applied sciences

MDPI

Review

Human–AI Cognitive Division of Labor: A Structured Narrative Review and Integrative Framework

Chen Chen *, Guili Li, Shuai Yuan, Qisheng Xu and Chenyang Cai

College of Furnishings and Industrial Design, Nanjing Forestry University, 159 Longpan Road, Xuanwu District, Nanjing 210037, China; guilia@njfu.edu.cn (G.L.)
* Correspondence: cchen1990@njfu.edu.cn

Abstract

AI-assisted systems redistribute cognitive work before final decisions are made, yet broad labels such as assistance, automation, and human involvement do not fully describe these arrangements. This theory-led structured narrative review connects Function Allocation with research on complementarity, appropriate reliance, cognitive offloading, delegation, mixed initiative, agency, and recovery. It draws on concept-driven retrieval, citation chaining, and selected theoretical foundations. Source selection and coding were AI-assisted, with verification boundaries reported separately. Framework A is a configuration-oriented integrative reporting framework and design space organized around allocation object, performer, allocation authority, change dynamics, and recovery architecture. Framework B provides a conditional temporal process lens on enactment, monitoring, recovery, and reallocation. Together, they support comparison of configuration descriptions across heterogeneous studies while separating configuration, operating conditions, process, and outcomes. The five dimensions build on established constructs. Coding reproducibility, comparative usefulness, and effects on interface design have not been empirically established.

Keywords: human–AI collaboration; cognitive division of labor; function allocation; appropriate reliance; intelligent interface design; allocation authority; mixed initiative; human-centered AI; human–computer interaction; control recovery

1. Introduction

Human–AI collaboration extends beyond discrete advice in the reviewed literature. AI systems retrieve and filter evidence, generate intermediate content, synthesize sources, decompose workflows, assign tasks, and initiate assistance. These functions redistribute cognitive work before a final decision, shaping the information people encounter, the representations they build, and the responsibilities they retain for checking and correction. Consequently, two conditions described as AI-assisted may embody different mechanisms: the observed effect may arise from the function assigned to AI, the authority structure, or the available recovery route rather than from assistance amount alone [1–3].

The established Human Factors tradition of Function Allocation provides an essential foundation for understanding this shift. Function Allocation treats the distribution of work between humans and technology as a sociotechnical design problem rather than as a static comparison of isolated capabilities. Effective allocation must support coherent joint-system performance while preserving manageable workload, observability, coordination, and the capacity to respond when automation fails. This tradition also cautions against assuming that greater automation produces uniformly better outcomes: an allocation that improves

Check for updates

Academic Editor: Francesco Zirilli

Received: 7 September 2026
Revised: 31 October 2026
Accepted: 2 November 2026
Published: 5 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

routine efficiency may weaken situation awareness, monitoring, or failure correction, while nominally leaving a function to the human does not ensure meaningful involvement if the person lacks the information or time needed to perform it [4]. Across the reviewed domains, automation stage, authority allocation, workload, and fault conditions have each been shown to affect performance and human–system interaction in different studies [3,5–7].

Different research traditions describe human–AI arrangements at different analytical units, including functions, information-processing stages, task assignments, and cognitive operations. Broad labels do not fully specify the operation assigned or its place in the workflow. Assistance amount and assistance function therefore answer different reporting questions. This review uses these distinctions to support cross-study comparison, without treating incomplete reporting or composite conditions as evidence of conceptual conflation.

Execution and allocation governance likewise require separate descriptions. The actor who performs an operation need not establish or revise its assignment; final approval also does not necessarily identify allocation authority. Shared execution does not establish reciprocal reassignment rights, and an available override does not demonstrate successful recovery. These are analytical distinctions, not findings that all or most studies confuse the dimensions.

Evaluation is similarly multidimensional. Improvement over an unaided human does not establish complementarity when AI alone performs better, and accuracy or speed does not show whether users relied appropriately, inspected evidence, experienced control, retained competence, or recovered after error. Workload, autonomy, perceived control, responsibility, satisfaction, and recovery behavior can diverge from performance [8–11].

Human–AI Cognitive Division of Labor is developed here as a theory-informed approach to comparing configurations across these studies. Framework A is a configuration-oriented integrative reporting framework and design space. It asks what is allocated, who performs it, who controls the allocation, when allocation changes, and how control is recovered. These questions distinguish arrangements with identical performers under different controllers or similar assistance amounts directed at different operations. The configuration fields are recorded separately from operating conditions and observed outcomes. Framework B provides a conditional temporal process lens on execution, monitoring, recovery, and reallocation. Neither framework is a statistically validated causal model or a prescription for an optimal configuration.

Review Objective, Questions, and Contributions

The review examines how cognitive work is distributed between humans and AI, how assignments are governed and revised, and what provisions support recovery after allocation failure. It integrates studies that use different analytical units and develops four research transitions. The contribution is a common reporting structure grounded in Function Allocation, with configuration distinguished from conditions, process, and outcomes.

2. Positioning Against Existing Reviews

Table 1 distinguishes two complementary bases for positioning the review: representative review programs and foundational frameworks for automation, function allocation, adaptation, and joint control. Panel A retains five representative reviews; Panel B compares four foundational frameworks and the present synthesis. The comparison identifies established constructs and differences in analytical purpose, rather than treating topical coverage as evidence of novelty.

Table 1. Representative reviews and foundational frameworks relevant to human–AI cognitive allocation.

Source	Primary Scope	Corpus/Method or Unit	Organizing Lens	Main Contribution/Established Coverage	Relation to Current Review and Boundary
Panel A. Representative review programs					
Gomez et al. [12]	AI-assisted decision making	Systematic review of 105 articles	Interaction patterns and interactivity	Taxonomy of interaction protocols and reciprocal patterns	A close complement for interaction and mixed-initiative elements; the current review uses interaction patterns as evidence relevant to allocation control and change dynamics.
Mehrotra et al. [13]	Appropriate trust in Human–AI interaction	PRISMA-guided systematic review	Beliefs–intentions–actions mapping of appropriate trust	Clarifies definitions, measures, tasks, and trust-fostering approaches	Supplies the appropriate-reliance and verification component literature; the current review retains appropriate reliance in the outcome layer.
Richardson et al. [14]	Levels of automation (LoA)/autonomy taxonomies	Systematic review of taxonomy critiques and recommendations	Taxonomy structure, interpretation, and automation–autonomy distinctions	Synthesizes critiques and guidance for LoA taxonomy development and use	Provides historical and conceptual continuity; the current review is not a replacement LoA taxonomy.
Kargarnovin et al. [15]	Human–AI teaming and human–robot teaming	PRISMA-guided systematic review and bibliometric analysis of 104 empirical studies	Application domains and teaming factors	Maps where teaming has been evaluated and which factors are assessed	Provides broad teaming context; the current review applies an allocation-centered lens across domains.
Li and Tian [16]	AI–human collaborative decision making	Systematic review and bibliometric analysis of 627 articles	AI–human dynamics and decision typologies	Conceptualizes collaborative decision-making paradigms	A complementary decision-making synthesis; the current review focuses on allocation, governance, change, and recovery within and beyond decision tasks.
Panel B. Foundational frameworks and the present synthesis					
Parasuraman et al. [17]	Types and levels of automation	Conceptual model; information-processing functions	Acquisition, analysis, decision selection, action	Human approval/veto; performance costs and error consequences	Established function and control distinctions; no advantage of Framework A over this model has been demonstrated.

Table 1. *Cont.*

Source	Primary Scope	Corpus/Method or Unit	Organizing Lens	Main Contribution/Established Coverage	Relation to Current Review and Boundary
Feigh and Pritchett [4]	Effective function allocation	Requirements for joint human–technical systems	Capabilities, coherent work, coordination, dynamics	Assignment, workload, observability, intervention and system demands	Grounds the joint-system perspective; authority terminology requires local interpretation. Allocation was not previously treated only as static.
Feigh et al. [18]	Adaptive systems	Adaptation targets, management and triggers	Assignment, responsibility, authority; allocation, scheduling, interaction, content	Adaptation targets, initiators, duration and disengagement	Already separates execution and assignment changes; its adaptation taxonomy is more detailed than the compact Framework A change field.
Lundberg and Johansson [19]	Joint Control Framework (JCF)	Processes, agents, control levels, temporal episodes	Cognitive control, monitoring, system joints, sequential/parallel developments	Agent relations, hierarchical control and allocation over time	Closest comparator. A five-field record may lose hierarchical and event-level detail; superiority over JCF has not been established.
Present review	Human–AI cognitive allocation	Theory-led structured narrative review	Operation- and condition-specific configuration; separate conditions, process, outcomes	Collects What, Who, Who controls, When and recovery provisions	Intended common comparison format. Constructs are established; coding reproducibility, comparative benefit and design effects require empirical evaluation.

Note. Panel A summarizes representative review programs; Panel B compares foundational frameworks and the present synthesis. Differences in analytical unit or reporting format do not imply that a framework lacks the underlying construct. The present-review row states intended use rather than validated superiority.

The proposed distinction concerns the reporting structure and analytical unit. Interaction-pattern, trust, automation, teaming, and decision-making syntheses supply complementary perspectives. Function Allocation, adaptive-systems frameworks, and the Joint Control Framework (JCF) already address combinations of work, agents, authority, coordination, monitoring, and change. Framework A collects selected configuration questions in a compact format; it does not replace richer representations within those traditions. Figure 1 is a conceptual genealogy of overlapping research emphases, not an empirical measure of their relative coverage.

For a single AI-assisted system, the reporting task is to identify the operation, performer, allocation rights, change rules, and recovery provisions. Existing theories describe many of these attributes; the configuration lens makes selected attributes explicit across reports. Whether this format reduces comparison effort or improves reproducibility requires evaluation.

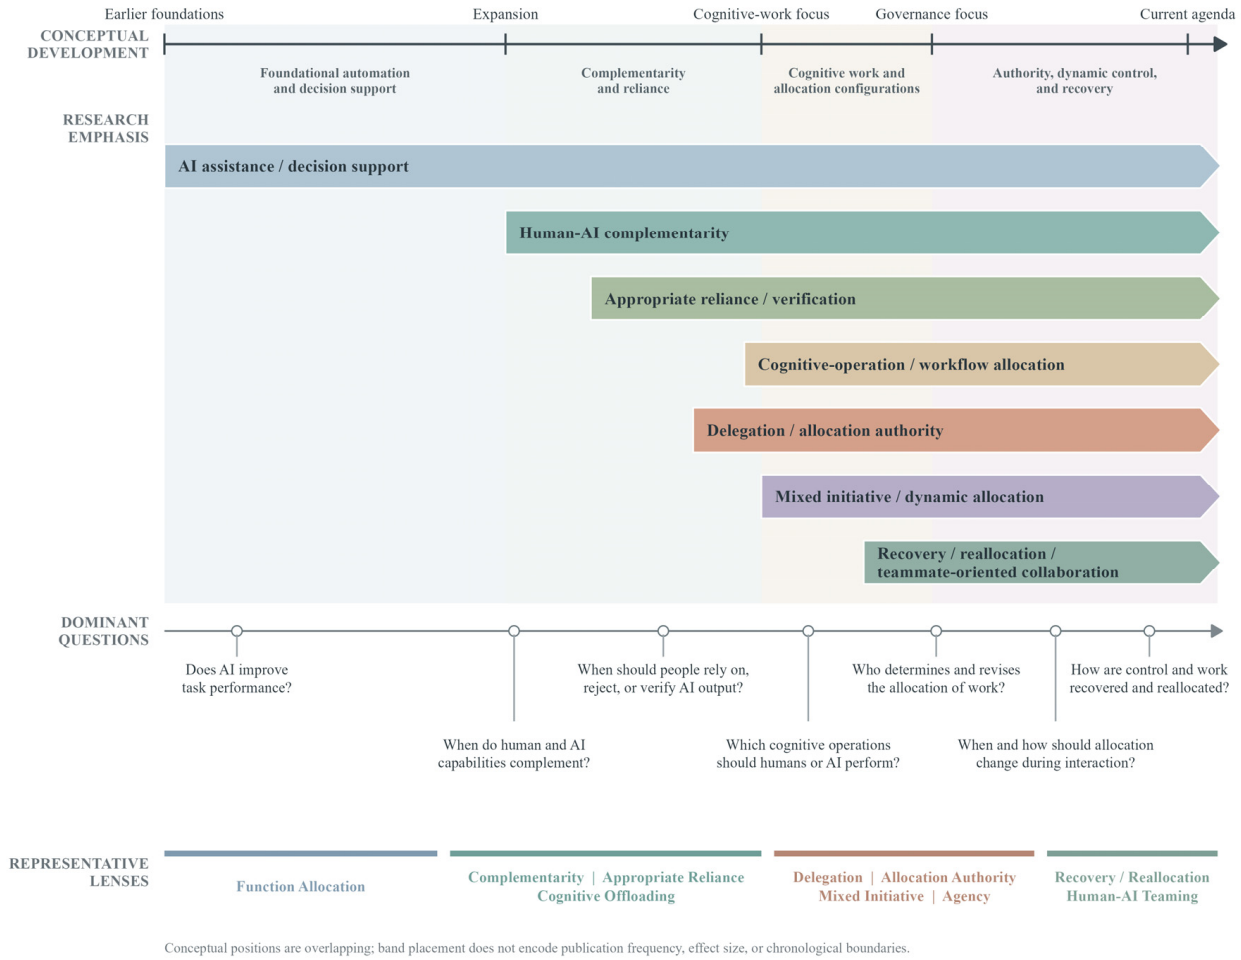


Figure 1. Conceptual genealogy of research emphases in human–AI cognitive division of labor. Horizontal placement represents a progression in organizing questions, and overlapping bands indicate coexisting research traditions. Band length is interpretive and carries no quantitative meaning for field size, duration, publication frequency, or evidential strength. Colors distinguish research emphases for visual clarity and do not encode quantitative values or evidential strength.

3. Review Methods

3.1. Review Design

This article uses a theory-led structured narrative review to connect mature Human Factors concepts with research in human–computer interaction (HCI), human–AI collaboration, organizational research, decision science, education, and human–robot interaction (HRI). Function Allocation provides the systems foundation; empirical studies and high-quality reviews clarify cognitive operations, authority, initiative, agency, and recovery. Literature identification was iterative and concept-driven, supporting cross-literature conceptual integration rather than exhaustive retrieval or field-wide prevalence estimation.

Initial searches from 29 to 31 August 2026 identified conceptual traditions, anchor studies, and terminology across adjacent studies. Subsequent targeted retrieval, backward and forward citation chaining, benchmark-review comparison, and publication-version verification refined the analytical distinctions. An update in OpenAlex, Crossref, and PubMed on 5 September 2026 supported this concept-driven development. The initial searches and later expansion served different purposes rather than applying a single uniform screening protocol. Retrieval was non-exhaustive: database interfaces returned capped relevance-ranked subsets, and later checks identified further eligible sources.

3.2. Scope and Literature Identification

The scope covers studies that explain how cognitive work is allocated, performed, controlled, changed, monitored, or recovered across human and artificial agents. The contemporary search window covered 2021–2026; earlier sources were retained selectively when they supplied established foundations in Function Allocation, automation, cognitive offloading, or mixed-initiative systems. The review includes human–AI complementarity, appropriate reliance, delegation, allocation authority, agency, and recovery. General adoption studies and attitude-only work were used only when they clarified one of these constructs.

Sources included OpenAlex, Crossref and digital object identifier (DOI) records, PubMed, the Association for Computing Machinery (ACM) Digital Library, publisher and institutional repositories, and authorized library access. Search modules combined human–AI terms with function allocation, delegation, authority, mixed initiative, complementarity, reliance, offloading, agency, override, and recovery. Supplementary Table S1 documents the historical search modules and dates. Revision-stage checks on 21 October 2026 were restricted to six sources with material citation or interpretation issues; they were not a renewed comprehensive search. On 22 October 2026, an additional AI-assisted citation check compared the identities and bibliographic metadata of all 86 references and examined 24 selected attribution statements. It was a targeted citation check, not independent human full-text verification of all cited studies.

A representative query combined ("human-AI" OR "human automation" OR "human-machine" OR "human-robot") with ("function allocation" OR "task allocation" OR delegation OR "allocation authority" OR "mixed initiative" OR complementarity OR "appropriate reliance" OR "cognitive offloading" OR agency OR override OR "control recovery"). Search terms were adapted to database syntax and to the terminology used in each component literature.

3.3. Source Selection and Verification

Selection was guided by conceptual relevance to an allocation dimension, temporal process, operating condition or outcome, theoretical foundation, or benchmark synthesis. Searches and candidate screening were assisted by Codex; the surviving records establish the source sets and later decisions, but do not preserve a complete contemporaneous screening history. Conceptual relevance describes the selection rationale, not a reconstructed exclusion decision. Preprints are identified by the examined version rather than assumed equivalent to proceedings versions.

The working database contained 107 records. A 60-paper working subset informed development of the analytic lens. A later concept-targeted expansion considered 30 candidates and retained 22 records, including one Dai study already represented in the core; it therefore added 21 new bibliography entries. The 86 cited sources comprise 55 core citations + 22 retained expansion citations − 1 overlapping study + 10 benchmark or theoretical sources. Five core members are not cited in the current bibliography. These are related source sets, not stages of a uniform systematic-review funnel (Supplementary Table S4).

The historical selection record is incomplete: the early screening snapshot has no populated exclusion reasons, and some record identifiers changed paper identity between snapshots. Later reconstruction cannot be treated as contemporaneous screening. The eight expansion candidates not retained comprise seven optional records and one do-not-add record, rather than eight documented full-text exclusions. No formal conceptual saturation, exhaustive coverage, or field-wide prevalence estimate is claimed.

The 86 citations are not a uniformly coded empirical sample. Supplementary Table S8 now indexes every bibliography entry by source-set membership, explicit evidence-use roles, its actual cited use and location in this manuscript, and the applicable checking boundary.

Membership follows the reconciled study identities, not historical identifiers alone. The index is a retrospective account of the present synthesis: it makes uses traceable but does not supply missing original exclusion reasons or independently validate every statement. Role labels distinguish theoretical/conceptual foundations, benchmark syntheses, empirical findings, technical/workflow illustrations, configuration/construct mappings, and research-agenda support; multiple uses may apply. They describe current citation functions, not study designs or evidence quality. Necessity rankings are not evidence types. Tables S5–S7 provide more detailed version, theoretical and condition-level anchors for selected sources.

3.4. Analytic Synthesis and Claim Calibration

The authors led literature discovery, screening, extraction, and five-dimensional coding, with Codex assistance. Revision-stage work linked interpretations to archived passages and flagged contradictions, ambiguity, and version issues. Original working codes and candidate corrections were kept separate. Seven papers were selected purposively for contrasting configurations and inspectable method passages. Their tasks, conditions, operations, performers, allocation rights, change mechanisms, and recovery provisions were checked against archived full texts for Table S7. Table S8 links reconciled source identities to current citation contexts. Its evidence-use labels are AI-assisted retrospective classifications, not independently human-validated coding. No dual-coder adjudication or inter-rater reliability estimate is reported.

The 60 × 5 working codes are not reported as validated findings or released as a human-validated dataset. Table S7 presents seven condition-specific comparisons using this review's interpretation of identified passages, rather than categories assigned by the original authors. "Not reported" marks an evidential limit in the inspected description, not the absence of a system property. Content editing is distinguished from operation reallocation, and intervention opportunities from observed restoration. The comparisons illustrate analytical distinctions without estimating their prevalence in the core or across the field.

4. Theoretical Synthesis of Component Studies

4.1. Function Allocation

Human–AI cognitive division of labor extends a longstanding Human Factors problem: distributing functions across people and technical systems while preserving joint-system effectiveness, observability, and recoverability. Feigh and Pritchett connect allocation to coherent responsibilities, agent capability, coordination, and intervention [4]. Roth et al. also emphasize operational demands, interdependencies, and alternatives under routine and off-nominal conditions [20]. Allocation was therefore not previously treated only as a static comparison of isolated capabilities. Adaptive-allocation and team-centered accounts extend these concerns to flexible authority and changing operational demands [21,22].

Fixed, operator-adjustable, system-adaptive, and negotiated arrangements distribute both work and authority. Adaptable automation preserves human-initiated change; adaptive automation lets the system revise allocation under specified conditions; hybrid and mixed-initiative arrangements combine routes. Studies of robotics, production, and multitasking show that the locus of change can alter workload, situation awareness, performance, and experienced influence [5,7,9]. Work on socially shared regulation extends this problem to collaborative cognition, where allocation must respond to changing group states while retaining intelligible human control [23,24]. This literature supplies the controller and change-dynamics distinctions, although it usually describes work at a broader level than individual cognitive operations.

4.2. Human–AI Complementarity

Human–AI complementarity is often invoked to justify combining human judgment with algorithmic prediction. Hybrid Intelligence presents the combination of complementary strengths as a design objective [25]; this does not guarantee superior joint performance. Improvement over an unaided human establishes assistance, whereas complementarity also requires the joint system to outperform AI alone. Meta-analytic evidence shows that human–AI combinations can fail to exceed their stronger component [26]. In contrast, Bansal et al. obtained complementary team performance, but explanation conditions did not significantly outperform a confidence-only baseline [27].

Allocation should therefore reflect comparative capability and error structure rather than a global ranking of human and AI performance. Category-specific competence can enable complementary expertise when users recognize where each agent performs well [28,29]. Learning systems can be optimized for complementary rather than stand-alone performance [30], while team studies show that generative AI changes both decision processes and decision quality [31]. Role research shows why an AI teammate label is insufficient: execution, responsibility, and adaptation still require specification [32,33]. Complementarity is therefore an outcome of a configuration, while comparative human and AI capabilities are conditions under which that configuration operates.

4.3. Appropriate Reliance and Verification

Appropriate reliance is a behavioral relation between the quality of an AI contribution and the user's response. Trust, agreement, and frequency of use may accompany that relation but do not define it. Accepting correct advice and rejecting incorrect advice are distinct achievements, and aggregate agreement can obscure failure on either component. High agreement can reflect indiscriminate acceptance, while low agreement can reflect either effective error rejection or rejection of valid advice. This separation prevents favorable trust ratings from being interpreted as calibrated behavior. Studies therefore operationalize appropriate reliance through behavior conditional on AI correctness rather than favorable attitudes alone [27,34–36].

An incorrect AI output is an event-level property, whereas system reliability describes performance across repeated outputs, histories, or operating conditions. The former guides a trial-level response; the latter shapes expectations about later recommendations. Experiments with changing automation accuracy indicate that people update reliability judgments from experience, with substantial individual variation. The study modeled those judgments rather than subsequent reliance or team performance [37]. In a multitask monitoring experiment, action-implementation automation was associated with fewer fixations on the automated task and poorer malfunction correction than other automation stages. The study did not find corresponding significant effects of reliability on these automated-task outcomes [3].

Verification requires users to inspect diagnostic information and decide whether an AI contribution should be revised. Explanations can expose evidence or substitute persuasion for checking [38–40]. Behavioral and cognitive studies distinguish selective adherence, observed reliance, and latent reliance strategies [41–43], while reviews and clinical experiments identify expertise, workload, explanation complexity, and control-related beliefs as relevant conditions [44,45]. Warren et al.'s fact-checking preprint reports evidence inspection alongside AI explanations [46]; its findings are task-specific and do not isolate a general effect of evidence access alone. Correctness and reliability history are operating conditions; reliance and verification are outcomes.

4.4. Cognitive Offloading and AI Scaffolding

Cognitive offloading uses external agents or representations to reduce or transform internally performed cognitive work [47]. Generative AI (GenAI) extends offloading from memory and calculation to planning, drafting, analysis, and evidence-related work. These forms cannot be summarized by assistance amount alone. Similar volumes of generated text may support different operations, while the same operation can be offered with different initiative and control.

Experimental evidence in learning shows that GenAI can change performance, motivation, and self-regulation, including conditions described as metacognitive laziness [48]. A longitudinal survey distinguishes dependent offloading, in which AI substitutes core thinking, from autonomous offloading that preserves user agency [49]; its design supports association rather than causal inference. Capability-oriented accounts likewise treat human agency and the ability to regulate automation as part of effective learning with AI [50]. Offloading research thus helps specify the allocation object and its cognitive consequences: it asks which operation is delegated away from the person, what oversight remains, and how the redistribution affects effort and agency.

4.5. Delegation and Task Allocation

Delegation separates the performer from the agent that decides who performs the work. A person can delegate a case to AI, an AI can route a case to a human, and both can contribute while an organizational rule controls allocation. Human delegation studies show that contextual information changes routing behavior [51]. Learning-to-defer models formalize case-level routing to an external decision-maker and evaluate system performance or fairness [52], while medical second-opinion models route cases by predicted expert disagreement rather than classification uncertainty alone [53]. Together, these studies make the performer–controller distinction explicit and demonstrate feasible routing mechanisms. Their main outcomes are system performance and fairness, leaving experienced authority, agency, and behavioral recovery to other research traditions.

4.6. Cognitive-Stage and Cognitive-Operation Allocation

A stage identifies position in a process, but its meaning varies across the literature. Automation research distinguishes information acquisition, analysis, decision selection, and action implementation [17]. Workflow studies use stages for editable sequences, advice studies distinguish pre- and post-judgment intervention, and cognitive-assistance studies refer to semantic operations such as evidence selection or integration. These structures are related but not interchangeable.

Cognitive operations name transformations applied to information: searching, selecting, generating, analyzing, integrating, verifying, recommending, and judging. Human–AI analysis systems based on large language models (LLMs) divide work across codebook application, coding, theme construction, iteration, and intervention [54–59]. The DeTAILS preprint illustrates an iterative coding-and-theme workflow with LLM support [60]. HICode separates label generation, hierarchical clustering, and human evaluation [61]. Studies of co-writing, user experience (UX) practice, and serendipitous discovery show that desired AI involvement varies across search, planning, drafting, revision, and evaluation [62–65]. A stage locates an operation within a process; operation-level analysis specifies the transformation assigned. Comparing operations requires comparable authority, information quality, interface friction, and performer options, rather than a stage label alone.

4.7. Allocation Authority

Allocation authority concerns the right to determine or revise the distribution of work. It differs from performer identity and final-decision authority. Scenario studies distribute decision weight across human-led, joint, and AI-delegated arrangements [66], but assigned weight does not by itself establish enacted organizational authority. Philosophical analysis also rejects a simple zero-sum relation between robot autonomy and human autonomy [67].

Authority is more informative when decomposed into proposal, approval, contestation, override, and reallocation rights. Human final approval may coexist with AI control over retrieval, routing, or workflow. Workload can affect preferences for decision authority [68], and human-control scholarship distinguishes supervision from more reciprocal teaming arrangements [69]. Autonomy-preserving design depends on competence, reflective engagement, and usable opportunities to recognize and interrupt faulty AI pathways [70]. Deployed clinical decision support shows how alerts can preserve clinicians' formal judgment while changing which cases receive attention [71]. This literature anchors allocation control as an architectural dimension and agency as a related outcome rather than treating the two as interchangeable.

4.8. Mixed Initiative and Dynamic Allocation

Dynamic allocation changes the distribution of work during performance. Humans may initiate adjustment, systems may adapt through rules or capability estimates, or both may negotiate. Shared execution alone is insufficient: mixed initiative requires reciprocal opportunities to propose or revise allocation. Human–robot studies demonstrate reciprocal initiative directly [72,73], while studies of proactive agents and teammate influence show how initiative and influence affect competence, satisfaction, and team processes [74,75]. Team-centered and socially shared regulation accounts connect these mechanisms to communication, state awareness, and changing task demands [22,23].

Communication is part of the mechanism because allocation changes must be intelligible and coordinated. Interview evidence from multiplayer human–AI teams identifies proactive information sharing and timing as important for situation awareness and trust [76]. These findings clarify interaction requirements but do not isolate mixed initiative as a causal factor. Design-science work in collaborative learning likewise shows that an agent organized around regulation stages may fail to facilitate the intended process [24], underscoring the need to evaluate enacted coordination rather than interface labels.

Chiou et al. compared human-initiative and mixed-initiative mode switching [72]. In the MICoBot preprint examined here, human and robot can formulate, accept, or reject proposals about who completes task steps [73]. These implemented capabilities illustrate reciprocal initiative. The examined study includes human-only and robot-only initiative ablations in simulation; however, the overall findings from the physical-robot evaluations concern the combined dialog, planning, capability-estimation, and allocation system and should not be attributed to mixed initiative alone. Reference [73] identifies the examined preprint revision, without assuming equivalence to conference proceedings.

4.9. Agency, Override, and Recovery

Agency-related outcomes include experienced influence, autonomy, competence, responsibility, and practical intervention capacity. Usable control also depends on information, timing, and the ability to act [8,69,70,77,78]. Empirical studies report task-specific effects or associations involving self-regulation, ownership, and autonomy [48,49,64,65,79]. Conceptual accounts offer complementary perspectives on agency and learning with AI [50,80]. Deployed clinical evidence illustrates how preserving professional judgment can coexist with AI-triggered attention shifts [71].

Override cancels or supersedes an AI contribution; takeover transfers operative control to a person. Recovery additionally requires error detection, state reconstruction, correction, and restoration of an acceptable state. An intervention mechanism creates an opportunity to recover rather than evidence that recovery succeeded. Relevant task-state information can support judgment [81], and driving studies distinguish automation deactivation from recovery of visual attention [82]. By analogy, recovery in evidence-intensive cognitive work may require reopening sources and reconstructing downstream reasoning; this is a proposed cross-domain extension, not a directly tested result of those studies.

More direct evidence on recovery comes from driving, robotics, and fault arbitration, where takeover, state updating, replanning, and control return are observable [6,82,83]. Cognitive-work studies more often examine individual components, such as source checking, correction, or retained authority, while autonomy-oriented design work identifies reflective engagement and override as requirements [70]. How these components combine when an AI error propagates through intermediate representations and requires both task repair and reallocation remains uncertain.

Supplementary Table S7 makes these distinctions concrete across seven source-located examples [1,6,8,27,35,73,82]. In Tausch and Kluge's vignette, the management and robot conditions prescribe the same four-plus-four task split but attribute the allocation decision to different agents [8]. In Ma et al., adaptive recommendation timing and visibility change while humans retain the final income judgment [35]; Bansal et al.'s adaptive explanations change displayed evidence rather than the identity of the final decision maker [27]. AI Chains permits editing intermediate results, which does not by itself demonstrate reassignment or complete recovery [1]. By contrast, Yu et al. explicitly represent human/robot step assignments and dialog-dependent constraints [73], while Bhardwaj et al. specify distinct takeover and return mechanisms [6]. Vogelpohl et al. measure motor and visual responses separately [82]. The comparison therefore identifies which property changes and what outcome would need measurement; it neither estimates conceptual confusion nor treats these heterogeneous interventions as isolated tests of five independent dimensions. Production and driving examples inform allocation governance and recovery; they do not establish equivalence between physical task steps and cognitive operations.

5. Integrated Human–AI Cognitive Division of Labor Framework

5.1. Rationale and Status

Framework A organizes constructs that are already established in automation, function allocation, adaptive systems, and joint-control research. Parasuraman, Sheridan and Wickens distinguish types and levels of automation, including arrangements involving human approval and veto [17]. Feigh and Pritchett connect allocation to agent capabilities, coherent work, coordination, and changing operational demands [4]. Feigh et al. distinguish assignment, responsibility, authority, and different targets and triggers of adaptation [18]. JCF describes processes and acting agents, cognitive control, monitoring, system joints, and changes in allocation over time [19]. Framework A therefore does not introduce five previously unrecognized constructs or supersede these accounts. Function-allocation research also addresses measurement [84].

Its proposed contribution is a common reporting structure for comparing how established constructs are configured in heterogeneous studies. The five fields identify the allocated operation, its performer, the rights to establish or revise allocation, the rules governing change, and the provisions available for recovery. These fields should be specified for the relevant operation and experimental or system condition; different conditions within a paper should not be combined into a configuration that was never implemented. Operating conditions, temporal enactment, and observed outcomes are recorded separately. This

separation is an analytical convention for reporting, not a claim that the dimensions are statistically independent or empirically orthogonal.

Consider the observed driving configuration in Bhardwaj et al. [6] (Supplementary Table S7; archived PDF pp. 16–18). In Autopilot, the driver monitors automation, presses a button to take full steering control and may press it again to re-engage automation. In Active Safety, automation takes over when it detects a collision risk, then switches off after obstacle avoidance; returning to the centerline is the driver's responsibility. These conditions share an obstacle-avoidance task and human–automation participants, but differ in who initiates the control transition and how control returns. A five-field report retains steering/avoidance as What, the condition-specific performer as Who, transition rights as Who controls, obstacle detection or driver action as When, and disengagement, human steering and re-engagement provisions as Recovery. The paper's collision and trajectory measures remain outcomes, not labels for the allocation architecture.

The same driving case is describable through all four established accounts. Types and levels of automation distinguish automated action implementation from monitoring and human intervention [17]. Effective-allocation requirements ask whether the agents can perform their assigned work and coordinate its timing and intervention [4]. The adaptive-systems taxonomy distinguishes changes in assignment or authority from changes to control content and identifies their triggers [18]. JCF can represent the driver and automation as acting agents, the relevant processes and control relations, monitoring, and their temporal reconfiguration [19]. Table S6 supplies the original theoretical anchors; Table S7 supplies the empirical case anchors. Framework A reorganizes selected attributes into common fields so that the same task and participants are not mistaken for the same control arrangement. This is an interpretive application of existing concepts, not evidence that those theories cannot express the distinction.

Relative to JCF, the compact record preserves the operation, performer, allocation rights, change trigger, and recovery provisions needed for this comparison. It omits explicit agent/process hierarchies, system-joint structure, detailed monitoring relations, and the temporal trajectory of interaction. Framework B summarizes enactment without recovering all of that detail. The proposed increment is a consistent reporting format across papers, with unknown or unreported attributes left visible. Table S7 illustrates the distinctions the record can display; it does not establish faster comparison, more reliable coding, better design decisions, or superior explanatory power. These potential benefits require empirical evaluation.

What is allocated × Who performs it × Who controls the allocation × When allocation changes × How control is recovered.

The multiplication sign denotes joint description rather than a statistical interaction.

5.2. Framework A: Five Architectural Dimensions

Figure 2 presents the architecture; Table 2 translates it into a design-reporting and comparison checklist; Table 3 traces each dimension to its component literature and clarifies the added integrative step.

An illustrative multi-source decision system shows how the reporting questions work together. AI may select evidence while a person integrates it into a judgment. The person may invoke or revise that arrangement, conflicting evidence may trigger reallocation, and the interface may preserve source access and correction routes. The same reporting questions can be applied across domains. In clinical decision support, the allocation object may be a case or evidence set, with recovery requiring access to source data and revision of a recommendation. Adaptive automation allocates monitoring or control operations and may require takeover information. Human–robot coordination allocates task steps and

may permit negotiated reassignment. Applying common questions does not make these interventions equivalent.

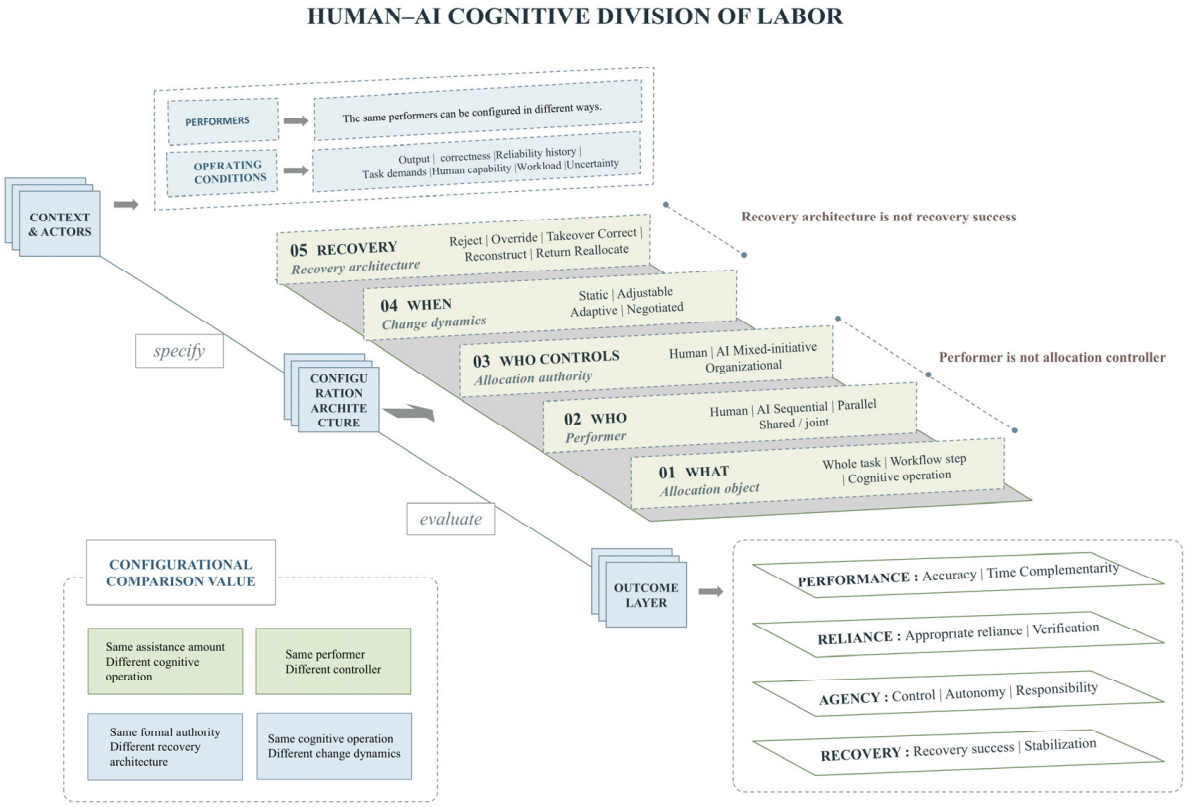


Figure 2. Framework A: an integrative design space for allocation object, performer, allocation controller, change dynamics, and recovery architecture. Operating conditions and outcomes remain separate analytical layers; the multiplication sign denotes joint description, not statistical interaction.

Table 2. Framework A design-reporting and comparison checklist.

Dimension	Required Specification	Comparison and Evaluation	Key Boundary
What	Allocation object and cognitive-operation granularity	Compare selection vs. integration or generation vs. verification; assess accuracy, verification, workload, and complementarity	Assistance amount does not identify assistance function
Who	Human, AI, sequential, parallel, or shared/joint performance	Compare the same operation under different performer arrangements; assess performance, time, workload, and complementarity	Performer does not identify allocation controller
Who controls	Proposal, approval, contestation, override, and revision rights	Compare the same performer under human-, AI-, or mixed-initiative control; assess perceived control, autonomy/agency, reliance, and responsibility	Formal authority differs from perceived control and final-decision authority
When	Change trigger, timing, initiator, and negotiation rule	Compare static with adjustable, adaptive, or negotiated allocation; assess adaptation burden, workload, agency, and recovery opportunity	A dynamic task does not necessarily use dynamic allocation

Table 2. *Cont.*

Dimension	Required Specification	Comparison and Evaluation	Key Boundary
Recovery	State access, intervention, reconstruction, correction, control return, and reallocation provisions	Compare the same authority with different recovery provisions; assess recovery success, performance, workload, and agency	Override or takeover does not establish successful recovery

Note. The table operationalizes Framework A as a design-reporting and comparison checklist. It does not prescribe an optimal configuration.

Table 3. Literature basis and integrative role of Framework A dimensions.

Dimension	Component Literature and Representative Evidence	Established Construct	Integrative Role in Framework A
What	Function Allocation [4]; automation-stage taxonomy [17]; chain and coding workflows [1,54]; operation-level workflow studies [24,61,64,65]	Functions, stages, workflow steps, and cognitive operations can be distinguished	Treats granularity and semantic operation as the allocation object
Who	Complementarity, delegation, and task-allocation research [5,28,29,52,85,86]	Human, AI, sequential, parallel, and joint performer arrangements are observable	Separates execution arrangement from authority over allocation
Who controls	Allocation-authority, meaningful-control, delegation, and mixed-initiative work [5,6,72,73,77]	Proposal, approval, contestation, override, and revision rights can be distributed	Combines these rights as an allocation-controller architecture
When	Adaptive automation and adaptation targets/triggers [18]; dynamic allocation and mixed-initiative interaction [21,22,72,73]	Allocation may be static, adjustable, adaptive, or negotiated	Makes triggers, timing, initiation, and negotiation jointly reportable
Recovery	Takeover, fault arbitration, interruption recovery, and replanning [6,82,83]; autonomy-preserving design principles [70]	Intervention, reconstruction, correction, control return, and reallocation are distinct elements	Treats available recovery provisions as an architectural property and their enactment as a process

Note. Representative references indicate the component studies and empirical traditions from which each Framework A dimension is synthesized.

What is allocated identifies the unit assigned within the joint system, from whole tasks and individual cases to workflow steps, automation stages, and cognitive operations. Operations such as search, selection, generation, analysis, integration, verification, recommendation, monitoring, and judgment specify the transformation performed on information. Who performs it records whether the operative output is produced by a human, AI, a sequence, parallel contributions, or joint execution. Reporting both dimensions helps distinguish equal amounts of assistance when they target different operations, and separates performer identity from governance.

Who controls the allocation records the person, system, or organizational rule authorized to determine, approve, contest, or revise who performs a specified operation or task. For cross-literature comparison, this review decomposes allocation control into proposal, approval, contestation, override, and reallocation rights. This is an authorial analytic decomposition informed by delegation, mixed-initiative, authority-allocation, and meaningful-human-control research rather than a scale already established in one source [5,6,51–53,72,73,77,78]. When allocation changes specifies whether the ar-

rangement is static, user-adjustable, system-adaptive, hybrid, or negotiated, together with the states or events that trigger revision.

Recovery architecture comprises the rights, information, interfaces, and system capabilities that support detection, intervention, state reconstruction, correction, stabilization, and control return. These elements are synthesized here across takeover, fault-authority arbitration, interruption, and replanning studies [6,82,83]; the full decomposition is therefore an authorial integration of established component processes, not a validated recovery scale or universal sequence. Recovery belongs to Framework A when it records provisions available before failure. The order in which those provisions are activated—and whether they restore an acceptable state—belongs to Framework B. Override availability marks an intervention opportunity, whereas recovery success requires an observed outcome.

The dimensions are distinguished for reporting, while their empirical discriminability and dependencies require evaluation. Mixed initiative spans allocation control and change dynamics because it concerns reciprocal rights exercised over time. Recovery may culminate in reallocation, linking recovery provisions to the operation, performer, and controller. Their relations should be specified rather than presumed orthogonal, additive, or fixed.

5.3. Architecture, Operating Conditions, and Outcomes

Framework A records configuration properties. Operating conditions describe the circumstances in which a configuration is used, whereas Framework B tracks its temporal enactment. Outcomes evaluate what follows. Output correctness, reliability history, task demands, human capability, workload, and uncertainty can shape performance; workload can also be measured as an outcome. Other outcomes include complementarity, reliance, verification, cognitive cost, perceived control, autonomy, competence, responsibility, satisfaction, and recovery success. Formal authority must be assessed separately from perceived control, and an available override separately from demonstrated recovery. Monitoring is not a sixth dimension because it can be the cognitive operation allocated or a process element during enactment.

5.4. Framework B: Conditional Temporal Process Lens

Figure 3 presents Framework B as a conditional temporal process lens.

Framework B is a conditional temporal process lens for organizing evidence about how a configuration is enacted, monitored, interrupted, repaired, and revised. Its purpose overlaps with joint-control analysis, including JCF accounts of cognitive control, monitoring, interaction at system joints, and changing allocations [19]. It aligns temporal evidence with Framework A without claiming a new process theory or a universal, independently validated sequence:

$$\text{Allocation} \rightarrow \text{Execution} \rightarrow \text{Monitoring} \rightarrow \text{Recovery} \rightarrow \text{Reallocation}$$

When disruption requires recovery, a conditional analytic decomposition is:

$$\text{Allocation} \rightarrow \text{Execution} \rightarrow \text{Monitoring} \rightarrow \text{Error Detection} \rightarrow \text{Override/Takeover} \rightarrow \text{State Reconstruction} \rightarrow \text{Correction/Replanning} \rightarrow \text{Stabilization} \rightarrow \text{Control Return} \rightarrow \text{Reallocation}$$

The extended chain integrates process elements reported in takeover, fault-arbitration, interruption, and replanning studies [6,82,83]. Individual studies support particular elements, not the complete sequence across domains. Elements may be skipped, repeated, or entered through different routes. The chain is therefore a conditional decomposition, with successful recovery requiring restoration of an acceptable state.

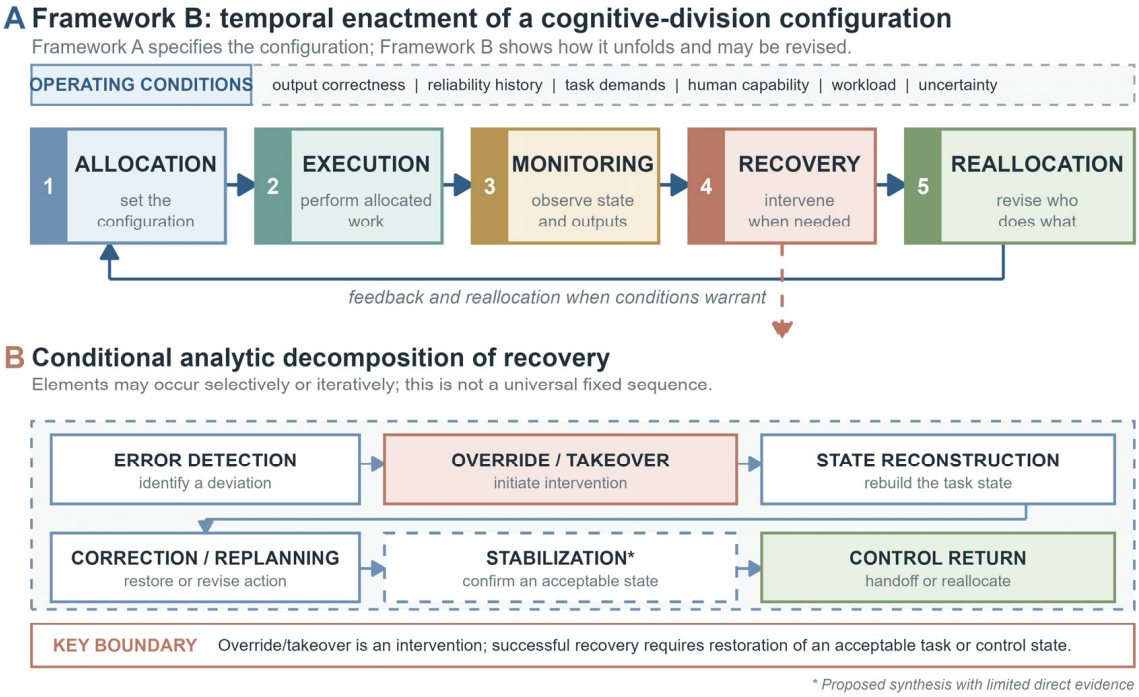


Figure 3. Framework B: Allocation → Execution → Monitoring → Recovery → Reallocation. The extended recovery chain is a conditional analytic decomposition; override or takeover does not itself demonstrate recovery success.

6. Discussion

The configuration lens directs attention to the design of cognitive work as well as the amount of AI involvement. Selection and synthesis identify different operations; performers, allocation controllers, and recovery provisions also require separate descriptions. Differences in these features are candidate explanations for apparently conflicting findings, but this synthesis does not isolate their causal contributions. Such explanations require designs that address changes in task, information, interface, and conditions.

The framework suggests testable design questions. Does allocation authority alter checking responsibility or intervention timing? Do selection, integration, generation, and judgment differ in the information available for detecting and correcting error? Do state access and correction routes support recovery under specified faults? Verification frequency, error-detection rate, correction accuracy, restoration quality, and time to recover are possible outcomes. Their measurement and causal interpretation must follow the study design.

Performance is necessary but insufficient when a design redistributes authority, checking, competence, or recoverability. A faster or more accurate configuration may impose verification costs, reduce experienced control, alter responsibility, or leave users unable to repair an error. Measures should follow the allocation problem: performance and complementarity for capability claims, reliance and verification for advice use, workload for coordination cost, agency and perceived control for governance, and demonstrated recovery for failure response. Interface evaluation should also assess whether users can understand allocation, intervene when necessary, and recover from AI error, alongside usability and task performance.

During an early usability audit or requirements review of a GenAI writing or qualitative-analysis tool, designers can apply Table 2 to one concrete workflow. They can identify the cognitive operations and performers, then record who may propose, approve, reject, or revise each assignment. For each change, the audit should specify its trigger and the information available to the user. A failure scenario can then examine source access, affected

intermediate outputs, intervention, correction, and rechecking. Unspecified rights or recovery steps become requirements to resolve. This use structures design questions; its effects on usability and recovery performance remain to be evaluated.

Experiments could isolate a focal contrast while holding other dimensions comparable, such as the same performer under different controllers or the same operation under different change rules. Such designs would test the analytical distinctions and their relations to monitoring, interruption, repair, and revision. This review does not establish that the five-field format itself improves causal attribution.

7. Future Research Agenda

Four recurring comparison problems define the next empirical steps: operation-level allocation, allocation authority, performance–agency alignment, and recoverable reallocation.

7.1. Transition 1. From Broad Automation Stages to Cognitively Specific Allocation

Research already distinguishes information–action stages, advice timing, partial automation, and assistance granularity; GenAI studies also map support across writing, UX, qualitative analysis, and learning [1,24,48,54,61,64,65]. A useful next step is a matched contrast between semantic operations under comparable information quality, authority, and interface friction. Experiments should isolate operations such as selection versus integration, search versus synthesis, generation versus verification, and recommendation versus judgment. Comparing joint performance with human-only and AI-only baselines can establish whether an advantage reflects complementarity, while the contrasts and outcomes in Table 4 specify what must remain comparable.

Table 4. Research transitions.

Research Transition	Testable Contrast	Key Outcomes	Boundary Conditions
Broad stages → cognitively specific allocation	Matched operation-level contrasts, including selection/integration, search/synthesis, generation/verification, and recommendation/judgment	Accuracy, verification, confidence, time, mental effort	Expertise, source conflict, task complexity, interface friction
Allocation outcomes → allocation authority architectures	Human-controlled, AI-controlled, and genuinely mixed-initiative allocation with comparable objects and performer options	Performance, perceived control, intervention, recovery success	Output correctness, reliability history, final-decision authority, stakes
Performance optimization → performance–agency alignment	Configurations matched on task while authority and participation rights vary	Performance, workload, verification, autonomy/agency, responsibility	Expertise, workload, error structure, repeated use
Static control → recoverable and reallocatable control	Recovery provisions that separately vary detection, intervention, reconstruction, correction, control return, and reallocation	Restoration quality, recovery time, residual error, perceived control	Fault source, state visibility, time pressure, system reliability history

Note. The contrasts are research priorities derived from the synthesis, not a claim that one sequence or configuration is universally preferable.

7.2. Transition 2. From Allocation Outcomes to Allocation Authority Architectures

Delegation, routing, scheduling, adjustable autonomy, and team design distribute allocation rights in different ways [5,6,52,53,66,85,86]. Their results do not isolate the effect of authority architecture when the allocation object, performer options, or final-decision authority also change. A direct test would hold those features comparable while varying human-controlled, AI-controlled, and genuinely mixed-initiative allocation. Crossing controller conditions with event-level output correctness and, where feasible, reliability history would connect performance and error response to perceived control, intervention, and recovery success. Proposal, approval, contestation, and reallocation rights should be reported explicitly.

7.3. Transition 3. From Performance Optimization to Performance–Agency Alignment

Performance–agency alignment is used here as an agenda label for examining objective performance alongside allocation authority, perceived control or agency, and recovery capacity. Prior empirical studies have examined subsets of these outcomes alongside workload, verification, competence, responsibility, and satisfaction [8–10,48,49,78,79]. Autonomy-preserving design accounts provide complementary conceptual guidance [70]. Future multi-outcome designs can test whether accuracy or speed gains coincide with usable authority and recovery, and when mixed initiative supports participation or adds coordination cost. Expertise, workload, stakes, and error structure may alter these relations.

7.4. Transition 4. From Static Control to Recoverable and Reallocatable Control

Adjustable autonomy, mixed initiative, takeover, and recovery provide precedents in automation, HRI, driving, and robotics [6,72,82,83]. Team-centered and adaptive-allocation work treats flexible authority and reallocation as design requirements [21,22]. In generative knowledge work, recovery may involve textual or conceptual artifacts rather than a discrete transfer of operative control. An upstream error can affect intermediate representations and downstream conclusions. Evaluation should distinguish opportunities to intervene, actual correction behavior, and restoration of an acceptable result. Local editing or rerunning does not by itself demonstrate complete recovery.

A correction to an upstream source interpretation may require identifying affected intermediate outputs, recomputing downstream products, and rechecking the final result. Candidate measures include residual error in downstream artifacts, traceability of corrections, and re-verification effort. Their suitability should be established for the task. These are evaluation requirements, not evidence that a complete recovery chain has been implemented or shown effective across GenAI systems.

Table 4 consolidates the focal contrasts, outcomes, and boundary conditions for these four priorities. Controlled tests of selected dimensions should precede claims about repeated use or transfer across domains.

8. Limitations

The search design supports conceptual coverage rather than exhaustive retrieval or prevalence estimation. Modular searches, citation chaining, benchmark reviews, and targeted retrieval were completed within the stated scope, but relevant work may use different terminology or appear outside the selected sources. Historical screening gaps and capped retrieval records limit reconstruction of the selection pathway; no formal saturation test or field-wide prevalence estimate is inferred from the corpus.

Operational heterogeneity limits direct comparison. Human Factors, HCI, organizational research, decision science, education, automated driving, and HRI define authority, autonomy, shared control, mixed initiative, reliability, override, and recovery differently.

Studies also vary automation level, interface, and information together, while technical, qualitative, conceptual, and experimental designs answer different questions. Concept-driven selection and classification depend on interpretive judgments.

The evidence remains domain-imbalanced. Authority, mixed-initiative, and recovery mechanisms are especially developed in driving, robotics, scheduling, monitoring, and industrial tasks, where control states and failures are observable. Evidence-intensive knowledge work involves less visible state, ambiguous correctness, and intermediate products whose errors may propagate without a discrete takeover event.

Recent GenAI studies contribute timely evidence on writing, qualitative analysis, UX work, and learning, but many rely on small samples, interviews, technical evaluations, or short-term tasks. Their results identify mechanisms, practices, and design concerns rather than stable cross-domain effects.

Screening, extraction, coding, and revision checks were AI-assisted; source-located passages do not establish independent human confirmation. The seven selected comparisons cannot establish coding reproducibility across the core or the frequency of conceptual conflation. S8 traces current citation uses but cannot recover missing historical screening decisions. Preprints are identified by the examined versions.

Evaluation of the five-dimensional reporting structure requires tests of construct discriminability, coding reproducibility, and comparative usefulness across tasks and domains. No statistical independence, causal validity, or improvement in design performance is established here. The temporal synthesis requires observations of how process elements combine during failure and recovery; component examples do not validate a universal sequence.

9. Conclusions

This review organizes established concepts into a common reporting language for human–AI cognitive allocation configurations. Its contribution is to make operations, performers, allocation rights, change rules, and recovery provisions explicit in cross-study comparisons. The synthesis does not establish that existing research generally conflates these attributes.

Framework A is a configuration-oriented integrative reporting framework and design space grounded in Function Allocation and adjacent traditions. It separates configuration from operating conditions, temporal enactment, and outcomes. The multiplication sign denotes joint description rather than statistical interaction, and no optimal configuration is prescribed.

Framework B provides a conditional temporal process lens on enactment, monitoring, recovery, and reallocation, with substantial continuity with joint-control analysis. Together, the frameworks structure questions for interface design and comparative research. Their practical usefulness, reproducibility, and empirical relations among the dimensions require evaluation.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/app162110970/s1, Table S1: Search Modules and Representative Executed Queries; Table S2: Analytic Lens for Cross-Literature Synthesis; Table S3: Boundary Decision Rules for Ambiguous Constructs; Table S4: Source Sets and Selection Boundaries; Table S5: Selected Evidence and Version Boundaries; Table S6: Source Anchors for the Theoretical Comparison; Table S7: Condition-Specific Cross-Study Configuration Comparison; Table S8: Complete Cited-Source Membership and Use Index.

Author Contributions: Conceptualization, C.C. (Chen Chen) and G.L.; methodology, C.C. (Chen Chen) and G.L.; investigation, G.L. and S.Y.; data curation, G.L. and S.Y.; writing—original draft preparation, G.L.; writing—review and editing, C.C. (Chen Chen), G.L., S.Y., Q.X. and C.C. (Chenyang Cai); visualization, G.L. and S.Y.; supervision, C.C. (Chen Chen) and C.C. (Chenyang Cai); project

administration, C.C. (Chen Chen). The authors are accountable for the interpretation and reporting of this work. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study is a literature review and did not involve human participants or animals.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new primary participant data were collected. The cited publications are the evidence sources. Supplementary Tables S1–S3 provide the search modules and proposed analytic boundary rules; Table S4 describes source-set relationships and historical selection limits; Tables S5–S8 provide selected source locators, evidence uses, theoretical mappings, and checking boundaries. The AI-assisted 60 × 5 working codes are not supplied as a validated dataset. The supplement does not contain copyrighted full texts or a complete retrospective screening log.

Acknowledgments: The authors led the research, source selection, interpretation, and coding. ChatGPT and Codex (OpenAI) assisted literature discovery, initial screening, information extraction, and five-dimensional coding, as well as language refinement, structural editing, consistency checking, and conceptual comparison. The exact software and model versions used during manuscript preparation and revision cannot be reliably verified. During revision, Codex connected working interpretations to archived passages and performed targeted source and version checks. These activities are AI-assisted checks, not independent human verification or a reliability assessment. The coding matrix is not presented as having undergone independent human validation of all 300 coding entries. The authors retain responsibility for source selection, interpretation, the disclosed limitations, and the final submission.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Wu, T.; Terry, M.; Cai, C.J. AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts. In *CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2022; pp. 1–22. [CrossRef]
2. Dhillon, P.S.; Molaei, S.; Li, J.; Golub, M.; Zheng, S.; Robert, L.P. Shaping Human-AI Collaboration: Varied Scaffolding Levels in Co-writing with Language Models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2024; pp. 1–18. [CrossRef]
3. Wang, S.; Liu, Y.; Wang, X.; Liu, Z.; You, X.; Li, Y. Where Is the Function Allocation Boundary? The Effect of Degree of Automation on Attention Allocation and Human Performance Under Different Reliabilities. *Hum. Factors J. Hum. Factors Ergon. Soc.* **2025**, *67*, 757–775. [CrossRef]
4. Feigh, K.M.; Pritchett, A.R. Requirements for Effective Function Allocation: A Critical Review. *J. Cogn. Eng. Decis. Mak.* **2014**, *8*, 23–32. [CrossRef]
5. Gombolay, M.; Bair, A.; Huang, C.; Shah, J. Computational design of mixed-initiative human–robot teaming that considers human factors: Situational awareness, workload, and workflow preferences. *Int. J. Robot. Res.* **2017**, *36*, 597–617. [CrossRef]
6. Bhardwaj, A.; Ghasemi, A.H.; Zheng, Y.; Febbo, H.; Jayakumar, P.; Ersal, T.; Stein, J.L.; Gillespie, R.B. Who's the boss? Arbitrating control authority between a human driver and automation system. *Transp. Res. Part F Traffic Psychol. Behav.* **2020**, *68*, 144–160. [CrossRef]
7. Ge, X.; Xu, H.; Zhang, K.; Ni, H.; Xu, J.; Song, X. Effects of Human–Automation Authority Allocation on Multitasking Performance under Workload Conditions. *Int. J. Hum.–Comput. Interact.* **2026**, 1–20. [CrossRef]
8. Tausch, A.; Kluge, A. The best task allocation process is to decide on one's own: Effects of the allocation agent in human–robot interaction on perceived work characteristics and satisfaction. *Cogn. Technol. Work* **2022**, *24*, 39–55. [CrossRef]
9. Berretta, S.; Tausch, A.; Glogowski, P. Maybe adaptive (not adaptable) automation in production: An experimental study comparing the locus of authority in work system dynamics. *Front. Organ. Psychol.* **2025**, *3*, 1685961. [CrossRef]
10. Kares, F.; König, C.J.; Polly, M.; Langer, M. Responsibility attributions in AI-assisted decisions: The interplay of the level of automation and system reliability. *Comput. Hum. Behav. Artif. Hum.* **2026**, *9*, 100336. [CrossRef]
11. Douglas, H.; Di Vincenzo, M.; Dossa, R.F.J.; Nunziante, L.; Sujit, S.; Arulkumaran, K. Levels of shared autonomy in brain-robot interfaces: Enabling multi-robot multi-human collaboration for activities of daily living. *Front. Hum. Neurosci.* **2026**, *19*, 1718713. [CrossRef]

12. Gomez, C.; Cho, S.M.; Ke, S.; Huang, C.M.; Unberath, M. Human-AI collaboration is not very collaborative yet: A taxonomy of interaction patterns in AI-assisted decision making from a systematic review. *Front. Comput. Sci.* **2025**, *6*, 1521066. [CrossRef]
13. Mehrotra, S.; Degachi, C.; Vereschak, O.; Jonker, C.M.; Tielman, M.L. A Systematic Review on Fostering Appropriate Trust in Human-AI Interaction: Trends, Opportunities and Challenges. *ACM J. Responsible Comput.* **2024**, *1*, 1–45. [CrossRef]
14. Richardson, L.S.; Fidock, J.; Gunawan, I. Systematic Literature Review of Levels of Automation (Autonomy) Taxonomy: Critiques and Recommendations. *Int. J. Hum.–Comput. Interact.* **2025**, *41*, 15824–15843. [CrossRef]
15. Kargarnovin, S.; Hernandez, C.I.; Reiners, D.; Cruz-Neira, C.; Bochenek, G.; Karwowski, W. From testbeds to high-stakes work: A review of Human-AI teaming domains and teaming factors. *Front. Robot. AI* **2026**, *13*, 1733942. [CrossRef]
16. Li, H.; Tian, F. Advancing Decision-Making through AI-Human Collaboration: A Systematic Review and Conceptual Framework. *Group Decis. Negot.* **2026**, *35*, 26. [CrossRef]
17. Parasuraman, R.; Sheridan, T.B.; Wickens, C.D. A Model for Types and Levels of Human Interaction with Automation. *IEEE Trans. Syst. Man Cybern.—Part A Syst. Hum.* **2000**, *30*, 286–297. [CrossRef]
18. Feigh, K.M.; Dorneich, M.C.; Hayes, C.C. Toward a Characterization of Adaptive Systems: A Framework for Researchers and System Designers. *Hum. Factors* **2012**, *54*, 1008–1024. [CrossRef]
19. Lundberg, J.; Johansson, B.J.E. A Framework for Describing Interaction between Human Operators and Autonomous, Automated, and Manual Control Systems. *Cogn. Technol. Work* **2021**, *23*, 381–401. [CrossRef]
20. Roth, E.M.; Sushereba, C.; Militello, L.G.; DiIulio, J.; Ernst, K. Function Allocation Considerations in the Era of Human Autonomy Teaming. *J. Cogn. Eng. Decis. Mak.* **2019**, *13*, 199–220. [CrossRef]
21. Abbass, H.A. Social Integration of Artificial Intelligence: Functions, Automation Allocation Logic and Human-Autonomy Trust. *Cogn. Comput.* **2019**, *11*, 159–171. [CrossRef]
22. Hagemann, V.; Rieth, M.; Suresh, A.; Kirchner, F. Human-AI teams—Challenges for a team-centered AI at work. *Front. Artif. Intell.* **2023**, *6*, 1252897. [CrossRef]
23. Järvelä, S.; Nguyen, A.; Hadwin, A. Human and artificial intelligence collaboration for socially shared regulation in learning. *Br. J. Educ. Technol.* **2023**, *54*, 1057–1076. [CrossRef]
24. Edwards, J.; Nguyen, A.; Lämsä, J.; Sobocinski, M.; Whitehead, R.; Dang, B.; Roberts, A.-S.; Järvelä, S. Human- AI collaboration: Designing artificial agents to facilitate socially shared regulation among learners. *Br. J. Educ. Technol.* **2025**, *56*, 712–733. [CrossRef]
25. Dellermann, D.; Ebel, P.; Söllner, M.; Leimeister, J.M. Hybrid Intelligence. *Bus. Inf. Syst. Eng.* **2019**, *61*, 637–643. [CrossRef]
26. Vaccaro, M.; Almaatouq, A.; Malone, T. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nat. Hum. Behav.* **2024**, *8*, 2293–2303. [CrossRef]
27. Bansal, G.; Wu, T.; Zhou, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M.T.; Weld, D. Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2021; pp. 1–16. [CrossRef]
28. Inkpen, K.; Chappidi, S.; Mallari, K.; Nushi, B.; Ramesh, D.; Michelucci, P.; Mandava, V.; Vepřek, L.H.; Quinn, G. Advancing Human-AI Complementarity: The Impact of User Expertise and Algorithmic Tuning on Joint Decision Making. *ACM Trans. Comput.-Hum. Interact.* **2023**, *30*, 1–29. [CrossRef]
29. Zhang, Q.; Lee, M.L.; Carter, S. You Complete Me: Human-AI Teams and Complementary Expertise. In *CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2022; pp. 1–28. [CrossRef]
30. Wilder, B.; Horvitz, E.; Kamar, E. Learning to Complement Humans. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*; IJCAI: Sydney, Australia, 2020; pp. 1526–1533. [CrossRef]
31. Zercher, D.; Jussupow, E.; Benke, I.; Heinzl, A. How Can Teams Benefit From AI Team Members? Exploring the Effect of Generative AI on Decision-Making Processes and Decision Quality in Team–AI Collaboration. *J. Organ. Behav.* **2026**, *47*, 1369–1396. [CrossRef]
32. Siemon, D. Elaborating Team Roles for Artificial Intelligence-based Teammates in Human-AI Collaboration. *Group Decis. Negot.* **2022**, *31*, 871–912. [CrossRef]
33. Flathmann, C.; Nguyen, H.; Andre, K.; Guynup, C. The Teammate Divide: How Humans Approach Teamwork Differently with Human and AI Teammates. *Proc. ACM Hum.-Comput. Interact.* **2026**, *10*, 1–25. [CrossRef]
34. Buçinca, Z.; Malaya, M.B.; Gajos, K.Z. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. *Proc. ACM Hum.-Comput. Interact.* **2021**, *5*, 1–21. [CrossRef]
35. Ma, S.; Lei, Y.; Wang, X.; Zheng, C.; Shi, C.; Yin, M.; Ma, X. Who Should I Trust: AI or Myself? Leveraging Human and AI Correctness Likelihood to Promote Appropriate Trust in AI-Assisted Decision-Making. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2023; pp. 1–19. [CrossRef]
36. Vasconcelos, H.; Jörke, M.; Grunde-McLaughlin, M.; Gerstenberg, T.; Bernstein, M.S.; Krishna, R. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. *Proc. ACM Hum.-Comput. Interact.* **2023**, *7*, 1–38. [CrossRef]
37. Strickland, L.; Farrell, S.; Wilson, M.K.; Hutchinson, J.; Loft, S. How do humans learn about the reliability of automation? *Cogn. Res. Princ. Implic.* **2024**, *9*, 8. [CrossRef]

38. Cabrera, Á.A.; Perer, A.; Hong, J.I. Improving Human-AI Collaboration With Descriptions of AI Behavior. *Proc. ACM Hum.-Comput. Interact.* **2023**, *7*, 1–21. [CrossRef]
39. Fok, R.; Weld, D.S. In search of verifiability: Explanations rarely enable complementary performance in AI-advised decision making. *AI Mag.* **2024**, *45*, 317–332. [CrossRef]
40. Poursabzi-Sangdeh, F.; Goldstein, D.G.; Hofman, J.M.; Wortman Vaughan, J.W.; Wallach, H. Manipulating and Measuring Model Interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2021; pp. 1–52. [CrossRef]
41. Alon-Barkat, S.; Busuioc, M. Human–AI Interactions in Public Sector Decision Making: "Automation Bias" and "Selective Adherence" to Algorithmic Advice. *J. Public Adm. Res. Theory* **2023**, *33*, 153–169. [CrossRef]
42. Cao, S.; Huang, C.M. Understanding User Reliance on AI in Assisted Decision-Making. *Proc. ACM Hum.-Comput. Interact.* **2022**, *6*, 1–23. [CrossRef]
43. Tejeda, H.; Kumar, A.; Smyth, P.; Steyvers, M. AI-Assisted Decision-making: A Cognitive Modeling Approach to Infer Latent Reliance Strategies. *Comput. Brain Behav.* **2022**, *5*, 491–508. [CrossRef]
44. Romeo, G.; Conti, D. Exploring automation bias in human–AI collaboration: A review and implications for explainable AI. *AI Soc.* **2026**, *41*, 259–278. [CrossRef]
45. Küper, A.; Lodde, G.C.; Livingstone, E.; Schadendorf, D.; Krämer, N. Psychological Factors Influencing Appropriate Reliance on AI-enabled Clinical Decision Support Systems: Experimental Web-Based Study Among Dermatologists. *J. Med. Internet Res.* **2025**, *27*, e58660. [CrossRef]
46. Warren, G.; Sun, J.; Shklovski, I.; Augenstein, I. Show me the evidence: Evaluating the role of evidence and natural language explanations in AI-supported fact-checking. *arXiv* **2026**, arXiv:2601.11387v1. [CrossRef]
47. Risko, E.F.; Gilbert, S.J. Cognitive Offloading. *Trends Cogn. Sci.* **2016**, *20*, 676–688. [CrossRef]
48. Fan, Y.; Tang, L.; Le, H.; Shen, K.; Tan, S.; Zhao, Y.; Shen, Y.; Li, X.; Gašević, D. Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *Br. J. Educ. Technol.* **2025**, *56*, 489–530. [CrossRef]
49. Zhu, Q.; Li, X.; Dong, Y.; Chang, P.; Fan, M. Not all cognitive offloading is equal: Distinguishing dependent and autonomous offloading to generative AI. *Front. Psychol.* **2026**, *17*, 1878629. [CrossRef]
50. Markauskaite, L.; Marrone, R.; Poquet, O.; Knight, S.; Martinez-Maldonado, R.; Howard, S.; Tondeur, J.; De Laat, M.; Buckingham Shum, S.; Gašević, D.; et al. Rethinking the entwinement between artificial intelligence and human learning: What capabilities do learners need for a world with AI? *Comput. Educ. Artif. Intell.* **2022**, *3*, 100056. [CrossRef]
51. Spitzer, P.; Holstein, J.; Hemmer, P.; Vössing, M.; Kühl, N.; Martin, D.; Satzger, G. Human Delegation Behavior in Human-AI Collaboration: The Effect of Contextual Information. *Proc. ACM Hum.-Comput. Interact.* **2025**, *9*, 1–28. [CrossRef]
52. Madras, D.; Pitassi, T.; Zemel, R. Predict Responsibly: Improving Fairness and Accuracy by Learning to Defer. In *Advances in Neural Information Processing Systems 31*; Neural Information Processing Systems Foundation: San Diego, CA, USA, 2018; Volume 31. Available online: https://proceedings.neurips.cc/paper/2018/hash/09d37c08f7b129e96277388757530c72-Abstract.html (accessed on 1 November 2026).
53. Raghu, M.; Blumer, K.; Sayres, R.; Obermeyer, Z.; Kleinberg, B.; Mullainathan, S.; Kleinberg, J. Direct Uncertainty Prediction for Medical Second Opinions. In *Proceedings of the 36th International Conference on Machine Learning (PMLR 97)*; ML Research Press: Cambridge, MA, USA, 2019; Volume 97, pp. 5281–5290. Available online: https://proceedings.mlr.press/v97/raghu19a.html (accessed on 1 November 2026).
54. Dai, S.C.; Xiong, A.; Ku, L.W. LLM-in-the-loop: Leveraging Large Language Model for Thematic Analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2023*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2023; pp. 9993–10001. [CrossRef]
55. De Paoli, S. Performing an Inductive Thematic Analysis of Semi-Structured Interviews With a Large Language Model: An Exploration and Provocation on the Limits of the Approach. *Soc. Sci. Comput. Rev.* **2024**, *42*, 997–1019. [CrossRef]
56. Parfenova, A.; Marfurt, A.; Pfeffer, J.; Denzler, A. Text Annotation via Inductive Coding: Comparing Human Experts to LLMs in Qualitative Data Analysis. In *Findings of the Association for Computational Linguistics: NAACL 2025*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2025; pp. 6456–6469. [CrossRef]
57. Xiao, Z.; Yuan, X.; Liao, Q.V.; Abdelghani, R.; Oudeyer, P.Y. Supporting Qualitative Analysis with Large Language Models: Combining Codebook with GPT-3 for Deductive Coding. In *28th International Conference on Intelligent User Interfaces*; Association for Computing Machinery: New York, NY, USA, 2023; pp. 75–78. [CrossRef]
58. Yan, L.; Echeverria, V.; Fernandez-Nieto, G.M.; Jin, Y.; Swiecki, Z.; Zhao, L.; Gašević, D.; Martinez-Maldonado, R. Human-AI Collaboration in Thematic Analysis using ChatGPT: A User Study and Design Recommendations. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2024; pp. 1–7. [CrossRef]

59. Zhao, C.; Tan, Z.; Wong, C.W.; Zhao, X.; Chen, T.; Liu, H. SCALE: Towards Collaborative Content Analysis in Social Science with Large Language Model Agents and Human Intervention. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2025; pp. 8473–8503. [CrossRef]
60. Sharma, A.; Cochrane, K.; Wallace, J.R. DeTAILS: Deep Thematic Analysis with Iterative LLM Support. *arXiv* **2025**, arXiv:2510.17575v2. [CrossRef]
61. Zhong, M.; Wang, P.; Field, A. HICode: Hierarchical Inductive Coding with LLMs. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2025; pp. 31060–31078. [CrossRef]
62. Jiang, J.A.; Wade, K.; Fiesler, C.; Brubaker, J.R. Supporting Serendipity: Opportunities and Challenges for Human-AI Collaboration in Qualitative Analysis. *Proc. ACM Hum.-Comput. Interact.* **2021**, *5*, 1–23. [CrossRef]
63. Lee, M.; Liang, P.; Yang, Q. CoAuthor: Designing a Human-AI Collaborative Writing Dataset for Exploring Language Model Capabilities. In *CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery: New York, NY, USA, 2022; pp. 1–19. [CrossRef]
64. Reza, M.; Thomas-Mitchell, J.; Dushniku, P.; Laundry, N.; Williams, J.J.; Kuzminykh, A. Co-Writing with AI, on Human Terms: Aligning Research with User Demands Across the Writing Process. *Proc. ACM Hum.-Comput. Interact.* **2025**, *9*, 1–37. [CrossRef]
65. Takaffoli, M.; Li, S.; Mäkelä, V. Generative AI in User Experience Design and Research: How Do UX Practitioners, Teams, and Companies Use GenAI in Industry? In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*; Association for Computing Machinery: New York, NY, USA, 2024; pp. 1579–1593. [CrossRef]
66. Wen, Y.; Wang, J.; Chen, X. Trust and AI weight: Human-AI collaboration in organizational management decision-making. *Front. Organ. Psychol.* **2025**, *3*, 1419403. [CrossRef]
67. Formosa, P. Robot Autonomy vs. Human Autonomy: Social Robots, Artificial Intelligence (AI), and the Nature of Autonomy. *Minds Mach.* **2021**, *31*, 595–616. [CrossRef]
68. Karakikes, M.; Nathanael, D. The effect of cognitive workload on decision authority assignment in human–robot collaboration. *Cogn. Technol. Work* **2023**, *25*, 31–43. [CrossRef]
69. Tsamados, A.; Floridi, L.; Taddeo, M. Human control of AI systems: From supervision to teaming. *AI Ethics* **2025**, *5*, 1535–1548. [CrossRef]
70. Buijsman, S.; Carter, S.E.; Bermúdez, J.-P. Autonomy by Design: Preserving Human Autonomy in AI Decision-Support. *Philos. Technol.* **2025**, *38*, 97. [CrossRef]
71. Henry, K.E.; Kornfield, R.; Sridharan, A.; Linton, R.C.; Groh, C.; Wang, T.; Wu, A.; Mutlu, B.; Saria, S. Human–machine teaming is key to AI adoption: Clinicians' experiences with a deployed machine learning system. *npj Digit. Med.* **2022**, *5*, 97. [CrossRef]
72. Chiou, M.; McCabe, F.; Grigoriou, M.; Stolkin, R. Trust, Shared Understanding and Locus of Control in Mixed-Initiative Robotic Systems. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*; IEEE: New York, NY, USA, 2021; pp. 684–691. [CrossRef]
73. Yu, A.; Li, C.; Macesanu, L.; Balaji, A.; Ray, R.; Mooney, R.; Martín-Martín, R. Mixed-Initiative Dialog for Human-Robot Collaborative Manipulation. *arXiv* **2026**, arXiv:2508.05535v2. [CrossRef]
74. Diebel, C.; Goutier, M.; Adam, M.; Benlian, A. When AI-Based Agents Are Proactive: Implications for Competence and System Satisfaction in Human–AI Collaboration. *Bus. Inf. Syst. Eng.* **2026**, *68*, 259–278. [CrossRef]
75. Flathmann, C.; Schelble, B.G.; Rosopa, P.J.; McNeese, N.J.; Mallick, R.; Madathil, K.C. Examining the impact of varying levels of AI teammate influence on human-AI teams. *Int. J. Hum.-Comput. Stud.* **2023**, *177*, 103061. [CrossRef]
76. Zhang, R.; Duan, W.; Flathmann, C.; McNeese, N.; Freeman, G.; Williams, A. Investigating AI Teammate Communication Strategies and Their Impact in Human-AI Teams for Effective Teamwork. *Proc. ACM Hum.-Comput. Interact.* **2023**, *7*, 1–31. [CrossRef]
77. Cavalcante Siebert, L.; Lupetti, M.L.; Aizenberg, E.; Beckers, N.; Zgonnikov, A.; Veluwenkamp, H.; Abbink, D.; Giaccardi, E.; Houben, G.J.; Jonker, C.M.; et al. Meaningful human control: Actionable properties for AI system development. *AI Ethics* **2023**, *3*, 241–255. [CrossRef]
78. Van den Bosch, K.; Van Diggelen, J.; Verdult, S.; Haije, T.; Van der Waa, J. Measuring meaningful human control in human–AI teaming: Effects of team design in AI-assisted pandemic triage. *AI Ethics* **2025**, *5*, 3329–3353. [CrossRef]
79. Darvishi, A.; Khosravi, H.; Sadiq, S.; Gašević, D.; Siemens, G. Impact of AI assistance on student agency. *Comput. Educ.* **2024**, *210*, 104967. [CrossRef]
80. Satyanarayan, A.; Jones, G.M. Intelligence as Agency: Evaluating the Capacity of Generative AI to Empower or Constrain Human Action. In *An MIT Exploration of Generative AI*; Massachusetts Institute of Technology: Cambridge, MA, USA, 2024. [CrossRef]
81. Srivastava, D.; Lilly, J.M.; Feigh, K.M. Exploring the role of judgement and shared situation awareness when working with AI recommender systems. *Cogn. Technol. Work* **2025**, *27*, 1–18. [CrossRef]
82. Vogelpohl, T.; Gehlmann, F.; Vollrath, M. Task Interruption and Control Recovery Strategies After Take-Over Requests Emphasize Need for Measures of Situation Awareness. *Hum. Factors J. Hum. Factors Ergon. Soc.* **2020**, *62*, 1190–1211. [CrossRef]

83. Kostolani, D.; Brenter, B.; Schlund, S. Human Takeover: Overriding Task Assignments Mid-Execution in Human-Robot Interaction. In *Annals of Scientific Society for Assembly, Handling and Industrial Robotics 2024*; Springer: Cham, Switzerland, 2026; pp. 51–61. [CrossRef]
84. Pritchett, A.R.; Kim, S.Y.; Feigh, K.M. Measuring Human-Automation Function Allocation. *J. Cogn. Eng. Decis. Mak.* **2014**, *8*, 52–77. [CrossRef]
85. Fügener, A.; Grahl, J.; Gupta, A.; Ketter, W. Cognitive Challenges in Human–Artificial Intelligence Collaboration: Investigating the Path Toward Productive Delegation. *Inf. Syst. Res.* **2022**, *33*, 678–696. [CrossRef]
86. Hemmer, P.; Westphal, M.; Schemmer, M.; Vetter, S.; Vössing, M.; Satzger, G. Human-AI Collaboration: The Effect of AI Delegation on Human Task Performance and Task Satisfaction. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*; Association for Computing Machinery: New York, NY, USA, 2023; pp. 453–463. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.