

A Review of Cognitive Assistants for Healthcare: Trends, Prospects, and Future Directions

SARAH MASUD PREUM, Department of Computer Science, University of Virginia

SIRAJUM MUNIR, Bosch Research and Technology Center

MEIYI MA, Department of Computer Science, University of Virginia

MOHAMMAD SAMIN YASAR, Department of Electrical and Computer Engineering, University of Virginia

DAVID J. STONE, Departments of Anesthesiology and Neurosurgery, and the Center for Advanced Medical Analytics, University of Virginia School of Medicine; MIT Critical Data, Laboratory for Computational Physiology, Harvard-MIT Health Sciences and Technology, Massachusetts Institute of Technology

RONALD WILLIAMS and HOMA ALEMZADEH, Department of Electrical and Computer Engineering, University of Virginia

JOHN A. STANKOVIC, Department of Computer Science, University of Virginia

Healthcare cognitive assistants (HCAs) are intelligent systems or agents that interact with users in a context-aware and adaptive manner to improve their health outcomes by augmenting their cognitive abilities or complementing a cognitive impairment. They assist a wide variety of users ranging from patients to their healthcare providers (e.g., general practitioner, specialist, surgeon) in several situations (e.g., remote patient monitoring, emergency response, robotic surgery). While HCAs are critical to ensure personalized, scalable, and efficient healthcare, there exists a knowledge gap in finding the emerging trends, key challenges, design guidelines, and state-of-the-art technologies suitable for developing HCAs. This survey aims to bridge this gap for researchers from multiple domains, including but not limited to cyber-physical systems, artificial intelligence, human-computer interaction, robotics, and smart health. It provides a comprehensive definition of HCAs and outlines a novel, practical categorization of existing HCAs according to their target user role and the underlying application goals. This survey summarizes and assesses existing HCAs based on their characteristic features (i.e., interactive, context-aware, and adaptive) and enabling technological aspects (i.e.,

Authors' addresses: S. M. Preum, Department of Computer Science, Dartmouth College, Hanover, NH, 03755, USA; email: spreum@dartmouth.edu; S. Munir, Bosch Research and Technology Center, Pittsburgh, PA, 15222, USA; email: sirajum.munir@us.bosch.com; M. Ma, Department of Computer Science, University of Virginia, Charlottesville, VA, 22904, USA; email: mm5tk@virginia.edu; M. S. Yasar, Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA, 22904, USA; email: msy9an@virginia.edu; D. J. Stone, Departments of Anesthesiology and Neurosurgery, and the Center for Advanced Medical Analytics, University of Virginia School of Medicine and MIT Critical Data, Laboratory for Computational Physiology, Harvard-MIT Health Sciences and Technology, Massachusetts Institute of Technology, Charlottesville, VA, 22904, USA; email: djs4v@hscmail.mcc.virginia.edu; R. Williams and H. Alemzadeh, Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA, 22904, USA; emails: {rdw, ha4d}@virginia.edu; J. A. Stankovic, Department of Computer Science, University of Virginia, Charlottesville, VA, 22904, USA; email: jas9f@virginia.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0360-0300/2021/02-ART130 \$15.00

<https://doi.org/10.1145/3419368>

sensing, actuation, control, and computation). Finally, it identifies critical research questions and design recommendations to accelerate the development of the next generation of cognitive assistants for healthcare.

CCS Concepts: • **Information systems** → **Decision support systems**; • **Human-centered computing** → **Ubiquitous and mobile computing systems and tools**; • **Computing methodologies** → **Artificial intelligence**; **Machine learning**; • **Computer systems organization** → **Embedded and cyber-physical systems**;

Additional Key Words and Phrases: Cognitive assistant, agent based systems for healthcare, smart health, intelligent agent, intelligent assistant, virtual assistant, virtual agent, personal assistant, healthcare application

ACM Reference format:

Sarah Masud Preum, Sirajum Munir, Meiyi Ma, Mohammad Samin Yasar, David J. Stone, Ronald Williams, Homa Alemzadeh, and John A. Stankovic. 2021. A Review of Cognitive Assistants for Healthcare: Trends, Prospects, and Future Directions. *ACM Comput. Surv.* 53, 6, Article 130 (February 2021), 37 pages. <https://doi.org/10.1145/3419368>

1 INTRODUCTION

The rapid digitization of healthcare, along with the advancement in ubiquitous computing technology, has accelerated the development of assistive technologies for healthcare. These assistive technologies aim to support different user groups in the healthcare domain ranging from patients to their healthcare providers. Although there are several existing surveys on assistive technologies for healthcare [4, 35, 42, 51, 61, 66], only a few of them focus on cognitive assistants [35, 51, 66]. Most of the existing surveys focus on reviewing the assistive technologies for healthcare from application domains rather than pointing out the key technological challenges and future directions to provide cognitive assistance in healthcare. Thus, there is a knowledge gap in finding the key challenges and state-of-the-art technologies suitable for developing capable cognitive assistants for healthcare.

However, cognitive assistant for healthcare is an emerging topic of current and future research. It poses several interesting challenges that should be addressed to create a significant impact on outcomes of individual and population-level health. To bridge this knowledge gap, we provide a comprehensive survey of the existing research and state-of-the-art healthcare cognitive assistants (HCAs) in this article. While there are different perspectives of assistive technology for healthcare ranging from neuroscience to robotics, we specifically focus on existing research on cognitive assistants for healthcare from the domains of robotics [1, 7, 26, 47, 66, 84, 93, 110, 123, 135, 135, 138, 142], artificial intelligence [33, 109, 110, 117, 124, 129], cyber-physical systems [18, 32, 33, 40, 66, 78, 84, 93, 103, 104, 128, 129], human-computer interaction [32, 33, 35, 40, 66, 93, 103, 104, 110], and smart and connected health [124, 129, 136, 137].

There is no standard definition of healthcare cognitive assistants (HCAs). Our definition of HCA is inspired by existing definitions of related systems, including general cognitive assistants, intelligent agents, assistive technology for cognition, and healthcare assistants or agents. The relevant existing definitions can be found in Section 1 of the online supplemental materials.¹ We define HCAs as follows: *A Cognitive Assistant for healthcare is an **interactive, contextual, and adaptive** system that possesses **computational capabilities** based on a large amount of data or explicit models of the environment and provides cognition power to improve health outcome by either **augmenting human intelligence** or providing complementary assistance for **cognitive impairment**.*

Here, an improved health outcome refers to any positive outcome for the physical, mental, and psychological health or well-being of an individual. The improvement can be achieved by providing cognitive support to (or augmenting the cognitive ability of) anyone involved in

¹There, we also compare HCAs with agents used in *robotics* and *reinforcement learning*.

healthcare, i.e., physicians, nurses, patients, in-home caregivers, or emergency responders. Thus, improving the efficiency of healthcare providers or augmenting their cognitive ability can be one of the goals of HCAs. Also, computational capabilities can be based on natural language processing, machine learning, computer vision, and reasoning and inference. The environment refers to the collection of situations, contexts, resources, and users (e.g., a caregiver, patient, or a human operator) that the cognitive assistant is used in or interacts with. The main contributions of this survey are the following:

- (1) We provide a **comprehensive definition** of HCAs and identify the **characteristic features of HCAs** that are suitable for the underlying application domains (Section 3). We also identify the **critical aspects** of HCAs that are relevant to multiple domains, including but not limited to **robotics, artificial intelligence, cyber-physical systems, human-computer interaction, and smart and connected health**.
- (2) We review and analyze **existing research and state-of-arts of HCAs** in terms of these key features and critical cyber-physical components. We create and consolidate **taxonomies of HCAs** according to these features (Section 3) and cyber-physical components (Section 4).
- (3) We also present the **application goals/objectives of HCAs** in terms of who they assist (e.g., patients or their care providers) and the types of assistance they provide (e.g., real-time decision support, or complement a cognitive impairment) (Section 2). We identify the potential application requirements for each of these application types.
- (4) We provide a set of **critical challenges, future research directions, and design guidelines** for the next generation of intelligent or cognitive healthcare assistants with respect to current and imminent pervasive technologies (Section 5).

A brief **outline of the scope of this article and the range of existing HCAs applications** are presented in Figure 1. It shows the different user groups of assistive healthcare applications and the variety of situations where such applications are used. The defining characteristics (i.e., interactive, adaptive, and context-aware) and cyber-physical aspects of HCAs (i.e., sensing, actuation, and control and computation) are also presented in this figure. We considered a wide array of research on cognitive assistants, including intelligent personal assistants, personal software agents, assistive robots, virtual assistants, virtual coach for healthcare, personalized assistants, and assistive technology used for healthcare. The goal is to identify relevant existing research even though different research communities use different terminologies to describe their works. However, only the works that at least partially satisfy the proposed definition of HCAs mentioned above are included in this survey. A **list of acronyms** used throughout the article is presented in **Table 8 in Section 6**.

2 APPLICATIONS OF COGNITIVE ASSISTANT FOR HEALTHCARE

Several existing surveys on healthcare assistants provide taxonomies of healthcare applications [42, 50] and present a detailed review for the surveyed applications. For instance, they cover categorization based on the intended users, (i.e., patient-centered, staff-centered, healthcare organization centered) [51], categorization based on the functionality of applications [50, 51], or categorization based on the input and output modalities [61]. While such categorizations provide an overview of healthcare applications targeted at different user roles, they do not highlight the patterns of design requirements and the technical challenges relevant to these applications. Hence, we categorize HCAs according to user roles, situations, and underlying objectives, as these factors determine the design requirements of HCAs from different application areas. We also characterize

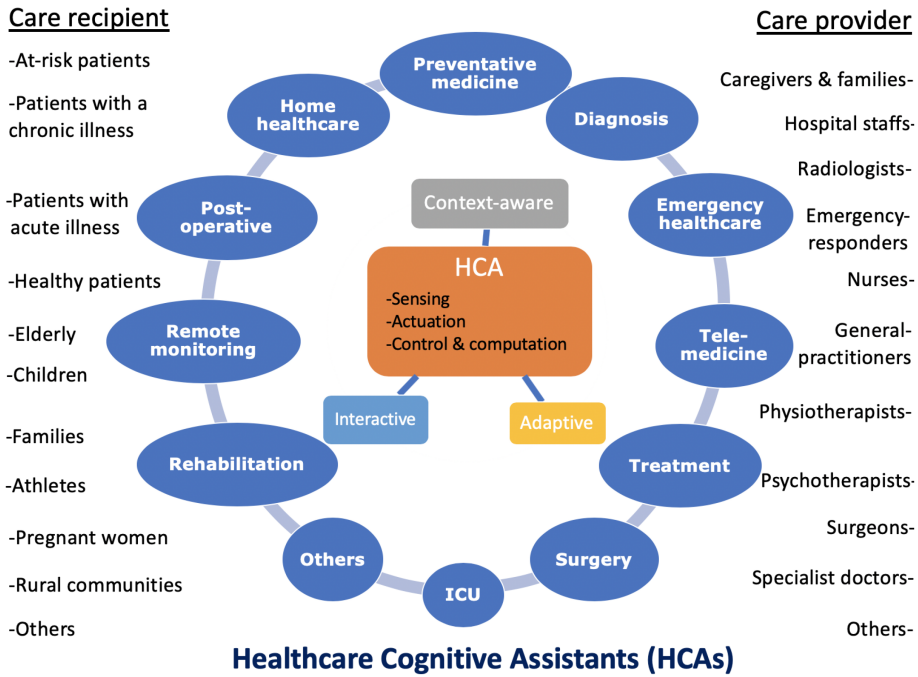


Fig. 1. This survey focuses on reviewing healthcare cognitive assistants (HCA). We show the **dual user roles** common in healthcare practice: **the care recipients** (on the left) and **the care providers** (on the right). In the blue circle, we list the different situations where HCAs are used, e.g., home healthcare, preventative medicine, and diagnostics. Inside the blue circle, we show the defining features and cyber-physical components of HCAs. The **defining features for an HCA are interactive, context-aware, and adaptive** (referring to Section 3). These features are implemented by **different cyber-physical components, including sensing, actuation, control, and computation** (referring to Section 4). Referring to our proposed definition of HCAs presented above, the goal of an HCA is to improve health outcome. This is achieved by either augmenting the user’s (i.e., a care provider or a care recipient) intelligence or providing complementary assistance for a cognitive impairment of a patient.

the essential features and cyber-physical systems aspects of the underlying technology of HCAs as described in Sections 3 and 4. Based on our review, we categorize HCAs in the following classes:

- **Patient-facing HCAs** to provide pervasive cognitive assistance even in the absence of professional healthcare providers (refer to Table 1).
- HCAs to provide **cognitive assistance to professional healthcare providers** for scalable, efficient, and effective care delivery (refer to Table 2).
- HCAs used for **training patients and professional healthcare providers** (refer to Table 3).

It should be noted that patients and professional healthcare providers mentioned above can include any of the categories depicted in Figure 1. We list HCAs for training as a separate category, since it has different application requirements than the other two categories as shown in Table 3.

3 FEATURES OF COGNITIVE ASSISTANT FOR HEALTHCARE

Upon reviewing the current research on cognitive assistants and assistive technologies for healthcare [12, 27, 55, 70, 85], we have identified three key features that enable an HCA to provide

Table 1. Different Types of Applications That Belong to the Set of Patient-facing HCAs

Application Type	Potential Application Requirement	Examples
Provide decision support	Natural interaction, explainable, user context, empathy	A smartphone-based conversational agent for self-care of heart failure patients [28].
		A desktop-based 3D virtual, empathetic agent for suggesting intervention to change drinking behavior [67].
Educate patients	Natural interaction, user context, explainable, multimodal actuation	A smartphone-based 3D conversational virtual agent for educating individuals with atrial fibrillation [9].
Provide diagnostic support	Natural interaction, user context, explainable, multimodal sensing	A desktop-based conversational agent for early dementia detection based on the standard protocol of questionnaire [2].
		A chatbot for checking symptoms and mapping them to diseases using medical knowledge bases [34].
Support activities of daily living (ADL)	Temporal, spatial, personal, and situational context, mobility or ubiquitous, adaptiveness, context-aware interaction, energy efficiency, embedded processing	A mobile, autonomous robotic assistant for generating reminders for routine activities, answering a limited set of questions and providing guided navigation of the user [93].
Support specific cognitive challenges	Personal, situational, and spatial context, multimodal interface, empathy, embedded processing	A smartphone-based personalized navigational guidance system for visually challenged individuals [3].
Provide companionship	Personal and situational context, empathy, emotion, appearance, multimodal interaction including verbal and nonverbal interaction, adaptive	An emotive companion robot for the elderly population that represents a pet cat in terms of physical appearance [26].
Provide counseling or psychotherapy	Personal and situational context, empathy, emotion, appearance, multimodal interaction including verbal and nonverbal interaction, adaptive	A robot-based anxiety management system for providing personalized therapies to reduce user's anxiety level [1].
Provide physical therapy	Multimodal actuation including haptic feedback, visualization, adaptive, personal and situational context	A wearable, physio-therapeutic system for post-surgery rehabilitation that utilizes haptic feedback to ensure safe and effective movement of a target body part or joint based on depth sensing [104].
Data collection & self-monitoring	Temporal, spatial, personal, and situational context, mobility or ubiquitous, adaptiveness, and energy efficiency	Providing support for (i) self-monitoring of patients or (ii) collection of longitudinal behavioral data for disease management or health risk assessment [61] [44].

Each application type poses some requirements as shown in the second column. For instance, patient-facing HCAs that provide decision support to patients should support natural interaction, provide explainable intervention, be aware of user context, and demonstrate empathy. The third column shows some examples of existing systems that belong to these application types.

cognitive support effectively. They are context-awareness, interactivity, and adaptiveness. In the following subsections, we review the existing research in terms of each of these features. Specifically, how existing HCAs implement these features, what are the emerging trends, and what are the potential challenges that are yet to be addressed. It should be noted that often these features are intertwined. For instance, to identify the **context** of a user, the system would need to **interact** with users or their environment. The outline of this section is also summarized in Figure 2.

Table 2. HCAs to Provide Cognitive Assistance to Professional Healthcare Providers

Application Type	Potential Application Requirement	Examples
Diagnostic decision support	Deep understanding of diseases and their interpretation in multiple modalities (X-ray, Ultrasound, CT, MRI, Clinical text); natural language understanding; multimodal anomaly detection; summaries imaging studies	MedicalSieve, a radiologist cognitive assistant for an image-guided informatics system to filter the essential clinical information for diagnosis and treatment planning. Integrates data from EHR, pharmacy, labs, hospital notes, and radiology/cardiology images and videos [58].
Patient screening and assessment, and visit documentation	User training for adaptation, real-time sensing and contextual information retrieval, intelligent visualization, multimodal perception, hands-free interface, domain-adapted seamless verbal interaction	Multimodal AR-based Cognitive assistant for screening patients, visualize patient records and direct data acquisition and control. Supports multimodal hands-free, real-time interaction through speech, eye tracking, head-mounted display, and large VR display [129].
Treatment decision support for physicians	Multimodal data analysis and perception, natural language processing, natural language understanding, domain knowledge, evidence based treatment decision, confidence score of recommendation	Integrate information and extract knowledge from relevant guidelines, best practices, and growing published medical research to suggest personalized treatment options for oncologists based on a patient's longitudinal medical records. The treatment options are ranked by level of confidence and contain supporting evidence. Supports 13 types of cancer [137].
ICU monitoring and analysis	Automatic integration of electronic medical records; integration of temporal, textual, image, and video data analytics; support multiple programming languages and reusability	IBM Artemis enables analytic developers within a medical institution to build and run real-time analytics on large streams of heterogeneous data from ICU patients. Used in neonatal and neurological ICUs [10, 58].
Psychotherapy assistant	Interface to simulate activities (e.g., driving), mobility and navigation (if applicable); haptic, audio, olfactory stimuli; customizable stimuli delivery	VR-based assistant for treatment of combat-related PTSD. It enables the physicians to personalize the therapy sessions according to their patients' individual needs and simulate VR scenarios to introduce and control a patient's real-time trigger stimuli (e.g., gunfire, vehicle crash, explosions, and insurgent attacks) [111].
Patient interview agent	Natural and realistic interface; Verbal and nonverbal interaction; Natural language understanding; Dialogue management; Contextual knowledge integration	Elic, virtual human interviewer agent to conduct semi-structured interviews to initiate interactions to enable automatic assessment of distress indicators (i.e., behaviors associated with PTSD, depression, or anxiety)[25].
Nurse assistant	User-friendly interface and appearance, mobile platform, posture stability, intelligent navigation control, 3D sensing and perception, real-time monitoring and safety control	A humanoid, mobile robotic nursing assistant for lifting and moving patients and heavy objects inside a hospital to increase patient and nurse safety and operational efficiency [47].
Documentation and alert assistant for emergency response	Knowledge of emergency response protocols; speech recognition and transcription in noisy environment; device and resource constraints;	Trauma Tracker, a smartphone-based assistant for accurate documentation of trauma resuscitation and providing visual alert to first responders on a display or wearable smart glass device. It uses belief-desire-intention model-based agent. It enables responders to automatically record dynamic vital signs of a patient, results of diagnostic procedures, conducted at scene, and validate drug dosage and administration details [22, 78].
Support responder for managing large-scale emergency	Knowledge processing; Knowledge of emergency response protocols; multimodal perception; autonomous navigation; device and resource constraints; speech recognition and transcription in noisy environment	Cognitive mobile robotic agents to assist search-and-rescue operations in large-scale emergency situations by completing tasks instructed by human rescuers autonomously and communicating even in low-bandwidth setting efficiently. It interprets human commands expressed in natural language and gestures, scans incident area for potential anomalies or hazards, detects objects of interest, and navigates through emergency area [142].
Support responder for managing emergency	Domain knowledge integration; speech recognition and transcription in noisy environment; device and resource constraints; natural language understanding and generation;	A voice based cognitive assistant for real-time, protocol-driven decision support and automatic documentation for emergency medical service (EMS) providers. This assistant extracts EMS protocol specific information from the spoken language collected at scene using a headset and uses that information to model and execute EMS protocols for real-time intervention suggestion to ensure safety of patients [94, 98, 102, 126].
Robotic surgery assistant	Trocar planning and placement, extracting surgical context and using knowledge of the present context, patient's anatomy and surgeon's profile to assist the surgeon's maneuvers; performing monotonous and recurring tasks that require less cognitive load	ARsist, an augmented reality application to assist the first assistant (FA) in passing necessary materials to the main surgeon and help him with trocar planning and placement, manipulate hand-held instruments, and perform robot docking. It supports real-time, 3D rendering of the robotic instruments, hand-held instruments and endoscope and real-time stereo endoscopy that is configurable to suit the FA's hand-eye coordination. It uses a head-mounted display [101].
Telemedicine and tele-presence assistants	Real-time interaction; degradable service; visualize test results and past medical history; configurable interface	A smartphone-based application that enables physicians to interact in real-time through video call interface with patients from rural or remote areas or with mobility issues. Physicians can access patient's medical records, schedule same-day appointment, perform initial diagnosis [5].

Each application type poses some requirements as shown in the second column. The third column shows examples of existing systems that belong to these application types.

Table 3. HCAs Used for Training Patients and Professional Healthcare Providers

Application Type	Potential Application Requirement	Examples
Training patients: enhance cognitive functionalities	Interface to experts and therapists to configure the games; adaptive to user's preference and behavior; natural interface	A mixed-reality gaming platform that uses a multi-touch touch-screen tabletop interface for preserving cognitive functionality of the elderly, including memory, reasoning, selective attention, divided attention, and categorization. It supports single-player and multi-player games and personalized content for each player [32].
Training patients: cognitive orthotics	Real-time, autonomous, mobile perception; longitudinal assessment of patient's condition; episodic memory retrieval	A mixed-reality training platform and storyboard to help people with dementia and declining memory to interact through pen gestures, eye tracker, video camera, microphone, and bio-sensors. It provides real-time contextual suggestions to perform instrumental activities of daily living, send reminders on what to do next and how to do it and relates this to active memory training [128].
Training first responders	Realistic simulation of dynamic environment of large-scale emergency events; responsive and natural interaction; assessment mechanism for performance evaluation [41, 64, 65]	Training emergency responders to find the optimal resource allocation to complete different tasks in a distributed search and rescue mission using a mixed-reality location-based game [106]. The game simulates real-world disaster response scenarios and enables human-agent collaboration.
Virtual patients for physicians	Realistic user interface, facial expression, emotion; accurate response to pain and treatment; reconfigurable to critical use cases	Pediatric Hal, a wireless and tetherless pediatric patient simulator, simulates lifelike emotions (e.g., anger, ongoing pain, crying, anxious, yawning) through dynamic facial expressions, movement, and speech. It supports providers of all levels to develop skills needed to diagnose, communicate, and treat patients in many clinical areas. The simulator supports (i) real patient monitoring, including SpO ₂ , EKG, capnography, defibrillation, and (ii) emergency interventions, including surgical airway, needle decompression, and chest tube [121].
Training surgeons: robotic surgery	Realistic simulation of the patient's anatomy, allowing surgeons to practice on a particular curricula of tasks; providing evaluation metrics with respective scores after task completion	Virtual simulation-based training such as SimNow [49] and dv-Trainer [76] have provided objective and scalable methods for evaluating surgeon's skills and improving their training. The simulation platforms allow surgeons to familiarize with the surgical robot and improve their hand-eye coordination by maneuvering the manipulators as well as the endoscopic camera when performing tasks based on a specified curricula. The physics engine allows for realistic tool-tissue interaction that ranges from burning tissues, bleeding, and cutting.

Each of these application types poses some requirements as shown in the second column. The third column shows examples of existing systems that belong to these application types.

3.1 Interactivity

One of the fundamental features of a CA is to interact with users. Also, the CA can interact with the physical environment, services, processors, devices, and other CAs. The goal of such an interaction is to (i) sense the user goal, intent, or requirement, (ii) resolve ambiguity and incompleteness, and (iii) provide cognitive assistance. An HCA can interact through adaptive multimodal interfaces and visualization techniques. In this section, we discuss several major aspects of interactivity in existing HCAs that shape the design of the underlying systems, including the entities that existing




Fig. 2. Key features of healthcare cognitive assistants (HCAs): (i) interactive, (ii) context-aware, and (iii) adaptiveness. This figure also summarizes the outline of Section 3.

HCAs interact with, modes of interaction, realism of interaction, and nature of the interaction (i.e., proactive or reactive interaction).

3.1.1 Entity of Interaction. Most HCAs directly interact with only the target user. For instance, an HCA for robotic surgery interacts with the first assistant of the human surgeon [101]. However, based on the design requirement, often HCAs can interact with multiple users. For example, RoNA [47] is a humanoid, mobile robotic nursing assistant for lifting and moving patients and heavy objects inside a hospital to increase patient and nurse safety and operational efficiency. In addition to a nurse or physician (i.e., who needs assistance to move a patient), it interacts with a telepresence operator through a visual interface where the operator can see and control movements to ensure safe operation. Similarly, often HCAs for ADL support interaction with the patient and their primary caregiver [66, 93] or professional healthcare provider [104].

The design issues with multi-user interaction include, but are not limited to, data flow between multiple users, maintaining the privacy of each user, providing personalized and consistent interventions/feedback to each user. In addition to human users, a single HCA can interact with other HCAs or assistive services or devices that operate in the same environment. For example, an HCA for ADL support can interact with the personal assistant of the corresponding user (e.g., Alexa or Google Home) for weather updates and schedule the user's daily routine accordingly. The potential challenges of such interactions are discussed in Section 5.

3.1.2 Mode of Interaction. Based on our literature review, the modes of interaction for HCAs can be primarily categorized as: **verbal** and **nonverbal**. Verbal mode of interaction includes textual [5, 77, 137], audio [2, 71, 77], and video [23]. Such interactions may take multiple rounds of information exchange to understand user intent or requirements, resolve ambiguity, and address incompleteness. **Nonverbal interaction** includes interaction through a **haptic interface** [48, 62, 63, 92, 104], **visual interface** [33, 47], **augmented reality** [18, 45, 69, 79], **virtual reality** [23,

Table 4. Examples of Modes of Interaction as Found in Existing HCAs: the first column lists the modes of interaction

Mode of interaction	Goal/objective	Example HCAs
Verbal	Identify the intent and context of the user through natural language understanding, context detection, and dialogue management [2, 23, 77, 85, 108].	Ellie conducts virtual interviews with humans to automatically assess distress indicators [23]. The indicators are behaviors associated with depression, post-traumatic stress disorder, or anxiety. It uses four classifiers for natural language understanding, context detection, and response generation.
Message	Provide reminder, intervention, or alerts and ask questions through short messages on a laptop, smartphone or other smart display interface [66, 71, 83, 117, 128].	EMMA is a smartphone-based virtual personal assistant [33] that interacts with the user through recurring ecological momentary assessments for tracking their level of energy, positivity, and overall well-being.
		A navigational assistant for visually impaired people provides auditory feedback to alert a user about the size of the object (obstacle) and the object's distance from the user [117].
Tactile and Kinesthetic	Deliver kinesthetic, or tactile, feedback in a smart wearable device or sensor-embedded object [64, 84, 104, 133].	GuideCane [133] supports visually impaired individuals through haptic feedback to their cane to indicate their next step while walking (e.g., go straight, stop). The user interacts through a thumb-operated joystick to control direction.
Expression and Emotion	Interact with users by showing emotions or expressions expressed through a 2D or 3D virtual avatar or robot [9, 67, 110].	EmIR [110] is an empathetic social robot to provide cognitive support to the elderly for activities of daily life. It displays seven emotions to generate an empathetic response to users: angry, afraid, disgusted, happy, neutral, sad, and surprised.

The second column contains goals of each mode of interactions. The third column demonstrates an example of the corresponding mode of interaction as found in existing HCAs. For instance, the first row demonstrates the goals and an example of **verbal interaction**. Rows 2–4 show different types of **nonverbal interaction**. In addition to such individual modes of interaction, many HCAs support multimodal interaction combining more than one form of verbal or nonverbal interaction.

111, 112, 134], **mixed reality** [128], or **through sensor embedded objects** [32, 132, 133]. Some HCAs also include **olfactory** interaction to enhance simulation of a situation or to trigger particular memory [111]. A lot of HCAs perform multi-modal interaction and often combine both verbal and nonverbal interactions [9, 23, 47, 83, 110, 111, 134]. Often virtual interaction is carried out through a **2D or 3D avatar** [23, 79]. Examples of different modes of interactions are presented in Table 4. Additional details about the sensors and actuators used in different modalities of interaction are discussed in Section 4. More examples of different modes of interaction are available in Section 2 of online supplementary materials.

3.1.3 Natural Interaction. It is desired that the HCAs interactions with their intended users is **natural** and **meaning based**. Natural interaction can enhance the user acceptance and usability of HCAs. Several examples of natural interaction as found in existing HCAs are presented below.

- **Complex daily activities:** Kognichef [83], a cognitive assistant for complex cooking activities, provides a hands-free interface for browsing a recipe when the user's hands are occupied (detected through a built-in camera) to ensure natural and effective interaction. The user can pause, alter, and skip steps of a recipe and thus always stays in charge.

- **Physiotherapy:** KinoHaptics [104], a cognitive assistant for physical therapy and rehabilitation, provides haptic feedback when a user performs an unsafe movement during an in-home physical therapy session. Because haptic feedback requires less attention and focus from the user than audio or visual feedback, patients can better focus on their physical therapy. It also provides intuitive feedback by showing a progress bar and a real-time animation of movement. Thus, it is more engaging and easy-to-interpret as reported by the users who participated in a user-study to evaluate the usability of KinoHaptics.
- **Navigational assistance:** An emerging trend among cognitive assistants for visually impaired individuals is providing auditory [37, 71, 117] or haptic [63] feedback for navigation or performing an activity. Another navigational assistant for visual impairment uses a smartphone application paired with a mobile robot [84]. This mobile navigational assistant robot provides a natural interface in the smartphone app. Specifically, each sub-window in a phone screen is mapped to a predefined destination in the corresponding indoor environment of the user (e.g., cafeteria, rest-room). As a user finds sub-windows on a phone and then taps them, the app sends a command to the mobile robot to go that destination.
- **Enhance cognitive ability through games:** A mixed reality (MR)-based HCA for training elderly individuals through interactive games [32] provides natural interaction through tabletop MR platform. Tabletop interfaces mainly use touchscreens and multi-touch technologies. They do not require using a mouse or a touchpad.

However, several existing HCAs lack natural interaction. For example, consider the rehabilitation HCA for post-stroke hand rehabilitation using spatial augmented reality [45] that simulates common hand movements using AR, including reaching, wrist-tilting, pointing, and grasping [45]. It helps a patient practice these hand movements. However, it is not clear how the assistant responds if the user makes any hand movement other than the supported movements. It can result in an interrupted and unsafe user experience. Sainarayanan et al. [117] present a blind navigation assistant that uses sonification (recognition of an object from sound). However, the user requires a significant amount of training before using the system to interpret the feedback from the system that comes through head-mounted gear. Also, the system does not deal with moving objects. Some HCAs provide snooze option for notification [103]. Such systems should ensure that snooze option does not annoy users or cause any discomfort.

3.1.4 Proactive or Reactive Interaction. Most of the HCAs are reactive in terms of initiating the interaction. The proactive assistive services are often referred to as “detect-assistant” and use a two-step approach [104, 105, 122, 128]. First, the assistant detects the deficit observed in an abnormal behavior or activity and then it proposes suitable assistance [122]. Here, we present a few examples of HCAs that support pro-active interaction. *Kinohaptics*, an HCA for physical therapy and post-injury rehabilitation, supports proactive interaction. It tracks movement and alerts users as soon as unsafe movement is detected [104]. In the surgical robot assistant *da Vinci*, proactive monitoring is implemented to prevent device malfunctions [105] from affecting the outcome of safety-critical events. *Kognit* provides proactive feedback for instrumental activities of daily living to remind the users of an activity and alert them if they make any mistake while doing an activity [128]. *FindIt* provides proactive reminders to its users if they leave behind any critical device, including phone or keys, while going outside [20]. Another cognitive assistant for monitoring and detecting early dementia initiates diagnostic conversations with the user proactively [2].

Existing HCAs vary drastically in terms of (i) entities of interaction, (ii) modes of interaction, and (iii) the nature of the interaction (i.e., proactive or reactive interaction) as discussed above. Although natural interaction is critical for usability and effectiveness of HCAs [13, 26, 50, 51, 61], several HCAs overlook this aspect or demonstrate low degree of natural interaction [45, 103, 117]. The

next generation HCAs should consider this issue from the design level and take an interdisciplinary approach to address it. We also cover the challenges regarding interaction among multiple HCAs in Section 5.2. Another relevant area of research emerging from the human-computer interaction research community is how users (i.e., patients, individuals, caregivers, and professional healthcare providers) interact with HCAs [13, 32, 33, 35, 40, 77, 93, 104, 110, 127].

3.2 Context-awareness

Context-awareness refers to the feature of a system that allows the system to react differently according to different contexts. The system usually has some underlying representation of contexts, and it learns the context automatically, semi-automatically, or from user feedback. A healthcare CA should be able to **understand, identify, and extract contextual** elements from the interaction with a user and environment. The set of contexts for HCAs can be categorized into four classes: temporal, spatial, the user or personal, and situational contexts [86].

HCAs are often designed to be **temporally** context-aware, i.e., they identify and respond according to time of the day, day of the week, or some other predefined time slots. HCAs often provide location-aware interventions and use **spatial context** for inferring user state. For HCAs, the **user context** includes a user's physiological, psychological, behavioral, and medical context. HCAs can provide interventions that are customized to one or more user contexts.

- Physiological context refers to a user's age, height, weight, and other physiological factors.
- Psychological context refers to a user's emotion, mood, personality, level of positivity, and other psychological factors.
- Behavioral context encompasses a user's behavior, action, predefined priority or preferences, level of skills, and professional training and certification.²
- Medical context refers to a user's past medical history, present medical condition, symptoms, diagnosis, medications, genetic profile, family history, and similar medical factors.

The **situational context** includes environmental context, process context, and operational context. Different situations for healthcare are presented in Figure 1. For instance, this includes home healthcare, remote monitoring, ICU, surgery, telemedicine, or emergency response. In addition to temporal, spatial, and user contexts, often HCAs are aware of situational context in terms of ongoing process, operations, or pre-defined protocols. Table 5 presents examples of context-awareness of existing HCAs. Additional review of context-awareness of existing HCAs can be found in Section 3 of online supplementary materials.

Identifying and considering a user's context are essential for HCAs. However, it also raises concern regarding privacy, security, safety, and confidentiality. Researchers and developers should address these challenges while designing and developing HCAs.

3.3 Adaptiveness

An adaptive system³ refers to a system that changes its behavior in response to its environment. A healthcare cognitive assistant should be adaptive so it can accommodate the dynamic behavior

²For professional care providers, including physicians, nurses, and emergency responders.

³It should be noted that context-awareness and adaptiveness are often used interchangeably in some of the existing literature. However, we distinguish between these features as follows: A system can react differently under the same context to provide a more adaptive response. To illustrate the point, a context-aware cognitive assistant for ADL support can generate specific reminders for an activity based on spatial, temporal, or situational context, e.g., suggesting outdoor exercise when the weather is nice, and the user is physically inactive for a long period [103]. However, the system can be adaptive if it adapts its reminder based on the user's response, e.g., when the user declines the reminder to perform outdoor exercise repeatedly, the system adapts to the user behavior and suggests an alternative exercise.

Table 5. Examples of Context-awareness of Existing HCAs: The first and second columns contain the names of different types of contexts and their sub-categories, respectively

Name of Context	Type of Context	Example HCAs
Temporal	Duration	Step Up Life, generates notifications and reminders when the user is physically inactive for a predefined period. The user can also set “no reminder” timeslots so the assistant does not generate any exercise reminder even if there is a long span of inactivity [103].
	Time of the day/day of the week	EMMA, an empathetic virtual assistant for well-being, uses temporal context (time of the day and the day of the week), in addition to other contexts, to infer mood [33].
Spatial	Inside/ outside	Gabriel, a Google Glass-based wearable assistant for individuals with cognitive decline, performs location-aware sensor control [40] and user-activity recognition. Such as, if the user falls asleep at home, it turns off the built-in camera to save battery life. It awakens users if they fall asleep while traveling in public transport.
	Exact and relative location	EMMA uses spatial context to infer user’s mood [33]. It uses the user’s exact and relative location (e.g., distance from work/home) as spatial features to predict the user’s mood.
	Landmarks	A navigational assistant for people with cognitive impairments uses augmented reality [43]. Instead of street names/distance, it focuses on user-friendly routes based on user-known landmarks.
User	Physiological	Quro, a conversational assistant to support symptom checking by patients, is aware of a subset of user’s physiological and medical contexts, e.g., gender, age, smoking history, and heart problems [34].
	Psychological	EMMA recommends activities according to the user’s mood to promote the emotional well-being of the user [33].
	Behavioral	Ellie, a virtual human interview agent, generates appropriate real-time nonverbal interaction/behavior based on the conversational context and user’s facial expression and gestures [23], e.g., facial expression, eye and head movement, blink, and gaze.
	Medical	Babylon, a chatbot to support self-diagnosis for patients, generates its response according to the user’s medical context [5].
	Multiple	ODVIC, a multimodal conversational assistant to deliver evidence-based interventions for behavior change, provides interventions that are customized to the user’s profile, past behavior, mood, and recent conversation [67]. Thus, it combines psychological, physiological, and behavioral contexts to provide customized interventions.
Situational	Environmental	A navigational assistant that supports path planning is aware of environmental and spatial context [88].
	Process	A surgical assistant monitors the current surgical task of a surgical procedure [141] and conducts context-specific safety checks.
	Operational	CognitiveEMS, a cognitive assistant for emergency response, suggests interventions to EMS responders that are specific to the context of standard EMS protocols [126].

For instance, user context can be physiological, psychological, behavioral, medical, or a combination of any of these user contexts. The third column presents examples of each of these different types of context.

of its environment (including the physical environment or ambiance of the system) and the user’s goals, needs, requirements, actions, and behaviors. While generating dynamic response is one of the essential characteristics of adaptive systems, additional characteristics include, but are not limited to, (i) resolving ambiguity, (ii) tolerating unpredictability, and (iii) learning from experience. The degree of adaptiveness can vary across systems. Based on our review of existing HCAs, there are four dimensions of adaptiveness.

Table 6. Examples of Adaptiveness of Existing HCAs: The first and second columns contain the dimension and degrees of adaptiveness, respectively

Dimension of adaptiveness	Degree of adaptiveness	Example HCAs
User's action	Velocity of movement	Pearl, a mobile robotic assistant acts as a navigational guide and ADL assistant for the elderly and adapts its velocity according to the user's velocity [93]. Pearl estimates the user's velocity and adjusts its speed accordingly.
User's behavior	Interventions performed by emergency medical services (EMS) responders	CognitiveEMS is a cognitive assistant to provide real-time decision support to EMS responders. It monitors a patient's condition and user's actions (i.e., the actions performed by an EMS responder to manage the emergency scene) and suggests intervention according to EMS protocols [94, 98, 99, 126].
User's need	Errors made by the user	A cognitive assistant to support visually challenged people for meal preparation provides vocal instructions as the user proceeds with preparing their meal [36]. It suggests adaptive interventions that are customized to the types of error the user most frequently makes (e.g., initiation, planning, attention, and memory deficit).
Environment	Network failure and computational resources	Gabriel, a Google Glass-based wearable assistant for ADL support is adaptive to network failures and unavailability of remote tiers [40]. It performs computation on server hardware when the network is available to save the device energy and increase the service speed. In case of network failure, the system offloads computation to a fallback device e.g., the user's smartphone.

The third column presents example of such adaptiveness. For instance, the first row describes Pearl, a mobile robotic assistant that adapts its mobility according to the user's action, i.e., their velocity.

- **User's action:** HCAs can be adaptive to users' actions, i.e., the system monitors the user's activities and adjusts intervention suggestions accordingly [8, 18, 83, 93, 94, 98, 99, 126].
- **User's behavior:** HCAs are often designed to be adaptive to users' verbal and nonverbal behaviors [23, 26, 88, 93].
- **User's need:** This refers to the feature of an HCA where the HCA adapts its response to a user's (e.g., patients, caregivers) cognitive need, as the user's condition (e.g., disease, stress level, psychological state) changes over time [8, 36, 104, 111].
- **Environment:** HCAs are designed to respond dynamically with the change in the surrounding environment. For example, many HCAs developed for navigational assistance are adaptive to the surrounding environment [69, 84, 133], e.g., obstacles, visibility, and illumination.

Table 6 demonstrates examples of different dimensions and degrees of adaptiveness as found in existing HCAs. Additional review of adaptiveness in existing HCAs can be found in Section 4 of online supplementary materials.

3.4 Limiting Features of Existing HCAs

- **Context-aware and adaptive assistance require memory:** The assistants need to *remember* past interactions and then respond according to the specific application in the current circumstances. Besides, the assistant should be able to distinguish anomalies from a new behavior pattern to provide accurate and satisfactory cognitive assistance. However, these two issues are overlooked in most existing HCAs.
- **Uncertainty and ambiguity:** HCAs should be adaptive to ambiguity and uncertainty, i.e., they should be able to "resolve ambiguity and tolerate unpredictability" [85]. They should define a problem specifically by asking additional questions or utilizing additional input

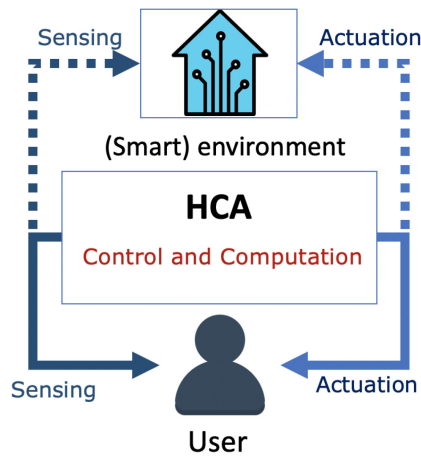


Fig. 3. Cyber-physical components of healthcare cognitive assistants (HCAs): (i) sensing, (ii) actuation, and (iii) control and computation. HCAs sense overall user’s state (e.g., user’s need, behavior, context) and actuate on the user to perform intervention. Optionally, some HCAs might sense and actuate on the surrounding environment of the user. The surrounding environment could be *smart*, e.g., smart home, virtual environment, or augmented reality–based environment.

sources to resolve ambiguity and incompleteness. Although the spectrum of resolving ambiguity and unpredictability can vary across different applications, the majority of the current HCAs do not address ambiguity and uncertainty.

- **Preserving privacy, confidentiality, and security:** Most of the HCAs we reviewed overlook privacy, confidentiality, or security issues. However, one of the most significant differences of HCA comparing to other cyber-physical systems is the actuator of the system, i.e., human beings. As we discussed in previous sections, HCAs are context-aware and adaptive, which, at the same time, also means that they require more personable information from humans and take *actions* on humans. For example, a smart reminder system gives reminders to patients based on their daily living habits, which also may detect a person’s other activities; a navigation system for visual impairment people knows where and when they go most of the time. Protecting this private information from leaking, or malicious usage, is a significant challenge for the HCAs. Also, malicious attacks can result in adverse or even fatal outcomes for users (error in robotic surgery, wrong or unsafe medication dosage) [139]. Medical devices and HCAs might be hacked and can cause harm or demand ransomware [29].

4 PHYSICAL AND CYBER COMPONENTS FOR COGNITIVE ASSISTANCE IN HEALTHCARE

In this section, we review the cyber-physical components (CPS) of HCAs, including (i) sensing modality (i.e., detection/perception), (ii) actuation (i.e., response and intervention), and (iii) control and computation. These CPS components are essential for HCAs, as they enable the key features of HCAs. For instance, the proper sensing and actuation modules enable an HCA to be *interactive*; or through sensing and control/computation an HCA can detect and identify different contexts. The interaction between a user and an HCA through different CPS components is shown in Figure 3. It also shows that often an HCA can also interact with the user’s surrounding environment (e.g., activity recognition in a smart home, sensing and actuating on a user’s virtual environment).

4.1 Sensing Modality: Detection and Perception

This section presents an overview of the sensing and perception technologies used in existing HCAs. The sensors can be roughly categorized into six classes:

- (1) Primitive sensors: These include, but are not limited to, PIR motion detectors, temperature sensors, contact sensors, light sensors, and humidity sensors.
- (2) Physiological sensors: While the primitive sensors obtain the environmental states, physiological sensors (e.g., pulse oximeter, blood glucose monitor, heart rate sensors, EKG, blood pressure, and skin conductance sensors) are applied to measure the patients' physiological states. Physiological sensors that are smaller in size and provide accurate sensing, wireless communication, and user-friendly interface are more suitable for usage in HCAs.
- (3) Acoustic and Ultrasonic Sensors: These are often used for environmental sensing and obstacle detection to support navigation of visually challenged individuals. In addition, several conversational HCAs use acoustic sensors (e.g., microphone array or a built-in microphone in smart devices) to recognize user's speech to detect emotions or semantics of their speech.
- (4) RGB Camera: These are often used in HCAs for (i) sensing the surrounding of a visually challenged individual (e.g., obstacles, object detection, people identification, localization) and support navigational assistance, (ii) detecting nonverbal interaction (e.g., facial expression, gaze, emotion, empathy), and (iii) fine-grained activity recognition.
- (5) RGB-D and Depth sensors: Depth sensors provide a more privacy-preserving approach for detecting objects and sensing the surroundings of an individual. Thus, they are an alternative to RGB cameras for navigational assistance, localization, and object detection. In addition, HCAs often use depth sensors for tracking body gesture and pose to support real-time monitoring of psycho-motor exercises and physiotherapy sessions.
- (6) GPS and Bluetooth low energy (BLE) beacons: These sensors are used in mobile HCAs to support people or object tracking, localization, and navigational assistance.

Table 7 demonstrates examples of usage of these six classes of sensors in existing HCAs. Additional examples of different sensing techniques used in existing HCAs are available in Section 5 of online supplementary materials. We also review usage of multimodal sensing in existing HCAs.

4.1.1 Multimodal Sensing. Several HCAs rely on multimodal sensing to provide multimodal interaction [83, 93], assist users in multiple cognitive functions [93, 135], act as a robotic surgery assistant [123, 138], and support augmented reality (AR), virtual reality (VR), or mixed reality (MR) interfaces [109]. For instance, the relative location and motion of the user's head needs to be determined to accurately adjust the projected image or holographs for headset-based VR or AR applications. It is achieved using an Inertial Measurement Unit (IMU) that combines an accelerometer, a gyroscope, and a magnetometer. By combining the relative positions information from the three sensors, the user's head position and movement are accurately tracked.

For navigational assistance. Ribeiro et al. [109] propose an auditory augmented reality system, where the system integrates acoustic virtual objects into the real-world to assist people with visual impairment. The goal is to allow the innate ability of individuals of sound source identification and source separation to determine nearby objects. The subject wears a helmet that is instrumented with an RGB-camera, an IMU, and a headphone. A 3D gyroscope is used to track the head, and a 3D accelerometer is used to infer the floor plane by estimating the gravity vector. RGB-D stream is used to infer high-level features of interest (face detection and recognition, floor mapping for navigation, and plane detection). The detected high-level features are conveyed to the user by

Table 7. Examples of Six Classes of Sensors as Found in Existing HCAs: the first column lists the classes of sensors. The second column contains the set of relevant tasks of each class of sensors. The third column demonstrates an example usage of relevant sensors in existing HCAs. In addition to individual modes of interaction, many HCAs support multimodal sensing as discussed in Section 4.1.1.

Types of sensors	Relevant tasks/usage	Example HCA
Primitive	Occupancy detection, event detection, activity recognition and monitoring [8, 36, 110]	A robot assistant uses environmental sensors such as CO_2 , humidity, temperature, propane and butane to trigger alarms for potentially risky events, e.g., leakage of propane gas, high level of CO_2 [110].
Physiological	Sense and measure physiological state, e.g., heart rate, blood pressure, blood glucose [9, 128, 130, 134]	A smartphone-based conversational assistant to promote self-care in people with atrial fibrillation (AF) [9], uses AliveCor Kardia mobile heart rhythm monitor, a sensor- monitor validated for AF screening. The device is attached to a smartphone and transmits data via Bluetooth.
		A virtual coach to improve exercise performance is proposed in Reference [134] that relies on a VR bike-frame and several physiological sensors to capture the user's brain activity and other vitals while using the bike. It uses Electroencephalogram to capture brain activity. It also captures heart rate, respiration rate, bike pedal rate, and power exerted by the user on the bike.
Acoustic and ultrasound	Speech recognition [9, 23, 83, 110, 111], obstacle detection for blind navigation [95, 125, 133], surrounding environment sensing [54]	GuideCane equips a cane with ultrasonic sensors for obstacle detection and to help visually challenged people to go around the obstacles [133].
RGB camera	Navigation [69, 84, 117], AD support [18, 75, 135], and nonverbal interaction [20, 110]	A mobile robot for visually challenged people uses on-board camera to navigate by tracking pre-deployed markers (or stickers) on a floor [84].
		A robot is equipped with a camera to capture facial images that are used for people identification and emotion classification [110].
RGB-D and depth sensors	Gesture and pose detection for psychomotor exercise and physiotherapy [11, 16, 17, 39, 57, 60, 79]	A cognitive assistant for remote physiotherapy monitors the patient's exercise session through Kinect [104] at home. It keeps track of movements of selected joints to provide haptic feedback through the patient's armband.
GPS	Navigation, tracking, and localization [3, 30, 71, 88, 103, 118]	Step Up Life, uses smartphone's GPS sensor, CELL ID, and Wi-Fi details for tracking user's location [103]. It tracks a user's activities using the phone's accelerometer and magnetometer to generate exercise reminders.

using pre-recorded wave files, a text to speech synthesizer after spatializing each sound. Lee et al. [63] mount an RGB-D sensor and IMU sensor into a pair of glasses instead of mounting them above a helmet [109] to build a navigational assistant. A smartphone is used to specify the destination.

Supporting cognitive decline. Vorobieva et al. develop a robotic system to assist people who are losing their autonomy, e.g., disabled, elderly [135]. The system has a gripper with a stereo camera (for visual servoing or vision-based robot control), pressure sensors, and optical barrier to detect when an object is in the gripper. The user can request the robot to find an object from a predefined list of objects. The robot then navigates through the environment to pick up the object and bring it to the user. The goal of the system is to stimulate the cognitive state of the user by playing games.

Supporting multiple functionalities and complex activities. Pollack et al. develop a robotic assistant for cognitive orthotic functions (i.e., providing context-aware and adaptive reminders for activities of daily living) and safe navigation for the elderly [66, 93]. The system utilizes SICK laser range finders and sonar sensors for navigation, microphones for user’s speech recognition, and touchscreen display to detect user needs. It utilizes a camera data stream for face detection, activity recognition, and object tracking and detection for navigation support. It deploys multimodal sensing for navigation by combining sensor data streams corresponding to user localization, object detection, and tracking. KogniChef [83] is a cognitive cooking assistant for preparing a meal. It uses a Kinect RGB-D sensor and a thermal camera to perform object detection, tracking, and grasp detection. A scale is used in addition to cameras to estimate fill-level for pouring ingredients. A microphone array is used for speech recognition, and a speaker is used to provide feedback.

For robotic surgery. Shademan et al. present “Smart Tissue Autonomous Robot (STAR)” for automating soft tissue surgical tasks and providing a collaborative platform for decision-making and execution of surgical tasks to surgeons [123]. The STAR system utilizes 3D plenoptic vision, near-infrared fluorescent (NIRF) imaging, sub-millimeter positioning, actuated surgical tools, and force sensing to construct and execute surgical tasks. The combination of “NIRF technology and 3D quantitative plenoptic imaging” addresses the problems of occlusion and target tissue recognition by observing “luminescent NIRF markers” [123].

4.2 Actuation Modality: Response and Interventions

Based on the type of tasks digital assistants perform, they can be categorized into three classes [70]: (i) personal assistant or butler that performs a task on behalf of the user, (ii) cognitive orthotic that provides adaptive and contextual feedback and reminders to people with cognitive impairment or decline, and (iii) mentor or coach. Based on our review, we found that most of the current HCAs mostly fall in the second category. However, this categorization does not consider the set of HCAs that enhances cognitive functionalities (e.g., HCAs for training healthcare providers or providing decision support to them). So, instead of following the taxonomy mentioned above, we present the different modalities of actuation in existing HCAs in this section.

4.2.1 Visual. A dashboard or display-based system is one of the most common forms of actuation and often provides visual guidance to users to perform a task properly. Pearl utilizes a touch-sensitive graphical display for ADL reminder and navigational instructions [93]. KogniChef [93], an ADL assistant specifically designed for complex cooking tasks, uses a display to inform a user about the current state of cooking through structured visual information to reduce cognitive load. A physical rehabilitation HCA provides visual feedback on a user’s specific physical movement during a physiotherapy session [104] to enable the user to visualize their movements. The STAR system [123] provides suture automation software that displays a geometrically optimized suture plan in real-time. If the placement of the suturing tool is problematic, the surgeon has the option of intervening and making adjustments.

Visual actuation often consists of contextual and adaptive textual interventions. EmIR [110], a cognitive assistant for emotional well-being uses textual messages to provide contextual reminders and recommendations and persuade the users in activities to lift their emotional state. Another cognitive orthotic HCA developed for meal preparation alerts the user when a missing or wrong step is detected [8]. It sends the message with instructions using not only text, but also figures to better explain the missing/wrong cooking steps to users.

4.2.2 Audio. Conversational HCAs often implement verbal communication through audio [2, 9, 23, 83, 110, 111] with an underlying text-to-speech conversion and transcription module, e.g.,

Google speech, IBM, or CMU sphinx. In addition to such verbal communication, audio mode is often used for nonverbal interaction [3, 69, 88, 109, 118] and cognitive orthotics [83, 93]. The most common form of such actuation is found in navigational assistants where the navigation system gives the user step-by-step instructions using earphones [3, 69, 88, 118]. As an example of cognitive orthotics, Pearl [93] utilizes built-in speakers for speech synthesis to answer user queries and provide ADL reminders.

4.2.3 Haptic. Haptic feedback can be kinesthetic or tactile or a combination of both. Kinesthetic feedback refers to the haptic sensation that is felt by the muscles, joints, or tendons. Usually, kinesthetic actuation includes weight and stretch. However, tactile actuation refers to the haptic sensation felt by the surface of our body and includes vibration, pressure, and texture. Most of the HCAs reviewed in this article that use haptic feedback use tactile feedback. Haptic feedback can address the issue of accessibility to some extent, since it can be more desired than audio and visual feedback for people with declining auditory and visual perception, respectively. It is also useful for implementing a hands-free interface, since the user may be engaged in some activity (e.g., physical exercise or therapy) and cannot hold any device.

Khademi et al. develop an augmented reality rehabilitation system that uses haptic feedback to enable patients with stroke to practice their hand and arm movements without the presence of a physical therapist [56]. Phamduy et al. develop a novel belt to provide tactile stimulation in the abdomen for situational awareness and obstacle avoidance by integrating micro fiber composites into the belt [92]. Nguyen et al. build a way-finding system, which is deployed on a mobile robot that a user would follow to navigate [84]. The user would use a smartphone to select a destination from a predefined set of destinations. While a user follows the robot, the feedback from the robot is encoded as tactile vibrations of the smartphone to notify the user. There are four types of vibrations to suggest “turn left,” “turn right,” “go straight,” and “stop.” The navigational assistants presented in References [62, 63] use tactile feedback through a vest. There are four vibration motors integrated into a vest that is controlled wirelessly to provide four navigation cues: straight (“no tactile sensors on”), stop and scan (“all tactile sensors on”), turn left (“top-left sensor on”), and turn right (“top-right sensor on”). The authors argue that the vest-type interface would reduce the cognitive burden of the user compared to audio-based navigation feedback.

KinoHaptics, an HCA for self-care and post-surgery rehabilitation, monitors the patient’s physiotherapy session and provides haptic feedback through the patient’s armband [104] in real-time. The vibro-haptic feedback is generated to make sure the user does not overdo or under-perform an exercise suggested for physiotherapy. The armband contains an array of vibration motors, and it connects to the feedback-generating server machine via Bluetooth connection. Step Up Life uses haptic feedback for cognitive orthotics, specifically for physical activity and movement. If it observes prolonged inactivity of a user, it notifies the user along with an exercise suggestion by generating haptic vibrations using the cell phone vibration motor. The duration of the haptic vibration depends on the number of times the user has snoozed a notification [103].

4.2.4 Multimodal Actuation. Several existing HCAs perform multimodal actuation to support multiple cognitive functionalities and to provide natural, realistic interaction, often through AR, VR, or MR interface.

Weede et al. presented a surgical robotic assistant that provides two interventions: (i) knowledge-based camera guidance that provides an optimal view of the surgical workspace and (ii) a port and setup planning to provide an optimal position to insert the endoscope and the two end-effectors into the patient’s body [138]. Rizzo et al. developed a virtual assistant for psychotherapy that simulates traumatic events based on a patient’s description [111]. It provides general navigation for driving in the simulated scenario using a standard gamepad. It also




Fig. 4. An emerging trend in HCAs is using AR, VR, and MR, as demonstrated in the HCAs above. (A) Shows a VR-based exposure therapy platform to simulate trauma-inducing events based on narration of people suffering from combat-related PTSD [111]. It supports simulating general navigation for driving, dismantled foot patrol, holding mock M4 gun, and generating audio, vibrotactile, and olfactory stimuli. (B) Presents a use case of medication management and tracking in Kognit, an MR-based assistant for elderly individuals [128]. (C) Shows *ElderGames*, an MR-based game to improve cognitive functions of elderly individuals [32]. It provides natural interaction through multi-touch technology where multiple players can play together using pens on the table top. Here, real objects (i.e., pens) are used to interact with virtual ones (i.e., virtual objects displayed on the touch-sensitive and interactive table top). (D) Demonstrates configurable visualization of a stereo endoscopy in ARssist [101], an HCA for real-time cognitive support for a first assistant (FA) in a robotic surgery: In the left figure, the endoscopy is shown in a virtual display to enable the FA to visualize both the surgical field and the endoscopy with minimal head rotation. In the right figure, the endoscopy and the instruments are rendered inside the patient's body. This enables the FA to intuitively operate instruments into the endoscopic field-of-view, even with an inconvenient docking configuration of the robotic arms. (E) Presents the “*da Vinci Si surgeon's console*” and the “*Skills Simulator backpack*” that uses VR for training and evaluating robot-assisted surgery skills [38]. [AR: *Augmented Reality*, VR: *Virtual Reality*, MR: *Mixed Reality*]

provides the option to simulate the context of dismantled foot patrol and a user-held mock M4 gun through a thumb mouse attachment. It provides audio, vibrotactile, and olfactory stimuli to users for realistic simulation of the traumatic event. It is shown in Figure 4(A).

ARCoach [18] is a task-reminder system to assist individuals with cognitive impairments that provides cues to complete tasks, detects incorrect steps on-the-fly, and helps to correct a task. Unlike other approaches that require users to match picture cues with reality, ARCoach overlays artificial information on real-world images captured through a webcam. The overlaid information can be in texts, sounds, pictures, or a combination of these. CARA [69] is a cognitive augmented reality assistant for the blind. Using head-mounted Microsoft HoloLens, CARA uses onboard video and infrared sensors to construct a 3D map of the surrounding space. Then each object in the scene

generates a voice that comes from the location of the object. As the object gets closer to the user, the voice’s pitch increases. It helps blind subjects to avoid obstacles, perform navigation, scene interpretation, and “formation and recall of spatial memories.”

Parsons et al. [91] use VR to train people with autistic spectrum disorders to enhance their social skills. The key idea is to provide a safe virtual environment to practice social events by performing role-play in different contexts. Kognit aims to help dementia patients by leveraging mixed reality [128]. The authors describe their approach as therapeutic, which enhances the cognitive abilities of dementia patients. Kognit produces new episodic memory visualizations by allowing physical and virtual objects to co-exist and interact with each other. An example use case is shown in Figure 4(B), where mixed reality is used to monitor the medication-taking behavior of an elderly person.

4.3 Control and Computation

In this section, we review different aspects of the control and computation component of existing HCAs. It should be noted that some computational models relevant to sensing or perception and actuation have already been discussed in the previous sections and subsections. Here, we mention additional interesting insights regarding control and computational models.

4.3.1 Underlying Control and Computational Model.

Data-driven Model. HCAs often use off-the-shelf trained machine learning models for different sub-tasks. To name a few, navigational assistants require object detection and scene interpretation [7, 20, 120]; conversational assistants require detecting facial expression [14, 110] and emotion [110], understanding natural language [2, 23, 28, 112] and dialogue management [2, 9, 23]; companion robots and ADL support HCAs require activity recognition [93] and person identification [110].

Most of the existing HCAs use off-the-shelf trained machine learning models for this. For instance, Jaime et al. [110] use a web service for person identification and emotion recognition from images due to the limited computation capacity of the robot assistant. The robot assistant captures facial images and sends the images to the web service where all images are processed, and then results are returned to the robot in real-time. A cognitive orthotic HCA for cooking [8] uses separate models for different underlying components, including activity recognition, human-machine interaction events, behavior or activity errors detection, errors characterization, and diagnosis regulation. Then the outputs from these models are integrated to monitor, detect, and assist in cooking. Another approach is developing application-specific data-driven models separately and integrating them into the HCAs. Consider Pearl [93], an HCA for the elderly that provides real-time navigational guidance and adaptive, context-aware reminders for ADL. Pearl uses a quantitative temporal Bayes net for activity modeling and inference. It adopts a hierarchical variant of a “partially observable Markov decision process” (POMDP) as the control architecture to mitigate the significant level of noise in the assistant’s perception, which is originated from the laser range-finder sensors and user input from microphone and touchscreen.

The models are often developed and adapted on-the-fly through user training and longitudinal usage, resulting in personalized models. NavCog [3] uses reinforcement learning to generate a step-by-step instruction personalized to the mobility skill of a user. It builds a user-specific behavior model to ensure successful navigation. However, when data is not enough for a new user, it uses transfer learning techniques to apply other users’ data to the new model. Mattos et al. develop a speaker-independent, language-independent model to assist hearing impairment patients to read lips using Generative Adversarial Networks (GANs) to learn mouth pictures [74]. It uses synthetic 3D models for training, and videos collected from real subjects for testing.

Knowledge-driven Model. Often HCAs use a knowledge-driven approach for control and computation. In CognitiveEMS [126], one of the proposed approaches for real-time decision support

through EMS protocol-specific intervention suggestion is modeling the EMS protocols using a Behavior Tree. It is a computational model for knowledge representation that uses a dynamic data structure that adapts to an ongoing process or incoming information flow. In KinoHaptics [104] a patient's physiotherapist develops a specific personalized exercise program. The exercise program contains critical information, including what should be the angle of elevation of a joint during physiotherapy sessions, how frequently it should be moved, and what should be the duration of an exercise session. This exercise configuration file is loaded on a local desktop or laptop at the patient's home. Then the Kinect-based monitoring system (that is integrated with the local machine) monitors the user's movement during a physiotherapy session and generates vibro-haptic feedback to notify the user if they are overdoing or under-performing a movement that involves the selected joint [104]. Weede et al. present a prototype of a cognitive system for minimally invasive surgery that leverages knowledge acquired on the workflow of surgical interventions through collecting trajectories of different surgical contexts [138]. Their implementation includes several control modes that can be called upon, depending on the surgical context, with the modes being teleoperation, hands-on mode, and autonomic camera guidance.

Knowledge can also be presented as a rule base. Emma [33], a virtual assistant to promote psychological and mental wellness, utilizes a rule base to generate (i) predefined response and (ii) intervention suggestions on emotionally appropriate micro-activities based on frequent ecological momentary assessment (EMA) survey collected from a user. The rule base is generated by professional care providers. Another virtual assistant for mental health [110] presents arguments to persuade users for an intervention based on analogy, popular practice, or expert opinion.

Control models are often developed based on the application requirement, safety requirement, and other constraints and thus embed domain knowledge. A humanoid, mobile robotic nursing assistant for lifting and moving patients inside a hospital achieves semi-autonomous and autonomous functionalities through a behavior-driven control model [47]. The control model is adjusted to ensure user safety (both patients and nurses) and operational efficiency.

Hybrid Model. HCAs that provide treatment-related suggestions often combine domain knowledge models with data-driven approaches. IBM Watson for oncology [137] combines both data-driven and knowledge-driven models to provide customized decision support to oncologists for diagnosis and treatment plan selection. Specifically, it provides an interactive, context-aware interface for information visualization and summarization using natural language inference and knowledge integration. Upon logging into the system, oncologists can view and browse through the relevant medical information for each of their patients, including but not limited to, medical history, family history, test results, suggested treatment options, and knowledge curated from recent and historical cases similar to the current patient. It generates treatment suggestions based on a model trained on prior data collected from Memorial Sloan Kettering Hospital oncology records. In addition, it combines knowledge extracted from over 300 medical journals and 200 textbooks and rationales from leading oncologists. It also shows relevant statistics from the curated literature for different treatment options. In cognitiveEMS [94, 98, 126] data-driven language models and distributional semantic models are used to extract safety-critical concepts that are relevant to standard emergency medical service (EMS) protocols in real-time from the spoken language collected at an emergency scene. In addition to that, domain knowledge from standard EMS protocols are integrated using behavior tree data structure to provide effective and safe intervention suggestions to the responders.

In an empathetic virtual assistant for changing drinking behavior of individuals, real-time data-driven models are combined with behavioral models from psychology and other domains [67]. The system is controlled based on (i) the perception of user state sensed from real-time text and video

datastreams and (ii) established psychometric instruments. Empathic reactions with intervention suggestions are generated based on predefined rule-base that captures the domain knowledge of experts. For instance, the behavior change assistant decides its facial expression, head movements, eyebrow expressions, and complex verbal reflections based on a user's perceived state and knowledge of predefined behavior protocols.

4.3.2 Device-level Computing. By following the four-tier computing model [119], where Tier-1 represents the cloud (e.g., data centers), Tier-2 represents cloudlets (e.g., high-end laptops, desktop PCs), Tier-3 represents embedded devices (e.g., smartphones, wearables), and Tier-4 represents energy-harvesting devices (e.g., RFID tags), a trend in existing HCAs is the usage of Tier-2 and Tier-3 level local computing to provide pervasive cognitive assistance to users even with low or no network connectivity and device constraints. Tian et al. present a navigational assistant for visually challenged people that requires users to carry a mini laptop that performs the entire computation locally [132]. It processes RGB-D sensor data mounted in the belt of a user to detect staircases and pedestrian crosswalks by using a Hough transformation and an SVM classifier to enable blind navigation. Similarly, local processing is performed [63, 117] in real-time, where a user needs to carry an entire computational unit. Specifically, a blind individual needs to carry headgear containing a digital video camera and wear a smart vest that contains the processing equipment, a Micro box PC-300 chassis [117]. The vest also contains rechargeable batteries.

In addition to navigational assistants, device-level computing (Tier-2, Tier-3) is also preferred in other pervasive, mobile HCAs, especially in HCAs for cognitive orthotics. A visually impaired user needs to follow a moving robot for indoor navigation that performs computation locally [84]. There is an offline phase to map the environment and travel routes. Neumann et al. propose KogniChef, a cognitive cooking assistant, where all the computations are performed on a "6-core Linux machine" [83]. González-Ortega et al. propose a real-time system that runs locally on a PC with an attached Kinect and asks the user sitting in front of it to perform different psychomotor exercises to assess neuropsychiatric disorders and mental illnesses [39].

Pollack et al. propose Pearl, a cognitive assistant for navigational guidance and adaptive reminders for ADL, that contains a differential drive system and two onboard Pentium PCs as the backend server [93]. It also uses Wi-Fi to access the internet for responding to user's query regarding weather and date/time. However, Pearl can provide support for the core cognitive functions (i.e., navigational guidance and adaptive ADL reminder) without the Wi-Fi connection. Its head unit is designed to mimic the appearance of a humanoid, which makes the visual interaction more natural.

4.3.3 Hybrid Computing. Some HCAs support both cloud and device-level computing (Tier-3) and adapt their performance according to the availability of computational and communication resources. A Google Glass-based cognitive assistant gracefully degrades services in offline settings [40]. There is a tradeoff between battery life and response time. It offloads computation to a cloud server when possible, since the wearable CA devices are resource-poor in terms of lifetime and computational capacity when compared to server-side devices. NavCog, a navigational assistant, applies a hybrid computing strategy using smartphones and cloud [118]. The system localizes users using the phone and sends the information to the map server on the cloud, where the server calculates the route combining the map information. It verbally interacts with the user through a conversational system. The conversational system combines: (1) a basic conversation script and (2) a recommendation engine. Vorobieva et al. present a robotic assistant for automatic object manipulation and object grasping using vision-based robot control [135]. In this case, object recognition is performed using a client-server architecture. For searching an object, the client sends an image to the network server. Then the server analyzes it and responds to the client.

They use a “SOAP-XML based communication protocol” [52] for this purpose. The robotic system performs local computation to determine how to pick up the item selected by the user using its gripper.

4.4 Emerging Trends and Limitations in HCAs: A Cyber-physical Systems Perspective

- The capacity of existing conversational HCAs for understanding verbal communication and generating accurate and personalized dialogue is often limited [23]. Advanced techniques for these tasks should be integrated into future HCAs for more adaptive, context-aware, realistic, and accurate interaction.
- Some emerging trends among recent HCAs are providing multimodal interaction, offering cognitive assistance for multiple functionalities, and using AR, VR, or MR interfaces. Some example HCAs that use AR, VR, or MR are shown in Figure 4. However, the current version of such HCAs still lack meaningful, frequent interactions with users. Also, they are more vulnerable to attacks that threaten user privacy, security, and safety [77].
- Several HCAs provide suggestions and alerts to professional healthcare providers, including physicians, surgeons, and emergency responders. Qualitative surveys often show that automatic alerts often result in alert fatigue and increased cognitive burden by sending irrelevant alerts. However, alerts can also prevent severe errors and problems. So, alerts should be context-aware, adaptive, and presented according to a hierarchy of severity [53].
- Most HCAs that provide haptic feedback use tactile or vibrotactile feedback. However, some HCAs might benefit from kinesthetic feedback for more natural and realistic actuation, such as the HCAs that track hand motions or focus on physical rehabilitation [45, 128].
- HCAs need be trained with knowledge of patient profiles, populations, guidelines, care pathways, or workflows used by caregivers for context-aware and personalized interventions. Modeling, encoding, and representation of such knowledge (e.g., clinical protocol guidelines and care pathways) are also critical.

5 DESIGN RECOMMENDATIONS AND FUTURE DIRECTIONS FOR INTELLIGENT/COGNITIVE ASSISTANTS IN HEALTHCARE

In this section, we discuss overall design recommendations and future directions for intelligent/cognitive healthcare assistants with respect to current and imminent technologies.

5.1 Enhancing the Cognitive Ability of HCAs

Enhancing cognitive ability is a core challenge for the future generation of HCAs. The ultimate goal of cognitive assistants is to mimic human cognition, which is not yet understood properly. However, cognitive processes can be categorized into different classes, and cognitive assistants are designed to mimic these processes. These include “attention, perception, memory, language, learning, and higher reasoning” [113]. The success of deep learning and reinforcement learning approaches in computer vision and speech processing has enabled current HCAs to provide reasonable performance for low-level cognitive processes such as attention (i.e., “process for selecting an object on which to concentrate”), perception (e.g., activity recognition, object recognition), and memory (i.e., process for storing, finding, and accessing knowledge). However, there is still a large gap in cognitive assistants when dealing with more complex cognitive processes, as presented below.

Language: This refers to the processes for understanding and communicating through natural and meaning-based language. Even the state-of-the-art assistive technologies (i.e., Alexa, Google personal assistant) often fail to **understand user interaction, resolve ambiguity, and deal with uncertainty**. The recent focus on natural language understanding (NLU) and natural language generation (NLG)-related research can help to reduce this gap. Some specific use

cases include, (i) communicating with users based on their health education level such as using easy-to-understand, jargon-free language, and layman’s terms while interacting with patients with **low health literacy**, (ii) choosing **level of details for interaction in a context-aware manner** so users are satisfied with their query and not overwhelmed by the information load, and (iii) **learning user preferences** explicitly and implicitly from their interaction with the system.

Learning: This refers to the “*process for synthesizing new knowledge and connecting new information and experiences with existing knowledge*” [113]. Since the underlying computational models used in HCAs have static representation, the process of learning is often bound by the representation. Dynamic, non-linear knowledge representation to link abstract concepts to concrete representations can be one of the many potential directions to accelerate the learning process of next-generation HCAs. Another major challenge is learning under uncertainty. Recent research points out that in machine learning-based approaches, **uncertainty can come from data** (e.g., noisy data, mis-measurement) and **model** (e.g., model structure and parameters) [31]. Additional uncertainty can be introduced in the context of HCAs from **human error, system malfunctions, and unpredictable or anomalous human behavior** [46, 100]. So, the HCA design principle should accommodate such uncertainty. Some additional issues are **personalized learning, distinguishing emerging patterns from anomalies or outliers, learning in real-time, un-learning the wrong or undesired patterns, and increasing situational awareness**.

Higher Reasoning: This encompasses any process that involves “*reflective cognition such as problem-solving, planning, reasoning, decision-making*” [113]. While the concept of autonomous HCAs that perform critical interventions on their own is a far-fetched idea, HCAs are being used at an increased rate to provide decision support at different steps of healthcare systems as described in Section 2. That demands future research to augment the reasoning capabilities of HCAs.

5.2 Interaction of Pervasive Assistants

With the increasing popularity and diversity of HCAs, a single user is highly likely to use multiple HCAs simultaneously. Thus, multiple HCAs are often likely to interact with each other. Besides, they can also interact with other pervasive assistive technologies and services that the users and their co-residents (i.e., patients and their caregiver or family members, elderly patients living in an assistive facility) use. Since HCAs are developed and deployed independently, often such interactions can result in unwarranted and unexpected situations as illustrated below. We identify this is a potential challenge and research direction for future assistive technologies.⁴ HCAs and assistive services interact at different stages, including sensing, actuation, and control and computational modality. Thus, those interactions result in unique challenges, as outlined below.

Sensing: Multiple HCAs can share data with each other for improved performance and efficiency. For example, an HCA for navigation may detect issues with the motor movement of the user and inform another HCA to select specific lessons of psychomotor exercises to alleviate neuropsychiatric disorder and mental illness of the user. In another example, an in-home HCA may detect agitation, stress, or irritation of the user and inform the HCA for navigation to avoid specific routes or play specific music to reduce stress. In another example, an overweight user is using an HCA to adjust his daily meal and exercise pattern for complying with the guideline suggested by his physician. Now, this HCA can interact with his existing HCA dedicated to ADL support to access current data (i.e., mealtime, duration, exercise duration). While such interaction should be supported by HCAs to decrease redundant user interaction (e.g., both HCAs asking a user about

⁴It should be noted that this interaction is different from the interaction in a multi-agent system [50, 51], since such systems are designed and developed with the multi-agent context in mind. It is not the case when multiple HCAs and other assistive services run simultaneously.

their mealtime and duration might be annoying for the user), the data sharing should also preserve user privacy and confidentiality. This also illustrates the tradeoff between **data re-usability** across applications and potential **privacy vulnerability**. This issue could be more severe for platforms like Alexa. As of now, Alexa skill supports about 1,700 health and fitness apps. In such a platform where each app is typically developed by different vendors, allowing data flow among multiple apps begs the question of how these data are shared and stored. Users should be aware of the terms and conditions of the applications and their permission.

Actuation: When multiple assistants (including HCAs) perform interventions, they may result in **conflicts** [72, 95, 96]. For example, an HCA monitoring the mood of a depression patient suggests the user take a walk outside to boost their mood while another HCA monitoring his ADL reminds him to perform a pre-scheduled task. While some conflicts can be resolved easily by the user and might have non-severe consequences, other conflicts might have serious effects on user health and overall wellness. For example, an HCA for one disease suggesting a food (e.g., grapefruit) or medicine to a patient suffering from multi-morbidity that interferes with his other diseases or medications (i.e., Lipitor or similar medications used for treating high cholesterol) and cause adverse side effects.

Another issue regarding the actuation modality is **coordinating the interventions** among multiple relevant HCAs. Consider a first responder wearing a smart vest that provides haptic feedback to alert about the condition of one or more patients in an emergency scene (e.g., sudden drop or rise in blood pressure or other vitals). If there are too many haptic feedbacks in his vest, it could be confusing and overwhelming for him. A similar issue may arise with navigational HCAs that provide haptic feedback to visually challenged individuals [84]. Some design challenges relevant to this issue include determining the **effective interface** and **frequency and duration** of actuation. The interventions from multiple HCAs should also be coordinated in real-time to make sure they are **adaptive to the dynamic state** of the user, e.g., prognosis of disease, change in behavior or skill level. For example, an HCA suggesting an intervention for a symptom a patient had recently, which is already being treated by another HCA or their care provider.

Control and Computation: Since multiple assistive services (including HCAs) can share resources, they can often interact in control and computational modality. For example, regular assistive apps and apps for assisting cognitive health running in an Alexa device can often run on the local device and perform local computing. Potential issues regarding **computational resource distribution** and **priority of applications** can stem from such situations.

Some design recommendations to address such interaction among multiple pervasive assistants (including HCAs) are as follows: (i) Maintaining a central system similar to home IoT control system for resource management and communication with the HCAs that interact with each other in one or more modalities. (ii) Developing **conflict detection**, degree, or severity detection of potential conflicts and **conflict resolution** systems for health applications [73, 80, 81, 95, 96]. (iii) Developing **synergistic** applications to exploit and transfer information and knowledge from one assistive service to another. For instance, an indoor HCA helps a mobile CA by detecting user's stress and mood, or an HCA for ADL support tracks a user's activities of daily living to increase the user's adherence to medical treatment.

5.3 Addressing Realism

Unrealistic Assumptions: Often the design of HCAs relies on unrealistic assumptions. For example, the increasingly popular modes of interaction through AR, VR, or MR often require the user to wear a head-mounted display, headset, or HoloLens. It may not be realistic for many applications, including navigational assistance and activities of daily living. Another assumption is providing repeated and continuous audio suggestions to users for navigational assistance. It

might annoy users, especially if they are already aware of suggested intervention. Another issue with navigational assistance is they often overlook the case where the user makes a mistake, e.g., the user misses a turn. It is not clear how such systems would resolve such issues. Several HCAs utilizing RGB cameras or depth sensors overlook lighting and background issues. This might result in poor performance in real deployment.

Adaptive Interaction of HCAs with Progression of a Disease: With the surge of voice-based HCAs and CAs, one potential point of failure is the effect of the progression of a disease on user's voice. For some diseases and conditions (e.g., throat cancer, alcohol abuse, thyroid disorders), a user's voice may change with onset and/or progression of the disease [24]. It can cause an error in speaker identification, one of the critical parts of a voice-based assistant, unless the assistant is adaptive to the user's voice. A similar issue can arise from changes in speech patterns (for voice-based assistants), facial expression, and gaits (for multimodal interaction-based assistants). So, one of the design recommendations is to develop patient-facing HCAs that are adaptive to a patient's disease prognosis.

Single User Assumption: Most of the HCAs are developed and tested in a single user setting, which is not often the case. These systems can often fail to recognize desired users correctly and result in an error. For example, most of the HCAs for ADL support and patient-facing decision support do not have any user identification mechanism. Thus, they cannot monitor a user's activity properly (e.g., cooking assistant may be confused when multiple people cooking at the same time or a smartwatch-based HCA for ADL support identifying wrong activities of a couple as they mistakenly switch their smartwatches) or may leak sensitive private information to other family members (e.g., a diagnostic chatbot revealing sensitive information of one user to another user who shares the same device). Researchers have been working on these issues, e.g., speaker ID recognition using voice, or daily activity detection with multiple users using individual smartwatches or smartphones. However, not all the HCAs have sensors such as smartwatches, microphones, or cameras to recognize speaker ID. So, for the next-generation HCAs, the context of usage (single user vs. multiple users) should be considered in the design.

Longitudinal, Real-world Testing: Most of the existing HCAs lack longitudinal, real-world testing and cannot address many realistic situations once they are deployed. For example, NavCogs [121] is tested in a large shopping mall with over 120 beacons, but it may not be feasible and **cost-effective** to deploy so many beacons for an HCA or deploy beacons everywhere in a city for the purpose of navigation for visually impaired people. Also, as HCAs intervene with patient's health, they can adversely affect patient's **health safety**. For example, many recent studies report how machine learning-based diagnostic tools that achieve state-of-the-art accuracy in research experiments often fail when deployed in real-world due to lack of data diversity and other issues [87, 107, 131]. One design recommendation is to develop robust and comprehensive **simulation platforms** for human physiological functionalities and other health environments, when real-world testing is prohibitively expensive. Another relevant challenge is to consider fairness, diversity, inclusion, and equity for minority groups (e.g., in terms of race, gender, and other aspects) during design, development, and test phases of HCAs. Especially for data-driven or model-driven HCAs, the training data should be representative of diverse user groups. This is critical for HCAs that directly impact health outcomes of an individual.

Real-time Embedded AI Challenges: The recent success of AI in solving challenging problems in computer vision, NLP, and speech recognition along with the advances in embedded device technology offers opportunities for developing pervasive and intelligent HCAs powered by embedded AI. The main challenge is to enable complex computations on resource-constrained embedded devices (with constraints on computing power, memory, battery power, and heat dissipation) while satisfying real-time requirements (e.g., for user interactions in HCA applications). As a

result, different research directions are explored to enable real-time AI, e.g., by reducing variance in the search space, changing the order by which the search space is explored, precomputing, incremental and approximate problem solving, focusing on representation of models that understand uncertainty and can reason based on recent data, as mentioned by Musliner et al. [82]. Computing power- and memory-related constraints are especially challenging for perception and control using computer vision and speech recognition. Low-power hardware accelerators for AI and deep learning (e.g., Intel Modivius, Google Edge TPU, and Coral toolkit) have shown promising results in automotive, robotics, and IoT applications. For example, when object detection is performed using the single-shot multibox detector (SSD) algorithm [68] on MobileNet architecture, it gets only 0.5 frames per second (FPS) when running on a Raspberry Pi 3 CPU and 3.5 FPS when Intel Modivius neural compute stick was used in addition to Raspberry PI 3 CPU [89]. While 3.5 FPS may be too low for automotive applications, it may be adequate for some HCAs. However, such solutions need to be a lot more lightweight to run continuously on battery power and enable portable HCA.

5.4 Domain Adaptation of Existing Technology for HCAs

With the recent advancement of machine learning in different domains (i.e., computer vision, speech processing, and natural language processing) general assistive technologies often achieve impressive performance for perception and actuation, e.g., identifying user's face, object, expression, tracking user's mood through their voice and facial expression, speech recognition, and speaker identification. However, due to domain-specific challenges, there is limited applicability of these technologies to HCAs. Some specific use cases are as follows:

Speech recognition: While Google speech recognition API achieves state-of-the-art performance for speech recognition, it has limited applicability in real-time assistants for emergency responders [94, 98, 102], depending on the (i) degree and type of noise present in the speech data collected at the scene and (ii) the seamless availability of network connectivity at the emergency scene. That highlights the need to develop an accurate, standalone, light-weight, open-source speech recognition tool that could be integrated into an HCA for real-time decision support for emergency response.

Textual information extraction for health data: Although there are several solutions for information extraction from medical texts, including EHR data, medical journals, and articles, often these techniques are not sufficient for other medical data [96–98], e.g., textual data related to EMS. This is because EMS data often contain shorthand, abbreviations, and vocabulary that are unique to the EMS domain, and EMS data can be much noisier than the EHR data. Thus, domain-specific techniques should be developed for textual information extraction from noisy, domain-specific health data.

Another related challenge is **processing temporal information in medical data**. It is already a challenging task for clinical records and EMR data [124] and would be even more challenging for the data collected by HCAs, since the latter data are noisier and less organized than the former. The tasks relevant to temporal information extraction are (i) medical NLP techniques for temporal information extraction, (ii) temporal ordering of events from different data streams, and (iii) modeling and representing of time from different data streams. These tasks are crucial for HCAs providing real-time decision support, as they often need to represent the sequence of events chronologically.

Multimodal machine learning: While multimodal machine learning is gaining increased popularity for various AI applications, traditional methods may have limitations when applied to HCAs due to fragmented data flow for healthcare applications and restricted access to health data.

5.5 Performance Metrics

From the perspective of human-computer interaction research, the performance metrics for HCAs can be categorized into the following classes:

- Efficiency of the assistive system, including response time, cost, and scalability.
- Effectiveness of the assistive system, including accuracy, reliability, confidence, explainability.
- User satisfaction with the assistive system [26, 61, 77], including trustworthiness, perceived ease of use, perceived need, perceived safety, data privacy, ethics, and security.

The performance metric can also vary depending on the context of the application: The performance metrics for chatbots [61] would be different than the performance metrics for a humanoid robotic assistant. However, the performance metric should be mapped with the **overarching objective** of the system to go beyond the traditional performance metrics listed above [61]. For example, consider a cognitive assistant to provide navigational support to elderly and cognitively challenged people for using public transport. Such system focuses on improving mobility independence for the target population. So, the performance metrics should reflect that people are more independent in terms of their mobility, e.g., they are taking more trips using public transports when using the navigational assistant.

Also, HCAs often require **application-specific performance metrics** [35, 50, 79]. For example, Shu et al. [126] introduce an average normalized risk score for evaluating the safety of interventions suggested by HCA to the first responders in emergency medical services. However, most of the existing research reviewed in this survey overlooks this critical issue.

5.6 Identifying Application Gap: New Application Areas for HCAs

We identify some gaps in application areas of HCAs and list some emerging areas for future HCAs. Current HCAs mostly focus on assisting repetitive but more straightforward tasks and increase scalability for effective delivery of healthcare, e.g., a virtual conversational agent for disease management and diagnosis, virtual interview agent, virtual therapy assistant. In the future, in addition to improving such HCAs, a new trend of developing HCAs for augmentation of cognitive capabilities might emerge.

5.6.1 Improving Cognition Ability of Users. In addition to assisting users to accomplish desired tasks, future HCAs could target improving the cognition ability of the users. Parsons et al. use Virtual Reality to treat people with autistic spectrum disorders. VR provides an interactive way to train social skills by performing role-playing in different social contexts, thus enhancing the mental stimulation of social events [91]. VR-based game [25] is used to train cognitive and motor skills of Parkinson's disease patients. The main cognitive demands involve planning movements to hit or dodge targets, distribute attention between task and gait to memorize maximum information to be able to answer at the end. Similarly, VR-based and AR-based virtual coaches are used to train users [79], or enhance their cognitive ability and performance [18, 134]. Chicchi et al. explore the potential of using AR to treat psychological disorders [19]. Since AR offers enormous opportunities to create a virtual environment mimicking realistic scenarios and forces users to interact in such an environment safely, it can be used to treat different kinds of phobia.

5.6.2 Training Healthcare Providers. As cognitive assistants become popular, one fundamental research challenge remains on how to provide feedback to users in a way to complete the subsequent tasks as well as to improve their cognition power. AR, VR, and MR platforms should be utilized to address these challenges to develop personalized, efficient, and cost-effective HCAs for

training patients and professional healthcare providers (see Table 3 for examples of such HCAs). VR-based training platforms, such as the SimNow from da Vinci, are widely deployed to train medical residents in a simulated surgical environment. Residents can train on guided or freehand surgical procedures [49] as well as specific sub-tasks [49, 76] and get scores based on their proficiency and safety. Such HCAs improve the learning curve of the residents while also saving hospital resources by not requiring a human teacher/observer or an actual surgical robot to train with.

5.6.3 Automatic Documentation and Summarization. Another potential new area for HCAs is assisting hospital staff in cognitively demanding and error-prone tasks, including recording and interpreting data correctly and validating data. Since human errors in these tasks are common, hospitals are understaffed, and the staff are overworked, such HCAs can reduce human errors and increase health safety. For example, some recent research focuses on automatic documentation of emergency response incidents and generation of patient care reports [21, 22, 102]. Padoy et al. [90] present a model for automatic surgical activity recognition that can be used for offline documentation/report generation of a given surgery and medical resource optimization. By predicting the current surgical workflow, the time remaining to complete the surgery can be estimated, and this can be used to provide reminders to surgical staff that are next-in-line to use the medical resource.

5.6.4 HCAs to Handle Uncertainties in Clinical Care. One of the central challenges of clinical care is the uncertainties of clinical data. Based on user role, uncertainties occur in at least three cases:

(i) First, while collecting and tracking patient's symptoms and medical history, physicians often receive uncertain information, because a patient may simulate, exaggerate, understate, or even forget their symptoms. Therefore, it becomes complicated to infer any concrete conclusions from such cases. Patients might have accurate but fragmented recall of the actual event and may describe things in a non-linear fashion. A patient-facing HCA can help in such cases to collect specific symptoms from the users/patients.

(ii) Second, while performing a physical exam, conducting a diagnostic test, or interpreting test results, physicians and other care providers can overlook a sub-task or crucial information/feature. A care provider facing HCA can "remember" things on their behalf. Cognitive assistance can also be provided in the detection and monitoring of physiological and laboratory anomalies, a major component of a clinician's time [115]. AI-assisted interpretation of images such as chest x-rays has come to be the most generally utilized assistive tool in real-time clinical practice and even presents a potential role to play in the recent COVID-19 viral pandemic crisis.⁵

(iii) Third, patients often do not conform to the guidelines before a diagnostic test (e.g., not taking certain medications before a specific blood test or dietary restrictions before a diagnostic test) and thus test results do not reflect the original condition. To prevent such uncertainties, a patient-facing HCA can track the user's activity and provide reminders to them to ensure they conform to the pre-test guidelines.

New HCAs can focus on handling such uncertainties in different stages of the healthcare workflow, ranging from patient self-care to diagnosis and treatment by physicians.

5.6.5 HCAs for Situational Awareness. Another gap in existing HCAs is lack of assistance for situational awareness. Potential examples include HCAs to monitor physicians and nurses in the ICU for critical procedure to predict and generate alert for potential unsafe or risky interventions (e.g., for tunnel vision during intubation); tracking patients to make sure they follow the suggestions and do not violate any safety measures (e.g., having a food or beverage that has contraindication

⁵<http://blog.qure.ai/notes/chest-xray-AI-qxr-for-covid-19>.

with their medication) for treatment adherence. This is particularly important for treatment adherence, since adherence to recommendations is often poor and typically degrades over time. A challenge is to develop adaptive and personalized recommendations that result in high adherence rates. Situational awareness can be augmented by clinical assistant devices that detect instability while interpreting a variety of channels of incoming, real-time data [114].

5.6.6 HCAs for the Intensive Care Unit. The intensive care unit (ICU) represents a particularly rich target for cognitive assistants in healthcare because of the dynamic and complex workflows involved that are complemented by a copious amount of data [15]. The development of such solutions is best approached with a team consisting of experts in a variety of pertinent domains. While critical care medicine already possesses cognitive assistive tools to determine severity and predict outcome, these would benefit from improvements that would make them more clinically useful [140]. The identification of actionable targets is an important process in optimizing outcomes: The determination of readiness for discharge has benefited from the development of cognitive assistants [6]. Ultimately, decision-making regarding diagnosis and treatment is at the core of any clinical process. In this case, reinforcement learning has been utilized to determine overall fluid requirements and to optimize outcomes in sepsis by determining whether fluids or vasopressors are the best choice at a particular juncture in time [59]. In general, cognitive assistants can be applied in this clinical context to provide the necessary degree of prevention, control, and repair that is required. Finally, these digital assistants must be utilized to prevent, rather than potentially augment, overdiagnosis and overtreatment, and work to reduce the negatives (increased lengths of stay, costs, adverse effects) associated with these errors in practice [116].

5.6.7 Life-long and Life-wide Assistants. Another emerging trend is developing life-long and life-wide personal assistants that interact with a user for not only health-specific activities but also other regular events and activities. One type of such applications is virtual coaches [18, 30, 79, 134]. Many current virtual coaches for behavior change are not *cognitive* yet, as they lack one or more essential features of a cognitive assistant, i.e., interactive, adaptive, and context-aware. In the future, they can be enhanced with the integration of other intelligent assistive services.

6 GLOSSARY OF ACRONYMS

The list of acronyms and shorthand forms used in this article are presented Table 8.

Table 8. List of Acronyms and Shorthand Forms Used in This Paper

Acronym	Full form
2D	2-Dimensional
3D	3-Dimensional
ADL	Activities of Daily Living
AI	Artificial Intelligence
AR	Augmented Reality
BLE	Bluetooth Low Energy
CA	Cognitive Assistant
CPS	Cyber-physical Systems
CPU	Central Processing Unit
CT	Computed Tomography
EHR	Electronic Health Record

(Continued)

Table 8. Continued

Acronym	Full form
EKG/ ECG	Electrocardiography
EMA	Ecological Momentary Assessments
EMR	Electronic Medical Record
EMS	Emergency Medical Services
FA	First Assistant
GPS	Global Positioning System
HCA	Cognitive Assistant for Healthcare/Healthcare Cognitive Assistant
ICU	Intensive Care Unit
IMU	Inertial Measurement Unit
ML	Machine Learning
MR	Mixed Reality
MRI	Magnetic Resonance Imaging
NIRF	Near-Infrared Fluorescent
NLG	Natural Language Generation
NLI	Natural Language Inference
NLP	Natural Language Processing
NLU	Natural Language Understanding
PTSD	Post-Traumatic Stress Disorder
RGB	Red Green Blue
RGB-D	RGB-Depth
STAR	Smart Tissue Autonomous Robot
VR	Virtual Reality

The left and right columns contain the shorthand form or the acronym and the full form, respectively.

7 CONCLUSION

Healthcare cognitive assistants (HCAs) are going to play a significant role in the near future to ensure evidence-driven, accurate, sustainable, and effective healthcare delivery on a large scale. The main objective of this survey article is to identify the technical scope of current HCAs, discover their characteristics and emerging trends, and identify critical research questions and application gaps to enable future research on cognitive assistant technology for healthcare. We review state-of-the-art HCAs from a wide variety of application areas that serve a range of users under different scenarios spanning from preventive medicine to emergency care. Specifically, it presents examples of HCAs from 26 types of applications that are categorized into three high-level classes according to the overarching application goals. It identifies potential application requirements for each type of application. It also identifies the characteristic features of HCAs (i.e., interactive, context-aware, and adaptive) and provides a taxonomy of each of these features as observed in existing HCAs. Then it reviews the critical CPS components of existing HCAs, including sensing or perception, actuation or response, and control and computation. This survey also identifies the emerging trends and potential challenges concerning the features and CPS aspects of HCAs. Finally, it identifies a novel set of challenges, research questions, and potential application gaps for future HCAs from the perspective of machine learning and cyber-physical systems.

REFERENCES

- [1] Azizi Ab Aziz, Ali Saad Fahad, and Faudziah Ahmad. 2017. CAKNA: A personalized robot-based platform for anxiety states therapy. In *Intelligent Environments (Workshops)*. IOS Press, 141–150.

- [2] Saleh Ahmed, Mahboob Qaosar, Rizka Wakhidatus Sholikah, and Yasuhiko Morimoto. 2018. Early dementia detection through conversations to virtual personal assistant. In *Proceedings of the AAAI Spring Symposium Series (AAAI'18)*.
- [3] Dragan Ahmetovic, Cole Gleason, Chengxiong Ruan, Kris Kitani, Hironobu Takagi, and Chieko Asakawa. 2016. NavCog: A navigational cognitive assistant for the blind. In *Proceedings of the 18th International Conference on Human-computer Interaction with Mobile Devices and Services*. ACM, 90–99.
- [4] Rodolfo S. Antunes, Lucas A. Seewald, Vinicius F. Rodrigues, Cristiano A. Da Costa, Rodrigo R. Righi, Andreas Maier, Björn Eskofier, Malte Ollenschläger, Farzad Naderi, Rebecca Fahrig, et al. 2018. A survey of sensors in healthcare workflow monitoring. *ACM Comput. Surv.* 51, 2 (2018), 42.
- [5] Babylon. 2019. Babylon Talk-to-a-Doctor: Our end-to-end services enable providers to deliver flexible, accessible healthcare. Retrieved from <https://www.babylonhealth.com/us/our-services/talk-to-a-doctor>.
- [6] Omar Badawi and Michael J. Breslow. 2012. Readmissions and death after ICU discharge: Development and validation of two predictive models. *PLoS One* 7, 11 (2012).
- [7] Shahram Bahadori, Amedeo Cesta, Giorgio Grisetti, Luca Iocchi, R. Leone, Daniele Nardi, Angelo Oddi, Federico Pecora, and Riccardo Rasconi. 2003. Robocare: An integrated robotic system for the domestic care of the elderly. In *Proceedings of the Workshop on Ambient Intelligence*. Citeseer.
- [8] Jérémy Bauchet, Hélène Pigot, Sylvain Giroux, Dany Lussier-Desrochers, Yves Lachapelle, and Mounir Mokhtari. 2009. Designing judicious interactions for cognitive assistance: The acts of assistance approach. In *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, 11–18.
- [9] Timothy W. Bickmore, Everlyne Kimani, Ha Trinh, Alexandra Pusateri, Michael K. Paasche-Orlow, and Jared W. Magnani. 2018. Managing chronic conditions with a smartphone-based conversational virtual agent. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*. 119–124.
- [10] Marion Blount, Maria R. Ebling, J. Mikael Eklund, Andrew G. James, Carolyn McGregor, Nathan Percival, Kathleen P. Smith, and Daby Sow. 2010. Real-time analysis for intensive care. *IEEE Eng. Med. Biol. Mag.* 29, 2 (2010), 110.
- [11] Antonio Padilha Lanari Bo, Mitsuhiko Hayashibe, and Philippe Poinet. 2011. Joint angle estimation in rehabilitation with inertial sensors and its integration with Kinect. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 3479–3483.
- [12] Nathalie Bricon-Souf and Conrad R. Newman. 2007. Context awareness in health care: A review. *Int. J. Med. Inf.* 76, 1 (2007), 2–12.
- [13] Elizabeth Broadbent, Rebecca Stafford, and Bruce MacDonald. 2009. Acceptance of healthcare robots for the older population: Review and future directions. *Int. J. Social Robot.* 1, 4 (2009), 319.
- [14] Manuel G. Calvo and Daniel Lundqvist. 2008. Facial expressions of emotion (KDEF): Identification under different display-duration conditions. *Behav. Res. Meth.* 40, 1 (2008), 109–115.
- [15] Leo Anthony Celi, Benjamin Fine, and David J. Stone. 2019. An awakening in medicine: The partnership of humanity and intelligent machines. *Lancet Dig. Health* 1, 6 (2019), e255–e257.
- [16] Chien-Yen Chang, Belinda Lange, Mi Zhang, Sebastian Koenig, Phil Requejo, Noom Somboon, Alexander A. Sawchuk, and Albert A. Rizzo. 2012. Towards pervasive physical rehabilitation using Microsoft Kinect. In *Proceedings of the 6th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth'12) and Workshops*. IEEE, 159–162.
- [17] Yao-Jen Chang, Shu-Fang Chen, and Jun-Da Huang. 2011. A Kinect-based system for physical rehabilitation: A pilot study for young adults with motor disabilities. *Res. Dev. Disab.* 32, 6 (2011), 2566–2570.
- [18] Yao-Jen Chang, Ya-Shu Kang, and Po-Chiao Huang. 2013. An augmented reality (AR)-based vocational task prompting system for people with cognitive impairments. *Res. Dev. Disab.* 34, 10 (2013), 3049–3056.
- [19] Irene Alice Chicchi Giglioli, Federica Pallavicini, Elisa Pedroli, Silvia Serino, and Giuseppe Riva. 2015. Augmented reality: A brand new challenge for the assessment and treatment of psychological disorders. *Comput. Math. Meth. Med.* 2015 (2015).
- [20] Ângelo Costa, Ester Martinez-Martin, Angel P. Del Pobol, Ricardo Simoes, and Paulo Novais. 2013. Find it—An assistant home agent. In *Trends in Practical Applications of Agents and Multiagent Systems*. Springer, 121–128.
- [21] Angelo Croatti, Sara Montagna, and Alessandro Ricci. 2017. A personal medical digital assistant agent for supporting human operators in emergency scenarios. In *Agents and Multi-agent Systems for Health Care*. Springer, 59–75.
- [22] Angelo Croatti, Sara Montagna, Alessandro Ricci, Emiliano Gamberini, Vittorio Albarello, and Vanni Agnoletti. 2019. BDI personal medical assistant agents: The case of trauma tracking and alerting. *Artif. Intell. Med.* 96 (2019), 187–197.
- [23] David DeVault, Ron Artstein, Grace Benn, Teresa Dey, Ed Fast, Alesia Gainer, Kallirroi Georgila, Jon Gratch, Arno Hartholt, Margaux Lhommet, et al. 2014. SimSensei Kiosk: A virtual human interviewer for healthcare decision support. In *Proceedings of the International Conference on Autonomous Agents and Multi-agent Systems*. 1061–1068.

- [24] NHS Digital. 2018. Overview—Laryngeal (larynx) cancer. Retrieved from <https://www.nhs.uk/conditions/laryngeal-cancer/>.
- [25] Felipe Augusto dos Santos Mendes, José Eduardo Pompeu, Alexandra Modenesi Lobo, Keyte Guedes da Silva, Tatiana de Paula Oliveira, Andrea Peterson Zomignani, and Maria Elisa Pimentel Piemonte. 2012. Motor learning, retention and transfer after virtual-reality-based training in Parkinson’s disease-Effect of motor and cognitive demands of games: A longitudinal, controlled clinical study. *Physiotherapy* 98, 3 (2012), 217–223.
- [26] Joanna Ejdyś and Katarzyna Halicka. 2018. Sustainable adaptation of new technology—The case of humanoids used for the care of older adults. *Sustainability* 10, 10 (2018), 3770.
- [27] Susan Feldman. 2000. The answer machine. *Searcher* 8, 1 (2000).
- [28] George Ferguson, Jill Quinn, Cecilia Horwitz, M. Swift, J. Allen, and Lucian Galescu. 2010. Towards a personal health management assistant. *J. Biomed. Inform.* 43, 5 (2010), S13–S16.
- [29] Lorenzo Fernandez Maimo, Alberto Huertas Celdran, Angel L. Perales Gomez, García Clemente, J. Félix, James Weimer, and Insup Lee. 2019. Intelligent and dynamic ransomware spread detection and mitigation in integrated clinical environments. *Sensors* 19, 5 (2019), 1114.
- [30] Brian French, Divya Tyamagundlu, Daniel P. Siewiorek, Asim Smailagic, and Dan Ding. 2008. Towards a virtual coach for manual wheelchair users. In *Proceedings of the 12th IEEE International Symposium on Wearable Computers*. IEEE, 77–80.
- [31] Yarin Gal. 2016. *Uncertainty in Deep Learning*. Ph.D. Dissertation. University of Cambridge.
- [32] Luciano Gamberini, Francesco Martino, Bruno Seraglia, Anna Spagnolli, Malena Fabregat, Francisco Ibanez, Mariano Alcaniz, and Javier Montesa Andrés. 2009. Eldergames project: An innovative mixed reality table-top solution to preserve cognitive functions in elderly people. In *Proceedings of the 2nd Conference on Human System Interactions*. IEEE, 164–169.
- [33] Asma Ghandeharioun, Daniel McDuff, Mary Czerwinski, and Kael Rowan. 2018. EMMA: An emotionally intelligent personal assistant for improving wellbeing. *arXiv preprint arXiv:1812.11423* (2018).
- [34] Shameek Ghosh, Sammi Bhatia, and Abhi Bhatia. 2018. Quro: Facilitating user symptom check using a personalised chatbot-oriented dialogue system. *Stud. Health Technol. Inform.* 252 (2018), 51–56.
- [35] Alex Gillespie, Catherine Best, and Brian O’Neill. 2012. Cognitive function and assistive technology for cognition: A systematic review. *J. Int. Neuropsychol. Society* 18, 1 (2012), 1–19.
- [36] Sylvain Giroux, Nathalie Bier, Hélène Pigot, Bruno Bouchard, Abdenour Bouzouane, Mélanie Levasseur, Mélanie Couture, Carolina Bottari, Bonnie Swaine, Pierre-Yves Therriault, et al. 2015. Cognitive assistance to meal preparation: Design, implementation, and assessment in a living lab. In *Proceedings of the AAAI Spring Symposium Series*.
- [37] Reginald G. Golledge et al. 1999. *Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes*. JHU press.
- [38] Pedro Pablo Gomez, Ross E. Willis, and Kent R. Van Sickle. 2015. Development of a virtual reality robotic surgical curriculum using the da Vinci SI surgical system. *Surg. Endosc.* 29, 8 (2015), 2171–2179.
- [39] D. González-Ortega, F. J. Díaz-Pernas, Mario Martínez-Zarzuela, and Miriam Antón-Rodríguez. 2014. A Kinect-based system for cognitive rehabilitation exercises monitoring. *Comput. Meth. Prog. Biomed.* 113, 2 (2014), 620–631.
- [40] Kiryong Ha, Zhuo Chen, Wenlu Hu, Wolfgang Richter, Padmanabhan Pillai, and Mahadev Satyanarayanan. 2014. Towards wearable cognitive assistance. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 68–81.
- [41] Glenn I. Hawe, Graham Coates, Duncan T. Wilson, and Roger S. Crouch. 2012. Agent-based simulation for large-scale emergency response: A survey of usage and implementation. *ACM Comput. Surv.* 45, 1 (2012), 8.
- [42] Marion Hersh. 2015. Overcoming barriers and increasing independence-Service robots for elderly and disabled people. *Int. J. Adv. Robot. Syst.* 12, 8 (2015), 114.
- [43] Ramón Hervás, José Bravo, and Jesús Fontecha. 2013. An assistive navigation system based on augmented reality and context awareness for people with mild cognitive impairments. *IEEE J. Biomed. Health Inform.* 18, 1 (2013), 368–374.
- [44] Jesse Hoey, Craig Boutilier, Pascal Poupart, Patrick Olivier, Andrew Monk, and Alex Mihailidis. 2012. People, sensors, decisions: Customizable and adaptive technologies for assistance in healthcare. *ACM Trans. Interact. Intell. Syst.* 2, 4 (2012), 20.
- [45] Hossein Mousavi Hondori, Maryam Khademi, Lucy Dodakian, Steven C. Cramer, and Cristina Videira Lopes. 2013. A spatial augmented reality rehab system for post-stroke hand rehabilitation. In *Stud. Health Technol. Inform.* 184, 279–285.
- [46] Enamul Hoque, Robert F. Dickerson, Sarah M. Preum, Mark Hanson, Adam Barth, and John A. Stankovic. 2015. Holmes: A comprehensive anomaly detection system for daily in-home activities. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems*. IEEE, 40–51.
- [47] John Hu, Aaron Edsinger, Yi-Je Lim, Nick Donaldson, Mario Solano, Aaron Solocheck, and Ronald Marchessault. 2011. An advanced medical robotic system augmenting healthcare capabilities-robotic nursing assistant. In *Proceedings of the IEEE International Conference on Robotics and Automation*. IEEE, 6264–6269.

- [48] Rajibul Huq, Patricia Kan, Robby Goetschalckx, Debbie Hébert, Jesse Hoey, and Alex Mihailidis. 2011. A decision-theoretic approach in the design of an adaptive upper-limb stroke rehabilitation robot. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics*. IEEE, 1–8.
- [49] Inc Intuitive Surgical. 2019. SimNow from da Vinci. Retrieved from <https://tinyurl.com/y5ubtgd>.
- [50] Sajid Iqbal, Wasif Altaf, Muhammad Aslam, Waqar Mahmood, and Muhammad Usman Ghani Khan. 2016. Application of intelligent agents in health-care. *Artif. Intell. Rev.* 46, 1 (2016), 83–112.
- [51] David Isern and Antonio Moreno. 2016. A systematic literature review of agents applied in healthcare. *J. Med. Syst.* 40, 2 (2016), 43.
- [52] François Jammes, Antoine Mensch, and Harm Smit. 2007. Service-oriented device communications using the devices profile for web services. In *Proceedings of the 21st International Conference on Advanced Information Networking and Applications Workshops (AINAW'07)*, Vol. 1. IEEE, 947–955.
- [53] Minoa Jung, A. Hoerbst, W. O. Hackl, F. Kirrane, D. Borbolla, M. W. Jaspers, M. Oertle, V. Koutkias, L. Ferret, P. Massari, et al. 2013. Attitude of physicians towards automatic alerting in computerized physician order entry systems. *Meth. Inf. Med.* 52, 02 (2013), 99–108.
- [54] Leslie Kay. 1974. A sonar aid to enhance spatial perception of the blind: Engineering design and evaluation. *Radio Electron. Eng.* 44, 11 (1974), 605–627.
- [55] John E. Kelly III and Steve Hamm. 2013. *Smart Machines: IBM's Watson and the Era of Cognitive Computing*. Columbia University Press.
- [56] Maryam Khademi, Hossein Mousavi Hondori, Cristina Videira Lopes, Lucy Dodakian, and Steve C. Cramer. 2012. Haptic augmented reality to monitor human arm's stiffness in rehabilitation. In *Proceedings of the IEEE-EMBS Conference on Biomedical Engineering and Sciences*. IEEE, 892–895.
- [57] Naofumi Kitsunezaki, Eijiro Adachi, Takashi Masuda, and Jun-ichi Mizusawa. 2013. KINECT applications for the physical rehabilitation. In *Proceedings of the IEEE International Symposium on Medical Measurements and Applications (MeMeA'13)*. IEEE, 294–299.
- [58] M. Sun Kohn, J. Sun, S. Knoop, A. Shabo, B. Carmeli, D. Sow, T. Syed-Mahmood, and W. Rapp. 2014. IBM's health analytics and clinical decision support. *Yearb. Med. Inform.* 23, 01 (2014), 154–162.
- [59] Matthieu Komorowski, Leo A. Celi, Omar Badawi, Anthony C. Gordon, and A. Aldo Faisal. 2018. The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nat. Med.* 24, 11 (2018), 1716–1720.
- [60] Belinda Lange, Chien-Yen Chang, Evan Suma, Bradley Newman, Albert Skip Rizzo, and Mark Bolas. 2011. Development and evaluation of low cost game-based balance rehabilitation tool using the Microsoft Kinect sensor. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 1831–1834.
- [61] Liliana Laranjo, Adam G. Dunn, Huong Ly Tong, Ahmet Baki Kocaballi, Jessica Chen, Rabia Bashir, Didi Surian, Blanca Gallego, Farah Magrabi, Annie Y. S. Lau, et al. 2018. Conversational agents in healthcare: A systematic review. *J. Amer. Med. Inform. Assoc.* 25, 9 (2018), 1248–1258.
- [62] Young Hoon Lee and Gérard Medioni. 2011. RGB-D camera based navigation for the visually impaired. In *Proceedings of the Robotics: Science and Systems Conference (RSS'11)*.
- [63] Young Hoon Lee and Gerard Medioni. 2014. Wearable RGBD indoor navigation system for the blind. In *Proceedings of the European Conference on Computer Vision*. Springer, 493–508.
- [64] Manning Li, Jie Hou, Henry Zhang, et al. 2014. A tale of two virtual advisors: An empirical study investigating the empowerment effect of mobile mental-health advisory systems on emergency rescuers. In *Proceedings of the Pacific Asia Conference on Information Systems (PACIS'14)*. 227.
- [65] Manning Li, Zhenhui Jack Jiang, Zhiping Fan, and Jie Hou. 2017. Expert or peer? Understanding the implications of virtual advisor identity on emergency rescuer empowerment in mobile psychological self-help services. *Inf. Manag.* 54, 7 (2017), 866–886.
- [66] Ruijiao Li, Bowen Lu, and Klaus D. McDonald-Maier. 2015. Cognitive assisted living ambient system: A survey. *Dig. Commun. Netw.* 1, 4 (2015), 229–252.
- [67] Christine Lisetti, Reza Amini, Ugan Yasavur, and Naphtali Rische. 2013. I can help you change! An empathic virtual agent delivers behavior change health interventions. *ACM Trans. Manag. Inf. Syst.* 4, 4 (2013), 19.
- [68] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. 2016. SSD: Single shot multibox detector. In *Proceedings of the European Conference on Computer Vision*. Springer, 21–37.
- [69] Yang Liu, Noelle R. B. Stiles, and Markus Meister. 2018. Augmented reality powers a cognitive assistant for the blind. *eLife* 7 (2018), e37841.
- [70] MultiMedia LLC. 2016. Intelligent Cognitive Assistants: Workshop Summary and Recommendations. Retrieved from https://www.nsf.gov/crssprgm/nano/reports/2016-1003_ICA_Workshop_Final_Report_2016.pdf.

- [71] Jack M. Loomis, Reginald G. Golledge, and Roberta L. Klatzky. 1998. Navigation system for the blind: Auditory display modes and guidance. *Presence* 7, 2 (1998), 193–203.
- [72] Meiyi Ma, Sarah Masud Preum, and John A. Stankovic. 2017. Cityguard: A watchdog for safety-aware conflict detection in smart cities. In *Proceedings of the 2nd International Conference on Internet-of-things Design and Implementation*. 259–270.
- [73] Meiyi Ma, John A. Stankovic, and Lu Feng. 2018. CityResolver: A decision support system for conflict resolution in smart cities. In *Proceedings of the ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPs'18)*. IEEE, 55–64.
- [74] Andrea Britto Mattos and Dario Augusto Borges Oliveira. 2018. Multi-view mouth rendering for assisting lip-reading. In *ACM Proc. Internet Access. Things*. 28 (Apr. 2018) 1–10.
- [75] Alex Mihailidis, Brent Carmichael, and Jennifer Boger. 2004. The use of computer vision in an intelligent environment to support aging-in-place, safety, and independence in the home. *IEEE Trans. Inf. Technol. Biomed.* 8, 3 (2004), 238–247.
- [76] Inc Mimic Technologies. 2019. dV-Trainer. Retrieved from <https://mimicsimulation.com/dv-trainer/>.
- [77] Adam S. Miner, Nigam Shah, Kim D. Bullock, Bruce A. Arnow, Jeremy Bailenson, and Jeff Hancock. 2019. Key considerations for incorporating conversational AI in psychotherapy. *Front. Psychi.* 10 (2019).
- [78] Sara Montagna, Angelo Croatti, Alessandro Ricci, Vanni Agnoletti, Vittorio Albarello, and Emiliano Gamberini. 2020. Real-time tracking and documentation in trauma management. *Health Inform. J.* 26, 1 (2020), 328–341.
- [79] Fariba Mostajeran, Nikolaos Katakis, Oscar Ariza, Jann Philipp Freiwald, and Frank Steinicke. 2019. Welcoming a holographic virtual coach for balance training at home: Two focus groups with older adults. In *Proceedings of the IEEE Conference on Virtual Reality and 3D User Interfaces (VR'19)*. IEEE, 1465–1470.
- [80] Sirajum Munir, Mohsin Ahmed, and John Stankovic. 2015. Eyephy: Detecting dependencies in cyber-physical system apps due to human-in-the-loop. In *Proceedings of the 12th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*. Citeseer, 170–179.
- [81] Sirajum Munir and John A. Stankovic. 2014. DepSys: Dependency aware integration of cyber-physical systems for smart homes. In *Proceedings of the ACM/IEEE International Conference on Cyber-Physical Systems (ICCPs'14)*. IEEE, 127–138.
- [82] David J. Musliner, James A. Hendler, Ashok K. Agrawala, Edmund H. Durfee, Jay K. Strosnider, and C. J. Paul. 1995. The challenges of real-time AI. *Computer* 28, 1 (1995), 58–66.
- [83] Alexander Neumann, Christof Elbrechter, Nadine Pfeiffer-Leßmann, Risto Kõiva, Birte Carlmeyer, Stefan Rütter, Michael Schade, André Ückermann, Sven Wachsmuth, and Helge J. Ritter. 2017. KogniChef: A cognitive cooking assistant. *KI-Künstl. Intell.* 31, 3 (2017), 273–281.
- [84] Quoc-Hung Nguyen, Hai Vu, Thanh-Hai Tran, and Quang-Hoan Nguyen. 2017. Developing a way-finding system on mobile robot assisting visually impaired people in an indoor environment. *Multimedia Tools Applic.* 76, 2 (2017), 2645–2669.
- [85] Ahmed K. Noor. 2015. Potential of cognitive computing and cognitive systems. *Open Eng.* 5, 1 (2015), 75–88.
- [86] John Oakley. 2018. Intelligent Cognitive Assistants: Workshop Summary and Recommendations. Retrieved from https://www.nsf.gov/crssprgm/nano/reports/ICA2_Workshop_Report_2018.pdf.
- [87] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [88] Eshed Ohn-Bar, Kris Kitani, and Chieko Asakawa. 2018. Personalized dynamics models for adaptive assistive navigation systems. In *Proceedings of the Conference on Robot Learning*. 16–39.
- [89] Nashwan Adnan Othman and Ilhan Aydin. 2018. A new deep learning application based on Movidius NCS for embedded object detection and recognition. In *Proceedings of the 2nd International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT'18)*. IEEE, 1–5.
- [90] Nicolas Padoy, Tobias Blum, Seyed-Ahmad Ahmadi, Hubertus Feussner, Marie-Odile Berger, and Nassir Navab. 2012. Statistical modeling and recognition of surgical workflow. *Med. Image Anal.* 16, 3 (2012), 632–641.
- [91] Sarah Parsons and Peter Mitchell. 2002. The potential of virtual reality in social skills training for people with autistic spectrum disorders. *J. Intell. Disab. Res.* 46, 5 (2002), 430–443.
- [92] Paul Phamduy, John-Ross Rizzo, Todd E. Hudson, Marina Torre, Kalle Levon, and Maurizio Porfiri. 2017. Communicating through touch: Macro fiber composites for tactile stimulation on the abdomen. *IEEE Trans. Haptics* 11, 2 (2017), 174–184.
- [93] Martha E. Pollack, Laura Brown, Dirk Colbry, Cheryl Orosz, Bart Peintner, Sailesh Ramakrishnan, Sandra Engberg, Judith T. Matthews, Jacqueline Dunbar-Jacob, Colleen E. McCarthy, et al. 2002. Pearl: A mobile robotic assistant for the elderly. In *Proceedings of the AAAI Workshop on Automation as Eldercare*, Vol. 2002. 85–91.
- [94] Sarah Preum, Sile Shu, Mustafa Hotaki, Ronald Williams, John Stankovic, and Homa Alemzadeh. 2019. CognitiveEMS: A cognitive assistant system for emergency medical services. *ACM SIGBED Rev.* 16, 2 (2019), 51–60.

- [95] Sarah Masud Preum, Abu Sayeed Mondol, Meiyi Ma, Hongning Wang, and John A. Stankovic. 2017. Preclude: Conflict detection in textual health advice. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom'17)*. IEEE, 286–296.
- [96] Sarah Masud Preum, Abu Sayeed Mondol, Meiyi Ma, Hongning Wang, and John A. Stankovic. 2017. Preclude2: Personalized conflict detection in heterogeneous health applications. *Pervas. Mob. Comput.* 42 (2017).
- [97] Sarah Masud Preum, Md Rizwan Parvez, Kai-Wei Chang, and John Stankovic. 2018. A corpus of drug usage guidelines annotated with type of advice. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC'18)*.
- [98] Sarah Masud Preum, Sile Shu, Homa Alemzadeh, and John A. Stankovic. 2020. EMSContExt: EMS protocol-driven concept extraction for cognitive assistance in emergency response. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 13350–13355.
- [99] Sarah Masud Preum, Sile Shu, Jonathan Ting, Vincent Lin, Ronald Williams, John Stankovic, and Homa Alemzadeh. 2018. Towards a cognitive assistant system for emergency response. In *Proceedings of the ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS'18)*. IEEE, 347–348.
- [100] Sarah Masud Preum, John A. Stankovic, and Yanjun Qi. 2015. MAPer: A multi-scale adaptive personalized model for temporal human behavior prediction. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. 433–442.
- [101] Long Qian, Anton Deguet, and Peter Kazanzides. 2018. ARssist: Augmented reality on a head-mounted display for the first assistant in robotic surgery. *Healthc. Technol. Lett.* 5, 5 (2018), 194–200.
- [102] M. Arif Rahman, Sarah Masud Preum, Ronald D. Williams, Homa Alemzadeh, and John A. Stankovic. 2020. GRACE: Generating summary reports automatically for cognitive assistance in emergency response. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 13356–13362.
- [103] Vijay Rajanna, Raniero Lara-Garduno, Dev Jyoti Behera, Karthic Madanagopal, Daniel Goldberg, and Tracy Hammond. 2014. Step Up Life: A context aware health assistant. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on the Use of GIS in Public Health*. ACM, 21–30.
- [104] Vijay Rajanna, Patrick Vo, Jerry Barth, Matthew Mjelde, Trevor Grey, Cassandra Oduola, and Tracy Hammond. 2016. KinoHaptics: An automated, wearable, haptic assisted, physio-therapeutic system for post-surgery rehabilitation and self-care. *J. Med. Syst.* 40, 3 (2016), 60.
- [105] Emad Rajih, Côme Tholomier, Beatrice Cormier, Vanessa Samouëlian, Thomas Warkus, Moïshe Liberman, Hugues Widmer, Jean-Baptiste Lattouf, Abdullah M. Alenizi, Malek Meskawi, et al. 2017. Error reporting from the da Vinci surgical system in robotic surgery: A Canadian multispecialty experience at a single academic centre. *Canad. Urolog. Assoc. J.* 11, 5 (2017), E197.
- [106] Sarvapali D. Ramchurn, Feng Wu, Wencho Jiang, Joel E. Fischer, Steve Reece, Stephen Roberts, Tom Rodden, Chris Greenhalgh, and Nicholas R. Jennings. 2016. Human-agent collaboration for disaster response. *Auton. Agents Multi-agent Syst.* 30, 1 (2016), 82–111.
- [107] Sandeep Reddy, Sonia Allan, Simon Coghlan, and Paul Cooper. 2020. A governance model for the application of AI in health care. *J. Amer. Med. Inform. Assoc.* 27, 3 (2020), 491–497.
- [108] Hadley Reynolds and Susan Feldman. 2014. Cognitive computing: Beyond the hype. *KM World* 27 (2014).
- [109] Flavio Ribeiro, Dinei Florencio, Philip A. Chou, and Zhengyou Zhang. 2012. Auditory augmented reality: Object sonification for the visually impaired. In *Proceedings of the IEEE 14th International Workshop on Multimedia Signal Processing (MMSp'12)*. IEEE, 319–324.
- [110] Jaime A. Rincon, Angelo Costa, Paulo Novais, Vicente Julian, and Carlos Carrascosa. 2019. A new emotional robot assistant that facilitates human interaction and persuasion. *Knowl. Inf. Syst.* 60, 1 (2019), 363–383.
- [111] Albert Rizzo, Greg Reger, Greg Gahm, JoAnn Difede, and Barbara O. Rothbaum. 2009. Virtual reality exposure therapy for combat-related PTSD. In *Post-traumatic Stress Disorder*. Springer, 375–399.
- [112] Albert Rizzo, Thomas Talbot, C. D. Combs, J. A. Sokolowski, and C. M. Banks. 2015. Virtual reality standardized patients for clinical training. In *The Digital Patient: Advancing Healthcare, Research, and Education*. Wiley, 257–272.
- [113] Y. Rogers, H. Sharp, and J. Preece. 2011. *Interaction Design: Beyond Human-computer Interaction*. Wiley & Sons.
- [114] Caroline M. Ruminski, Matthew T. Clark, Douglas E. Lake, Rebecca R. Kitzmiller, Jessica Keim-Malpass, Matthew P. Robertson, Theresa R. Simons, J. Randall Moorman, and J. Forrest Calland. 2019. Impact of predictive analytics based on continuous cardiorespiratory monitoring in a surgical and trauma intensive care unit. *J. Clin. Monit. Comput.* 33, 4 (2019), 703–711.
- [115] Barret Rush, Leo Anthony Celi, and David J. Stone. 2019. Applying machine learning to continuously monitored physiological data. *J. Clin. Monit. Comput.* 33, 5 (2019), 887–893.
- [116] Barret Rush, David J. Stone, and Leo Anthony Celi. 2018. From big data to artificial intelligence: Harnessing data routinely collected in the process of care. *Crit. Care Med.* 46, 2 (2018), 345.
- [117] Gopala Sainarayanan, R. Nagarajan, and Szazali Yaacob. 2007. Fuzzy image processing scheme for autonomous navigation of human blind. *Appl. Soft Comput.* 7, 1 (2007), 257–264.

- [118] Daisuke Sato, Uran Oh, Kakuya Naito, Hironobu Takagi, Kris Kitani, and Chieko Asakawa. 2017. NavCog3: An evaluation of a smartphone-based blind indoor navigation assistant with semantic features in a large-scale environment. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, 270–279.
- [119] Mahadev Satyanarayanan, Wei Gao, and Brandon Lucia. 2019. The computing landscape of the 21st century. In *Proceedings of the 20th International Workshop on Mobile Computing Systems and Applications*. 45–50.
- [120] Tobias Schwarze, Martin Lauer, Manuel Schwaab, Michailas Romanovas, Sandra Böhm, and Thomas Jürgensohn. 2016. A camera-based mobility aid for visually impaired people. *KI-Künstl. Intell.* 30, 1 (2016), 29–36.
- [121] Gaumard Scientific. 2019. Pediatric & Neonatal Care Simulators. Retrieved from <https://www.gaumard.com/products/pediatric-neonatal>.
- [122] Audrey Serna, Hélène Pigot, and Vincent Rialle. 2007. Modeling the progression of Alzheimer’s disease for cognitive assistance in smart homes. *User Model. User-adapt. Interact.* 17, 4 (2007), 415–438.
- [123] Azad Shademan, Ryan S. Decker, Justin D. Opfermann, Simon Leonard, Axel Krieger, and Peter C. W. Kim. 2016. Supervised autonomous robotic soft tissue surgery. *Sci. Translat. Medicine* 8, 337 (2016), 337ra64–337ra64.
- [124] Afzal Hussain Shahid and M. P. Singh. 2019. Computational intelligence techniques for medical diagnosis and prognosis: Problems and current developments. *Biocyber. Biomed. Eng.* 39, 3 (2019), 638–672.
- [125] Shraga Shoval, Johann Borenstein, and Yoram Koren. 1998. The Navbelt-A computerized travel aid for the blind based on mobile robotics technology. *IEEE Trans. Biomed. Eng.* 45, 11 (1998), 1376–1386.
- [126] Sile Shu, Sarah Preum, Haydon Pitchford, Ronald Williams, John Stankovic, and Homa Alemzadeh. 2019. A behavior tree cognitive assistant system for emergency medical services. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’19)*. IEEE.
- [127] Daniel Siewiorek, Asim Smalagic, and Anind Dey. 2012. Architecture and applications of virtual coaches. *Proc. IEEE* 100, 8 (2012), 2472–2488.
- [128] Daniel Sonntag. 2015. Kognit: Intelligent cognitive enhancement technology by cognitive models and mixed reality for dementia patients. In *Proceedings of the AAAI Fall Symposium Series*.
- [129] Daniel Sonntag, Sonja Zillner, Christian Schulz, Markus Weber, and Takumi Toyama. 2013. Towards medical cyber-physical systems: Multimodal augmented reality for doctors and knowledge discovery about patients. In *Proceedings of the International Conference of Design, User Experience, and Usability*. Springer, 401–410.
- [130] Mirjana Spasojevic, Stefan Marti, Davide Di Censo, and Jaime Elliot Nahman. 2018. Cognitive load driving assistant. US Patent App. 15/541,458.
- [131] Eliza Strickland. 2019. IBM Watson, heal thyself: How IBM overpromised and underdelivered on AI health care. *IEEE Spect.* 56, 4 (2019), 24–31.
- [132] Yingli Tian. 2014. RGB-D sensor-based computer vision assistive technology for visually impaired persons. In *Computer Vision and Machine Learning with RGB-D Sensors*. Springer, 173–194.
- [133] Iwan Ulrich and Johann Borenstein. 2001. The GuideCane-applying mobile robot technologies to assist the visually impaired. *IEEE Trans. Syst. Man Cyber. Part A: Syst. Hum.* 31, 2 (2001), 131–136.
- [134] Bram Van de Laar, Anton Nijholt, and Job Zwiers. 2010. Monitoring user’s brain activity for a virtual coach. In *Proceedings of the International Conference on Entertainment Computing*. Springer, 511–513.
- [135] Hélène Vorobieva, Mariette Soury, Patrick Hède, Christophe Leroux, and Philippe Morignot. 2010. Object recognition and ontology for manipulation with an assistant robot. In *Proceedings of the International Conference on Smart Homes and Health Telematics*. Springer, 178–185.
- [136] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: The PANAS scales. *J. Person. Social Psychol.* 54, 6 (1988), 1063.
- [137] IBM Watson. 2019. IBM Watson for Oncology. Retrieved from <https://tinyurl.com/y5etnmrd>.
- [138] Oliver Weede, Andreas Bihlmaier, Jessica Hutzl, Beat P. Müller-Stich, and Heinz Wörn. 2013. Towards cognitive medical robotics in minimal invasive surgery. In *Proceedings of the Conference on Advances in Robotics*. ACM, 1–8.
- [139] Patricia A. H. Williams and Andrew J. Woodward. 2015. Cybersecurity vulnerabilities in medical devices: A complex environment and multifaceted problem. *Med. Dev.* 8 (2015), 305.
- [140] Laure Wynants, Ben Van Calster, Marc M. J. Bonten, Gary S. Collins, Thomas P. A. Debray, Maarten De Vos, Maria C. Haller, Georg Heinze, Karel G. M. Moons, Richard D. Riley, et al. 2020. Prediction models for diagnosis and prognosis of covid-19 infection: Systematic review and critical appraisal. *BMJ* 369 (2020).
- [141] Mohammad Samin Yasar and Homa Alemzadeh. 2020. Real-time context-aware detection of unsafe events in robot-assisted surgery. In *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN’20)*. IEEE Computer Society, 385–397.
- [142] Fereshtha Yazdani, Gayane Kazhoyan, Asil Kaan Bozcuoğlu, Andrei Haidu, Ferenc Bálint-Benczédi, Daniel Beßler, Mihai Pomarlan, and Michael Beetz. 2018. Cognition-enabled framework for mixed human-robot rescue teams. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’18)*. IEEE, 1421–1428.

Received October 2019; revised August 2020; accepted August 2020