilities of citizens to participate in the design of the urban environment and how these systems of technological mediation are embedded into the urban context, especially with respect to their local expertise, and stimulating their active participation (=> participatory design).
- Motivating citizens to get involved, to understand themselves as part of the urban community, to be actively engaged by contributing to the public

good and welfare (=> collective intelligence).

This implies the provision of techno-social systems that may support bottom-up creative, participatory, co-operative processes for appropriating the technologically mediated city experience.

- Enabling citizens to exploit their individual, creative, social and economic potential and to live a self-determined life, and thus
- Meeting some of the challenges of the urban age by enabling people to experience and enjoy a satisfying life and work.

This list of actions and requirements applied to future urban environments point to a promising prospect, but only if they are taken into account and affect the manner in which these systems are structured and realized in the implementations. One must be aware that there are severe risks caused by different goals and value systems of the different stakeholders in our society - requiring a discussion of pros and cons - on the way to a humane and cooperative smart urban society. Therefore it is important that the proliferation of AmI systems in the urban and social realm is proactively evaluated by a meticulous and adaptable approach at the level of policy making and governed by an appropriate legal framework that will safeguard these policies.

7.2. Claims for future developments

While the application domains of cities and urban societies will play an increasingly important role due to living in the urban age, we can also abstract and formulate some claims in a more general fashion. So, we anticipate the following developments:

- AmI is 20 years old. It might not survive another 20 years. As its sibling UbiComp, it might fall prey to a change in fashion as the Internet of Things movement has shown. But AmI’s early focus on developing and deploying technology based on a human-oriented and social responsible approach to increase life quality is timeless.
- We can observe, the more the computer disappears and becomes “invisible” in smart AmI environments, the more it determines our lives. The world around us is the “interface” and provides a rich bouquet of offerings and services – some that we need and want, some that are offered unsolicited without our approval.
- What kind of next-generation “interfaces” will be able to communicate intuitively a new dimension

of complexities to people? How must future affordances for interaction and communication be designed in order to cope with smart materials constituting the AmI environments? Machine learning and its opaque internal operations will make new forms of interfaces necessary. Text and image might become less and less relevant or even quaint objects of past, although there is also the position that their semiotic value is indisputable, and gestures and speech alone will not be the solution to the intricate issues we are confronted with.

- There is an eminent need to redefine the ‘smart-everything’ paradigm to avoid that people are losing control and are at the mercy of non-transparent, error-prone and rigid and at some point, even autonomous algorithms. Efforts and appropriate design trade-offs are needed to prioritize “people-empowering smartness” and control over autonomous automation so that “smart spaces make people smarter”.
For example, one could imagine a new class of algorithms that recognize when their actions might have adverse effects and actively seek council with human beings. We will have to design machines that want to share with us as we are asked to share with them. No doubt, this will lead to new complications. If an AmI system can help a neighbor by sharing her/his fire alarm data with me, it will violate privacy protocols but save the house.
- Data will increasingly be collected and processed by private companies and public/state institutions, often with dubious justifications and for inappropriate usage scenarios. In commercial contexts, privacy will become a commodity and thus a privilege unless we do something against this trend. Assuring privacy by supporting and/or demanding an appropriate design approach (‘privacy by design and by default’) combined with supportive legislation and regulations (as, e.g., the EU-GDPR), could result in a USP for companies meeting the concerns of privacy-aware customers and a benefit for all citizens.
- As AmI environments react to, and shape, their surroundings, we risk the introduction of biases (or self-fulfilling prophecies) in the data used to fuel the system intelligence. So, there is a growing need to understand the limits of machine intelligence. In order to design for such situations, we will have to revisit old assumptions. What

kind of data streams do we really need? Will sequential data allow for really accurate predictive actions?

- Automated decision making based on vast amounts of data will be ubiquitous. Therefore, we need to understand data collection, processing, prediction, and exploitation much better and to integrate an inherent way of providing transparency.
- Transparency and traceability of intelligent systems and their algorithms is already and will be even more in the future a recurring theme at different levels and a wide range of implications. Providing transparency has the potential of being a relevant condition for acceptance by people in their roles as users and citizens. Like privacy, it can be a USP for companies and a benefit for society at large. Thus, the need for transparency must be addressed and AmI objectives and methods can play a constructive role here. Thus, people-oriented design is needed for “keeping people in the loop and in control”, being transformed into citizen-centered design when applied to cities.
- Ethically aligned design within AmI must make daily life better and more just. Ethically aligned design can only become meaningful for society when designed and implemented to improve life quality for many people. For example: What happens with the savings produced by the smart city’s efficient energy systems? Support other cities and populations more vulnerable to climate change dynamics? If AmI is to be relevant, it must consider the larger economic and political dimensions of technical design.
- Artificial General Intelligence (AGI) will change the rules of engagement between people and computers much more radically than previous computing advances. It will impact every applied computing field, AmI included. Finding creative solutions to managing AGI might be key to surviving (and then living well) with systems superior to ourselves. Fear abounds. Maybe a future community of Mars dwellers living under the harshest of conditions will volunteer to subject themselves to the AGI systems?

Ambient Intelligence values, objectives and methods can play a major role in achieving the goal of reconciling people and technology in a future ‘smart’ society, hopefully a beyond ‘smart-only’ society. An im-

portant aspect, but no guarantee, is the actual implementation of the design trade-offs and the ethical considerations described and discussed before. But we must keep in mind, that the AmI approach is only one perspective and not at all a comprehensive solution for all the problems cities and society are facing today and will be so even more in the future. Beyond the role of AmI-based technologies, there is a wide range of important issues, including socio-economic, ecological, sustainability and political aspects.

References

- [1] E. Aarts, B. de Ruyter (2009). New research perspectives on Ambient Intelligence. *Journal of Ambient Intelligence and Smart Environments*, 1(1), pp. 5–14.
- [2] E. Aarts, J. Encarnação (Eds.) (2006). *True Visions: The Emergence of Ambient Intelligence*. Springer-Verlag
- [3] A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, R. E. Schapire (2014). Taming the Monster: A Fast and Simple Algorithm for Contextual Bandits. Proceedings of the 31st International Conference on Machine Learning.
- [4] AI NOW Institute (2017). https://ainowinstitute.org/AI_Now_2017_Report.pdf
- [5] AI NOW Institute (2018). Algorithmic Impact Assessments: Toward Accountable Automation in Public Agencies.
- [6] M. Alirezaie, A. Loutfi (2015). Reasoning for sensor data interpretation: An application to air quality monitoring. *Journal of Ambient Intelligence and Smart Environments*, 7(4), pp. 579-597.
- [7] D. Amodi, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, D. Mane (2016). Concrete Problems in AI Safety.
- [8] M. Argyle. *Bodily Communication* (2nd edition). London and New York: Routledge. 1988
- [9] R. Arkin (2008). Governing Lethal Behavior: Embedding Ethics in a Hybrid Deliberative/ Reactive Robot Architecture. Proceedings of the 2008 Conference on Human-Robot Interaction, pp. 121–128.
- [10] A. Bietti, A. Agarwal, J. Langford (2018). Practical Evaluation and Optimization of Contextual Bandit Algorithms. pp. 1–28.
- [11] F. Biocca, B. Delaney. Immersive virtual reality technology. In F. Biocca & M. R. Levy (Eds.), *Communication in the age of virtual reality* (pp. 57–124). Hillsdale, NJ: Lawrence Erlbaum Associates. 1995
- [12] J. Bohn, V. Coroamă, M. Langheinrich, F. Mattern, M. Rohs (2004). Living in a World of Smart Everyday Objects – Social, Economic, and Ethical Implications. *Human and Ecological Risk Assessment*, Vol. 10, No. 5.
- [13] M. Böhlen (1999). A robot in a cage - exploring interactions between animals and robots. Proceedings 1999 IEEE International Symposium on Computational Intelligence in Robotics and Automation. CIRA'99 (Cat. No.99EX375)
- [14] M. Böhlen, J.T. Rinker (2005). Experiments with Whistling Machines. *Leonardo Music Journal*, 15: 45-52. MIT Press.
- [15] N. Bostrom (2003). Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence, Vol. 2, ed. I. Smit et al., Int. Institute of Advanced Studies in Systems Research and Cybernetics, pp. 12-17.
- [16] danah boyd (2016). Transparency! = Accountability. EU Parliament Event. Algorithmic Accountability and Transparency. Brussels, November 7.
- [17] D. Charitos. Precedents for the design of Locative Media. In: Saariluoma, P. & Isomäki, H. (eds.) “Future Interaction Design II”, London: Springer Verlag. 2008
- [18] D. Charitos, I. Theona (2015). Placemaking by mediated urban spatial experiences in the era of the Internet of Things. Proceedings of the 5th Media City International Conference. University of Plymouth.
- [19] J. Chin, V. Callaghan, S.B. Allouch (2019). The Internet of Things: Reflections on the Past, Present and Future from a User Centered and Smart Environments Perspective. Tenth Anniversary Issue, *Journal of Ambient Intelligence and Smart Environments* (11) 1. IOS Press.
- [20] N. Diakopoulos, S. Friedler (2016). How to Hold Algorithms Accountable. *MIT Technology Review*. Nov 17.
- [21] C. Dimitrakakis (2009). Complexity of stochastic branch and bound methods for belief tree search in Bayesian reinforcement learning arXiv : 0912 . 5029v1,
- [22] P. Duquenoey (2004). Intelligent ethics. In: Building the information society: IFIP 18th World Computer Congress. Toulouse, France. R. Jacquart (Ed). International Federation for Information Processing (156). Kluwer Academic Publishers, Boston, pp. 597-602.
- [23] European General Data Protection Regulation (GDPR) (EU) 2016/679.
- [24] EU-GDPR (2016). <http://www.eugdpr.org/> and more specific <http://www.privacy-regulation.eu/en/13.htm> and <https://www.privacy-regulation.eu/en/22.htm> (last check May 2018)
- [25] K. Eykholt, I. Evtimov, E. Fernandes, B. Li, A. Rahmati, C. Xiao, A. Prakash, T. Kohno, D. Song (2018). Robust Physical-World Attacks on Deep Learning Models. Cornell University Library <https://arxiv.org/abs/1707.08945> (accessed June 2018).
- [26] N. Fukasawa (2005). Design dissolving in behavior. Keynote speech at Seventh International Conference on Ubiquitous Computing (UbiComp 2005), Tokyo, Japan. ubicomp.org/ubicomp2005/programs/pdfs/keynote-2.pdf
- [27] M. J. Funk, D. Westreich, C. Wiesen, T. Stürmer, M. A. Brookhart, M. Davidian. (2011). Doubly robust estimation of causal effects. *American Journal of Epidemiology*, 173(7):761–767.
- [28] Future of Life Institute (2015, ongoing). Research Priorities for Robust and Beneficial Artificial Intelligence: An Open Letter. <https://futureoflife.org/ai-open-letter/?cn-reloaded=1>
- [29] M. Gams, I. Yu-Hua Gu, A. Härmä, A. Muñoz, V. Tam (2019). Artificial Intelligence and Ambient Intelligence. Tenth Anniversary Issue, *Journal of Ambient Intelligence and Smart Environments* (11) 1. IOS Press.

Acknowledgements

The authors like to thank the anonymous reviewers for their feedback and suggestions and Boris de Ruyter for his supporting role in the initial phase of this paper and for his detailed comments on the submitted paper.

- [30] W. Gaver (1991). Technology affordances. Proceedings of CHI'91, (New Orleans, Louisiana). ACM Press. pp. 79-84.
- [31] J.J. Gibson (1966). The senses considered as perceptual systems. Allen and Unwin, London.
- [32] J. J. Gibson. The Ecological Approach to Visual Perception, Lawrence Erlbaum Associates, London, 1986 (first published in 1979).
- [33] J. J. Goldberger, A. E. Buxton (2013). Personalized medicine vs guideline-based medicine. *JAMA: the Journal of the American Medical Association*, 309(24):2559–60.
- [34] C. Gomez, S. Chessa, A. Fleury, G. Roussos, D. Preuvenciers (2019). Internet of Things for Enabling Smart Environments: A Technology-Centric Perspective. Tenth Anniversary Issue, *Journal of Ambient Intelligence and Smart Environments* (11) 1. IOS Press.
- [35] S. Hajian, J. Domingo-Ferrer, A. Monreale, D. Pedreschi, F. Giannotti. Discrimination- and privacy-aware patterns (2015) *Data Mining and Knowledge Discovery*, 29(6):1733–1782.
- [36] E. T. Hall (1966). The Hidden Dimension. Anchor Books.
- [37] Y. N. Harari (2018). An Algorithm Will Be Your Best Therapist, But It Can Be Hacked Too. WEF Davos.
- [38] S. Harrison (2009). Media Space: 20+ Years of Mediated Life, Springer-Verlag: London.
- [39] K. Hook, D. Benyon. A.J. Munro (2003). Designing Information Spaces: The Social Navigation Approach, Computer Supported Cooperative Work series, Springer-Verlag: London.
- [40] M. Hutson (2018). Missing data hinder replication of artificial intelligence studies. Science (15. February 2018). Report about AAAI Conference, New Orleans. 2018. <https://www.sciencemag.org/news/2018/02/missing-data-hinder-replication-artificial-intelligence-studies> (last check February 2018)
- [41] S. Jones, S. Hara, J. C. Augusto (2014). eFRIEND: An ethical framework for intelligent environments development. *Ethics and Information Technology*, 17 (1). Springer Verlag, pp. 11-27.
- [42] M. C. Kaptein, P. Markopoulos, B. de Ruyter, E. Aarts (2009). Persuasion in ambient intelligence. *Journal of Ambient Intelligence and Humanized Computing*, 1(1):43–56.
- [43] M. Kaptein, P. Markopoulos, B. de Ruyter, E. Aarts (2015). Personalizing persuasive technologies: Explicit and implicit personalization using persuasion profiles. *International Journal of Human-Computer Studies*, 77, 38-51
- [44] M. C. Kaptein (2018). Computational Personalization. Data science methods for personalized health. Technical report, Inaugural address at the University of Tilburg.
- [45] H. Kerdegari, S. Mokaram, K. Samsudin, A.R. Ramli (2015). A pervasive neural network-based fall detection system on smart phone. *Journal of Ambient Intelligence and Smart Environments*, 7(2), pp. 221-230.
- [46] S. Kuivakari, S. Kangas (2005). Pleasure Platforms and Sensorimotoric Interfaces: Notes from a Preliminary Survey of Adaptive User Interface Design. In: M. Ylä-Kotola, S. Inkinen, H. Isomäki (Eds), The Integrated Media Machine: Aspects of Future Interfaces and Cross-Media Culture, University of Lapland, Rovaniemi, pp. 77-99.
- [47] W. Mark (1999), Turning pervasive computing into mediated spaces, *IBM Systems Journal*, Vol 38, No4, 677-678.
- [48] M. McCullough (2004), Digital Ground: Architecture, Pervasive Computing, and Environmental Knowing, MIT Press, Cambridge MA, p. 49.
- [49] M. McCullough (2013). Ambient Commons: Attention in the Age of Embodied Information. MIT Press. Cambridge, MA.
- [50] W. G. Macready, D. H. Wolpert (1998). Bandit problems and the exploration/exploitation tradeoff. *IEEE Transactions on Evolutionary Computation*, 2(1):2–22, 1998.
- [51] R.E. McInerney, S.J. Roberts, I. Rezek (2010). Sequential Bayesian Decision Making for Multi-Armed Bandit. Sequential Decision Making.
- [52] MIT Technology Review (2018). What Uber’s fatal accident could mean for the autonomous-car industry. March 19.
- [53] A. Nijholt (2004). Where Computers Disappear, Virtual Humans Appear. *Computers and Graphics*, 28(4), pp. 467-476.
- [54] D.A. Norman (1988). The Psychology of Everyday Things (POET). Basic Books, New York. (revised version published as The Design of Everyday Things.)
- [55] D.A. Norman (1999). Affordance, Conventions and Design. *Interactions* 6(3):38-43, May 1999, ACM Press.
- [56] D.A. Norman (1999) The Invisible Computer, MIT Press, Cambridge: Massachusetts.
- [57] J. Pearl (2003). Statistics and Causal Inference: A Review. 1(2),
- [58] A. Prati, C. Shan, K. Wang (2019). Sensors, Vision and Networks: from video surveillance to activity recognition and health monitoring. Tenth Anniversary Issue, *Journal of Ambient Intelligence and Smart Environments* (11) 1. IOS Press.
- [59] M. T. Pratola et al. (2016). Efficient Metropolis–Hastings proposal mechanisms for Bayesian regression tree models. *Bayesian Analysis*, 11(3):885–911.
- [60] T. Prante, R. Stenzel, C. Röcker, N. Streitz, C. Magerkurth (2004). Ambient Agoras: InfoRiver, SIAM, Hello.Wall. In: CHI'04 Extended Abstracts on Human Factors in Computing Systems, pp. 763-764.
- [61] C. Price (1966). Technology is the answer, but what was the question? Title of Lecture.
- [62] B. Reeves, C. Nass (2002). The Media Equation: How people treat computers, television and new media like real people and places. Centre for the Study of Language and Information Publications.
- [63] G. Riva, C. Galimberti (2001). Virtual Communication: Social Interaction and Identity in an Electronic Environment. In: G. Riva and F. Davide (Eds), Communication Through Virtual Technology: Identity, Community and Technology in the Internet Age. IOS Press, Amsterdam, pp. 23-46.
- [64] C. Rizopoulos (2017). Cognitive and environmental factors of communication between humans and intelligent environments, PhD Thesis, accepted at the Department of Communication and Media Studies, National and Kapodistrian University of Athens.
- [65] C. Rizopoulos, D. Charitos (2007). How do we communicate with(in) intelligent spaces? Convivio Webzine, issue 6, 2007. URL: [http://daisy.cti.gr/webzine/Issues/Issue%206/Articles/How%20do%20we%20communicate%20with\(in\)%20intelligent%20spaces/index.html](http://daisy.cti.gr/webzine/Issues/Issue%206/Articles/How%20do%20we%20communicate%20with(in)%20intelligent%20spaces/index.html)
- [66] D.M. Russell, N.A. Streitz, T. Winograd (2005). Building Disappearing Computers. *Communications of the ACM*, 48(3), pp. 42-48.
- [67] SAE (2014). Levels of Driving Automation defined by the new SAE Standard J3016. https://www.sae.org/misc/pdfs/automated_driving.pdf Revised in 2016: https://www.sae.org/standards/content/j3016_201609/
- [68] C. Sandvig, K. Hamilton, K. Karahalios, C. Langbort (2014). Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms. 64th Annual Meeting of the International Communication Association. May 22, Seattle.

- [69] A. Schmidt (2005). Interactive Context-Aware Systems Interacting with Ambient Intelligence. In G. Riva, F. Vatalaro, F. Davide and M. Alcañiz (Eds.), *Ambient Intelligence: The Evolution of Technology, Communication and Cognition Towards the Future of Human-Computer Interaction*, IOS Press, Amsterdam, pp. 159-178.
- [70] R. Schroeder (Ed.) (2002). *The Social Life of Avatars: Presence and Interaction in Shared Virtual Environments*. London: Springer Verlag
- [71] A. Selbst (2017). Disparate Impact in Big Data Policing. 52 *Georgia Law Review* 109.
- [72] B. Sterling. (2018). Stop Saying ‘Smart Cities’. *The Atlantic*, February 12.
- [73] N. Streitz (2001). Augmented Reality and the Disappearing Computer. In: M. Smith, G. Salvendy, D. Harris, R. Koubek (Eds.), *Cognitive Engineering, Intelligent Agents and Virtual Reality*. Lawrence Erlbaum. pp. 738-742.
- [74] N. Streitz (2008). The Disappearing Computer. In: T. Erickson, D.W. McDonald (Eds.), *HCI Remixed: Reflections on Works that have Influenced the HCI Community*. MIT Press, pp 55 – 60.
- [75] N. Streitz (2011). Smart Cities, Ambient Intelligence and Universal Access. In: C. Stephanidis (Ed.) *Universal Access in Human-Computer Interaction*. Lecture Notes in Computer Science (LNCS 6767). Springer. pp. 425–432.
- [76] N. Streitz (2015). Citizen-Centered Design for Humane and Sociable Hybrid Cities. In: I. Theona, D. Charitos (Eds.), *Hybrid City 2015 - Data to the People*. University of Athens, Greece. pp. 17-20.
- [77] N. Streitz (2016). Smart Cities Need Privacy by Design for Being Humane. In: S. Pop, T. Tofit, N. Calvillo, M. Wright (Eds.), *What Urban Media Art Can Do - Why When Where and How?* Verlag avedition. pp. 268 – 274.
- [78] N. Streitz (2017). Reconciling Humans and Technology: The Role of Ambient Intelligence. Keynote Paper. In: A. Braun, R. Wichert, A. Mana (Eds.), *Proceedings of the 2017 European Conference on Ambient Intelligence*. Lecture Notes in Computer Science (LNCS 10217), Springer (pp. 1-16). https://doi.org/10.1007/978-3-319-56997-0_1
- [79] N. Streitz (2018). Beyond 'Smart-Only' Cities: Redefining the 'Smart-Everything' Paradigm. *Journal on Ambient Intelligence and Humanized Computing*. Springer. <https://doi.org/10.1007/s12652-018-0824-1>
- [80] N. Streitz, A. Kameas, I. Mavrommati (Eds.) (2007), *The Disappearing Computer: Interaction Design, System Infrastructures and Applications for Smart Environments*. Springer-Verlag, LNCS 4500.
- [81] N. Streitz, T. Prante, C. Röcker, D. van Alphen, R. Stenzel, C. Magerkurth, S. Lahlou, V. Nosulenko, F. Jegou, F. Sonder, D. Plewe (2007). Smart artefacts as affordances for awareness in distributed teams. In N. Streitz, A. Kameas, I. Mavrommati (Eds.), *The Disappearing Computer*. Springer-Verlag, LNCS 4500. (pp. 3-29).
- [82] N. Streitz, G. Privat (2009). Ambient Intelligence. Final section "Looking to the Future". In: C. Stephanidis (Ed.), *The Universal Access Handbook*. CRC Press. pp. 60.1 - 60.17.
- [83] N. Streitz, C. Röcker, T. Prante, D. van Alphen, R. Stenzel, C. Magerkurth (2005). Designing smart artifacts for smart environments. *IEEE Computer*, 38 (3), pp. 41-49
- [84] L.A. Suchman. *Plans and Situated Actions: The Problem of Human Machine Communication*, Cambridge University Press, Cambridge. 1987.
- [85] SWAMI – Safeguards in a World of Ambient Intelligence. 2005. European Union. Sixth Framework Project.
- [86] P. Tandler, N. Streitz, T. Prante (2002). Roomware - Moving Toward Ubiquitous Computers. *IEEE Micro*, 22 (6) pp. 36-47.
- [87] J. van de Ven, D. Anastasiou, F. Dylla, S. Boil, C. Freksa (2015). The SOCIAL project: Approaching Spontaneous Communication in Distributed Work Groups. In: B. de Ruyter, A. Kameas, P. Chatzimisios & I. Mavrommati (Eds.) *Proceedings of the 12th European Conference on Ambient Intelligence, National and Kapodistrian University of Athens, Greece*. Lecture Notes in Computer Science 9425. p. 173. Springer. November 11-13.
- [88] M. Weiser (1991). The Computer for the 21st Century. *Scientific American* (September 1991) pp. 66-75
- [89] C. Wisneski, H. Ishii, A. Dahley, M. Gorbet, S. Brave, B. Ullmer, P. Yarin (1998). Ambient Displays: Turning Architectural Space into an Interface between People and Digital Information. In: N. Streitz, S. Konomi, H. Burkhardt (Eds.), *Proceedings of CoBuild '98*, LNCS Vol. 1370, Springer. pp. 22-32.
- [90] B. Zhang, A.A. Tsiatis, E.B. Laber, M. Davidian (2012). A Robust Method for Estimating Optimal Treatment Regimes. *Biometrics*, May 2012.
- [91] N. Χρηστάκης. Το Πρόσωπο και οι Άλλοι: Θέματα Επικοινωνίας και Κοινωνικής Ψυχολογίας. Αθήνα: Εκδ. Παπαζήση. 2010.