

The Internet of Us

The European Journal on Artificial
Intelligence
1–12
© The Author(s) 2026
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/30504554251410835
journals.sagepub.com/home/eai



Loizos Michael^{1,2}, Ivano Bison³ , Matteo Busso³,
Luca Cernuzzi⁴ , Amalia de Götzen⁵, Shyam Diwakar⁶ ,
Kobi Gal⁷ , Amarsanaa Ganbold⁸, George Gaskell⁹,
Daniel Gatica-Perez¹⁰, Jessica Heesen¹¹, Daniele Miorandi¹²,
Salvador Ruiz-Correa¹³, Laura Schelenz⁷, Avi Segal¹¹,
Carles Sierra¹⁴, Hao Xu¹⁵ and Fausto Giunchiglia³ 

Abstract

Social Media and the Internet have catalyzed an unprecedented potential for exposure to human diversity in terms of demographics, language, culture, knowledge, opinions, talents and the like. Access to people's diversity gives us the possibility of exploiting skills and competences that we do not have, that we may not even know they exist, some the so-called *unknown unknowns*, but which, however, could be exactly what we need when *looking for help* in the solution to the our current problem. However, this potential has not come with new, much needed, instruments and skills to harness the complications which rise when trying to exploit the diversity of people. This paper presents our vision of the “*Internet of Us (IOU)*,” a new type of online *diversity-aware social interactions* capable of promoting richer and deeper social relations. We discuss the multiple facets of diversity in social settings as well as the multidisciplinary work that is required to reap the benefits of the IOU, toward a IOU-enabled diversity-aware hybrid human–AI society.

Keywords

The internet of Us, WeNet, diversity, social interactions, AI mediation

I A Vision for the Internet of Us

No person is an island. No one knows or is skilled in everything. Human societies empower us to reach out to people with skills and information complementary to ours; we rely on others, and they rely on us. This diversity is something we

¹Open University of Cyprus, Latsia, Cyprus

²CYENS Center of Excellence, Nicosia, Cyprus

³University of Trento, Italy

⁴Universidad Católica “Nuestra Señora de la Asunción”, Asuncion, Paraguay

⁵Aalborg University, Copenhagen, Denmark

⁶Amrita Vishwa Vidyapeetham, Kollam, Kerala, India

⁷Ben-Gurion University, Beer-Sheva, Israel

⁸National University of Mongolia, Ulaanbaatar, Mongolia

⁹London School of Economics and Political Science, London, UK

¹⁰Idiap Research Institute and EPFL, Switzerland

¹¹University of Tuebingen, Tuebingen, Germany

¹²U-Hopper, Trento, Italy

¹³IIIA-CSIC, Barcelona, Spain

¹⁴IPICYT, San Luis Potos, Mexico

¹⁵Jilin University, Changchun, China

This is revised version of the ArXiv paper (Michael et al., 2025).

Corresponding Author:

Fausto Giunchiglia, University of Trento, Italy.

Email: fausto.giunchiglia@unitn.it

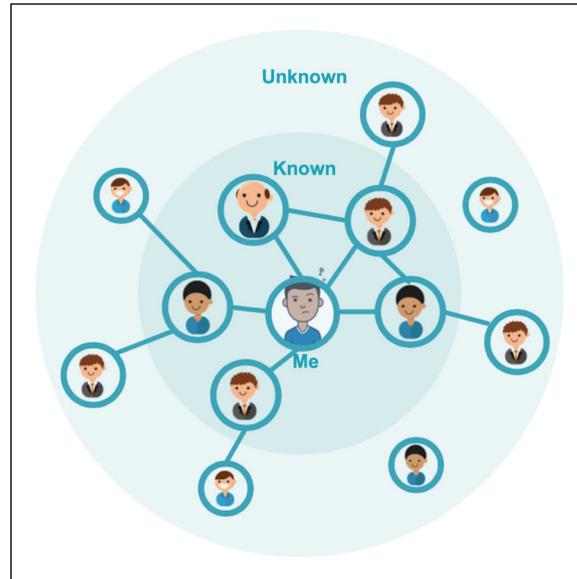


Figure 1. I o U - Us, reaching out to possibly unknown people.

exploit in our everyday life, often without even realizing it. We ask a doctor for a diagnosis and treatment when we are ill, we call up a plumber if a pipe leaks in our house, and look for someone who speaks our language and Chinese if we need to arrange a stay in Beijing. Embracing and exploiting diversity can help us improve. *Diversity*—our “differences” in language (Bella et al., 2024; Giunchiglia et al., 2017a, 2023; Koch et al., 2024), knowledge (Achón et al., 2024; Giunchiglia & Bagchi, 2022; Giunchiglia & Fumagalli, 2020), routines, social practices, personality traits, competencies, skills (Girardini et al., 2023; Hume et al., 2022; Mercado et al., 2023; Schelenz et al., 2021), and so on—is a *desired feature that should be treasured and sought after*. Over the past decades, the Internet, globalization, and the emergence of global digital platforms have transformed our lives and transcended geographical and cultural borders. Thanks to the Internet, the Web and social media, it is easy to connect to anybody in the world, regardless of their culture, social context, norms, habits, or language. The amount of diversity we have access to has increased exponentially. We get exposed daily to a seemingly unbound amount of diversity, most of which is unexpected and represents the “*unknown unknowns*.” That is, an essentially unbound source of knowledge, that we all can potentially exploit. Yet, evolution operates on a slower timescale than that of technology. We have basically the same instruments and skills that our grandparents had 50 years ago. Technology has provided us with an increased access to diversity, but has failed to provide us with the instruments for individuals and communities to cope with the social challenges that arise with diversity. As social interactions increasingly shift from the physical to the virtual world, where “the same” is reinforced via biased algorithmic recommendations (Orphanou et al., 2022), this *diversity in information access* is inevitably restricted (Cinelli et al., 2021), and often *viewed as an undesirable bug to be avoided and eliminated*, or delegated to opaque systems, for example, GenAI systems, with all the obvious risks, see, for example, (Jo, 2023; Mueller, 2024).

Not unlike the *Internet of Things*, this paper puts forward the vision of an “*Internet of Us* (I o U),” a hybrid human–AI enabled social interaction modality where diverse communities and individuals can co-exist and interact (Osman et al., 2021, 2020). The I o U will enable social interactions that reach beyond the limited social sphere of individuals, thus empowering them to explore more fully the *unknown unknowns* of their social search space, and help them capitalize on the range of their diverse collective interests and competencies. The situation is well represented in Figure 1. The inner circle represents the type of interactions enabled by mainstream social platforms, as extensively studied in the work on EGO-networks, see for example, Arnaboldi et al. (2017). Instead, the outer circle represents the possibility enabled by the I o U of reaching out to people beyond one’s EGO-network. The key intuition, as also represented by the uniformity of the links in Figure 1, is that the social interactions between people are independent of whether a person is inside or outside one’s EGO-network. What matters is the type of diversity and the consequent type of help that it enables. To draw a parallel, the service that we need from a plumber is independent of whether we know the person. The level of acquaintance may influence the level of trust of the person seeking the service and the motivation of the person providing the service, but not the expected service. At the same time, trust and motivation are two key elements in this type of social relation and, as a consequence, also at the core of the design of the I o U.

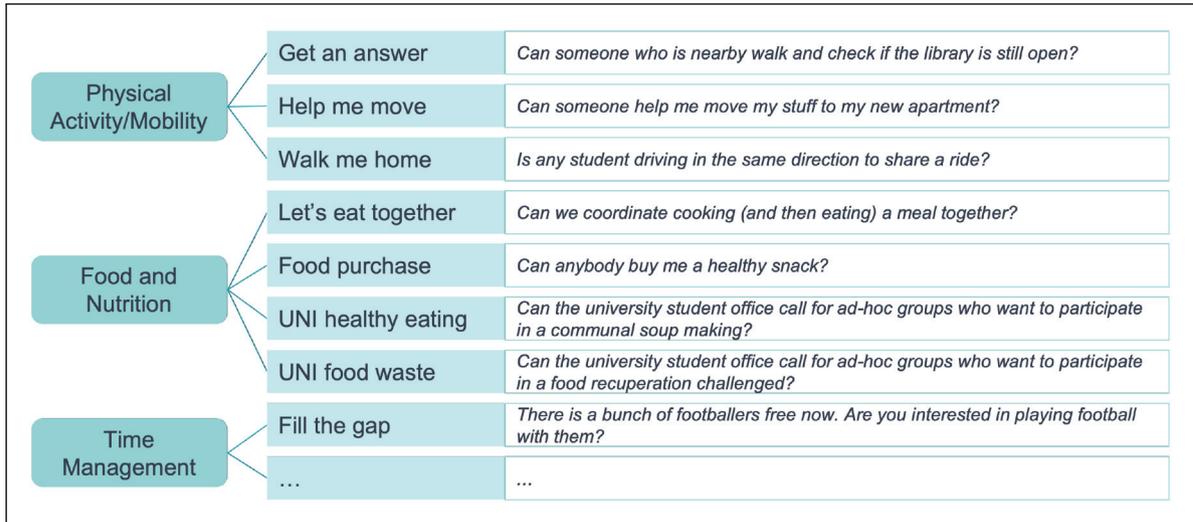


Figure 2. IoU - asking possibly unknown people for help, about an unknown solution to a possibly unknown problem.

The underlying idea is that IoU will allow each and everyone to have access to the knowledge, skills and competences of the person, most of the time unknown, most suitable for dealing with the current problem. The current situation where one types a query to a search engine or, most recently, to his/her favorite GenAI application, for example, Gemini and ChatGPT, with no human intermediation and no or little knowledge of the current context, is expected to be the minimal level of service provided by the IoU. At the same time, the IoU is expected to support a much broader spectrum of more person-centric solutions, where the upper bound is a fully context-aware interaction (in presence or online), *here* and *now*, between the person in need of help and the person providing such assistance. Precompiled online queries will become the default solution when no human help is available. The ultimate goal is to make it possible for everybody “to exploit, when trying to achieve his/her objectives, the best possible expertise available on the planet. Each person would be augmented with all the knowledge available in the world” (Giunchiglia, 2017). Figure 2 reports just some examples of possible services enabled by the IoU.

The remainder of this paper is organized as follows. First, we provide an overview of the main functionalities that the IoU should offer, illustrated through a motivating example (Section 2). Next, we present a detailed description of our current understanding of these functionalities based on the work developed thus far. We begin by sharing the results of a series of experiments conducted in real-world settings.¹ The goal of these experiments is to explore crucial aspects that can be used to model the diversity of individuals (Section 3, which discusses the results from the diversity pilots) as well as the needs of those seeking support (Section 4, which reports findings from the application pilots). Following this, we discuss the key mechanisms that facilitate and encourage the formation of communities (Section 5). We then address issues related to modeling and understanding the personal and social contexts of individuals (Section 6). Finally, we offer some reflections on how ethical considerations can play a vital role in establishing the “Internet of Us” (Section 7).

2 Designing the Internet of Us

A platform has been developed which implements a first version of the IoU. This platform has been used and tested in several pilot studies (Busso et al., 2025; Giunchiglia et al., 2022; Kun et al., 2022), engaging with thousands of students from across the world² over the span of four years, and exploring the benefits that a diversity-aware chat application could bring to people, and how people interpret diversity when being presented with it in social interaction settings. Our empirical investigation utilized two different typologies of pilots. The so-called “*diversity pilots*” sought to collect data about the students’ characteristics and behaviors through online surveys and mobile phone sensors towards refining the model of diversity implemented in the platform. They were carried out using the *iLog* APP tuned to the need of the IoU (Kayongo et al., 2025; Zeni et al., 2014). The so-called “*application pilots*” sought to explore the needs and interactions of the students in real-life scenarios, leveraging their diversity through the integration in the platform of research-driven technological solutions. The data collected through the diversity pilots was used to develop and test measurement instruments on human behavior, and to study diversity from a sociological perspective, while also informing the structure of participant profiles in the platform. At the same time, the specific information provided by the users was used to populate

their profiles when they participated in the application pilots and to support research on developing machine learning and interaction algorithms that utilize diversity in the behavioral routines of users.

The current prototype of the platform is highly modular, consisting of various components that address specific issues related to the problem solution. We provide below a short description of the different dimensions of the problem, and the components solving them, via a motivating example.

A relevant, somewhat under-specified problem. Jamal is a second-year undergraduate at the London School of Economics, curious about how artificial intelligence (AI) is shaping learning and society across cultures. Feeling that one's local networks sometimes reinforce the same perspectives, Jamal joins the IOU Q/A APP, an application designed to connect students worldwide by mediating diversity-aware interactions according to their interests. He hopes the app will not only provide answers but also broaden his horizons, exposing him to perspectives he did not know he was missing.

Exploring the Unknown. One evening, Jamal posts a question in IOU Q/A APP: "How do students in other countries learn and debate about AI in class?" Within minutes, replies begin to arrive: a student in Beijing describes technical labs where group work structure the learning process; a peer in Delhi emphasizes the debates about AI ethics taking place in student unions; and another student in Mexico City explains how community groups discuss the social impacts of AI, particularly in relation to the evolving skills needed in the workforce. Jamal is surprised by how different these practices are compared to his own seminars in London. In those discussions, the focus often revolves around the differences between natural and artificial intelligence, as well as the issue of control and regulation, which Jamal has always found to be quite Orwellian. This confirms the APP's promise: by connecting him to peers he would otherwise never meet, Jamal explores the *unknown unknowns*, perspectives and routines (see Figure 1), far outside his immediate social sphere (Cinelli et al., 2021; Giunchiglia & Bagchi, 2022).

Staying Motivated through Incentives. At first, Jamal mostly asks questions, but he soon notices that the platform rewards him for answering as well. After explaining how UK classes often focus on critical debate about AI regulation rather than purely technical training, he receives a digital badge recognizing his contribution. The recognition feels meaningful: not only has he shared something of value, but he also learns that incentives are designed to encourage active participation, both in giving and receiving support. This motivates him to contribute more frequently, echoing research showing that diversity-aware incentive mechanisms (such as badges and personalized recommendations) can sustain engagement without imposing a one-size-fits-all approach (Botta et al., 2023a, 2023b; Yanovsky et al., 2019).

Experiencing Interaction Norms. Over time, Jamal realizes that the app gently guides his behavior. When he tries to direct multiple follow-up questions to Maria from Mexico, the APP reminds him that he has already contacted her several times this week and suggests reaching out to someone new. At first, he is surprised by this restriction, but he comes to appreciate it: Maria is not overwhelmed, and Jamal's own network expands as he meets more students. Here, he sees norms in action - rules of engagement built into the system that reflect community values of fairness and reciprocity (García-Camino et al., 2009; Osman et al., 2021, 2020). For Jamal, this makes the platform feel more like a self-regulating community than a random chat service.

Adapting to the Situational Context. As Jamal becomes more active, he notices that the APP adapts to his context, recognizing the different information and data that compose it (see e.g., Busso et al., 2025). When he logs in from the library late at night, his queries are routed to peers in different time zones who are more likely to be online and responsive. During the day, however, the app prioritizes responses from nearby universities in Europe. This subtle adjustment helps Jamal always find engaged interlocutors, illustrating the importance of situational awareness in AI systems, and showing how context-aware models need localized data, since habits and schedules differ across countries (Assi et al., 2023). Without such sensitivity, the system would risk misaligning support and leaving questions unanswered.

Building the Conversational Context. As weeks pass, Jamal notices that the app not only adapts to his situation but also remembers his past interactions. It suggests reconnecting with Priya from India, who once offered him a nuanced take on fairness in machine learning, while also introducing him to new voices. This balance between continuity and novelty makes his conversations richer: he can build trust with familiar peers while still being challenged by new perspectives. In doing so, the app draws on both tie strength estimation (Perikos & Michael, 2022) and argumentation-based policy frameworks (Dietz et al., 2022; Markos et al., 2022; Michael, 2019, 2023a; Tubella et al., 2020), ensuring that recommendations are not only data-driven but also aligned with norms of diversity and fairness.

Discovering Diversity in Practice. When Jamal posts a question about "AI in healthcare", he is amazed by the range of responses. The app deliberately diversifies who answers: younger and older students, men and women, but also - and perhaps more importantly - students with very different competencies. A medical student describes clinical uses, an engineer explains algorithmic design, and a philosophy student reflects on ethics. Jamal realizes that while visible diversity (age, gender) matters, invisible diversity (skills, values, disciplinary knowledge) often adds even greater value to his research project. This reflects the IOU's approach of allowing communities to decide which diversity dimensions to prioritize (Harrison et al., 1998; Markos & Michael, 2022a; Tsui et al., 1992).

Balancing Diversity with Ethics. Eventually, Jamal feels ready to ask about a sensitive issue: “bias and discrimination in AI systems”. He hesitates at first, worried about receiving hostile responses, but the app allows him to post anonymously and restrict his question to a trusted pool of peers. The responses are constructive and respectful, and Jamal realizes how crucial these safeguards are. IoU’s design treats diversity as an instrumental value (Zimmerman & Bradley, 2019): it should foster inclusion and learning, but never at the expense of protection. By balancing exposure with safety, the app enacts the tension between inclusion and security that is central to ethical debates about diversity in technology (Goffman, 2009; Helm et al., 2022; Schelenz et al., 2024). For Jamal, this ethical mediation turns a potentially risky question into a meaningful dialogue.

The solution to an unknown problem. By the semester’s end, Jamal has collected enough material to write his paper on “AI-Mediated Dialogue Across Cultures.” He even organizes a cross-country student group on IoU Q/A APP to continue discussing AI ethics, applications, and governance. Reflecting on his experience, Jamal sees that each feature of the app—exploration, incentives, norms, situational and conversational context, diversity, and ethics—was not isolated, but worked together to create a safe and stimulating environment. His story shows how AI can mediate cross-cultural learning by carefully weaving together motivation, fairness, personalization, and protection into everyday social interactions.

3 Understanding and Modeling Human Diversity

A key prerequisite for achieving the goal of our platform is the ability to capture diversity within a community and to represent it in a machine-readable fashion. We have identified two main classes of diversity dimensions across which diversity is measured (Harrison et al., 1998; Tsui et al., 1991, 1992):

- *visible* (or “observable”) demographic characteristics, for example, sex and gender, skin color/ethnicity, physical (dis)abilities, and age; and
- *invisible* (or “deep”) characteristics, for example, ethnic/cultural/socioeconomic background, personality characteristics, cognitive and interaction attitudes, social/human values, competencies and education, type of work organization, and role.

The complete list of diversity dimensions, along with descriptions of the experiments conducted to gather this data and the collected datasets, is provided in Busso et al. (2025), Giunchiglia et al. (2022). Among the many participant dimensions that we have modeled, a few key invisible characteristics can be grouped into three abstract dimensions, usually considered as part of the general theory of the *Communities of Practice: material, competence, and meaning* (Røpke, 2009; Shove & Pantzar, 2005). Material relates to the tangible assets that an individual has access to. Competencies incorporate an individual’s endowed skills, (background) knowledge, as well as social and relational skills required to perform a practice. Meaning incorporates understandings, beliefs, values, lifestyle, emotions, and social and symbolic significances, which are not transferable but can be learned through socialization (Berger & Luckmann, 2016). Meaning enables the connection of material and competencies, giving rise to “behavioral routine” actions that become practices at a social level. It also acts as a “nexus,” not only within a single practice, but in linking different practices, and even in creating new ones.

Both visible and invisible characteristics are operationalized in the profile of each community member. The values associated with these characteristics position each individual within a multi-dimensional space, which can be explored to establish social relationships among community members. By imposing a distance metric on this space, we can formalize the notion of variability between individuals.

However, when developing specific applications for certain communities on the platform, some characteristics may become less relevant. For instance, a community centered on the ethical implications of AI may consider “material” aspects, such as owning a state-of-the-art computer, to be less critical. In contrast, a group focused on developing new algorithms for human-aware AI will view this as essential. The general approach is that the platform supports all the diversity dimensions we have identified, allowing the various applications built on top of it to decide which dimensions to utilize. The study referenced in Meegahapola et al. (2023) describes initial research on the explanatory power of the diversity dimensions we have chosen to model.

The richness of the profiles comes from a descriptive view of diversity, which captures the variability within a community (or a subset chosen to receive a request), but also between the person seeking support and each potential respondent to that request. Being able to compute degrees for these two types of diversity is a step towards making the members of a community aware of the diversity that exists within their ranks and files. A community member seeking support can then explicitly make use of this awareness by identifying which dimensions of the search space to restrict, and in what way, so that the exploration of the search space can become feasible, but not overly restrictive.

4 Understanding the Needs and Constraints of Humans Seeking Support

To study the perception of users when presented with diversity, we developed a technology probe (Hutchinson et al., 2003), in the form of a QA chat APP built on top of the IoU platform, with a lightweight interface (avoiding the technical complexity of a full-fledged conversational agent with natural language processing capabilities) that connects people in a community. In the application pilots, users could submit a support query in written text. They had the option to select the domain of their query, such as hobbies, arts, career, or studies. Additionally, users could specify who should be able to view their query and indicate whether the query was sensitive in nature, in which case it would be submitted anonymously. This information was then used to create a ranking of users, who were subsequently invited to respond to the query.

The chat app was evaluated favorably in the pilots. It was appreciated for its ability to connect students from different departments, educational and cultural backgrounds, and with various skills and competencies. While the possibility of getting support from “expert users” was an acknowledged feature of the app, a major attraction for some users was contacting other students, making social connections, and creating the feeling of a community while exploring its diversity. Most of the expectations towards diversity carried an exploratory quality, where people were driven by curiosity or wished for serendipity in posing queries to people different from themselves. This need was particularly apparent in the pilots conducted during the COVID-19 pandemic. Exit surveys and focus groups conducted with students after the most recent pilot programs reinforced the finding that many users were eager to leverage diversity. They indicated that giving advice was slightly more satisfying than receiving it. Users appreciated the absence of photographs and valued the opportunity to participate anonymously. Additionally, marking a question as sensitive was seen positively and reflected the ethical principles that guided the design of the chat app.

Beyond the aforementioned pilot studies and our developed QA chat APP, the platform was also utilized by several external third parties who were interested in developing alternative use cases using our technology. One such external third party was the Spanish NGO *Fundación Cibervoluntarios*, which coordinates a network of more than 1,800 cybervolunteers helping citizens to level out digital gaps. In March 2020, during the lockdown caused by the COVID-19 pandemic, *Fundación Cibervoluntarios* set up a helpline to provide support to citizens having doubts or problems related to technology, computing devices, and their connectivity, configuration, usage, or cybersecurity. *Fundación Cibervoluntarios* piloted our technology to empower cybervolunteers to effectively help out citizens in difficulty. Another external third party was the Greek social cooperative enterprise CommonsLab, which leveraged our technology to implement a chat application (MaTSE: a Matchmaking Tool for Social Entrepreneurs) that brings together different constituencies in the social innovation landscape: entrepreneurs with a business idea for a social innovation venture, potential investors, and other actors who may help in the process (e.g., established entrepreneurs who could provide early feedback, etc.).

As part of our effort to utilize (and evaluate) our technology in diverse contexts, we have been involved in two additional studies. The first study sought to elicit the perspective of Afghan refugees in Germany on the QA chat APP. Afghan refugees are particularly marginalized technology users because of their experiences of political, economic, and social violence in Afghanistan and abroad. The study spanned four months in 2022 and was conducted in cooperation with an NGO for the rights of Afghan women. In conversations with 14 Afghan refugee women, it became clear that empowerment through technology is not inevitable. The QA chat APP presented significant barriers for their participation due to the “complicated” registration process and usage commands.

The second study aimed to utilize our technology to alleviate the diminished social interactions between the Greek-Cypriot and Turkish-Cypriot communities of Cyprus, particularly those affected by the language barrier. A simplified version of the QA chat APP was extended with the extra functionality of allowing each user to interact (with the APP and others) in their native language, while supporting the automatic translation of messages exchanged between the two communities. The effort itself has been met with interest from members of both communities. Still, it has also highlighted the need to be wary of unforeseen pitfalls, such as the potential of mistranslation that may be deemed offensive by the recipient of the message. The two studies have highlighted that an *Internet of Us* must consider diversity in digital literacy and language skills, and ensure that the app design is inclusive of low-literate and socio-economically marginalized groups, and non-homogeneous groups that might not exhibit a shared trust.

5 Supporting and Incentivizing the Self-Determination of Communities

Awareness of descriptive diversity and the dimensions along which it is measured supports the self-determination of communities by providing the primitive concepts against which a community can differentiate itself from other communities. A community could, for instance, self-determine as including only residents of Spain who are wine aficionados, effectively utilizing the “country of residence” and the “hobbies and leisure” features of its members as the properties that make it distinct from other communities.

The same primitive concepts support a second path to *community self-determination*, by providing the basis over which a community can define the *norms* that apply during the interaction between its members. Sociologically speaking, norms capture the rules of engagement, or the social rules of behavior of a community and its members, and are the direct application of the social values of that community. They capture community-wide provisions (e.g., do not ask the same person for support more than three times within any single week) or individual provisions (e.g., do not share a request for support with members blocked by an individual). Norms are also meant to interpret the search directives provided by members of a community. The same cue (e.g., seeking support from people nearby) can be interpreted differently if seeking support within one’s “neighborhood watch” community, or within one’s “alma mater graduates” community. Technologically speaking, norms define the protocol of exploring the search space when seeking to solve the search problem of identifying appropriate respondents to a sequence of queries.

Norms come in the form of natural language expressions in human societies, and, traditionally, in the form of logic-based expressions in formal machine-readable representations (e.g., constraint-based García-Camino et al., 2009, time-based Ågotnes et al., 2009(@, or event-based Kowalski & Sergot, 2005). In our work, we have adopted a declarative representation of norms through conditional logic-based rules (Osman et al., 2021). The rule premises refer to profile features (such as gender, competencies, etc.) or other state properties of the platform (such as the number of active community members, the total number of queries, the level of user satisfaction, etc.) that can encode information about a community, a user, or even a task or interaction. The rule conclusions refer to the possible actions that can be performed, which include message-passing actions (such as sending a message to another user, or sending a message to an app’s user-facing interface to enable a given button), and actions that change profile features or state properties (such as the privacy setting of one’s location, or the rating for a given task). A norm engine is then designed to mediate interactions by checking their compliance with the relevant norms in place, be they community-wide norms or individual norms.

Of course, having community members aligned in their understanding of norms is crucial, but difficult in open systems. For instance, if a norm prohibits hate speech, not all members might have the same understanding of what constitutes hate speech. We have developed learning mechanisms that can learn the meaning of a norm from a community’s interactions, explain this meaning to community members, and adapt over time based on members’ feedback (dos Santos et al., 2022). Furthermore, working with norms allows us to formally analyze the outcomes of interactions (Montes et al., 2022). For instance, we can analyze whether certain norms promote/demote certain human values (Sierra et al., 2021), as well as synthesize norms to ensure the maximum promotion of given human values (Montes & Sierra, 2021).

How can we promote the growth of active, inclusive, and diverse communities, taking also into account that the platform allows for interactions with people who are *unknown unknowns*? In the spirit of community self-determination, the key consideration is that communities may not wish for external guidelines on inclusiveness and active participation to be applied in a paternalizing and all-knowing manner. Consequently, instead of imposing such guidelines, our platform seeks to incentivize their adoption, taking into account the diverse characteristics of different participants and communities, and avoiding a “one-size-fits-all” approach common in current systems. Past works have recognized user diversity and harnessed it for generating incentives using badges (Yanovsky et al., 2019), text messages (Segal et al., 2018), and recommended items (Ben Zaken et al., 2022). Recently, we proposed a diversity-aware multi-armed bandit approach (Botta et al., 2023a, 2023b), which reasons about users’ diversity characteristics when making action recommendations. Such reasoning was shown to lead to improved outcomes compared to a “static” approach, which ignored participants’ rich and changing profiles (cf. Section 3).

The platform’s incentivization mechanism directs participants to engage with others by seeking support, responding to support queries, and rating the responses they receive, towards becoming more active members of their community. This same mechanism could be used to incentivize participants not to apply search filters unnecessarily, and to actively request that the recipients of their query be diversified (cf. Section 6).

6 Modeling the Personal and the Social Context of Individual Participants

One of the aspects that makes individuals diverse, and that affects their social interactions through the application of norms, is their personal context (Bontempelli et al., 2022; Giunchiglia et al., 2017b). Due to its inherently dynamic nature, the personal context of any given platform participant cannot be populated in their profile when they first join the platform, but needs to be learned. The worldwide deployment of diversity pilots over our platform presented a unique opportunity to investigate how personal contexts can be learned across diverse communities. Short surveys about the participants’ current place, activities, social situation, and mood were collected in parallel with data from smartphone sensors and mobile APP logs. The datasets were used to investigate how diversity attributes, such as the country of residence of participants, can be adopted to build ML-based inference models from mobile data. Such an investigation allows for the systematic comparison of country-specific and country-agnostic models for tasks like activity recognition (Assi et al., 2023) or mood

recognition (Meegahapola et al., 2023), and facilitates the empirical evaluation of whether model transfer is possible across communities.

The results indicate that considering country-level diversity in model learning is crucial for recognition performance, and transferring models between countries is not a simple task. In this sense, the results contribute towards highlighting the fundamental need for locally-valid data that represents the reality of the world beyond economically advantaged regions (Gatica-Perez et al., 2019), and for documenting the resulting models with respect to these issues (Gebru et al., 2021). More broadly, our empirical results emphasize the need for machine learning models to be aware of the diversity that exists between communities, and by extension, the need for this diversity to be reflected in the learning data that is available (in the participant profiles, in the case of our work).

Beyond their personal context, a profile also includes an individual's social context, which in its simplest form comprises a list of learnable social tie strengths (Perikos & Michael, 2022) that capture the frequency and manner in which the individual has interacted with other platform participants. A more advanced form of social context comes through a set of norms that specify an individual's policy on whether / how to interact with other participants. This policy can be thought to be more persistent than an individual's social tie strengths, but can still be revised over time. Our work has focused on developing a formal declarative policy-representation language (Markos & Michael, 2022b) that acknowledges the key role of argumentation as a calculus for human-centric AI (Dietz et al., 2022). A policy elicitation paradigm accompanies the language (Michael, 2019, 2023a) that is cognitively light and computationally efficient (Markos et al., 2022). To support the utilization of this formal policy language, we have investigated how policies can be elicited through natural language cues (Ioannou & Michael, 2021), and how they can be integrated with machine learning models (Tsamoura et al., 2021), towards making them explainable (Michael, 2023b) and contestable (Tubella et al., 2020).

The use of individual norms to condition a user's interactions on their personal and social context, and the dependence of the user's context on their interactions, may lead to a vicious spiral of diversity-reducing interactions. To compensate for this possibility, we adopt a second view of diversity as a prescriptive value to be promoted, and the platform diversifies the ranking of participants that receive any given query. One might think that diversification should be applied maximally on all visible characteristics, but not on any of the invisible ones, which relate to the abilities of participants to respond to queries and should, therefore, be allowed to be constrained. This view is understandable, given that human societies have practically equated the notion of diversity with that of inclusiveness based on visible characteristics. But this, by itself, does not make this a valid view.

In certain scenarios, one may wish to select participants based on their visible characteristics, to target, for instance, a particular minority demographic group. In other scenarios, one may want to diversify based on the participants' invisible characteristics, not necessarily as a way to avoid discrimination, but as a way to get a more varied perspective (or serendipity, cf. Section 4). In the spirit of self-determination, our diversification process allows each community to choose the profile characteristics (one or many, visible or invisible) against which to diversify, and the extent to which the diversity requirement will override a user's own preferences (Markos & Michael, 2022a). Empirical studies on how humans perceive diversity in participant rankings have been carried out to validate the process.

7 Diversity As An Instrumental Value That Supports Fundamental Values

Any attempt to operationalize and mediate diversity in social interactions must consider the possible ethical ramifications. Technological mediation is not value-neutral, and the development of technology is never done without prioritizing different values and advancing a certain "agenda" (Brey, 2010; Friedman & Nissenbaum, 1996; Schelenz et al., 2024). Being mindful of the inevitably value-ridden nature of the technology development process, the $\text{I}\circ\text{U}$ platform seeks to promote the participants' exposure to diversity to the extent that it does not pose a threat to their safety. We, thus, have to carefully weigh the values of inclusion and protection to enable a meaningful exposure to diversity (Schelenz, 2022).

Our rationale follows the assumption that diversity is not an end, but a means; it has instrumental rather than intrinsic value (Zimmerman & Bradley, 2019). In the same way that diversity may promote the value of inclusiveness, it may also, if abused, demote the value of protection. This relates to the classic discussion in content moderation between the need to balance the freedom of speech and the prevention of hate speech and violence (Brown, 2021; Gillespie, 2018). On the one hand, public discourse should include as many perspectives as possible. On the other hand, public discourse should be free from abusive communication (such as, for instance, racist or sexist attacks) and protect minority positions.

Values like inclusiveness and protection are of a more fundamental nature vis à vis diversity (Helm et al., 2022), and they ought to be promoted for their own sake, even at the expense of diversity. Think of a request about a topic that is associated with stigma (e.g., mental health problems), where an individual may risk social exclusion just for seeking support (Goffman, 2009). Somewhat counter-intuitively, restricting the diversity of the respondents ends up enhancing the

individual's access to a diverse section of the social search space, by protecting them from hurtful and derogatory comments, and by creating a safe space within which to seek support. Accordingly, our platform implements algorithms and offers interface functionalities to interact anonymously and to restrict the respondent pool. At the same time, transparency is essential to enable individual control over personal data and to trust in the potential of a diversity-aware platform like WeNet (Schelenz et al., 2024).

8 Conclusion

We have sought to make progress on an ambitious goal: to design the I_oU and the prototypical implementation of an AI-mediated and diversity-aware platform for social interactions enabling it. In attempting to fulfill that goal, we have explored various facets of diversity while, notably, staying grounded in empirical work across a diverse set of participating communities. Among our contributions is the gathering and eventual sharing of social interaction data at a scale that goes beyond past attempts. This is only a first step, and there is still a long way to go. But we believe that the direction of rethinking social interactions under a diversity paradigm is promising, and it does raise several relevant research questions. Ultimately, the main take-home message is the importance of a holistic, empirically-grounded, and multidisciplinary perspective in pushing towards a future of diversity-aware hybrid human–AI societies.

It goes without saying that realizing the I_oU vision raises several key issues that this paper has not even mentioned. A key issue is how to make people searchable and findable while maintaining their privacy. This is crucial because people must disclose specific properties, such as their location and skills, which may ultimately make them recognizable. A second issue is security — how to manage malicious behavior. A third is how to keep people motivated in providing help to others. The work done so far is only a very first preliminary step. A fourth related issue is how to minimize disturbance. People would definitely lose motivation if they were asked for intervention too often. A fifth issue is data and profile quality. Data that changes frequently is costly to maintain. And so on. No I_oU can be enabled without a solid solution to these issues. At the same time, this work provides a first insight into the issue of how to understand and manage diversity, an issue which we believe is foundational and that, we believe, had not been previously studied with the necessary depth.

Acknowledgments

The authors wish to thank all our collaborators for their fantastic support; a most likely incomplete list is: Marina Bidoglia, Andrea Bontempelli, Maria Chiara Campodonico, Carlo Caprini, Ronald Chenu Abente, Galileo Disperati, Paula Helm, Alethia Hume, Peter Kun, Meegahapaola Lakmal, Vasileios Theodoros Markos, Chaitanya Nutakki, Marcelo Dario Rodas Brites, and Donglei Song.

ORCID iDs

Ivano Bison  <https://orcid.org/0000-0002-9645-8627>
 Luca Cernuzzi  <https://orcid.org/0000-0001-7803-1067>
 Shyam Diwakar  <https://orcid.org/0000-0003-1546-0184>
 Kobi Gal  <https://orcid.org/0000-0001-7187-8572>
 Fausto Giunchiglia  <https://orcid.org/0000-0002-5903-6150>

Funding

This research has received funding from the European Union's Horizon project "WeNet - The Internet of Us", under grant agreement no. 823783.

Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

References

- Achón, L., De Souza, A., Hume, A., & Cernuzzi, L. (2024). A diversity-aware recommendation system for tutoring. *AI Communications*, 37(4), 711–733. <https://doi.org/10.3233/AIC-230434>
- Ågotnes, T., van der Hoek, W., Rodríguez-Aguilar, J. A., Sierra, C., & Wooldridge, M. J. (2009). Multi-modal CTL: Completeness, complexity, and an application. *Studia Logica*, 92(1), 1–26. <https://doi.org/10.1007/s11225-009-9184-3>
- Arnaboldi, V., Conti, M., Passarella, A., & Dunbar, R. I. (2017). Online social networks and information diffusion: The role of ego networks. *Online Social Networks and Media*, 1, 44–55.
- Assi, K., Meegahapola, L., Droz, W., Kun, P., De Götzen, A., Bidoglia, M., Stares, S., Gaskell, G., Chagnaa, A., Ganbold, A., Zundui, T., Caprini, C., Miorandi, D., Zarza, J. L., Hume, A., Cernuzzi, L., Bison, I., Rodas Brites, M. D., Busso, M., & Chenu-Abente, R., ... Gatica-Perez, D. (2023). Complex daily activities, country-level diversity, and smartphone sensing: A study in denmark, italy,

- mongolia, paraguay, and uk. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, CHI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9781450394215. <https://doi.org/10.1145/3544548.3581190>.
- Bella, G., Helm, P., Koch, G., & Giunchiglia, F. (2024). Tackling language modelling bias in support of linguistic diversity. In *The 2024 ACM conference on fairness, accountability, and transparency*. (pp. 562–572).
- Ben Zaken, D., Segal, A., Cavalier, D., Shani, G., & Gal, K. (2022). Generating Recommendations with Post-Hoc Explanations for Citizen Science. In *Proceedings of the 30th ACM conference on user modeling, adaptation and personalization*. (pp. 69–78).
- Berger, P., & Luckmann, T. (2016). The social construction of reality. In “*Social Theory re-wired*”. Routledge.
- Bontempelli, A., Britez, M. R., Li, X., Zhao, H., Erculiani, L., Teso, S., Passerini, A., & Giunchiglia, F. (2022). Lifelong personal context recognition. *arXiv preprint arXiv:2205.10123*.
- Botta, C., Segal, A., & Gal, K. (2023a). Personalizing interventions with diversity aware bandits. *HHAI-WS 2023*.
- Botta, C., Segal, A., & Gal, K. (2023b). Sequencing educational content using diversity aware bandits. In *EDM*.
- Brey, P. (2010). Values in Technology and Disclosive Computer Ethics. In L. Floridi (Ed.) *The cambridge handbook of information and computer ethics* (pp. 41–58). Cambridge: Cambridge University Press. ISBN 0511845235. <https://doi.org/10.1017/CBO9780511845239.004>
- Brown, A. (2021). What is Hate Speech? Presentation of Hard Cases.
- Busso, M., Bontempelli, A., Malcotti, L. J., Meegahapola, L., Kun, P., Diwakar, S., Nutakki, C., Britez, M. D. R., Xu, H., Song, D., Correa, S. R., Mendoza-Lara, A. R., Gaskell, G., Stares, S., Bidoglia, M., Ganbold, A., Chagnaa, A., Cernuzzi, L., Hume, A., & Chenu-Abente, R. (Giunchiglia, F. (2025). Diversityone: A multi-country smartphone sensor dataset for everyday life behavior modeling. *Proceeding of the ACM Interactive Mobile Wearable Ubiquitous Technologies*, 9(1), 1–49. <https://doi.org/10.1145/3712289>
- Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociochi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9), e2023301118.
- Dietz, E., Kakas, A., & Michael, L. (2022). Argumentation: A calculus for human-centric AI. *Frontiers in Artificial Intelligence*, 5.
- dos Santos, T. F., Osman, N., & Schorlemmer, M. (2022). Ensemble and Incremental Learning for Norm Violation Detection. In P. Faliszewski, V. Mascardi, C. Pelachaud & M. E. Taylor (Eds.), *21st International conference on autonomous agents and multiagent systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022* (pp.427–435). International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS). <https://doi.org/10.5555/3535850.3535899>
- Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Transactions on Information Systems*, 14(3), 330–347. <https://doi.org/10.1145/230538.230561>
- García-Camino, A., Rodríguez-Aguilar, J. A., Sierra, C., & Vasconcelos, W. W. (2009). Constraint rule-based programming of norms for electronic institutions. *Autonomous Agents and Multi-Agent Systems*, 18(1), 186–217. <https://doi.org/10.1007/s10458-008-9059-4>
- Gatica-Perez, D., Santani, D., Biel, J. I., & Phan, T. (2019). Social multimedia, diversity, and global south cities: A double blind side. In *in Proc. ACM workshop on fairness, accountability, and transparency in multimedia (FAT/MM)*.
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- Gillespie, T. (2018). *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.
- Girardini, N. A., Centellegher, S., Passerini, A., Bison, I., Giunchiglia, F., & Lepri, B. (2023). Adaptation of student behavioural routines during COVID-19: A multimodal approach. *EPJ Data Science*, 12(1), 55.
- Giunchiglia, F. (2017). Computational humanism. Presentation final review meeting, EC FP7 project. *Smart Society*.
- Giunchiglia, F., & Bagchi, M. (2022). Representation heterogeneity. *arXiv preprint arXiv:2207.01091*.
- Giunchiglia, F., Batsuren, K., & Bella, G. (2017a). Understanding and Exploiting Language Diversity. *Proceedings of the international Joint Conference on Artificial Intelligence (IJCAI)*. (pp. 4009–4017).
- Giunchiglia, F., Bella, G., Nair, N. C., Chi, Y., & Xu, H. (2023). Representing interlingual meaning in lexical databases. *Artificial Intelligence Review*, 56(10), 11053–11069.
- Giunchiglia, F., Bignotti, E., & Zeni, M. (2017b). Personal Context Modelling and Annotation. In *2017 IEEE international conference on pervasive computing and communications workshops (PerCom Workshops)*. (pp. 117–122). IEEE.
- Giunchiglia, F., Bison, I., Busso, M., Chenu-Abente, R., Rodas, M., Zeni, M., Can, G., Giuseppe, V., Amalia, D., Peter, K., Amarsanaa, G., Altangerel, C., George, G., Sally, S., Miriam, B., Luca, C., Alethia, H., Jose Luis, Z., Hao, X., & Donglei, S., ... Daniel, G. P. (2022). A worldwide diversity pilot on daily routines and social practices (2020-2021). See <https://datascientia.disi.unitn.it/>. University of Trento Technical Report—DataScientia dataset descriptors.
- Giunchiglia, F., & Fumagalli, M. (2020). Entity Type Recognition – dealing with the diversity of knowledge. In *Knowledge representation conference (KR)*. Rhodes, Greece.
- Goffman, E. (2009). *Stigma: Notes on the management of spoiled identity*. Simon and Schuster.
- Harrison, D. A., Price, K. H., & Bell, M. P. (1998). Beyond relational demography: Time and the effects of surface- and deep-level diversity on work group cohesion. *Academy of Management Journal*, 41(1), 96–107.

- Helm, P., Michael, L., & Schelenz, L. (2022). “Diversity by Design”? balancing the inclusion and protection of users in an online social platform. In *Proceedings of the 2022 AAAI/ACM conference on artificial intelligence, ethics and society (AIES '22)*. (pp. 324–334). <https://doi.org/10.1145/3514094.3534149>
- Hume, A., Zarza, J. L., Gatica-Perez, D., Bison, I., & Cernuzzi, L. (2022). Analysis of the big-five personality traits in the chatbot “UC—paraguay”. *CLEI*, 25(2).
- Hutchinson, H., Mackay, W., Westerlund, B., Bederson, B. B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H., Roussel, N., & Eiderbäck, B. (2003). Technology Probes: Inspiring design for and with families. In *Proceedings of the SIGCHI conference on human factors in computing systems*, CHI '03, (pp. 17–24). New York, NY, USA: Association for Computing Machinery. ISBN 1581136307. <https://doi.org/10.1145/642611.642616>
- Ioannou, C., & Michael, L. (2021). Knowledge-based translation of natural language into formal logic. In *7th IJCAI Workshop on linguistic and cognitive approaches to dialog agents (LaCATODA)*.
- Jo, A. (2023). The promise and peril of generative ai. *Nature*, 614(1), 214–216.
- Kayongo, I., Malcotti, L., Zhao, H., & Giunchiglia, F. (2025). A methodology and a platform for high-quality rich personal data collection#. *The European Journal on Artificial Intelligence*, 38(2), 30504554251333615.
- Koch, G., Bella, G., Helm, P., & Giunchiglia, F. (2024). Layers of technology in pluriversal design decolonising language technology with the live language initiative. *CoDesign*, 20(1), 77–90.
- Kowalski, R. A., & Sergot, M. J. (2005). A Logic-based calculus of events. In I. Mani, J. Pustejovsky & R. J. Gaizauskas (Eds.), *The Language of Time—A Reader* (pp. 217–240). Oxford University Press.
- Kun, P., de Götzen, A., Bidoglia, M., Gommessen, N., & Gaskell, G. (2022). exploring diversity perceptions in a community through a Q&A Chatbot. In *in Proc. DRS022: Bilbao*.
- Markos, V., & Michael, L. (2022a). Post-hoc diversity-aware curation of rankings. In *Proceedings of the 14th international conference on agents and artificial intelligence (ICAART'22)*. virtual.
- Markos, V. T., & Michael, L. (2022b). Prudens: An argumentation-based language for cognitive assistants. In *Proceedings of the 6th International joint conference on rules and reasoning (RuleML+RR'22)*. virtual / Berlin, Germany.
- Markos, V. T., Thoma, M., & Michael, L. (2022). Machine coaching with proxy coaches. In *Proceedings of the COMMA 2022 international workshop on argumentation and machine learning (ArgML'22)*. Cardiff, Wales, U.K.
- Meegahapola, L., Droz, W., Kun, P., de Götzen, A., Nutakki, C., Diwakar, S., Correa, S. R., Song, D., Xu, H., Bidoglia, M., Gaskell, G., Chagnaa, A., Ganbold, A., Zundui, T., Caprini, C., Miorandi, D., Hume, A., Zarza, J. L., Cernuzzi, L., & Bison, I., ... Gatica-Perez, D. (2023). generalization and personalization of mobile sensing-based mood inference models: An analysis of college students in eight countries. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(4). <https://doi.org/10.1145/3569483>
- Mercado, A., Hume, A., Bison, I., Giunchiglia, F., Ganbold, A., & Cernuzzi, L. (2023). Social interactions mediated by the internet and the big-five: A cross-country analysis. *arXiv preprint arXiv:2309.10681*.
- Michael, L. (2019). Machine Coaching. In *Proceedings of the IJCAI 2019 workshop on explainable artificial intelligence (XAI'19)*. S.A.R. Macau, P.R. China.
- Michael, L. (2023a). Coachable AI. In *ERCIM News No. 132, special theme: Cognitive AI & Cobots*.
- Michael, L. (2023b). Explainability and the Fourth AI Revolution. In E. G. Carayannis & E. Grigoroudis (Eds.), *Handbook of research on artificial intelligence, innovation and entrepreneurship*. Edward Elgar Publishing.
- Michael, L., Bison, I., Busso, M., Cernuzzi, L., De Götzen, A., Diwakar, S., Gal, K., Ganbold, A., Gaskell, G., Gatica-Perez, D., Heesen, J., Miorandi, D., Ruiz-Correa, S., Schelenz, L., Segla, A., Sierra, C., Xu, H., & Giunchiglia, F. (2025). Towards open diversity-aware social interactions. *arXiv preprint arXiv:2503.16448*.
- Montes, N., Osman, N., & Sierra, C. (2022). A computational model of ostrom’s institutional analysis and development framework. *Artificial Intelligence*, 311, 103756. <https://doi.org/10.1016/j.artint.2022.103756>
- Montes, N., & Sierra, C. (2021). Value-guided synthesis of parametric normative systems. In F. Dignum, A. Lomuscio, U. Endriss & A. Nowé (Eds.), *AAMAS '21: 20th international conference on autonomous agents and multiagent systems, Virtual Event, United Kingdom, May 3-7, 2021*. (pp. 907–915). ACM. <https://doi.org/10.5555/3463952.3464060>
- Mueller, M. (2024). The myth of agi. *Internet Governance Project*.
- Orphanou, K., Otterbacher, J., Kleantous, S., Batsuren, K., Giunchiglia, F., Bogina, V., Tal, A. S., Hartman, A., & Kuflik, T. (2022). Mitigating bias in algorithmic systems—A fish-eye view. *ACM Computing Surveys*, 55(5), 1–37.
- Osman, N., Chenu-Abente, R., Shen, Q., Sierra, C., & Giunchiglia, F. (2021). Empowering users in online open communities. *SN Computer Science*, 2(4), 338. <https://doi.org/10.1007/s42979-021-00714-5>
- Osman, N., Sierra, C., Chenu-Abente, R., Qiang, S., & Giunchiglia, F. (2020). Open Social Systems. In *17th European conference on multi-agent systems (EUMAS)*. Thessaloniki, Greece.

- Perikos, I., & Michael, L. (2022). A Survey on Tie Strength Estimation Methods in Online Social Networks. In *Proceedings of the 14th international conference on agents and artificial intelligence (ICAART2022)*. (pp. 484–491). <https://doi.org/10.5220/0010845100003116>
- Røpke, I. (2009). Theories of practice—New inspiration for ecological economic studies on consumption. *Ecological Economics*, 68(10), 2490–2497.
- Schelenz, L. (2022). Diversity concepts in computer science and technology development: A critique. *Science, Technology, & Human Values*, 48(5), 1054–1079. <https://doi.org/10.1177/01622439221122549>
- Schelenz, L., Bison, I., Busso, M., De Götzen, A., Gatica-Perez, D., Giunchiglia, F., Meegahapola, L., & Ruiz-Correa, S. (2021). The Theory, Practice, and Ethical Challenges of Designing a Diversity-Aware Platform for Social Relations. In *Proceedings of the 2021 AAAI/ACM conference on ai, ethics, and society*. (pp. 905–915).
- Schelenz, L., Segal, A., Adelio, O., & Gal, K. (2024). Transparency-check: An instrument for the study and design of transparency in ai-based personalization systems. *ACM Journal on Responsible Computing*, 1(1), 1–18.
- Segal, A., Gal, K., Kamar, E., Horvitz, E., & Miller, G. (2018). Optimizing interventions via offline policy evaluation: Studies in Citizen Science. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 32).
- Shove, E., & Pantzar, M. (2005). Consumers, producers and practices: Understanding the invention and reinvention of nordic walking. *Journal of Consumer Culture*, 5(1), 43–64.
- Sierra, C., Osman, N., Noriega, P., Sabater-Mir, J., & Perelló, A. (2021). Value Alignment: A Formal Approach. *CoRR* abs/2110.09240. <https://arxiv.org/abs/2110.09240>.
- Tsamoura, E., Hospedales, T., & Michael, L. (2021). Neural-Symbolic Integration: A Compositional Perspective. In *Proceedings of the 35th AAAI conference on artificial intelligence (AAAI'21)*. virtual.
- Tsui, A., Terri, E., & O'Reilly, C. (1991). Being different: Relational demography and organizational attachment. *Academy of Management Proceedings*, 1, 183–187.
- Tsui, A., Terri, E., & O'Reilly, C. (1992). Being different: Relational demography and organizational attachment. *Administrative Science Quarterly*, 37(4), 549–579.
- Tubella, A. A., Theodorou, A., Dignum, V., & Michael, L. (2020). Contestable Black Boxes. In *Proceedings of the 4th international joint conference on rules and reasoning (RuleML+RR'20)*. Oslo, Norway.
- Yanovsky, S., Hoernle, N., Lev, O., & Gal, K. (2019). One Size Does Not Fit All: Badge Behavior in Q&A Sites. In *Proceedings of the 27th ACM conference on user modeling, adaptation and personalization*. (pp. 113–120).
- Zeni, M., Zaihrayeu, I., & Giunchiglia, F. (2014). Multi-device activity logging. In *Proc. of the ACM PERCOM conference—Adjunct publication* (pp. 299–302).
- Zimmerman, M. J., & Bradley, B. (2019). Intrinsic vs. extrinsic Value. <https://plato.stanford.edu/entries/value-intrinsic-extrinsic/>.