

Leveraging Diversity in Online Interactions

Nardine OSMAN, Bruno ROSELL I GUI and Carles SIERRA

Artificial Intelligence Research Institute (IIIA-CSIC), Barcelona

Abstract. This paper addresses the issue of connecting people online to help them find support with their day to day problems. We make use of declarative norms for mediating online interactions, and we specifically focus on the issue of leveraging diversity when connecting people. We run pilots at different university sites, and the results show relative success in the diversity of the selected profiles, backed by high user satisfaction.

Keywords. norms, interaction models, profiles, diversity, user experiments

1. Introduction

Despite the growing use of generative AI in various domains [7], connecting people to help them solve their problems and enrich their interactions will always be imperative. [2] shows that Stack Overflow’s traffic went down 14% in March 2023, due to people moving to ChatGPT for finding answers. However, this is not sustainable, as ChatGPT finds ‘good’ answers only because it trains on ‘good’ data, such as that provided by Stack Overflow. We argue that data flow from us humans will always be key. Similarly, finding support in fellow humans will always be indispensable.

The WeNet project attempts to address the issue of connecting people to help them find support with their day to day problems. While there are many platforms that do so today, leveraging AI to find suitable people to connect to is still lagging. In this paper we focus on the issue of mediating humans’ online interactions to ensure online behaviour unfolds according to expectations. As we focus on humans and their needs, our work has been motivated by: 1) empowering users, by giving them control over some of the technology’s functionalities, yet ensuring crucial basic requirements are not broken; 2) connecting people, by developing profile matching mechanisms that leverage profile diversity; and 3) aligning users’ understanding of the norms to ensure community members have a clear understanding of the norms mediating their interactions. We thus contribute to the hybrid machine-artificial intelligence field through diversity-enriching interaction-mediating technologies.

The following three sections (Sections 2–4) describe our approach for addressing the above three issues. Section 5 provides some initial results from our pilots. Finally, we conclude with Section 6 that analyses the results of the pilots and suggests future work.

2. A Declarative Approach for Modelling Interactions

The proposed WeNet interaction model [12,11] tackles the issue of empowering users by considering both, user requirements and system (or organisational) requirements. This gives users control over some functionalities (such as specifying what profiles are deemed relevant), while ensuring crucial basic requirements are not violated (such as ethical requirements and avoiding biases). These requirements that are mediating interactions are specified through declarative norms [13,1,3].

Following a declarative approach for norm representation allows for dynamic norms that can change at runtime. This has proven to be useful in our pilots as we could continue to tweak online behaviour (through norms) even after the app's deployment into production. For our plans for future work (Section 6), we plan to learn and adapt the norms at runtime. This highlights further the need for a declarative approach.

The language used to specify norms is based on a simple *if then* statement structure. The conditions are selected from the developed application's variables, allowing the norms to check issues such as if a question has been created, answered, etc., or whether certain profile attributes satisfy certain requirements. Similarly, the consequences are also based on the the application's acceptable actions, such as sending messages, updating a profile, etc.

3. Leveraging Diversity when Selecting Profiles

Connecting people is another of the main objectives of WeNet. We especially focus on leveraging diversity when selecting the profiles to tackle a given issue[10,8,9]. However, we ask the person making a question, the questioner, to decide which diversity dimensions are relevant for their question. We believe that different diversity dimensions might be suited for different contexts.

We choose to focus on four diversity dimensions: selecting profiles based on their domain interests, beliefs and values, social closeness, and physical closeness. Notice that we focus on what we refer to as deep features to avoid potential biases from arising when allowing the questioner to filter profiles based on shallow features like gender or nationality. Nevertheless, the WeNet project developed a diversification mechanism based on shallow features, to ensure that selected profiles are sufficiently diverse with respect to a given shallow feature, like gender.

How the above diversity dimensions are used in the different pilots to leverage diversity is domain dependent, and it is the norms that dictate how those are used.

4. Norm Alignment through User Feedback

Ensuring community members have a clear and aligned understanding of the norms that mediate their interactions is another issue that we have focused on. Say there is a norm that prohibits hate speech. We argue that what is considered hate speech might change from one community to another, or even change over time. As such, we have developed learning mechanisms [5,6,4] that allow the system to understand the meaning of norms (say what constitutes hate speech) from people's interaction and their feedback. It can then use what it learned to guide users' actions by highlighting when they violate certain norms and why.

5. Pilots

We have developed a number of pilots across different sites: London School of Economics (LSE), Aalborg University (AAU), and the National University of Mongolia (NUM). The objective was for students to use our developed application to find other students that can help answer their questions.

The diversity dimensions were defined differently for each of the pilot sites, and how the norms leveraged all diversity dimensions also varied from one pilot site to another, based on each site's requirements. For example, LSE gave different weights to the different dimensions while AAU and NUM did not, although they did specify that some diversity dimensions were primary while others were secondary, where requirements on primary dimensions should be satisfied first. Furthermore, social closeness was interpreted and implemented differently at each site.

Figure 1 presents some of the results of the AAU pilot. The results of the other pilots are aligned with those of AAU. The y-axis presents the number of selected profiles with the corresponding matching score (x-axis), or the score that describes the degree of a profile of satisfying the diversity requirements. The graphs follow a logarithmic scale.

The results highlight the capability of diversifying selected profiles using declarative norms. Different iterations of the pilots illustrated that the larger the communities, the easier it was to find profiles that fit the diversity requirements. We are confident that even larger community sizes would result in even better profile matches. (Community size in the last iteration whose results are presented here were at 105 users at LSE, 51 users at AAU, and 115 users at NUM.) Finally, the exit survey highlighted the users' satisfaction with the diversity dimensions and appreciated their value opportunity.

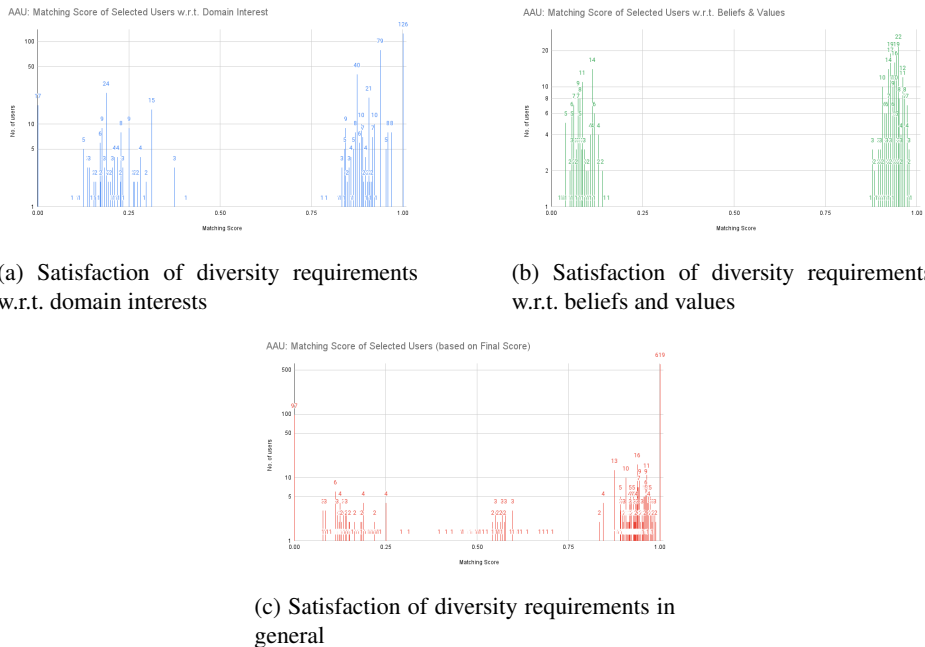


Figure 1. The selected profiles at the AAU pilot site and their satisfaction of diversity requirements

6. Conclusions and Future Work

The pilots, first and foremost, showcased the power of a declarative approach to norm specification. Even when the deadline has passed for developing the app and moving it to production, we continued to be able to adapt some functionalities through declarative norms that can change at runtime.

The norms opened the door for users to specify their diversity requirements. For a given question to be put forward in a community, the questioner had the option to choose whether they were interested in looking for similar/diverse profiles with respect to domain interests and beliefs and values, as well as close/distance profiles with respect to social and physical distance. The norms also implemented some fixed requirements, such as sending a question to five users only, and diversifying the final list of potential responders with respect to gender. The results of these pilots showed that the larger the communities, the easier it is to leverage diversity.

Explainability has emerged as a key requirement for users to understand how the norms and the diversity requirements function. The last pilot iteration demonstrated the success of explainability, though it also highlighted the challenges of explaining complex algorithms to users. A trade-off needs to be made when it comes to the details versus the clarity of the explanations.

The interest in leveraging diversity was limited, despite the fact that the exit survey has illustrated that users did in fact find value in those diversity requirements. As such, we strongly believe that diversity requirements are very much context dependent. When applicable, their value is appreciated. Though further research is needed to analyse the applicability per context.

Last, but not least, the pilots have highlighted the difficulty in predicting which norm interpretation is applicable for different contexts. For example, when should social closeness depend on past relations, and when should it depend on coming from similar academic circles? How many users should receive a question to guarantee good enough answers are collected, yet users are not bombarded with numerous questions? We believe that feedback is key in answering these questions. Furthermore, addressing those questions may be automated.

We argue that our work on developing mechanisms that allows the system to understand the meaning of norms (say what hate speech means) from people's interaction and their feedback can provide the basis for the next steps in our research, which we believe should focus on using such mechanisms to learn the most suitable norm for a given community. For example, analysing interaction data and user feedback can help learn the best number of people to forward a question to in a given community. And this number can be adapted to different communities. One may also learn which meaning of social closeness do users prefer and under what context. This can help to automatically adapt the interpretation and specification of norms for different contexts.

Acknowledgements

This research is supported by the EU funded VALAWAI (# 101070930) and WeNet (# 823783) projects, and the Spanish funded VAE (# TED2021-131295B-C31) and Rhy-mas (# PID2020-113594RB-100) projects.

References

- [1] Giulia Andrighetto, Guido Governatori, Pablo Noriega, and Leendert W. N. van der Torre. *Normative Multi-Agent Systems*, volume 4 of *Dagstuhl Follow-Ups*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2013.
- [2] David F. Carr. Stack Overflow is ChatGPT casualty: Traffic down 14% in March. *similarweb*, April 2023. Last accessed on 21-05-2023.
- [3] Mark D’Inverno, Michael Luck, Pablo Noriega, Juan A. Rodríguez-Aguilar, and Carles Sierra. Communicating open systems. *Artificial Intelligence*, 186:38–94, 2012.
- [4] Thiago Freitas dos Santos, Stephen Cranefield, Bastin Tony Roy Savarimuthu, Nardine Osman, and Marco Schorlemmer. Cross-community adapter learning (cal) to understand the evolving meanings of norm violation. In *32nd International Joint Conference on Artificial Intelligence, IJCAI 2023, Macao, China, August 19-25, 2023*, 2023.
- [5] Thiago Freitas dos Santos, Nardine Osman, and Marco Schorlemmer. Learning for detecting norm violation in online communities. In Andreas Theodorou, Juan Carlos Nieves, and Marina De Vos, editors, *Coordination, Organizations, Institutions, Norms, and Ethics for Governance of Multi-Agent Systems XIV - International Workshop, COINE 2021, London, UK, May 3, 2021, Revised Selected Papers*, volume 13239 of *Lecture Notes in Computer Science*, pages 127–142. Springer, 2021.
- [6] Thiago Freitas dos Santos, Nardine Osman, and Marco Schorlemmer. Ensemble and incremental learning for norm violation detection. In Piotr Faliszewski, Viviana Mascardi, Catherine Pelachaud, and Matthew E. Taylor, editors, *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*, pages 427–435. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2022.
- [7] Yogesh K. Dwivedi, Nir Kshetri, Laurie Hughes, Emma Louise Slade, Anand Jeyaraj, Arpan Kumar Kar, Abdullah M. Baabdullah, Alex Koohang, Vishnupriya Raghavan, Manju Ahuja, Hanaa Albanna, Mousa Ahmad Albashrawi, Adil S. Al-Busaidi, Janarthanan Balakrishnan, Yves Barlette, Sriparna Basu, Indranil Bose, Laurence Brooks, Dimitrios Buhalis, Lemuria Carter, Soumyadeb Chowdhury, Tom Crick, Scott W. Cunningham, Gareth H. Davies, Robert M. Davison, Rahul Dé, Denis Dennehy, Yanqing Duan, Rameshwar Dubey, Rohita Dwivedi, John S. Edwards, Carlos Flavián, Robin Gauld, Varun Grover, Mei-Chih Hu, Marijn Janssen, Paul Jones, Iris Junglas, Sangeeta Khorana, Sascha Kraus, Kai R. Larsen, Paul Latreille, Sven Laumer, F. Tegwen Malik, Abbas Mardani, Marcello Mariani, Sunil Mithas, Emmanuel Mogaji, Jeretta Horn Nord, Siobhan O’Connor, Fevzi Okumus, Margherita Pagani, Neeraj Pandey, Savvas Papagiannidis, Ilias O. Pappas, Nishith Pathak, Jan Pries-Heje, Ramakrishnan Raman, Nripendra P. Rana, Sven-Volker Rehm, Samuel Ribeiro-Navarrete, Alexander Richter, Frantz Rowe, Suprateek Sarker, Bernd Carsten Stahl, Manoj Kumar Tiwari, Wil van der Aalst, Viswanath Venkatesh, Giampaolo Viglia, Michael Wade, Paul Walton, Jochen Wirtz, and Ryan Wright. “so what if chatgpt wrote it?” multidisciplinary perspectives on opportunities, challenges and implications of generative conversational ai for research, practice and policy. *International Journal of Information Management*, 71:102642, 2023.
- [8] Athina Georgara, Juan A. Rodríguez-Aguilar, and Carles Sierra. Towards a competence-based approach to allocate teams to tasks. In Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé, editors, *AAMAS ’21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, pages 1504–1506. ACM, 2021.
- [9] Athina Georgara, Juan A. Rodríguez-Aguilar, Carles Sierra, Ornella Mich, Raman Kazhamiakin, Alessio Palmero Aprosio, and Jean-Christophe R. Pazzaglia. An anytime heuristic algorithm for allocating many teams to many tasks. In Piotr Faliszewski, Viviana Mascardi, Catherine Pelachaud, and Matthew E. Taylor, editors, *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*, pages 1598–1600. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2022.
- [10] Athina Georgara, Carles Sierra, and Juan A. Rodríguez-Aguilar. TAIP: an anytime algorithm for allocating student teams to internship programs. *CoRR*, abs/2005.09331, 2020.
- [11] Nardine Osman, Ronald Chenu-Abente, Qiang Shen, Carles Sierra, and Fausto Giunchiglia. Empowering users in online open communities. *SN Comput. Sci.*, 2(4):338, 2021.
- [12] Nardine Osman, Carles Sierra, Ronald Chenu-Abente, Qiang Shen, and Fausto Giunchiglia. Open social systems. In Nick Bassiliades, Georgios Chalkiadakis, and Dave de Jonge, editors, *Multi-Agent Systems and Agreement Technologies - 17th European Conference, EUMAS 2020, and 7th International Confer-*

ence, AT 2020, Thessaloniki, Greece, September 14-15, 2020, Revised Selected Papers, volume 12520 of *Lecture Notes in Computer Science*, pages 132–142. Springer, 2020.

- [13] Harko Verhagen, Martin Neumann, and Munindar P. Singh. Normative multiagent systems: Foundations and history. In Amit Chopra, Leendert van der Torre, Harko Verhagen, and Serena Villata, editors, *Handbook of Normative Multiagent Systems*, page 3–25. College Publications, 2018.