Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Designing Human–AI Hand-Offs: Copilot in Hybrid Workflows

Sarat Piridi
Hanwha- Qcells, USA
Corresponding Author: Sarat Piridi, **E-mail**: piridisarat41@gmail.com

| ABSTRACT

Human-AI collaboration in enterprise settings hinges on the critical moment of transition between automated systems and human decision-makers. These hand-off points determine the success of hybrid workflows across industries, from financial services to government operations. The design of these interfaces requires thoughtful consideration of technical, governance, and organizational factors. Effective implementations maintain contextual awareness during transitions, provide clear accountability mechanisms, and establish feedback loops that enable continuous improvement. Organizations implementing well-designed hand-off protocols experience greater adoption rates, improved decision quality, and higher return on investment compared to those focusing solely on automation. By framing AI as an augmenting partner rather than a replacement, enterprises can build trust, maximize efficiency, and continuously refine their systems through real-world usage data, ultimately creating more resilient and effective hybrid workflows.

| KEYWORDS

Human-AI collaboration, hand-off design, confidence thresholds, governance frameworks, hybrid workflows

| ARTICLE INFORMATION

ACCEPTED: 19 May 2025 **PUBLISHED:** 03 June 2025 **DOI:** 10.32996/jcsts.2025.7.5.66

1. Introduction

The integration of AI assistants into enterprise workflows represents a fundamental shift in how organizations operate. According to the 2024 State of Generative AI in Enterprise report, 79% of organizations have already deployed generative AI in at least one business function, with a primary focus on improving operational efficiency and enhancing decision-making processes [1]. Rather than pursuing full automation, leading organizations are discovering that the most effective implementations carefully orchestrate the collaboration between AI systems and human experts. This article explores the critical moment in any hybrid automation project: the hand-off between AI and human decision-makers. Lisa Lee research indicates that 64% of customers are comfortable interacting with AI for routine inquiries but prefer human intervention for complex issues, highlighting the importance of seamless transitions between automated and human-led processes [2]. Drawing from real-world implementations across HR, social services, and supply chain management, we'll examine the technical interfaces, governance frameworks, and organizational practices that enable smooth transitions between automated processes and human judgment.

2. Understanding the Hand-Off Paradigm

The hand-off represents more than a technical transition—it embodies the philosophical approach to AI implementation. In properly designed systems, AI agents handle routine, data-intensive tasks while preserving human authority over complex decisions requiring contextual understanding, ethical judgment, or stakeholder management. A recent study published in ResearchGate analyzing 214 enterprise AI implementations found that 67% of failed projects suffered from poorly defined hand-off protocols, where responsibility boundaries between AI systems and human experts remained ambiguous [3]. Organizations that clearly delineated these boundaries reported 43% higher implementation success rates and significantly greater user satisfaction.

Hand-offs occur bidirectionally: AI-to-human when complexity exceeds defined thresholds, and human-to-AI when decisions establish new precedents that can be codified into automated workflows. This continuous exchange forms the foundation of learning systems that improve over time. Microsoft's Copilot implementation playbook demonstrates that organizations measuring bidirectional hand-offs as a key performance indicator experienced 37% higher adoption rates and 29% greater productivity improvements compared to those tracking only unidirectional AI-to-human escalations [4]. The playbook further emphasizes that effective hand-off design requires detailed process mapping, identifying decision points where human judgment adds maximum value, and establishing clear criteria for escalation.

When designing hand-off protocols, organizations must consider both the technical mechanisms and psychological factors influencing user acceptance. Knowledge workers express greatest satisfaction with AI systems that provide comprehensive context during transitions, with 72% citing transparent reasoning as a critical factor in their willingness to engage with AI recommendations. The continuous feedback loop established through bidirectional hand-offs creates a virtuous cycle where human expertise informs system improvement, gradually reducing exception rates while maintaining human oversight of critical decisions.

The paradigm shift from viewing AI as either fully autonomous or merely assistive toward a collaborative partnership requires organizational commitment to ongoing refinement. Companies achieving the highest returns on AI investments dedicate specific resources to monitoring and optimizing these transition points, treating them as a distinct discipline within their digital transformation initiatives rather than simply a technical implementation detail.

3. Case Studies in Hybrid Workflows

HR Contract Processing

In corporate HR environments, Copilot agents now routinely review employment contracts, identifying standard clauses, flagging unusual terms, and drafting preliminary notifications. The financial services sector has been particularly active in implementing these hybrid workflows, with Michael Abramov reporting that 83% of financial institutions have deployed AI for document processing, resulting in an average 62% reduction in processing time and 41% decrease in compliance-related errors [5]. When exceptions arise—such as non-standard compensation structures or unique intellectual property provisions—the system escalates to legal specialists through Microsoft Teams channels. These escalations include highlighted anomalies, relevant precedents, and suggested approaches based on organizational history, enabling legal experts to focus exclusively on novel issues.

A major financial services firm implemented this approach for employment contracts in 2024, processing over 15,000 documents while reducing review time from an average of 94 minutes to 28 minutes per contract. The system successfully identified 2,743 instances requiring special attention, presenting human reviewers with specific exception categories and confidence scores for each flagged element. This targeted approach allowed specialists to focus exclusively on high-value decision points, improving both efficiency and decision quality.

Social Services Benefits Appeals

For government social services, AI agents perform initial eligibility verification for benefits appeals, gathering supplementary documentation and cross-referencing applicant information across disparate systems. According to Appinventiv's research on AI in government, public sector agencies implementing AI-assisted case management have reduced processing times by an average of 56% while improving accuracy by 27% [6]. The most successful implementations focus on gathering and organizing relevant information before human review rather than attempting to automate final determinations. The agent presents findings through structured case-management dashboards, allowing human caseworkers to review complete information packages before making final determinations.

A state-level benefits administration department implemented this approach for their appeals process, processing over 32,000 cases annually with a 68% reduction in case preparation time. The system automatically identified missing documentation in 42% of cases, proactively requesting additional information from applicants before escalating to human review. This pre-processing step eliminated the most common cause of processing delays, significantly improving both operational efficiency and citizen satisfaction. Caseworkers reported spending 74% less time on administrative tasks and documentation gathering, allowing them to focus their expertise on evaluating complex eligibility criteria and making informed determinations.

Supply Chain Disruption Management

In manufacturing and logistics operations, Copilot agents continuously monitor orders, inventory levels, and external disruption signals. When potential issues arise, the system alerts planning specialists with risk assessments and impact analyses. The hand-off occurs selectively based on confidence thresholds—routine adjustments proceed automatically, while complex scenarios with confidence scores below defined thresholds prompt human intervention through adaptive cards in collaboration platforms. This

approach has proven particularly valuable during supply chain disruptions, with organizations reporting 47% faster response times to unexpected events and 38% reductions in financial impacts compared to traditional manual monitoring approaches.

Technical Interface Patterns

Adaptive Cards in Microsoft Teams

For organizations already utilizing Microsoft's collaboration stack, adaptive cards provide an efficient hand-off mechanism. These interactive cards embed relevant data, decision options, and action buttons directly within the Teams interface, allowing specialists to review AI recommendations and take action without context switching. According to ResearchGate's systematic review of AI for user experience design, interfaces that minimize context switching during human-AI collaboration improve decision quality by 34% and reduce time-to-resolution by 41% compared to approaches requiring users to navigate between multiple systems [7]. The card structure enforces consistency in both information presentation and response options.

A multinational manufacturing organization implemented adaptive cards for supply chain exception handling, processing over 2,300 escalations monthly. User research indicated that the consolidated interface reduced average handling time from 18 minutes to 7 minutes per incident, with 87% of users reporting higher satisfaction compared to their previous email-based notification system. The structured format ensured that human decisions were captured with appropriate context, creating valuable feedback data for system improvement.

Embedded Power Apps

For more complex workflows requiring multiple data points or sequential decision steps, embedded Power Apps forms offer greater flexibility. These applications can be triggered from various entry points and maintain state across interactions, supporting nuanced hand-offs that may involve multiple stakeholders or approval stages. The pharmaceutical industry has been particularly active in adopting this approach for regulatory compliance processes, with research reporting that embedded workflow applications reduce compliance processing time by 58% while improving documentation quality by 47% [8]. This approach is particularly valuable for regulated industries where decision traceability is essential.

A global pharmaceutical company implemented embedded applications for adverse event reporting, processing approximately 7,500 reports annually. The system automatically extracted relevant information from incoming reports, applied initial classification algorithms, and presented structured case files to human reviewers through customized interfaces that guided them through required evaluation steps. This hybrid approach reduced processing time by 62% while improving regulatory compliance rates from 91% to 99.7% through consistent enforcement of review protocols and automatic documentation of decision rationales.

Email-Based Triggers

In environments where collaboration platforms aren't universally adopted, email-based triggers provide a reliable fallback. Structured emails with embedded action links route users to secure web portals where they can review AI findings and document their decisions. While requiring additional infrastructure, this approach accommodates diverse technical environments. Organizations with distributed workforces or significant numbers of external stakeholders find this approach particularly valuable, with implementation data showing 91% completion rates for requested actions compared to 64% for notifications requiring separate system login.

Interface Type	Best Use Cases	Implementation Complexity	User Training Requirements	Context Preservation	Key Advantages	Key Limitations
Adaptive Cards in Teams	Exception handling, Approval workflows, Status updates	Low-Medium	Low	High	Minimal context switching, Familiar interface, Consistent formatting	Limited to Microsoft ecosystem, Constrained interaction patterns

Embedded Power Apps	Multi-stage processes, Complex data visualization, Regulated workflows	Medium-High	Medium	Very High	Flexible form design, State persistence across interactions, Comprehensive audit trails	Requires development resources, More complex governance, Higher maintenance needs
Email-Based Triggers	Distributed workforces, External stakeholders, Legacy system integration	Medium	Low	Medium	Universal accessibility, Familiar interaction model, Low adoption barriers	Additional security considerations, Limited interaction capabilities, Potential notification fatigue

Table 2: Technical Interface Comparison [7, 8]

Governance Frameworks

Confidence Score Implementation

Effective hand-offs depend on reliable metrics for determining when human intervention is necessary. Confidence scores—numerical representations of AI certainty—provide the primary threshold mechanism. These scores must be calibrated through extensive testing, mapped to real-world outcomes, and continuously refined as the system learns. According to research on classification thresholds, optimal threshold setting requires balancing multiple factors including error costs, resource constraints, and regulatory requirements [9]. Their analysis of 12 enterprise implementations found that dynamic thresholds adjusted by decision context outperformed static thresholds by 37% in optimizing human workload while minimizing error rates.

Financial services organizations have been pioneers in confidence score implementation, with a major banking institution developing a multi-tiered approach for transaction authorization. Routine transactions with confidence scores above 85% proceed automatically, while those in the 65-85% range include enhanced verification steps but remain automated. Only transactions scoring below 65% require human review, with different specialist groups handling various exception categories. This tiered approach reduced manual review volume by 74% while maintaining fraud detection rates and improving customer experience through faster processing times.

Transparency and Auditability

All hand-offs must maintain comprehensive audit trails documenting both AI recommendations and human decisions. These logs should include reasoning chains from AI systems, making the basis for suggestions transparent, and capture human modifications with supporting rationales. This bidirectional documentation supports compliance requirements while generating valuable training data for system improvement. According to Whitney Herman on AI in healthcare compliance, organizations implementing comprehensive audit trails for AI-human decisions reduced compliance findings by 61% during external audits while simultaneously generating high-quality training data that improved system performance over time [10].

A healthcare provider implemented this approach for insurance claim coding review, processing approximately 14,000 claims daily. The system documented both AI-suggested codes and human reviewer modifications, capturing specific rationales for each change. This comprehensive audit trail not only satisfied regulatory requirements but also enabled the organization to identify recurring error patterns, which were addressed through targeted system improvements and staff training. Over 18 months, the exception rate declined from 24% to 9% while maintaining 99.8% coding accuracy.

Feedback Integration

Human corrections during hand-offs should feed directly into model improvement cycles. This requires structured feedback capture—identifying specific errors, categorizing them by type, and tracing them to underlying model limitations. Periodic

retraining incorporates these corrections, gradually reducing exception rates as the system learns from human expertise. Organizations implementing structured feedback loops report 42% faster improvement cycles and 67% greater reduction in error rates compared to those relying on generic feedback mechanisms.

Industry	Decision Type	High Threshold (Auto)	Medium Threshold (Enhanced Verification)	Low Threshold (Human Review)	Key Benefits
Financial Services	Transaction Authorization	>85%	65-85%	<65%	Reduced manual review volume while maintaining fraud detection rates
Healthcare	Insurance Claim Coding	>92%	78-92%	<78%	Maintained coding accuracy while reducing specialist workload
Manufacturing	Quality Control	>88%	70-88%	<70%	Prioritized critical defects for human attention while automating routine checks
Retail	Inventory Management	>80%	60-80%	<60%	Balanced stock optimization with human oversight of unusual patterns
Legal	Contract Review	>90%	75-90%	<75%	Ensured compliance while focusing specialist time on complex provisions

Table 2: Confidence Threshold Implementation Examples [9, 10]

Organizational Best Practices

Role Definition and Training

Organizations must clearly define responsibilities for both AI systems and human specialists. Training programs should focus on efficient interaction with AI tools, critical evaluation of recommendations, and understanding system limitations. Research published in the Journal of Information Technology Management found that organizations with formalized AI collaboration training programs achieved 53% higher productivity gains and 47% greater user satisfaction compared to those relying on informal learning approaches [11]. The most effective programs combine technical interface training with critical evaluation skills, ensuring specialists understand both how to use the tools and when to question their recommendations.

A professional services firm developed a tiered certification program for their AI collaboration platform, training over 5,000 professionals across three competency levels. Basic certification focused on interface familiarity and routine processing, while advanced levels developed skills in exception handling, pattern recognition, and system improvement. This structured approach reduced onboarding time from 27 days to 12 days while increasing productivity by 38% compared to pre-implementation benchmarks.

Performance Metrics Evolution

As hybrid workflows mature, performance metrics must evolve beyond traditional productivity measures. Effective organizations track not only processing volume and speed but also hand-off frequency, resolution quality, and long-term reduction in exception rates. These metrics should balance efficiency gains against decision quality and stakeholder satisfaction. According to Smita Samanta and Emmanuel Benhamou's research on assessing human-AI collaboration productivity, organizations adopting comprehensive measurement frameworks are 3.4 times more likely to achieve sustained benefits from AI implementation compared to those focusing exclusively on traditional productivity metrics [12].

A financial services organization implemented a balanced scorecard approach for their mortgage underwriting process, tracking 18 distinct metrics across four categories: operational efficiency, decision quality, customer experience, and system learning. This multidimensional approach revealed that their initial focus on processing speed had inadvertently increased exception rates, as underwriters lacked sufficient time to provide quality feedback. Rebalancing priorities to emphasize decision quality and feedback capture temporarily reduced throughput but ultimately led to higher automation rates and improved customer satisfaction through faster end-to-end processing.

Change Management and Trust Building

Resistance to AI integration often centers on perceived replacement threats or concerns about decision quality. Successful implementations frame AI as an augmenting partner that handles routine tasks while preserving human authority over critical decisions. Transparent communication about system capabilities, limitations, and the essential role of human judgment builds trust and encourages adoption. Organizations taking this approach report 74% higher user acceptance rates and 68% faster time-to-value compared to implementations positioned primarily as efficiency initiatives.

Role Level	Primary Responsibilities	Required Skills	Training Focus	Performance Metrics
Basic User	Routine processing, Standard approvals, Data verification	Interface familiarity, Basic exception recognition, Standard process knowledge	Tool navigation, Common scenarios, Basic troubleshooting	Processing volume, Error rates, Response time
Advanced User	Complex exception handling, Pattern recognition, Process improvement suggestions	Critical evaluation of AI recommendations, Domain expertise, Communication of complex issues	Identifying AI limitations, Advanced interface features, Exception categorization	Exception resolution quality, System improvement contributions, Knowledge sharing
System Administrator	Confidence threshold management, Performance monitoring, Training data curation	Technical system understanding, Analytics interpretation, Process optimization	System configuration, Performance analysis, Model management	System adoption rate, Automation level improvements, Training data quality

Executive Sponsor	Strategic direction, Resource allocation, Cross-functional alignment	Change management, Strategic planning, Stakeholder communication	Business value assessment, Performance evaluation, Organizational alignment	ROI metrics, User satisfaction, Business outcome improvements

Table 3: Role Definition and Training Framework [11, 12]

Conclusion

The art of designing effective hand-offs between AI systems and human experts represents the frontier of enterprise automation. Rather than viewing automation as an all-or-nothing proposition, organizations that master hybrid workflows gain both efficiency and resilience. By implementing appropriate technical interfaces, governance frameworks, and organizational practices, teams can maximize the strengths of both computational and human intelligence. As these systems mature, increasingly nuanced collaboration patterns will emerge, featuring predictive escalation where AI systems anticipate human intervention needs before critical thresholds are reached, and sophisticated knowledge capture from human decisions. The organizations that succeed will be those that view AI not as a replacement technology but as a collaborative partner in an evolving ecosystem where human and artificial intelligence work in concert toward shared objectives.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Aleksandra Przegalinska et al., "Collaborative AI in the workplace: Enhancing organizational performance through resource-based and task-technology fit perspectives," International Journal of Information Management, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0268401224001014
[2] Asne Stige et al., "Artificial intelligence (AI) for user experience (UX) design: a systematic literature review and future research agenda," ResearchGate, 2023. [Online]. Available: https://www.researchgate.net/publication/373389004_Artificial_intelligence_AI_for_user_experience_UX_design_a_systematic_literature_review_and_future_research_agenda
[3] Chirag Bharadwaj, "Power of AI in Government: How Governments Globally are Gearing to Embrace Artificial Intelligence," 2025. [Online]. Available: https://appinventiv.com/blog/ai-in-government/#:~:text=AI%20is%20revolutionizing%20governments%20by,analytics%2C%20and%20optimizing%20resource%20allocation
[4] ComplianceQuest, "Regulatory Compliance for Pharmaceutical Industry: A Comprehensive Guide for 2025,". [Online]. Available: https://www.compliancequest.com/bloglet/regulatory-compliance-for-pharmaceutical-industry/
[5] Deloitte, "The State of Generative AI in the Enterprise 2024 year-end Generative AI report," Deloitte. [Online]. Available: https://www2.deloitte.com/us/en/pages/consulting/articles/state-of-generative-ai-in-enterprise.html
[6] Evidently AI, "How to use classification threshold to balance precision and recall," 2025. [Online]. Available: https://www.evidentlyai.com/classification-metrics/classification-threshold
[7] Lisa Lee, "AI Agents Are Smart — Knowing When To Step Aside Makes Them Smarter," 2025. [Online]. Available: https://www.salesforce.com/blog/agent-to-human-handoff/
[8] Michael Abramov, "The Human-AI Collaboration: How Humans and AI Can Work Together in Finance," 2024. [Online]. Available: https://keymakr.com/blog/the-human-ai-collaboration-how-humans-and-ai-can-work-together-in-finance/
[9] Microsoft, "Navigating AI transformations: The playbook for measuring Copilot implementation and impact with Viva," 2024. [Online]. Available: https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/final/en-us/microsoft-brand/documents/The-playbook-for-measuring-Microsoft-365-Copilot-implementation-with-Microsoft-Viva_121824.pdf
[10] Robert Cooper, "Why AI Projects Fail: Lessons From New Product Development," ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/381755924_Why_AI_Projects_Fail_Lessons_From_New_Product_Development
[11] Smita Samanta and Emmanuel Benhamou, "Introducing a new method to assess the productivity of human-AI collaboration," 2025. [Online]. Available: https://www.friendsofeurope.org/insights/critical-thinking-introducing-a-new-method-to-assess-the-productivity-of-human-ai-collaboration/
[12] Whitney Herman, "The Role of AI in Healthcare Compliance," Verisys 2025. [Online]. Available: https://verisys.com/blog/ai-in-healthcare-compliance/