Intra-Organizational Communication 2.0

Savvas Charalambides[2], Loizos Michael[1,2]

[1]Open University of Cyprus, Nicosia, Cyprus
[2]CYENS Center of Excellence, Nicosia, Cyprus

Abstract

We put forward an architecture for the next generation of AI-mediated intra-organizational communication (IOC), towards enhanced team productivity and satisfaction. Our proposal rests on three key principles: *hybrid human-AI collaboration* via natural language interactions; *diversity-aware dissemination* of queries within the organization; *incremental and participatory development* of the IOC policy. We briefly discuss our ongoing work towards realizing the proposed IOC architecture.

Keywords

productivity, human-AI interaction, diversity, sustainability

1. Introduction

Communication resembles an organization's "central nervous system", and is the "social glue" that holds it together [1]. Effective communication leads to teamwork and cooperation, and is crucial for the realization of an organization's goals [2].

Achieving effective communication is not always straightforward, especially within large and diverse organizations with geographically-distributed teams and groups, which tend to communicate with others that happen to be geographically closer [3], rather than with those that would, in fact, be formally responsible and competent for a certain task. This problem can be exacerbated when it involves newly-hired / relocated employees, who might be unfamiliar with the organization's structure and their new coworkers' business responsibilities and know-how.

Towards tackling these intra-organizational communication (IOC) challenges, we present a proposal for an AI-mediated platform that matches employee queries to colleagues who can provide answers. Our proposal rests on three key principles:

(Principle 1) Support of cognitively-light and user-friendly hybrid human-AI collaboration, through the interaction of employees with others and with the AI-mediation modules in (a specified form of controlled) natural language.

(Principle 2) Matching of employee queries to competent colleagues in a diversity-aware manner, so that social and organizational biases / preconceptions are not reinforced, and a diverse perspective on subjective matters is promoted.

HHAI-WS 2023: Workshops at the Second International Conference on Hybrid Human-Artificial Intelligence (HHAI), June 26–27, 2023, Munich, Germany

charalambides.savvas@gmail.com (S. Charalambides); loizos@ouc.ac.cy (L. Michael)

(Principle 3) Incremental improvement of the AI-mediation policy in a participatory manner that involves the employees themselves, as the domain experts.

In the remainder of this paper, we first overview past work that relates to our proposed IOC platform (Section 2), we discuss the platform's constituent modules and how they relate to the three key principles discussed above (Section 3), and we briefly report on our ongoing work for implementing the platform (Section 4).

2. Overview of Related Work

The problem of supporting online collaboration by identifying appropriate members from a group to work on a given project or task has received attention in recent years [4]. HelpExchange is an early example of building an over-the-Internet "help network" [5], which employs a taxonomy to match posed queries to experts who might be able to answer them. The conversational agent INDIGO (Individual Differences for Group Optimization) is an example of more recent work [6], which is able to suggest suitable team members for optimal teaming by identifying and utilizing the participating users' preferences, expectations, and personality traits.

u-Help is a further relevant example of a distributed collaboration platform [7, 8], which maintains a local community of people helping each other with their everyday tasks, and identifies the most appropriate community member that can provide assistance while taking into account a trust-based scoring mechanism.

Closer to the knowledge-sharing aspect of our IOC platform is a mobile application by Mourtzis et al. [9], which uses natural language processing to understand problem statements, their causes, and their solutions. A repository of such triplets constitutes the accumulated knowledge of an organization's workforce, which supports the identification of solutions to future problem statements.

With the exception of the last reviewed work above, the remaining works do not focus on the IOC problem, but rather on broad day-to-day tasks or particular communities. More importantly however, none of the reviewed works attempt to consider diversity in their recommendations, nor do they propose a self-sustainable process for the improvement of the communication and recommendation policies.

Among the commercially-available IOC products that we were able to identify, closer to satisfying our proposed principles is Starmind's expert identification tool [10, 11]. The tool addresses Principle 1 by employing natural language processing to comprehend user queries. Those are, then, matched to experts through a principled approach that minimizes recommendation bias [12, 13], without, however, addressing fully Principle 2 by returning diverse perspectives on each query. The matching process improves over time through a form of semi-supervised learning from user feedback and external information, but the participation of users in line with Principle 3 is limited to evaluating the quality of responses, failing to capitalize on the organizational knowledge held by an organization's management in terms of which employees are (or should be) able to respond to some query.

Some considerations surrounding the role of diversity in an online social platform are touched upon by Helm et al. [14]. When viewed under a descriptive lens, diversity can help capture

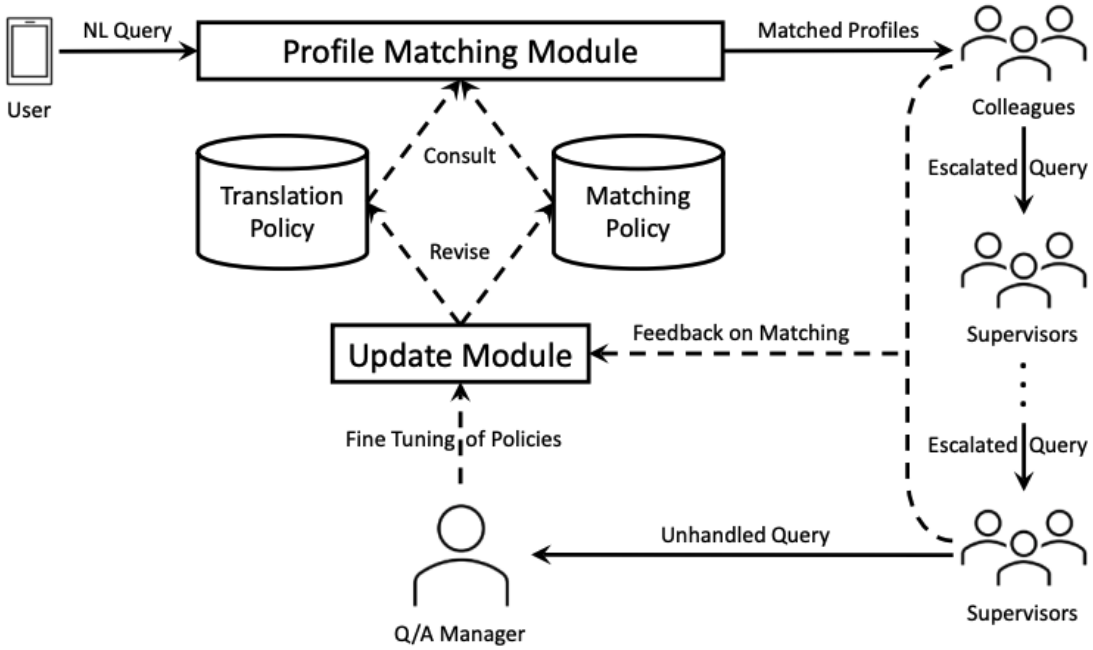

Figure 1: In our proposed IOC architecture, an event is initiated after a user poses a natural language query, which is translated into a machine-readable form and matched to a diverse set of colleagues who might help. If this is not the case, the recipients may escalate the query to their supervisors and/or offer feedback to the AI-mediation module on how to improve its matching policy. In case the query goes completely unhandled, it is forwarded to a Quality Assurance Manager, who can fine-tune the matching policy, or, if needed, improve the translation policy.

differences between users across their skills, practices, and personality traits. When viewed under a normative lens, it acts as an instrumental value to be promoted in support of certain fundamental values (e.g., productivity and inclusion), while being mindful of its ramifications on other fundamental values (e.g., protection and privacy). Both views are readily applicable to our IOC platform, towards ensuring that like-minded employees can work on tasks of joint interest, while making diverse employees feel comfortable to contribute and complement other employees within the frame of a difficult task [15].

3. Proposed IOC Architecture

Figure 1 illustrates our proposed IOC architecture, following the three principles.

The architecture foresees an interaction between employees and the platform through natural language, as per Principle 1. A translation policy, customized to the particularities and communication nuances of each organization, fine-tuned as needed by a Quality Assurance Manager, ensures that user queries and feedback are appropriately considered for the operation and improvement of the platform.

Once comprehended, a user query is forwarded, through the application of a matching

policy, to a subset of the user's colleagues, taking into account each employee's organizational profile, including their duties and responsibilities (and other deep-level characteristics [15]), as determined, for instance, either by their job scope and mandate, or by their supervisor's evaluation of their business know-how. In alignment with Principle 2, the matching above is counterbalanced with diversity-awareness measures according to legal, social, or ethical norms, which might typically relate to surface-level characteristics [15] like gender and age.

In the simplest scenario, employees can self-report whether a query they have received was appropriately forwarded to them, or whether a query they had posed was adequately addressed by their colleagues. Beyond that, however, the architecture provides for an escalation mechanism for unhandled queries to higher levels of the organizational hierarchy. Supervisors are assumed to be aware of whether their subordinates are (or should be) able to respond to a particular query, and can thus provide direct feedback for the improvement of the matching policy.

Queries that go completely unhandled are forwarded to a Quality Assurance Manager, who can investigate the existence of an out-of-scope query, the presence of a gap in organizational expertise, or other eventualities. A more mundane, but perhaps more likely, eventuality is the inappropriate parsing of a query or of some previously offered piece of feedback. In such cases, the Quality Assurance Manager can (instruct a technical expert to) intervene and fine-tune the platform's policies.

The provision for explicit translation and matching policies comes in support of the alignment of the architecture with Principle 3. The feedback coming from the employees and the Quality Assurance Manager can be much richer and direct compared to a labeled data-point, as typically used in supervised learning. Consequently, a policy can improve much more efficiently, and be explainable by design, by having the feedback become an integral part of its updated version. Quality-wise, this feedback leads to a less arbitrary update of the policy, compared to the often circumstantial generalizations that come from few-shot supervised learning.

4. Towards an Implementation

Our proposed architecture does not commit to any particular implementation. In our ongoing work, however, we have identified certain tools and techniques that we believe are good fit, and which we are using to realize our envisioned platform.

For the incremental and participatory improvement of the matching policy we adopt *machine coaching* [16, 17, 18], a paradigm that dictates how a machine can elicit policies in a principled and efficient manner via an argumentative dialogue with humans. The translation from natural language to the logic-based arguments required for machine coaching leverages elements from NESTOR [19], a system that itself adopts the machine coaching paradigm, at a meta level, to incrementally elicit (from the Quality Assurance Manager) the translation policy that it utilizes.

Central in our implementation plan is the use of the WeNet infrastructure [20] for the efficient development of machine-mediated social interaction applications, while protecting the users' privacy from potential surveillance of their communications, and balancing the promotion of diversity against other matters [21]. Once developed, we hope to evaluate our IOC platform within our host organizations.

Acknowledgments

This work was supported by funding from the EU's Horizon 2020 Research and Innovation Programme under grant agreements no. 739578 and no. 823783, and from the Government of the Republic of Cyprus through the Deputy Ministry of Research, Innovation, and Digital Policy.

References

[1] K. Cacciattolo, Defining Organisational Communication, European Scientific Journal 11 (2015) 79–87. URL: https://eujournal.org/index.php/esj/article/view/5954.

[2] S. R. Fussell, R. J. Kreuz, Social and Cognitive Approaches to Interpersonal Communication, Psychology Press, New York, U.S.A., 1998. doi:10.4324/9781315805917.

[3] J. M. Jackson, The Organization and its Communication Problems, Journal of Communication 9 (1959) 158–167. URL: https://doi.org/10.1111/j.1460-2466.1959.tb00313.x.

[4] M. Poblet, C. Sierra, Understanding Help as a Commons, The Commons Journal 14 (2020) 481–493. URL: https://thecommonsjournal.org/articles/10.5334/ijc.1029.

[5] P. Losleben, A. S. Tull, E. Zhou, Help-Exchange: An Arbitrated System for a Help Network, Journal of Network and Computer Applications 20 (1997) 123–134. URL: https://www.sciencedirect.com/science/article/pii/S108480459690034X.

[6] Z. Xiao, M. X. Zhou, W.-T. Fu, Who Should Be My Teammates: Using a Conversational Agent to Understand Individuals and Help Teaming, in: Proceedings of the 24th International Conference on Intelligent User Interfaces (IUI), 2019, pp. 437–447. URL: https://doi.org/10.1145/3301275.3302264.

[7] N. Osman, B. Rosell, C. Sierra, M. Schorlemmer, J. Sabater-Mir, L. Lemus, u-Help: Intelligent Volunteer Search for Mutual Help Communities, ArXiv abs/2301.11112 (2023). URL: https://arxiv.org/abs/2301.11112.

[8] A. Koster, J. Madrenas, N. Osman, M. Schorlemmer, J. Sabater-Mir, C. Sierra, A. Fabregues, D. de Jonge, J. Puyol-Gruart, P. Garcia-Calvés, u-Help: Supporting Helpful Communities with Information Technology, in: Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2013, pp. 1109–1110. URL: https://ceur-ws.org/Vol-918/111110378.pdf.

[9] D. Mourtzis, M. Doukas, N. Milas, A Knowledge-Based Social Networking App for Collaborative Problem-Solving in Manufacturing, Manufacturing Letters 10 (2016) 1–5. URL: https://www.sciencedirect.com/science/article/pii/S2213846316300104.

[10] Starmind, https://www.starmind.ai, ???? Accessed: 2023-05-07.

[11] S. Schafheitle, A. Weibel, I. Ebert, G. Kasper, C. Schank, U. Leicht-Deobald, No Stone Left Unturned? Toward a Framework for the Impact of Datafication Technologies on Organizational Control, Academy of Management Discoveries 6 (2020) 455–487. URL: https://doi.org/10.5465/amd.2019.0002.

[12] J. Stilgoe, R. Owen, P. Macnaghten, Developing a Framework for Responsible Innovation, Research Policy 42 (2013) 1568–1580. URL: https://www.sciencedirect.com/science/article/pii/S0048733313000930.

[13] A. Krishnakumar, Assessing the Fairness of AI Recruitment Systems, 2019. URL: http://resolver.tudelft.nl/uuid:1ce06e89-72a7-47fe-bdbd-93775732a30c.

[14] P. Helm, L. Michael, L. Schelenz, Diversity by Design? Balancing the Inclusion and Protection of Users in an Online Social Platform, in: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES), 2022, pp. 324–334. URL: https://doi.org/10.1145/3514094.3534149.

[15] L. Schelenz, I. Bison, M. Busso, A. de Götzen, D. Gatica-Perez, F. Giunchiglia, L. Meegahapola, S. Ruiz-Correa, The Theory, Practice, and Ethical Challenges of Designing a Diversity-Aware Platform for Social Relations, in: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES), 2021, pp. 905–915. URL: https://doi.org/10.1145/3461702.3462595.

[16] L. Michael, Machine Coaching, in: Proceedings of the IJCAI Workshop on Explainable Artificial Intelligence (XAI), 2019, pp. 80–86. URL: https://doi.org/10.5281/zenodo.3931266.

[17] V. Markos, M. Thoma, L. Michael, Machine Coaching with Proxy Coaches, in: Proceedings of the 1st COMMA Workshop on Argumentation and Machine Learning (ArgML), 2022, pp. 45–64. URL: https://ceur-ws.org/Vol-3208/paper4.pdf.

[18] V. Markos, L. Michael, Prudens: An Argumentation-Based Language for Cognitive Assistants, in: Proceedings of the 6th International Joint Conference on Rules and Reasoning (RuleML+RR), 2022, pp. 296–304. URL: https://doi.org/10.1007/978-3-031-21541-4_19.

[19] C. Ioannou, L. Michael, Knowledge-Based Translation of Natural Language into Symbolic Form, in: Proceedings of the 7th Linguistic and Cognitive Approaches To Dialog Agents Workshop (LaCATODA), 2021, pp. 24–32. URL: https://ceur-ws.org/Vol-2935/paper3.pdf.

[20] WeNet: The Internet of Us, https://www.internetofus.eu/, ???? Accessed: 2023-04-30.

[21] V. Markos, L. Michael, Post-hoc Diversity-Aware Curation of Rankings, in: Proceedings of the 14th International Conference on Agents and Artificial Intelligence (ICAART), 2022, pp. 323–334. URL: https://www.scitepress.org/Link.aspx?doi=10.5220/0010839600003116.