

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

A Chatbot for Asylum-Seeking Migrants in Europe

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Fazzinga, B., Palmieri, E., Vestoso, M., Bolognini, L., Galassi, A., Furfaro, F., et al. (2024). A Chatbot for Asylum-Seeking Migrants in Europe [10.1109/ICTAI62512.2024.00104].

Availability:

This version is available at: <https://hdl.handle.net/11585/1003796> since: 2025-02-17

Published:

DOI: <http://doi.org/10.1109/ICTAI62512.2024.00104>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

A Chatbot for Asylum-Seeking Migrants in Europe

Bettina Fazzinga^{*✉}, Elena Palmieri^{†✉}, Margherita Vestoso^{‡✉}, Luca Bolognini^{†✉},

Andrea Galassi^{†✉}, Filippo Furfaro^{*✉}, Paolo Torroni^{†✉}

^{*}University of Calabria, Rende (CS), Italy

Email: {bettina.fazzinga, filippo.furfaro}@unical.it

[†]DISI, University of Bologna, Bologna, Italy

Email: {e.palmieri, luca.bolognini6, a.galassi, p.torroni}@unibo.it

[‡]University of Naples Federico II, Naples, Italy

Email: margherita.vestoso@unina.it

Abstract—We present ACME: A Chatbot for asylum-seeking Migrants in Europe. ACME relies on computational argumentation and aims to help migrants identify the highest level of protection they can apply for. This would contribute to a more sustainable migration by reducing the load on territorial commissions, Courts, and humanitarian organizations supporting asylum applicants. We describe the background context, system architecture, underlying technologies, and a case study used to validate the tool with domain experts.

I. INTRODUCTION

Addressing migratory phenomena sustainably and equally is universally understood as a crucial global challenge. So much so, that the United Nations' 2030 Agenda for Sustainable Development includes, among its targets, to *facilitate orderly, safe, regular and responsible migration and mobility of people*.¹ Not surprisingly, migration is one of the top current concerns of the European Union too, and there's every indication that the phenomenon is expanding. In 2023 there were almost one million asylum applications in the Member States of the European Union,² a 21.3% increase compared to the previous year and the highest number since 2016. Because of recent conflicts, the number of irregular arrivals and asylum applications is estimated to grow further.³ To handle the phenomenon more effectively and humanly [1], the EU Commission has set different reforms. Among them is the recent Pact on Migration and Asylum, i.e., a set of new rules managing migration and establishing a common asylum system at the EU level. Such a normative effort was made

necessary by the complexity of current asylum rules. The application of such rules requires a non-trivial evaluation of legal and factual data concerning the applicant's living conditions, both in the host country and in the country of origin, which makes it difficult for the territorial commissions, Courts, and even for migrants, to achieve a preliminary overview on their chances to obtain protection.

In this context, a preliminary assessment and support tool would be beneficial in orienting applicants and helping them understand the possibilities and the requirements before they begin the formal asylum request process. The simplest tool could be a checklist. However, that would help with a well-identified simple procedure, but it could hardly tackle the additional, more complex problem of understanding which procedures are relevant. Indeed, there is not only one type of protection but several ones. Importantly, since applicants may be political refugees and victims of abuse, discrimination, and persecution, the collection and processing of their personal data for immigration purposes must satisfy requirements of privacy. Finally, the legal domain requires transparency and auditability concerning the application of the reference normative framework.

In the last years, Large Language Model (LLM)-based chatbots have become a popular choice for non-trivial question-answering systems of all sorts, and there's ample choice of technologies and tools for building systems with an unprecedented level of fluency and a seeming ability to perform basic inference steps. However, it also became apparent that LLM-based chatbots have significant limitations in terms of hallucinations, biases, jailbreaks and anonymity [2], [3]. These limitations undermine their reliability and usability in domains with strong requirements in terms of transparency, auditability, privacy, and where the factuality of the system's answers must be guaranteed. A popular approach to mitigating part of these problems is resorting to Retrieval-Augmented Generation (RAG) architectures [4], whereby LLMs are prompted to generate outputs only based on verified, traceable sources. Nevertheless, even state-of-the-art RAG systems are hampered by their apparent inability to perform complex (multi-hop) reasoning and integrate information from multiple documents [5]. This is certainly an issue in the domain at hand, where the complex legislation on international protection cannot

Bettina Fazzinga, Elena Palmieri, and Margherita Vestoso contributed equally. Corresponding author: Andrea Galassi. This work was partially supported by the following projects: European Commission's NextGeneration EU programme, PNRR – M4C2 – Inv. 1.3, PE00000013 - "FAIR - Future Artificial Intelligence Research" – Spoke 8 "Pervasive AI"; EU H2020's ICT48 project "Humane AI Net" Grant #952026; MUR project PRIN 2022 EPICA "Enhancing Public Interest Communication with Argumentation" (CUP H53D23003660006) funded by Next Generation EU.

¹<https://www.migrationdataportal.org/sdgs?node=10>

²<https://ec.europa.eu/eurostat/data/database>

³International Centre for Migration Policy Development (ICMPD), Migration Outlook 2024, <https://www.icmpd.org/about-us/icmpd-migration-outlook>.

be broken down into simple independent fragments where answers can be readily found but, on the contrary, answers require complex reasoning from multiple texts.

An alternative approach to building such a tool could rely on computational argumentation. The tool we propose builds on [6] and is named ACME: A Chatbot for asylum-seeking Migrants in Europe. We shall stress that, like [6], whose aim is to inform about COVID-19 vaccines and not play the role of a medical doctor, the purpose of ACME is to support, inform, and guide asylum applicants, not to assist or replace a judiciary expert in a “predictive justice” fashion, a topic imbued with ethical issues [7] which is outside our scope. Compared to [6], ACME has a similar system architecture and preserves its properties, but it offers more modularity, an enhanced natural language interface, and can handle more complex domains. Relevant properties ACME exhibits include: data governance and privacy thanks to its modular architecture; transparency and explainability thanks to argumentative reasoning; and the ability to integrate and reasoning with explicit, expert-made, formalized knowledge, ensuring auditability.

In the next section, we provide the essential background. We then present the chatbot architecture and technologies, and discuss its validation via a small-scale study with experts.

II. BACKGROUND

A. Legal background

The *Refugee Status* is the strongest level of protection that can be granted. It is given to anyone who fears persecution for reasons of race, religion, nationality, membership to a particular social group or political opinion and is unable or unwilling to avail himself of the protection of his country. *Subsidiary Protection Status* is a second-level strong guarantee which takes place mostly for the same reasons as the refugee status but only if the latter cannot be recognized. Given the similarity of the requirements, for the purposes of this case study, we consider them to be the same. Finally, *Special (or Humanitarian) Protection* is a National type of protection; this is a weaker safeguard that is guaranteed to applicants who are not eligible as refugees or beneficiaries of subsidiary protection but worthy of protection due to special humanitarian reasons. For each kind of protection, the application is grounded on requirements specified by the law, considered also in terms of living law. In most cases, migrants do not know the regulatory framework of international protection; therefore, they ignore which kind of protection they can apply for. Most of the time, they apply for all types although they can only get one, or none. ACME is meant to help the applicant understand the level of protection they could realistically apply for. To provide such information, ACME uses criteria defined by legal experts that consider both positive discipline and relevant jurisprudence. This kind of data is usually unavailable online but requires lawyers to use on-licence software and databases to get it and then make their evaluations. The knowledge arising from such a process is thus inaccessible to traditional LLMs.

B. Argumentation-enabled chatbots

Recently there have been some proposals for argumentation-enabled chatbots supporting persuasive interaction [8] and privacy-preserving medical advice [6]. Argumentation is particularly suited to this sort of application due to its amenability to capture rich reasoning patterns and interconnections between alternative options and supporting elements, as well as dialogical interactions [9]. Computational argumentation frameworks are ideal candidates to support reasoning tools that must exhibit transparency and auditability requirements, because they are symbolic and thus interpretable [10].

Chatbots can benefit from computational argumentation in terms of explainability and auditability, improving user trust, and exploiting a reasoning engine to guide the conversation. The application of argumentative chatbots spans several topics such as healthcare [11], [12] and COVID-19 [6], [8], [13], [14], ethics [15], university fees [16], [17] and education [18]. In [16], the authors utilize a crowd-sourced argument graph as a knowledge base to support a chatbot persuading the user, considering the UK’s university fees as case study [16], and then extend the same approach to COVID-19 vaccines [8]. The chatbots described in these studies address the user’s concerns by associating them with an argument in the graph and replying with its counterargument. Aiming at broadening the user’s viewpoint, the authors of [15] have developed a German argumentative chatbot to converse about AI-related ethically sensitive topics. The system determines the user’s stance on a certain topic and challenges it. In their study to implement a healthcare chatbot, the authors of [11] employ an Explanation-Question-Reply (ERQ) argument scheme modelling the agents’ interactions and clinically specialized argument schemes (ASPT). The chatbot aims to deliver patients’ treatment plans and answer the user’s questions. Our work builds on [6], [13], where the authors propose a modular architecture based on similarity between embeddings and computational argumentation to address a case study on COVID-19 vaccines. One of the shortcomings of such work is that natural language understanding is addressed only through sentence embedding similarity. The ability of the system to understand is therefore tied to the richness of the sentences in the knowledge base and may hinder the comprehension of concepts that are not explicitly expressed in a similar way. Besides the natural language module, other important differences from previous work are the domain, extent, and complexity of the case study, since the COVID-19 case study was much simpler and self-contained, and the implementation, which [6], [13] don’t have, whereas ACME has been implemented into an open-source tool.⁴

C. Chatbots for immigration

As far as chatbots in the immigration domain are concerned, the authors of [19] leverage a knowledge graph containing information about the UK’s immigration laws to aid the user in

⁴<https://github.com/It-nlp-lab-unibo/ACME-A-Chatbot-for-Migrants-in-Europe>

preparing a visa request using question templates as a language interface, but does not use argumentation in any way. In [20], the authors conducted a series of activities with migrants and other stakeholders to co-design a personality-driven chatbot aimed at favor integration, focusing on aspects such as the characteristic of the avatar and how to motivate users to train the chatbot.

III. ARCHITECTURE AND TECHNOLOGY

ACME is a modular system, with several interfaces to fit different contexts. For example, for practical reasons, we believe a speech-to-text interface will be part of it. However, this paper focuses on the core components of ACME responsible for reasoning and interaction, and the interplay between computational argumentation and natural language interface.

Like [6], ACME has a neuro-symbolic architecture, where a neural module dedicated to understanding the user is combined with a symbolic module dedicated to reasoning. It includes:

- a **knowledge base** (KB) consisting of a text file representing an argument graph and a JSON file that associates a set of natural language sentences with each node;
- a neural **natural language** module whose task is to interact with the user, understanding what they write;
- a symbolic **argumentation** module to reason from all the available data, compute answers and provide explanations.

The user interacts with the natural language module providing relevant information and answering questions. The information thus elaborated is mapped to KB concepts. The argumentation module performs reasoning on the KB and finds the appropriate answer. In case a definite answer is not possible, the system identifies a node that would help determine an answer and asks the user to provide such information.

The architecture design is focused on decoupling between the natural language module and the argumentation module. The latter has access only to the information mapped into the concepts of the KB, which by definition are relevant for inferring an answer. Any other irrelevant information provided by the user is discarded, guaranteeing their privacy.

A. User Interface

The user interface consists of two windows: the main chat window and the status window. The main window, represented in Figure 1a, has a standard chatbot layout, with the main dialogue area above and the user input area below. The status window, in Figure 1b, contains a list of all the argumentative nodes of the graph. The nodes are highlighted in real time while the user shares new information. If the node tallies with the user input, the interface will highlight it in green; if it contradicts it, the interface will highlight it in red. The UI and the other ACME modules are all implemented in Python.

B. Natural Language Module

The focus of this module is to elicit information from the user and then perform Natural Language Understanding (NLU) to map them into the nodes of the KB. In particular,

the matching of the user answer with KB nodes is performed in several ways.

If the chatbot is inquiring about the information represented in a specific node and the answer of the user is a direct “yes” or “no”, the corresponding node is considered matched.

Each node in the KB is associated with a set of natural language strings representing that concept. The user’s input and the KB strings are embedded using SentenceTransformers [21]. Then, the input is compared against each string and if their similarity is above a fixed threshold, the sentence’s node is considered matched.

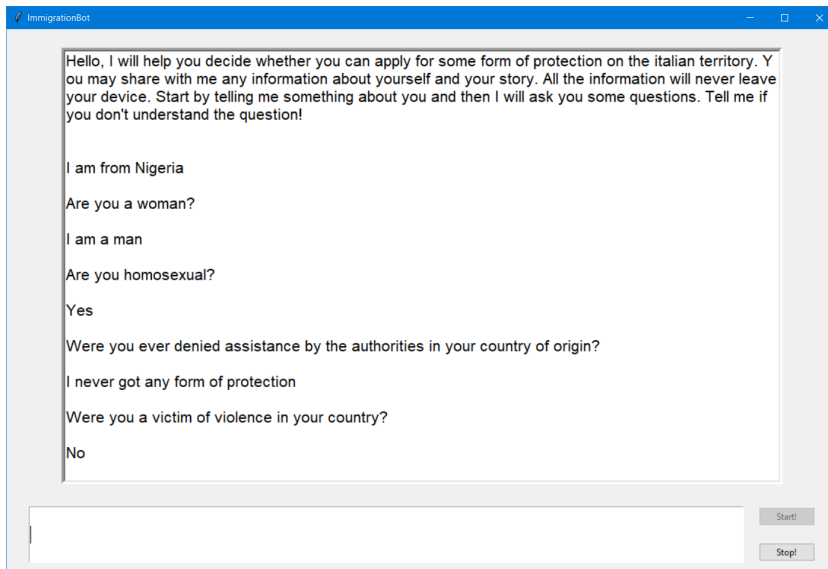
Otherwise, if the similarity does not reach the threshold, the input is given to a Large Language Model. Through the use of LLM, the system can understand concepts that are not explicitly presented as they are in the knowledge base, but can still be inferred. Presently, ACME integrates Llama 2 [22] 70B-chat; however, it is safe to assume that realistically field applications will rely on much smaller language models such as Microsoft’s Phi-3. The LLM is given a prompt that includes the sentence provided by the user, a list of all the nodes of the graph with their textual description, and instructions to return the relevant node. If all three methods do not find any corresponding node, the chatbot asks the user for clarification.

C. Argumentation Module

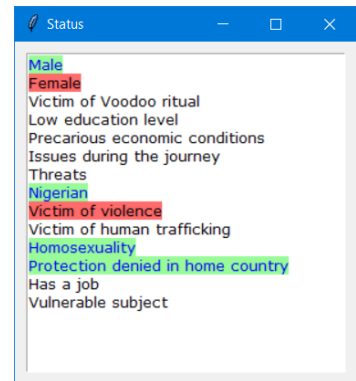
The Argumentation Module is in charge of selecting the proper questions to be asked to the user and of computing the replies. This module relies on a Knowledge Base on the form of an Abstract Argumentation Framework [10], where all the facts needed to obtain the different protections are encoded as special arguments named *status arguments*, while the different form of protections are encoded as special arguments named *reply arguments*. Figure 2 show an excerpt of our Knowledge Base. Arguments P_1 and P_2 , depicted in blue in the figure, are the *reply arguments*, as P_1 and P_2 encode the different forms of protections. All the other arguments, depicted in white, are the *status arguments*.

Each status argument is linked to one or more reply arguments it *endorses*. Status nodes may also *attack* other status or reply arguments, typically because the facts they represent are incompatible with one another. In Figure 2, the endorsement relation is depicted in green, while the attack relation is depicted in red. For example, argument *woman* mutual attacks argument *man*, endorses P_1 and attacks P_2 , because P_1 is a form of protection intended for women and P_2 is a form of protection destined for men.

Each dialogue session relies on dynamically acquired knowledge, expressed as a set of status arguments S , that encode user information. Basically, S contains the status nodes of the KB *activated* so far, that is corresponding to the information the user has communicated to the system since the start of the dialogue session. Differently from other proposals, at each turn, our system does not simply select a reply endorsed by S . On the contrary, the aim of the dialogue strategy is to provide the user with a reply that is both endorsed and *defended* by S from every attack. In other words, the



(a) Main chat window



(b) Status window

Fig. 1: Overview of the interface.

system works to provide only robust replies, possibly delaying replies that need further fact-checking. In fact, our system distinguishes between *consistent* and *potentially consistent* reply. The former can be given to the user right away, as it can not possibly be proven wrong in the future. The latter, albeit consistent with the current known facts, may still be defeated by future user input (that means that it is attacked by some argument and it is not currently defended by what it is in S), and therefore it should be delayed until a successful elicitation process is completed. Basically, at each turn, the module searches for a consistent reply, and if it does not exist, a potentially consistent reply is selected and all the arguments that can lead to its defense are tried to be added to S by asking the user the proper questions. For example, supposing that the chatbot has acquired the information that the user is a woman, S contains the *woman* argument, but ACME can not give right away the P_1 outcome to the user, as P_1 is attacked by the *others* argument (and by many others in the actual KB). Thus, P_1 is a *potentially consistent* reply, and the chatbot has to ask more questions to the user in order to promote P_1 to a *consistent* reply. To this end, ACME has to get to know, as an example, whether the user comes from Nigeria: in this case, the attack to P_1 from argument *others* is “neutralized” by argument *Nigeria*, as the latter attacks *others* and, thus, defends P_1 .

Besides providing replies, the system can also provide *explanations* for the given replies. An explanation of a reply r consists of two parts. The first one contains the arguments leading to r , i.e., those belonging to the set S that endorses r . The second one encodes the *why nots*, to explain why the system did not give other replies.




Fig. 2: An excerpt of our Knowledge Base

IV. VALIDATION

We implemented an asylum request case study in collaboration with immigration lawyers and other domain experts, based on an analysis of data on asylum applications in Italy over the last 10 years. We focus on the possible use of ACME when applying for international protection in Italy by people whose country of origin is Nigeria, considering the conditions required by two protection classes: the Refugee/Subsidiary Protection Status, and the Special Protection.

A. Modeling

We defined 13 argument nodes, each with its opposite, for a total of 26, and a node for each possible reply, namely Refugee Status, Special Protection and no protection. The argumentative nodes are all relevant to understand the person’s background and potentially grant them one of the possible protections. They include gender and country of origin, and a collection of nodes representing all the possible reasons that motivate and support some form of protection. These include being a victim of human trafficking, violence, or threats, being employed in the hosting country, being homosexual, having

a precarious economic condition, a low level of instruction, being vulnerable. For example, someone who has to provide for their family or has children born in the hosting country is considered vulnerable.

Each node attacks its opposite node. For example, if a person states that he is a man, he cannot state that he is a woman. If he states that he is from Nigeria, he cannot state that he is from another country. If something like that happens during the conversation, ACME will recognize the contradiction thanks to the attack relation between the opposite nodes and ask the user for clarification.

Other attack relations are between argument nodes and reply nodes, for example, if someone lacks the criteria of vulnerability, it is unlikely that they will get the Special Protection, so there is an attack relation between the argumentative node “No Vulnerability” and the “Special Protection” reply node. On the other hand, the vulnerability condition can lead to the Special Protection, so there is a support relation between the two nodes. Another support relation exists between the non-existence of conditions that can lead to the strongest form of protection, i.e., the Refugee Status, and the weaker Special Protection node.

B. Experiment

To validate our tool on this case study, we run a double-blind experiment involving a human expert on the domain. First, we generated 10 cases that simulate an asylum-seeking person (profiles). We did so, by assigning a random value to each relevant element considered in our knowledge base (such as country of origin, employment situation, etc.). For each case, we interacted with the chatbot by impersonating a possible applicant, coherently with the randomly generated profile, describing the case through natural language and eventually obtaining an answer. We then provided the same facts in the profile to the legal expert, and asked them to determine the highest degree of asylum the applicant could request. For the sake of this experiment, the expert was instructed to consider all the provided elements as true, independently of how plausible they might be. Finally, we compared expert answers with ACME answers.

Our tool and the legal expert agreed on all 10 cases, giving the same answer, thus validating our work for this specific case study.

For comparison, we run the same experiment with a state-of-the-art LLM with zero-shot prompting. We instructed GPT-4o to impersonate an expert in Italian immigration law tasked with aiding an asylum-seeking migrant to identify the strongest form of protection the migrant can apply to. We gave the migrant’s statements as they were listed in the relevant profile, and instructed the model to choose between the options Special Protection, Refugee Status and no protection, and to justify the answer. We tried with two different prompts. In both cases, GPT-4o returned the correct answer only 3 times out of 10. This corroborates our hypothesis that even the most advanced LLMs do not have access to the relevant jurisprudence, or are unable to use the facts they know meaningfully in this domain.

In two cases, the same questions asked using different prompts yielded different answers, showing a fundamental lack of robustness. We also found it concerning that the argumentation used by the LLM to justify the (mostly random) answers did sound like expert advice.

V. DISCUSSION

It has often been observed that legal knowledge is difficult, if not impossible, to access for the individuals who should benefit from it [23]. The point has been made in particular for consumer law. However, it also holds for the regulatory framework of international protection: a system whose beneficiaries, the migrants, are usually unable to navigate alone. In this work we described a system, ACME, intended to facilitate a sustainable approach to migration by guiding and supporting a migrant or someone acting on their behalf towards identifying the highest level of protection they can apply for.

We are aware of possible ethical issues arising from the adoption of AI tools in the legal domain [7]. However, ACME is not meant to assist or replace international protection judges. Indeed, a Court’s decision to approve or reject an asylum application is grounded on evaluations, like assessing the credibility and coherence of the applicant’s narration or the usefulness of the proof provided, which go well beyond the scope of our tool. For example, during the validation of the case study, the legal expert pointed out that one of the cases presented a suspicious degree of incoherence. An asylum application based on the same elements would have been probably regarded as not credible and eventually rejected by a court. These considerations regarding the truthfulness of the applicant’s declarations are outside the scope of this work. Our purpose is to support migrants before they present asylum applications to inform and guide them with such a complex procedure, and as such, it has to be regarded as an effort to make migration more sustainable. When used in the field, ACME will prominently exhibit a disclaimer to that effect.

In its present realization, ACME is only a proof-of-concept prototype. However, it is the first one of its kind: a necessary step towards exploring the suitability and potential of the interplay between computational argumentation and natural language processing in such a complex domain. We hope that our work can inspire and motivate further research in this domain. In the future, we plan to extend our proof-of-concept to cover a larger variety of situations and to include further modules, such as a speech-to-text interface and automatic translation.

ACKNOWLEDGMENT

We thank Avv. Mariapia Frisina for the valuable insights and discussions.

REFERENCES

- [1] K. Hailbronner, *Immigration and asylum law and policy of the European Union*. BRILL, 2024.
- [2] A. Casheekar, A. Lahiri, K. Rath, K. S. Prabhakar, and K. Srinivasan, “A contemporary review on chatbots, AI-powered virtual conversational agents, ChatGPT: Applications, open challenges and future research directions,” *Computer Science Review*, vol. 52, p. 100632, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1574013724000169>

- [3] A. Azaria, "ChatGPT: More human-like than computer-like, but not necessarily in a good way," in *ICTAI*. IEEE, 2023, pp. 468–473. [Online]. Available: <https://doi.org/10.1109/ICTAI59109.2023.00074>
- [4] K. Guu, K. Lee, Z. Tung, P. Pasupat, and M. Chang, "Retrieval augmented language model pre-training," in *ICML*, ser. Proceedings of Machine Learning Research, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 3929–3938. [Online]. Available: <https://proceedings.mlr.press/v119/guu20a.html>
- [5] J. Chen, H. Lin, X. Han, and L. Sun, "Benchmarking large language models in retrieval-augmented generation," in *AAAI*, vol. 38, no. 16, 2024, pp. 17 754–17 762.
- [6] B. Fazzinga, A. Galassi, and P. Torroni, "A privacy-preserving dialogue system based on argumentation," *Intell. Syst. Appl.*, vol. 16, p. 200113, 2022. [Online]. Available: <https://doi.org/10.1016/j.iswa.2022.200113>
- [7] S. Lazar and J. Stone, "On the site of predictive justice," *Noûs*, 2023.
- [8] L. A. Chalaguine and A. Hunter, "Addressing popular concerns regarding COVID-19 vaccination with natural language argumentation dialogues," in *ECSQARU*, ser. Lecture Notes in Computer Science, vol. 12897. Springer, 2021, pp. 59–73.
- [9] S. Modgil, F. Toni, F. Bex, I. Bratko, C. I. Chesñevar, W. Dvořák, M. A. Falappa, X. Fan, S. A. Gaggi, A. J. García, M. P. González, T. F. Gordon, J. Leite, M. Možina, C. Reed, G. R. Simari, S. Szeider, P. Torroni, and S. Woltran, *The Added Value of Argumentation*. Dordrecht: Springer Netherlands, 2013, pp. 357–403. [Online]. Available: https://doi.org/10.1007/978-94-007-5583-3_21
- [10] I. Rahwan and G. R. Simari, *Argumentation in Artificial Intelligence*, 1st ed. Springer Publishing Company, Incorporated, 2009.
- [11] F. Castagna, A. Garton, P. McBurney, S. Parsons, I. Sassoön, and E. I. Sklar, "Eqrbot: A chatbot delivering eqr argument-based explanations," *Frontiers in Artificial Intelligence*, vol. 6, 2023. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/frai.2023.1045614>
- [12] I. Sassoön, N. Kökciyan, E. Sklar, and S. Parsons, "Explainable argumentation for wellness consultation," in *Explainable, Transparent Autonomous Agents and Multi-Agent Systems*, D. Calvaresi, A. Najjar, M. Schumacher, and K. Främling, Eds. Cham: Springer International Publishing, 2019, pp. 186–202.
- [13] B. Fazzinga, A. Galassi, and P. Torroni, "An argumentative dialogue system for COVID-19 vaccine information," in *CLAR*, ser. Lecture Notes in Computer Science, vol. 13040. Springer, 2021, pp. 477–485.
- [14] —, "A preliminary evaluation of a privacy-preserving dialogue system," in *NL4AI@AI*IA*, ser. CEUR Workshop Proceedings, vol. 3015. CEUR-WS.org, 2021.
- [15] C. Hauptmann, A. Krenzer, J. Völkel, and F. Puppe, "Argumentation effect of a chatbot for ethical discussions about autonomous ai scenarios," *Knowledge and Information Systems*, Mar 2024. [Online]. Available: <https://doi.org/10.1007/s10115-024-02074-x>
- [16] L. A. Chalaguine and A. Hunter, "A persuasive chatbot using a crowd-sourced argument graph and concerns," in *COMMA*, ser. Frontiers in Artificial Intelligence and Applications, vol. 326. IOS Press, 2020, pp. 9–20.
- [17] E. Hadoux, A. Hunter, and S. Polberg, "Strategic argumentation dialogues for persuasion: Framework and experiments based on modelling the beliefs and concerns of the persuadee," *Argument & Computation*, vol. 14, pp. 1–53, 12 2022.
- [18] K. Guo, J. Wang, and S. K. W. Chu, "Using chatbots to scaffold efl students' argumentative writing," *Assessing Writing*, vol. 54, p. 100666, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1075293522000629>
- [19] S. Kotiyal, D. Thakker, and Y. H. Dubal, "Knowledge graph based intelligent conversational agent for uk immigration case work," in *NLIWoD@ESWC*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:254073040>
- [20] Z. Chen, Y. Lu, M. P. Nieminen, and A. Lucero, "Creating a chatbot for and with migrants: Chatbot personality drives co-design activities," in *DIS '20: Designing Interactive Systems Conference 2020, Eindhoven, The Netherlands, July 6-10, 2020*, R. Wakkary, K. Andersen, W. Odom, A. Desjardins, and M. G. Petersen, Eds. ACM, 2020, pp. 219–230. [Online]. Available: <https://doi.org/10.1145/3357236.3395495>
- [21] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *EMNLP-IJCNLP*, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3982–3992. [Online]. Available: <https://aclanthology.org/D19-1410>
- [22] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [23] M. Lippi, G. Contissa, A. Jablonowska, F. Lagioia, H. Micklitz, P. Palka, G. Sartor, and P. Torroni, "The force awakens: Artificial intelligence for consumer law," *J. Artif. Intell. Res.*, vol. 67, pp. 169–190, 2020. [Online]. Available: <https://doi.org/10.1613/jair.1.11519>