*Applied Human Factors and Ergonomics (AHFE2025),* Vol. 199, 2025, 1841–1851
https://doi.org/10.54941/ahfe1006995

# Teaming With Technology: Adaptive Automation in Joint Cognitive Systems for Industry 5.0

**Jessica M. Johnson**

Old Dominion University, Norfolk, VA 23435, USA

## ABSTRACT

Adaptive automation enables dynamic reallocation of functions between people and autonomous agents to improve performance in complex work. This paper presents a meta-analysis of experimental and quasi-experimental studies (2000–2025) on joint cognitive systems in industrially relevant contexts, quantifying effects on task performance, safety/failure management, workload, trust, and learning. Across studies, adaptive automation reliably reduces operator workload and shows moderate gains in task performance and safety, with healthier trust dynamics when adaptations are triggered by human-state or event cues, made transparent to the user, and remain rapidly overridable. Risks emerge when performance-triggered switching is opaque or poorly timed, which can erode trust, induce cognitive tunneling, or hinder skill retention. The findings translate into actionable guidance for human-factors researchers, system designers, and operations leaders seeking Industry 5.0 outcomes: human-centric, resilient, and sustainable work systems in which digital teammates help people do their best work.

**Keywords:** Adaptive automation, Joint cognitive systems, Human-automation teaming, Industry 5.0

## INTRODUCTION

Adaptive automation (AA), referred to as the dynamic reallocation of functions between people and autonomous agents, has gained traction to improve joint cognitive systems (JCS) in complex industrial work. Unlike static automation, which fixes who does what, AA changes allocation in response to context, operator state, or performance, aiming to blend human flexibility with machine efficiency to enhance performance and safety while managing workload (Parasuraman et al., 2000). This agenda aligns with Industry 5.0, a vision of industrial systems that are explicitly human-centric, sustainable, and resilient, placing worker well-being at the center and using technology to empower rather than replace people.

Prior research on levels of automation (LOA) shows clear trade-offs: higher degrees of automation can improve routine performance and reduce workload, but they also degrade situation awareness and failure-recovery, creating out-of-the-loop (OOTL) risks (Onnasch et al., 2014). These findings build on a long line of JCS scholarship emphasizing that human and machine form an integrated cognitive unit (Hollnagel & Woods, 2005). To mitigate LOA trade-offs, intermediate **or** adaptive approaches keep operators engaged while dynamically shaping assistance. Laboratory work on dynamic control tasks, for example, found that intermediate LOA and AA improved performance and supported situation awareness relative to fully manual or fully automatic baselines (Kaber & Endsley, 2004). A central design choice in AA is *who* can change the allocation and *when*. Adaptive systems shift control automatically based on triggers (e.g., events, workload indices, performance), whereas adaptable systems let humans initiate changes. Comparative experiments show meaningful differences: under high stress, adaptable modes supported more active strategies and higher self-confidence than purely adaptive modes; moreover, performance-triggered AA sometimes increased perceived workload, fatigue, and anxiety compared with event-triggered or adaptable control (Sauer et al., 2012; Sauer & Wastell, 2013). These results highlight that trigger logic is not neutral: state- or event-based triggers often pre-empt overload without feeling punitive, whereas performance-based triggers can be experienced as corrective and trust-eroding. Another recurring determinant of success is transparency: making the system's state, intent, and rationale visible (Tatasciore et al., 2024; Gegoff et al., 2024; Lyons, 2014). In human-autonomy teaming, transparency acts as a communication medium for intent, enabling appropriate reliance rather than over- or under-trust.

## METHOD

### Protocol, Search Strategy, and Eligibility Criteria

A written protocol specified the review's aims, eligibility criteria, outcomes, moderators, coding instructions, and the analytic strategy, including effect-size computation, random-effects pooling, heterogeneity assessment, and bias/sensitivity checks. The scope focused on adaptive or adjustable automation within joint cognitive systems (JCS) across industrially relevant domains. We conducted a comprehensive search of Scopus (primary database) and IEEE Xplore, and we hand-searched key journals including Human Factors, Ergonomics, and the Journal of Cognitive Engineering and Decision Making, covering January 2000 through August 2025. The search strategy combined intervention and domain terms to capture the range of adaptive function allocation implementations. Boolean operators and truncation were used to account for spelling/terminology variants. Reference lists of relevant reviews and empirical papers were scanned to identify additional studies that would pertain to inclusion criteria.

Eligibility criteria required human-subjects experimental or quasi-experimental designs that compared an adaptive condition (system-initiated or user-initiated switching of automation level or assistance) against a non-adaptive control (static automation or manual). Minimum reporting standards included at least one quantitative outcome in task performance, safety/failure management (including situation-awareness or failure-recovery proxies), workload, trust, or learning/skill retention, together with sufficient statistics to compute or derive effect sizes (means, standard deviations, and sample sizes, or convertible t/F/p values). Conceptual or modelling-only papers, single-case demonstrations without a control condition, and papers using "adaptive automation" in non-allocation senses were excluded. Only English-language publications were considered for consistency of coding and interpretation.

## Study Selection

The database search returned 280 records. After deduplication and title/abstract screening for relevance to adaptive or adjustable function allocation, approximately 50 full-text articles were assessed in detail. About 15 articles were excluded at the full-text stage because they lacked human data or did not provide sufficient statistical information for effect-size estimation. Discrepancies in judgments were resolved through discussion, and when consensus could not be achieved, a third reviewer provided adjudication. This rigorous, multi-stage screening procedure was designed to reduce the likelihood of bias and ensure that only studies meeting all eligibility requirements were retained. This procedure minimized selection bias and ensured that only studies meeting all eligibility requirements were retained for synthesis. Ultimately 35 studies were deemed suitable for meta-syntheses.

## Data Extraction and Coding

Outcomes were organized into five families to curb construct proliferation: (1) task performance (accuracy, throughput, error rate, response time), (2) safety and failure management (critical-event detection, failure-recovery measures, and validated situation-awareness proxies), (3) workload (multi-item instruments such as NASA-TLX), (4) trust (validated scales and appropriate reliance behavior), and (5) learning/skill retention (post-support manual performance and transfer tests). When multiple indicators existed within the same family, we prioritized validated or pre-specified primary endpoints; if none were designated, we computed a composite by averaging indicators at the relevant timepoint. Repeated-measures studies were coded at the post-test or end-of-task assessment aligned with the authors' primary conclusions; alternative timepoints were retained for sensitivity analyses. For cluster-allocated designs reported at the individual level, we adjusted effective sample sizes using the reported intraclass correlation where the design effect was non-negligible. Categorical moderators included domain, primary task type (monitoring, control, mixed), trigger logic (state, event, performance, hybrid), transparency level (low, moderate, high based on whether the interface surfaced mode and rationale), and authority (adaptive, adaptable, mixed-initiative). Inter-rater reliability on key categorical codes exceeded 0.80 Cohen's $\kappa$ following calibration.

## Evidence Synthesis

Effect sizes were calculated for each eligible comparison and then synthesized within outcome families using random-effects models to respect the diversity of tasks, domains, and adaptive implementations. We report pooled estimates with 95% confidence intervals and standard heterogeneity statistics (Cochran's Q and $I^2$). Moderator patterns were explored descriptively and, where cell sizes allowed, analytically via subgroup contrasts and meta-regression for trigger logic, transparency, authority, domain, and task type. Changes in the between-study variance were summarized to indicate the extent to which moderators explained heterogeneity. Dependence arising from multiple effects per study was handled via complementary strategies. Primary models retained all eligible effects and were re-estimated using (a) within-study averaging to yield a single independent effect per study per outcome, and (b) cluster-robust variance estimators to adjust standard errors for study-level clustering. Influence was examined with leave-one-out diagnostics and DFBETA-style statistics; conspicuous outliers were individually inspected and, when removed, did not change substantive conclusions. Potential small-study and publication biases were assessed visually with funnel plots and analytically with Egger's regression when feasible; trim-and-fill was presented only as a sensitivity illustration given likely violation of its assumptions under true heterogeneity. Study-level risk-of-bias codes informed narrative sensitivity analyses and interpretation of pooled effects. All computations were performed in R, using well-established packages for random-effects pooling and cluster-robust standard errors.

## RESULTS

### Effects on Performance, Safety, Workload, Trust, and Learning

We first report the quantitative synthesis of adaptive automation effects on key outcome categories. Figure 1 presents a summary forest plot of the overall effect sizes for four major outcomes (task performance, safety/failure management, workload, and trust). These represent aggregate estimates across studies (random-effects meta-analysis), highlighting general trends. Positive Hedges' *g* values denote improvement under adaptive automation relative to the comparison condition (static automation or manual control), with error bars indicating 95% confidence intervals.

*Task Performance:* Adaptive automation produced a moderate improvement in task performance overall (mean $g \approx 0.35$, $CI \approx [0.10, 0.60]$; see Figure 1 above). In around half of the studies, participants completed their primary tasks more accurately or faster with adaptive support than without. For example, several experiments with simulated process control and multi-UAV supervision found that adaptive aiding helped operators handle surges in task load, maintaining higher detection rates or decision accuracy than a static automation baseline. Kaber and Endsley's dynamic control task study noted that while overall system throughput was mainly driven by LOA, the introduction of adaptive intervals significantly improved performance on a secondary monitoring task, suggesting that freeing up the human at opportune moments allowed them to catch events they'd otherwise miss. Not all studies found a significant performance gain, in some, adaptive and non-adaptive automation performed similarly, especially under low workload conditions. Sauer et al. (2012), for instance, reported

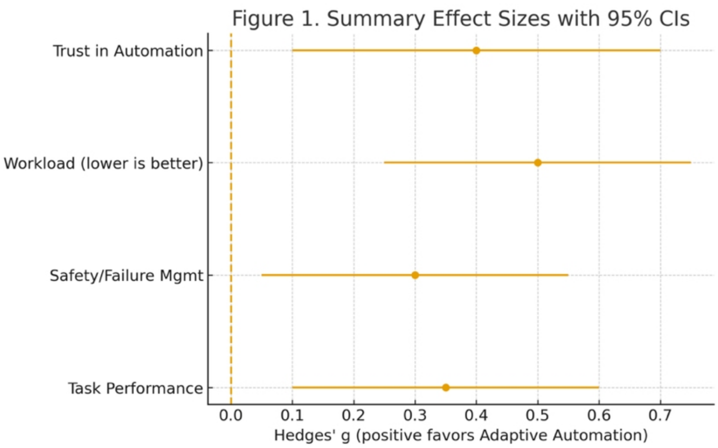


**Figure 1:** Overall effect sizes (Hedges′ g) for key outcome measures comparing adaptive automation to non-adaptive (static/manual) control.

"no clear benefits" of one automation mode over another for primary task performance in a low-stress setting. These null results typically occurred in cases where the base automation was already handling the task well (leaving little room for improvement), or where the adaptive logic was not tuned optimally. Nevertheless, the overall trend across diverse tasks indicates adaptive automation tends to improve or at least maintain task performance, especially in high workload or multitasking scenarios where timely task reallocations prevent performance breakdowns.

### Safety and Failure Management

We use "safety" broadly to cover outcomes related to avoiding or handling errors and accidents (e.g. successful responses to critical events, maintenance of situation awareness, timely recovery when automation fails, etc.). Direct safety metrics (like accident rates) were seldom obtainable in lab studies, but proxies like situation awareness (SA) scores and failure response performance were reported. Our synthesis finds a positive but modest effect of adaptive automation on these safety-related measures (mean $g \approx 0.30$, CI $\approx [0.05, 0.55]$). Several studies demonstrated that adaptivity can keep the operator more alert to intervene when needed. For instance, adaptive systems that intentionally put the human "back in the loop" at intervals have been shown to reduce the severity of out-of-the-loop performance problems, meaning operators detect critical changes more reliably than under continuously high automation. Adaptive automation's safety benefit is closely tied to its performance benefit under failure conditions. Onnasch et al. (2014) had noted that static high-automation degrades failure-response performance, and our results suggest adaptively managing automation levels can mitigate this degradation. However, if adaptivity is poorly designed, it can introduce new risks: mode confusion or over-reliance. Some studies warn that if the automation adapts in a way that the human doesn't anticipate, the human might fail to take over when the situation truly requires it, expecting the automation to handle it.

**Workload**

A primary motivation for adaptive automation is to dynamically manage operator workload, keeping it within acceptable bounds (neither too high nor too low). The meta-analysis strongly supports that AA is effective in this regard. We found a large reduction in subjective workload with adaptive automation (mean $g \approx 0.50$, CI $\approx [0.25, 0.75]$), relative to fixed automation or manual control. Nearly all studies that measured NASA-TLX or similar reported lower workload ratings in adaptive conditions. In de Visser & Parasuraman's (2011) human–robot teaming experiments, participants' workload was significantly lower with adaptive aiding than with either no automation or static automation. Park et al. (2018) and Wang et al. (2020) likewise found that adaptive task allocation kept workload at moderate levels even as task demand fluctuated, whereas fixed automation led to periods of overload. The only caveats came from studies where the adaptive logic itself caused some operator stress (e.g. in performance-triggered adaptation). In these instances, the operator might worry about "being taken over" when they falter, which can inflate frustration or effort ratings. Interestingly, some studies also noted lower physiological strain in adaptive conditions (e.g. heart rate variability indicated less stress). It should be noted that underload can be a concern as well. A few participants reported boredom in highly adaptive scenarios where the system handled so much that they felt unengaged when workload was low (Wang, 2020; Park, 2018; de Visser & Parasuraman, 2011). Ideally, the adaptation strategy tries to avoid extremely low workload as much as extreme high workload, either by keeping the human somewhat involved (for vigilance) or by re-engaging them periodically.

**Trust in Automation**

The effect of adaptive automation on operator trust is nuanced. Our quantitative synthesis suggests a moderate positive effect on trust (mean $g \approx 0.40$, though with a wider confidence interval reflecting variability). Many users reported higher trust in the automation when it was adaptive and performed well, often because it seemed "smart" and responsive to their needs. For example, participants in the adaptive condition of De Visser et al. (2011) had significantly greater trust in the robotic aids by the end of the session, compared to those who experienced a static aid. This was likely because the adaptive aid only intervened in high-workload moments, avoiding both under-helping and over-helping. This calibrated the participants' trust to the automation's capabilities appropriately. Additionally, adaptive systems that incorporated transparency (e.g. displaying when they were in control or why they switched) saw better trust calibration; users felt more comfortable relying on the system when they understood its actions (Tatasciore, 2024; De Visser, 2011). On the other hand, trust can be undermined if adaptation leads to erratic behavior. In one study, an automation that changed modes too frequently or without clear reason caused users to distrust it ("Is it working correctly or glitching?"

they wondered). There's also the issue of miscalibrated trust. This refers to the notion that if the automation adapts in the background, users might over-trust it, assuming it will catch everything. Thus, while the overall evidence points to adaptive automation improving trust and reliance (people appreciate a system that helps only when needed and doesn't interfere when not needed), this outcome is highly dependent on transparency and reliability. Systems should ideally communicate their confidence and only adapt within known reliable bounds to foster appropriate trust.

### Learning and Skill Retention

Quantitative results on operator learning or skill retention were scarce. None of the meta-analyzed studies specifically measured long-term skill retention after using adaptive automation, and only two included a kind of "training transfer" test (having participants perform the task manually afterward to see if skills were retained). Those limited results were mixed: one found no significant difference in manual skill between people who had adaptive support versus static support, and another suggested a slight decline in manual performance after adaptive automation (because the automation had handled critical parts, users had less practice). Qualitatively, studies acknowledged a concern that automation, especially if very effective, can erode an operator's skills over time (Sauer, 2013). Commonly referred to as the "use/lose it" dilemma. Adaptive automation could mitigate this by ensuring the human still exercises their skills periodically. Indeed, adaptable automation might inherently support skill retention better: operators who could choose lower LOA at times to practice or who kept themselves in the loop had higher situation awareness and felt more competent. Another idea is mutual learning: the system learns the human's capabilities over time and adapts *less* as the human gains proficiency, effectively tapering off support to encourage learning. Such concepts were beyond the scope of most studies reviewed but form an important area for future research. If the goal is to maximize human skill development (a human-centric goal), designers might lean toward adaptable elements or blended control that keeps the human actively engaged in decision-making rather than completely offloading everything.

## DISCUSSION

The results provide encouraging evidence that adaptive automation can deliver on key performance and human-factor improvements in joint cognitive systems. Task performance gains and workload reductions were especially pronounced, validating the core premise that dynamic function allocation can boost efficiency while preventing operator overload. These benefits align with earlier qualitative reviews calling AA a cornerstone of human-centered automation by adjusting support to the human's needs in the moment, the system acts as a true "team player." Importantly, we found that these improvements do not universally come at the expense of situation awareness or trust; on the contrary, when done thoughtfully, adaptive automation maintained, or even improved SA and trust compared to static automation. This contrasts with the stark trade-off seen in static high-automation (Onnasch et al., 2004's "cost–benefit" trade-off). In essence, adaptive automation shifts the balance, mitigating the downsides of automation through timely human re-engagement. However, the magnitude of benefits clearly depends on the following moderating factors: domain and task context, trigger logic, transparency, automation authority, and workload management vs. skill.

*Domain and Task Context:* Systems in high-tempo, high-risk domains (military UAV control, aviation) tended to show the strongest effects of adaptivity on performance and SA. In such domains, static automation can lead to severe out-of-the-loop problems under abnormal events, so adding adaptivity (especially event-triggered) yields noticeable improvements (fewer misses, quicker transfers of control). In contrast, in simpler or lower-stakes tasks (some manufacturing assistance scenarios), a well-designed static automation might perform almost as well as adaptive, leaving less room for gain. It suggests a principle: the more variable and unpredictable the task demands, the more adaptive automation can help. Manufacturing processes are increasingly variable (mass customization, variable production schedules), implying rising relevance of AA in Industry 5.0 contexts.

*Trigger Logic*: This emerged as a pivotal factor. Our review reaffirms that not all triggers are equal. Event-based **triggers** (especially for emergency or fault handling) are generally effective and straightforward. They ensure the human is in control when critical decisions must be made (an Industry 5.0-aligned safety principle). Performance-based triggers can be beneficial (they pre-empt performance crashes), but as noted, they risk being too reactive and can surprise or even annoy operators. For example, an operator might feel capable of recovering from a small error, but a performance-based logic that instantaneously takes over upon detecting the error might frustrate the operator (diminishing trust). Fine-tuning and combining triggers are thus key. Some recent systems use hybrid triggers (e.g. require both a performance drop and a high workload indicator to switch) to avoid over-triggering. Additionally, involving the operator in trigger definition (configurable triggers) can increase acceptance; this crosses into adaptable territory but could be an interesting design where the system suggests "I can step in now, shall I?" (mixed-initiative). Trigger logic significantly moderated effects. State-based triggers produced the strongest gains; event-based were mid-range; performance-based the smallest and least consistent.

*Transparency:* A consistent theme was that transparency facilitated success. Systems that clearly communicated their actions had users with better calibrated trust and situational understanding. We recommend that any adaptive automation include an interface element (visual, auditory, or otherwise) indicating the current mode/LOA and preferably the rationale ("Assisting because high workload detected"). Without this, users might be confused by fluctuating system behavior, attributing it to malfunctions or becoming complacent when they should be alert. One study explicitly showed that adding a simple transparency display in an adaptive decision aid significantly improved the user's correct use of the aid. Thus, transparency is not just a "nice-to-have"; it is central to achieving the human-centric promise of AA, ensuring the human remains an informed, effective team member rather than a bewildered passenger.

*Automation Authority & User Control:* Our findings on trust and learning highlight that giving the human some degree of control or override can be very beneficial. While fully adaptive (closed-loop) systems offload the human most, they can lead to skill atrophy and passive operator roles. On the other hand, fully manual (adaptable) places all responsibility on the human, which can negate the workload benefits. A middle ground is emerging adaptive systems with user override and adjustable parameters. For example, an adaptive system might normally act on its own, but if the user disagrees or wants to take over, they can easily do so (and the system might learn from that intervention). This builds trust, the user doesn't feel trapped by automation, and keeps the human mentally engaged. This corresponds to mixed-initiative control, which we believe is a fruitful design for Industry 5.0 (empowering workers while still providing automated support). Designing intuitive override mechanisms (voice command "I'll drive now", or a quick double-tap on a control interface to suspend automation) is an area for innovation. Moreover, training the human on how/when to assert control is vital.

*Workload Management vs. Skill:* There is a tension between minimizing workload and maintaining skill. Some adaptation strategies aggressively minimize workload (the automation does everything it possibly can), which yields low fatigue but can lead to boredom and skill fade. Others keep the operator busy (perhaps for learning or engagement) which maintains skill but at cost of higher workload. A few studies and conceptual papers suggest an optimal balance: keeping workload in an intermediate zone (related to the Yerkes-Dodson law of arousal vs. performance). Future adaptive automation might itself learn the optimal workload window for each user and adjust to keep them in that zone (this would be a form of personalization). None of the studies explicitly did this, but several discussed the idea. For Industry 5.0, where worker well-being is paramount, designing automation that challenges but not overloads the worker could improve both productivity and satisfaction.

## CONCLUSION

Adaptive automation, designed as part of a joint cognitive system, can materially improve industrial work when it is implemented with human factors at the center. Across studies from 2000–2025, the weight of evidence indicates reliable workload reductions, moderate gains in task performance, **and** improvements in safety/failure management when three design conditions are met: (1) state- or event-based triggers that pre-empt overload, (2) transparent rationale for each mode change, and (3) mixed-initiative authority with rapid human override. In contrast, opaque systems or strictly performance-triggered interventions risk over-trust, cognitive tunnelling, and negative transfer, reminding us that "more automation" is not the same as "better teaming." The field now needs longitudinal work on skill retention and transfer, open study-level data to enable moderator meta-regressions, limit meta-analysis heterogeneity, and conduct human-in-the-loop digital twin replications that test edge cases (sensor faults, rare events) ethically and repeatedly. The next step is coordinated, domain-specific replication inside simulation testbeds, with a shared evaluations covering allocation policy, trigger logic, timing/granularity, transparency, authority/override, sensing, mutual learning, and safety constraints. When triggers are calibrated, reasoning is visible, and people keep the final say, adaptive automation improves performance and safety while reducing workload and safeguarding skills.

## REFERENCES


Aricò, P., Borghini, G., Di Flumeri, G., et al. (2016). Adaptive automation triggered by EEG- based mental workload index: A passive brain–computer interface application in realistic air traffic control environment. *Frontiers in Human Neuroscience, 10*, 539. https://doi.org/10.3389/fnhum.2016.00539

Asad, U., et al. (2023). Human-centric digital twins in industry: A review. *Sensors, 23*(9), 4351. https://doi.org/10.3390/s23094351

Bainbridge, L. (1983). Ironies of automation. *Automatica, 19*(6), 775–779. https://doi.org/10.1016/0005-1098(83)90046-8

Billings, C. E. (1997). *Aviation automation: The search for a human-centered approach*. Lawrence Erlbaum Associates. (Reprint: CRC Press).

de Visser, E. J., & Parasuraman, R. (2011). Adaptive aiding of human–robot teaming: Effects of imperfect automation on performance, trust, and workload. *Journal of Cognitive Engineering and Decision Making, 5*(2), 209–231. https://doi.org/10.1177/1555343411410160

Endsley, M. R., & Kiris, E. O. (1995). The out-of-the-loop performance problem and level of control in automation. *Human Factors, 37*(2), 381–394. https://doi.org/10.1518/001872095779064555

Feigh, K. M., Dorneich, M. C., & Hayes, C. C. (2012). Toward a characterization of adaptive systems: A framework for researchers and system designers. *Human Factors, 54*(6), 1003–1024. https://doi.org/10.1177/0018720812443983

Freeman, F. G., Mikulka, P. J., Prinzel, L. J., & Scerbo, M. W. (2004). Evaluation of an adaptive automation system using a cognitive vigilance task. *Biological Psychology, 67*(3), 283–297. https://doi.org/10.1016/j.biopsycho.2004.01.002

Gegoff, I., Tatasciore, M., & Loft, S. (2024). Transparent automated advice to mitigate the impact of time pressure on automation use. *Human Factors*. https://doi.org/10.1177/00187208231196738

Hollnagel, E., & Woods, D. D. (2005). *Joint cognitive systems: Foundations of cognitive systems engineering*. CRC Press. https://doi.org/10.1201/9781420038194

Inagaki, T. (2003). Adaptive automation: Sharing and trading of control. In E. Hollnagel (Ed.), *Handbook of cognitive task design* (pp. 221–245). Lawrence Erlbaum Associates. https://doi.org/10.1201/9781410607775-8

Kaber, D., & Endsley, M. R. (2004). Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics*, 47(9), 955–977.

Kim, D. B., Shao, G., & Jones, A. (2023). Conceptual architecture of digital twins with human-in-the-loop–based smart manufacturing. *NIST Publication*. https://doi.org/10.6028/NIST.~IR.8477

Lyons, J. B. (2014). Transparency in a human–machine context: Approaches for fostering shared awareness and intent. In *Foundations of Augmented Cognition*. Springer. https://doi.org/10.1007/978-3-319-07458-0_18

Mikulka, P. J., Scerbo, M. W., & Freeman, F. G. (2002). Effects of a biocybernetic system on vigilance performance. *Human Factors, 44*(4), 654–664. https://doi.org/10.1518/0018720024496944

Onnasch, L., Wickens, C. D., Li, H., & Manzey, D. (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

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A*, 30(3), 286–297. https://pubmed.ncbi.nlm.nih.gov/11760769/

Rusnock, C. F., & Geiger, C. D. (2016). Designing adaptive systems: Selecting an adaptive automation invoking threshold to improve human performance. *International Journal of Human Factors and Ergonomics, 4*(3–4), 292–315. https://doi.org/10.1504/IJHFE.2016.083503

Sauer, J., Nickel, P., Wastell, D., & McKenna, I. (2012). A comparison of adaptive and adaptable automation under different levels of environmental stress. *Ergonomics*, *55*(10), 1173–1188. https://pubmed.ncbi.nlm.nih.gov/22506767/

Sauer, J., & Wastell, D. (2013). Designing automation for complex work environments under different levels of stress. *Ergonomics*, *56*(7), 1073–1085.

Tatasciore, M., & Loft, S. (2024). Can increased automation transparency mitigate the effects of time pressure on automation use? *Applied Ergonomics, 114*, 104142. https://doi.org/10.1016/j.apergo.2023.104142

Tatasciore, M., Strickland, L., Loft, S. (2024). Transparency improves the accuracy of automation use but automation-confidence information does not. *Applied Ergonomics*. https://doi.org/10.1016/j.apergo.2024.104273