Beyond Automation: Embracing Symbiotic AI for Meaningful Human Agency and Skill Preservation

GIUSEPPE DESOLDA, ANDREA ESPOSITO, ROSA LANZILOTTI, ANTONIO PICCINNO, and MARIA FRANCESCA COSTABILE, University of Bari Aldo Moro, Italy

As Artificial Intelligence systems evolve into proactive, multi-agent entities, the traditional paradigm of automation experience faces unprecedented challenges regarding transparency, human agency, and long-term skill degradation. This position paper argues that moving from general Human-Centered AI (HCAI) toward Symbiotic AI (SAI) is essential for designing effective and safe agentic systems. SAI has been recently proposed as a sub-field of HCAI: it addresses systems letting humans and AI collaborate and co-evolve through mutual learning and continuous interaction. By synthesizing findings from a prior comprehensive survey and two medical case studies (on rhinocytology and on brain oncology), in this paper we discuss how SAI addresses the core themes of AutomationXP26: (1) enhancing multi-agent transparency through explainable interfaces, (2) ensuring humans remain 'in control' even in complex automated workflows, and (3) sustaining human skills by fostering a collaborative environment where AI acts as a partner rather than a black-box oracle. In doing so, this paper positions SAI as a design paradigm for agentic automation, offering workshop participants actionable concepts, design tensions, and empirical case studies to reason about how agentic automation can preserve human agency, transparency, and skills in real-world settings.

CCS Concepts: • **Human-centered computing** → **HCI theory, concepts and models**; • **Computing methodologies** → *Artificial intelligence*; • **General and reference** → Surveys and overviews.

Additional Key Words and Phrases: Symbiotic AI, Human-Centred AI, Preserving Human Agency, Preserving Human Skills, Humans in Control

ACM Reference Format:
Giuseppe Desolda, Andrea Esposito, Rosa Lanzilotti, Antonio Piccinno, and Maria Francesca Costabile. 2026. Beyond Automation: Embracing Symbiotic AI for Meaningful Human Agency and Skill Preservation. In *Proceedings of AutomationXP26 Workshop of the 2026 CHI Conference on Human Factors in Computing Systems, April 14, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 6 pages.

1 Introduction

The rapid proliferation of Artificial Intelligence (AI) is transforming the landscape of automation, shifting from static support tools to proactive "agentic" systems. Unlike traditional tools, these systems coordinate multiple specialized agents to proactively solve complex problems with minimal human intervention [14]. Consider, for instance, a medical support system where distinct AI agents independently segment an MRI (Magnetic Resonance Imaging), retrieve patient history, and cross-reference clinical guidelines, all collaborating behind the scenes to propose a synthesis. The opacity of such multi-agent coordination threatens to obscure the attribution of outcomes, undermining human trust. In this

Authors' Contact Information: Giuseppe Desolda, giuseppe.desolda@uniba.it; Andrea Esposito, andrea.esposito@uniba.it; Rosa Lanzilotti, rosa.lanzilotti@uniba.it; Antonio Piccinno, antonio.piccinno@uniba.it; Maria Francesca Costabile, mariafrancesca.costabile@gmail.com, Department of Computer Science, University of Bari Aldo Moro, Bari, Italy.

context, we refer to *human agency* not merely as the ability to initiate a task, but as the user's capacity to meaningfully understand, intervene in, and steer the system's actions, preventing them from being reduced to passive bystanders.

To address these concerns, the field of Human-Centered AI (HCAI) has emerged at the intersection of Human-Computer Interaction (HCI) and AI, aiming to design systems that augment rather than replace human expertise [13]. HCAI promotes the creation of ethically aligned, reliable, safe, and trustworthy systems that reflect human factors [5, 12]. However, our research suggests that for complex, high-autonomy scenarios, standard HCAI is a necessary but insufficient foundation. A key finding emerging from our comprehensive survey is the need to move toward *Symbiotic AI (SAI)* [1, 5, 8]. SAI systems are defined as systems that permit a "symbiosis" of human intelligence and artificial intelligence, where both parties augment each other's capabilities through a collaboration that balances their respective strengths and weaknesses [5]. Thus, AI systems function as *cognitive orthotics* that support and extend existing human expertise, rather than *cognitive prostheses* that merely substitute human abilities and increase risk-inducing passivity [5, 8].

This position paper leverages a prior review over the landscape of SAI and (more generally) HCAI [5], to highlight how the SAI paradigm addresses the core themes of the *AutomationXP26* workshop [14]:

- *Enabling Transparency and Attribution:* We discuss how SAI moves beyond static explanations to interactive interpretability, essential for understanding multi-agent contributions.
- *Balancing Human Agency:* We propose intervention mechanisms that ensure humans remain "in control," transforming passive oversight into active collaboration [5].
- *Sustaining Human Skills:* We demonstrate, through medical case studies, how symbiotic interaction prevents "de-skilling" by fostering continuous mutual learning.

More specifically, this paper contributes to this workshop by articulating a shared vocabulary and set of design concerns to help the community collectively reason about agency, transparency, and skill preservation in agentic automation experiences. This reasoning is instrumental in proposing finalized solutions.

2 The Symbiotic AI Paradigm

The conceptual foundation of our proposal lies in the evolution from HCAI to the more specialized and integrated paradigm of SAI. As identified in our recent multidisciplinary survey [5], HCAI is not a monolithic field but a convergence of several disciplines—namely HCI, AI (including Machine Learning and eXplainable AI), Software Engineering, and Ethics. While HCAI focuses on creating systems that are reliable, safe, and trustworthy by involving users in the design and evaluation process, the emergence of agentic automation requires a deeper level of integration.

2.1 Symbiotic AI Systems

Symbiotic AI systems are HCAI systems powered by a continuing and deeper collaboration between humans and AI that mutually augment their capabilities—i.e., such systems allow a kind of symbiosis of human intelligence and artificial intelligence [5]. This transition implies a shift in the role of the AI agent. Currently, AI systems are designed as *cognitive prostheses*, since traditional automation often acts as a prosthesis designed to *replace* human abilities, potentially leading to dependency and deskilling [8]. The goal of Symbiotic AI is to shift toward *cognitive orthosis*, since SAI systems are designed to support, correct, and collaborate with the human, reinforcing their capabilities rather than supplanting them [5, 8]. This distinction is crucial for agentic systems: an autonomous agent should not aim to bypass humans but should scaffold their decision-making process.

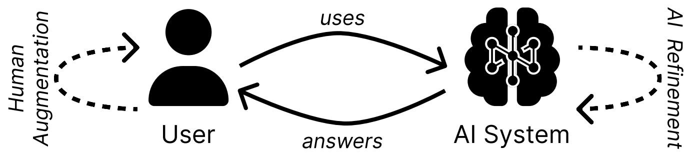

Fig. 1. The main idea behind Symbiotic AI

2.2 The Mechanism of Co-evolution

The heart of the SAI paradigm is the principle of *co-evolution*, which transforms the interaction from a transactional exchange to a bidirectional growth process [2, 5]. This is characterized by two synergetic loops (see Figure 1):

(1) *The AI Refinement Loop:* The agent benefits from a continuous stream of user-provided data and interventions. In a medical context, for instance, every time a physician corrects an agent's classification, the system refines its internal model, mitigating risks such as "concept drifting" over time.
(2) *The Human Augmentation Loop:* Humans benefit from the agent's proactive insights. The AI's explanations and pattern recognition capabilities serve as a form of "machine teaching," enhancing the user's domain knowledge and decision-making objectivity.

2.3 SAI for Agentic Systems

In the context of the AutomationXP26 workshop [14], SAI provides the theoretical scaffolding to manage multi-agent coordination. Unlike standard autonomous agents that might execute tasks in isolation (black-box), symbiotic agents are designed to maintain a "Human-on-the-loop" (HOTL) or "Human-in-the-loop" (HITL) arrangement [5]. This ensures that the human partner can intervene to deal with situational exceptions or reconfigure the AI behavior at run-time [7, 11]. By fostering this symbiotic relationship, we address the challenge of preserving meaningful human agency in an increasingly proactive agentic landscape.

3 Balancing Human Agency and Sustaining Skill Preservation

A critical challenge addressed by the AutomationXP26 workshop is the risk of "de-skilling," where the proactive nature of agentic systems leads to human passivity and the gradual erosion of professional expertise [14]. The SAI paradigm mitigates this by shifting the design philosophy from full automation to a "Human-in-control" philosophy. This not only helps guaranteeing human agency and sustaining human skills, but it also helps creating reliable, trustworthy, and safe systems by reducing the amount of times an AI systems "surprises" its users (e.g., by providing unexpected outcomes) [12]. Through the lens of two medical case studies analyzed in our research [3–5], this section briefly shows how SAI maintains human agency and skills.

3.1 Rhino-Cyt: From Automated Counting to Interactive Learning

Rhinocytology, the study of nasal mucosa cells, traditionally requires labor-intensive microscopic observation to identify nine distinct cytotypes (e.g., neutrophils, eosinophils, lymphocytes) [5]. While an AI agent (Rhino-Cyt) was initially developed to automate cell counting using a Convolutional Neural Network (CNN) [6], its evolution into an SAI system highlights the benefits of symbiosis for skill preservation:

- *Intervention Mechanisms:* Instead of acting as a "black-box oracle," the system employs a "Human-on-the-loop" (HOTL) approach. It presents a preliminary classification which the physician must validate or correct via direct manipulation. This ensures the professional remains the final authority and actively manages the agent's performance against "concept drifting" [5].
- *Knowledge Reinforcement (Machine Teaching):* Through integrated explanations, the agent facilitates a "machine teaching" dynamic. By highlighting specific cellular features, the agent helps physicians recognize rare or ambiguous cytotypes they might otherwise misclassify, thereby actively reinforcing—rather than replacing—their diagnostic skills [5].

3.2 Tumor Detection: Augmenting Agency in High-Risk Scenarios

The analysis of Brain Magnetic Resonance Imaging (MRI) represents a high-risk scenario where the costs of error are critical. Here, agency is preserved by designing the agent to support the surgeon's spatial reasoning rather than dictating a diagnosis:

- *Contextual Augmentation:* The SAI paradigm moves beyond binary diagnostic outputs (tumor/no tumor). Instead, it answers specific clinical questions: "Where is the anomaly?", "Is it a tumor or a cyst?", and "What is the confidence level?" [5]. This nuanced feedback loop forces the physician to engage cognitively with the data.
- *Interactive Verification:* By engaging the surgeon in an interactive verification process—where the human must confirm the tumor boundaries suggested by the AI—the system ensures that the expert maintains situational awareness and meaningful agency over the surgical planning process.

3.3 Implications for Agentic Automation

These two case studies demonstrate that SAI transforms the human-AI dynamic from a zero-sum game of autonomy into a co-evolutionary partnership. By designing agentic systems that require periodic "meaningful human intervention," organizations can leverage the efficiency of AI agents while ensuring that human competence is not only preserved but actively enhanced through the collaboration.

4 Enabling Multi-Agent Transparency

A defining characteristic of agentic automation is the complexity of attribution: when multiple specialized agents coordinate to produce an outcome, it becomes increasingly difficult for the user to discern which entity is responsible for specific actions. This opacity creates a "black box" that threatens user trust and accountability. Within SAI, we propose that transparency must transcend static post-hoc explanations to become a dynamic, interactive process of "interpretability." A first proposal toward this direction is in [7].

4.1 From Explainability to Interactive Attribution

Our survey identifies a crucial distinction between *explainability* (providing reasons for a decision) and *interpretability* (the ability to understand the model's inner workings). For multi-agent systems, simple explanations are often insufficient. We advocate for "interactive attribution" interfaces where users can query the system to isolate the contributions of individual agents. Briefly, we consider the main techniques for providing post-hoc explanations:

- *Feature Importance Techniques:* SAI systems should utilize model-agnostic techniques like SHAP (Shapley Additive Explanations) [9] or LIME (Local Interpretable Model-Agnostic Explanations) [10]. In an agentic context, these can be adapted to quantify the “weight” of a specific agent’s input on the final collective decision, providing granular traceability.

- *Visual Saliency and Attention:* For agents dealing with visual data, such as the MRI analysis in our tumor detection case, saliency maps are vital. These tools visually highlight the specific pixels or regions an agent focused on. This allows the human expert to immediately validate whether the agent is actually considering clinically relevant features or confounding artifacts, thereby building trust through verification.

4.2 Calibrating Trust via Confidence Feedback

Transparency is not merely about explaining “why,” but also about communicating “how sure.” A key pillar of symbiotic transparency is the agent’s ability to communicate uncertainty. By providing explicit feedback on confidence levels and operational limitations, the agent helps users calibrate their trust. This prevents two common failure modes in automation: *over-reliance* (blindly trusting a low-confidence agent) and *under-utilization* (distrusting a capable agent).

4.3 Balancing Transparency and Cognitive Load

Finally, transparency must be balanced against human cognitive limits. As highlighted in our analysis of HCAI challenges [5], providing excessive information about every agent interaction can lead to cognitive overload, paradoxically reducing decision quality. Therefore, SAI interfaces must be *context-sensitive*: providing high-level status updates during routine operations while reserving detailed, granular attribution layers for moments when the user explicitly requests them or when the system detects high uncertainty or risk.

5 Conclusion and Future Research Agenda

As agentic systems become increasingly autonomous and proactive, the risk of marginalizing the human component grows. Within the context of this workshop, this paper is an invitation to collectively rethink agentic automation experiences—not in terms of how much autonomy to delegate, but in terms of how to design enduring human–AI collaboration that support agency, learning, and accountability over time. We have argued that *Symbiotic AI* offers a robust paradigm to mitigate the risk of marginalizing humans by transforming automation from a unidirectional process into a bidirectional interaction, that fosters the co-evolution of humans and AI agents. By synthesizing the findings of our extensive survey and validating them through medical case studies, we have demonstrated that keeping the human “in control” is not a hindrance to efficiency but a prerequisite for long-term reliability and skill preservation.

Based on our analysis, we propose the following research agenda for the Automation Experience community:

- *Standardization of Symbiotic Metrics:* Traditional UX metrics are insufficient for measuring the quality of a human-agent collaboration. Future research should define quantitative measures for mutual learning rates (co-evolution) and the effectiveness of trust calibration over time.
- *Adaptive Transparency Interfaces:* Recognizing the limits of human attention, we need interfaces that dynamically scale the granularity of attribution based on the user’s real-time cognitive load and the criticality of the multi-agent task.
- *Long-term Impact Studies:* It is crucial to conduct longitudinal studies to assess whether agentic SAI effectively sustains professional expertise over years of operation, ensuring that AI-assisted workflows lead to skill enhancement rather than the de-skilling feared in automation scenarios.

In conclusion, the path toward effective “Agentic Automation Experiences” lies in the seamless integration of human and machine intelligence. Moving from HCAI to SAI is not just a terminological shift, but a necessary evolution to ensure that humans remain the protagonists—and not merely the supervisors—of an increasingly automated world. This main-actor role for humans is essential in specific domains, like medicine and e-health, where human control is also required from a legal and an accountability standpoint.

References

[1] Paolo Buono, Nadia Berthouze, Maria Francesca Costabile, Adela Grando, and Andreas Holzinger. 2024. Special Issue on Human-Centered Artificial Intelligence for One Health. *Artificial Intelligence in Medicine* 156 (Oct. 2024), 102946. doi:10.1016/j.artmed.2024.102946

[2] Maria Francesca Costabile, Antonio Piccinno, Daniela Fogli, and Andrea Marcante. 2006. Supporting Interaction and Co-Evolution of Users and Systems. In *Proceedings of the Working Conference on Advanced Visual Interfaces - AVI '06*. ACM Press, Venezia, Italy, 143. doi:10.1145/1133265.1133294

[3] Antonio Curci and Andrea Esposito. 2024. Detecting Brain Tumors Through Multimodal Neural Networks. In *Proceedings of the 13th International Conference on Pattern Recognition Applications and Methods*. SCITEPRESS - Science and Technology Publications, Rome, Italy, 995–1000. doi:10.5220/0012608600003654

[4] Giuseppe Desolda, Giovanni Dimauro, Andrea Esposito, Rosa Lanzilotti, Maristella Matera, and Massimo Zancanaro. 2024. A Human–AI Interaction Paradigm and Its Application to Rhinocytology. *Artificial Intelligence in Medicine* 155 (Aug. 2024), 102933. doi:10.1016/j.artmed.2024.102933

[5] Giuseppe Desolda, Andrea Esposito, Rosa Lanzilotti, Antonio Piccinno, and Maria F. Costabile. 2024. From human-centered to symbiotic artificial intelligence: a focus on medical applications. *Multimedia Tools and Applications* 84 (2024), 32109—32150. doi:10.1007/s11042-024-20414-5

[6] Giovanni Dimauro, Francesco Girardi, Matteo Gelardi, Vitoantonio Bevilacqua, and Danilo Caivano. 2018. Rhino-Cyt: A System for Supporting the Rhinologist in the Analysis of Nasal Cytology. In *Intelligent Computing Theories and Application*, De-Shuang Huang, Kang-Hyun Jo, and Xiao-Long Zhang (Eds.). Lecture Notes in Computer Science, Vol. 10955. Springer International Publishing, Cham, 619–630. doi:10.1007/978-3-319-95933-7_71

[7] Andrea Esposito, Miriana Calvano, Antonio Curci, Francesco Greco, Rosa Lanzilotti, and Antonio Piccinno. 2025. Explanation-Driven Interventions for Artificial Intelligence Model Customization: Empowering End-Users to Tailor Black-Box AI in Rhinocytology. In *End-User Development*, Carmen Santoro, Albrecht Schmidt, Maristella Matera, and Andrea Bellucci (Eds.). Lecture Notes in Computer Science, Vol. 15713. Springer Nature Switzerland, Cham, 161–170. doi:10.1007/978-3-031-95452-8_10

[8] Scott S. Grigsby. 2018. Artificial Intelligence for Advanced Human-Machine Symbiosis. In *Augmented Cognition: Intelligent Technologies*, Dylan D. Schmorrow and Cali M. Fidopiastis (Eds.). Lecture Notes in Computer Science, Vol. 10915. Springer International Publishing, Cham, 255–266. doi:10.1007/978-3-319-91470-1_22

[9] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 4768–4777.

[10] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, San Francisco California USA, 1135–1144. doi:10.1145/2939672.2939778

[11] Albrecht Schmidt and Thomas Herrmann. 2017. Intervention User Interfaces: A New Interaction Paradigm for Automated Systems. *Interactions* 24, 5 (Aug. 2017), 40–45. doi:10.1145/3121357

[12] Ben Shneiderman. 2020. Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. *International Journal of Human–Computer Interaction* 36, 6 (March 2020), 495–504. doi:10.1080/10447318.2020.1741118

[13] Ben Shneiderman. 2022. *Human-Centered AI*. Oxford University Press.

[14] Philipp Spitzer, Matthias Baldauf, Philippe Palanque, Virpi Roto, et al. 2026. AutomationXP26: Agentic Automation Experiences. https://matthiasbaldauf.com/automationxp26/. Accessed: 2026-02-13.