Gaze-based Opportunistic Privacy-preserving Human-Agent Collaboration

Jan Grau
jan.grau@student.unisg.ch
University of St. Gallen
St.Gallen, Switzerland

Simon Mayer
simon.mayer@unisg.ch
University of St. Gallen
St.Gallen, Switzerland

Jannis Strecker
jannisrene.strecker@unisg.ch
University of St. Gallen
St.Gallen, Switzerland

Kimberly Garcia
kimberly.garcia@unisg.ch
University of St. Gallen
St.Gallen, Switzerland

Kenan Bektaş
kenan.bektas@unisg.ch
University of St. Gallen
St.Gallen, Switzerland

Figure 1: An opportunistic notification from a context-aware, privacy-preserving notification system based on gaze data.

ABSTRACT

This paper introduces a novel system to enhance the spatiotemporal alignment of human abilities in agent-based workflows. This optimization is realized through the application of Linked Data and Semantic Web technologies and the system makes use of gaze data and contextual information. The showcased prototype demonstrates the feasibility of implementing such a system, where we specifically emphasize the system's ability to constrain the dissemination of privacy-relevant information.

CCS CONCEPTS

• **Human-centered computing → Collaborative interaction**; *Mixed / augmented reality*; Ubiquitous and mobile computing systems and tools; Synchronous editors; • **Security and privacy →** *Usability in security and privacy*; *Privacy protections*; • **Computer systems organization →** *Distributed architectures.*

KEYWORDS

Privacy-Preserving, Human-Agent-Collaboration, Solid, Koreografeye

ACM Reference Format:
Jan Grau, Simon Mayer, Jannis Strecker, Kimberly Garcia, and Kenan Bektaş. 2024. Gaze-based Opportunistic Privacy-preserving Human-Agent Collaboration. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24), May 11–16, 2024, Honolulu, HI, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3613905.3651066

1 INTRODUCTION

Enhancing the interaction and collaboration between humans and artificial software agents has received much attention over the last decades (cf. [7]). In such collaborative settings, humans exhibit a broad range of abilities that could support artificial agent systems when these are missing specific actuators or sensors for a task—and vice versa. Consider an industrial company equipped with an agent system tasked with the maintenance of a machinery fleet. The agent system might be capable of accessing and analyzing sensor data from the machines, enabling it to assess immediate failure risks. Additionally, it may take action on various actuators to ensure the smooth production of the final product. However, for the system to run continuously, an annual routine check is required for the machines, necessitating the expertise of a skilled employee capable of maintaining and inspecting the machinery. As the agent system

CHI EA '24, May 11–16, 2024, Honolulu, HI, USA

ACM ISBN 979-8-4007-0331-7/24/05
https://doi.org/10.1145/3613905.3651066

lacks the capability to perform this task independently, human-agent collaboration is essential for the system to prevail.

However, the maintenance process demands the employee to bring specialized tools and be physically present on-site. A possible solution to achieve the maintenance process includes scheduling dedicated time of the employee and using planning techniques to ensure that the required equipment is prepared in advance. A more efficient approach would be to instruct the employee *while they are already engaged in working on the machine*, to procure the necessary equipment and carry out the maintenance task. However, this approach might require the employee to transition from their ongoing task (working at the machine) to fetch the needed equipment for performing the maintenance task. Such task switches and transitioning between activities incur high cost, as shown in previous research [23, 24, 26, 30, 34, 35].

To minimize such transitions, we instead propose that the employee is *opportunistically nudged in a contextually adequate moment* to procure the necessary tools and carry out the maintenance task. For instance, this nudge would be appropriate when they are already on their way to the machine and—coincidentally—pass by the required tools. To enable such opportunistic nudging, it is necessary to identify the appropriate moment to notify the employee about the potential task, based on contextual information. Possibly, for instance, the employee's current gaze behavior might reveal relevant information; utilizing gaze as a proxy for intent has been found to be highly suitable [5, 28]. Specifically, if the employee is already looking at the necessary specialized tool, we can infer that they could potentially transport the tools to the machine where they are required. We note that the inclusion of gaze data however poses a significant privacy risk due to the wealth of information that our gaze behavior might reveal about us, including information about our "biometric identity, gender, age, ethnicity, body weight, personality traits, drug consumption habits, emotional state, skills and abilities, fears, interests, and sexual preferences" [22]. Hence, it is imperative for the employee in our scenario to retain control over the information-sharing process, deciding what information should be disclosed to whom, and at what time. In addition to the attention of the human, further contextual parameters need to align with the situation. For instance, it should be likely that the employee has enough time to carry out the desired task, be adequately qualified, and be headed towards the to-be-maintained machine; otherwise, the system should refrain from nudging them to avoid overly cluttering their perception.

The industrial scenario described above is one use case among the many in which our opportunistic human-agent collaboration system could be useful. It was chosen to illustrate that our system could indeed reduce time, effort, and hence cost by contextually supporting humans performing complex tasks that might require preparation to be done (e.g., procure equipment). Other application domains for our system are hospitals, schools, and smart homes.

This paper presents a prototype of a novel opportunistic human-agent collaboration system that determines what nudges a person should receive based on different contextual information sources, including live gaze data. Our design facilitates the precise timing of notifications to humans, ultimately aiming to enhance human-agent collaboration in a manner that closely aligns with human preferences and workflows. To counteract privacy challenges that arise from the use of gaze data in our system, the prototype includes a mechanism that permits people to select which gaze-derived information should be shared, and which should remain private.

2 RELATED WORK

Significant research has been conducted in the domains of human-agent-collaboration, opportunistic notifications, and using gaze data with a privacy focus. In the following, we provide a brief overview of relevant research across these domains.

Human in the Loop: Human-Agent Collaboration. Since its inception, the field of human-computer interaction has been investigating the efficient collaboration of people and automated systems. Recent work on human-in-the-loop systems in industry presents a structured review of key technologies employed to enhance human-mediated automated production scenarios [31]. This survey highlights that, in addition to explainability (especially of Artificial Intelligence components) and automated audit trails, context-aware computing, visualization, interaction, and collaborative robotics play pivotal roles in advancing the research field. An illustrative example is found in the work of Garrido-Hidalgo et al. [15], who advocate for human-machine collaboration in an Internet of Things (IoT) environment. This collaboration involves incorporating sensors into industrial clothing, creating a human-in-the-loop system tailored to human needs. This system was shown to enhance security and prevent some workplace hazards. In the context of human-computer interaction, also *opportunistic notifications* have been investigated. For instance, [14] explores optimal timing for delivering notifications on mobile phones. This study shows that higher situational appropriateness of the interruption significantly reduces delay and increases task completion rate.

Interruptions at the Workplace. The adverse effect of interruptions while people engage in tasks and workflows has also been well-documented more broadly: O'Conaill and Frohlich [27] demonstrated that in over 40% of workflow interruptions, the recipient did not resume the initially intended work, highlighting the necessity for more efficient strategies to address interruptions. Zijlstra et al. [36] found that, while interruptions did not directly degrade performance (thanks to effective coping strategies), they did have a negative impact on emotion and well-being, consequently leading to increased effort.

Gaze as a Proxy for Intent. To determine when an interruption might be most appropriate, the current intent of the targeted user likely provides valuable information. In addition to the recently demonstrated possibility of using gaze data to determine an individual's current activity on an Augmented Reality (AR) headset [2], already a 2005 study of Qvarfordt and Zhai [28] confirms the potential benefits of utilizing eye-gaze data for understanding human intent. For example, to improve the mutual understanding and sense of co-presence in remote collaboration environments, Bektaş et al. [4] recently developed an AR interface that allows a pair of human users to see each other's gaze behavior and make eye contact. Ajanki et al. [1] document similar insights by integrating gaze data into an AR application. Taking a further step, Kandemir and Kaski [19] integrate machine learning with gaze tracking to learn specific gaze patterns exhibited by individuals. Gaze information might not only permit estimating user intent, but can furthermore be used to better communicate with the user: Gaze-contingent displays can dynamically adjust the level of detail according to the user's needs (e.g., attention and cognitive load) and provide them with personalized assistance in solving visual tasks [3]. Klauck et al. [21] employed gaze information to reduce the frequency of notifications, including chat alerts, new emails, or error messages on a screen. In a separate study in 2023, Hüber et al. [17] successfully used gaze information to strategically place notification banners in less distracting locations. Commercially available AR headsets can be used to recognize activities from user eye movements and provide them with AR feedback that is relevant to their current activity [6]. While eye tracking data hence has many potential uses in the context of an opportunistic human-agent interaction system, its usage is also problematic: In their comprehensive literature review, Kröger et al. [22] show that eye tracking data has the potential to reveal a wide range of personal information, including biometric identity, gender, age, ethnicity, body weight, personality traits, drug consumption habits, emotional state, skills and abilities, fears, interests, and sexual preferences. This underscores the need to treat eye tracking data as highly sensitive. To address privacy concerns, Bozkir et al. [12] proposes a privacy-preserving gaze estimation method, utilizing synthetic images generated through a randomized encoding-based framework.

3 A PRIVACY-PRESERVING GAZE-BASED OPPORTUNISTIC NOTIFICATION SYSTEM FOR HUMAN-AGENT COLLABORATION

Based on the surveyed related work, we present an opportunistic human-agent collaboration system that makes use of contextual information from gaze and additional sources to determine when human users should be nudged towards a given task. Towards preserving their privacy, in our system, the user remains in control of the processing of their contextual data (including gaze).

An overview of our system is shown in Figure 2: In addition to a head-mounted display (HMD) used for gaze tracking, we use the *Koreografeye* [16] orchestration engine to manage context data that is distributed across several personal decentralized data stores—to this end, our system makes use of the *Solid*[1] specification (cf. [6]). By configuring access control lists that are associated with these data stores, users (and agents) can exert fine-grained control about when and with what systems the data in their pods is shared. The orchestration engine is configured with a collection of rules; it monitors the data stores and reacts to specific data changes according to these rules by processing this data and writing the results to other Solid data stores. The orchestration engine, specifically, processes the contained data to infer currently possible activities of the user and informs the agent that it may nudge the user towards one of these.

3.1 Technologies Used

Our prototype system consists of a large number of interacting components. To enable the reader to better understand the overall system, we first give an overview of each utilized technology.

[1]See https://solidproject.org/

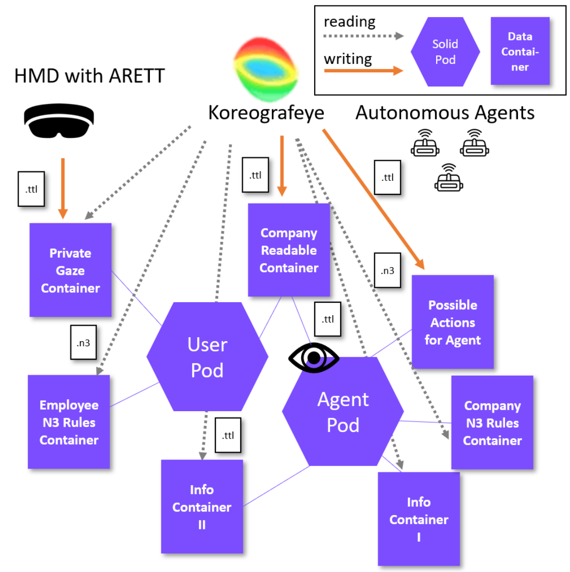

Figure 2: Demonstration Setup - Simplified. A user equipped with a Head Mounted Display (HMD) in conjunction with their Solid pod, along with a Koreografeye orchestrator and an agent system comprising a Solid pod, as well as reasoning and acting software.

Solid. Solid (formerly *Social Linked Data*) pods are decentralized data-sharing platforms on the basis of the World Wide Web Consortium's (W3C) Solid specification [29]. These pods enable users to create *Linked Data* [10] containers for integrating personal data into the Semantic Web while maintaining fine-grained access control. Users can specify access permissions for each container, ensuring a secure and controlled approach to data sharing.

Orchestration on the Semantic Web. The incorporation of N3-based rules [8] as exemplified in Listing 1 represents a Semantic-Web-based approach to seamlessly integrate rules and reasoning. Grounded in the N3 format [9], these rules facilitate reasoning processes reliant on existing queryable data, aligning seamlessly with the Solid ecosystem [29]. An orchestrator leverages these rules to dynamically generate or modify data. An example orchestrator in this domain, which we make use of in our system, is Koreografeye [16]. This orchestrator employs the EYE reasoner [32] to execute rules in conjunction with specified data.

Listing 1: A sample N3 query [8]

```
{
    ?x a ex:TestData .
    ( ?x + 1 ) AS ?s .
} => {
    ?s a ex:Result .
}
```

Web Embodied Agents. The W3C Web of Things (WoT) [18] addresses the challenge of application-layer fragmentation on the Internet of Things (IoT) through a common Web-based application layer and interface definition language called W3C WoT Thing Description (TD). This facilitates seamless integrations and recent research has proposed and demonstrated integration of the WoT with core technologies and approaches from the field of Autonomous Agents and Multiagent Systems (cf. [13]). This approach further is compatible with a recent integration of the WoT with Solid as data platform (see [37]), which demonstrates the successful embodiment of agents into the Web.

JaCaMo Web-based Autonomous Agents. The JaCaMo project [11] is dedicated to advancing the Multi-Agent Oriented Programming (MAOP) paradigm by offering a comprehensive development platform. This platform seamlessly integrates a suite of tools and programming languages tailored for addressing the design dimensions (agent, environment, and organization) of Multi-Agent Systems (MAS).

3.2 Demonstrator

Our system integrates these technologies to permit opportunistic human-agent collaboration. As the system could potentially be used in a variety of setups, we first present the generic setup of the system to be followed by the detailed implementation of the real-world use case illustrated in the introduction to prove its functionalities.

The overall configuration is depicted in Figure 2. Our system comprises a gaze-enabled HMD that captures gaze data as well as the focused objects. To enable users to exert fine-grained control over their data as well as derived information our system is capable of writing data to a private Solid pod of the user. The same pod includes a container with N3-based rules that specify what data should be made public. The orchestration engine at the same time monitors changes in the private gaze data pods, executes the specified N3-based rules, and writes the results from processing the data using these rules to another gaze data container. This container may be made fully public, or the user may configure it to permit access to specific entities (e.g., software programs or other people).

On the receiving end of the information system is the agent, whose architecture is flexible and presented as a suggestion. Other agent systems can be employed due to the high interoperability of the WoT interface [33]. In our prototype system, the agent also uses a Solid pod. Leveraging the `WebhookChannel`[2] capability of the Solid specification, the agent monitors changes in the public gaze data. It incorporates N3-based rules that process the gaze data to determine whether the system should expose an ability of the user that the agent might nudge them towards. In the proposed setup, WoT Thing Description[3] data is used for the transmission of essential information to a JaCaMo agent. This equips the JaCaMo agent, as a part of a multi-agent system based on Multi-Agent-Oriented-Programming (MAOP), to integrate the nudging possibility into its own reasoning and possible actions. As already mentioned, the JaCaMo agent should be seen as a suggestion. Due to the high interoperability of the WoT interface[33] a variety of different software-based agents could be used to leverage the benefits of the system presented.

To demonstrate the effectiveness of the proposed system in more detail, the system is applied to the machine maintenance use case as

[2]https://solid.github.io/notifications/webhook-channel-2023
[3]https://www.w3.org/TR/wot-thing-description/

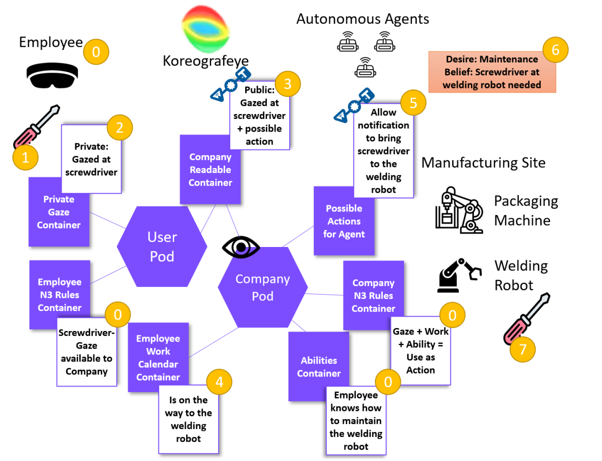

Figure 3: Real World Scenario. Illustration of all the relevant steps included in the demonstrator.

described in the introduction. The machine maintenance scenario involves seven steps, depicted in Figure 3. The maintenance task in this scenario is typically carried out by experts who regularly engage with the machine and is required in a regular interval (e.g. once a year). Additionally, a specific tool, in this case a screwdriver, is necessary for the maintenance process. The seven steps are explained in more detail below describing a possible happy path in the machine maintenance scenario to better illustrate the chain of events required in the proposed system.

Precondition In the outlined scenario, the foundational prerequisites involve the employee using an HMD (i.e. Microsoft Hololens 2) equipped with the Augmented Reality Eye Tracking Toolkit (ARETT) [20] for tracking the gaze data of the employee. Additionally, the employee utilizes a Solid pod where N3-based rules are defined, specifying what information should be made public and what should remain private. JaCaMo-based autonomous agents are responsible for managing the manufacturing site and have their own Solid pod. This pod provides additional contextual data, such as an employee's eligibility for machine maintenance and details about the employee's work agenda. The company can also establish dynamic rules, updating possible actions based on changes in the public gaze data of the employee. Through Solid WebhookChannels, the agent monitors changes in the public gaze data of the employee's pod.

Step 1 The employee gazes at a screwdriver and the HMD records the raw gaze data.

Step 2 The HMD converts the gaze data along with information about the location of scene objects (i.e., the screwdriver) to semantic triples. In this case, a triple is created that expresses that the user gazes at the screwdriver. This information is stored in the private gaze container of the employee.

Step 3 The reasoner assesses whether the gaze data should be made public, adhering to the provided N3-based rules. For instance, these rules can specify that information about gazing at *tools* may be published while information about gazing at *other people* or must remain private.

Step 4 Due to the configured WebhookChannels the agent's Solid pod reacts to the change in the user's gaze information. The hook triggers the Koreografeye orchestrator to run the reasoning based on the new data and the provided N3-based rules. Because the context information available in the agent's data container indicates that the employee who looked at the screwdriver is not only on the way to the welding robot, but also eligible to maintain the welding robot, a possible action to bring the screwdriver to the welding robot is inferred.

Step 5 This identified possible action (bringing the screwdriver to the welding robot) is transmitted to the JaCaMo agent.

Step 6 The JaCaMo agent incorporates the new possible action into its own ruleset. Since the agent is determined to maintain the welding robot, the agent is able to *opportunistically nudge* the employee, because of the *contextually adequate moment*. Thus, the agent invokes the WoT Thing Description integrated in the possible action and notifies the employee to take the screwdriver. The possible nudge is illustrated in Figure 1.

Step 7 The employee receives the notification of bringing the screwdriver to the welding robot. Since it is still a *nudge* the employee, the employee could still neglect the notification. However, because the context for bringing the screwdriver to the welding robot is ideal, the additional effort on the employee is fully minimized and therefore the employee seizes the chance and brings the screwdriver to the welding robot for maintenance.

Opportunistic Collaboration. In the aforementioned scenario, the employee who was looking at the screwdriver was coincidentally on the way to the welding robot and eligible of maintaining the welding robot. This scenario was used to demonstrate the happy path of the system presented. However, to see the context fit the optimal circumstances will occur less often then not. For instance, if the employee would have not been on the way to the welding robot when looking at the screwdriver, the reasoning would not forward the possible action to bring the screwdriver to the welding robot. The same is also true for the case when the user is not eligible of maintaining the welding robot. The system presented is therefore able to react opportunistically and thus only nudging in opportune moments. But foremost, it is the employee who is profiting from this presented demonstrator, as the N3-based rules allow the employee to change the rules during runtime and could always configure all gazes to private (for example during a break), which would lead to no notifications at all.

4 DISCUSSION

Our prototype system effectively showcases the feasibility of utilizing gaze data to establish a privacy-preserving notification system for integrating human and agent collaboration. The system exhibits high scalability of privacy and context awareness, permitting the inclusion of additional contextual or privacy information, leveraging the capabilities of Solid and its ecosystem. This scalability facilitates the creation of an exceptionally optimized system for human-agent collaboration, enhancing the capabilities of autonomous agents while safeguarding the privacy and workflow of the human user.

Having such a system in place allows for more solutions in the problem domain of better integrating human abilities into already existing agent systems while maximising both human and agent system efficiency, then just the motivating scenario of machine maintenance used as a demonstrator in this paper. The proposed system could find usage not only in domain of healthcare where it could be a support tool for healthcare workers providing services to patients while optimizing their daily schedule to become even more efficient, but also in the domain of sustainable living/working, nudging people to include climate-friendly actions into their habits and personal workflows.

Nonetheless, several aspects remain open for further exploration: First, the current integration of gaze data could be enhanced by adopting more nuanced criteria. Our implemented system is sensitive to even very short gaze duration; however, individual gaze patterns vary significantly from one person to another. This could be accounted for in an enhanced version of this system, where individual/personal factors are taken into account when determining whether a registered gaze point actually implies that the user has gazed at an object. A suitable way to include this aspect could be the integration of an ontology that describes the user-specific relationship between 3D gaze coordinates and semantic gaze relationships (i.e., does a 300ms fixation in an area of interest imply that the user has gazed at the respective object?).

Another avenue for improvement is the incorporation of gaze prediction techniques, such as Gazeformer [25]. This would not only enable the agent to interact more seamlessly with the human but also anticipate probable user actions within the context of the system presented in this paper. Third, considering the findings of Zijlstra et al. [36], which indicate that not all interruptions are harmful, there is an opportunity to individualize the interruption pattern and tailor it to the specific characteristics of the human user. This could enhance the overall adaptability and user-centricity of the notification system.

5 CONCLUSION

In this paper, we present a privacy-preserving context-aware opportunistic notification system for human-agent collaboration utilizing gaze data: Notification occurs only when the human is situated in a context where nudging the user is appropriate. The autonomy of the individual is further emphasized through the ability to prevent specific agents from accessing their data using the Solid protocol's authorization system. This approach facilitates the integration of human capabilities with autonomous software systems, mitigating the costs associated with task switching for the human. Moreover, it affords the individual complete control over their data, allowing for the specification of preferences. The system, therefore, serves as an effective collaborative tool, showcasing a harmonious balance between human agency and autonomous functionality. Our implemented prototype serves as a demonstration of the general principle of employing gaze-based, privacy-preserving, and context-aware methods using Linked Data and Semantic Web technologies to determine beneficial timing for interacting within human-agent workflows.

ACKNOWLEDGMENTS

To Adrian, Luka and Lukas for accompany me on this journey. ChatGPT v3.5 was used for copy editing.

REFERENCES

[1] Antti Ajanki, Mark Billinghurst, Hannes Gamper, Toni Järvenpää, Melih Kandemir, Samuel Kaski, Markus Koskela, Mikko Kurimo, Jorma Laaksonen, Kai Puolamäki, Teemu Ruokolainen, and Timo Tossavainen. 2011. An Augmented Reality Interface to Contextual Information. *Virtual Reality* 15, 2 (01 06 2011), 161–173. https://doi.org/10.1007/s10055-010-0183-5

[2] Kenan Bektaş, Jannis Strecker, Simon Mayer, Kimberly Garcia, Jonas Hermann, Kay Erik Jenß, Yasmine Sheila Antille, and Marc Solèr. 2023. GEAR: Gaze-enabled augmented reality for human activity recognition. In *Proceedings of the 2023 Symposium on Eye Tracking Research and Applications* (Tubingen, Germany) *(ETRA '23)*. Association for Computing Machinery, New York, NY, USA, Article 9, 9 pages. https://doi.org/10.1145/3588015.3588402

[3] Kenan Bektas. 2020. Toward A Pervasive Gaze-Contingent Assistance System: Attention and Context-Awareness in Augmented Reality. In *ACM Symposium on Eye Tracking Research and Applications* (Stuttgart, Germany) *(ETRA '20 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 36, 3 pages. https://doi.org/10.1145/3379157.3391657

[4] Kenan Bektaş, Adrian Pandjaitan, Jannis Strecker, and Simon Mayer. 2024. GlassBoARd: A Gaze-Enabled AR Interface for CollaborativeWork. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3613905.3650965

[5] Kenan Bektaş, Jannis Rene Strecker, Simon Mayer, and Markus Stolze. 2022. EToS-1: Eye Tracking on Shopfloors for User Engagement with Automation. In *AutomationXP22: Engaging with Automation, CHI'22*. CEUR Workshop Proceedings. http://www.alexandria.unisg.ch/266339/

[6] Kenan Bektaş, Jannis Strecker, Simon Mayer, and Kimberly Garcia. 2024. Gaze-enabled activity recognition for augmented reality feedback. *Computers & Graphics* (2024), 103909. https://doi.org/10.1016/j.cag.2024.103909

[7] Rachel K. E. Bellamy, Sean Andrist, Timothy Bickmore, Elizabeth F. Churchill, and Thomas Erickson. 2017. Human-Agent Collaboration: Can an Agent be a Partner?. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI EA '17)*. Association for Computing Machinery, New York, NY, USA, 1289–1294. https://doi.org/10.1145/3027063.3051138

[8] Tim Berners-Lee, Dan Connolly, and Sandro Hawke. 2003. Semantic web tutorial using n3. In *Twelfth International World Wide Web Conference*.

[9] Tim Berners-Lee, Dan Connolly, Lalana Kagal, Yosi Scharf, and Jim Hendler. 2007. N3Logic: A Logical Framework For the World Wide Web. *CoRR* abs/0711.1533 (2007). arXiv:0711.1533 http://arxiv.org/abs/0711.1533

[10] Christian Bizer, Tom Heath, and Tim Berners-Lee. 2011. Linked data: The story so far. In *Semantic services, interoperability and web applications: emerging concepts*. IGI global, 205–227.

[11] Olivier Boissier, Rafael H. Bordini, Jomi F. Hübner, and Alessandro Ricci. 2019. Dimensions in programming multi-agent systems. *The Knowledge Engineering Review* 34 (2019), e2. https://doi.org/10.1017/S026988891800005X

[12] Efe Bozkir, Ali Burak Ünal, Mete Akgün, Enkelejda Kasneci, and Nico Pfeifer. 2020. Privacy Preserving Gaze Estimation using Synthetic Images via a Randomized Encoding Based Framework. In *ACM Symposium on Eye Tracking Research and Applications* (Stuttgart, Germany) *(ETRA '20 Short Papers)*. Association for Computing Machinery, New York, NY, USA, Article 21, 5 pages. https://doi.org/10.1145/3379156.3391364

[13] Andrei Ciortea, Simon Mayer, Fabien Gandon, Olivier Boissier, Alessandro Ricci, and Antoine Zimmermann. 2019. A Decade in Hindsight: The Missing Bridge Between Multi-Agent Systems and the World Wide Web. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems* (Montreal QC, Canada) *(AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1659–1663.

[14] Joel E. Fischer, Chris Greenhalgh, and Steve Benford. 2011. Investigating episodes of mobile phone activity as indicators of opportune moments to deliver notifications. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services* (Stockholm, Sweden) *(MobileHCI '11)*. Association for Computing Machinery, New York, NY, USA, 181–190. https://doi.org/10.1145/2037373.2037402

[15] Celia Garrido-Hidalgo, Diego Hortelano, Luis Roda-Sanchez, Teresa Olivares, M. Carmen Ruiz, and Vicente Lopez. 2018. IoT Heterogeneous Mesh Network Deployment for Human-in-the-Loop Challenges Towards a Social and Sustainable Industry 4.0. *IEEE Access* 6 (2018), 28417–28437. https://doi.org/10.1109/ACCESS.2018.2836677

[16] Hochstenbach, Patrick and Verborgh, Ruben and Van de Sompel, Herbert. 2023. Using event notifications, solid and orchestration for decentralizing and decoupling scholarly communication. *CODE4LIB JOURNAL* 58, Article 17823 (2023). https://journal.code4lib.org/articles/17823

[17] Sebastian Hueber, Eunae Jang, and Jan Borchers. 2023. Attentive Notifications: Minimizing Distractions of Mobile Notifications through Gaze Tracking. In *Proceedings of the 25th International Conference on Mobile Human-Computer Interaction* (Athens, Greece) *(MobileHCI '23 Companion)*. Association for Computing Machinery, New York, NY, USA, Article 1, 7 pages. https://doi.org/10.1145/3565066.3608695

[18] Sebastian Kaebisch, Takeshi Kamiya, Michael McCool, Valentina Charpenay, and Matthias Kovatsch. 2020. *Web of Things (WoT) Thing Description*. Technical Report. World Wide Web Consortium. http://www.w3.org/TR/2020/REC-wot-thing-description-20200409/

[19] Melih Kandemir and Samuel Kaski. 2012. Learning Relevance from Natural Eye Movements in Pervasive Interfaces. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction* (Santa Monica, California, USA) *(ICMI '12)*. Association for Computing Machinery, New York, NY, USA, 85–92. https://doi.org/10.1145/2388676.2388700

[20] Sebastian Kapp, Michael Barz, Sergey Mukhametov, Daniel Sonntag, and Jochen Kuhn. 2021. ARETT: Augmented Reality Eye Tracking Toolkit for Head Mounted Displays. *Sensors* 21, 6 (2021). https://doi.org/10.3390/s21062234

[21] Michaela Klauck, Yusuke Sugano, and Andreas Bulling. 2017. Noticeable or Distractive? A Design Space for Gaze-Contingent User Interface Notifications. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI EA '17)*. Association for Computing Machinery, New York, NY, USA, 1779–1786. https://doi.org/10.1145/3027063.3053085

[22] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. *What Does Your Gaze Reveal About You? On the Privacy Implications of Eye Tracking*. Springer International Publishing, Cham, 226–241. https://doi.org/10.1007/978-3-030-42504-3_15

[23] Baptist Liefooghe, Pierre Barrouillet, André Vandierendonck, and Valérie Camos. 2008. Working memory costs of task switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 34, 3 (2008), 478.

[24] Gloria Mark, Daniela Gudith, and Ulrich Klocke. 2008. The Cost of Interrupted Work: More Speed and Stress. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) *(CHI '08)*. Association for Computing Machinery, New York, NY, USA, 107–110. https://doi.org/10.1145/1357054.1357072

[25] Sounak Mondal, Zhibo Yang, Seoyoung Ahn, Dimitris Samaras, Gregory Zelinsky, and Minh Hoai. 2023. Gazeformer: Scalable, Effective and Fast Prediction of Goal-Directed Human Attention. arXiv:2303.15274 [cs.CV]

[26] Stephen Monsell. 2003. Task switching. *Trends in cognitive sciences* 7, 3 (2003), 134–140.

[27] Brid O'Conaill and David Frohlich. 1995. Timespace in the workplace: Dealing with interruptions. In *Conference companion on Human factors in computing systems*. 262–263.

[28] Pernilla Qvarfordt and Shumin Zhai. 2005. Conversing with the user based on eye-gaze patterns. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, Oregon, USA) *(CHI '05)*. Association for Computing Machinery, New York, NY, USA, 221–230. https://doi.org/10.1145/1054972.1055004

[29] Andrei Vlad Sambra, Essam Mansour, Sandro Hawke, Maged Zereba, Nicola Greco, Abdurrahman Ghanem, Dmitri Zagidulin, Ashraf Aboulnaga, and Tim Berners-Lee. 2016. Solid: a platform for decentralized social applications based on linked data. *MIT CSAIL & Qatar Computing Research Institute, Tech. Rep.* (2016).

[30] Maureen Schmitter-Edgecombe and Chad Sanders. 2009. Task switching in mild cognitive impairment: switch and nonswitch costs. *Journal of the International Neuropsychological Society* 15, 1 (2009), 103–111.

[31] Christopher J. Turner, Ruidong Ma, Jingyu Chen, and John Oyekan. 2021. Human in the Loop: Industry 4.0 Technologies and Scenarios for Worker Mediation of Automated Manufacturing. *IEEE Access* 9 (2021), 103950–103966. https://doi.org/10.1109/ACCESS.2021.3099311

[32] Ruben Verborgh and Jos De Roo. 2015. Drawing Conclusions from Linked Data on the Web: The EYE Reasoner. *IEEE Software* 32, 3 (2015), 23–27. https://doi.org/10.1109/MS.2015.63

[33] World Wide Web Consortium (W3C). n.d.. *Web of Things (WoT) Developers*. https://www.w3.org/WoT/developers/#td-tooling Accessed: 2024-01-23.

[34] Glenn Wylie and Alan Allport. 2000. Task switching and the measurement of "switch costs". *Psychological research* 63 (2000), 212–233.

[35] Bianca Zickerick, Sven Thönes, S Oliver Kobald, Edmund Wascher, Daniel Schneider, and Kristina Küper. 2020. Differential effects of interruptions and distractions on working memory processes in an ERP study. *Frontiers in Human Neuroscience* 14 (2020), 84.

[36] Fred RH Zijlstra, Robert A Roe, Anna B Leonora, and Irene Krediet. 1999. Temporal factors in mental work: Effects of interrupted activities. *Journal of Occupational and Organizational Psychology* 72, 2 (1999), 163–185.

[37] Antoine Zimmermann, Andrei Ciortea, Catherine Faron, Eoin O'Neill, and María Poveda-Villalón. 2023. Pody: a Solid-based approach to embody agents in Web-based Multi-Agent-Systems. In *International Workshop on Engineering Multi-Agent Systems*. Springer, 220–229.