Human Factors and Simulation, Vol. 83, 2023, 12–24
https://doi.org/10.54941/ahfe1003558

Beyond the Tool vs. Teammate Debate: Exploring the Sidekick Metaphor in Designing Human-AI Dyads

Sylvain Bruni[1], Mary Freiman[2], and Kenyon Riddle[3]

[1]Aptima, Johnson City, TN 37604, USA
[2]Booz Allen Hamilton, Norfolk, VA 23502, USA
[3]Aptima, Grand Junction, CO 81507, USA

ABSTRACT

From symbiosis to copilot, a wide range of metaphors have been employed to characterize cooperative and collaborative relationships between human and non-human agents (be they software, robots, algorithms, or automated agents of any kind) in support of designing such advanced technologies. Recently, the emergence and rapid commoditization of artificial intelligence (AI) and machine learning (ML) algorithms have driven a highly bimodal debate on what metaphor is best to account for AI's and ML's new capabilities, particularly when those closely mimic humans': Is AI a tool or a teammate for humans using the technology? This debate, however, occludes critical elements necessary to practitioners in the fields of human system design. To move past the "tool vs. teammate debate," we propose an orthogonal metaphor, that of a sidekick, inspired by popular and literary culture, which can both accomplish and facilitate work (i.e., they do, and they help do). The sidekick metaphor was applied to a variety of efforts where it yielded novel design considerations which would have otherwise been unattainable by previous approaches. In this contribution, we report on the debate, describe the sidekick metaphor, and exemplify its application to real-world use cases, in domains such as intelligence analysis, aircraft maintenance, and missile defense.

Keywords: Sidekick, Human-ai teaming, System design, Design metaphor, Design thinking

INTRODUCTION

Defined as "figures of speech in which a word or phrase is applied to an object or action to which it is not literally applicable" (Oxford Languages), metaphors have been widely studied and employed for research and development purposes, particularly regarding interactive computer and software systems (Saffer 2005, Blackwell 2006). Although sometimes decried as unhelpful at best and harmful at worst (Cooper 1995) or even deviant (Ortony 1993), metaphors ostensibly reduce training and learning time, increase familiarity and affect, improve memorization, or reduce complexity (Marcus 1998). Beyond supporting design activities, metaphors are deemed an asset in *communicating* about design, particularly in agile development (Corral & Fronza 2018): Metaphors contribute to effectiveness by facilitating stakeholders' engagement, buy-in, and alignment as well as product alignment with end-users' mental models.

A variety of metaphors have been explored in the domain of human-machine interaction and collaboration: symbiosis (Licklider 1960), copilot (Mehdi 2023), wingman (Cummings & Morales 2005; Winnefeld & Kendall 2011), or teammate (Seeber et al 2020). Our team investigates, designs, and prototypes technologies leveraging Artificial Intelligence and Machine Learning (AI/ML) with advanced interactive modalities in High Intensity Vital Environments (HIVEs) such as defense, healthcare, law enforcement, or finance. In so doing, we have found ourselves unsatisfied by these common metaphors. Similarly, the ongoing debate in academia and industry whether "AI/ML should be considered a tool or a teammate" to human users has seemed too limited. Through the application of design thinking (Brown 2008, Spool 2018), we identified common principles that led us to devise and apply a new metaphor, that of a sidekick.

In this contribution, we report on (1) the observed limitations of the tool vs. teammate approach in HIVEs; (2) the sidekick metaphor and its alignment with designing and building human-machine systems in HIVEs; and (3) applications of the sidekick metaphor to real-world use cases, spanning a variety of domains such as intelligence analysis, aircraft maintenance, and missile defense.

THE TOOL VS. TEAMMATE DEBATE

Early research into cooperative relationships between humans and technology explored a targeted state of symbiosis (Licklider 1960; Gill 1996), seeking to achieve better effectiveness, efficiency, and communication (Gerber et al. 2020, Myers et al. 2019) together compared to on one's own. A vast body of academic research further explored and sought to characterize so-called "man-machine systems" along a variety of dimensions. These included, to name a few, what each member of the dyad contributes, how the system agents (human or digital) work with one another, or how system artifacts (such as information) evolve. This ongoing work has led to a plethora of concepts and frameworks, such as, amongst many, MABA-MABA (Fitts 1951; Dekker & Woods 2002), stages and levels of automation (Sheridan & Verplank 1978; Wickens et al. 2010; Save, Feuerberg & Avia 2012), and interaction models (Parasuraman, Sheridan & Wickens 2000; Allen & Ferguson 2002).

In parallel, constructive criticism of these concepts and frameworks has yielded an impressive secondary market for evolving theories and models (Reason 1987; Hancock & Scallen 1996; Sheridan 2000; Miller & Parasuraman 2003; de Winter & Hancock 2015). These theories and models have crystalized more recently into a dichotomous debate, about whether machine members of these systems are "tools" or "teammates," seen as a key question for any researcher or practitioner in the evolving fields of human-machine teaming (HMT) and AI/ML (Seeber et al. 2020; Zhang et al. 2021; Shneiderman 2022).

While arguing for "tool" or "teammate" is a worthy academic pursuit in that it supports the identification of new research streams and lenses through which to analyze human-machine systems, our experience as practitioners tasked with building such systems has shown this debate to be quite sterile, as it fails to account for real-world characteristics. Rather, we propose to employ the metaphor of the sidekick, as an orthogonal approach to the tool-teammate spectrum, in service to the design, development, and deployment of human-machine systems. Although it may be argued that "sidekicks" are either part of the spectrum or yet another framework, we have found it useful, simple, and elegant enough to drive innovation in the human-centered system space.

Figure 1: Robin, one of Batman's acolytes and most famous sidekick in popular comics. (Public Domain – CC0).

THE SIDEKICK METAPHOR

Overview

Our sidekick approach is modeled on popular culture and literary tropes of characters that both accomplish and facilitate work (i.e., they do, and they help do) through the anticipation of the hero's needs, in a way that matches their preference, and without getting in the way. Well-known examples of fictitious sidekicks include Hermione Granger and Ron Weasley (who support Harry Potter without stealing his star role or inconveniencing him too much), Alfred Pennyworth and Robin/Dick Grayson (who assist Batman by providing anticipatory and proactive support at home and on mission), or Chewbacca (who, as Han Solo's acolyte in the Star Wars series, knows how to help him in a *-the-* way Solo wants to be helped).

Alignment With Human-Machine Teams

Throughout a series of prototyping efforts in the defense domain, our team has identified the need or desire for similar characteristics in human-machine systems, particularly in High Intensity Vital Environments (HIVEs): The machine needs to reduce the human's overload (i.e., do) and the overhead burdening the human (i.e., help do), while being cognizant of the current and future contexts (anticipation and adaptation), the human's knowledge, skills, and capabilities (such as habits and patterns of work), or the risks it introduces (e.g., local optima, distractions, complacency, errors of omission and commission). Beyond a systems engineering application of the sidekick metaphor, we have found that it serves well the interaction with subject matter experts and stakeholders and facilitates the testing and evaluation of human-machine systems.

Components of the Sidekick Metaphor

Our work on HMT in HIVEs focused first on intelligence analysis by teams of distributed officers whose primary responsibility is to issue intel reports for rapid consumption by command staff and dismounted soldiers alike. From that initial foray emerged three critical principles, which we named BUTLER principles from their common underlying technological component, a "Behavior and Updates Tracker for Learning Expectations and Relevance." The three BUTLER principles (Bruni 2018) were summarized with the following saying: "*Good butlers learn their employers' (users') preferences, anticipate their needs, and remain discreet*," which would not be out of character for Alfred Pennyworth.

As our work expanded to other domains (e.g., aircraft maintenance, mission planning, missile defense simulation), refinement of these three principles became necessary to further guide our design and engineering team. Beyond general concepts, our team need to move towards *actionable design principles* that could be used as *decision criteria* to select what to design or implement amongst various options. To that end, we systematically employed the Short-form Creative Brief approach (SCB; Spool 2018) to identify key design principles (Parts 4 of SCBs) in over a dozen research efforts. A retrospective analysis combined with a consensus-based voting by those team members across all efforts and all disciplines (product ownership, scrum mastering, cognitive sciences, industrial/organizational psychology, human factors, design, software architecture/engineering/quality assurance, and customer engagement). Six core design principles emerged as common components for our sidekicks:

1. **Automate to remove overhead** – When users' brains are not absolutely needed to perform an activity, task, or workflow successfully, sidekicks do it for them. That way, cognitive burden is reduced, and users can focus on what requires human cognition.
2. **Go with my flow** – Automated support that does not fit the workflow creates impediments rather than solve them. Sidekicks adapt to their users' workflow. They exploit opportunities and augment the work to

fit. Ultimately, users remain in control and can decide to turn automated augmentations on or off as they please.

3. **Design for trust** – Working in HIVEs is hard, both for humans and machines. Thus, digital sidekicks support transparency, are reliable and forthright with information communicated to users, and know their limits so they can stay within their realm of expertise.
4. **Design for feedback** – Users and sidekicks are in it together, for the long run. Sidekicks are built to collect usage data (automatically) and feedback data (on demand, from users) so they can improve both user experience and system performance.
5. **Simplicity over cleverness** – In critical environments with high volumes of complex data, users must benefit from a clear and understandable experience. Sidekicks avoid overwhelming users with data and features, even if that requires foregoing neat and novel capabilities.
6. **Tailored transparency over uniform power** – Data literacy varies across users. Sidekicks tailor data views to their users and their context, rather than provide all bells and whistles to everyone in all things.

A post-hoc mapping of how the emerging six design principles support, apply to, or are relevant to the original BUTLER principles (Table 1) revealed imperfect coverage. In other words, individual BUTLER principles alone are not enough to address the needs for designing and implementing the sidekick metaphor; and, similarly, not all individual design principles on their own can cover all BUTLER principles.

Table 1. Mapping the six preliminary design principles to the BUTLER principles.

#	Emerging Design Principles	Learn Preferences	Anticipate Needs	Remain Discreet
1	Automate to remove overhead		X	X
2	Go with my flow	X	X	X
3	Design for trust	X	X	
4	Design for feedback			X
5	Simplicity over cleverness	X	X	
6	Tailored transparency over uniform power	X	X	X

APPLICATIONS

The theoretical considerations for the sidekick metaphor and how they align with human-machine teams are the product of iterative and incremental usage of the metaphor in a range of research and development efforts aimed at creating advanced prototypes of human-AI dyads. We describe in detail in the following sections the practical approach taken to leverage the sidekick metaphor in concrete examples in three different domains considered as HIVEs.

A Sidekick for Intelligence Analysis

Our first example is our original sidekick, "Analytic for Federated Data Tool for Human Efficiency: Behavior and Updates Tracker for Learning Expectations and Relevance," or ALFRED THE BUTLER, a sidekick for Army intel analysts.

In this use case, the problem to be solved is that analysts, working on their own or as part of teams, synchronously or asynchronously, in a collocated or distributed environment, do not have enough time to sift through all the data and sources they have at their disposal. This problem yields incomplete analyses or reports, both in breadth and depth of coverage, thereby hampering downstream decision making, planning, or other warfighting activities.

Our SCB described ALFRED as a desktop-based digital sidekick that searches, identifies, prioritizes, and summarizes key intel bits of information out of hundreds of thousands of source documents, and adapts its work based on current and expected future needs of analysts. ALFRED would do its work automatically, in the background, without interrupting but with constant adaptation of its activities and outputs based on the analysts' progress in the workflow and their immediate vs. long-term needs ("multi-horizon").

The principles of Table 1 were applied to this use case using the Spiraled Agile Design Sprinting (SPADES; Bruni 2020) approach. An example of the resulting interactive interface for ALFRED is presented in Figure 2. How the principles apply to ALFRED are detailed in Table 2.

A Sidekick for Aircraft Maintenance

Our second example is the "Collaborative Heuristic-based Engineering Workbench with a Behavior-Aware Classifier for Cognitive Assistance," or CHEWBACCA, a sidekick for Navy aircraft maintainers.

In this use case, the problem to be solved is the competency gap in aircraft maintenance: Due to a shortage of maintainers, junior maintainers are expected to conduct difficult maintenance tasks at a high level of proficiency, like

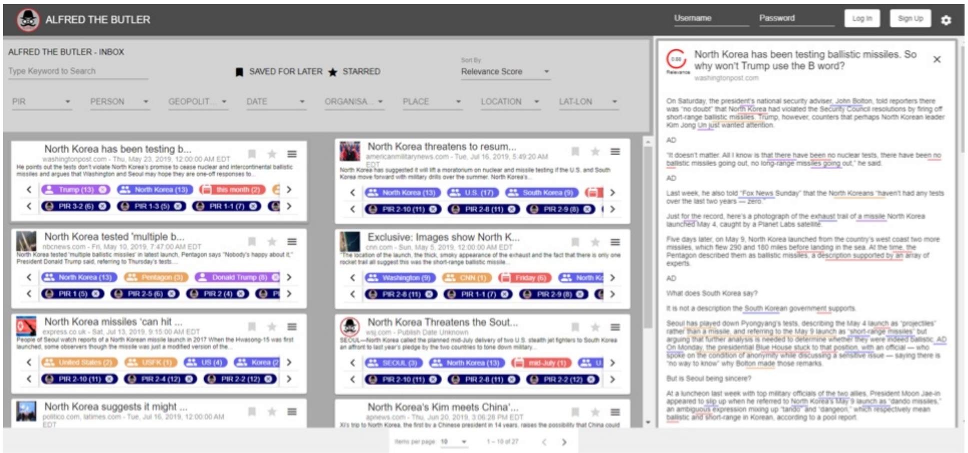

Figure 2: Screenshot of ALFRED's desktop interface for intelligence analysis.

Table 2. The six preliminary sidekick principles applied to ALFRED.

#	Principle	Description
1	Automate to remove overhead	Most overhead activities (searching, retrieving, pruning, identifying entities and relations, formatting into a report) are fully or quasi-fully automated.
2	Go with my flow	Recommendations for what document to consult are work-flow-dependent, over multiple horizons (now vs. later).
3	Design for trust	Direct feedback is immediately and visibly accounted for in the interface; transparency into ALFRED's inner workings is afforded in an admin panel.
4	Design for feedback	All user interactions, whether directly feedback related (e.g., thumb up on a document or source) or not (e.g., starring for later) are treated as feedback that helps ALFRED learn, to improve performance.
5	Simplicity over cleverness	ALFRED's interface only shows four data types (document, document card/metadata, entities, and requests) to ease usage and accelerate processes rather than displaying more data types in more complex visualizations.
6	Tailored transparency over uniform power	Entity recognition and disambiguation is directly accessible and matched to the contextual queries submitted by the user (as opposed to all encompassing).

experienced maintainers would. This problem yields maintenance delays, longer grounding of platforms, and thus, greatly reduced readiness of air assets and capabilities.

Our SCB described CHEWBACCA as a mobile digital sidekick embedded in augmented reality glasses that provides contextually relevant, step-by-step instructions and tips to junior maintainers so they can perform like experts.

The principles of Table 1 were applied to this use case using the Pareto Analysis for Technology Integration (PATI; Bruni 2022) and a 7-dimensional data framework (Weiss et al. 2021). An example of the resulting interactive interface for CHEWBACCA is presented in Figure 3. How the principles apply to CHEWBACCA are detailed in Table 3.

A Sidekick for Missile Defense

Our third example is the "Mixed-Automation Visualizer for Emerging Relationships & Insights in Complex Knowledge," or MAVERICK, a sidekick for missile defense analysts.

In this use case, the problem to be solved is simulation analysts in the missile defense space do not have enough time and computing capacity to process and analyze in a reasonable timeframe the petabytes of numerical

Figure 3: Screenshot of CHEWBACCA's AR interface for aircraft maintenance.

Table 3. The six preliminary sidekick principles applied to CHEWBACCA.

#	Principle	Description
1	Automate to remove overhead	Maintainers can access ("see") instructions right as they do the work, without having to go get the right physical maintenance manual or flipping through pages.
2	Go with my flow	Step-by-step instructions are calibrated and aligned with the work being performed and contextually detected.
3	Design for trust	CHEWBACCA embeds S1000D diagrams (the core "source-of-truth documentation") as well as tips or best practices provided by other maintainers.
4	Design for feedback	Maintainers can share what they are doing, so CHEWBACCA updates its state estimates and tracks progress against the workflow.
5	Simplicity over cleverness	In operational mode, CHEWBACCA only displays two kinds of information: step-by-step instructions with tips, and the list of tools needed. A third optional info element (on-demand) can be added: opportunistic job suggestions.
6	Tailored transparency over uniform power	Transparency is based on the current and next immediate tasks as per the instructions (rather than the full workflow, part ordering, or reporting).

data output by simulators on a weekly basis. This problem yields suboptimal, local analyses that focus on outlier behavior or a downgraded scope of inquiry.

Our SCB described MAVERICK as a desktop-based digital sidekick that efficiently processes highly dimensional, dense, high-volume, numerical data with no-loss and controlled-loss of data entropy, and visually surfaces insights for effective and efficient data analysis and exploration.

The principles of Table 1 were applied to this use case using the SPADES approach. An example of the resulting interactive interface for MAVERICK is presented in Figure 4. How the principles apply to MAVERICK are detailed in Table 4.

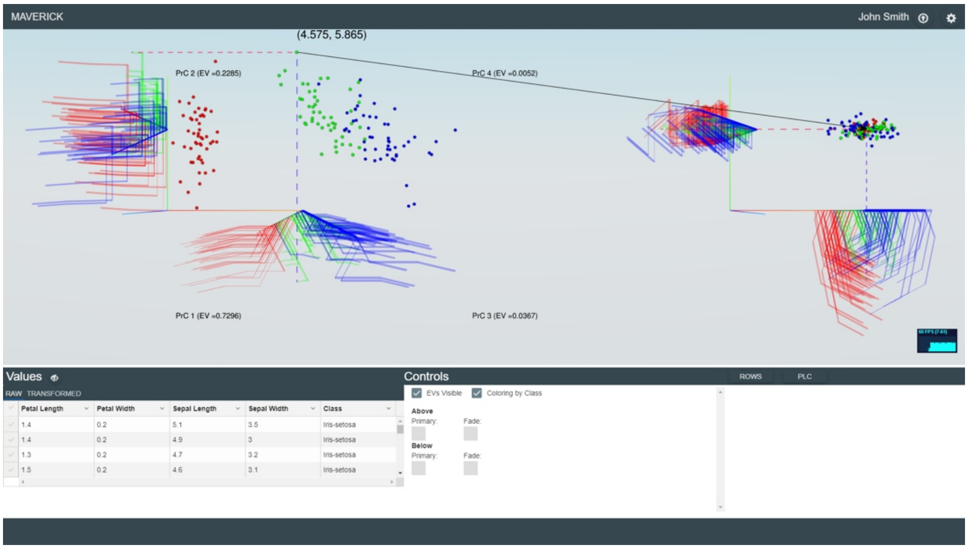

Figure 4: Screenshot of MAVERICK's desktop interface for data analysis.

CONCLUSION

In this contribution, we reported on the practical limitations of current metaphors in human-AI dyads, particularly the "tool vs. teammate" debate. We introduced the sidekick metaphor as an alternative and described its alignment with designing and building human-machine systems in HIVEs. We proposed six preliminary design principles to guide the implementation of technology as sidekicks in HMT. We finally detailed a sample set of applications of the sidekick metaphor to real-world use cases, spanning a variety of domains such as intelligence analysis, aircraft maintenance, and missile defense.

Our focus thus far has been limited, however, to HIVEs and dyads. We have not formally explored cases in non-intense or informal environments, nor those in settings where multiple humans collaborate with one or multiple sidekicks.

Further work will need to focus on those two aspects. Additionally, we plan on specifically investigating how the sidekick metaphor fares against other

Table 4. The six preliminary sidekick principles applied to MAVERICK.

#	Principle	Description
1	Automate to remove overhead	Data ingestion is a drag-and-drop interaction, all data pre-processing and visualization is automated, dimensionality reduction can be fully automated. Insights are surfaced automatically, with little to no interactions.
2	Go with my flow	MAVERICK accelerates specific elements of the workflow in cross-tabbed (linked) visualizations and tables.
3	Design for trust	MAVERICK supports lossless and controlled-loss dimensionality reduction, giving control over capabilities and how they are used that increase trust in the system.
4	Design for feedback	MAVERICK tracks how users operate the simple parameterization of its visualizations so it can learn what best reveals data insights or the analyst's internal mental model of the data.
5	Simplicity over cleverness	Three visualizations are available with simple controls to elicit insights from various approaches without overcomplicating or combining conceptually complex methodologies.
6	Tailored transparency over uniform power	Algorithm transparency is centered on its outputs (tied to insights) rather than its model/process (tied to data) to provide "task x analyst"-centric view.

systems engineering frameworks for human-automation collaboration, such as the Human Automation Collaboration Taxonomy (HACT; Bruni et al. 2007) and the Trust Engineering Framework (Ezer et al. 2019).

Finally, we encourage the application of the sidekick metaphor by practioners in the HMT field and welcome feedback.

ACKNOWLEDGMENT

The authors would like to acknowledge our colleagues at Aptima for their work on the related research and development efforts on which this contribution is based. This material is based upon work supported by the U.S. Army C5ISR Center under Contracts No W56KGU-17-C-0049 and No. W56KGU-19-C-0004., by the U.S. Navy Naval Air Warfare Center Aircraft Division under Contract No N68335-19-C-0765, and by the U.S. Missile Defense Agency under Contracts No HQ0147-19-C-7057 and No HQ086020C7136. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. Army C5ISR Center, the U.S. Navy Naval Air Warfare Center Aircraft Division, or the U.S. Missile Defense Agency.

Sidekick™ is a trademark of Aptima, Inc., used with permission.

REFERENCES

Allen, J. and Ferguson, G., 2002, October. Human-machine collaborative planning. In Proceedings of the Third International NASA Workshop on Planning and Scheduling for Space (pp. 27–29).

Blackwell, A. F., 2006. The reification of metaphor as a design tool. ACM Transactions on Computer-Human Interaction (TOCHI), 13(4), pp. 490–530.

Brown, T., 2008. Design thinking. Harvard business review, 86(6), p. 84.

Bruni, S., Marquez, J. J., Brzezinski, A., Nehme, C. and Boussemart, Y., 2007, June. Introducing a human-automation collaboration taxonomy (HACT) in command and control decision-support systems. In International Command and Control Research and Technology Symposium (pp. 1–13).

Bruni, S., 2018. Why Automated Agents Should Be More Like Alfred Pennyworth. In Proceedings of the 2018 Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC), Orlando, FL.

Bruni, S., 2020. Balancing Innovation and Tangibility Using the Spiraled Agile Design Sprinting Approach. In Advances in Creativity, Innovation, Entrepreneurship and Communication of Design: Proceedings of the AHFE 2020 Virtual Conferences on Creativity, Innovation and Entrepreneurship, and Human Factors in Communication of Design, July 16-20, 2020, USA (pp. 17-23). Springer International Publishing.

Bruni, S., Freiman, M., Weiss, C., Ward, D., Lynch, S. and Kay, K., 2020, August. A 7-dimensional framework for technical data in high-intensity vital environments and its application to aircraft maintenance. In 2020 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA) (pp. 211–215). IEEE.

Bruni, S., 2022. Introducing PATI: The Pareto Analysis for Technology Insertion-A Human-Centered Methodology to Identify and Prioritize Innovation in Complex Systems. Creativity, Innovation and Entrepreneurship, 31, p. 64.

Cooper, A., 1995. The Myth of Metaphor, Visual Basic Programmer's Journal, July 1995.

Corral, L. and Fronza, I., 2018, September. Design thinking and agile practices for software engineering: an opportunity for innovation. In Proceedings of the 19th Annual SIG Conference on Information Technology Education (pp. 26–31).

Cummings, M. L. and Morales, D., 2005. UAVs as tactical wingmen: Control methods and pilots' perceptions. Unmanned Systems, 23(1), pp. 25–27.

de Winter, J. C. and Hancock, P. A., 2015. Reflections on the 1951 Fitts list: Do humans believe now that machines surpass them?. Procedia Manufacturing, 3, pp. 5334–5341.

Dekker, S. W. and Woods, D. D., 2002. MABA-MABA or abracadabra? Progress on human–automation co-ordination. Cognition, Technology & Work, 4, pp. 240–244.

Ezer, N., Bruni, S., Cai, Y., Hepenstal, S. J., Miller, C. A. and 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.

Fitts, P. M. 1951. Human Engineering for an Effective Air-Navigation and Traffic-Control System. Oxford, England: Division of National Research Council.

Gerber, A., Derckx, P., Döppner, D. A. and Schoder, D., 2020. Conceptualization of the human-machine symbiosis–A literature review. Proceedings of the 53rd Hawaii International Conference on System Sciences.

Gill, K. S., 1996. The foundations of human-centred systems. Human machine symbiosis: The foundations of human-centred systems design, pp. 1–68.

Hancock, P. A. and Scallen, S. F., 1996. The future of function allocation. Ergonomics in design, 4(4), pp. 24–29.

Licklider, J. C., 1960. Man-computer symbiosis. IRE transactions on human factors in electronics, (1), pp. 4–11.

Marcus, A., 1998. Metaphor design in user interfaces. ACM SIGDOC Asterisk Journal of Computer Documentation, 22(2), pp. 43–57.

Mehdi, Yusuf. (February 7, 2023) Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web. Microsoft Website: https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web/.

Miller, C. A. and Parasuraman, R., 2003, October. Beyond levels of automation: An architecture for more flexible human-automation collaboration. In Proceedings of the human factors and ergonomics society annual meeting (Vol. 47, No. 1, pp. 182–186). Sage CA: Los Angeles, CA: SAGE Publications.

Myers, C., Ball, J., Cooke, N., Freiman, M., Caisse, M., Rodgers, S., Demir, M. and McNeese, N., 2018. Autonomous intelligent agents for team training. IEEE Intelligent Systems, 34(2), pp. 3–14.

Ortony, A. E., 1993. Metaphor and thought. Cambridge University Press.

Parasuraman, R., Sheridan, T. B. and Wickens, C. D., 2000. A model for types and levels of human interaction with automation. IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans, 30(3), pp. 286–297.

Reason, J., 1987. Cognitive aids in process environments: prostheses or tools?. International journal of man-machine studies, 27(5-6), pp. 463–470.

Saffer, D., 2005. The role of metaphor in interaction design. Information Architecture Summit, 6.

Save, L., Feuerberg, B. and Avia, E., 2012. Designing human-automation interaction: a new level of automation taxonomy. Proc. Human Factors of Systems and Technology, 2012.

Seeber, I., Bittner, E., Briggs, R. O., De Vreede, T., De Vreede, G. J., Elkins, A., Maier, R., Merz, A. B., Oeste-Reiß, S., Randrup, N. and Schwabe, G., 2020. Machines as teammates: A research agenda on AI in team collaboration. Information & management, 57(2), p. 103174.

Sheridan, T. B., 2000. Function allocation: algorithm, alchemy or apostasy?. International Journal of Human-Computer Studies, 52(2), pp. 203–216.

Sheridan, T. B. and Verplank, W. L., 1978. Human and computer control of undersea teleoperators. Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab.

Shneiderman, B., 2022. Human-centered AI. Oxford University Press.

Spool, J. M., 2018. The Magical Short-Form Creative Brief. Online UIE Article. Viewed 10 February 2022.

Weiss, C., Freiman, M., Bruni, S., Ward, D., Lynch, S. and Kay, K., 2021, June. Information Ergonomics: Expediting Maintenance Workflow Using Mixed Reality. In Advances in Simulation and Digital Human Modeling: Proceedings of the AHFE 2021 Virtual Conferences on Human Factors and Simulation, and Digital Human Modeling and Applied Optimization, July 25-29, 2021, USA (pp. 119–126). Cham: Springer International Publishing.

Wickens, C. D., Li, H., Santamaria, A., Sebok, A. and Sarter, N. B., 2010, September. Stages and levels of automation: An integrated meta-analysis. In Proceedings of the human factors and ergonomics society annual meeting (Vol. 54, No. 4, pp. 389–393). Sage CA: Los Angeles, CA: Sage Publications.

Winnefeld, J. A. and Kendall, F., 2011. Unmanned systems integrated roadmap FY 2011–2036. US Department of Defense.
Zhang, R., McNeese, N. J., Freeman, G. and Musick, G., 2021. "An ideal human" expectations of AI teammates in human-AI teaming. Proceedings of the ACM on Human-Computer Interaction, 4(CSCW3), pp. 1–25.