

An Expert Digital Companion for Working Environments

Janick Spirig
janick.spirig@student.unisg.ch
University of St.Gallen
St. Gallen, Switzerland

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

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

ABSTRACT

Research in proactive agents capable of anticipating users' needs has been conducted in different application areas, from agents aiming to help users accomplish their next task while using a text editor, to digital assistants that provide support to medical professionals in hospitals. Considering these works and following the rapid development in machine learning, the growing adoption of the IoT, the prevalence of pervasive computing, and the various innovative user interfaces that are becoming mainstream, we propose the creation of *expert* digital companions. *Expert* digital companions are agents that know the specifics of an environment, the available physical equipment in a space, as well as software components that might be reachable in such an environment. These *expert* digital companions can interact with *personal* digital companions, which know the preferences of a specific user. Thus, *expert* and *personal* digital companions collaborate to provide customized assistance to a user given her environment. In this paper, we present an implementation of an *expert* digital companion for employees working in an office and a shopfloor environment that takes advantages of: the decoupling that the W3C WoT Thing Description provides to interact with and control devices and other tools; a Knowledge Graph to provide richer descriptions of elements in the environment; a Computer Vision Algorithm to perceive the physical world; Mixed Reality as the medium to deliver assistance to workers.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing**; • **Computing methodologies** → **Computer vision tasks; Knowledge representation and reasoning**; • **Information systems** → **Semantic web description languages; Personalization; Service discovery and interfaces**.

KEYWORDS

Digital Companions, Web of Things, Computer Vision, Mixed Reality, Knowledge Graph, Ubiquitous Computing

1 INTRODUCTION

This paper follows the vision of Digital Companions (DCs) that assist human users to navigate an increasingly complex world [4], in which a plethora of IoT devices and their services are available, but can remain underutilized by people or can even interfere with their activities by adding steps to accomplish them. Unlike popular digital assistants (e.g., Siri and Alexa), a DC is capable of actively perceiving the environment, it knows about the available services and devices, and it can assess a person's current context in order to proactively compute and deliver relevant and personalized assistance for her [3]. To achieve this, we propose two types of DCs,

namely personal and experts. A personal DC is in charge of assisting and looking after its user, it has learnt relevant information about her, including preferences and relationships with other users. In contrast, an expert DC manages a physical or virtual environment, it knows about the different roles a person can play in such an environment and it is responsible for communicating the code of conduct that should be followed. Moreover, an expert DC communicates with personal DCs to provide them with information, access to services, and devices in the environment it manages.

DCs operate on the Web, given its ability to interconnect documents, devices, virtual services, and even abstract concepts. Furthermore, DCs take advantage of technology fields that have traditionally evolved separately, but which can be merged to create powerful systems that assist humans. Those fields are: the Web of Things (WoT), Knowledge Graph (KG), Deep Learning (DL), Multi-Agent Systems (MAS), and Innovative User Interfaces e.g., Mixed Reality (MR). In this paper, we present an expert DC that takes advantages of the strengths of such technologies to assist workers in an office and a shopfloor environment. Thus, workers are supported, informed and made aware of available services in such spaces, and of actions that can be taken to achieve their tasks, as well as to make their working spaces more comfortable.

2 RELATED WORK

Autonomous agents capable of supporting people in performing routine, time-consuming, and difficult tasks have been a subject of research for several decades. In the mid-90s, Maes [21] identified the need for competent and trustworthy agents capable of supporting people in environments that were becoming increasingly complex due to the popularization of personal computers and the imminent digitization of several industries. In [8] a system of agents capable of supporting a human-organization in performing tasks such as scheduling meetings and monitoring events is presented. Moreover, York-Smith [40] takes an office desktop use case to create proactive agents for assisting users perform common tasks. These and other research on the agents field have not only demonstrated that software agents can work for and collaborate among each other to help on tasks that are simple for humans, but which require constant communication, negotiation, and organization among several software agents. These works have also produced research on the importance of behavioral characteristics that must be considered when creating agents, such as proactivity [40].

In parallel to the agents field, other research areas have called out for a human-centric evolution of technologies to enhance human capabilities and assist people on diverse aspects of their lives. As proposed in [30], this could be achieved by making technologies converge and build on the strengths of one another. Wilks [38] defined an artificial companion as a helpful cognitive agent that knows the habits of a user, talks to her, entertains her, and assists

her in simple tasks. Wilks also predicted that artificial companions were going to be deeply researched within the language and speech research areas. As Wilks pointed out, with the raise in popularization of digital assistants (e.g., Siri, Alexa, and Google Assistant), research in conversational AI has had huge advancement on the recent years [17]. In [27] an artificial companion that considers the implications of empathy and theory of mind is proposed, a companion that learns by imitation. The concept of empathy and emotional connection between humans and companions has been carried over to recent years [24]. Moreover, in [9] an assistant for medical personnel is proposed, in which agent technologies and Mixed Reality are used to present relevant and contextual information to physicians in charge of critical time-dependent pathologies.

In order for assistance systems to evolve into smarter agents, such as DCs [4], they should be provided with capabilities to perceive the environment a user is in. In this way, DCs will be able to find opportunities to proactively help humans, as well as to react to sudden events. Towards perceiving the physical environment, the Web of Things [13] provides standardized means to interact with and control devices available on the environment a DC operates. Moreover, the computer vision field, can provide DCs the ability to understand the physical space by detecting objects located in it [28] and understanding the relationships with one another [18].

3 APPROACH

As previously mentioned, the creation of Digital Companions that offer assistance to people and keep them safe in diverse settings requires building systems that take advantages of the advancements of several technology fields (see Figure 1). On the one hand, for an expert DC to provide users with relevant information or assistance, it needs to be able to perceive the environment, which includes the physical space a person is located in as well as the digital services and devices available in such an environment. Thus, to perceive the physical space, we use computer vision. Well-known object detection algorithms can provide a DC with dynamic information of what might be relevant for a user at a specific moment (e.g., user looking at a tool or a robot). Moreover, advanced algorithms in scene understanding [41] [20] could be incorporated to give a DC a larger and more integrated scope of a user’s current context. On the other hand, for the expert DC to be aware of the state of digital services available in an environment, we take advantage of the WoT, since its objective is to bring any type and size of devices (referred to as *Things* in the World Wide Web Consortium’s Web of Things Architecture¹) to the Web by creating machine-readable representations of their interfaces known as Thing Descriptions (TDs)². Through a TD, an expert DC can gather the means to obtain information about the current status of a device and use any other services offered by it. It furthermore gains access to support services, for instance for discovery and search in dynamic WoT environments [2]. Given that TDs use semantic technologies to describe a device, it is straightforward to connect them in a larger KG to other relevant descriptions of an environment, to go beyond the functional aspects of a device. Thus, a sensor can be a Thing that measures CO₂, which in high concentrations is detrimental to an

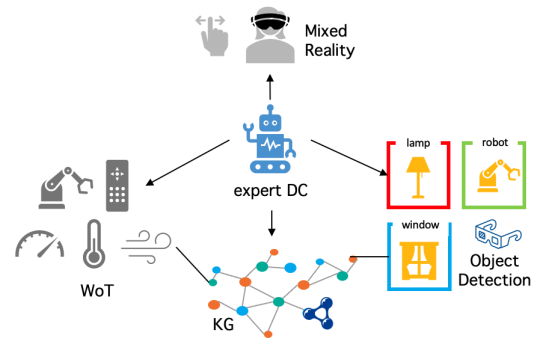


Figure 1: Expert DC Technologies Overview

individual’s health. However, constant ventilation can reduce CO₂ concentration. Such well-known knowledge that in many cases has already been agreed-upon by standardization bodies, can be used to fuel the knowledge of a DC so it can compute suitable assistance for a user. Moreover, the usage of innovative user interfaces, such as Mixed Reality, allows the DC to take advantage of the physical world, and use it as a canvas to provide information, alerts and warnings to a user. Additionally, devices are augmented with information about their current status and running processes, and can even be controlled through intuitive user interfaces made available to the DC users.

4 WORKING ENVIRONMENTS USE CASE

More than two decades ago, Shen and Norrie [31] predicted manufacturing enterprises to operate in an environment characterized by frequently shifting markets, new emerging technologies and globally operating competitors. Today, the manufacturing and industrial sectors are driving the Industry 4.0, in which the increasing adoption of Cyber Physical Systems and the integration of the IoT are escalating the number of involved components and processes. Thus, increasing the complexity of the environments for employees.

We have considered such increased complexity to demonstrate an expert DC capable of assisting employees in two different environments, namely an office and a shopfloor. Following, we introduce some relevant conditions that an expert DC can be constantly monitoring to ensure employees’ comfort, safety, and efficiency.

4.1 Office Environment

In [19], Kohll concludes that for employees to be most productive, they need to feel comfortable in their physical working environment. The level of comfort is influenced by people’s psychological (e.g. gender, age) and psychical (e.g. personality, expectations) characteristics, as well as conditions of the environment they work in. In particular, Indoor Air Quality (IAQ) and lighting levels have the greatest impact on workers comfort and thus eventually on their productivity [14].

4.1.1 Indoor Air Quality (IAQ). In a study Wargocki and Wyon [39] show that poor IAQ decreases productivity levels by six to ten percent and can cause unpleasant symptoms such as headaches.

¹<https://www.w3.org/TR/wot-architecture/>

²<https://www.w3.org/TR/wot-thing-description/>

Moreover, headaches, dizziness, concentration difficulties and fatigue are indicators of the Sick Building Syndrome (SBS) [15], which can occur on building occupants depending on the time they spend in such a building. Therefore, investing in high IAQ can improve working conditions, account towards employees health, and be economically beneficial, since productivity benefits can be up to 60 times higher than the investments required [37]. To achieve a high IAQ, guidelines regarding carbon dioxide (CO₂), humidity, indoor temperature, and ventilation must be met [10] [34] [11]. Specifically, Pettenkofer [34] concluded that a concentration that does not exceed 1000 ppm of CO₂ indicates good IAQ. Moreover, the SECO [10] advises to keep humidity levels in the range of 30% to 50% during winter and in the range of 40% to 60% during summer. The advised indoor temperature is 20 C° to 24 C° in winter and 23.5 C° to 26.5 C° in summer.

In case the recommended CO₂ threshold is surpassed, one mean of intervention is to reduce air pollution by increasing the supply of outdoor air [36], either with natural ventilation (e.g., opening windows and doors) or through the buildings' equipment. Whenever natural ventilation is used, it is important to ensure that the outside air is of good quality [22] as it could often contain pollutants e.g., fine particulate matters (PM2.5) that can be unhealthy for humans. However, a study conducted by Stazi [32] shows that the primary stimulus for people increasing natural ventilation is to satisfy thermal perception, not to improve air quality.

4.1.2 Lighting Conditions. People working in a space with high-quality lighting conditions perceive their office as more attractive, embody a more pleasant mood and overall well-being [33]. On the contrary, a negative perception on lighting conditions impairs the mood of employees and leads to low motivation, eventually affecting their productivity negatively [6]. Additionally, an experiment [23] has shown that since lighting preferences amongst individual workers can vary greatly, allowing them to control the lighting themselves can improve their mood, their satisfaction, and their self-assessed productivity. Moreover, lighting controls should be easily accessible, easy to understand and capable of changing conditions substantially.

4.2 Shopfloor Environment

As mentioned before, the increased complexity of working environments can be burdensome for workers that need to be informed about automated processes, as well as remain aware of the dynamic conditions of the environment they work in. Following, we describe two aspects in which a DC could assist a person working on a shopfloor.

4.2.1 Human in the Loop. The Industry 4.0 has brought a large number of intelligent manufacturing systems which operate increasingly autonomously. Although this is great for automation, it can cause humans to be unaware of current processes and even to be pushed out of the production loop [26]. As a result, workers might face difficulties to intervene when anomalies occur. This concern is often referred to as keeping the »human-in-the-loop«. Jwo, Lin and Lee [16] propose an approach that considers three different angles to keep the human-in-the-loop: (i) intellect, (ii) interact, and (iii) interface. Intellect refers to the human possessing knowledge

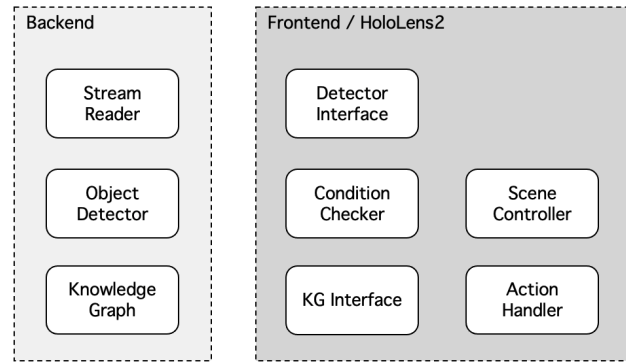


Figure 2: DC's Software Components. The backend is implemented in python running on a GPU/CPU. The frontend is a Unity application implemented in C#, deployed and run on a Microsoft HoloLens2.

from which the IMS can learn from. Interact focuses on the supervisory and collaborative roles of humans and machines. Finally, the interface aspect proposes that humans should not only request data from an IMS but act as an interface to aggregate data with their own knowledge. Approaching the human-in-the-loop challenge from different angles, envisions to achieve a more positive and symbiotic relationship between humans and machines [16].

4.2.2 Employees' Safety. For 2018 the EU-27³ reported that most of work injuries happen in construction and third most in manufacturing [25]. In construction, the majority (56%) of accidents are caused by lack of Personal Protection Equipment (PPE). In [35], it is shown that PPE detection is a valid possibility to measure if employees comply with the health and safety regulation. Moreover, Al-karawi [1] shows that the costs of complete PPE is much lower compared to the cost of an accident.

5 SYSTEM ARCHITECTURE

The implementation of the expert DC for the employees consists of a front- and a backend. The frontend runs on a Microsoft HoloLens2 which facilitates interaction between the DC and the user. Thus, the user wearing the HoloLens2 interacts with the physical world through holograms presented by the frontend, which in turn uses the software components on the backend. Figure 2 shows the general DC architecture and the developed software components.

5.1 Backend

The backend of our system consists of three components:

- A *Stream Reader*. This component reads the frames captured by the Microsoft HoloLens2 PV camera to give the backend access to real-time video stream of the user's field of view. Such stream is recorded at a rate of 30 FPS. This module is largely based on IntelligentEdgeHOL⁴, which converts

³EU-27 refers to the 27 member states of the European union: https://europa.eu/european-union/about-eu/countries_en

⁴<https://github.com/Azure/IntelligentEdgeHOL>




Figure 3: YOLOv4 has been trained to detect 14 different objects and achieves a mean average precision (mAP) of 99.02% at an Intersection over Union (IoU) threshold of 0.5.

the video stream frames into a numerical vector using the OpenCV⁵ library.

- An *Object Detector*. This component receives a numerical vector as an input and uses a Tensorflow implementation⁶ of the YOLOv4 object detection algorithm [5] that detects objects in an image and classifies them according to the images it was trained on. YOLOv4 advantage over other object detection algorithms, such as Faster R-CNN [29], lies in its ability to perform object detection in almost real time. The algorithm has been trained to detect 14 different custom objects with a train and test ratio of 80/20. As Figure 3 shows, the algorithm achieves a mAP of 99.02% on the test images.
- A *Knowledge Graph (KG)*. This component holds machine readable representations of the elements present in the environment. Such representations include Thing Descriptions⁷ (TDs) of sensors, actuators, and equipment available in an office and shopfloor environment. Moreover, the Building Topology Ontology⁸ (BOT) is used to describe spaces, and the OEMA External Factors Ontology⁹ to represent weather and environmental conditions (e.g., rain, humidity, and CO₂), as well as their recommended levels according to the different seasons of the year.

⁵<https://opencv.org/>

⁶<https://github.com/theAIGuysCode/tensorflow-yolov4-tflite>

⁷<https://www.w3.org/TR/wot-thing-description/>

⁸<https://w3c-lbd-cg.github.io/bot/>

⁹<http://www.purl.org/oema/externalfactors>

5.2 Frontend

The frontend application is deployed on a Microsoft HoloLens2. It consists of several components (see Figure 2):

- *Scene controller*. It manages the hologram that the DC is currently displaying for the user. A hologram can provide information, allow controlling equipment, or display warnings. Information holograms provide users with insights regarding the environment, e.g., information about the current process running on a robot. Control holograms enable the user to interact with equipment in the environment, e.g., displaying a button to turn on and off the lights. Warning holograms are displayed when a condition that deserves attention is met e.g., the CO₂ is higher than a normal index. Warning holograms are used to inform a user to take action or to inform them about an action taken by the DC.
- *Detector Interface*. This component is a lightweight Web service¹⁰ that acts as an interface with the backend. When the object detector finds a known object, it informs the frontend through this component.
- *KG interface*. It enables communication with the KG to obtain the TDs of the sensors and equipment that will be controlled or queried. Moreover, the KG interface provides the DC with knowledge about acceptable thresholds to keep workers comfort.
- *Condition Checker*. This component monitors the current conditions of the environment that a user is exploring, e.g., it checks every X minutes if the CO₂ level in the office is acceptable. To do so, the condition checker obtains the current CO₂ sensor reading and uses the query results from the KG interface component, to compare the sensor readings with the known acceptable CO₂ levels.
- *Action Handler*. It is activated when a user interacts with hologram to execute the desired action, e.g. pressing a button to move a robot.

5.3 Flow of Operation

Upon initialization of the back- and frontend, the *Stream Reader* establishes a connection to the *HoloLens2 PV Camera* as shown in Figure 4 (*arrow 1*). The *Stream Reader* then passes each frame as an input into the object detector (*arrow 2*), which runs the YOLOv4 object detection algorithm on the frame to find objects of interest, e.g., a robot, or a lamp. Then, for each object found with a > 70% confidence, a HTTP request is sent to the *Detector Interface* (*arrow 3*) indicating the presence of a known object (a Thing). Once a request has been sent for a specific object, it is remembered by the object detector to reduce time and prevent network overloading. However, when an object is no longer present and later re-appears, a new request is sent to the *Detector Interface*, which communicates with the *Scene Controller* (*arrow 4*).

In turn, the *Scene Controller* talks to the *Condition Checker* (*arrow 5*) which uses the *KG Interface* (*arrow 6*) to query the *Knowledge Graph* (*arrow 7*) for obtaining baseline values to compare with the current environmental conditions. The *Scene Controller* directly influences the hologram displayed for the user wearing the HoloLens2 (*arrow 8*). Thus, the *User* interacts with the hologram (*arrow 9*),

¹⁰<https://gist.github.com/amimaro/10e879ccb54b2caca4e4b81abea455b10>

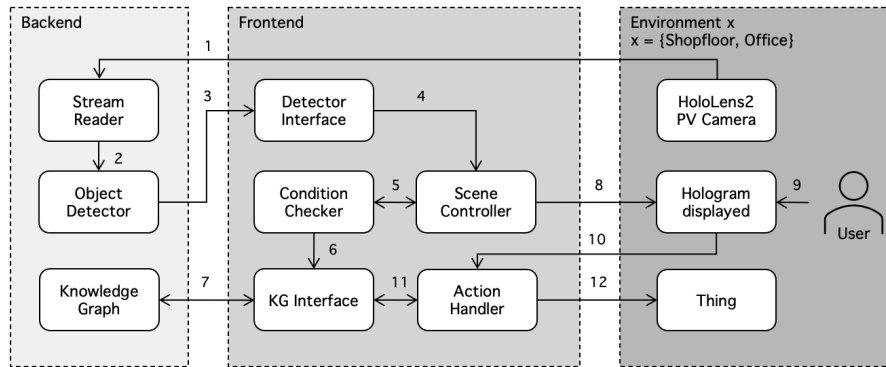


Figure 4: The flow between the different components and the environment

which can either be of type information, control or a warning. In case the interaction with the hologram triggers an action (e.g. the user wants to control one of the robots), the *Action Handler* is called (arrow 10) which communicates with the *KG Interface* in order to gather the means to communicate with the desired *Thing* (arrow 11). Upon getting such information, the *Action Handler* sends out the corresponding API request to the *Thing* and executes the requested action (arrow 12). Moreover, an action can also be triggered directly by the *Condition Checker* (arrow 5) without involving the object detection process, e.g., in case it is found out that the CO₂ concentration in the air returns a value that exceeds 1000 ppm, the *Condition Checker* redirects this information to the *Scene Controller* so that the corresponding warning hologram is displayed to the user. The user then acknowledges the warning and acts accordingly.

6 DEMONSTRATION SCENARIO

To demonstrate that the DC is capable of assisting employees, look after their comfort and keep them in the loop, we have selected and implemented a demonstrator for an office and a laboratory/shopfloor environment¹¹.

The office space used for our demonstrator consists of four desks, it is equipped with an array of sensors that include a CO₂ and a humidity sensor for measuring Indoor Air Quality (IAQ). Moreover, a Philips hue lamp is placed on one of the desks, and the room is enabled with blinds that follow the KNX standard¹². All the sensors and equipment in our scenario have been enabled with a TD which allows the DC to discover Things in the environment and to interact with them. In this environment, the sensors are used by the DC to monitor the IAQ, while the smart lamp and the blinds have been chosen to let workers have control over their lighting conditions preferences.

We have simulated a shopfloor in our university laboratory, which is equipped with ceiling-lights, a UFACTORY xArm⁷¹³ and a PhantomX Reactor Robot Arm¹⁴. Additionally, a Mirocard [12] with an integrated temperature sensor, among others, was placed in this space. In this environment, the ceiling-lights are used to

¹¹Videos of these demonstrators are available at <https://tinyurl.com/5h2p7st3> (office environment) and <https://tinyurl.com/5pnmk5f> (shopfloor environment)

¹²<https://www.knx.org/knx-en/for-professionals/What-is-KNX/A-brief-introduction>

¹³<https://www.ufactory.cc/products/xarm-7-2020>

¹⁴<https://www.trossenrobotics.com/p/phantomx-ax-12-reactor-robot-arm.aspx>

show that the DC can help achieving sufficient lighting, the robots demonstrate that the DC helps to keep the human in the loop about current processes, and to enable the worker to be part of such processes by allowing them to control the robot in a simplified manner. Finally, the Mirocard allows the DC monitoring the room temperature to keep it at a comfortable state [10].

6.1 Usability Tests

The usability test has been conducted in-person with a total of nine participants, mostly students with an average age of 25.56 years old. Eight participants are students and one is a research assistant. Hence, all of them with an academic background. Moreover, 88.80% of the participants had no prior experience using a Microsoft HoloLens.

The test consisted of three parts: familiarization with the hardware, experiencing the expert DC, and collecting impressions. The goal of the first part was to ensure that participants knew how to interact with holograms (e.g., pressing buttons, grabbing, and moving holograms around) while wearing a Microsoft HoloLens2. To achieve this, each participant was asked to played around (for around 3 minutes) with the the Hand Interactions Example scene that is part of the MRTK Examples Hub¹⁵. The second part consisted in experiencing the shopfloor followed by the office environment. Although our work envisions to let users explore an area freely while being assisted by the DC, in this test set-up the DC guides the user through different steps by instructing them on where to go next and what action to execute. The usability test was designed in this guided format to guarantee reproducibility. During the final part of the test, each participant was asked to fill out a short questionnaire followed by an interview of around 30 minutes. During the interview, participants were asked about their general experience with the expert DC, the HoloLens2, and their general view on digital assistants.

6.1.1 Shopfloor Environment. As Figure 5 shows, when the user approaches the door (which has a visual indicator), the DC displays a welcome message informing him about the environment he is entering. The user acknowledges the message and enters the room.

¹⁵<https://docs.microsoft.com/en-us/windows/mixed-reality/mrtk-unity/features/example-scenes/example-hub?view=mrtkunity-2021-05>

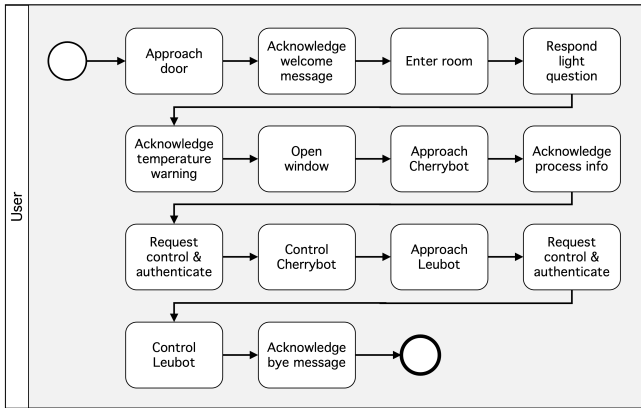


Figure 5: Guided procedure in the usability test to explore the shopfloor environment.

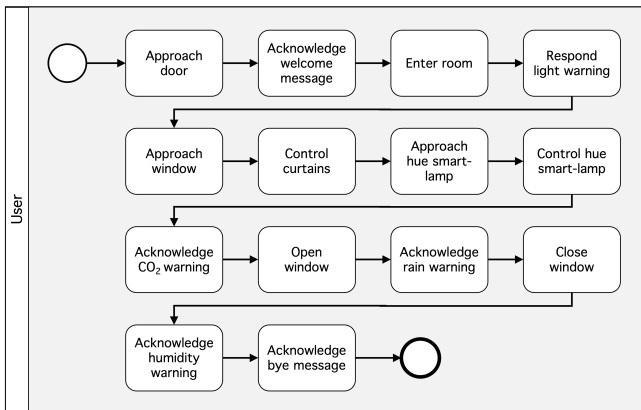


Figure 6: Guided procedure in the usability test to explore the office environment.

As the light is off, the DC asks the user if it should turn on the light for him. If the user confirms the DC switches it on. Following, the user is presented with a temperature warning, the DC recommends to open the windows. Next, the user approaches a UFACTORY xArm7 robot, called *Cherrybot* which is currently placing tennis balls in a can. The user gets information about the robot process and requests to control the robot. After successful authentication, the user helps the *Cherrybot* to finalize the process using holograms the DC displays to him. When the process is complete, the user moves towards a PhantomX Reactor Robot Arm robot, called *Leubot*, he authenticates himself, and controls the *Leubot* freely with the available buttons displayed in a hologram. Finally, a goodbye message is displayed to the user.

6.1.2 Office Environment. Similar to the shopfloor environment, Figure 6 shows that when the user approaches the door, which has a visual indicator, a welcome message is displayed. The user acknowledges the message and enters the room. As the DC has noticed that the user forgot to turn off the light in the shopfloor, the system displays a warning message and asks the user if it should switch the lights off. Upon user confirmation, the DC turns

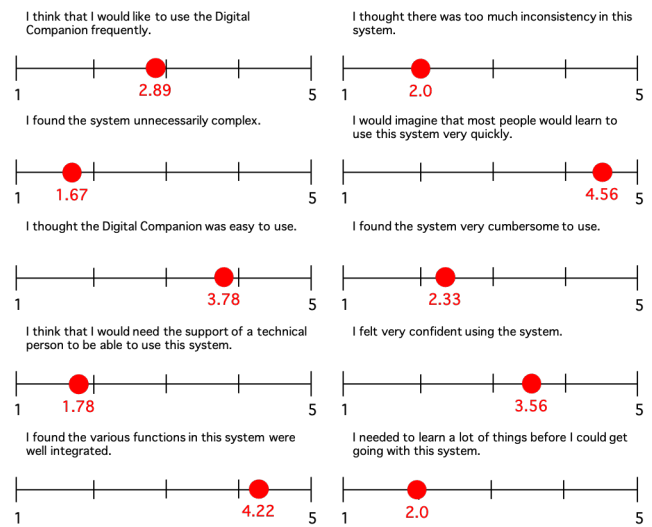


Figure 7: Mean results from the 10 System Usability Scale questions that the 9 participants were asked to respond after having experienced the DC in both environments. Likert scale 1=strongly disagree to 5=strongly agree.

off the light in the shopfloor. Following, the user approaches the window where a hologram to control the blinds is displayed. The user adjusts the blinds according to his preferences. Next, the user approaches the hue smart-lamp which is placed on one of the desks. The DC displays information about the smart-lamp together with an option to adjust its color and brightness. The user freely changes the color and brightness of the lamp. The users continues walking around the office until a warning is displayed about the current CO₂ concentration, indicating that it is too high and suggest to open the window to increase natural ventilation. Suddenly, a second warning is displayed informing that the probabilities of rain have increased considerably. Thus, the user should consider closing the windows. After a while, a final warning is displayed that tells the user that the current humidity level is too high and thus the mechanical ventilation has already been adjusted by the DC. Once the user acknowledges the last warning, a goodbye message is displayed.

6.2 Observations and findings

During the final part of the test, participants were asked to answer some demographic questions, as well as, to fill out a System Usability Scale (SUS) [7] questionnaire, which consists of 10 items with five response options (from strongly agree to strongly disagree). Figure 7 summarizes the participants' responses, which indicate that participants generally have a positive attitude towards the DC as they perceived little complexity, ease of use, and low cumbersome when using the DC.

During the personal interview (conducted after filling out the questionnaire), each participants gave additional insights into their experience with the DC. Moreover, participants were asked about their experience using currently available digital assistants, and their general thoughts regarding proactiveness in digital assistants. Three common topics were identified from such interviews: (i)

physical device constraints, (ii) low usage of conventional digital assistants and (iii) openness towards proactive DCs.

Several participants mentioned that the HoloLens2 became uncomfortable to wear after a while. Thus, they felt relief when the hologram with the bye message appeared. The HoloLens2 weighs 566gm and it has an adjustable band to fit on the users' head¹⁶. Furthermore, some participants reported difficulties combining the real world with the displayed holograms. They mentioned that it requires constant attention to keep perceiving both worlds while not neglecting one or the other. This could be explained by the fact that most participants had no previous experience with Mixed Reality, or because, in order to simplify interactions with the holograms, they were purposely designed to be rather big and are prominently displayed in front of the user's eyes. A few participants compared the use of the HoloLens2 with fully immersive experiences, such as Virtual Reality.

Participants were asked about their current usage and perception of reactive digital assistants such as Siri and Alexa. Most participants reported very low to zero usage of such assistants. The disclosed usage was to perform simple tasks such as setting an alarm, playing music, or calling somebody. However, almost nobody uses an assistant to look up information on the Web. A reason mentioned for this limited usage is that participants think they can execute the task manually faster than when using a digital assistant, since an instruction often has to be repeated several times until the assistant "understands" it correctly. This is perceived as inefficient and cumbersome; therefore more reliable interactions with a system, such as typing are preferred. In summary, participants reported only usage of assistants for tasks they know the assistants are capable of performing. Also, no participant reported usage of conventional digital assistants in public due to social inhibition.

The last notable topic addressed in the interview was how the participants envision the future usage of proactive DCs based on their experience with the tested implementation. Many participants mentioned that they would be open to use a DC with a Mixed Reality user interface given a notable hardware shrink (regular glasses size). Participants imagine making use of a DC in both, their professional and personal lives. While some participants mentioned that it would be useful to have an assistant watching out for them permanently, others addressed the concern that they would not want to use a DC at all times, they would prefer to use it in selected environments, and upon explicit request. Furthermore, the proactive aspect of the DC (e.g., showing alerts about the current environment conditions, showing controls when looking at a robot for some time) was perceived positively, many participants see great potential in this respect. They mentioned that a DC could help them navigate unfamiliar environments, obtain additional information of places, alert them of environmental conditions to take action, and help them control such conditions. Some concerns expressed regarding proactiveness is that information provided without explicitly asking for it should be seamlessly integrated in the current user context. Users should not be disturbed or be presented with unnecessary information. Thus, a balance between proactiveness and reactivity should be found. Additionally, a DC should not conflict with a user's decision making. Users should remain in charge and be

able to alter any decision. Participants mentioned that they would like a DC to know and be able to execute unimportant, boring, and predictable tasks.

7 DC IN HYPERMEDIA ENVIRONMENTS

Our implementation of an expert Digital Companion is built on a combination of co-evolving technology fields that permit it to stay aware and interact with devices and services in its environment on behalf of the user. In an effort to decouple the companion's implementation from the specific devices and services in this environment – and their provided functions – we have equipped all these entities with W3C WoT Thing Descriptions stored in a triplestore. However, these TDs will be made discoverable and searchable via a hypermedia search engine that is geared towards WoT environments [2]. In this way, digital companions that live in and interact with hypermedia environments can discover devices, interaction controls, workflows, and plans that have been made available in an environment through exposing their TDs. Similar to how people navigate the Web using browsers, this will be possible starting from single entry points, which will enable DCs to mediate the interaction of users also in *previously unknown* environments. With the increasing prevalence of W3C WoT TD-described devices¹⁷ and appropriate discovery mechanisms, this has the potential for our proposed DCs to become pervasive across different domains and be able to assist people in different aspects of their lives.

8 CONCLUSIONS AND FUTURE WORK

In this paper, we presented a demonstrator of an expert digital companion that is watching out for employees working in an office and a shopfloor environment, in order to improve their comfort, productivity and keep them safe, by monitoring the current environmental conditions, providing them with means to control and interact with equipment and devices available in such environments, and by keeping them in the loop of autonomous processes. Such an expert digital companion has been built by combining technology fields that have evolved independently, but which strengths have been combined. That is, thanks to the machine-readable descriptions of interfaces of Things (i.e., TDs) available in both environments, the expert DC is able to get readings from sensors, obtain information about current processes on devices, and even control them. Those TDs are part of a Knowledge Graph able to provide richer descriptions of the Things. Moreover, other concepts are represented in the KG, such as pollution, rain, and their thresholds. This provides the DC with the knowledge to inform about abnormal environmental conditions. Moreover, by using a Mixed Reality head-mounted display and an object detection algorithm, the expert DC is able to provide contextualized assistance according to the object a user is looking at. Additionally, Mixed Reality is used by the DC to communicate with the user in an intuitive manner by overlaying information on the physical world that otherwise would be hidden.

The future work of this research is plentiful on making the different research areas converge. To enable higher autonomy levels on DCs, we will bring them to a Multi-Agent System ecosystem. This will entail looking into agents and artifact models, and agents

¹⁶<https://docs.microsoft.com/en-us/hololens/hololens2-hardware>

¹⁷These are, for instance, entering the market as part of Siemens' Desigo Building Automation components

communication protocols. To provide a richer representation of environments and relevant aspects of an application domain, we will keep scouting and utilizing standardized knowledge models, to enable interoperability with future implementations of DCs. In this way, expert DCs made available in different locations could seamlessly work with unknown personal DCs. Regarding personal DCs, means to learn user's habits and preferences will be implemented. Additionally, more sophisticated means to perceive the environment through computer vision will be explored [41]. Finally, ways for the DC to interact with users through Mixed Reality and other innovative user interfaces will be investigated in more depth.

REFERENCES

- [1] Sultan Noori Al-karawi. 2014. *Comparison of costs of personal protective equipment for all workers to avoid costs of fall accidents*. Ph.D. Dissertation. Missouri University of Science and Technology, Rolla, Missouri. https://scholarsmine.mst.edu/cgi/viewcontent.cgi?article=8260&context=masters_theses
- [2] Authors. 2020. Autonomous Search in a Social and Ubiquitous Web. *Personal and Ubiquitous Computing* (Jun 2020). <https://doi.org/10.1007/s00779-020-01415-1>
- [3] Authors. 2020. A Digital Companion for Air Travelers. In *22nd International Conference on Human-Computer Interaction with Mobile Devices and Services (Oldenburg, Germany) (MobileHCI '20)*. Association for Computing Machinery, New York, NY, USA, Article 35, 6 pages. <https://doi.org/10.1145/3406324.3410544>
- [4] Authors. 2020. Proactive Digital Companions in Pervasive Hypermedia Environments. In *2020 IEEE 6th International Conference on Collaboration and Internet Computing (CIC)*. 54–59. <https://doi.org/10.1109/CIC50333.2020.00017>
- [5] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv:2004.10934 [cs, eess]* (April 2020). <http://arxiv.org/abs/2004.10934> arXiv: 2004.10934.
- [6] P R Boyce, J A Veitch, G R Newsham, C C Jones, J Heerwagen, M Myer, and C M Hunter. 2006. Lighting quality and office work: two field simulation experiments. *Lighting Research & Technology* 38, 3 (Sept. 2006), 191–223. <https://doi.org/10.1191/1365782806lrt1610a>
- [7] John Brooke. 1995. SUS: A quick and dirty usability scale. *Usability Eval. Ind.* 189 (11 1995).
- [8] Hans Chalupsky, Yolanda Gil, Craig A Knoblock, Kristina Lerman, Jean Oh, David V Pynadath, Thomas A Russ, and Milind Tambe. 2002. Electric Elves: Agent technology for supporting human organizations. *AI magazine* 23, 2 (2002), 11–11.
- [9] Angelo Croattini, Manuel Bottazzi, and Alessandro Ricci. 2020. Agent-Based Mixed Reality Environments in Healthcare: The Smart Shock Room Project. In *Advances in Practical Applications of Agents, Multi-Agent Systems, and Trustworthiness. The PAAMS Collection*, Yves Demazeau, Tom Holvoet, Juan M. Corchado, and Stefania Costantini (Eds.). Springer International Publishing, Cham, 398–402.
- [10] Suisse Secrétariat d'Etat à l'économie (SECO). 2020. Commentaire de l'ordonnance 3 relative à la loi sur le travail, Art. 16. <https://www.seco.admin.ch/seco/fr/home/Arbeit/Arbeitsbedingungen/Arbeitsgesetz-und-Verordnungen/Wegleitungen/wegleitung-zur-argv-3.html#-1755811567>
- [11] Lei Fang, David Peter Wyon, Geo Clausen, and Povl Ole Fanger. 2004. Impact of indoor air temperature and humidity in an office on perceived air quality, SBS symptoms and performance. *Indoor Air* 14, s7 (Aug. 2004), 74–81. <https://doi.org/10.1111/j.1600-0668.2004.00276.x>
- [12] Andres Gomez. 2020. On-demand communication with the batteryless MiroCard: demo abstract. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 629–630.
- [13] Dominique Guinard, Vlad Trifa, Friedemann Mattern, and Erik Wilde. 2011. From the internet of things to the web of things: Resource-oriented architecture and best practices. In *Architecting the Internet of things*. Springer, 97–129.
- [14] Barry P. Haynes. 2008. The impact of office comfort on productivity. *Journal of Facilities Management* 6, 1 (Feb. 2008), 37–51. <https://doi.org/10.1108/14725960810847459>
- [15] Sumedha Joshi. 2008. The sick building syndrome. *Indian journal of occupational and environmental medicine* 12 (08 2008), 61–4. <https://doi.org/10.4103/0019-5278.43262>
- [16] Jung-Sing Jwo, Ching-Sheng Lin, and Cheng-Hsiung Lee. 2021. Smart technology-driven aspects for human-in-the-loop smart manufacturing. *The International Journal of Advanced Manufacturing Technology* 114, 5-6 (May 2021), 1741–1752. <https://doi.org/10.1007/s00170-021-06977-9>
- [17] Chandra Khatri, Behnam Hedayatnia, Anu Venkatesh, Jeff Nunn, Yi Pan, Qing Liu, Han Song, Anna Gottardi, Sanjeev Kwatra, Sanju Pancholi, et al. 2018. Advancing the state of the art in open domain dialog systems through the alexa prize. *arXiv preprint arXiv:1812.10757* (2018).
- [18] Ue-Hwan Kim, Jin-Man Park, Taek-Jin Song, and Jong-Hwan Kim. 2019. 3-d scene graph: A sparse and semantic representation of physical environments for intelligent agents. *IEEE transactions on cybernetics* 50, 12 (2019), 4921–4933.
- [19] Alan Kohll. 2019. How Your Office Space Impacts Employee Well-Being. <https://www.forbes.com/sites/alankohll/2019/01/24/how-your-office-space-impacts-employee-wellbeing/> Section: Leadership Strategy.
- [20] Yikang Li, Wanli Ouyang, Bolei Zhou, Jianping Shi, Chao Zhang, and Xiaogang Wang. 2018. Factorizable net: an efficient subgraph-based framework for scene graph generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 335–351.
- [21] Pattie Maes. 1995. Agents that reduce work and information overload. In *Readings in human-computer interaction*. Elsevier, 811–821.
- [22] William W Nazaroff. 2013. Four principles for achieving good indoor air quality. *Indoor Air* 23, 5 (Oct. 2013), 353–356. <https://doi.org/10.1111/ina.12062>
- [23] Guy Newsham, Jennifer Veitch, Chantal Arsenault, and Cara Duval. 2004. Effect of dimming control on office worker satisfaction and performance. In *Proceedings of the IESNA annual conference*. Citeseer, 19–41.
- [24] Jasmin Niess and Pawel W Woźniak. 2009. Embracing companion technologies. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–11.
- [25] European Statistics on Accidents at Work (ESAW). 2020. Accidents at work statistics. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Accidents_at_work_statistics
- [26] Marie-Pierre Pacaux-Lemoine, Damien Trentesaux, Gabriel Zambrano Rey, and Patrick Millot. 2017. Designing intelligent manufacturing systems through Human-Machine Cooperation principles: A human-centered approach. *Computers & Industrial Engineering* 111 (Sept. 2017), 581–595. <https://doi.org/10.1016/j.cie.2017.05.014>
- [27] Malcolm Peltu and Yorick Wilks. 2008. Close engagements with artificial companions: Key social, psychological, ethical and design issues. OII Forum Discussion Paper.
- [28] Joseph Redmon and Ali Farhadi. 2017. YOLO9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7263–7271.
- [29] Shaoping Ren, Kaiming He, Ross Girshick, and Jian Sun. 2016. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *arXiv:1506.01497 [cs]* (Jan. 2016). <http://arxiv.org/abs/1506.01497> arXiv: 1506.01497.
- [30] Mihail C Roco and William S Bainbridge. 2002. Converging technologies for improving human performance: Integrating from the nanoscale. *Journal of nanoparticle research* 4, 4 (2002), 281–295.
- [31] Weiming Shen and Douglas H. Norrie. 1999. Agent-Based Systems for Intelligent Manufacturing: A State-of-the-Art Survey. *Knowledge and Information Systems* 1, 2 (May 1999), 129–156. <https://doi.org/10.1007/BF03325096>
- [32] Francesca Stazi, Federica Naspi, Giulia Ulpiani, and Costanzo Di Perna. 2017. Indoor air quality and thermal comfort optimization in classrooms developing an automatic system for windows opening and closing. *Energy and Buildings* (2017), 15.
- [33] Ja Veitch, Gr Newsham, Pr Boyce, and Cc Jones. 2008. Lighting appraisal, well-being and performance in open-plan offices: A linked mechanisms approach. *Lighting Research & Technology* 40, 2 (June 2008), 133–151. <https://doi.org/10.1177/1477153507086279>
- [34] Max von Pettenkofer. 1858. *Über den Luftwechsel in Wohngebäuden*. Cotta.
- [35] Zijian Wang, Yimin Wu, Lichao Yang, Arjun Thirunavukarasu, Colin Evison, and Yifan Zhao. 2021. Fast Personal Protective Equipment Detection for Real Construction Sites Using Deep Learning Approaches. *Sensors* 21, 10 (May 2021), 3478. <https://doi.org/10.3390/s21103478>
- [36] Pawel Wargocki. 2008. Improving Indoor Air Quality Improves The Performance Of Office Work And School Work. (2008), 7.
- [37] Pawel Wargocki, K. W. Tham, S. C. Sekhar, and D. Cheong. 2003. Estimate of an economic benefit from investment in improved indoor air quality in an office building. In *Proceedings of Healthy Buildings 2003*, Vol. 3. University of Singapore, Department of Buildings, 382–387. <http://hb2003.nus.edu.sg/>
- [38] Yorick Wilks. 2006. Artificial companions as a new kind of interface to the future internet. (2006).
- [39] David P Wyon and Pawel Wargocki. 2005. Indoor Air Quality Effects On Office Work. (2005), 13.
- [40] Neil Yorke-Smith, Shahin Saadati, Karen L Myers, and David N Morley. 2012. The design of a proactive personal agent for task management. *International Journal on Artificial Intelligence Tools* 21, 01 (2012), 1250004.
- [41] Alireza Zareian, Svebor Karaman, and Shih-Fu Chang. 2020. Bridging knowledge graphs to generate scene graphs. In *European Conference on Computer Vision*. Springer, 606–623.