

Proactive Digital Companions in Pervasive Hypermedia Environments

Kimberly García
University of St.Gallen
St.Gallen, Switzerland
kimberly.garcia@unisg.ch

Simon Mayer
University of St.Gallen
St.Gallen, Switzerland
simon.mayer@unisg.ch

Alessandro Ricci
University of Bologna
Cesena, Italy
a.ricci@unibo.it

Andrei Ciortea
University of St.Gallen
St.Gallen, Switzerland
Inria, Université Côte d'Azur, CNRS
Sophia Antipolis, France
andrei.ciortea@unisg.ch

Abstract—Artificial companions and digital assistants have been investigated for several decades, from research in the autonomous agents and social robots areas to the highly popular voice-enabled digital assistants that are already in widespread use (e.g., Siri and Alexa). Although these companions provide valuable information and services to people, they remain *reactive* entities that operate in isolated environments waiting to be asked for help. The Web is now emerging as a uniform hypermedia fabric that interconnects everything (e.g., devices, physical objects, abstract concepts, digital services), thereby enabling unprecedented levels of automation and comfort in our professional and private lives. However, this also results in increasingly complex environments that are becoming unintelligible to everyday users. To ameliorate this situation, we envision *proactive* Digital Companions that take advantage of this new generation of *pervasive hypermedia environments* to provide assistance and protection to people. In addition to Digital Companions perceiving a person's environment through vision and sound, pervasive hypermedia environments provide them with means to further contextualize the situation by exploiting information from available connected devices, and give them access to rich knowledge bases that allow to derive relevant actions and recommendations.

Index Terms—Digital Companion, Pervasive Computing, Web of Things, Hypermedia Agents, Knowledge Graph.

I. INTRODUCTION

The idea of Digital Companions (DC) supporting people to perform tasks that relieve them from routine, time-consuming, and difficult jobs goes back at least to the mid-90s—when Maes [1] had already identified the need for competent and trustworthy agents that could support people in a world that was becoming more complex due to the increasing availability of computers and portable devices. Twenty six years later, given the pervasiveness of connected devices, the digitization of some industries, and the surge of completely new digital ones, it has become rather common that technologies alienate people. Moreover, the availability of digital assistants (e.g., Alexa, Siri, and Google Assistant) has accustomed people to clumsy interactions with their environment. To exacerbate this, the Internet of Things (IoT) is unfolding with the promise to enable more interconnected, personalized, and sophisticated services. “It’s not you—the world has become more complex” also begins a 2011 article in the Harvard Business Review¹,

and traces the increasing complexity of today’s decision-making processes to the digitization of massive amounts of information and to ever greater interdependencies between smart systems.

The recent standardization of the Web of Things (WoT) at the IETF and the W3C is a key factor in this evolution, since it further facilitates application-layer interoperability in the IoT. The Web is now emerging as the middleware of choice for most distributed systems—a homogeneous hypermedia fabric that interconnects everything: devices, physical objects, information resources, abstract concepts, and even holographic resources [2]. The latest standards allow software clients not only to *browse* and *query*, but also to *observe* and *act* on this hypermedia fabric. The Web is now evolving from a collaborative space for people, to a collaborative space for both people and increasingly sophisticated software clients.

II. BRINGING DIGITAL COMPANIONS TO THE WEB

We envision people and DCs co-existing and interacting in pervasive hypermedia environments. The hypermedia environments are pervasive in two fundamental ways: (i) they provide DCs with uniform access to devices in their physical environment, and (ii) they pervade all application areas—since the Web provides a flexible underpinning that can interconnect all these pervasive hypermedia environments in one world-wide system. DC can then navigate the Web (individually or using hypermedia search engines [3]) to discover and interact with digital services, devices, people, other agents, or any other resources on the Web. In addition to uniform access, pervasive hypermedia environments can also be designed and programmed to support interaction, coordination, and stigmergy among people and DCs (see also [4]), or they can act as shared knowledge repositories—e.g., a place where both people and DCs can create and share Web annotations [5].

Hypermedia environments will enable DCs to support people in their professional and private lives online as well as offline. To this end, we picture three types of DCs: (i) *guardians* capable of protecting and keeping people safe in the physical, and virtual world by remaining constantly aware of people’s surroundings, and acting for their benefit, (ii) *assistants* capable of performing tasks that enhance people’s experience either by relieving them from tedious, monotonous,

¹<https://hbr.org/2011/08/the-world-really-is-more-compl>

or dangerous jobs, or by providing them with useful information and services, and (iii) *experts* of a particular physical or virtual environment (e.g., hospital, factory, school) capable of setting the code of conduct, orchestrating adequate interactions between guardians and assistants DCs, and providing people with suitable assistance and protection.

A key characteristic of DCs is to anticipate people’s needs in terms of assistance or protection and to be proactive in meeting those needs. Yorke-Smith [6] proposed several principles of proactive behaviour for personal assistant agents such as: pertinent, competent, anticipatory, safe, unobtrusive, and transparent. Indeed, *unobtrusiveness* and *transparency* are two core features that DCs should be able to achieve in order to manage the increasing complexity of the environments and to be perceived and accepted by people as trustworthy companions [7]. To this end, DCs should be able to either: (i) adapt the hypermedia environment to people, e.g., by showing contextualized and personalized information, (ii) mediate elements of the hypermedia environment by partially adapting the environment to the people, but also asking people to adapt their behaviour through normative cues, (iii) shape the hypermedia environment to decrease the possibility of hazardous behavior.

We believe this vision of proactive DCs is ready to be born given the current and expected advancements in five areas of research, which we discuss in the following.

A. Web of Things (WoT)

Over the past two decades, the IoT has evolved from silos built around custom protocol stacks into a system of systems built around standards. Recent standardization efforts for the WoT allow constrained Web servers to target devices with as little as 10 KiB of RAM and 100 KiB of ROM [8], which means that sensors and actuators can be abstracted behind embedded Web services. Going further, the WoT aims to provide uniform access to IoT devices through the Web—by hiding the protocols and interfaces used to access the devices behind abstract interaction patterns and hypermedia controls (e.g., see the W3C WoT Architecture [9]). From the edge of the network to the cloud, the Web is now emerging as a uniform hypermedia fabric that interconnects IoT devices and digital services—and the W3C WoT Thing Description [10] allows to represent this hypermedia fabric as a world-wide knowledge graph.

B. Knowledge Graphs

The term Knowledge Graph (KG) became popular in 2012 when Google announced their semantic search approach, which allows to look for *things* instead of *strings*. Since then, KGs have been associated to a broad variety of knowledge-based approaches, implementations, and technologies (e.g., DBpedia, Wikidata, and Facebook’s social graph), with the common objective of attempting to stop treating data in isolation, but instead adopting a semantic approach to interconnect such data. However, knowledge-based systems (KBS) have been proposed and implemented since the first wave of AI [11]. KBS’ main component is a knowledge base; in

most cases an ontology—recently re-branded as a KG. A KG’s general objective is to provide a common understanding of a specific domain, allowing applications and systems to seamlessly communicate and inter-operate. Moreover, KGs are able to provide data with meaning by linking it to other related and relevant concepts. One of the most popular applications of a semantic approach is the Semantic Web, which came about due to the need for transforming a Web that was written for human consumption (i.e., with mostly unstructured documents), into a distributed knowledge base and application platform that could allow for better exploitation of information and services by software. Furthermore, the Linked Data principles were proposed [12] to make a Web-scale semantic approach accessible to practitioners. The IoT also benefits from a semantic approach, given that the WoT consists of digital representations of physical objects that are embedded in the Web’s information fabric as resources. This facilitates the interoperation of sensors, actuators, and other physical objects to provide end users with composed services. The WoT is currently being aligned with agent technologies to enable a new generation of dynamic, open, and long-lived systems on the Web.

C. Hypermedia Multi-Agent Systems

The recent advancements in the WoT, Linked Data, and KGs have brought renewed interest in systems of *autonomous agents* on the Web [13]²—a vision dating back to the early days of the Web³ and articulated in the seminal paper on the Semantic Web [14]. In AI research, an *autonomous agent* is commonly defined as “a computer system, situated in some environment, that is capable of flexible autonomous action in order to meet its design objectives” [15]. A *Multi-Agent System (MAS)* is then a system conceptualized in terms of agents that are situated in a *shared environment* and interact with one another to meet their design objectives [16]. MAS research could therefore provide the concepts and tooling for designing, programming, and deploying systems of DCs and other agents on the Web—and in particular in pervasive hypermedia environments. The Web has already raised a lot of interest in MAS research with the advent of service-oriented computing in the early 2000s [17], but the use of Linked Data and hypermedia services in MASs is a recent and active topic of research (e.g., [18], [19]). The latter allows for the conceptual integration of MASs with distributed hypermedia environments.

D. Deep Learning

Driven by the advent of general-purpose computation on GPUs and the large amounts of data being produced daily,

²See also the Dagstuhl Seminar on *Autonomous Agents on the Web* (to be held on February 14-19, 2021): <https://www.dagstuhl.de/en/program/calendar/semhp/?semnr=21072>.

³At the *First International Conference on the World Wide Web* (1994), Sir Tim Berners-Lee was noting that documents on the Web describe real objects and relationships among them, and if the semantics of these objects are represented explicitly then machines can browse through and manipulate reality: <https://videos.cern.ch/record/2671957>

recent years have brought great advancements in deep neural network applications. Of particular interest for DCs are those applications that allow machines to perceive the physical world and to interact with people. To the former point, Computer Vision (CV) can provide the sense of sights to machine, given that CV applications are vast, ranging from simple object detection to image and video captioning, and even image generation. Moreover, recent research on scene understanding is focusing on ways to represent and exploit the semantics of a visual scene [20] [21]. On its side, a clear example of deep learning applications that allow for interacting with people are in the Natural Language Processing (NLP) field. Such is the case of conversational AI [22], which state of the art has greatly advanced thanks to digital assistants such as Alexa and Google Assistant.

Going further, deep learning applications need appropriate channels to interact with their users.

E. Human interfaces

Many such conduits that support the intuitive interaction of individuals and their physical and virtual environments have appeared over the past decades. Particularly promising in our perspective are projected and head-worn Augmented Reality (AR) and Mixed Reality (MR) user interfaces, recently with adaptive retinal projection technology that permits to fuse digital content with natural light in a physically correct way. These are currently still obtrusive, expensive, and not socially accepted; however, over the coming decade, we expect them to become almost indistinguishable from regular glasses, which will be enabled by progress regarding the projection technology (e.g., based on metalenses [23]), battery technology, and (low-power) Internet and Web connectivity. We furthermore expect research breakthroughs regarding the real-time intermixing of virtual and physical *audio* contexts [24], thereby opening up Augmented Audio as the technology of choice for a wide range of scenarios, including overriding of a user's primary visual interface with high-priority audio notifications as well as overlaying lower-priority context information in a discreet way. We argue that, through hypermedia and mediated by proactive DCs, this progress in human interfaces will enable individuals to better monitor, understand, control, and manage environments that are populated by more and more heterogeneous pervasive connected objects—ranging in scale from low-power, energy-harvesting, and even passive devices that can be spread in the environment as "information fountains" [25] and even low-power actuators [26] to mains-connected industrial robots, and beyond.

III. USE CASES

To illustrate our vision, we explore proactive Digital Companions scenarios that address three aspects of people's life.

A. Industrial - Work Safety

In industrial settings, in which workers often find themselves in potentially unsafe situations, a proactive DC that is able to understand the context a person is in by looking at

what is currently happening, reasoning about it, and providing assistance or protection can be of great value [27].

Consider a Building Automation field engineer performing a commission tasks. This engineer handles tools, follows manuals, engages in calls with remote colleagues for instructions or troubleshooting, and documents her actions; all this while being exposed to an increasingly complex and more connected environment that could put the engineer in danger if she were to get distracted from her surroundings. However, a proactive DC capable of looking at a scene and reacting to it will be able to understand the engineer's interactions with the perceived objects, and how those interactions and the environment affect her. Thus, the DC will be capable of anticipating how to support her by: a) using AR and MR to help the engineer follow installation instructions, b) being able to automatically call colleagues for remote support, c) making her aware of the safety measures to take e.g., wearing a helmet or a mask, and d) physically protecting her from moving physical objects endangering her while she is focused on installing a device. The DC will in this way become a trusted entity that is literally watching out for the engineer.

B. Societal - Combating Online Disinformation

Online social networks have undermined the gate-keeping function of traditional news outlets: on the one hand they provide their users with a more diverse, immediate, and unfiltered access to information, but on the other hand users are left with the difficult task of assessing the credibility of online information. Ensuring the quality of information on the open Web is both a societal and a technical challenge that has already attracted diverse initiatives (e.g., EUFACTCHECK⁴, the W3C Credible Web Community Group⁵). On the technical side, recent standards such as W3C Web Annotations [5] allow the weaving of uniform hypermedia-based conversation layers on the Web: people and autonomous agents can use Web Annotations to enrich statements within Web pages with metadata, such as provenance information, truth assessments by others, or opposing views from other sources. DCs can then navigate and reason upon these hypermedia-based conversation layers to present their users with aggregated summaries of credibility assessments (e.g., see [28]). In doing so, DCs can reclaim the gate-keeping function for ensuring a safer online environment for their users—one in which their users are provided with more objective information that allows them to form more informed opinions.

C. Personal - Scalable Health Interventions

Finally, we highlight the utility of DCs to improve micro-decisions in everyday lives. This carries particular relevance regarding the development of effective and scalable strategies to improve individual diets—in light of diet-related non-communicable diseases, which represent the leading cause of mortality worldwide [29]. Such strategies, which are typically

⁴A fact-checking initiative of the European Journalism Training Association: <https://eufactcheck.eu/>

⁵<https://www.w3.org/community/credibility/>

not tailored to the individual [30] are traditionally enabled by dietary self-tracking, which is often discontinued due to the high effort of the manual logging [31]. DCs would enable the automatic tracking of dietary habits while using the hypermedia fabric to associate identified meals with nutrition metadata in the background, and at the same time proactively produce personalized suggestions for individuals, for instance to motivate healthier food choices when grocery shopping.

Tracking and documentation also poses difficulties in formal healthcare environments, such as hospitals. For instance, acquiring documentation during the process of trauma resuscitation is particularly challenging due to several factors pertaining to context, such as multitasking of the person in charge of acquiring the data, parallel activities of the different team members, the plurality of quickly oscillating vital signs to be monitored, or retrospective documentation from collective memory [32]. In such settings, medical staff can benefit from assistance from DCs capable of reasoning on local contextual information, past procedures or events, or global access to knowledge repositories and medical expertise.

IV. CHALLENGES TO OVERCOME

In Section II we discussed recent advancements in several areas we consider central for bringing proactive DCs into pervasive hypermedia environments. Each of these areas also raises a number of research challenges to be addressed.

A. Web of Things (WoT)

The WoT aims at providing a uniform interface for accessing sensors and actuators in physical environments. However, this promise depends on achieving semantic interoperability across connected devices. The recently proposed W3C WoT Thing Description ontology [33] takes advantage of the standardized Semantic Sensor Network Ontology⁶ to provide the first layer of semantics. However, continuous efforts are necessary to align other well-established ontologies with the WoT. Moreover, adopters of the WoT should be made aware of the interoperability and flexibility benefits that the usage of the proposed semantic layers bring to their implementations.

B. Knowledge Graph

Due to the large and intertwined search space that is induced by KGs, querying and reasoning is computationally expensive. However, in recent years, the implementation of KGs to add semantics to large data repositories has been gaining momentum in the industry, given their need for interoperability that allows cross-querying data repositories with heterogeneous information models that have been developed in silos [34]. KGs are being used as a common data model for homogenizing such repositories, and for enhancing query results with richer descriptions. Such spotlight is driving research on more scalable management of RDF data and accelerated query processing [35] and it is advancing the development of scalable infrastructures and tools that can make KG a more accessible technology. However, those advancements are still

on the early stages, and commercial/out-of-the-box solutions that hit all the right usability and computational metrics are not ready yet. Moreover, Linked Data and the WoT are pushing the more traditional knowledge representation research community towards adopting a *distributed* KG approach capable of operating at Web scale⁷. However, the adaptation of more advanced capabilities in distributed KG such as reasoning is being currently researched [36].

C. Hypermedia MAS

Pervasive hypermedia environments provide DCs with uniform access not only to their local deployment context (e.g., sensor readings), but also to global knowledge and services. Proactive DCs then have the ability to make decisions based on their own perception and knowledge about the local environment (e.g., what people are experiencing in that context at that time) as well as cognitive services discovered through the Web [37]. This raises the need for agent architectures that provide a proper level of abstraction to (i) support working in (world-wide) pervasive hypermedia environments, and (ii) balance local and global access to knowledge in real-time. Furthermore, the pervasive hypermedia environments themselves can be designed to support: efficient discovery of context-relevant knowledge and services; interaction and coordination among people, DCs, and any other agents on the Web; and monitoring, interpreting, and regulating the behavior of DCs.

D. Deep learning

Although extremely fast paced advancements are being made in this area, deep neural networks require large amounts of data for training; this poses a specific challenge in non-common scenarios to which not every community has access to, leading to fewer domain-specific reusable data sets (i.e., “small data” contexts). This narrows the scope in which DCs can take advantage of the deep learning techniques that allow perceiving the environment. Current research in zero-shot learning is working to overcome this problem. Moreover, making solutions based on deep learning models transparent and understandable to end-users is still an open challenge.

E. Human interfaces

Notwithstanding the expected technological process regarding human interface technologies, to enable the outlined use cases it is in principle necessary for DCs to be able to perceive their relevant virtual surroundings as well as the immediate, physical, context of the user in real time. Our experience with MR head-mounted displays, even with early commercial products e.g., Google Glass, shows that these do not enjoy a high level of social acceptance in public settings. This effect seems subdued in more formal settings, as the recent success of head-worn MR devices in industry demonstrates. Contrary to visual MR, *augmented audio* technologies might be able to overcome this given the widespread acceptance of wearable audio headsets, which are perceived as less intrusive

⁶<https://www.w3.org/TR/vocab-ssn/>

⁷Distributed Knowledge Graphs EU cost action <https://www.cost.eu/actions/CA19134/>

to the privacy of bystanders even if they include microphones (e.g., noise-cancelling headsets and hearing aids) – still, given the dependency of current MR designs on vision, the social acceptance of DCs will remain a challenge.

V. A CALL TO ACTION

As hinted in the previous sections, there is a lot of potential for meaningful advancement within and across the presented research areas. The Deep Learning and WoT communities are already exploring ways to incorporate knowledge-based approaches to provide meaning to the inputs and outputs of device and services and to enable interoperability among them. All discussed domains require innovative human interfaces for their end-user applications to be widely accepted, used, and to keep acquiring new data and knowledge that can reinforce such technologies. Hypermedia Multi-Agent Systems can provide the flexible underpinning to integrate results from all the other communities—WoT, KG, Deep Learning, Human Interfaces—in a collaborative environment for both people and DCs. Thus, it is essential that these research communities work together to make the vision of proactive pervasive Digital Companions a reality.

REFERENCES

- [1] P. Maes, "Agents that reduce work and information overload," in *Readings in human-computer interaction*. Elsevier, 1995, pp. 811–821.
- [2] S. Mayer, A. Ciorcea, A. Ricci, M. I. Robles, M. Kovatsch, and A. Croatti, "Hypermedia to connect them all - autonomous hypermedia agents and socio-technical interactions," *Internet Technology Letters*, vol. 1, no. 4, April 2018. [Online]. Available: <https://www.alexandria.unisg.ch/255750/>
- [3] S. Bienz, A. Ciorcea, S. Mayer, F. Gandon, and O. Corby, "Escaping the streetlight effect: Semantic hypermedia search enhances autonomous behavior in the web of things," in *Proceedings of the 9th International Conference on the Internet of Things*, ser. IoT 2019. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: <https://doi.org/10.1145/3365871.3365901>
- [4] A. Ricci, M. Piunti, and M. Viroli, "Environment programming in multi-agent systems: an artifact-based perspective," *Autonomous Agents and Multi-Agent Systems*, vol. 23, no. 2, pp. 158–192, 2011. [Online]. Available: <https://doi.org/10.1007/s10458-010-9140-7>
- [5] R. Sanderson, P. Ciccarese, and B. Young, "Web annotation data model," World Wide Web Consortium (W3C), W3C Recommendation, 2017. [Online]. Available: <https://www.w3.org/TR/annotation-model/>
- [6] N. Yorke-Smith, S. Saadati, K. L. Myers, and D. N. Morley, "The design of a proactive personal agent for task management," *International Journal on Artificial Intelligence Tools*, vol. 21, no. 01, p. 1250004, 2012.
- [7] D. A. Norman, "How might people interact with agents," *Communications of the ACM*, vol. 37, no. 7, pp. 68–71, 1994.
- [8] C. Bormann, A. P. Castellani, and Z. Shelby, "Coap: An application protocol for billions of tiny internet nodes," *IEEE Internet Computing*, vol. 16, no. 2, pp. 62–67, 2012.
- [9] M. Kovatsch, R. Matsukura, M. Lagally, T. Kawaguchi, K. Toumura, and K. Kajimoto, "Web of Things (WoT) Architecture, W3C Recommendation 9 April 2020," World Wide Web Consortium (W3C), W3C Recommendation, Apr. 9 2020. [Online]. Available: <https://www.w3.org/TR/2020/REC-wot-architecture-20200409/>
- [10] S. Kaebisch, T. Kamiya, M. McCool, V. Charpenay, and M. Kovatsch, "Web of Things (WoT) Thing Description, W3C Proposed Recommendation 9 April 2020," World Wide Web Consortium (W3C), W3C Recommendation, Apr. 9 2020. [Online]. Available: <https://www.w3.org/TR/2020/REC-wot-thing-description-20200409/>
- [11] R. Akerkar and P. Sajja, "Knowledge-based system, capítulo natural language interface: Question answering system," Jones and Barlett Publishers, pp. 323–330, 2010.
- [12] C. Bizer, T. Heath, and T. Berners-Lee, "Linked data: The story so far," in *Semantic Services, Interoperability and Web Applications: Emerging Concepts*, A. Sheth, Ed. IGI Global, 2011, pp. 205–227.
- [13] A. Ciorcea, S. Mayer, F. Gandon, O. Boissier, A. Ricci, and A. Zimmermann, "A decade in hindsight: The missing bridge between multi-agent systems and the world wide web," in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '19. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 1659–1663. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3306127.3331893>
- [14] T. Berners-Lee, J. Hendler, and O. Lassila, "The semantic web," *Scientific american*, vol. 284, no. 5, pp. 34–43, 2001.
- [15] N. R. Jennings and M. Wooldridge, *Applications of Intelligent Agents*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, pp. 3–28. [Online]. Available: https://doi.org/10.1007/978-3-662-03678-5_1
- [16] G. Weiss, *Multiagent systems: A Modern Approach to Distributed Artificial Intelligence*. MIT press, 2000.
- [17] M. P. Singh and M. N. Huhns, *Service-Oriented Computing: Semantics, Processes, Agents*. John Wiley & Sons, 2006.
- [18] O. Dikenelli, O. Alatlı, and R. C. Erdur, "Where are all the semantic web agents: Establishing links between agent and linked data web through environment abstraction," in *Agent Environments for Multi-Agent Systems IV*, D. Weyns and F. Michel, Eds. Cham: Springer International Publishing, 2015, pp. 41–51.
- [19] A. Ciorcea, O. Boissier, and A. Ricci, "Engineering world-wide multi-agent systems with hypermedia," in *Engineering Multi-Agent Systems*, D. Weyns, V. Mascardi, and A. Ricci, Eds. Cham: Springer International Publishing, 2019, pp. 285–301.
- [20] R. Zellers, M. Yatskar, S. Thomson, and Y. Choi, "Neural motifs: Scene graph parsing with global context," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5831–5840.
- [21] P. Anderson, B. Fernando, M. Johnson, and S. Gould, "Spice: Semantic propositional image caption evaluation," in *European Conference on Computer Vision*. Springer, 2016, pp. 382–398.
- [22] C. Khatri, B. Hedayatnia, A. Venkatesh, J. Nunn, Y. Pan, Q. Liu, H. Song, A. Gottardi, S. Kwatra, S. Pancholi et al., "Advancing the state of the art in open domain dialog systems through the alexa prize," *arXiv preprint:1812.10757*, 2018.
- [23] W. T. Chen, A. Y. Zhu, J. Sisler, Z. Bharwani, and F. Capasso, "A broadband achromatic polarization-insensitive metalens consisting of anisotropic nanostructures," *Nature Communications*, vol. 10, no. 1, pp. 1–7, 2019.
- [24] J. Yang and F. Mattern, "Audio augmented reality for human-object interactions," in *Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*, ser. UbiComp/ISWC '19, ACM. London, UK: ACM, Sep. 2019, pp. 408 – 412.
- [25] L. Sigrist, R. Ahmed, A. Gomez, and L. Thiele, "Harvesting-aware optimal communication scheme for infrastructure-less sensing," *ACM Transactions on Internet of Things*, 2020.
- [26] M. Miskin, A. Cortese, K. Dorsey, E. Esposito, M. Reynolds, Q. Liu, M. Cao, D. Muller, P. Mceuen, and I. Cohen, "Electronically integrated, mass-manufactured, microscopic robots," *Nature*, vol. 584, pp. 557–561, 08 2020.
- [27] A. Shafei, J. Hodges, and S. Mayer, "Ensuring workplace safety in goal-based industrial manufacturing systems," *Procedia Computer Science*, vol. 137, pp. 90–101, 2018.
- [28] A. Wild, A. Ciorcea, and S. Mayer, "Designing social machines for tackling online disinformation," in *Companion Proceedings of the Web Conference 2020*, ser. WWW '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 650–654. [Online]. Available: <https://doi.org/10.1145/3366424.3385770>
- [29] M. Forouzanfar, L. Alexander, H. Anderson, M. M. Tellez Rojo Solis, V. Bachman, S. Biryukov, M. Brauer, R. Burnett, D. Casey, M. Coates, A. Cohen, K. Delwiche, K. Estep, J. Frostad, A. Kc, H. Kyu, M. Moradi-Lakeh, M. Ng, E. Slepak, and C. Murray, "Global, regional, and national comparative risk assessment of 79 behavioural, environmental and occupational, and metabolic risks or clusters of risks in 188 countries, 1990–2013: A systematic analysis for the global burden of disease study 2013," *The Lancet*, vol. 386, p. 2287–2323, 12 2015.

- [30] J. Brug, A. Oenema, and M. Campbell, "Past, present, and future of computer-tailored nutrition education," The American journal of clinical nutrition, vol. 77, no. 4, pp. 1028S–1034S, 2003.
- [31] K. L. Fuchs, M. Haldimann, D. Vuckovac, and A. Ilic, "Automation of data collection techniques for recording food intake: a review of publicly available and well-adopted diet apps," in 2018 International Conference on Information and Communication Technology Convergence (ICTC). IEEE, 2018, pp. 58–65.
- [32] A. Croatti, S. Montagna, A. Ricci, E. Gamberini, V. Albarello, and V. Agnoletti, "Bdi personal medical assistant agents: The case of trauma tracking and alerting," Artificial Intelligence in Medicine, vol. 96, pp. 187 – 197, 2019. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0933365717306000>
- [33] V. Charpenay and S. Käbisch, "On modeling the physical world as a collection of things: The w3c thing description ontology," in The Semantic Web, A. Harth, S. Kirrane, A.-C. Ngonga Ngomo, H. Paulheim, A. Rula, A. L. Gentile, P. Haase, and M. Cochez, Eds. Cham: Springer International Publishing, 2020, pp. 599–615.
- [34] H. Kim, A. Bhattacharyya, and K. Anyanwu, "Semantic query transformations for increased parallelization in distributed knowledge graph query processing," in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 2019, pp. 1–14.
- [35] B. Liu, X. Wang, Y. Yang, and Y. Chai, "Ontods: An ontology-aware distributed storage scheme for rdf graphs," in International Conference on Web Information Systems Engineering. Springer, 2019, pp. 645–659.
- [36] G. Antoniou, S. Batsakis, R. Mutharaju, J. Z. Pan, G. Qi, I. Tachmazidis, J. Urbani, and Z. Zhou, "A survey of large-scale reasoning on the web of data," The Knowledge Engineering Review, vol. 33, 2018.
- [37] S. Montagna, S. Mariani, E. Gamberini, A. Ricci, and F. Zambonelli, "Complementing agents with cognitive services: A case study in healthcare," Journal of Medical Systems, vol. 44, no. 10, p. 188, 2020. [Online]. Available: <https://doi.org/10.1007/s10916-020-01621-7>