A Utilitarian Approach to the Structure of Human-AI Teams

Thom Hawkins
PM Mission Command
jeffrey.t.hawkins10.civ@army.mil

Daniel N. Cassenti
DEVCOM Army Research Laboratory
daniel.n.cassenti.civ@army.mil

Abstract

With the current progress and limitations of artificial intelligence (AI), optimal outcomes in any complex problem-solving scenario will involve humans and AI working together. Yet, there is no accepted standard or guidance to configure the structure of human-AI teams for optimizing performance. We propose to use principles of economics and Fitts' list[1] to provide this guidance. The study of the allocation of scarce resources provides the ideal model for mapping out the allocation of team members to accomplish an objective. Tasks can be allocated on an absolute basis, or relative to other tasks and their allocation. Some economic measures to consider are error rate, output time, and cost per unit time, in addition to conceptual measures, including trust or confidence, AI processing power, and human cognitive capacity. Although there may be more specific considerations for any human-AI team structure, two dimensions emerge as critical for an understanding of how to achieve optimality following a cognitive task analysis. First, cognitive task analysis will help separate the distinct tasks from one another. Those constituent parts should lead to an assessment of which tasks are best performed by humans and AI after consulting Fitts' list[1]. Second, an overall understanding of the task's computational complexity will help determine the ideal level of automation (i.e., the overall amount of the task allocated to the human teammates versus the AI teammates). We conclude that economic principles can help guide optimal human-AI teaming configurations. If done successfully, team-structure assessments can aid AI developers and save money and resources from pursuing AI capabilities that can remain with human teammates.

1 Introduction

1.1 Best Outcomes: Humans and AI Together

As artificial intelligence (AI) develops, it has become clear that optimal outcomes in complex problem-solving scenarios require humans and AI to work together. However, there is no widely accepted standard or guidance for configuring the structure of human-AI teams to optimize outcomes. The tasks involved in human-AI teaming are often complex and dynamic, making it difficult to determine the optimal allocation of tasks between humans and machines. Additionally, the characteristics of humans and machines can vary widely, adding to the complexity of the task allocation process.

In his *Treatise of Human Nature*[2], Hume discusses the division of labor, noting the efficiency of specialization across a group of laborers rather than the homogenized labor of a single worker completing various tasks. Enter the machine age: where machines, computers, robots, and AI have been primarily specialized. While we may yet see more general AI and robotic features capable of novel task completion, we now see the division of labor between humans and artificial agents (computers, robots, AI).

1.2 How Much Should Humans and AI Contribute?

One of the first questions we should ask about humans and AI teaming together is how much each should contribute. Although later we will argue that the components of the particular task change this answer, we can derive approximations depending on different factors. Perconti[3] constructed such a general answer for human-AI collaboration for military scenarios focusing on maneuvering through an environment and derived such a response in Figure 1.

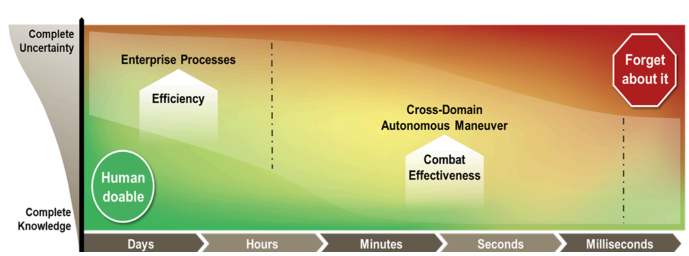

Distribution Statement A: Approved for public release. Distribution is unlimited.

Figure 1: How uncertainty and timescale affect optimal human-AI structure in the battlespace[3]

On the x-axis is timescale, represented categorically as days, hours, minutes, seconds, and milliseconds. The y-axis represents the amount of certainty (e.g., the highest value is complete uncertainty, lowest is complete knowledge) in the environment where the team operates. The intersection of timescale and certainty is coded to a specific color. The green hue represents the extent to which human teammates should contribute for optimal outcomes, while the red hue represents optimal AI contributions. This proportion is represented by *Levels of Automation*[4] (LOA), a critical concept for the optimal structuring of teams.

The area in the figure starts at large timescales with little need for AI intervention, except at higher levels of uncertainty. In the middle, most of the space requires both humans and AI. Finally, only with high levels of uncertainty and the shortest timescales, human cognitive capacity is exceeded, and only AI has hope of contributing. The Perconti[3] figure recommends that for the vast majority of the space, there should be a mix of human-AI (i.e., mid-LOA) with limited cases of one or the other alone for optimal outcomes.

1.3 Team Structure Guidance

While the media often focuses on how and where AI overtakes humans in more and more tasks (e.g., playing chess and generating malicious content), human-machine teaming remains a viable and versatile path forward. Evidence shows that humans and AI working together outperform AI or humans alone[5,6,7].

Parasuraman and Riley[8] have documented the suboptimal results of a purely functional allocation (e.g., loss of situational awareness and a lack of vigilance). Burkhard et al.[9] notes the advantages of variable task allocation. The interactive effects of teaming complicate the ability to predict in advance an optimal structure based on its components (individual humans and artificial agents); Cummings and Bruni[10] propose using sensitivity analysis tools to measure the balance of these effects on the outcome. Finally, Mingyue Ma et al.[11] and Groom and Nass[12] argue that humans will ultimately find artificial agents untrustworthy, to the team's detriment.

1.4 Fitts' List

One resource that takes particular importance for our goal of deriving guidance for the optimal structure of human-AI teams is Fitts' List[1]. In 1951, Paul Fitts led an inter-governmental panel to discuss how to understand and develop the burgeoning field of cybernetics (i.e., the study of closed feedback between biological and technical systems). The result of this project turned into Fitts' List, a double-sided list of skills that humans have, which far exceed machines, and skills the machines have, which exceed humans. Table 1 represents Fitts' List.

Table 1: Fitts' List[1]

Humans Better	Machines Better
Detecting inconsistencies	In extreme environments
Diverse sensing	Sensing particular stimuli
Perceiving patterns	Response speed
Attention to relevancy	Processing speed
Creative thinking	Precision in repetition
Strategic task allocation	Multitasking
Flexibility	Smooth force exertion
Learning from experience	Accurate performance
Low-chance events	Impervious to distraction
Induction	Deduction

The most remarkable feature of Fitts' List is that it has been robust to years of further research[13] and technological advances such that it can cover AI as well Gupta. We will cover this more in Section 3 below. Still, Fitts' List allows us to divide roles and responsibilities, not just a general division of labor, but specified tasks. When a human performs a task better, it should be assigned to humans, and when AI is better at a task, it should be assigned to AI.

1.5 Economics: Scarce Resource Allocation

Economics is the study of how individuals and societies choose to allocate scarce resources, why they allocate them that way, and the consequences of those decisions [15]. One of the central concepts in economics is scarcity. Scarcity refers to limited resources with wants and needs that exceed available supply. This means individuals and societies must choose how to allocate these limited resources. For example, a society may have limited food-production resources while people have increasing wants and needs for food. Thus, choices must be made about allocating these resources, such as whether to produce more food or use those resources to produce something else. Various factors drive economic decisions, including individual preferences, market forces, and government policies. Individuals make decisions based on their preferences and constraints, such as their income or the prices of goods and services. Market forces, such as supply and demand, play a role in determining prices and the allocation of resources in a market economy. Government policies like taxes and regulations can also affect economic outcomes.

2 Economics As a Guide for Team Structure

2.1 Why Economics?

No one would ever buy or sell anything if we all placed the same value on an asset. If one person had a coat and another needed a coat, it would be a zero-sum situation whereby, if someone sold the coat, it would be the same situation in reverse—one person needs a coat, and one person has a coat. If one person has two coats, they will place less value on one coat than someone with no coats. This is what drives economic activity.

Task allocation in a human-machine teaming situation is a market. While a machine may not value a task the way a shivering person might a coat, both tasks and agents have varying attributes that drive activity. This activity seeks a goal state, entropy or equilibrium, where all needs have been satisfied. The needs here satisfy the activity's outcome rather than individual desires.

The scarcity of resources is one echelon of need. Once everyone has a coat, there is another layer of value in the various attributes of the coat—color, style, material, lining, length, etc.—and the preferences of those who might wear them. Measuring the relative difference in tasks and the human or AI agent's ability to execute the tasks is crucial to the allocation process.

Similarly, in allocating tasks for human-machine teaming, functional assessments based on scarcity, such as the cognitive capacity for a human or processing power for a computer, are constrained and limit an asset. For example, you may have two clothes dryers, one much better than the other, but if you have a lot of clothes to dry and only a limited amount of time, you cannot dry them all in the better machine. Quality assessments like error rate, time to output, cost per time unit, and trust or confidence in output also create points of arbitrage that affect allocation decisions and optimization.

2.2 Measurement factors

2.2.1 Factors of Scarcity or Constraint

As with the example of the two coats, some factors represent limited resources, with scarcity being a key driver of economic exchange. For example, humans and artificial agents have separate, comparable limitations in human cognitive capacity and computer processing power. Cognitive capacity determines what a human teammate can do and is defined by the limits of divided attention. Divided attention does not burden an artificial agent much, depending on the extent of computing infrastructure. Still, calculation complexity may drive the need for processing power and cognitive capacity.

Cognitive capacity (human)

As the calculation complexity increases, the cognitive burden on humans tends to increase[16,17]. Complex calculations require more attention, focus, and mental effort, which can result in cognitive overload or fatigue. When performing complex calculations, humans may need to break down the problem into smaller parts, use heuristics or shortcuts, or rely on external aids such as calculators or spreadsheets. However, even with these aids, complex calculations can still be mentally demanding, leading to errors or slower performance.

Processing power (computer)

As the calculation complexity increases, the processing power required for computers also increases[18]. Complex calculations may require more memory, processing cycles, and sophisticated algorithms. For example, a simple arithmetic operation such as addition can be performed quickly by a computer. Still, complex mathematical operations, such as matrix multiplication or numerical integration, can require much more processing power. Depending on the available hardware and software resources, a computer may handle complex calculations efficiently, slow down or crash.

Available time

The time available for a task is usually time-bound by need. Computers have a speed advantage over humans in tasks computers can do, but if human interaction is required, then available time may remain a constraint for the team.

Environment

The decision-making environment may impact the ability of either a human or an artificial agent to operate; for example, temperature, atmosphere, radiation, weather, landscape, among other factors.

2.2.2 Factors of Quality

Another set of factors measures the output quality from a single participant or the entire team, such as error rate, time to output, cost per time or output, and trust or confidence in the output. These are the points of arbitrage where humans and artificial agents differ and therefore are most viable for analyzing trade-offs.

Error rate

The error rate is the number of errors a human or an automated agent makes divided by some output (time, units, etc.).

Time to output

Given that time is often constrained, the comparative time for a human and an automated agent to produce a given output is a likely factor for trade-off.

Cost per time or output

Another quality factor is the cost of a team member to perform for a period of time to produce a specified output. This could be pay for a human, or cloud service consumption fees for an artificial agent, for example. Parasuraman and Riley[8] noted that "designers tend to automate everything that leads to an economic benefit and leave the operator to manage the resulting system" (i.e., cost is often the primary basis for task allocation).

Trust or confidence in output

Humans tend to trust other humans and distrust automation[8]. Engineering trust in a teaming relationship can be complicated[5], and a lack of trust can be apparent in the output itself, or the user's acceptance of the output. A computer may complete a task more quickly, but at the expense of transparency, which could mean more time is spent by human teammates or by users trying to understand the actions taken by the automation than for humans to execute the tasks themselves.

2.2.3 Human-Centric Factors

Some factors are only a concern from the human perspective. Unlike robots (for now), psychological and sociological factors, such as level of engagement, the feeling of competition, and mortality, impact human productivity from an individual perspective and a societal concern for safety.

Engagement

Marx's[19] "theory of alienation" claims that inserting a layer of automation between a worker and their work has an alienating effect. Engagement or interest in one's work can improve productivity[20]. However, engagement is not only a productivity concern; it can also impact other outcomes, such as employee morale and retention.

Competition

Awareness of the progress or even the presence of other competitors in the marketplace can drive productivity[21]. Although an artificial agent could be programmed to adjust rates based on external factors such as competition, these factors are usually transparent to the agent. In addition, following on from Marx, separating work from its context, including a sense of competition, makes people less like humans and more like machines.

Although time to output is considered a quality output above, and time itself as a constraint, it should be noted that time pressure can hurt human performance[10].

Mortality

The mortality of humans could be considered an environmental constraint (see section 2.2.1); for example, humans can only survive in a specific temperature range, with oxygen, limited radiation, etc. While a computer may share some environmental constraints, a human has awareness of the risk of injury or death, which may impact not only their performance, but also the decision to place them in a particular environment.

2.3 Relative Versus Absolute Team Mapping

Ljesnjanin and Velagic[22] provide an overview of market-based approaches to task allocation, but crucially, these are auction-based—that is, microeconomic—mechanisms. Relative or comparative advantage is more often associated with macroeconomics and trade.

Absolute advantage refers to the ability of a person, group, or machine to produce a particular good or service more efficiently than another person, group, or machine. Relative advantage, on the other hand, refers to the ability of a person, group, or machine to produce a good or service at a lower opportunity cost than another person, group, or machine. In task allocation between humans and computers, decision-makers can use absolute and relative advantage to determine which tasks best suit each entity.

An example of absolute advantage for task allocation between humans and computers is using computer vision for image recognition tasks. Computers are typically better at processing large amounts of visual data quickly and accurately than humans. Therefore, tasks that involve image recognition, such as identifying objects in a photograph, can be allocated to computers with an absolute advantage.

Customer service tasks are an example of relative advantage for task allocation between humans and computers. While computers may perform certain customer service tasks, such as responding to frequently asked questions, humans often have a relative advantage due to their ability to understand complex human emotions and provide personalized solutions. Therefore, tasks that require a high level of empathy and interpersonal communication, such as handling customer complaints, can be allocated to humans with a relative advantage.

Figure 2 depicts a notional analysis of absolute versus relative advantage for two tasks concerning time to output. While the computer retains an absolute advantage over the human for both tasks, if the computer could not complete both tasks in parallel, then the human would be relatively efficient in completing task B while the computer completes task A.

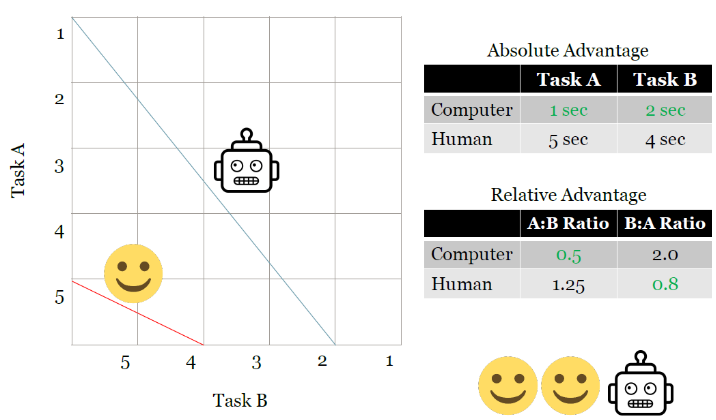

Figure 2. Absolute versus relative advantage

In most cases, the time-to-output differential between a computer and a human would overwhelmingly favor the computer, obviating the relative advantage. Other factors measuring advantage, such as error rate or confidence in output, may present a more competitive trade-off. Both Ranz et al.[23] and Burkhard et al.[9] argue for variable task allocation based on efficiency considerations, and this framework supports that perspective. Burkhard et al.[9] even mention comparative advantage in passing but do not take the idea further. Considering the measurement of multiple factors, the selection for a given task can be determined using multi-objective pareto optimization[24]. Use of Pareto optimization also subsumes, to an extent, the sensitivity analysis proposed by Cummings and Bruni[10].

3 Levels of Automation and Task Allocation

The two most prominent factors in structuring a team are the relative amount of resources allocated to the team members and what roles each will take. For our purposes here, the team members reduce to two types—humans and AI. An account of how that mix of skills best match tasks would need to be produced for scenarios with given configurations of tasks and skills among the resources available. However, this is beyond the scope of the guidance we would like to recommend here. We will start with a more detailed account of Fitts' List than in Section 1, then review cognitive task analysis to show how to break down a task into constituent sub-tasks. Following this, we will account for the use of LOA, then review how to create an ideal task breakdown to promote optimal outcomes.

3.1 Importance of Fitts' List

A panel of researchers developing a list in 1951 would seem unlikely to be relevant to human and artificial intelligence team planning in 2023. However, that is our proposal here. When organizing a team of human beings, a leader would first consider the skills of each team member for assignment to specific roles. This is no different when dividing roles among human and AI agents. We need to know how human and AI skills differ, which is what Fitts' List provides.

One could acknowledge that Fitts' List may have applied to machines in 1951, but ask whether it applies decades later, even with enormous advances in the development of AI. Cummings[14] argued that although dissent over the robustness of Fitts' List is prevalent, Fitts' List has proven consistent throughout developments in technology (at least by 2014 with the publication of her article). A survey of engineering students showed that the students largely agreed with the claims of Fitts' List for the automation of 2015.

de Winter and Dodou[25] explain that the design of Fitts' List includes six properties that make it ideal for the flexible inclusion of new technology developments. These properties are Plausibility, Explanatory adequacy, Interpretability, Simplicity, Descriptive adequacy, and Generalizability. Though we agree with this in principle, one potential path could prevent Fitts' List from remaining accurate into the future. Cassenti, Veksler, and Ritter[26] argue that the most common arguments against AI achieving human levels of cognitive acumen may not continue to prevent progress. Though this argument is beyond the scope of this paper, it is worth noting that AI may infringe on the human side of Fitts' List, thus breaking the robustness of the List. However, at this time, AI is not there.

Looking over Table 1, some general principles emerge. Humans have clear advantages over AI in dealing with anomalies, making inferences, flexibility, creativity, and deriving solutions to complex problems. AI has faster information processing, the ability to handle greater amounts of data simultaneously, accurate and precise outcomes, and immunity to distraction and stress. These divergent sets of skills cover a general landscape of settings, problem scenarios, and task sets. For example, Figure 1 shows that we should generally lean on human skills with longer timescales and greater certainty. With smaller timescales and greater uncertainty, AI needs to take a greater share of the task set to work fast and process vast amounts of data to reduce uncertainty, respectively.

So, general principles can be derived for the differential use of AI versus human teammates. In the following sub-section, we will discuss cognitive task analysis, a tool leaders can use to break down problem scenarios into component tasks and assign team members with given skills to those tasks.

3.2 Cognitive Task Analysis

Cognitive task analysis (CTA) is the breakdown of a task into the various cognitive skills required to perform it[27]. CTA has many[26] functions, but we will focus here on its ability to help assign team members to roles. To do this, team members must undergo a skills assessment, then be paired by matching the skills they perform best in the group with the task requirements derived from the CTA. In their paper, Brenner et al.[28] focused on human teams, but we posit that the same can be done with human and AI mixed teams.

As Fitts' List indicates, AI senses only particular stimuli, and humans are better at working with low-probability events. As such, a skills analysis does not need to be conducted to differentiate the AI agents from one another. Instead, AI developers design skills to achieve specific purposes so those purposes would be ascribed with the AI assigned accordingly.

For humans, there are numerous methods of assessing skills, including aptitude testing, practical exercises, educational reports, and self-reports. Forming ideal teams by differentiating humans from one another and assigning by sub-task can draw from any of these sources.

We merely need to use Fitts' List to differentiate humans from AI. Suppose the CTA indicates that computational skills are required to analyze vast information, perform multiple tasks simultaneously, complete repetitive tasks accurately or precisely, and avoid distraction or stress. In that case, AI should be assigned to those sub-tasks. For humans, any sub-tasks requiring generalization, dealing with anomalies, higher-order cognition (e.g., decision-making, problem-solving), creativity, or flexible and creative thinking need to be assigned to humans. From this, it is clear that humans and AI cover the broad gamut of skills required for almost any problem scenario in a complex environment.

For example, Cassenti and Kaplan[29] describe the division of labor for decision-making under uncertainty. They argue that AI should focus on the information-processing requirements that lead to situational awareness and summarize this data for human consumption. The human must act as the decision maker to use the reduced data to derive relevancy to the decision and select the outcome. In their review of the human-agent teaming literature, Chen and Barnes[30], agree that in human-agent teaming, agents should keep human teammates updated on situational awareness and not diminish human decision authority. Although this is one example, we posit that using cognitive task analysis and Fitts' list (adjusted with advances in AI) ought to provide other general principles for the practical construction of human-agent teams for optimal outcomes.

3.3 LOA: Determining a Task's Balancing Point

As mentioned in the introduction, the LOA of a human-AI-teaming task is the ideal proportion of automated to human activity during task completion. Generally, it is best to avoid situations where humans do too much work (Level 1) because humans may view AI as inadequately useful in reducing cognitive burden[31]. Similarly, when automation is doing too much of the task (Level 5 or 7, depending on the LOA model), human teammates may develop the feeling that they have lost control of the task and distrust the automation[32] (i.e., engagement as discussed above). This latter situation can be particularly harmful when leaders feel undermined by AI[33]. Cassenti, Roy, Hawkins, and Thomson[34] varied LOA empirically and validated that configurations at the extreme ends of LOA resulted in worse performance and lower trust in automation for a cyber-security task.

Settling on the notion that middle LOA values are optimal merely rules out the extremes of LOA. Even in a five-level LOA model[30], three additional levels could still be used (e.g., Levels 2, 3, and 4). To determine which middle level is best, a researcher has two methods for deriving an answer. First, one could use CTA to calculate the balance of skills required for the task by analyzing the task and determining what percentage of skills are needed on both sides of Fitts' List, then assign human and AI team members accordingly. The other method requires time, planning, and resources if the answer does not already exist in the literature to run human-subjects experiments testing LOA. Cassenti et al.[34] recommend this approach when possible because empirical data is grounded in reality as opposed to LOA, which is theoretical.

3.4 Ideal Task Segmentation

The ideal assignment of human and AI teammates is not just deriving and creating the best LOA and assigning sub-tasks for best performance outcomes. Subjective factors are at least equally important as performance. If human teammates do not trust AI teammates, then they will not rely on them and aim to undermine the AI's contributions. Similarly, humans find working with AI difficult if subjective usability ratings are low. Human teammates may give up on the AI if that sentiment goes too high. Li and Lee[35] recommend ensuring that AI agents project a sense of working altruistically (i.e., contributing to shared goals). It is beyond the scope of this paper to expound on this too much, but those who wish to use this paper as guidance should attend to subjective factors. After all, a human teammate can diminish or eliminate AI participation in task completion no matter how well the AI performs its sub-tasks.

4 Conclusion

4.1 Applying Economics to Teaming

The collaboration between humans and AI has become a key area of research in recent years as AI technologies advance and become more prevalent in various domains. Several studies suggest that to achieve maximally beneficial outcomes, humans and AI must work together and complement each other's strengths and weaknesses.

Fitts' List remains a solid foundation, but as AI progresses and encroaches on territory once thought to be human domain, it will be necessary to look beyond absolute advantage to use comparative advantage to dynamically allocate tasks. An economic model such as comparative advantage provides a tool for analyzing the trade-offs between human and artificial agents for various qualitative factors, and then multi-objective Pareto optimality can be used to combine those factors to determine the task allocation to agent. Determining overall team composition may require another layer of multi-objective Pareto optimization across potential team members considering other factors such as availability and resource-loading.

Using Fitts' List and CTA allows role assignment based on expertise, thus offering a practical solution for specialization of labor, one of the fundamental topics of the first treatise on economics, Adam Smith's *Wealth of Nations* [36]. Although the absolute division between human and machine labor represented in Fitts' list will erode over time, the use of comparative advantage as an analytical tool, combined with CTA, will provide a framework for future dynamic and variable task allocation.

4.2 Benefits of Our Approach

An economic approach derives as much as possible from limited resources and thus uses up fewer resources. Instead of using research and development dollars (including AI developers' time) to advance AI that can do what humans do, why not save that money and plug human teammates into the equation instead? If humans struggle to keep up with information-processing requirements, why not have AI teammates do those parts and leave the humans to what Fitts' List indicates we do better?

Recently, AI development has made impressive leaps in capabilities (see[37,38,39]). Despite these advances, AI incursions into the human side of Fitts' List are still woefully behind respective human skills. We still need humans and AI to work together, and with this paper, we offer guidelines for structuring those teams. We hope this will benefit us as we need humans and AI to work together. We have not reached the singularity[40] yet, and it is likely still a long time from now when it will happen. Until then, this guide may put forward the most economically sound path for our future.

References

[1] Fitts, P. M. (1951). Human engineering for an effective air navigation and traffic control system. Washington, DC: National Research Council.

[2] Hume, D. (1896). *A treatise of human nature*. Clarendon Press.

[3] Perconti, P., (2018). Autonomous maneuver for multi domain operations." Presented at *International Conference on Science+Innovation for Land Power*.

[4] Sheridan, T.B. & Verplank, W.L.. (1978). Human and computer control of undersea teleoperators. Cambridge, MA: Massachusetts Institute of Technology: Man-Machine Systems Lab.

[5] Ezer, N., Bruni, S., Cai, Y., Hepenstal, S. J., Miller, C. A., & Schmorrow, D. D. (2019, November). Trust engineering for human-AI teams. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 63, No. 1, pp. 322-326). Sage CA: Los Angeles, CA: SAGE Publications.

[6] Walliser, J. C., de Visser, E. J., Wiese, E., & Shaw, T. H. (2019). Team structure and team building improve human–machine teaming with autonomous agents. *Journal of Cognitive Engineering and Decision Making*, *13*(4), 258-278.

[7] Schelble, B., Flathmann, C., Canonico, L. B., & Mcneese, N. (2021, January). Understanding human-AI cooperation through game-theory and reinforcement learning models. In *Proceedings of the Annual Hawaii International Conference on System Sciences*.

[8] Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human factors*, *39*(2), 230-253.

[9] Burkhard, M., Guggemos, J., & Seufert, S. (2022). Who is Best Suited for the Job? Task Allocation Process Between Teachers and Smart Machines Based on Comparative Strengths. In *Computer Supported Education: 13th International Conference, CSEDU 2021, Virtual Event, April 23–25, 2021, Revised Selected Papers* (pp. 3-23). Cham: Springer International Publishing.

[10] Cummings, M. L., & Bruni, S. (2010). Human-automated planner collaboration in complex resource allocation decision support systems. *Intelligent Decision Technologies*, *4*(2), 101-114.

[11] Mingyue Ma, L., Fong, T., Micire, M. J., Kim, Y. K., & Feigh, K. (2018). Human-robot teaming: Concepts and components for design. In *Field and Service Robotics: Results of the 11th International Conference* (pp. 649-663). Springer International Publishing.

[12] Groom, V., & Nass, C. (2007). Can robots be teammates?: Benchmarks in human–robot teams. *Interaction studies*, *8*(3), 483-500.

[13] de Winter, J. C. F., & Hancock, P. A. (2015). Reflections on the 1951 Fitts list [sic]: Do humans believe now that machines surpass them? Procedia Manufacturing, 3, 5334–5341.

[14] Cummings, M. M. (2014). Man versus machine or man + machine? IEEE Intelligent Systems, 29, 62–69.

[15] Mankiw N. G. (2021). Principles of economics (9th ed.). Cengage Learning.

[16] Wirzberger, M., Bijarsari, S. E., & Rey, G. D. (2017). Embedded interruptions and task complexity influence schema-related cognitive load progression in an abstract learning task. *Acta Psychologica*, *179*, 30-41.

[17] Haji, F. A., Rojas, D., Childs, R., de Ribaupierre, S., & Dubrowski, A. (2015). Measuring cognitive load: performance, mental effort and simulation task complexity. *Medical education*, *49*(8), 815-827.

[18] Fortnow, L., & Homer, S. (2003). A short history of computational complexity. *Bulletin of the EATCS*, *80*(01), 2003.

[19] Marx, K. (2013). Economic and philosophic manuscripts of 1844: Estranged labour. In *Social Theory Re-Wired* (pp. 158-164). Routledge.

[20] Bélanger, J. (2000). *The influence of employee involvement on productivity: a review of research*. Hull, Que: Applied Research Branch, Human Resources Development Canada.

[21] Zlatcu, I. T., & Clodnitchi, R. (2018). The Impact of Competition on Productivity. Theory and Evidence. *Revista de Management Comparat International*, *19*(4), 410.

[22] Ljesnjanin, M., & Velagic, J. (2009). A market based approach for complex task allocation for wireless network based multi-robot system. In *2009 XXII International Symposium on Information, Communication and Automation Technologies* (pp. 1-5). IEEE.

[23] Ranz, F., Hummel, V., & Sihn, W. (2017). Capability-based task allocation in human-robot collaboration. *Procedia Manufacturing*, *9*, 182-189.

[24] Deb, K., & Gupta, H. (2005). Searching for robust Pareto-optimal solutions in multi-objective optimization. In Evolutionary Multi-Criterion Optimization: *Third International Conference, EMO 2005, Guanajuato, Mexico, March 9-11, 2005. Proceedings 3* (pp. 150-164). Springer Berlin Heidelberg.

[25] De Winter, J.C., Dodou, D. (2014). Why the Fitts list has persisted throughout the history of function allocation. *Cognition, Technology & Work, 16, 1-11*

[26] Cassenti, D.N., Veksler, V.D., & Ritter, F.E. (2022). Editor's Review and Introduction: Cognition-Inspired Artificial Intelligence. *Topics in Cognitive Science, 14*, 652-664.

[27] Schraagen, J. M., Chipman, S. F., & Shalin, V. L. (Eds.). (2000). *Cognitive task analysis*. Psychology Press.

[28] Brenner, T., Sheehan, K., Arthur Jr, W., & Bennett Jr, W. (1998). *Behavioral and cognitive task analysis integration for assessing individual and team work activities*. Pensacola FL: Navy Advancement Center.

[29] Cassenti, D. N., & Kaplan, L. M. (2021). Robust uncertainty representation in human-AI collaboration. In Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications III (Vol. 11746, p. 1174613). International Society for Optics and Photonics.

[30] Chen, J. Y., & Barnes, M. J. (2014). Human–agent teaming for multirobot control: A review of human factors issues. *IEEE Transactions on Human-Machine Systems, 44*, 13-29.

[31] Dwork, C., & Minow, M. (2022). Distrust of Artificial Intelligence: Sources & Responses from Computer Science & Law. Daedalus, 151, 309-321.

[32] Endsley, M.R. & Kaber, D.B. (1999). Level of automation effects on performance, situation awareness and workload in a dynamic control task." In *Ergonomics, Volume 42*, pp. 462–492.

[33] Hoadley, D. S., & Lucas, N. J. (2018). Artificial intelligence and national security. Congressional Research Service.

[34] Cassenti, D.N, Roy, A., Hawkins, T., & Thomson, R. (2022). The Effect of Varying Levels of Automation during Initial Triage of Intrusion Detection. In T. Ahram, J. Kalra & W. Karwowski (Eds.) Artificial Intelligence and Social Computing. AHFE (2022) International Conference. New York: AHFE International.

[35] Li, M., & Lee, J. D. (2022, September). Modeling Goal Alignment in Human-AI Teaming: A Dynamic Game Theory Approach. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Vol. 66, No. 1, pp. 1538-1542). Sage CA: Los Angeles, CA: SAGE Publications.

[36] Smith, A. (1776). An Inquiry into the Nature and Causes of the Wealth of Nations. Retrieved from https://www.ibiblio.org/ml/libri/s/SmithA_WealthNations_p.pdf

[37] Roscoe, R. D. (2023). Building Inclusive and Equitable Artificial Intelligence for Education. XRDS: Crossroads, The ACM Magazine for Students, 29(3), 22-25.

[38] Grimaldi, G., & Ehrler, B. (2023). AI et al.: Machines Are About to Change Scientific Publishing Forever. ACS Energy Letters, 8, 878-880.

[39] Biswas, S. S. (2023). Potential use of Chat GPT in global warming. Annals of biomedical engineering, 1-2.

[40] Kurzweil R. (2006). *The singularity is near: When humans transcend biology*. Penguin Books.