

User-Like Bots for Cognitive Automation: A Survey

Habtom Kahsay Gidey¹, Peter Hillmann¹, Andreas Karcher¹, and Alois Knoll²

¹ Universität der Bundeswehr München, Germany

{habtom.gidey, peter.hillmann, andreas.karcher}@unibw.de

² Technische Universität München, München, Germany

{knoll}@in.tum.de

Abstract. Software bots have attracted increasing interest and popularity in both research and society. Their contributions span automation, digital twins, game characters with conscious-like behavior, and social media. However, there is still a lack of intelligent bots that can adapt to the variability and dynamic nature of digital web environments. Unlike human users, they have difficulty understanding and exploiting the affordances across multiple virtual environments.

Despite the hype, bots with human user-like cognition do not currently exist. Chatbots, for instance, lack situational awareness on the digital platforms where they operate, preventing them from enacting meaningful and autonomous intelligent behavior similar to human users.

In this survey, we aim to explore the role of cognitive architectures in supporting efforts towards engineering software bots with advanced general intelligence. We discuss how cognitive architectures can contribute to creating intelligent software bots. Furthermore, we highlight key architectural recommendations for the future development of autonomous, user-like cognitive bots.

Keywords: software bot · cognitive architecture · cognitive automation.

1 Introduction

Software bots are becoming an integral part of automation and social computing. Digital platforms, including software ecosystems and cyber-physical systems, are growing increasingly complex. The complexity of diverse digital systems can overwhelm even expert human users [1]. In such scenarios, software agents acting as autonomous users can assist in automating human user activities. As a result, there is a growing interest in augmenting bots into software-intensive business, social, or industrial environments for cognitive automation [2]. In Industry 4.0 (I4.0) and digital twins (DTs), software agents are playing a crucial role in enabling smart factories with higher flexibility, efficiency, and safety [3–5]. Additionally, in the service industry, robotic process automation (RPA) leverages bots to automate business processes [6].

Software development bots, also known as DevBots, are making their mark in automated software engineering [7, 8]. It has become common to see bots assisting in code review and bug-fixing on platforms like GitHub [8, 9]. Wessel et al. [8] identified 48 bots used for this purpose. From committing code to coordinating open-source projects, bots are increasingly becoming a part of the software development life cycle [8, 9]. Their impact on development can significantly affect how future digital innovation ecosystems are managed and governed [10–12]. Social platforms and games also serve as environments for social bots and virtual avatars [13, 14]. Although claims of political intent and the influence of social bots on social media are exaggerated, bots are also prevalent on social platforms nowadays [15]. The diverse applications of bots highlight the desiderata and requirements for advanced cognitive agents [16].

However, the reality falls short of the expectations, and advanced social bots do not exist today [15, 17]. First, the level of autonomy in industrial software agents and robotic process automation (RPA) is minimal [18, 19]. Agents often have architectures that are tightly coupled with specific service platforms. As a result, bots are designed and optimized for these specific platforms, limiting their adaptability. They lack the ability to autonomously recognize the variability and then function effectively across diverse web or service environments. Consequently, they lack a sense of awareness and the ability to adapt to different contexts. Second, bots have very limited or no autonomy in their behavior. Addressing these challenges requires a focus on identifying and resolving architectural concerns [20–23]. It is important to note that bots in social media, games, industry, and business use cases may differ, but the software systems architecture and engineering challenges can span domains [24, 25].

Thus, by reviewing and synthesizing existing works, this study aims to explore two architectural aspects of bots through two research questions. The first question investigates the role of cognitive architectures in bot behavior, while the second question examines the strict separation of bots and their operational environment. In our context, the second question is operationalized by evaluating *user-likeness* or similarity. A *user-like bot* refers to the level of similarity a bot possesses compared to a human user [26].

The subsequent sections are structured as follows: Section 2 provides the foundational background on bots and cognition, Section 3 describes the approach employed in surveying bots, Section 4 presents the evaluated results, Section 5 discusses the implications of the obtained results, and finally, Section 6 concludes the study.

2 Bots and Cognition

2.1 Software Bots

Software bot, or simply bot, is an umbrella term for diverse software agents [27]. The term is used loosely in domains such as RPA. RPA, robotic process automation, is an automation approach that employs software bots, sometimes referred

to as digital workers [6]. RPAs and digital twins utilize various toolsets and development paradigms from agent-based systems [5, 28].

Lebeuf [29] conducted a comprehensive study proposing a broad definition and general taxonomy of bots. She defines software bots as interfaces that connect users to software services and describes them as a “*new [user] interface paradigm*”. According to this perspective, users can access software services through a bot, where the user interface takes the form of a conversational interface. The user is typically human, although other programs and systems can also utilize the bot. Software services, in this context, refer to applications or digital platforms that provide additional functionalities. While these services are typically external, they can also be integrated as internal components of the bot.

Furthermore, Lebeuf [29] classifies bots based on their observable properties and behaviors. The taxonomy defines three dimensions: environment, intrinsic, and interactions. The environment dimension refers to the software service properties that bots operate on. Ideally, a software bot is separate from a specific platform and can operate on multiple platforms. However, most bots are tailored to a specific platform, such as a Twitter bot. The intrinsic dimension describes the internal abilities of the bot. Lebeuf also puts anthropomorphism as an intrinsic dimension, which specifies if the bot has human *user-like* features, such as name, visualization, and persona. The interaction dimension specifies how the bot accesses and interacts with the environment.

Lebeuf’s taxonomy is based solely on observable properties and does not consider the architectural aspects and components of bots [29]. In our context, the interaction dimension concerns perception, action, and autonomy, which impact the system architecture and environment of the bot. Therefore, in this study, we introduce another term to restrict the scope to bots that are similar to users.

User-like Software Bots Today, the web is an operational environment for human users and software user agents [26]. Unlike virtual reality environments, where agents are typically manifested as conscious-like avatars, the web represents a different kind of mixed-reality environment. In this environment, both human users and software bots interact, tying virtual elements with real-world extensions, as exemplified by the Web of Things (WoT) [30–32]. A diverse set of software agents interact with the web as their environment [29, 32].

In this study, we refer to these classes of mixed-reality software agents as *user-like bots* [26]. The term *user-like* implies that these bots exhibit similarities to human users. User-like bots use graphical user interfaces to perceive and act within a software service environment. They interact with keyboard and mouse operations.

As Gibson’s theory of visual perception suggests, perception is not merely about passive observation; it’s also about actively distinguishing the potential actions or “*affordances*” that an environment offers to an agent [33]. The environment, in this context, presents affordances [34]. Affordances refer to possibilities for interaction and action [26]. Platform services or features can be analogous

to affordances in the real world. User-like bots are expected to understand or perceive these affordances and act on them similar to human users.

While the web is the primary context, a software service environment can also refer to any desktop application.

2.2 Cognitive Architectures

Research in engineering machine intelligence, particularly Artificial General Intelligence (AGI), aims to endow software systems with cognitive abilities that enable machines to think at or beyond the level of humans [23]. To achieve this goal, efforts to understand the brain from disciplines such as cognitive and neuroscience have led to various promising approaches. One such approach is the study of cognitive architectures, which focuses on designing high-level cognitive functions [26, 35].

Cognitive architectures serve as essential architectural design foundations for artificial intelligence research. The ultimate objective of cognitive architectures is to enable software with cognitive abilities equal to or greater than human-like intelligence. Cognitive architectures can be described as a set of specifications or theories of cognition that outline the essential structural elements and capabilities of a cognitive system [36, 37]. Metzler and Shea [38] compiled a list of cognitive functions as components that constitute a cognitive architecture, such as learning, reasoning, decision-making, perception, planning, and acting.

Cognitive architectures are designed to handle a broader set of cognitive tasks or cognitive functions. They can enable perpetual learning from the environment, adapting to changes, and reasoning based on available information. Due to this universal approach, an agent implementing a cognitive architecture may operate successively and simultaneously in various applications. While most cognitive architectures remain theoretical specifications, some, such as ACT-R, LIDA, and SOAR, have implementations and active communities [36]. Duch et al. [39] have conducted an in-depth comparison on a technical level of these architectures, among others.




Fig. 1: Schematic Structure of the Standard Model of Cognition [40].

Currently, efforts are being made in the cognitive modeling research community to establish a comprehensive understanding of the architectural assumptions

that define aspects of human-like cognition, whether natural or artificial [40]. Fig. 1 shows a schematic structure of the Standard Model of Cognition [40]. The Standard Model of Cognition aims to consolidate knowledge from all existing cognitive architectures and establish a unified understanding of the generic aspects of cognition, such as perception and motor functions, common to all cognitive architectures.

2.3 Cognitive Automation

Traditional automation methods and approaches, such as process automation and RPA, have played crucial roles in automating repetitive tasks and workflows [2, 41, 42]. However, as organizations strive for increased agility, efficiency, and even hyper-automation, the increased complexity of cyber-physical systems necessitates additional layers of automation, i.e., cognitive automation [2].

Cognitive automation aims to advance the capabilities of traditional automation and RPA by combining them with other technologies in artificial intelligence. This integrated approach enables the automation of more complex and cognitive tasks that traditionally require human intervention. Cognitive automation aims to automate knowledge and service work that involves decision-making, problem-solving, and other cognitive activities [2, 43]. It focuses on alleviating the burden of cognitive tasks on humans by automating their roles or enhancing mixed reality collaborations [2, 42, 44].

Similarly, advanced user-like bots or software agents can be employed to achieve generalizable intelligence or cognitive capabilities on digital platforms comparable to human users [26]. As autonomous users, software agents can then assist in automating human user activities. Consequently, one prominent application of cognitive automation can be utilizing user-like bots. Bots with generalizable intelligence, matching human users, are one way to address the augmentation of autonomous digital workers to cyber-physical systems. These bots can enhance efficiency in various domains, such as RPA, knowledge platforms, digital twins, and smart factories [5, 26].

3 Research Approach

In this section, we describe the study method and its execution.

3.1 Planning

We conducted a literature review to examine the integration of cognitive architectures into the engineering of user-like software bots. This approach allows us to investigate and challenge existing architectural assumptions of software bot development, which fail to enact meaningful, intelligent behavior. The method also uses specific criteria to narrow the review’s scope and establish a clear and replicable methodology. The survey is executed by systematically searching, selecting, and evaluating relevant works.

3.2 Initial Selection

Initially, we collected various bots, agents, software tools, and personal digital assistants to get an overview. Since terms such as agent, bot, software bot, user-agent, and chatbot are interchangeably used in literature, the initial efforts resulted in an extensive collection of sources. However, this initial process helped explore and map the software bots already developed throughout all time.

We then limited the collection to the bots that use some form of cognitive architecture or cognitive model. Furthermore, we determined the search to include works that distinctively show some similarity to a way a human user would access and with software service platforms, hence only user-like software bots.

First, we systematically selected relevant works from comprehensive databases, such as Scopus. To that end, specific keywords, including variations of “software bot,” “cognitive architecture,” and “user-like or user-agent,” were devised and conducted. Next, we used forward and backward snowball sampling [45] to find citation chains of relevant studies that claim the development of bots with cognitive capabilities.

Our initial collection resulted in approximately 190 works. With a closer look, we found that many were unrelated to the distinct interest. The initial results are then narrowed to a representative selection of software bots for further investigation and evaluation.

3.3 Selection and Evaluation Criteria

The representative selection, also called candidates, is evaluated with four essential selection criteria. The four criteria used to evaluate the selection are derived to ensure the proper implementation of the interest and objective of the study. Each candidate bot selected is evaluated against all four criteria and the corresponding sub-criteria. Fig. 2 depicts the four criteria and subcriteria of the first two: Software Bot and User Similarity.

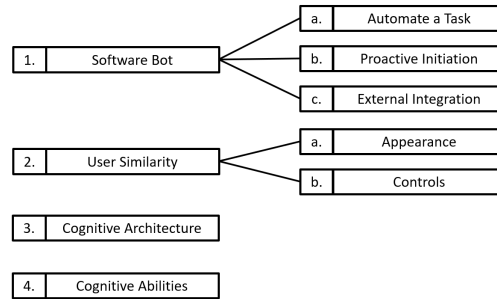


Fig. 2: Criteria used to evaluate the bots applying cognitive models.

(1) *Software bot*: this criterion evaluates bot capability. It helps exclude conversational interfaces or chatbots that do little or no autonomous action on

service platforms. The criterion has three subcriteria: a. Automate a Task, b. Proactive Initiation, and c. External integration. One of the three aspects is the necessity that the bots automate a task in a digital environment (a. Automate a Task). In this context, the scope of the environment and the complexity of the task are neglected. To differentiate the candidate from other programs and scripts, as a second aspect, independence from the user is required (b. Proactive Initiation). Changes to other elements of the environment, combined with intrinsic motivations, should trigger the behavior and actions the bot performs. Indirect reactions to the actions of the user are permitted. For the third aspect, the candidate bot should operate externally in the targeted digital environment (c. External integration). Direct integration into the environment, as a fixed part of a system, is not considered user-like.

(2) *User Similarity*: Equal to the first criterion, this criterion stems from the distinction set regarding user-like software bots in Sect. 2. User-like is used to denote agents who either resemble human users or act on behalf of users. We focus on drawing a clear distinction between the bots that autonomously enact user-like behavior and others.

(3) *Cognitive Architecture*: While claiming the implementation of a cognitive architecture is relatively easy, its practical application poses a challenge.

(4) *Cognitive Abilities*: Implementing a cognitive architecture does not necessarily result in cognitive behavior. Consequently, this criterion unites various aspects regarding a candidate’s cognitive abilities. The list aims at excluding narrowly set low-level heuristics from the selection while considering the desired high-level cognitive functions or generalized cognitive capabilities.

4 Results

This section presents the selected agents, personal assistants, and applied cognitive models. Their evaluation is based on the criteria outlined in Sect.3. Furthermore, descriptions of notable selected results from the corresponding review are provided. Table 1 shows the main results of the software bots analyzed and evaluated. The entries are ordered by the year of the publication. However, other arrangements are also feasible. The meaning of the encoded column captions is provided in Sect. 3.3 and Fig. 2. The elements of the selection cover diverse research domains. In the following, we highlight the significant bases of the elements in Table 1. It is worth noting that none of the candidates in the related publications are described as “*user-like software bots*” as termed in this study. They were included because their implicit descriptions in the literature align with the previously set criteria and guidelines.

4.1 IDA and Virtual Mattie

The agent Virtual Mattie was developed by Franklin et al. in 1996 to perform clerical tasks related to the organization of seminar schedules [46]. Human seminar organizers are contacted via email. The primary cognitive abilities of Virtual

Table 1: Overview of selected bots, agents, or assistants, with evaluation results.

Study Ref.	Description	Software Bot		User Similarity		Cognitive Architecture	Cognitive Abilities
		a.	b.	a.	b.		
Franklin et al. 1996 [46]	Virtual Mattie is a clerical agent that collects and shares weekly email information with a designated group.	✓	✓	✓	✓	✓. Extended versions of Maes' behavior net [47] and Hofstadter and Mitchell's Copycat architecture [48].	X
Franklin et al. 1998 [49]	IDA: Intelligent distribution agent for US Navy sailor billet assignments via email.	✓	✓	✓	✓	✓. Architecture based on global workspace theory. The basis for LIDA [50].	✓
Knoblock et al. 2003 [51]	A personal information agent that gathers information from sources on the internet and links them to the specific task.	✓	✓	✓	X	✓. Employs, in combination with others, the rule-based approach, Planning by Rewriting [52].	✓
Berry et al. 2004 [53]	PCaM is a personalized calendar agent that manages the calendar of an individual and coordinates as well as schedules meetings.	✓	✓	✓	X	✓. Utilizes the Open Agent Architecture to connect the cognitive components directly [54].	✓
Freed et al. 2008 [55]	RADAR is a personal assistant that reduces email overload by identifying task-relevant content and managing emails.	✓	✓	✓	X	X. Interconnected structure of task-specific, AI-enhanced modules/agents.	X
Berry et al. 2011 [56]	PTIME is a learning cognitive assistant agent that represents a personalized time manager.	✓	✓	✓	X	✓. Employs a preference module, which is connected to various schedulers.	✓
SRI International [57]	CALO is an extensive cognitive agent that learns and organizes. It is a predecessor of Apple Siri.	✓	✓	✓	X	✓. Cooperating cognitive agents including, for example, PTIME for language processing, see [58].	✓
Strain et al. 2014 [59]	Medical Agent X, an agent for clinical diagnostics.	✓	X	X	X	✓. Employs the cognitive architecture LIDA [50].	✓
Lebriere et al. 2015 [60]	An automated approach to determine the task of a piece of malware.	✓	X	X	X	✓. Partially employs the cognitive architecture ACT-R [61].	X
Shi et al. 2017 [62]	An experimental platform where agents learn to interact with web environments using only keyboard and mouse operations.	✓	✓	✓	✓	X. Training models based on supervised and reinforcement learning [62].	X
Wendt et al. 2018 [63]	A cognitive agent for building energy management.	✓	✓	✓	X	✓. Employs the KORE cognitive architecture for building automation [64].	✓

Mattie are aimed at understanding the free-form and probably incomplete messages. While Virtual Mattie implements mechanisms for goals and attention, the small number of cognitive abilities does not meet the respective criteria. The other criteria are satisfied.

IDA, the intelligent distribution agent designed by Franklin et al. in 1998 for the US Navy, is the successor of Virtual Mattie and the predecessor of the cognitive architecture LIDA [65]. The agent assigns new long-term tasks to sailors who have finished their current ones. A notable feature of the architecture is its working memory, based on the global workspace theory [66]. In comparison to Virtual Mattie, IDA possesses more enhanced cognitive abilities. In addition to decision-making and attention modules, IDA also includes a module for emotions. Due to the communication by email and similar related access to human coworkers, IDA, at a basic level, satisfies the criteria set that evaluates the systemic architectural perspective of integrating behavior and user similarity.

4.2 The CALO program (CALO, PTIME)

The Cognitive Assistant that Learns and Organizes (CALO) project brought together researchers from 22 organizations to advance research in cognitive software systems by developing a long-lasting, personalized cognitive agent [67]. The Defense Advanced Research Projects Agency (DARPA) funded it under the Perceptive Assistant that Learns (PAL) program. This project led to the CALO meeting assistant as well as a variety of related cognitive agents. One example developed to build on the foundation of CALO is PTIME [56]. With its ability to reason, plan schedules, and learn user preferences, PTIME meets the cognitive ability criteria set out in our methodology at a basic level. However, due to the implementation of the system with a user interface, it lacks the requested user similarity. The CALO meeting assistant integrates multiple previously developed cognitive concepts and possesses various cognitive abilities. It is the predecessor of Apple Siri. Nevertheless, in this study context, such systems are classified as conversational user interfaces rather than autonomous user agents.

4.3 Lebiere et al. 2015 and Wendt et al. 2018

Cognitive architectures can be directly applied to solve cognitive tasks previously performed by humans. In Lebiere et al. 2015, the cognitive architecture ACT-R is used to identify malware tasks [60]. The corresponding model was trained with historical malware data. While the agent system can perform its task by employing ACT-R, no other cognitive abilities were apparent. Additionally, the system must be initiated manually on the data, not meeting the respective criteria. Wendt et al. 2018 present a cognitive model capable of building energy management [63]. To achieve that, they employed the KORE cognitive architecture [64]. The agent is capable of being integrated externally into a building management interface. It processes data from physical sensors and establishes management rules based on that data. Like Lebiere et al.'s system, the agent does not employ all of the cognitive abilities in the criteria. However, some criteria, such as user

similarity, are satisfied with the active perception of the environment and the decision-making.

4.4 World of Bits (WoB)

WoB is an experimental learning platform to train software bots in open-domain web environments [62]. Agents in WoB perceive the web environments in the form of the Document Object Model (DOM) and rendered pixels. Interestingly, though benchmark results are low compared to human users, the approach accomplished some web tasks by sending mouse and keyboard actions. Furthermore, agents are separated from the environments in which they interact. Tasks and activities are low-level operations and lack high-level cognitive capabilities. As a result, it falls short of the evaluation criteria.

5 Discussion

The study’s results have important implications for the field of cognitive architectures and user-like software bots. Of the eleven agents and applied cognitive architectures examined in the previous section, Sec.4, only one satisfied the established criteria. Interestingly, the agent, IDA by Franklin et al. [49], was developed a little over two decades ago. Except for the first two candidates, which both precede the year 2000, all the systems featured missed the second criterion of user similarity. The two systems satisfying the criteria were limited by their time and would have been implemented as standalone systems if they had been designed a few years later. Presumably, due to a lack of alternatives, they both use email to communicate with their environment. This limitation likely led to them satisfying the user similarity criteria, not necessarily through the designer’s intent. The trend shown by the other candidates moved towards applications without user similarity.

5.1 Bots and Cognitive Architectures

One central question arising from the results is the absence of software bots currently employing cognitive architectures. To the best of our knowledge, it appears highly unlikely to find an agent or bot in current use that meets all the criteria established for this study. Despite significant advancements in other areas of artificial intelligence, the fact that only one bot, IDA by Franklin et al. [49], satisfied all the criteria suggests that the field of user-like intelligent bots is still nascent and in its early stages. These findings and observations align with existing literature, which indicates that integrating cognitive architectures into software bots is a complex and challenging endeavor [26, 68].

However, these results also underscore the potential that cognitive architectures hold for enhancing the capabilities of software bots. From a simplified architectural perspective, the example of IDA, which met all the criteria, demonstrates how a cognitive architecture can enable advanced autonomous behavior and user similarity in software bots.

These findings also have practical implications. These findings highlight the need for developers and researchers to focus their efforts on integrating cognitive architectures into software bots. For industry professionals, the results can guide the development of more advanced software bots that align with the requirements of cognitive automation, enabling more effective automation of various tasks across different domains.

5.2 Autonomy and Behavior: Architectural Perspectives

The architectural aspects that can address the observed shortcomings in this study, particularly regarding the level of autonomy and generalized behavior, can be viewed from two perspectives: the strict separation of bots and their environment and the integration of a separate behavior model [26].

First, as mentioned earlier, the strict separation of agents and their environment facilitates architectural possibilities for autonomy [26]. By decoupling bots from a single environment, they can gain the ability to interact with multiple environments dynamically, achieving higher levels of variability and adaptability. This separation also enables the design of software bots with a user-like orientation, treating bots as if they were human users. Consequently, this orientation opens up possibilities for alternative architectural design patterns, allowing software bots with user-like characteristics to seamlessly integrate into existing user interfaces or operate on digital platforms without the need for APIs or integration protocols. The separation already places the bots in a position where they can potentially establish their own intentions, goals, and deliberate interactions, access, and actions.

Second, an agent’s capacity for self-awareness, contextual awareness, and other high-level cognitive functions arises from holistic behavior models [26]. These behavior models are typically formulated using cognitive architectures or similar integrative models and architectures. Consequently, the componentization and integration of cognitive architectures into the system architecture of software bots become essential architectural considerations. The scientific understanding of machine intelligence, along with related models and principles, is still evolving and has not yet reached a crystallized state [17]. As a result, models of intelligent behavior will evolve over time, and older ones may require changes and modifications. Componentization can facilitate separation and change. Additionally, through this componentization, bots can potentially dynamically integrate separate behavior models in real-time.

Adopting these architectural perspectives can contribute to addressing the challenges related to autonomy and generalized behavior in bots, enabling them to operate with greater flexibility, adaptability, and user-like characteristics.

5.3 Study Limitations

Despite the study’s important findings and insights into integrating cognitive architectures into user-like software bots, the study has limitations. The review was limited to publicly available literature and may not capture all existing software

bots that employ cognitive architectures. Furthermore, the criteria for evaluating software bots may not encompass all possible features and capabilities. Future research could expand on our work by investigating databases further, employing different evaluation criteria, or examining the evolution of cognitive architectures in software bots over time.

6 Conclusion

The review highlights the importance of developing user-like software bots. These bots integrate cognitive architectures, enabling advanced autonomous behavior and user similarity. These requirements have significant implications for the engineering aspects of cognitive bots, particularly in cognitive automation.

Through the analysis of architectural recommendations and perspectives, this study has provided insights into achieving these goals at a design level. The distinctive view of bots and their environment, along with the dynamic integration of componentized behavior models, arise as key approaches to support both desiderata.

Consequently, implementing these architectural approaches can lead to increased autonomy and adaptability in software bots. This opens up new possibilities for developing autonomous user-like bots that can effectively perform a wide range of tasks. These findings contribute to the growing understanding of how to enhance the cognitive capabilities of software bots through architectural references.

Acknowledgement

We would like to thank Lorenz Bobber for his valuable support in the initial stages of the survey.

References

1. Jianjun Jiang, Yiqun Wang, Li Zhang, Daqing Wu, Min Li, Tian Xie, Pengcheng Li, Licao Dai, Peiyao Li, Xianyun Shi, et al. A cognitive reliability model research for complex digital human-computer interface of industrial system. *Safety science*, 108:196–202, 2018.
2. Christian Engel, Philipp Ebel, and Jan Marco Leimeister. Cognitive automation. *Electronic Markets*, 32(1):339–350, 2022.
3. Paulo Leitao, Stamatis Karnouskos, Luis Ribeiro, Jay Lee, Thomas Strasser, and Armando W Colombo. Smart agents in industrial cyber-physical systems. *Proceedings of the IEEE*, 104(5):1086–1101, 2016.
4. Jay Lee, Hossein Davari, Jaskaran Singh, and Vibhor Pandhare. Industrial artificial intelligence for industry 4.0-based manufacturing systems. *Manufacturing letters*, 18:20–23, 2018.
5. Stamatis Karnouskos, Paulo Leitao, Luis Ribeiro, and Armando Walter Colombo. Industrial agents as a key enabler for realizing industrial cyber-physical systems: Multiagent systems entering industry 4.0. *IEEE Industrial Electronics Magazine*, 14(3):18–32, 2020.

6. Lucija Ivančić, Dalia Suša Vugec, and Vesna Bosilj Vukšić. Robotic process automation: systematic literature review. In *BPM*, pages 280–295. Springer, 2019.
7. Linda Erlenhov, Francisco Gomes de Oliveira Neto, Riccardo Scandariato, and Philipp Leitner. Current and future bots in software development. In *2019 IEEE/ACM 1st International Workshop on Bots in Software Engineering (BotSE)*, pages 7–11. IEEE, 2019.
8. Mairieli Wessel, Bruno Mendes De Souza, Igor Steinmacher, Igor S Wiese, Ivanilton Polato, Ana Paula Chaves, and Marco A Gerosa. The power of bots: Characterizing and Understanding Bots in OSS Projects. *PACM:HCI*, 2(CSCW):1–19, 2018.
9. Martin Monperrus. Explainable software bot contributions: Case study of automated bug fixes. In *2019 IEEE/ACM 1st international workshop on bots in software engineering (BotSE)*, pages 12–15. IEEE, 2019.
10. Olivia B Newton, Samaneh Saadat, Jihye Song, Stephen M Fiore, and Gita Sukthankar. Everybody counts: Examining human–machine teams in open source software development. *Topics in Cognitive Science*, 2022.
11. Dimitrios Platis. Software development bot ecosystems. 2021.
12. Likoebe Maruping and Yukun Yang. Governance in digital open innovation platforms. In *Oxford Research Encyclopedia of Business and Management*. 2020.
13. Jim Hendler and Tim Berners-Lee. From the semantic web to social machines: A research challenge for ai on the world wide web. *Artificial intelligence*, 174(2):156–161, 2010.
14. Raúl Arrabales, Agapito Ledezma, and Araceli Sanchis. Towards conscious-like behavior in computer game characters. In *2009 IEEE symposium on computational intelligence and games*, pages 217–224. IEEE, 2009.
15. Florian Gallwitz and Michael Kreil. The Rise and Fall of ‘Social Bot’ Research. *SSRN*, 2021.
16. Kelsey R McDonald and John M Pearson. Cognitive bots and algorithmic humans: toward a shared understanding of social intelligence. *Current Opinion in Behavioral Sciences*, 29:55–62, 2019.
17. Patrick Butlin, Robert Long, Eric Elmoznino, Yoshua Bengio, Jonathan Birch, Axel Constant, George Deane, Stephen M Fleming, Chris Frith, Xu Ji, et al. Consciousness in artificial intelligence: Insights from the science of consciousness. *arXiv preprint arXiv:2308.08708*, 2023.
18. Stefan Kugele, Ana Petrovska, and Ilias Gerostathopoulos. Towards a taxonomy of autonomous systems. In *Software Architecture: 15th European Conference, ECSA 2021, Sweden, 2021*, pages 37–45. Springer, 2021.
19. Marialena Vagia, Aksel A Transeth, and Sigurd A Fjerdings. A literature review on the levels of automation during the years. what are the different taxonomies that have been proposed? *Applied ergonomics*, 53:190–202, 2016.
20. Gianluca Baldassarre, Vieri Giuliano Santucci, Emilio Cartoni, and Daniele Caligiore. The architecture challenge: Future artificial-intelligence systems will require sophisticated architectures, and knowledge of the brain might guide their construction. *Behavioral and Brain Sciences*, 40:e254, 2017.
21. Gerhard Kraetzschmar and D Deliverable. Software engineering factors for cognitive robotics. 2018.
22. Jan Bosch, Helena Holmström Olsson, and Ivica Crnkovic. Engineering ai systems: A research agenda. *Artificial Intelligence Paradigms for Smart Cyber-Physical Systems*, pages 1–19, 2021.
23. Ben Goertzel, Cassio Pennachin, and Nil Geisweiller. Engineering General Intelligence, Part 1. *Atlantis Thinking Machines*, 5, 2014.
24. Michael Wooldridge and Nicholas R Jennings. Pitfalls of agent-oriented development. In *Proceedings of the second international conference on Autonomous agents*, pages 385–391, 1998.
25. Silverio Martínez-Fernández, Justus Bogner, Xavier Franch, Marc Oriol, Julien Siebert, Adam Trendowicz, Anna Maria Vollmer, and Stefan Wagner. Software engineering for ai-based systems: a survey. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 31(2):1–59, 2022.

26. Habtom Kahsay Gidey, Peter Hillmann, Andreas Karcher, and Alois Knoll. Towards cognitive bots: Architectural research challenges. In *International Conference on Artificial General Intelligence*, pages 105–114. Springer, 2023.
27. Carlene Lebeuf, Margaret-Anne Storey, and Alexey Zagalsky. Software bots. *IEEE Software*, 35(1):18–23, 2017.
28. Michael Wooldridge and Nicholas R Jennings. Agent theories, architectures, and languages: a survey. In *Intelligent Agents: ECAI-94 Workshop on Agent Theories, Architectures, and Languages Amsterdam, The Netherlands August 8–9, 1994 Proceedings 1*, pages 1–39. Springer, 1995.
29. Carlene R. Lebeuf. *A taxonomy of software bots: towards a deeper understanding of software bot characteristics*. PhD thesis, UVic, 2018.
30. Thomas Holz, Abraham G Campbell, Gregory MP O’Hare, John W Stafford, Alan Martin, and Mauro Dragone. Mira—mixed reality agents. *IJHC*, 69(4), 2011.
31. Paul Milgram and Fumio Kishino. A taxonomy of mixed reality visual displays. *IEICE TRANSACTIONS on Information and Systems*, 77(12):1321–1329, 1994.
32. Victor Charpenay, Tobias Käfer, and Andreas Harth. A unifying framework for agency in hypermedia environments. In *Engineering Multi-Agent Systems: 9th International Workshop, EMAS 2021, Virtual Event, May 3–4, 2021, Revised Selected Papers*, pages 42–61. Springer, 2022.
33. James J Gibson. The theory of affordances. *Hilldale, USA*, 1(2):67–82, 1977.
34. Benjamin D Nye and Barry G Silverman. Affordances in ai, 2012.
35. David Vernon. Cognitive Architectures. In *Cognitive Robotics*. MIT Press, 2022.
36. David Vernon. *Artificial cognitive systems: A primer*. MIT Press, 2014.
37. Iuliia Kotseruba, Oscar J Avella Gonzalez, and John K Tsotsos. A review of 40 years of cognitive architecture research: Focus on perception, attention, learning and applications. *arXiv*, pages 1–74, 2016.
38. Torsten Metzler, Kristina Shea, et al. Taxonomy of cognitive functions. In *DS 68-7: Proceedings of the 18th International Conference on Engineering Design (ICED 11), Impacting Society through Engineering Design, Vol. 7: Human Behaviour in Design, Lyngby/Copenhagen, Denmark, 15.-19.08. 2011*, pages 330–341, 2011.
39. Wlodzislaw Duch, Richard Jayadi Oentaryo, and Michel Pasquier. Cognitive Architectures: Where do we go from here? In *AGI*, volume 171, pages 122–136, 2008.
40. John E Laird, Christian Lebiere, and Paul S Rosenbloom. A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *A Magazine*, 38(4):13–26, 2017.
41. Wil Van Der Aalst and Kees Max Van Hee. *Workflow management: models, methods, and systems*. 2004.
42. Wil MP Van der Aalst, Martin Bichler, and Armin Heinzl. Robotic process automation, 2018.
43. Dietmar Bruckner, Heimo Zeilinger, and Dietmar Dietrich. Cognitive automation—survey of novel artificial general intelligence methods for the automation of human technical environments. *IEEE Transactions on Industrial Informatics*, 8(2):206–215, 2011.
44. Santiago Aguirre and Alejandro Rodriguez. Automation of a business process using robotic process automation (rpa): A case study. In *Applied Computer Sciences in Engineering: 4th Workshop on Engineering Applications, WEA 2017, Cartagena, Colombia, September 27-29, 2017, Proceedings 4*, pages 65–71. Springer, 2017.
45. Claes Wohlin. Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the 18th international conference on evaluation and assessment in software engineering*, pages 1–10, 2014.
46. Stan Franklin, Art Graesser, Brent Olde, Hongjun Song, and Aregahegn Negatu. Virtual mattie—an intelligent clerical agent. In *AAAI Symposium on Embodied Cognition and Action*, Cambridge MA, 1996.
47. Pattie Maes. Situated Agents can have Goals. *Robotics and autonomous systems*, 6(1-2):49–70, 1990.
48. Douglas R Hofstadter and Melanie Mitchell. The Copycat Project: A model of mental fluidity and analogy-making. 1994.

49. Stan Franklin, Arpad Kelemen, and Lee McCauley. Ida: A cognitive agent architecture. In *SMC'98 Conference Proceedings. IEEE Conference on Systems, Man, and Cybernetics*, volume 3, pages 2646–2651. IEEE, 1998.
50. Stan Franklin, Steve Strain, Javier Snaider, Ryan McCall, and Usef Faghihi. Global workspace theory, its lida model and the underlying neuroscience. *Biologically Inspired Cognitive Architectures*, 1:32–43, 2012.
51. Craig A Knoblock. Deploying Information Agents on the Web. In *IJCAI*, pages 1580–1586. Citeseer, 2003.
52. José Luis Ambite and Craig A Knoblock. Planning by rewriting. *Journal of Artificial Intelligence Research*, 15:207–261, 2001.
53. Pauline Berry, Melinda Gervasio, Tomas Uribe, Karen Myers, and Ken Nitz. A Personalized Calendar Assistant. In *Working notes of the AAAI Spring Symposium Series*, volume 76, 2004.
54. Adam Cheyer and David Martin. The Open Agent Architecture. *Autonomous Agents and Multi-Agent Systems*, 4(1):143–148, 2001.
55. Michael Freed, Jaime G Carbonell, Geoffrey J Gordon, Jordan Hayes, Brad A Myers, Daniel P Siewiorek, Stephen F Smith, Aaron Steinfeld, and Anthony Tomasic. RADAR: A personal assistant that learns to reduce email overload. In *AAAI*, volume 8, pages 1287–1293, 2008.
56. Pauline M Berry, Melinda Gervasio, Bart Peintner, and Neil Yorke-Smith. PTIME: Personalized assistance for calendaring. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(4):1–22, 2011.
57. Apple Inc. Siri Intelligent Personal Assistant. <https://www.apple.com/uk/siri/>. Accessed: 18.06.2021.
58. Gokhan Tur, Andreas Stolcke, Lynn Voss, Stanley Peters, Dilek Hakkani-Tur, John Dowding, Benoit Favre, Raquel Fernández, Matthew Frampton, Mike Frandsen, et al. The CALO Meeting Assistant System. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(6):1601–1611, 2010.
59. Steve Strain, Sean Kugele, and Stan Franklin. The learning intelligent distribution agent (lida) and Medical Agent X (MAX): Computational intelligence for medical diagnosis. In *IEEE Symposium on Computational Intelligence for Human-like Intelligence (CIHLI)*, pages 1–8. IEEE, 2014.
60. Christian Lebiere, Stefano Bennati, Robert Thomson, Paulo Shakarian, and Eric Nunes. Functional cognitive models of malware identification. *Proceedings of ICCM, ICCM*, pages 9–11, 2015.
61. Frank E Ritter, Farnaz Tehranchi, and Jacob D Oury. ACT-R: A cognitive architecture for modeling cognition. *Wiley Reviews: Cognitive Science*, 10(3), 2019.
62. Tianlin Shi, Andrej Karpathy, Linxi Fan, Jonathan Hernandez, and Percy Liang. World of bits: An open-domain platform for web-based agents. In *ICML*, pages 3135–3144. PMLR, 2017.
63. Alexander Wendt, Stefan Kollmann, Lydia Siafara, and Yevgen Biletskiy. Usage of cognitive architectures in the development of industrial applications. In *proceedings of ICAART 2018*, 2018.
64. Gerhard Zucker, Alexander Wendt, Lydia Siafara, and Samer Schaat. A cognitive architecture for building automation. In *IECON 2016-42nd Annual Conference of the IEEE Industrial Electronics Society*, pages 6919–6924. IEEE, 2016.
65. Stan Franklin and FG Patterson Jr. The LIDA architecture: Adding new modes of learning to an intelligent, autonomous, software agent. *pat*, 703:764–1004, 2006.
66. Bernard J Baars. Global workspace theory of consciousness: Toward a cognitive neuroscience of human experience. *Progress in brain research*, 150:45–53, 2005.
67. Sri international: Artificial Intelligence Center - The CALO project. <http://www.ai.sri.com/project/CALO>. Accessed: 18.06.2021.
68. David Vernon. The architect’s dilemmas. *Cognitive Architectures*, 2019.