Towards Computer-Using Personal Agents

PIERO A. BONATTI, University of Naples Federico II, Italy
JOHN DOMINGUE, Knowledge Media Institute, The Open University, UK
ANNA LISA GENTILE, IBM Research, USA
ANDREAS HARTH, Friedrich-Alexander-Universität Erlangen-Nürnberg & Fraunhofer Institute for Integrated Circuits IIS, Germany
OLAF HARTIG, Linköping University, Sweden
AIDAN HOGAN, DCC, Universidad de Chile & IMFD, Chile
KATJA HOSE, TU Wien, Austria
ERNESTO JIMENEZ-RUIZ, City St George's, University of London, UK
DEBORAH L. MCGUINNESS, Rensselaer Polytechnic Institute, USA
CHANG SUN, Maastricht University, The Netherlands
RUBEN VERBORGH, IDLab, ELIS, Ghent University – imec, Belgium
JESSE WRIGHT, Department of Computer Science, University of Oxford, UK

Computer-Using Agents (CUA) enable users to automate increasingly-complex tasks using graphical interfaces such as browsers. As many potential tasks require personal data, we propose Computer-Using Personal Agents (CUPAs) that have access to an external repository of the user's personal data. Compared with CUAs, CUPAs offer users better control of their personal data, the potential to automate more tasks involving personal data, better interoperability with external sources of data, and better capabilities to coordinate with other CUPAs in order to solve collaborative tasks involving the personal data of multiple users.

1 Introduction

Advances in Generative AI, and particularly Large Language Models (LLMs), have led to the recent release of various *Computer-Using Agents* (*CUAs*) that automatically operate a user's computer on their behalf. These agents use multimodal capabilities to interact with graphical interfaces via simulated mouse and keyboard inputs. Prominent commercial examples of CUAs include OpenAI's Operator, Google's Jarvis, and new functionalities in Anthropic's Claude.

Potential use cases for CUAs involve personal and often sensitive data, such as credit card details for purchases, passport numbers for flight booking, addresses for deliveries, and allergy information for dinner reservations. While modern browsers sometimes store personal data to autocomplete web forms, CUAs could additionally take context into account (e.g., selecting between a home or work address, depending on the purchase) and go beyond simple autocompletion.

Passing personal data to CUAs raises valid concerns about how such data might be (mis)used. Currently, OpenAI's Operator invokes a *takeover mode* for tasks involving sensitive data (e.g., log-in or payment details): the user is required to fill the details in manually [25]. Such measures target users' concerns about how their personal information will be used by CUAs. OpenAI themselves

Authors' Contact Information: Piero A. Bonatti, University of Naples Federico II, Naples, Italy, pab@unina.it; John Domingue, Knowledge Media Institute, The Open University, Milton Keynes, UK, john.domingue@open.ac.uk; Anna Lisa Gentile, IBM Research, San Jose, CA, USA, annalisa.gentile@ibm.com; Andreas Harth, Friedrich-Alexander-Universität Erlangen-Nürnberg & Fraunhofer Institute for Integrated Circuits IIS, Nürnberg, Germany, andreas.harth@fau.de; Olaf Hartig, Linköping University, Linköping, Sweden, olaf.hartig@liu.se; Aidan Hogan, DCC, Universidad de Chile & IMFD, Santiago, Chile, ahogan@dcc.uchile.cl; Katja Hose, TU Wien, Vienna, Austria, katja.hose@tuwien.ac.at; Ernesto Jimenez-Ruiz, City St George's, University of London, London, UK, ernesto.jimenez-ruiz@city.ac.uk; Deborah L. McGuinness, dlm@cs.rpi.edu, Rensselaer Polytechnic Institute, Troy, NY, USA; Chang Sun, Maastricht University, Maastricht, The Netherlands, chang.sun@maastrichtuniversity.nl; Ruben Verborgh, IDLab, ELIS, Ghent University – imec, Ghent, Belgium, ruben.verborgh@ugent.be; Jesse Wright, Department of Computer Science, University of Oxford, Oxford, UK, jesse.wright@cs.ox.ac.uk.

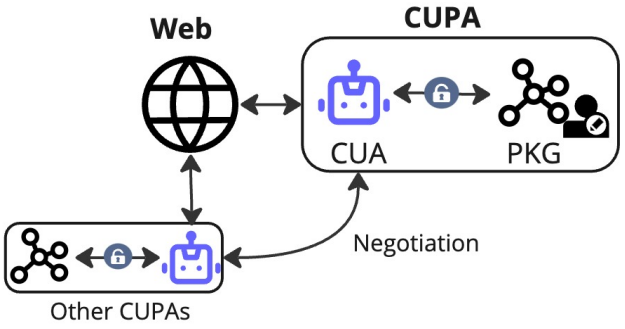

Fig. 1. Computer-Using Personal Agent

state that Operator is *"still learning, evolving and may make mistakes"* [25]. There are thus many open questions relating to the use of personal user data by CUAs.

Conversely, there are many potential benefits to users if CUAs are empowered with personal data. CUAs could autofill forms with personal data for users in a context-aware and potentially generative manner, automating a tedious task. CUAs could potentially enrich personal data with public data to better solve tasks. The CUAs of multiple users could negotiate to achieve a mutually beneficial result based on their users' personal context and preferences.

Towards providing users more oversight over their personal data while enabling higher levels of automation for complex tasks, we propose **Computer-Using Personal Agents** (**CUPAs**): *a Computer-Using Agent (CUA) that has controlled access to a structured repository of private information relating to a user.* This concept is illustrated in Figure 1. Specifically, we propose to instantiate the repository as a *Personal Knowledge Graph* (*PKG*) representing the user's personal data, which would facilitate the specification by users on how the CUA can access and use these data. This PKG can collect more personal data over time, with policies also evolving to reflect the user's fluctuating trust in the system [2]. Looking further forward, one can then imagine a scenario where CUPAs interact with websites and services via the underlying Web APIs instead of through a vision model, where CUPAs can assist in recommendations and negotiations based also on interactions with similar users and/or users' CUPAs.

We provide a road-map towards realising this vision of CUPAs, discussing what is achievable now with current technology, and what gaps must be addressed via further research and development.

2 User Scenario

Sam is expecting Jane over for dinner at 8pm, and is thinking about preparing Thai food. Sam is pre-diabetic, while Jane has a shellfish allergy. Sam requests that his CUPA to generates some suggestions of Thai recipes for the occasion. Consulting Sam's schedule, the CUPA recommends to filter recipes requiring more than an hour to prepare based on when he finishes work and his commute time. Sam agrees, and the CUPA starts to retrieve and present shellfish-free recipes of Thai food that are quick to prepare. Upon consulting external sources of nutritional information and recipes on the Web, the CUPA flags some recipes as being above the postprandial glucose threshold recommended by Sam's doctor (<180 mg/dL), or as having high glycaemic indices (>70).

Sam asks his CUPA to find out what recipe Jane might like. As Sam and Jane are friends, Sam's CUPA can send the candidate recipes to Jane's CUPA to see what she might like. Jane's CUPA suggests to avoid some recipes that include coriander (listed in some recipes as cilantro), which Jane hates. Sam's agent enforces his glucose thresholds and flags ingredients with high glycaemic indices, using external food and recipe knowledge graphs (e.g., the FoodKG [16]) to find the alternative

ingredients. Of the remaining options, Sam's agent suggests a tofu green curry recipe that catches Sam's eye. Since the recipe is flagged for having a high glycaemic index (78), the agent asks Sam if he might consider replacing jasmine rice with cauliflower rice as a healthier option. Sam refuses the substitution as it is a special occasion.

Sam requests his CUPA to order the ingredients from a local supermarket. Since green bell peppers are out of stock, the CUPA suggests to replace them with yellow bell peppers. Sam agrees, and the CUPA prepares the order for delivery to Sam's home address, soliciting Sam's confirmation. Later that night, Sam and Jane enjoy their dinner of Thai green curry. After Jane leaves, Sam suffers some slight heartburn. He requests his CUPA to order antacids and additionally registers the fact that green curry dishes may cause Sam heartburn for future reference.

3 State of the Art

Personal data play an increasingly important role in modern life [6, 24]. Early works [18] characterise such data based on the *concept of six senses*: owned by me, about me, directed to me, sent by me, already experienced by me, and useful to me. More restrictive definitions only include data created by the individual [3], or that the individual cares for/about [11, 12].

Much literature has been dedicated to Personal Information Management systems (PIMs), which deal with the acquisition, organisation, maintenance, retrieval and sharing of personal data [19]. Notable PIM technologies include blockchain systems [14, 36], systems capturing user behaviour on multiple user devices [22, 26], and end-user prototypes [9, 20, 24]. Personal Knowledge Graphs (PKGs) [7, 8, 27] further apply a graph abstraction to personal data, opening up possibilities for declarative access policies, deductive inference, and integration with external Knowledge Graphs.

Towards taking fuller advantage of such data, AI-powered agents show much promise, particularly those that can automate tasks currently performed by the user. Robotic Process Automation (RPA) [13, 29] automates interactions with human interfaces. However, such approaches are hard-coded, brittle to changes in the interface, and incapable of generalising to unseen interfaces. Conversely, AI-based agents are capable of learning and generalising. LLM-based agents have been proposed to operate in diverse environments using recursion, feedback, and careful prompt engineering [33]. Such LLM-based agents are capable of solving computer tasks – despite the limited reasoning capabilities in LLMs [21] – paving the way for CUAs such as Operator [25].

Regarding works unifying LLM-based agents with PKGs, AGENTIGraph [34] heads in this direction, but rather focuses on question answering. Closer to the idea of CUPAs is Charlie: a brief proposal by Berners-Lee [4] on combining LLM-based agents with PKGs instantiated by Solid pods using Semantic Web standards. This proposal, and the user scenario presented previously, echo the (yet unrealised) vision laid out by Berners-Lee et al. [5] for the Semantic Web 24 years ago. Wright [31] presents a "discuss then transact" model of LLM-interaction in support of this vision for LLM-based personal agents that represent legal entities.

4 Added Value

Societal and legal debates on personal data emphasise *protection* from the harm that they could inflict, and understandably so. Yet people voluntarily exchange personal data with others in their every-day lives in the pursuit of mutual benefit. People can decide to leverage more personal data, or different kinds of personal data, to achieve a desired outcome. For instance, patients might prefer to share fitness-tracker data with their doctor if this improves their treatment, or consumers might want to divulge allergies and dietary needs to streamline online shopping and avoid nasty surprises.

A dangerous assumption is that companies are more capable of distilling value from people's personal data than the people the data describe. A company certainly has advantages over individuals in this respect, such as the ability to aggregate over a great many users. But personal data about

a particular individual in isolation has much greater potential to empower that individual than a company they interact with, especially when the individual is coached by an agent such as a CUPA. CUPAs representing different parties could even negotiate a better outcome for *all* parties involved.

Considering the added value of CUPAs, and more generally of providing AI-based agents access to personal data, we highlight:

Multi-dimensional negotiation. CUPAs can help users to strike sweet-spots between multiple dimensions, such as the cost and duration of multi-hop flights, the deliciousness and healthiness of meal options, etc.

Increased granularity. Humans struggle to negotiate on a fine-grained level, and may thus prefer broad policies that reduce cognitive load (e.g., to always accept all cookies) [32]. CUPAs can help to reach fine-grained agreements that improve outcomes and honour party preferences.

Improved risk/reward assessment. CUPAs can help users simulate and analyse a variety of hypothetical data exchange scenarios, and warn users of a particular risk, for example that the supermarket – if informed of a condition of a severe allergy – could sell this information to third parties, leading to an increase in life assurance premiums.

Auditing and follow-up. CUPAs could automatically perform audits to assess whether the data were treated as agreed during the negotiation process, evaluate the benefit to the user, and improve for future interactions.

Such added value is, of course, dependent on the value outweighing the potential harms caused. This can be addressed via AI alignment, which ensures that artificial intelligence systems act in accordance with human intentions, values, and societal norms. It involves *outer alignment*, where an AI's objectives accurately reflect human goals, and *inner alignment*, ensuring learned behaviours remain aligned in novel scenarios. Machine-readable policies on how personal data from the PKG can or should be used by the AI-based agent can also help to avoid harm. Representing personal data as PKGs allows standards such as the Open Digital Rights Language (ODRL) [17] and policy engines implementing formal semantics [15] to specify and automate the processing of policies about how personal data are used, in what contexts, and under what conditions.

5 CUPA Capabilities

Computer-using personal agents must be able to *interact with diverse websites and APIs*. This allows them to book flights and hotels, search for job openings, and even schedule appointments. Moreover, they must possess the ability to *interact with other such agents*, such as coordinating travel arrangements with a travel agent or collaborating with a financial agent to manage expenses.

In addition to being able to *generate and adapt content* (e.g., personalised summaries and creative text), a computer-using personal agent must be able to *combine private data* from the user's personal knowledge graph (PKG) with external information. For example, when searching for a new apartment, the agent should combine the user's preferred neighbourhood from their PKG with data from real estate websites and local amenities databases to find the most suitable options. When utilising the knowledge stored within the PKG, the agent must also be able to *adapt the knowledge from the PKG for the current task*. For instance, when filling out a job application form, the agent should selectively use information from the user's CV and work history stored in the PKG, tailoring the presentation to the specific requirements of each application. This adaptability is crucial for ensuring that agent actions are relevant and effective in the given context.

CUPAs must continuously *collect and enrich user information* to effectively assist them. This involves gathering data from various sources, including interactions with websites and APIs, user inputs, and external sources. By continuously *learning about user preferences*, these agents

can personalise their assistance, such as recommending travel options that align with the user's preferences or suggesting recipes that cater to specific dietary restrictions or tastes. However, it is also crucial for such agents to *avoid learning one-off or irrelevant patterns*, for example, to assume that Sam will always suffer heartburn after eating Thai food and should thus avoid it.

Computer-using personal agents must exhibit a high degree of autonomy. They should ideally *act maximally autonomously*, including the ability to *proactively anticipate and address user needs*. For example, an agent could proactively remind users of upcoming appointments or suggest relevant articles based on their recent reading history. However, this autonomy must always be balanced with the ability to *be guided and controlled by the user*, allowing users to provide instructions, adjust preferences, and maintain control over agentic actions.

While acting largely autonomously, it is crucial that a computer-using personal agent *acts in alignment with the user*, ensuring that tasks are completed as desired. This is essential in scenarios like recipe searches where the agent must accurately reflect dietary restrictions and preferences. Moreover, such an agent should always act in the user's interests, even when *dealing with potentially conflicting goals*. For example, an agent helping a user plan a trip should consider factors like budget, travel time, and personal preferences, even if these factors may lead to a slightly more expensive or less convenient option. The agent should avoid *acting in an unethical or illegal manner* even if it potentially maximises a users immediate interests, e.g., via tax evasion.

To maintain user trust and ensure responsible behaviour, it is also crucial that agents *do not overstep bounds*, respecting user privacy and only acting within explicitly granted permissions. Finally, the repeated offering of *clear explanations of all actions* will aid in the fostering of trust and allow users to understand and verify agent behaviour.

6 Technical Challenges

The aforementioned desired capabilities for CUPAs, based on our vision of a trusted, accountable and largely autonomous agent acting with personal data for user benefit, raises a number of technical challenges.

Accountability and Liability In the case of undesired, illegal, or unethical acts involving CUPAs, it is important to determine who – or what – is responsible, who should be held accountable, and where the liability lies.

Explainability, Traceability, and Provenance Provenance techniques are required to trace and explain how personal and external data led to specific answers or actions being derived or carried out by the CUPA. These provenance techniques would need to support diverse data models, machine learning processes, user inputs and policies.

Data Interoperability Data interoperability is a key challenge towards implementing CUPAs. Being able to draw on and integrate more sources of data will improve the CUPAs performance. This is particularly challenging for new sources discovered on the fly.

Inter-Agent Communication, Negotiation and Coordination Agents must communicate effectively in the context of multi-agent systems to achieve shared goals, requiring both a shared conceptual understanding and a means of encoding and decoding messages [30]. The same challenge applies to networks of CUPAs who coordinate to solve a particular set of goals for users.

Security, Privacy, and Policies The sensitive nature of data processed by a CUPA calls for security, privacy, and usage control mechanisms, and the ability of the CUPA to understand and correctly apply the access/usage/action control policies of the user. In some countries, this would even be a legal requirement (e.g., under GDPR in the E.U.).

Trust, Delegation, and Action Control Achieving agent autonomy requires trust modelling, delegation mechanisms, and structured action control policies [28]. Trust models must be adaptable to different contexts, from rigid policies applicable in government agencies to more flexible, reputation-based approaches for personal agents [10].

User-in-the-Loop CUPAs will require input, guidance, permission and confirmation from the user. But to increase automation, the CUPA must avoid unnecessary interactions with the user. This creates the challenge of *when* to call upon the user, and how.

Self-Improvement The CUPA should leverage its experience with the user in order to improve the services it provides over time, leading to greater automation, and actions/results that better benefit the user. This raises questions about how such a history can be captured, represented, stored and leveraged.

Self-Determination and Alignment There are many ways an agent could be considered 'aligned' to a user. Naive approaches include ensuring that CUPA decision making always takes place within rules-based bounds - such as action controls set by a user - or doing a best effort to match user *intent* or *decision making*. There is a field of research discussing 'beyond preference matching' alignment which proposes that AI systems should be aligned to broader concepts such as value-based alignment, or prioritise user welfare over emulating decision making [35].

7 Roadmap

We envisage that moving from the current state of the art to fully addressing the above technical challenges will occur in three stages. These levels represent varying degrees of trust, accountability and autonomy.

CUAs enhanced with personal data In the first instance, we foresee extensions of CUAs – in the style of OpenAI's Operator [25] in a commercial setting and Agent-E [1] in a research setting – such that they use a PKG in order to access knowledge personal to the user. This would safely enable higher levels of automation, obviating the need to pass control back to the user in scenarios of the user's choosing that involve personal data.

Web-aware CUPAs CUAs currently rely on existing browser implementations to render an HTML page and then make use of vision models to interact with the page. An agent could rather observe HTTP requests made to a particular website, as well as the HTML forms present on a page, to invoke requests and actions directly via HTTP.

Networks of CUPAs We envision networks of CUPAs interacting in order to complete tasks involving multiple users. This may involve structured service descriptions [23], or a mix of natural language and structured communication per a "discuss then transact" model [31] whereby agents use natural language to first negotiate about a transaction they wish to perform, and then confirm this transaction using structured data.

8 Conclusion

Computer-Using Agents (CUAs) have the potential to transform how users interact with their computers, their browsers and amongst themselves. Not having access to personal data limits such interactions. Giving CUAs unfettered access to the personal (and most sensitive) data of a user seems unwise, as does providing CUAs no access to personal data. We thus argue for CUPAs as a configurable middle-ground, where a Personal Knowledge Graph (PKG) is used to represent, store and control access to the personal data of the user. As a starting point, the data that a user fills into web forms can be captured in the PKG, and enriched by an AI-based agent. These data can then be used, if the user so wishes, by CUAs to automate further tasks. In a next step, CUPAs can learn to

interact with websites via HTTP APIs rather than though visual interfaces. Finally, we envisage further into the future a network of CUPAs collaborating to address users' tasks.

Acknowledgments

This report is a result of Dagstuhl Seminar 25051 "Trust and Accountability in Knowledge Graph-Based AI for Self Determination", which took place in January 2025.

References

[1] Abuelsaad, T., Akkil, D., Dey, P., Jagmohan, A., Vempaty, A., and Kokku, R. Agent-e: From autonomous web navigation to foundational design principles in agentic systems. *arXiv preprint arXiv:2407.13032* (2024).

[2] Afroogh, S., Akbari, A., Malone, E., Kargar, M., and Alambeigi, H. Trust in AI: progress, challenges, and future directions. *Humanities and Social Sciences Communications 11*, 1 (2024), 1–30.

[3] Bergman, O., and Whittaker, S. *The science of managing our digital stuff*. MIT Press, 2016.

[4] Berners-Lee, T. Charlie Works. Design Issues, https://www.w3.org/DesignIssues/Works.html, 2025.

[5] Berners-Lee, T., Hendler, J., and Lassila, O. The semantic web. *Scientific American 284*, 5 (2001), 34–43.

[6] Birch, K., Cochrane, D. T., and Ward, C. Data as asset? the measurement, governance, and valuation of digital personal data by big tech. *Big Data and Society 8* (2021).

[7] Chakraborty, P., Dutta, S., and Sanyal, D. K. Personal research knowledge graphs. In *WWW 2022 - Companion Proceedings of the Web Conference 2022* (4 2022), Association for Computing Machinery, Inc, pp. 763–768.

[8] Chakraborty, P., and Sanyal, D. K. A comprehensive survey of personal knowledge graphs, 11 2023.

[9] Chaudhry, A., Crowcroft, J., Howard, H., Madhavapeddy, A., Mortier, R., Haddadi, H., and McAuley, D. Personal data: Thinking inside the box. *Aarhus Series on Human Centered Computing 1* (2015), 4.

[10] Chen, R., Bao, F., and Guo, J. Trust-based service management for social internet of things systems. *IEEE transactions on dependable and secure computing 13*, 6 (2015), 684–696.

[11] Cushing, A. L. PIM as a caring: using ethics of care to explore personal information management as a caring process. *Journal of the Association for Information Science and Technology 74*, 11 (2023), 1282–1292.

[12] Cushing, A. L., and Kerrigan, P. Personal information management burden: A framework for describing nonwork personal information management in the context of inequality. *Journal of the Association for Information Science and Technology 73* (11 2022), 1543–1558.

[13] da Silva Costa, D. A., Mamede, H. S., and da Silva, M. M. Robotic Process Automation (RPA) adoption: a systematic literature review, 6 2022.

[14] Faber, B., Michelet, G., Weidmann, N., Mukkamala, R. R., and Vatrapu, R. Bpdims:a blockchain-based personal data and identity management system. *Proceedings of the Annual Hawaii International Conference on System Sciences 2019-Janua* (2019), 6855–6864.

[15] Fornara, N., Rodríguez-Doncel, V., Esteves, B., Steyskal, S., and Smith, B. W. ODRL Formal Semantics, May 2024.

[16] Haussmann, S., Seneviratne, O., Chen, Y., Ne'eman, Y., Codella, J., Chen, C.-H., McGuinness, D. L., and Zaki, M. J. FoodKG: A semantics-driven knowledge graph for food recommendation. In *The Semantic Web – ISWC 2019* (Cham, 2019), C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois, and F. Gandon, Eds., Springer International Publishing, pp. 146–162.

[17] Iannella, R., and Villata, S. ODRL Information Model 2.2, Feb 2023.

[18] Jones, W. *The future of personal information management, part 1: Our information, always and forever*. Morgan & Claypool Publishers, 2012.

[19] Jones, W. P., and Teevan, J. *Personal information management*, vol. 14. University of Washington Press Seattle, WA, 2007.

[20] Kalokyri, V., Borgida, A., and lie Marian, A. YourDigitalSelf: a personal digital trace integration tool. *International Conference on Information and Knowledge Management, Proceedings* (2018), 1963–1966.

[21] Kim, G., Baldi, P., and McAleer, S. Language Models can Solve Computer Tasks. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023* (2023), A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds.

[22] Lin, J., and Wang, M. PKG : A Personal Knowledge Graph for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22), July 11â•fi15, 2022, Madrid, Spain* (2022), vol. 1, Association for Computing Machinery, pp. 3334–3338.

[23] Martin, D., Paolucci, M., McIlraith, S., Burstein, M., McDermott, D., McGuinness, D., Parsia, B., Payne, T., Sabou, M., Solanki, M., Srinivasan, N., and Sycara, K. Bringing semantics to web services: The OWL-S approach. In *Semantic Web Services and Web Process Composition* (Berlin, Heidelberg, 2005), J. Cardoso and A. Sheth, Eds., Springer Berlin Heidelberg, pp. 26–42.

[24] Mortier, R., Zhao, J., Crowcroft, J., Wang, L., Li, Q., Haddadi, H., Amar, Y., Crabtree, A., Colley, J., Lodge, T., Brown, T., McAuley, D., and Greenhalgh, C. Personal data management with the databox: What's inside the box? In *Proceedings of the 2016 ACM Workshop on Cloud-Assisted Networking* (New York, NY, USA, 2016), CAN '16, Association for Computing Machinery, p. 49–54.

[25] OpenAI Team. Introducing Operator. OpenAI Blog https://openai.com/index/introducing-operator/, published 2025-01-23, accessed 2025-01-29, 2025.

[26] Schröder, M., Jilek, C., and Dengel, A. A Human-in-the-Loop Approach for Personal Knowledge Graph Construction from File Names. In *Knowledge Graph Construction* (2022), vol. 3141, CEUR Workshop Proceedings.

[27] Skjæveland, M. G., Balog, K., Bernard, N., Łajewska, W., and Linjordet, T. An ecosystem for personal knowledge graphs: A survey and research roadmap. *AI Open 5* (2024), 55–69.

[28] South, T., Marro, S., Hardjono, T., Mahari, R., Whitney, C. D., Greenwood, D., Chan, A., and Pentland, A. Authenticated delegation and authorized ai agents. *arXiv preprint arXiv:2501.09674* (2025).

[29] van der Aalst, W. M., Bichler, M., and Heinzl, A. Robotic process automation. *Business and Information Systems Engineering 60* (8 2018), 269–272.

[30] Wooldridge, M. *An introduction to multiagent systems*. Wiley, 2009.

[31] Wright, J. Here's Charlie! Realising the semantic web vision of agents in the age of LLMs. *CoRR abs/2409.04465* (2024).

[32] Wright, J., Esteves, B., and Zhao, R. Me want cookie! Towards automated and transparent data governance on the Web, 2024.

[33] Yang, J., Prabhakar, A., Narasimhan, K., and Yao, S. InterCode: Standardizing and Benchmarking Interactive Coding with Execution Feedback. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023* (2023), A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds.

[34] Zhao, X., Blum, M., Yang, R., Yang, B., Carpintero, L. M., Pina-Navarro, M., Wang, T., Li, X., Li, H., Fu, Y., Wang, R., Zhang, J., and Li, I. AGENTiGraph: an interactive knowledge graph platform for LLM-based chatbots utilizing private data, 2024.

[35] Zhi-Xuan, T., Carroll, M., Franklin, M., and Ashton, H. Beyond preferences in ai alignment. *Philosophical Studies* (Nov. 2024).

[36] Zyskind, G., Nathan, O., and Pentland, A. S. Decentralizing privacy: Using blockchain to protect personal data. *Proceedings - 2015 IEEE Security and Privacy Workshops, SPW 2015* (2015), 180–184.