THE CYBER DEFENSE REVIEW
2025, VOL. 10, NO. 2, 225-240
https://doi.org/10.55682/cdr/25ae-dftk

RESEARCH ARTICLE

A Human-AI Teaming Approach to Closing the Talent Gap in Critical Infrastructure

Allyson Hauptman*

Army Cyber Institute, West Point, NY, USA

Many critical infrastructure sectors are facing significant talent gaps among their workforce. The Industrial Internet of Things revolution has introduced new technologies and requirements for workers to understand while continuing to perform the duties for which they were hired, and the introduction of these data-driven technologies has concurrently created the need for new team roles with their own sets of capabilities. One possible solution for overcoming these talent gaps is the integration of artificially intelligent teammates. Research suggests human-AI teaming could potentially offload tedious, repetitive, or dangerous human work and accomplish tasks that, while difficult for a human to complete, cater well to what computers do best. This paper proposes a simple 3-steps guiding framework for teams in critical infrastructure organizations to determine a) the gaps on their team, by distinguishing between gaps caused by insufficient personnel (capacity) and those driven by new technological demands (capability), b) which roles are well-suited for an AI teammate, based on the match between task demands and AI capabilities, and c) the human-centered design considerations, including presence, explainability, autonomy management, and ethical alignment, that are essential to its integration as an effective teammate.

Keywords: human-AI teaming, AI teammates, critical infrastructure, talent management, cyber defense workforce, human-centered design, Internet of Things, IoT

* Corresponding author: allyson.hauptman@westpoint.edu

1 INTRODUCTION

The increased use of and reliance on technology has created a more connected and efficient society than ever before; yet, these technologies simultaneously strain and challenge our critical infrastructure sectors to keep up. Many critical infrastructure sectors, such as energy (Fiume et al. 2006) and the Defense Industrial Base (DIB) (Goure 2018), face significant talent gaps among their workforce. There are multiple reasons behind this talent gap. First, some sectors are experiencing a mass wave of employee turnover and are struggling to backfill these positions (Fiume et al. 2006). Second, the Industrial Internet of Things (IIoT) revolution has introduced new technologies and requirements for workers to understand and operate while continuing to perform the duties for which they were hired (O'Heir 2017). Finally, the introduction of new data-driven technologies, such as artificial intelligence (AI), has created the need for the development of brand-new team roles with their own sets of capabilities. For industries already struggling to hire and retain enough workers to operate, these additional challenges are causes for immediate concern (Lewis 2019).

While the increased use of technology in various critical infrastructure sectors has created these problems, it may also be a solution. Human-AI teams (HATs) are a recent focus of AI research due to the possibilities they present to offload human workload and accomplish tasks that, while difficult for a human to complete, cater well to what computers do best. AI can be programmed to complete some tasks that teams do not have the manpower (O'Heir 2017) and/or capabilities to complete. An essential part of this construct's success is that the AI is not treated as just another tool that a human must manage; rather, the AI is an autonomous agent that fulfills a completely interdependent team role (O'Neill et al. 2022). In this way, the team's overall workload is reduced while its ability to tackle new challenges is enhanced.

While HATs have been explored as a performance-enhancing construct in a variety of environments, there is a significant lack of research on their role in critical infrastructure, particularly in addressing the growing workforce challenges. This paper explains the need for such a solution and presents a simple guiding framework for critical infrastructure teams to determine what gaps might be filled by an AI teammate, the scope of the AI teammate's role, and the human-centered design considerations the team needs to address in integrating the AI as an effective teammate.

The development of this framework was guided by the following research question:

How can a human-AI teaming guiding framework help critical infrastructure organizations identify, scope, and fill talent gaps essential for cyber defense?

2 BACKGROUND

2.1 The Development of a Talent Gap in Critical Infrastructure

Critical infrastructure encompasses the systems and assets vital to a nation's well-being and security, and can be divided into multiple sectors, including energy, transportation, financial services, and the Defense Industrial Base (Wagner 2021). Most critical infrastructure sectors, while overseen by various government agencies, are predominantly privately owned and operated and are highly dependent on uninterrupted operations to be successful (Durkovich 2020). This condition of uninterrupted operations creates a very high demand and workload for the professionals who operate in these sectors. Critical infrastructure sectors are interdependent, and thus when the sector is unable to provide continuous service, it affects the operations of various other sectors (Wagner 2021). For example, if a power company is unable to provide its promised output to the city it serves, that would directly affect the city's transportation services, communication services, healthcare services, and more are deeply affected.

This interdependency has increased with the emergence of cyber critical infrastructure and the vital computer services that keep all sectors running. As sectors introduce IoT devices to connect operational technologies to information technologies, both for increased efficiency and convenience, they simultaneously introduce vulnerabilities to cyber attacks (Malatji, Marnewick, and von Solms 2022). These vulnerabilities have several important implications for critical infrastructure and its workforce.

First, cybersecurity is now an essential additional layer that all critical infrastructure organizations need to build into their talent pools (Malatji, Marnewick, and von Solms 2022). While simple firewalls and scans may have been sufficient for systems that were segmented from outside threats, increased connectivity substantially increases an organization's risk of exploitation and thus requires dedicated manpower operating on dedicated equipment.

Second, organizations must be able to recognize and remediate issues caused by a cyber attack. Many critical infrastructure systems rely on operational technologies that require highly specialized expertise. Often, organizations lack these skills in-house. As a result, they frequently depend on contracted specialists, and only when they can first detect that something is wrong. A cyber attack can easily undermine this model if the organization is unable to identify that a system has been compromised or is operating abnormally (Thomas and Chothia 2021). These limitations compound an already serious challenge, leaving many organizations vulnerable to prolonged disruption and delayed response.

Research has found that a talent gap exists in critical infrastructure sectors involving both technical capabilities and sector-specific managerial skills needed to manage, secure, and implement technologies. The literature reveals a number of drivers of this gap, including an aging workforce (Ashworth 2006), diminishing vendor-locked knowledge (Sandborn and

Prabhakar 2015), and a younger generation seeking better working conditions and pay (Ayodele, Chang-Richards, and Gonzalez 2020). As it stands, many sectors are nearly stretched to the breaking point of human resources, facing urgent shortages in cybersecurity and operational technology roles.

In response, several initiatives are exploring how to rebuild the cyber talent pipeline. One promising development is the growth of cybersecurity clinics: university-based, service-learning programs that pair students with local governments, utilities, hospitals, and nonprofits to provide real cyber defense support while training the next generation of practitioners (Asare et al. 2025). These clinics provide students with hands-on experience working with real systems and threats, while offering critical infrastructure organizations a much-needed cyber capacity at a low cost. Recent national investments are expanding these clinics nationwide to address the chronic shortage of entry-level talent. Beyond clinics, sectors have experimented with a range of alternative training-focused strategies. A recent study sought to understand how tailored training could help overcome this gap and found that various sectors have experimented with simulation-based learning and virtual training environments to accelerate workforce preparation (Olonilua and Aliu 2025). In parallel, gamified training approaches have been developed to both attract new talent and enhance retention among the critical infrastructure workforce (Ashley et al. 2022). These efforts represent meaningful progress, yet they remain focused on mitigating the talent gap through human training alone. Given the scale and persistence of the problem, it is worth expanding the question: instead of training humans alone, why not explore whether AI agents can fill some of these gaps directly?

2.2 Human-AI Teaming

The emergence of AI as an everyday technology, sometimes called the Artificial Intelligence Revolution, has quickly reversed fears that AI would never succeed in applied settings (Li 2020). The ability to design AI to accomplish complex tasks with increasing levels of autonomy enables the use of a new team construct, where the AI is viewed as a teammate rather than a tool. A HAT consists of at least one human and one AI agent working interdependently in pursuit of shared goals (O'Neill et al. 2022). In other words, the AI agent must take on a fully independent team role that is integral to the team's accomplishment of overarching goals. For this to occur, the AI agent needs to operate with a minimum level of autonomy. HAT research has shown this minimum level to be that of partial autonomy, where the AI can fully execute a task once approved by a human (O'Neill et al. 2022). As such, it is essential to note that the AI referred to in this paper encompasses a wide range of technologies capable of acting as autonomous agents.

HAT constructs can provide significant advantages to teams, particularly those oriented around technical goals to which AI capabilities are naturally suited. HATs can work together to overcome challenges that all human teams struggle to overcome, such as tasks requiring complex data operations (Nyre-Yu, Gutzwiller, and Caldwell 2019) and those spread over widespread geography (Chen 2023). Teaming is a unique type of collaboration, and experiments pitting all-human teams and all-AI teams against HATs have shown that HATs are more likely to underperform due to an inability to establish basic teaming components such as shared situational awareness and intra-team trust (McNeese et al. 2021). This means that AI teammates must incorporate human-centered design principles.

AI agents that require too much human oversight or act in an unexpected manner can actually increase the workload and stress of human workers (Hauptman et al. 2023). HAT research on situational autonomy has shown drastic increases in team performance when the team can modify an AI agent's autonomy levels (Salikutluk et al. 2024). Empirical research within the cybersecurity and medical fields have shown that too much AI autonomy can result in catastrophic failure states if the AI makes an error and a human teammate is unable to respond to correct it (Hauptman, Schelble, Flathmann, et al. 2024). This is because as AI operates with more autonomy, humans lose situational awareness of the AI's actions and their consequences (Onnasch et al. 2014). This is an important finding in the HAT community, as it slightly revises the need for an AI teammate to always operate with partial or full autonomy. This paper's position is that, in order to function as an AI teammate, the AI agent must be minimally capable of operating with partial autonomy. As humans are sometimes restricted from making independent decisions, an AI teammate could also be restricted from operating at higher autonomy levels under certain conditions.

2.3 AI Teammates in the Workforce

Much of the existing HAT research focuses on the gaming community due to the ability to customize AI agents within gaming platforms, as well as the recruitment pool of players used to partnering with virtual teammates (Zhang et al. 2021).

As researchers have sought to study the role of AI agents in applied settings, multiple recent studies have explored the promise of AI teammates in cybersecurity and cyber defense (Hauptman et al. 2023; Hauptman, Mallick, et al. 2024; Maennel and Maennel 2024; Malatji 2024). The security and defense of increasingly complex systems is an ideal space to investigate the use of AI teammates, as they would be much more capable of analyzing and adapting to emerging threats than their human counterparts (Malatji 2024). For instance, interviews with cyber incident responders suggested that AI teammates could be used to identify and operationalize indicators of compromise (IOCs) within an intrusion prevention system during incident response in real-time (Hauptman et al. 2023). This is an important consideration for critical infrastructure, where the ability to predict, detect, and respond to cyber threats is a new, essential capability that most organizations do not currently possess.

Developing an effective AI teammate for distinct purposes, such as cybersecurity, requires AI designers to consider a variety of human-AI interaction factors. Teams utilize several mechanisms to coordinate and collaborate in pursuit of shared goals, mechanisms that need to be part of an AI teammate's design. For example, one study that incorporated an AI network infrastructure agent found that the way AI teammates communicate with human teammates should seamlessly integrate into its existing channels, as opposed to asynchronously providing humans with updates. Study participants noted that asynchronous updates would add mental load to already demanding jobs and be far more likely to be viewed as simply another tool (Hauptman, Schelble, Duan, et al. 2024).

Sometimes, these mechanisms cannot be replicated, as AI is inept at participating in behavioral communication, and failures of AI to replicate it have proven to produce worse team performance than if they tried not to replicate it at all (Demir, McNeese, and Cooke 2016). HAT research suggests that to overcome this, teams need to identify what essential information AI teammates need to communicate and develop the best ways for the AI to push that information in a manner that will be received positively by human teammates (Demir, McNeese, and Cooke 2016). Thus, an AI teammate that is effective in one environment may be ineffective if placed in a new type of team. For instance, an AI teammate designed to maintain a campus network with a small group of IT experts may be highly efficient and valued. However, if the same agent were deployed on a more secure network with focused cybersecurity professionals without modifications, it might be immediately disregarded as a liability. For this reason, it is essential that organizations determine which aspects of teaming are most crucial for an AI agent to replicate if it is to assume a full, interdependent team role.

3 PROCESS FOR DETERMINING IF/HOW TO EMPLOY AN AI TEAMMATE

Based on our research expertise in human–AI teaming, we developed a three-step process to help organizations assess whether, when, and how an AI teammate should be integrated into their workforce. The process presented in this paper is a conversation starter- it is intended to get critical infrastructure sectors thinking about the sources of their talent gaps, if AI is a probable solution, and, if so, identify the human-centered design considerations necessary for a successful human–AI collaboration.

The three steps are articulated as guiding questions that organizations should ask themselves when considering the integration of AI teammates: Is it a capacity or capability gap? What would be the role of the AI on the team? And what are some important human-centered design considerations? (Figure 1). Framing the process through these questions anchors the discussion in organizational needs, technological suitability, and human factors. The following subsections elaborate on each step in turn.

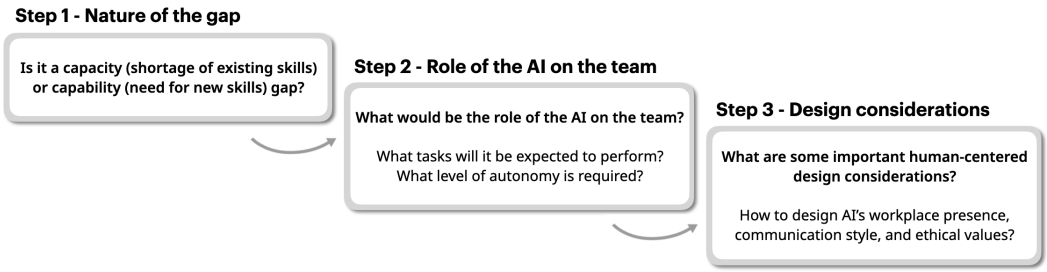

Figure 1. A 3-step process to determine if and how an AI teammate should be integrated.

3.1 Question 1: Is it a Capacity or Capability Gap?

The gaps that critical infrastructure sectors are struggling to fill can be classified by either an issue of *capacity* or an issue of *capability*. A capacity gap reflects skills that the team has organically possessed in the past but is currently struggling to maintain. This loss or shortage of skills is often due to retirement, turnover, or insufficient recruitment pipelines. In contrast, a capability gap concerns the emergence of new needs that existing human workers are ill-equipped to address. It arises when new operational demands, technologies, or threat vectors create requirements that the existing workforce is not equipped to meet. In the first case, an AI teammate is only beneficial in a small set of circumstances; whereas, in the case of a capability gap, an AI teammate is almost always value-added to the organization.

Understanding whether a workforce challenge reflects a loss of past capacity or the emergence of new capability demands is, therefore, a critical first step in determining whether, and how, an AI teammate should be integrated.

If you answered capacity gap:

What does a capacity gap mean for AI integration? Critical infrastructure sectors are encountering shallow hiring pools. Understanding the underlying causes of these shortages is essential for determining whether an organization should (a) invest in strategies to attract and retain human talent or (b) consider whether an AI agent could meaningfully supplement the team.

In some cases, the shallow pool is due to rising worker expectations around pay, benefits, and flexible work arrangements (Ashworth 2006). Over the years, sectors have utilized short-term fixes to fill critical roles. However, with sectors such as energy seeking to hire over 200,000 employees over the next three years (Hauser et al. 2025), it is increasingly clear that sustainable, long-term workforce solutions are necessary. For the most part, these positions remain inherently well-suited to human workers, and it behooves the sector to adapt to the changing demands of the workforce, as they are not going to disappear by filling one or two team roles with AI. Instead, replacing humans in work roles that suit their capabilities and interests is more likely to breed fear of replacement and increase the rate of turnover among workers who experience negative emotional responses to the AI (Verma and Singh 2022). In such circumstances, an AI teammate is not an appropriate solution.

On the other hand, in many cases, the shallow hiring pool stems from a growing mismatch between the jobs available and the skills pursued by younger workers (Becker 2010). This poses a significant challenge for critical infrastructure organizations that require very specific skill sets that, for a variety of reasons, are considered unattractive to today's workforce. These reasons include lengthy or demanding educational pathways, lower entry-level salaries, rigid work schedules, and perceived status in the social hierarchy (Becker 2010). Because these issues typically lie beyond an individual organization's control, AI may offer a viable means of addressing persistent capacity shortages. In fact, strategically deploying AI in positions that are less attractive to the workforce may redirect human workers toward developing new skills needed to address emerging capability gaps.

If you answered capability gap:

What does a capability gap mean for AI integration? The capability gap refers to the emergence of new, largely technologically driven, needs in critical infrastructure sectors. Big data and AI-driven technologies are emerging as vital components of a sector's growth and efficient operations (Johnson et al. 2021). The incorporation of these technologies is generating new skill requirements in teams where workers must understand how to operate and leverage the technology (Johnson et al. 2021). As of 2024, the United States alone faced more than 500,000 unfilled cybersecurity-related positions (Dubov and Dubova 2025). Some of these skill gaps can, and should, be filled by humans with the right education and experience. As AI can fill repetitive, automatable job roles, humans are free to fill those that require greater degrees of creativity and reason (Daugherty and Wilson 2018).

However, some emerging roles require levels of speed, precision, pattern recognition, or computational breadth that surpass human cognitive limits to truly leverage the technology for the team. For instance, cyber incident responders report the need for AI teammates that can collect, analyze, and make sense of network data more quickly and accurately than a human analyst (Hauptman et al. 2023). This could be particularly useful for critical infrastructure sectors involving Supervisory Control and Data Acquisition (SCADA) systems, as AI solutions can overcome the challenges of having to monitor, understand, and compare various protocols and languages simultaneously (Aldossary, Ali, and Alasaadi 2021). Thus, in the case of a capability gap, an AI teammate is a promising solution to an organization's talent gap.

3.2 Question 2: What would be the Role of the AI on the Team?

Once the gaps that an organization wants to address with an AI teammate are identified, the next step is to define the scope of the AI teammate's role clearly. This scoping process mirrors how one would scope a position for a human hire. It involves two complementary decisions. First, the organization must determine which tasks the AI agent will be expected to perform. Second, it needs to understand and define how much autonomy the AI agent should be given to accomplish those tasks efficiently and safely. Together, these decisions ensure that the AI teammate is positioned to contribute meaningfully without introducing unnecessary risks or ambiguity in its role.

Scope the AI role with the necessary tasks.

The most common use of AI is as a tool or assistive technology. In contrast, an AI agent functioning as a teammate is intended to occupy a defined, independent role rather than merely augment one already filled by a human (O'Neill et al., 2020). Human-AI teaming research has extensively investigated what roles would be best suited for an AI agent, based on its unique capabilities and realistic limitations. Studies show that roles involving the collection, integration, and analysis of vast amounts of data (Enskovitch et al. 2021), as well as the strict application of predefined actions (Hauptman et al. 2023) are well suited for AI teammates.

Conversely, roles requiring nuanced judgment, frequent adaptation to novel situations, or precise physical interaction with the environment remain better matched to human workers (Hauptman, Schelble, Duan, et al. 2024). For instance, studies have highlighted the potential of AI teammates as network security analysts, responsible for collecting and analyzing vast amounts of distributed data, or as healthcare assistants capable of storing and utilizing a patient's historical data and symptom patterns to develop a diagnosis (Hauptman, Schelble, Duan, et al. 2024).

Scope the AI role with the right level(s) of autonomy.

Once a team has determined that a specific role is suited for an AI teammate, the next step is to specify the level of autonomy that AI should have in executing its assigned tasks. This decision is analogous to determining a human worker's place in the management hierarchy, clarifying not only what the agent does but also how independently it is allowed to act. Human-AI teaming research has attempted to adapt and apply the levels of automation into the levels of autonomy to design and classify autonomous agents (O'Neill et al. 2022). As noted in Section 2, to be considered a teammate, an AI agent must minimally operate with the ability to fully execute its team role.

Further considerations include that the AI possesses enough autonomy to perform its duties without unduly increasing the workload or supervisory burden on human teammates (Hauptman et al. 2023). At the same time, its autonomy should not exceed the point at which it begins making decisions for which it lacks the necessary programming, contextual understanding, or reasoning capabilities (Hauptman, Schelble, Flathmann, et al. 2024). This ladder concept is of particular importance to teams that operate in high-risk environments, as the *ironies of autonomy* show that the more autonomy that AI has, the less capable human teammates are of detecting and responding to AI failures (Ganesh 2020). This balance is essential for teams that frequently navigate ethically ambiguous situations, a topic further explored in the third question.

3.3 Question 3: What are Important Human-centered Design Considerations Necessary to Ensure Effective Teaming?

Competent AI agents are not guaranteed to be either effective or accepted by a team (Flathmann et al. 2023). Empirical studies have shown that humans are more likely to trust and positively perceive an AI teammate when they perceive more similarities between themselves and the AI (Georganta and Ulfert 2024). In fact, several human-centered design considerations have been explored for generating trust in human-AI teams in the workplace (Hauptman, Duan, and McNeese 2022). This section highlights three such considerations particularly relevant for the design of an AI teammate operating in critical infrastructure sectors: workplace presence, explainability and communication style, and ethical ideology.

Design for the AI teammate's workplace presence.

An essential part of customizing an AI teammate is designing the type of presence the AI will have in the workplace. In many instances, an AI teammate may be required to have no physical presence at all, operating purely as a software agent embedded within a system or network resource. Even so, the nature of its presence—and how that presence is communicated to the team—often requires additional design attention. A recent study has shown that AI teammates should possess a degree of presence that is on par with how human team members are used to engaging with one another. For example, teams that collaborate in a physical workplace every day would more readily trust and accept an AI teammate that exhibits some form of physical embodiment. Conversely, teams that are geographically dispersed and communicate over video conferencing would prefer an AI represented through an avatar or visual agent consistent with their communication norms (Hauptman et al. 2023). This is an important design consideration for critical infrastructure sectors, where trust, clarity, and shared situational awareness are central to safe operations. As such, AI teammates may require additional representational layers and embodiments beyond what is necessary for them to function and complete their tasks. For instance, while an AI teammate operating as a network analyst may not require a physical presence to perform its duties, providing a dedicated interface or "presence point" within the team's workspace can substantially improve how its alerts, findings, and recommendations are perceived and acted upon.

Future research should investigate how these presence-related design choices function in real-world deployments, as empirical work in operational settings remains limited. The field could also benefit from drawing more systematically on insights from adjacent domains such as human–robot interaction (Putlitz and Roesler 2024), Computer-Supported Collaborative Work (CSCW), and digital twin research, all of which offer insightful perspectives on embodiment, co-presence, and mediated interaction. Additionally, design scholarship on data physicalization (Bae et al. 2022; Jansen et al. 2015) provides a rich repertoire of approaches for making computational agents and their outputs more legible, tangible, and meaningfully integrated into human work practices.

Design the AI teammate to communicate with the team.

An important component of an AI teammate's presence is how it communicates. Explainable AI (XAI) is an extremely active research area, and is commonly defined as "an interface between humans and a decision maker that is, at the same time, both an accurate proxy of the decision maker and comprehensible to humans" (Chatila et al. 2021). In essence, effective explanations from an AI to a human must be relevant to the decision at hand, accurate in representing the AI's reasoning, and communicated in a way that is understandable and usable by the human receiving it. While research generally supports AI explanations as increasing trust in AI systems (Schoenherr et al. 2023), human-AI teaming research has shown that overly detailed or frequent explanations can have the opposite effect (Hauptman, Mallick, et al. 2024). Evidence also suggests that the method by which an AI agent's explanations are communicated is more important to human teammates. It is essential that an AI teammate's explanations reflect how the team generally communicates. This extends to preferred communication platforms (email, messaging platforms, reports), modalities (written, verbal, imagery) and frequency (Hauptman, Mallick, et al. 2024). Aligning with these norms helps ensure that explanations feel natural, appropriately timed, and easy to integrate into existing workflows.

Design the AI teammate to reflect the team's ethical values.

AI teammates need to communicate more than just their actions. An integral component of AI teammate acceptance is how well the AI seems to understand and exhibit adherence to the team's shared ethical code. This shared code is an essential part of effective teaming (Ouakouak and Ouedraogo 2019). Shared codes of ethics are vital to professional organizations in critical infrastructure sectors that make decisions affecting the safety and welfare of workers, clients, and society on a daily basis (Nichols, Nichols, and Nichols 2007). HAT research has thus proposed that AI teammates that are designed to incorporate and consider the shared ethical ideology of their team are more likely to be trusted and accepted as teammates (Flathmann et al. 2021).

That humans can trust an AI agent to consider and act in accordance with the team's ethical code is integral to long-term teaming in the workplace, as studies show that classic trust repair strategies are ineffective when an AI agent commits an ethical violation (Textor et al. 2022). This will be harder to design for in some sectors than others, depending on the extent to which the sector has a set of codified ethical principles on which everyone trains, such as the medical sector's Hippocratic Oath (Nicolaides 2014). In addition to guiding behavior, ethical codes can also be used to design AI teammates that adapt autonomy levels in response to ethical dilemmas, thereby preventing violations of the team's ethical code (Hauptman, Schelble, and McNeese 2022). This represents a promising human-centered design feature, as it acknowledges the inherently variable nature of ethics in complex, dynamic environments, and promotes responsible AI behavior before violations occur.

4 THE WAY FORWARD

As several critical infrastructure sectors begin to address growing talent gaps in their workforces, it is essential to recognize not only the challenges but also the opportunities presented by the increasing capabilities and adoption of AI technologies. Many of the demands faced by these sectors are due to the creation of new team roles required to collect, understand, and utilize large data sets, tasks that lend themselves well to AI agents. Consequently, human-AI teaming constructs are a promising solution to overcoming current and forthcoming talent gaps.

To do this effectively, this paper introduced a question-based framework for identifying and scoping the roles of AI agents designed to work as teammates in critical infrastructure organizations. This framework should be understood as a starting point rather than a mature model, and the public and private sectors need to take multiple next steps to test and refine it. Future work is necessary to further develop and validate the framework. While the present article provides conceptual guidance, it does not include case studies that demonstrate how the framework operates in real operational environments. To fully operationalize the "how," future research should translate these questions into concrete organizational workflows, develop maturity models or decision trees, and create actionable checklists tailored to sector needs. These steps will be crucial for transforming the framework into a practical tool for critical infrastructure cyber defense.

Critical infrastructure organizations should apply the guiding framework proposed in this paper to clarify their talent gaps, understand the sources of those gaps, and determine which roles in their organizations are most suited for AI teammates. As noted in Section 2, most

HAT research to date has been conducted in low-risk or gaming settings. Governments and industry partners will need to invest in pilot programs that allow organizations to experiment with these roles in operational environments, helping them identify the human-centered design features most critical for effective integration into sector-specific teams.

Ultimately, while AI holds enormous potential for these sectors, it is also true that their integration comes with nuances and associated risks. Research on human–AI teaming is still recent, and many questions about real-world implementation, organizational fit, and long-term impacts remain open. Poorly designed systems, incorrect autonomy levels, and vulnerability to cyber attacks could all derail the employment of AI and exacerbate, rather than alleviate, the talent gap. Accordingly, iterative evaluations and field testing are essential and researchers must continue to refine and adapt this framework to include sector-specific considerations and lessons learned. Doing so will enable the development of AI teammates that critical infrastructure organizations can adopt, accept, and utilize in the real world.

ABOUT THE AUTHOR

MAJ Allyson Hauptman is a Cyber Warfare Officer in the United States Army and a senior researcher at the Army Cyber Institute at West Point. She holds a master's degree in Cybersecurity from Tallinn Technical University and a doctorate in Human-Centered Computing from Clemson University, where her dissertation focused on adaptive autonomous teammates. Her research areas include human-AI teaming, adaptive autonomy, and cyber defense policy.

ACKNOWLEDGMENTS

Thank you to the participants at the 2025 Jack Voltaic Workshop on Cyber Resilience and National Power Projection for your insights and contributions to the framework presented in this article.

REFERENCES

Aldossary, Lina Abdulaziz, Mazen Ali, and Abdulla Alasaadi. 2021. "Securing SCADA systems against cyberattacks using artificial intelligence." In *2021 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT),* 739–745. IEEE. https://doi.org/10.1109/3ICT53449.2021.9581394.

Asare, Isak Nti, Scott Shackelford, Jungwoo Chun, and Sarah Powazek. 2025. "Protecting Communities while Training Future Cybersecurity Professionals: Lessons from the Consortium of Cybersecurity Clinics." *The Cyber Defense Review* 10 (2). https://doi.org/10.556582/cdr/1dm8-wekj.

Ashley, Travis D., Roger Kwon, Sri Nikhil Gupta Gourisetti, Charalampos Katsis, Christopher A. Bonebrake, and Paul A. Boyd. 2022. "Gamification of cybersecurity for workforce development in critical infrastructure." *IEEE Access* 10:112487–112501. https://doi.org/10.1109/ACCESS.2022.3216711.

Ashworth, Michael J. 2006. "Preserving knowledge legacies: workforce aging, turnover and human resource issues in the US electric power industry." *The International Journal of Human Resource Management* 17 (9): 1659–1688. https://doi.org/10.1080/09585190600878600.

Ayodele, Olabode Adekunle, Alice Chang-Richards, and Vicente Gonzalez. 2020. “Factors affecting workforce turnover in the construction sector: A systematic review.” *Journal of Construction Engineering and Management* 146 (2): 04019112. https://doi.org/10.1061/(ASCE)CO.1943-7862.0001725.

Bae, S. Sandra, Clement Zheng, Mary Etta West, Ellen Yi-Luen Do, Samuel Huron, and Danielle Albers Szafir. 2022. “Making Data Tangible: A Cross-disciplinary Design Space for Data Physicalization.” In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems.* New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3491102.3501939.

Becker, Frank Stefan. 2010. “Why don’t young people want to become engineers? Rational reasons for disappointing decisions.” *European Journal of Engineering Education* 35 (4): 349–366. https://doi.org/10.1080/03043797.2010.489941.

Chatila, Raja, Virginia Dignum, Michael Fisher, Fosca Giannotti, Katharina Morik, Stuart Russell, and Karen Yeung. 2021. “Trustworthy AI.” In *Reflections on Artificial Intelligence for Humanity,* edited by Bertrand Braunschweig and Malik Ghallab, 13–39. Cham: Springer International. https://doi.org/10.1007/978-3-030-69128-8_2.

Chen, Zhisheng. 2023. “Collaboration among recruiters and artificial intelligence: removing human prejudices in employment.” *Cognition, Technology and Work* 25 (1): 135–149. https://doi.org/10.1007/s10111-022-00716-0.

Daugherty, Paul R., and H. James Wilson. 2018. *Human + Machine: Reimagining Work in the Age of AI.* Boston, MA: Harvard Business Review Press.

Demir, Mustafa, 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. IEEE. https://doi.org/10.1109/COGSIMA.2016.7497782.

Dubov, Dmytro, and S. Dubova. 2025. “A Global Shortfall of Cyber Workforce: Evaluating the US Strategy and Approach and Ukraine’s Emerging Challenges.” *Actual Problems of International Relations* 1 (162): 55–62. https://doi.org/10.17721/apmv.2025.162.1.55-62.

Durkovich, Caitlin. 2020. “Protecting Critical Infrastructure.” In *Beyond 9/11: Homeland Security for the Twenty-First Century,* edited by Chappell Lawson, Alan Bersin, and Juliette N. Kayyem. MIT Press. https://doi.org/10.7551/mitpress/13831.003.0012.

Enskovitch, John, Corey Fallon, Kate Miller, and Aritra Dasgupta. 2021. “Beyond visual analytics: Human-machine teaming for AI-driven data sensemaking.” In *2021 IEEE Workshop on Trust and Expertise in Visual Analytics (TREX),* 40–44. IEEE. https://doi.org/10.1109/TREX53765.2021.00012.

Fiume, Leonard J., Ilya Grinberg, Mohammed Safiuddin, and Robert F. Zahm. 2006. “A partnership between the electrical power industry and academia to address the technical talent gap.” In *2006 IEEE PES Power Systems Conference and Exposition,* 655–659. IEEE. https://doi.org/10.1109/PSCE.2006.296396.

Flathmann, Christopher, Beau G. Schelble, Nathan J. McNeese, Bart Knijnenburg, Anand K. Gramopadhye, and Kapil Chalil Madathil. 2023. “The purposeful presentation of AI teammates: Impacts on human acceptance and perception.” *International Journal of Human-Computer Interaction* 40 (20): 6510–6527. https://doi.org/10.1080/10447318.2023.2254984.

Flathmann, Christopher, Beau G. Schelble, Rui Zhang, and Nathan J. McNeese. 2021. “Modeling and guiding the creation of ethical human-AI teams.” In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society,* 469–479. https://doi.org/10.1145/3461702.3462573.

Ganesh, Maya Indira. 2020. “The ironies of autonomy.” *Humanities and Social Sciences Communications* 7 (1): 1–10. https://doi.org/10.1057/s41599-020-00646-0.

Georganta, Eleni, and Anna-Sophie Ulfert. 2024. “Would you trust an AI team member? Team trust in human-AI teams.” *Journal of Occupational and Organizational Psychology* 97 (3): 1212–1241. https://doi.org/10.1111/joop.12504.

Goure, Daniel. 2018. “Winning future wars: Modernization and a 21st century defense industrial base.” *The Heritage Foundation, 2019 Essays,* no. 4, 61–92. https://www.heritage.org/military-strength-topical-essays/2019-essays/winning-future-wars-modernization-and-21st-century.

Hauptman, Allyson I., Wen Duan, and Nathan J. McNeese. 2022. “The Components of Trust for Collaborating With AI Colleagues.” In *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social Computing,* 72–75. CSCW’22 Companion. Virtual Event, Taiwan: Association for Computing Machinery. https://doi.org/10.1145/3500868.3559450.

Hauptman, Allyson I., Rohit Mallick, Christopher Flathmann, and Nathan J. McNeese. 2024. "Human factors considerations for the context-aware design of adaptive autonomous teammates." *Ergonomics* 68 (4): 571–587. https://doi.org/10.1080/00140139.2024.2380341.

Hauptman, Allyson I., Beau G. Schelble, Wen Duan, Christopher Flathmann, and Nathan J. McNeese. 2024. "Understanding the influence of AI autonomy on AI explainability levels in human-AI teams using a mixed methods approach." *Cognition, Technology and Work* 26:435–455. https://doi.org/10.1007/s10111-024-00765-7.

Hauptman, Allyson I., Beau G. Schelble, Christopher Flathmann, and Nathan J. McNeese. 2024. "The Role of Autonomy Levels and Contextual Risk in Designing Safer AI Teammates." In *2024 IEEE 4th International Conference on Human-Machine Systems (ICHMS).* IEEE. https://doi.org/10.1109/ICHMS59971.2024.10555844.

Hauptman, Allyson I., Beau G. Schelble, and Nathan J. McNeese. 2022. "Adaptive Autonomy as a Means for Implementing Shared Ethics in Human-AI Teams." In *Proceedings of the AAAI Spring Symposium on AI Engineering.*

Hauptman, Allyson I., Beau G. Schelble, Nathan J. McNeese, and Kapil Chalil Madathil. 2023. "Adapt and overcome: Perceptions of adaptive autonomous agents for human-AI teaming." *Computers in Human Behavior* 138. https://doi.org/10.1016/j.chb.2022.107451.

Hauser, Michael, Donald Paul, Jim Crompton, and Iraj Ershaghi. 2025. "Workforce Development for the Energy Industry." Paper presented at the SPE Western Regional Meeting, Garden Grove, California, USA, April 2025. SPE. https://doi.org/10.2118/224208-MS.

Jansen, Yvonne, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. 2015. "Opportunities and Challenges for Data Physicalization." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems,* 3227–3236. CHI '15. Seoul, Republic of Korea: Association for Computing Machinery. https://doi.org/10.1145/2702123.2702180.

Johnson, Marina, Rashmi Jain, Peggy Brennan-Tonetta, Ethne Swartz, Deborah Silver, Jessica Paolini, Stanislav Mamonov, and Chelsey Hill. 2021. "Impact of big data and artificial intelligence on industry: developing a workforce roadmap for a data driven economy." *Global Journal of Flexible Systems Management* 22 (3): 197–217. https://doi.org/10.1007/s40171-021-00272-y.

Lewis, Ted G. 2019. *Critical Infrastructure Protection in Homeland Security: Defending a Networked Nation.* John Wiley & Sons.

Li, Robin. 2020. *Artificial intelligence revolution: How AI will change our society, economy, and culture.* Skyhorse Publishing.

Maennel, Kaie, and Olaf M. Maennel. 2024. "Human-AI Collaboration and Cyber Security Training: Learning Analytics Opportunities and Challenges." In *2024 17th International Conference on Security of Information and Networks (SIN),* 01–08. IEEE.

Malatji, Masike. 2024. "Human-Artificial Intelligence Teaming Model in Cybersecurity." In *2024 International Conference on Intelligent Cybernetics Technology and Applications (ICICyTA),* 216–220. IEEE. https://doi.org/10.1109/ICICYTA64807.2024.10913351.

Malatji, Masike, Annlize L. Marnewick, and Sune von Solms. 2022. "Cybersecurity capabilities for critical infrastructure resilience." *Information and Computer Security* 30 (2): 255–279. https://doi.org/10.1108/ICS-06-2021-0091.

McNeese, Nathan J., Beau G. Schelble, Lorenzo Barberis Canonico, and Mustafa Demir. 2021. "Who/what is my teammate? Team composition considerations in human-AI teaming." *IEEE Transactions on Human-Machine Systems* 51 (4): 288–299. https://doi.org/10.1109/THMS.2021.3086018.

Nichols, Nick, George V. Nichols, and Patsy A. Nichols. 2007. "Professional Ethics: The Importance of Teaching Ethics to Future Professionals." *Professional Safety* 52 (7).

Nicolaides, Angelo. 2014. "The Critical Role of Ethics Training in Medical Education." *African Journal of Hospitality, Tourism and Leisure* 3 (1): 1–12.

Nyre-Yu, Megan, Robert S. Gutzwiller, and Barrett S. Caldwell. 2019. "Observing cyber security incident response: qualitative themes from field research." In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting,* 63:437–441. 1. SAGE Publications. https://doi.org/10.1177/1071181319631016.

O'Heir, Jeff. 2017. "Filling the talent gap." *Mechanical Engineering* 139 (1): 28–33. https://doi.org/10.1115/1.2017-JAN-1.

O'Neill, Thomas, Nathan J. McNeese, Amy Barron, and Beau Schelble. 2022. "Human-autonomy teaming: A review and analysis of the empirical literature." *Human Factors* 64 (5): 904–938. https://doi.org/10.1177/0018720820960865.

Olonilua, Oluponmile, and John Ogbeleakhu Aliu. 2025. "Assessing Workforce Training Strategies in Critical Infrastructure: Insights and Recommendations (Report No. IHS-2025-1000)." *Institute for Homeland Security,* https://hdl.handle.net/20.500.11875/5062.

Onnasch, Linda, Christopher D. Wickens, Huiyang Li, and Dietrich Manzey. 2014. "Human performance consequences of stages and levels of automation: An integrated meta-analysis." *Human Factors* 56 (3): 476–488. https://doi.org/10.1177/0018720813501549.

Ouakouak, Mohammed Laid, and Noufou Ouedraogo. 2019. "Fostering knowledge sharing and knowledge utilization: The impact of organizational commitment and trust." *Business Process Management Journal* 25 (4): 757–779. https://doi.org/10.1108/BPMJ-05-2017-0107.

Putlitz, Johanna zu, and Eileen Roesler. 2024. "Let's get physical: The influence of embodiment on industrial human-robot interaction." *Proc. Hum. Factors Ergon. Soc. Annu. Meet.* 68, no. 1 (September): 437–443.

Salikutluk, Vildan, Janik Schöpper, Franziska Herbert, Katrin Scheuermann, Eric Frodl, Dirk Balfanz, Frank Jäkel, and Dorothea Koert. 2024. "An Evaluation of Situational Autonomy for Human-AI Collaboration in a Shared Workspace Setting." In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems.* CHI '24. Honolulu, HI, USA: Association for Computing Machinery. https://doi.org/10.1145/3613904.3642564.

Sandborn, Peter A., and Varun J. Prabhakar. 2015. "The forecasting and impact of the loss of critical human skills necessary for supporting legacy systems." *IEEE Transactions on Engineering Management* 62 (3): 361–371. https://doi.org/10.1109/TEM.2015.2438820.

Schoenherr, Jordan Richard, Roba Abbas, Katina Michael, Pablo Rivas, and Theresa Dirndorfer Anderson. 2023. "Designing AI Using a Human-Centered Approach: Explainability and Accuracy Toward Trustworthiness." *IEEE Transactions on Technology and Society* 4 (1): 9–23. https://doi.org/10.1109/TTS.2023.3257627.

Textor, Claire, Rui Zhang, Jeremy Lopez, Beau G. Schelble, Nathan J. McNeese, Guo Freeman, Richard Pak, Chad Tossell, and Ewart J. de Visser. 2022. "Exploring the relationship between ethics and trust in human-artificial intelligence teaming: A mixed methods approach." *Journal of Cognitive Engineering and Decision Making* 16 (4): 252–281. https://doi.org/10.1177/15553434221113964.

Thomas, Richard J., and Tom Chothia. 2021. "Learning from vulnerabilities-categorizing, understanding and detecting weaknesses in industrial control systems." In *Computer Security: ESORICS 2020 International Workshops, CyberICPS, SECPRE, and ADIoT, Revised Selected Papers,* 100–116. Springer International Publishing. https://doi.org/10.1007/978-3-030-64330-0_7.

Verma, Surabhi, and Vibhav Singh. 2022. "Impact of artificial intelligence-enabled job characteristics and perceived substitution crisis on innovative work behavior of employees from high-tech firms." *Computers in Human Behavior* 131:107215. https://doi.org/10.1016/j.chb.2022.107215.

Wagner, Paul. 2021. *Critical Infrastructure Security.* SSRN 3762693. https://doi.org/10.2139/ssrn.3762693.

Zhang, Rui, Nathan J. McNeese, Guo Freeman, and Geoff Musick. 2021. ""An ideal human": expectations of AI teammates in human-AI teaming." *Proceedings of the ACM on Human-Computer Interaction* 4 (CSCW3): 1–25. https://doi.org/10.1145/3432945.

Received 13 March 2025; Revised 5 November 2025; Accepted 20 November 2025