Invoking Principles of Groupware to Develop and Evaluate Present and Future Human-Agent Teams

Christopher Flathmann
cflathm@g.clemson.edu
Clemson University

Beau Schelble
bschelb@clemson.edu
Clemson University

Brock Tubre
btubre@clemson.edu
Clemson University

Nathan McNeese
mcneese@clemson.edu
Clemson University

Paige Rodeghero
prodegh@clemson.edu
Clemson University

ABSTRACT

Advances in artificial intelligence are constantly increasing its validity as a team member enabling it to effectively work alongside humans and other artificial teammates. Unfortunately, the digital nature of artificial teammates and their restrictive communication and coordination requirements complicate the interaction patterns that exist. In light of this challenge, we create a theoretical framework that details the possible interactions in human-agent teams, emphasizing interactions through groupware, which is based on literature regarding groupware and human-agent teamwork. As artificial intelligence changes and advances, the interaction in human-agent teams will also advance, meaning interaction frameworks and groupware must adapt to these changes. We provide examples and a discussion of the frameworks ability to adapt based on advancements in relevant research areas like natural language processing and artificial general intelligence. The results are a framework that detail human-agent interaction throughout the coming years, which can be used to guide groupware development.

CCS CONCEPTS

• **Human-centered computing → HCI theory, concepts and models**; • **Computing methodologies → Artificial intelligence**; **Cooperation and coordination**.

KEYWORDS

human-AI; human-agent interaction; human-agent teamwork; artificial intelligence; groupware

ACM Reference Format:

Christopher Flathmann, Beau Schelble, Brock Tubre, Nathan McNeese, and Paige Rodeghero. 2020. Invoking Principles of Groupware to Develop and Evaluate Present and Future Human-Agent Teams. In *Proceedings of the 8th International Conference on Human-Agent Interaction (HAI '20), November 10–13, 2020, Virtual Event, NSW, Australia.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3406499.3415072

HAI '20, November 10–13, 2020, Virtual Event, NSW, Australia

ACM ISBN 978-1-4503-8054-6/20/11…$15.00
https://doi.org/10.1145/3406499.3415072

1 INTRODUCTION

The field of human-agent teamwork research and the potential for implementing human-agent teams (HATs) in various domains advances; however, if these teams are going to function correctly, software tools that promote efficiency and productive team interactions need to be built with an understand of how human-agent team's interactions in software. Advancements in artificial intelligence (AI) have allowed machines to transition from simple tools to interactive teammates working alongside humans where artificial and human team members have specific roles, interdependent work, and common goals, resulting in human-agent interaction changes from static directional commands to more flexible and interactive communication [42]. With the potential for HATs to exist in a variety of different contexts and domains, such as medical [63] and military [13], the benefits for the introduction of human-agent teaming (HAT) have never been higher. This transition will be influenced by the software these teams use to collaborate. Groupware, the software designed to support teams and their interactions, has historically been challenging to develop due to the complexity that arises when multiple people are introduced into a collaborative context [27]. Limitations in the ability of modern AI teammates, chiefly communication deficiencies [14], introduce even more complexities to the creation of groupware. Features in human-agent interaction, such as cognitive monitoring, team organization, and communication aid, would become vital features in human-agent groupware. These groupware features need to be designed to improve interactions in human-agent teams. Thus, researchers in human-agent teaming need to collectively ensure that interactions are viewed from a technology and systems perspective; this view is necessary for ensuring that groupware systems are designed specifically for contextualized and changing human-agent interaction patterns.

The current paper combines important research in the areas of human-agent teaming and team interactions through groupware technology to create a high fidelity framework for human-agent interaction. This framework's unique design allows human-agent interaction to be better contextualized by ensuring that human-agent interaction is grounded in the tools that mediate interaction. This perspective and methodology has been shown as an effective way to evaluate human-agent interaction [35], and this paper extends such concepts by integrating important teamwork functions alongside the evolving field of AI. Our contextualization requires the fusion of both HAT research and groupware research to ensure the contribution of this framework advances both fields. Research in the areas of groupware and HATs is reviewed to ensure the framework considers the requirements and challenges noted and demonstrated by both domains. A theoretical interaction framework, with an emphasis on the software used for interaction, is designed to specifically detail how HAT's interact, which should be used as a guide when creating groupware.

Designed for modern HATs, the framework is then viewed in relation to the changing state of AI and artificial teammates, which are set to change as AI is further developed, and therefore the framework is set to advance alongside it. Natural language processing (NLP) and artificial general intelligence (AGI) are reviewed specifically to exemplify how their development and introduction would impact the design of the framework presented. The changes brought about by these technologies are detailed, alongside discussion as to why these technologies have the effects they do. Lastly, universal design features between all frameworks are highlighted to demonstrate the base requirements for any HAT groupware system. The core framework and the future technology discussion sections serve as a guide to developers and researchers that will focus on building software for the future of human-agent interactions in HATs.

2 BACKGROUND AND RELATED WORK

The development of HAT groupware requires a solid foundation in the research and principles of groupware technologies and HAT. This section details those two foundations and discusses the challenges and requirements of technology seeking to support teams in the manner presented in this paper. The details of historical development of groupware technology are examined, as HAT interactions will be influenced by these developments. Past literature on HATs regarding the development around these particular teams is reviewed, along with previous research and attempts in group support systems for HATs. The intentional intersection and collaboration between these two fields will be required to ensure a complete view of groupware is used to observe and define human-agent interaction in HATs.

2.1 Background In Groupware

2.1.1 Groupware History and Development. Groupware, defined in the early 1990s, is any computer-based system that supports individuals associated with one another through a common task or goal, providing them with an interface simulating a shared environment [18]. Software such as Lotus Notes exemplified early groupware, serving as a medium through which users could communicate, coordinate, and share documents [47]. Modern iterations of groupware like Slack, Microsoft Teams, and Discord offer multiple features relevant to communication, document sharing, and offer custom API integration for third party applications [10]. These groupware systems allow those individuals with a common goal to operate in a shared environment wherever they have access to a computational device and network connection. This connectivity introduced the possibility of distributed and remote work, addressed in the early 1990s with the introduction of the time-space taxonomy [18]. This time-space taxonomy was broken down into a 2x2 matrix that details the different interactions types shown in Figure 1.

The time-space taxonomy illustrates how groupware utilization in team contexts exists with varying degrees of time and distance

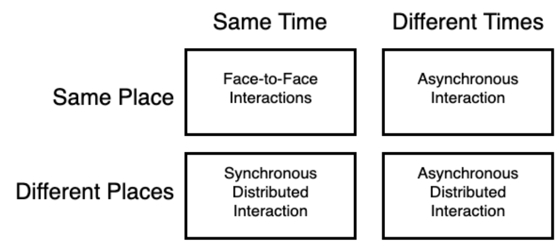

Figure 1: Time space taxonomy as described in [18].

barriers, which is a driving force for the existence of groupware. For example, software engineering teams that operate in a variety of different time zones around the world are subject to this matrix and become significantly enhanced by modern groupware systems [10]. Jonathan Grudin's 1994 paper on groupware and social dynamics outlined the early groupware development and research, stemming from the growth of technology in organizations that need to focus on organizational goals [27]. This focus on organizational goals alongside the creation of distributed personal computing through Ethernet, courtesy of Xerox PARC [32], led to the creation of the terms computer-supported cooperative work (CSCW) and computer human interaction (CHI). These new research areas were a direct outcome of the increased interest regarding human's enhancement through interaction with technology. Grudin identified eight challenges for groupware developers that represent roadblocks to groupware adoption within an organization, including: failure to have enough users adopt the system, adoption process mistakes, and the disruption of social processes. The disruption of social processes is a key problem many groupware systems have trouble overcoming, and the case study by Orlikowski on Lotus Notes highlighted many of these issues in practice [27, 47]. However, groupware has come a long way since Lotus Notes, e-mail, and electronic bulletin boards. Where groupware features were once fragmented across multiple programs, many modern day groupware systems effectively encompass all of the features once scattered across multiple different systems in the 1990s (ie video chat, text chat, document sharing, live document editing, file storage). The importance and ability of groupware systems have also never been more noticeable than during the 2020 SARS-CoV-2 global pandemic, which has seen millions of individuals across the globe working from home and utilizing a variety of groupware systems [9]. While modern groupware is highly advanced, with features such as video conferencing, calendar sharing, and text-based chat, it is also highly customizable with the use of custom API integration. This customizability can push groupware forward with more advanced features, such as the accommodation of AI teammates.

2.1.2 The Challenges and Requirements of Groupware. The requirements and challenges of groupware serve an important purpose in understanding team interaction as common interactions used to define software's requirements. These challenges become even more complex when viewed from the perspective of computers being social actors where the existence of groupware and the existence of AI teammates could modify the actual social interactions present in a team [45]. Thus, defining the requirements for groupware is a complex and difficult challenge, and errors in the requirements can lead to costly failures [7, 49]. Different parties that have a part in the development process play key roles in defining the interactions groupware needs to account for [7]. There are numerous considerations to make when starting the requirements gathering process for interactions covered by human-computer groupware systems. These include "communication, collaboration, cooperation, coordination, time, space, and awareness" [49]. Many research approaches, such as TRIDENT [6], CTT [48], Wisdom [46], UsiXML [37], FlowXML [24], focused on one or a combination of these, but rarely are all approaches considered simultaneously.

Penichet *et al.* presented a requirements gathering process called TOUCHE (Task-Oriented and User-Centred Process Model for Developing Interfaces for Human–Computer–Human Environments) which is a process model and methodology with a user interface to aid in the development of groupware applications [49]. TOUCHE helped throughout the development cycle of groupware, considering the human-computer interactions from the beginning of the development process [49]. During mobile groupware development, Huvnrylf *et al.* identified many of the requirements gathered actually come from the software developer's personal experience leading to creative uncertainty instead of an engineered process [29]. Huvnrylf proposed MCM (Mobile Collaboration Modelling), which is a role interaction notation that represented interactions among the participants in a human-driven mobile collaborative process, helping identify groupware requirements to be used in a model collaborative system [29].

Challenges in developing groupware is a highly researched area and recognized by Grudin's 1994 piece [27]. Although AI, development processes, and groupware requirements have changed since Grudin's writing, many of the challenges persist today. Bellika *et al.* published a 2008 piece with eight challenges when developing telemedicine groupware applications [4]. In 2016, research by Lazarin *et al.* focused on the challenges of developing web-based groupware and gather requirements using DPD (Distributed Participatory Design) practices [36]. Their review included studies based on co-located participants who were highly interactive and collaborative. They recognized the challenges of future interactions using the web-based groupware for distributed and remotely located collaborations [36]. The changes in these interactions and challenges become even more apparent when considering the media equation, where the integration of AI and new groupware technologies will see complex social perspectives formed upon their integration and advancement [54]. It is clear that as AI grows, so too with its interaction, which means the requirements and challenges of groupware will also evolve.

2.2 Background In Human-Agent Teamwork

Recent advancements in AI technology, along with the prevalence of teams in people's daily lives, have led to a surge in research around humans working in teams with AI. These teams consist of humans actively interacting with artificial teammates built with AI, which is an evolution upon previous integrations of more simplistic and less dynamic automation as a tool rather than a teammate [42]. While the potential for HATs is excellent in many fields, current limitations exist preventing modern HATs from reaching their potential. One of the most notable limitations involving human-agent interaction is AI's inability to communicate as well as humans, which is of vital importance to the efficient implementation of a HAT [12, 14, 34]. In addition to communication, other challenges face the implementation of HATs, such as status/intent monitoring and the directable ability of agents [34]. Adjoining human-agent teamwork is the more applied field of human-robot teaming and while human-robot teams can be seen as a type of HAT, their autonomous agents exist in a cyber-physcial domain allowing for different physical interactions not directly comparable to purely cyber agents [31]. Due to this key difference human-robot teams are outside the scope of the current research. This research is specifically focusing on the challenges that exist within HATs that utilize AI within a digital environment. In addition to these challenges, the future of AI research will determine other challenges that face AI's integration alongside humans.

While the transition to HATs will see additional requirements to properly function, foundational requirements of teamwork that have been identified in human-human teaming should be considered. Similar to human-human teaming, one of the cornerstones to proper teamwork for HATs is the explicit integration of team member roles [56, 59]. In addition to the definition of these roles, the workflow must be appropriately organized with roles in mind to ensure that time and energy are not repeatedly wasted when assigning work. This organization generally involves distributing functionalities across multiple AI as the specific intelligence of modern AI makes it more adept at individual tasks and roles [23]. As AI continues to develop, the roles that AI fill change as well, so HATs, and interactions within HATs, will need to change to account for the evolving roles of humans and agents [28]. In addition to role variability, context variability must also be considered, current roles and domains for AI integration have focused on its participation in medical [63] and military [13] teams; however, predictions see AI occupying a predominant role in other varied contexts, such as software development teams [61]. Developers and researchers need to be aware of these contexts if they intend to have a successful impact in understanding the interactions in HATs that need to be mediated by groupware.

While interactions in HATs can be diverse and complex, it is still important for groupware systems to cover the entire spectrum of HAT interaction. Past attempts at mediating these interactions through software have utilized AI as a tool to facilitate coordination and collaboration [53]. Other attempts have facilitated designs where human-agent collaboration is the primary concern, which is a critical research area, but human-human collaboration is also a key component to HAT interaction [65]. More recent designs implemented agents at both the support level and the teammate level in addition to the integration of multiple human teammates, but their communication design was still restrictive for a highly efficient team [38]. While the design of these new systems is essential, other researchers have identified the importance, validity, and appropriateness of standard software, such as Slack, to facilitate human-agent interaction [58]. This previous research illustrates the importance of considering how modern platforms can be modified and used to adequately meet the needs of these teams. Regardless of whether new software is going to be created or current software modified, each approach must ensure that the interaction requirements of HATs, both present, and future, are met. Without these requirements, the implemented groupware will not properly facilitate interactions within HATs.

3 A FRAMEWORK FOR HUMAN-AGENT INTERACTION THROUGH SOFTWARE

This paper creates a framework that details interactions in HATS, which can be utilized by future groupware developers and research to ensure groupware systems fully support HATs. The framework needs to consider and mediate the interactions between teammates, regardless of whether they are human. Other team support functionality should also be considered to ensure a multi-dimensional approach is taken to HAT interaction, rather than only viewing the single dimension of communication. These requirements have resulted in the creation of the framework shown in Figure 2. This framework provides a detailed view of the possible interactions in human-agent teams and how these interactions manifest through the use of groupware technologies.

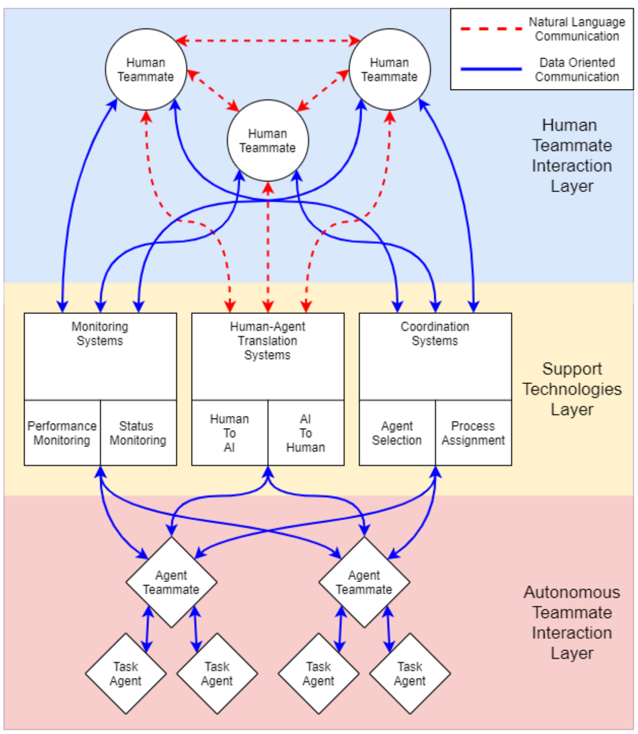

Figure 2: A framework for HAT interactions mediated through groupware. Separated into a layer for human integration, agent interaction, and a layer that supports the collaboration between the other two layers. Model design partially derived from [38] and [23].

The created framework has split team functionality between human functionality, agent functionality, and support functionality. This design was chosen because humans and agents have different requirements for how they interact in HATs. These differences, created by functional roles and AI limitations, have resulted in the need for a complex support technologies layer that promotes translation, coordination, and monitoring. It is also important to note that AI serves multiple roles within this framework. While explicit AI agents are outlined and detailed, the Support Technologies Layer could encorporate AI, strictly as a tool, to create robust support for the entire HAT. The following sections outline the current state and design of the layers in this framework and how each layer contributes to detailing the interactions in HATs from the perspective of the technologies they will use to interact. These interactions are then viewed in relation to modern groupware to see how well these technologies support human-agent teaming.

3.1 Human Teammate Layer

Human to human communication is a complex research area as many different modalities and contexts can be used to convey information. Additionally, teams are not all the same and the interactions they have can actually exist on a spectrum of different shared goals, shared mental frameworks, team-oriented motivation, trust with other teammates, and striving towards team objects are all relevant [26]. These teams move towards shared cognitive constructs through communication, which may take place in both physical and virtual settings using verbal and or nonverbal methods. Non-verbal communication includes gestures, smiling or frowning, widening eyes, and many more [16]. The level of complexity in human communication and successful teamwork makes integrating the two in HAT a considerable challenge. However, there is a great deal of research in human to AI communications, HATs, and modern software as groupware [5, 17, 21, 33]. With this prior research, along with the heightened interest in applying HATs to industry settings [66], the gap in knowledge is closing. The support technologies layer of the framework thus stands to mitigate any current difficulties in HAT and augment the team's capabilities. During the next sections, we discuss the various support technologies needed for possible interactions in HAT.

3.2 Support Technologies Layer

The Support Technologies Layer realizes the meaningful impact of the framework, as shown in the middle of Figure 2. This layer details the interactions that groupware needs to facilitate at the current stage of human-agent teaming. As outlined in the Human Teammate Layer, humans within the team are working together and communicating as usual through modalities offered by the groupware system. These functions, which are essential to humans, may not be the most effective method of collaboration between humans and the current state of artificial teammates. While this collaboration is made possible through the Support Technologies Layer, limitations in modern systems still limit these interactions, primarily in the way of a bottleneck that could occur by trafficking interaction through shared systems.

3.2.1 Monitoring & Coordination Systems. Through the Support Technologies Layer, the humans and AI members of the team take advantage of one another to enhance the team as a whole. Additionally, this layer is where specific technologies and features of a groupware system would mediate HAT interactions. It's important to retain a human-centered perspective when viewing these technologies as human accommodation still needs to take priority. Therefore, it is important that the processes of monitoring and coordination do not provide any extra cognitive load to the humans. As a result, the technologies used to interpret and utilize these interactions need to be hidden and implicit, focusing on understanding interactions that already exist rather than demanding new interactions be made. An excellent example of such a feature integrating human teamwork with the capabilities offered by AI is shown in Fan's 2010 work on modeling cognitive load for human team members through artificial agents [20]. Fan and Yen's implementation is an excellent example of what a monitoring and coordination system should accomplish in a groupware enabled HAT. Their system was able to work with a series of human-agent pairs to predict the human partner's current cognitive load status while supporting a graphical user interface (GUI) to enable a shared information space for team members. This is done through the creation of an architecture that integrates context, team member differences, cognitive loads, and available work to assign that work to AI that best empowers humans. As a result, a dynamic design is created where changes in monitoring are realized through changes in coordination, thus supporting the team without the need for explicit interaction, which allows for more practical and effective work assignment [20]. The system was also found to enhance the development of teams shared mental models, which are a significant predictor of team success and performance in a variety of contexts [40, 43]. However, the benefits of coordinating and monitoring systems would be handicapped without the ability for humans and autonomous agents to effectively communicate through the groupware.

3.2.2 Human Agent Translation Systems. The Human Agent Translation System serves as the bridge between the Human Teammate Interaction Layer and the Autonomous Teammate Interaction Layer, and the Monitoring System and the Coordinating System. The general goal of this support technology is to provide a more natural path for human-agent collaboration, which then helps human collaboration as it is natural to treat these technologies as social entities [54]. The result would be a system that humans naturally communicate with, which would in turn create agents that are naturally more similar to humans, a goal of human-agent teaming [42]. Additionally, a wealth of information about team interaction could be interpreted from these data without interrupting preexisting team processes. Specifically, this would be done by taking existing human level communication and translating it to agent level instructions, while also taking agent results and providing them to humans in an understandable form. Thus, new communication can be interpreted and utilized elsewhere in a transparent manner, but humans would still perceive a natural communication pattern, which would provide additional benefits to the HAT [8]. This actual process would be accomplished through NLP methods; however, the current state of NLP and the design of this technology could create a bottleneck that partially slows down communication. Future developments in NLP could help alleviate this bottleneck however, allowing for more efficient communication.

While Fan 2010 [20] is an example of how AI team members can benefit teams working together through groupware, it is not the only one. Other examples include multi-agent systems [68, 69], or HATs employing AI at different levels to support cognitive functions [19]. Based on examples from the literature, and current practice in human-AI interaction architectures, the framework presented in this paper utilizes the compartmentalization seen in the Support Technologies Layer within Figure 2. This approach allows the humans and agents to work together seamlessly [23], as well as allowing the implementation of various features offered by the groupware. Data-oriented communication created by both human and artificial teammates is leveraged to support the team as a whole through Monitoring and Coordination systems. Both of these systems serve as supporting technology and enhance the interaction between human and artificial teammates.

3.3 Autonomous Agent Layer

Modern HATs do not see AI teammates function or communicate in the same manner as human-human teammates, so it is essential to design around these differences to ensure they are properly accommodated. In general, these differences have led to a shift towards explainable systems that may use multiple agents rather than a single, black-box system [55]. These more transparent designs, along with the technical limitations of modern AI, have led to the creation of networks consisting of particular agents that specialize in their smaller individual tasks. These functional agents are then grouped to create a "single" AI teammate that completes the role it was assigned [23]. The design shown in Figure 2 accounts for this distributed method of design by allowing human teammates to be abstracted from a collection of single worker agents. For instance, in regards to a team that manages a UAV system, individual worker agents may be in charge of monitoring a single quadrant of a map, and the large agent would interpret the four quadrant worker agents to complete their goal of interpreting the entire map. Thus, It is also not enough for groupware to simply allow this clustering, the software must also allow for the dynamic nature of these clusters as their composition, and therefor interactions, will change based on the workload required [23]. The various agents imply the fact that groupware will need to adapt to a variety of different agent configurations without adding any cognitive burden to the humans in the team.

The model presented allows this dynamic nature while also promoting consistent communication between other agents and humans through the support technologies layer [38]. While communication does not happen directly between agents, due to the potential number of artificial teammates, communication is indirectly mediated through the designed support technologies layer. This layer will allow agents to learn from other agents as well as humans, making them better team members throughout their team's existence [11]. The human interaction could also contribute to transparency, a critical factor in building trust [64], and vital to efficient HAT [41]. Without these agent designs, trust in artificial teammates, and as a result, overall team performance could see significant drops. While groupware may not have the potential to eliminate the distrust in teammates, it must not cause the trust gap between teammates to grow. Groupware will need to assure the prioritization of human factors is allowed by agents if it is not going to hinder HAT performance.

3.4 Interactions With Modern Software

While the previous sections have outlined the interactions in HATs in regards to technology use, this section specifically demonstrates how the interaction model, shown in Figure 2, will apply to existing groupware systems. Interactions people have with groupware in modern software can be grouped similar to the structure of Figure 2: human-to-human, human-to-AI, and AI-to-AI [44]. Some modern groupware will use one of these communication mechanisms while others use a combination. For example, in 2010, Ball *et al.* developed groupware that was used to simulate the actions of a pilot for an Unmanned Aerial Vehicle (UAV) to match human behavior as closely as possible. This human-to-AI agent was integrated into a three-person team; however, this technology was heavily dependant on others for situational awareness when requesting or providing information to and from human teammates. For example, if the agent were continuously requesting images from the photography as she was actively taking pictures, this would interfere with the teams' success. This delay means that the agent must have a good understanding of situational awareness on its current tasks as well as its teammates' tasks [2]. This scenario also further highlights the necessity of the Support Technologies Layer to coordinate this vital flow of information and workload.

A more modern and widely used example that encapsulates all three ways groupware communication occurs is Slack [60]. Slack is similar to other messaging applications such as Skype, Google Hangouts, WhatsApp, Facebook Messenger, and many more [15]. Slack can be integrated into a team for human-to-human communication with text-based chat, video, and audio. Additionally, slack can also be integrated with other AI functionality for a team's specific needs. These can come in many different forms, but most are known as Slack bots. For example, AI agents exist to accompany team functions like task reminders, goal tracking, continuous integration and development monitoring, bug and issue tracking, and countless more. While these features are important to modern human-human teams, the manifestation of these functions will have to look different in HATs. An important aspect of the framework shown is the hidden aspect of the Support Technologies Layer. While modern iterations of Slack have bots that can aid in team functionality and interaction, the features outlined above are going to need to be ubiquitous to HATs. Modern groupware, such as Slack, won't have the functionality to create this ubiquitous support functionality due to its design for human-human teams. Another issue that is important in the framework is the clear distinction between human and artificial teammates, which will be especially important in the early stages of HATs where AI performance is not similar to humans. While Slack could represent artificial teammates similarly to human teammates but with AI specific names, this representation may be too similar to human teammates to create a clear distinction. Representing agents as simple bots may also prevent humans from seeing AI agents as teammates rather than tools. While new groupware systems may not be made from scratch for HATs, it is clear that the increased complexity of HAT interactions will not be fully facilitated by modern software. The detailed interactions outlined in Figure 2 need to be fully considered to create all-encompassing HAT groupware.

4 DISCUSSING A FUTURE ALONGSIDE ADVANCING AI

The rapidly advancing capabilities of AI due to increased research attention is a factor considered in the development of the proposed framework. To ensure the framework in this paper is viable in such a dynamic field, it is necessary to account for and predict how more advanced technologies will affect interactions in human-agent teams. This evaluation will changes in HAT interaction do not result in foundational changes to groupware. This section looks at two technologies that have the potential to impact HATs significantly: NLP and AGI. The framework in Figure 2 is modified to fit HATs with a consideration for interactions involving these more advanced AI. This evaluation of future HATs and their requirements shows the adaptability of this framework and how understanding HAT interactions is key to creating intelligent groupware

4.1 Natural Language Processing Impacts on Communication in Human-Agent Teams

One of the most prominent and rapidly developing fields in computer science with the potential to advance AI's potential in HATs is NLP [3]. In addition to being an essential factor in human-autonomy teaming, NLP is identified as an influential technology in a variety of other fields, including medical [30, 67], and financial [22]. The variety and multitude of domains and contexts in which NLP can be applied have led to substantial research, which is rapidly advancing it as a technology. One of the most notable advancements in this field is the release of OpenAi's powerful GPT-2 language generation model in 2019, a substantial improvement from the GPT model released only a year prior [51]; this advancement marks a critical milestone towards human-agent communication [52]. As these models continue to advance, human-agent communication, and thus HAT potential, will increase [12]. These advancements, along with the future path that NLP will take in human-autonomy teaming, have led to the modifications being made to the framework presented in Figure 2 to create the framework shown in Figure 3.

The first change that needs to be made to the framework is the increase in natural language used by humans towards agents. These continued advancements in NLP will lead to levels of communication similar to that between humans, and this human-like ability in AI agents can have a significant impact on HATs [42]. As NLP communication improves, explicit translation technologies will have a reduced need. Rather than having to go through a central layer devoted to supporting the conversion between human and agent communication, which would require less natural forms of communication, humans can simply converse directly with artificial teammates. These artificial teammates will each be equipped with these natural language processing abilities. In addition to the communication becoming more natural, this may also contribute to the personification of these agents, which enhances team communication between humans and their artificial teammates [57]. While this personification can be beneficial, this also increases the importance of preserving separation between human and artificial teammates to ensure humans remain in control and prioritized, which is key to ensuring humans properly utilize AI systems [1]. Additionally, this direct communication design would aid in reducing the bottleneck created by using single, centralized translation system.

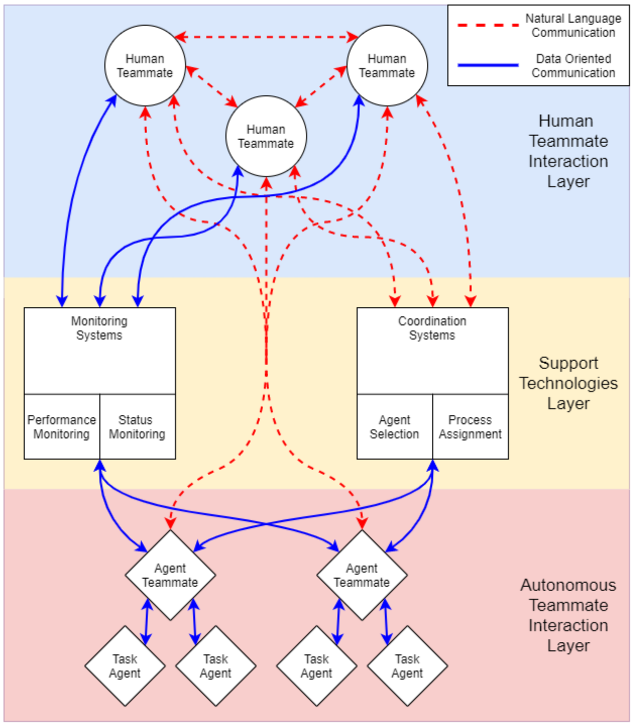

Figure 3: A framework for HAT interactions using groupware with a consideration for the future of NLP in AI. Explicit communication support systems are replaced with more traditional communication dynamics. Model design partially derived from [38] and [23].

While encouraging and aiding direct communication between agents and teammates is important, the preservation of the explicit coordination mechanism is still vital as it is an important mechanism in HATs. The use of these explicit coordination systems will be necessary to ensure the proper utilization of artificial teammates [23, 55]. However, the communication between this functionality and humans will be able to benefit from NLP. This is made possible due to the use of AI in these Support Technology Layer systems, which would allow the coordination system to possess NLP capabilities. While both the coordination systems and the artificial teammates both possess NLP capabilities, this communication design remains data-oriented as the data utilized is mostly derived from team interactions and actions. This new design in coordination will benefit human interaction with these systems while ensuring system-level performance is not compromised. Overall, advancement in NLP has the potential to be extremely beneficial to HATs. By utilizing the frameworks shown in Figure 3, a HAT can better transition to using agents with excellent NLP.

4.2 Artificial General Intelligence's Effect on Human Team Member's Perceptions

As it currently stands, AI is only capable of achieving a type of specific intelligence, imposing limitations on AI integration with teams. This specific intelligence is intelligence that only applies to the specific task the AI was trained for [39]. A negative side effect of this is the lack of adaptability in the artificial agent, leading artificial teammates to improperly function in unfamiliar tasks. For example, Deep Blue's knowledge would be specifically associated with the ability to play chess [62]. This inability to adapt to all tasks the team may undertake has the potential to undermine human team members' perception of the artificial teammate, contributing to negative affect and poor team dynamics [50]. For the framework in Figure 2 this adaptability is created by adjusting the individual agents that comprise a single artificial teammate. The downside to this workaround is the coordination and development of this large variety of agents, which still does come close replicating the benefits of AGI, or AI that have an expert level ability in all tasks [25]. However, at the moment, the literature has no single concrete solution or path towards AGI. Reviews such as [25] covered the current general intelligence research, and showcase its wide-ranging nature, with no one or two methodologies proving dominant. While the arrival of AGI is some time away, the framework presented in this paper should be able to adapt to such an important technology, leading to the creation of the framework in Figure 4.

Due to their high level of adaptability, the need for specific task agents group together by a teammate AI is not longer needed, as seen in Figure 4. With the loss of the individual task agents, the overall complexity of the system is greatly reduced, along with points of failure. The addition of general intelligence also allows the entire team to communicate through natural language, meaning all of the benefits identified in the NLP discussion would apply here as well. The enhanced ability to conduct any task allows the team to take on a much wider range of goals, no longer held back by the specific limitations of the task agents. The loss of the specific task agent limitations ensures the human team members are able to see the artificial teammate as a full member of the team, even in tasks that the agent is not specifically trained to complete, which is important for reducing negative attitudes towards the artificial agent, enhancing team dynamics [50].

Based on the review by [25], AGI stands to significantly alter the way team members within HATs interact with one another. The incredibly enhanced adaptability engenders greater responsibilities and a more diverse task load for the autonomous agents, ensuring they are given a larger role within the team. The outcome is realized as a more adaptable teammate that, from the surface, functions far closer to a human than a machine. This also allows teaming aspects like coordination to occur more naturally, while still being partially aided by support technologies. Due to the large impacts, AGI has the potential to create, the framework in this paper is designed to adapt to these more advanced artificial teammates.

4.3 A Summation of Overlapping Design Goals and Challenges

While differences exist between the three frameworks presented in this paper, it is still important to outline their similarities. These similarities are the core functionality of HATs. Without these similarities, HAT interaction will functionally change when AI advances, and groupware would also need to functionally change. This section covers the similarities between design goals and challenges created by these frameworks.

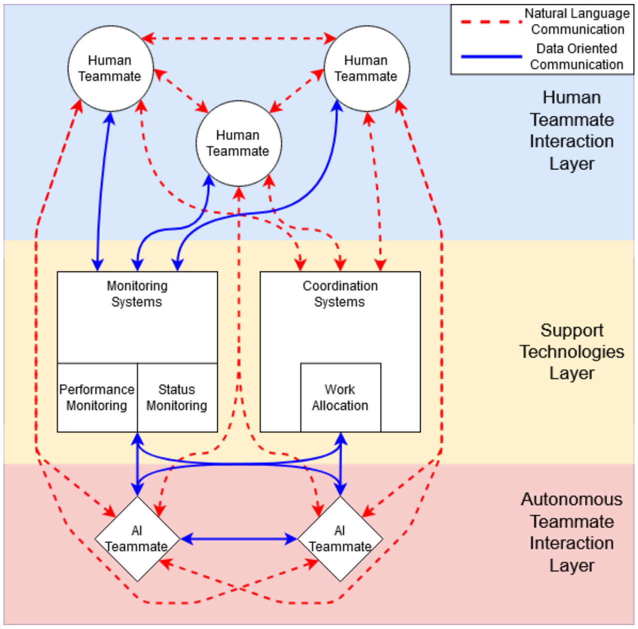

Figure 4: A framework for HAT interactions using groupware when AI agents have progressed to levels of AGI. In addition to communication becoming more natural, coordination is also more closely representative of human-human teamwork groupware. Model design partially derived from [38] and [23].

As communication advances through NLP and AGI, the support technologies follow AI personifications and mediums of communication. Areas of HCI with user experience concentrations will emerge to ensure the usability, utility, and interaction aesthetics give the best user experiences allowing humans to converse directly with artificial teammates in simple terms. Areas in visual and auditory interaction will need to be taken into considerations during the design phases of the future systems. In line with human-centered design, it is also improtant that the integration of systems do not disrupt pre-existing interaction. Additionally, even if NLP and AGI advancements are enough to where the AI is indistinguishable from human teammates, the groupware systems still require human peer analyzers, human and AI coordinators, and utilization metrics to ensure the AI groupware is integrated and work together seamlessly. The monitoring systems integrated as part of the support technology layer should also be as unobtrusive as possible and employ mechanisms to handle unexpected behavior with little to no human coordination. The roles within a team will also change, and this brings up challenges of deciding which roles are better for AI systems and which roles are better suited for human teammates. These role transitions should not only be possible, but aided by groupware. While positions such as leadership, management, HR, and many more will most likely be integrated as human teammates, AI's ability to aid in these roles still needs to be considered. This point also brings to light the priority that humans will need a level of global control [1], even as AI teammates become increasingly human-like. The challenge for designers will be ensuring all of these designs are met while still making groupware systems as ubiquitous as possible.

5 CONCLUSION

As the introduction and utilization of AI as a teammate continuously evolves, the technologies that support these teams must also evolve. To ensure this evolution is guided and purposeful, frameworks for HAT interaction need to be used during the development of groupware. These frameworks need to dynamically evolve based on multiple factors, including the state of artificial intelligence and the context in which they will function. Ensuring the adaptability of the frameworks proposed, a purposeful and interdisciplinary collaboration is required to ensure a multi-dimensional approach is used when evaluating HAT interactions. Without proper collaboration, ill-defined interaction designs will lead to ill-defined groupware. Due to the interdisciplinary nature of and the variety of contexts affected by HAT, the effects of this groupware, whether good or bad, will be impactful.

This paper established a framework to act as part of the guiding foundations for interactions in HATs that utilize technology for collaboration. The framework is created by synthesizing research from both groupware and HAT to ensure the foundations' interactions are considered when utilizing groupware technology. Additionally, this article not only explains why the dynamic nature of AI means that frameworks around HAT also need to be dynamic, but it demonstrates how the presented framework would be adapted appropriately to accommodate future HATs. This adaptation is key to ensuring the evolution of HAT interaction is met with the progressive evolution of groupware technology. Similar to the framework presented, this adaptation needs to be continuous and iterative to prevent massive redesigns whenever AI advances. Without this design, HATs will either be left with either inferior groupware that does not meet their specific requirements or outdated groupware that prevents the use of state-of-the-art AI. Either way, HATs will be heavily undermined if HAT interaction is not explicitly considered when creating groupware.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1829008.

REFERENCES

[1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.

[2] Jerry Ball, Christopher Myers, Andrea Heiberg, Nancy J Cooke, Michael Matessa, Mary Freiman, and Stuart Rodgers. 2010. The synthetic teammate project. *Computational and Mathematical Organization Theory* 16, 3 (2010), 271–299.

[3] Michael J. Barnes, Shan Lakhmani, Eric Holder, and J. Y. Chen. 2019. *Issues in Human-Agent Communication*. Technical Report ARL-TR-8636. Army Research Laboratory Aberdeen Proving Ground United States. https://apps.dtic.mil/docs/citations/AD1067050

[4] PHJ Bellika, Monika Johansen, Per Hasvold, Johanna Nystad, Per-Atle Bakkevoll, and Gunnar Hartvigsen. 2008. Eight Challenges for Developing Telemedicine Applications. *The Journal on Information Technology in Healthcare* 6, 4 (2008), 295–302.

[5] Mehmet Ilker Berkman, Dilek Karahoca, and Adem Karahoca. 2018. A measurement and structural model for usability evaluation of shared workspace groupware. *International Journal of Human–Computer Interaction* 34, 1 (2018), 35–56.

[6] F Bodart. 1994. alli. A Model-Based Approach to Presentation: A Continuum from task Analysis to Prototype. *Proc. of DSVIS'94* (1994).
[7] Barry Boehm, Paul Grunbacher, and Robert O Briggs. 2001. Developing groupware for requirements negotiation: lessons learned. *IEEE software* 18, 3 (2001), 46–55.
[8] Michael W Boyce, Jessie YC Chen, Anthony R Selkowitz, and Shan G Lakhmani. 2015. Effects of agent transparency on operator trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*. 179–180.
[9] Erik Brynjolfsson, J Horton, Adam Ozimek, Daniel Rock, Garima Sharma, and Hong Yi Tu Ye. 2020. COVID-19 and Remote Work: An Early Look at US Data.
[10] Fabio Calefato and Christof Ebert. 2019. Agile Collaboration for Distributed Teams [Software Technology]. *IEEE Software* 36, 1 (2019), 72–78.
[11] Micah Carroll, Rohin Shah, Mark K Ho, Tom Griffiths, Sanjit Seshia, Pieter Abbeel, and Anca Dragan. 2019. On the Utility of Learning about Humans for Human-AI Coordination. In *Advances in Neural Information Processing Systems*. 5175–5186.
[12] Mustafa Demir, Aaron D. Likens, Nancy J. Cooke, Polemnia G. Amazeen, and Nathan J. McNeese. 2019. Team Coordination and Effectiveness in Human-Autonomy Teaming. *IEEE Transactions on Human-Machine Systems* 49, 2 (April 2019), 150–159. https://doi.org/10.1109/THMS.2018.2877482 Conference Name: IEEE Transactions on Human-Machine Systems.
[13] Mustafa Demir, Nathan J. McNeese, and Nancy J. Cooke. 2016. Team communication behaviors of the human-automation teaming. In *2016 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. 28–34. https://doi.org/10.1109/COGSIMA.2016.7497782 ISSN: 2379-1675.
[14] Mustafa Demir, Nathan J. McNeese, and Nancy J. Cooke. 2017. Team situation awareness within the context of human-autonomy teaming. *Cognitive Systems Research* 46 (Dec. 2017), 3–12. https://doi.org/10.1016/j.cogsys.2016.11.003
[15] Sebastian Dennerlein, Robert Gutounig, Eva Goldgruber, and Stefan Schweiger. 2016. Web 2.0 messaging tools for knowledge management? Exploring the potentials of slack. In *European Conference on Knowledge Management*. Academic Conferences International Limited, 225.
[16] Joseph A DeVito, Susan O'Rourke, and Linda O'Neill. 2000. *Human communication.* Longman.
[17] Guillaume Dubuisson Duplessis, Chloé Clavel, and Frédéric Landragin. 2017. Automatic measures to characterise verbal alignment in human-agent interaction.
[18] Clarence A Ellis, Simon J Gibbs, and Gail Rein. 1991. Groupware: some issues and experiences. *Commun. ACM* 34, 1 (1991), 39–58.
[19] Xiaocong Fan, Bingjun Sun, Shuang Sun, Michael McNeese, and John Yen. 2006. RPD-enabled agents teaming with humans for multi-context decision making. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. 34–41.
[20] Xiaocong Fan and John Yen. 2010. Modeling cognitive loads for evolving shared mental models in human–agent collaboration. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 41, 2 (2010), 354–367.
[21] Sara Feldman, Ozge Nilay Yalcin, and Steve DiPaola. 2017. Engagement with artificial intelligence through natural interaction models. *Electronic Visualisation and the Arts (EVA 2017)* (2017), 296–303.
[22] Ingrid E. Fisher, Margaret R. Garnsey, and Mark E. Hughes. 2016. Natural Language Processing in Accounting, Auditing and Finance: A Synthesis of the Literature with a Roadmap for Future Research. *Intelligent Systems in Accounting, Finance and Management* 23, 3 (2016), 157–214. https://doi.org/10.1002/isaf.1386 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/isaf.1386.
[23] Christopher Flathmann, Nathan McNeese, and Lorenzo Barberis Canonico. 2019. Using Human-Agent Teams to Purposefully Design Multi-Agent Systems. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 63, 1 (Nov. 2019), 1425–1429. https://doi.org/10.1177/1071181319631238 Publisher: SAGE Publications Inc.
[24] Josefina Guerrero Garcia, Jean Vanderdonckt, and Juan Manuel Gonzalez Calleros. 2008. FlowiXML: a step towards designing workflow management systems. *International Journal of Web Engineering and Technology* 4, 2 (2008), 163.
[25] Ben Goertzel. 2014. Artificial general intelligence: concept, state of the art, and future prospects. *Journal of Artificial General Intelligence* 5, 1 (2014), 1–48.
[26] Victoria Groom and Clifford Nass. 2007. Can robots be teammates?: Benchmarks in human–robot teams. *Interaction Studies* 8, 3 (2007), 483–500.
[27] Jonathan Grudin. 1994. Groupware and social dynamics: Eight challenges for developers. *Commun. ACM* 37, 1 (1994), 92–105.
[28] William Halal, Jonathan Kolber, Owen Davies, and T Global. 2017. Forecasts of AI and future jobs in 2030: Muddling through likely, with two alternative scenarios. *Journal of Futures Studies* 21, 2 (2017), 83–96.
[29] Valeria Herskovic, Sergio F Ochoa, and José A Pino. 2019. Identifying groupware requirements in people-driven mobile collaborative processes. *J Univers Comput Sci* 25, 8 (2019), 988–1017.
[30] Jason K. Hou, Timothy D. Imler, and Thomas F. Imperiale. 2014. Current and Future Applications of Natural Language Processing in the Field of Digestive Diseases. *Clinical Gastroenterology and Hepatology* 12, 8 (Aug. 2014), 1257–1261. https://doi.org/10.1016/j.cgh.2014.05.013
[31] Chien-Ming Huang and Bilge Mutlu. 2016. Anticipatory robot control for efficient human-robot collaboration. In *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 83–90.
[32] Jeff Johnson, Teresa L. Roberts, William Verplank, David Canfield Smith, Charles H. Irby, Marian Beard, and Kevin Mackey. 1989. The xerox star: A retrospective. *Computer* 22, 9 (1989), 11–26.
[33] Hanbyul Joo, Tomas Simon, Mina Cikara, and Yaser Sheikh. 2019. Towards Social Artificial Intelligence: Nonverbal Social Signal Prediction in A Triadic Interaction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 10873–10883.
[34] G. Klien, D.D. Woods, J.M. Bradshaw, R.R. Hoffman, and P.J. Feltovich. 2004. Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems* 19, 6 (Nov. 2004), 91–95. https://doi.org/10.1109/MIS.2004.74 Conference Name: IEEE Intelligent Systems.
[35] Matti Krüger, Christiane B Wiebel, and Heiko Wersing. 2017. From tools towards cooperative assistants. In *Proceedings of the 5th International Conference on Human Agent Interaction*. 287–294.
[36] Carlos Alberto Joia Lazarin and Leonelo Dell Anhol Almeida. 2016. Distributed participatory design web-based groupware: gathering requirements through braindraw. In *Proceedings of the 15th Brazilian Symposium on Human Factors in Computing Systems*. 1–10.
[37] Quentin Limbourg, Jean Vanderdonckt, Benjamin Michotte, Laurent Bouillon, Murielle Florins, and Daniela Trevisan. 2004. Usixml: A user interface description language for context-sensitive user interfaces. In *Proceedings of the ACM AVI'2004 Workshop" Developing User Interfaces with XML: Advances on User Interface Description Languages"(Gallipoli*. 55–62.
[38] Charles Z. Liu and Manolya Kavakli. 2018. s — Part B: Agent-based collaborative information processing and coordination. In *2018 13th IEEE Conference on Industrial Electronics and Applications (ICIEA)*. 633–638. https://doi.org/10.1109/ICIEA.2018.8397792 ISSN: 2158-2297.
[39] Huimin Lu, Yujie Li, Min Chen, Hyoungseop Kim, and Seiichi Serikawa. 2018. Brain intelligence: go beyond artificial intelligence. *Mobile Networks and Applications* 23, 2 (2018), 368–375.
[40] John E Mathieu, Tonia S Heffner, Gerald F Goodwin, Eduardo Salas, and Janis A Cannon-Bowers. 2000. The influence of shared mental models on team process and performance. *Journal of applied psychology* 85, 2 (2000), 273.
[41] Nathan McNeese, Mustafa Demir, Erin Chiou, Nancy Cooke, and Giovanni Yanikian. 2019. Understanding the role of trust in human-autonomy teaming. In *Proceedings of the 52nd Hawaii international conference on system sciences.*
[42] Nathan J. McNeese, Mustafa Demir, Nancy J. Cooke, and Christopher Myers. 2018. Teaming With a Synthetic Teammate: Insights into Human-Autonomy Teaming. *Human Factors* 60, 2 (March 2018), 262–273. https://doi.org/10.1177/0018720817743223 Publisher: SAGE Publications Inc.
[43] Susan Mohammed, Lori Ferzandi, and Katherine Hamilton. 2010. Metaphor no more: A 15-year review of the team mental model construct. *Journal of management* 36, 4 (2010), 876–910.
[44] Yi Mou and Kun Xu. 2017. The media inequality: Comparing the initial human-human and human-AI social interactions. *Computers in Human Behavior* 72 (2017), 432–440.
[45] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 72–78.
[46] Nuno Jardim Nunes and João Falcão e Cunha. 2000. Wisdom—A UML based architecture for interactive systems. In *International Workshop on Design, Specification, and Verification of Interactive Systems*. Springer, 191–205.
[47] Wanda J Orlikowski. 1992. Learning from notes: Organizational issues in groupware implementation. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*. 362–369.
[48] Fabio Paternò, Cristiano Mancini, and Silvia Meniconi. 1997. ConcurTaskTrees: A diagrammatic notation for specifying task models. In *Human-computer interaction INTERACT'97*. Springer, 362–369.
[49] Victor MR Penichet, Maria D Lozano, José A Gallud, and Ricardo Tesoriero. 2010. Requirement-based approach for groupware environments design. *Journal of Systems and Software* 83, 8 (2010), 1478–1488.
[50] Mark S Pfaff. 2012. Negative affect reduces team awareness: The effects of mood and stress on computer-mediated team communication. *Human Factors* 54, 4 (2012), 560–571.
[51] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. *URL https://s3-us-west-2. amazonaws. com/openai-assets/researchcovers/languageunsupervised/language understanding paper. pdf* (2018).
[52] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. (2019).
[53] Heri Ramampiaro, Monica Divitini, and Sobah Abbas Petersen. 1999. Agent-based groupware: Challenges for cooperative transaction models. In *Proc. of the International Process Technology Workshop*.

[54] Byron Reeves and Clifford Ivar Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places.* Cambridge university press.

[55] Karina A Roundtree, Michael A Goodrich, and Julie A Adams. 2019. Transparency: Transitioning from human–machine systems to human-swarm systems. *Journal of Cognitive Engineering and Decision Making* 13, 3 (2019), 171–195.

[56] Eduardo Salas, Dana E. Sims, and C. Shawn Burke. 2005. Is there a "Big Five" in Teamwork? *Small Group Research* 36, 5 (Oct. 2005), 555–599. https://doi.org/10.1177/1046496405277134 Publisher: SAGE Publications Inc.

[57] Shruti Sannon, Brett Stoll, Dominic DiFranzo, Malte Jung, and Natalya N. Bazarova. 2018. How Personification and Interactivity Influence Stress-Related Disclosures to Conversational Agents. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '18)*. Association for Computing Machinery, Jersey City, NJ, USA, 285–288. https://doi.org/10.1145/3272973.3274076

[58] Kazuto Sasai, Takuya Hoshino, Hiroki Nagasawa, Ryota Fukutani, and Tetsuo Kinoshita. 2018. A Flexible Collaboration Mechanism Between Humans and Agents on Disaster Recovery of Network Systems. In *2018 IEEE 7th Global Conference on Consumer Electronics (GCCE)*. 219–220. https://doi.org/10.1109/GCCE.2018.8574872 ISSN: 2378-8143.

[59] Maarten Sierhuis, Jeffrey M Bradshaw, Alessandro Acquisti, Ron Van Hoof, Renia Jeffers, and Andrzej Uszok. 2003. Human-agent teamwork and adjustable autonomy in practice. In *Proceedings of the seventh international symposium on artificial intelligence, robotics and automation in space (I-SAIRAS)*.

[60] Slack. [n.d.]. Where work happens. https://slack.com/

[61] Neel Sundaresan. 2019. From Code to Data: AI at Scale for Developer Productivity. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. Association for Computing Machinery, Anchorage, AK, USA, 3175. https://doi.org/10.1145/3292500.3340410

[62] Peter Voss. 2007. Essentials of general intelligence: The direct path to artificial general intelligence. In *Artificial general intelligence*. Springer, 131–157.

[63] Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H. Beck. 2016. Deep Learning for Identifying Metastatic Breast Cancer. (June 2016). https://arxiv.org/abs/1606.05718v1

[64] Ning Wang, David V Pynadath, and Susan G Hill. 2016. Trust calibration within a human-robot team: Comparing automatically generated explanations. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 109–116.

[65] Justin D. Weisz, Mohit Jain, Narendra Nath Joshi, James Johnson, and Ingrid Lange. 2019. BigBlueBot: teaching strategies for successful human-agent interactions. In *Proceedings of the 24th International Conference on Intelligent User Interfaces (IUI '19)*. Association for Computing Machinery, Marina del Ray, California, 448–459. https://doi.org/10.1145/3301275.3302290

[66] H James Wilson and Paul R Daugherty. 2018. Collaborative intelligence: humans and AI are joining forces. *Harvard Business Review* 96, 4 (2018), 114–123.

[67] Adrian Wong, Joseph M. Plasek, Steven P. Montecalvo, and Li Zhou. 2018. Natural Language Processing and Its Implications for the Future of Medication Safety: A Narrative Review of Recent Advances and Challenges. *Pharmacotherapy: The Journal of Human Pharmacology and Drug Therapy* 38, 8 (2018), 822–841. https://doi.org/10.1002/phar.2151 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/phar.2151.

[68] Xiao-Feng Xie and Jiming Liu. 2008. Multiagent optimization system for solving the traveling salesman problem (TSP). *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 39, 2 (2008), 489–502.

[69] John Yen, Xiaocong Fan, Shuang Sun, Timothy Hanratty, and John Dumer. 2006. Agents with shared mental models for enhancing team decision makings. *Decision Support Systems* 41, 3 (2006), 634–653.