Health Equity in AI Development and Policy: An AI-enabled Study of International, National and Intra-national AI Infrastructures

Manpriya Dua[1], J.P. Singh[2], Amarda Shehu[3]
[1,3]Department of Computer Science, School of Computing, George Mason University,
Virginia, USA
[2]Schar School of Policy and Government, George Mason University,
Virginia, USA
[1]mdua@gmu.edu, [2]jsingh19 @gmu.edu, [3]amarda@gmu.edu

Abstract

This study examines how concerns related to equity in AI for health are reflected at the international, national, and intra-national level. Utilizing unsupervised learning over corpora of published AI policy documents and graph structurization and analysis, the research identifies and visualizes the presence and variation of these concerns across different geopolitical contexts. The findings reveal interesting differences in how these issues are prioritized and addressed, highlighting the influence of local policies and cultural factors. The study underscores the importance of tailored approaches to AI governance in healthcare, advocating for increased global collaboration and knowledge sharing to ensure equitable and ethical AI deployment. By providing a comprehensive analysis of policy documents, this research contributes to a deeper understanding of the global landscape of AI in health, potentially offering insights for policymakers and stakeholders.

Introduction

As AI technologies promise to transform healthcare access and outcomes worldwide, concerns around equitable global health in development and policy for such technologies are increasingly prominent. A recent white paper by the World Economic Forum[1] is enthusiastic about the potential of AI within healthcare and refers to it as a potential game-changer that can revolutionize immunization programs, supply chains, referrals, diagnoses, drug safety, and overall healthcare system efficiency. However, the report also warns about potential dangers and the critical need for ethical considerations in AI development.

AI is increasingly being employed in healthcare (Lee and Yoon 2021), whether for internal efficiencies, diagnoses, or treatments. Concerns around health equity are now coming into focus. AI adoption in healthcare can perpetuate inequities, such as health disparities, catalyzing conversations around the need for regulation to oversee equity dimensions in AI solutions. There is a growing understanding that we need effective policies to ensure that AI solutions have an equitable impact on global health. Better understanding of data and algorithm bias and development of mitigation strategies are paramount for equitable health outcomes. These include the need for transparency, reliability, and accountability in AI solutions and processes. Work in (Thomasian, Eickhoff, and Adashi 2021) looks at regulatory strategies for uprooting algorithmic bias in healthcare AI to help inform for federal oversight. Work in (Abràmoff, Tarver, and Loyo-Berrios 2023) describes the sources and impacts of bias in AI systems on health equity and proposes approaches for potential mitigation across the AI's Total Product Lifecycle.

As we collectively better understand the various dimensions and concerns around equitable AI for health, important questions arise on who is driving concerns and policy in this space, which nations, which international agencies, which intra-national agencies, if any? This study answers these questions via an AI-enabled approach carried over published AI policy documents. The analysis is carried out at three levels: international, national, and intra-national. The study exposes which international agencies are driving the conversation around global health equity and on which aspects of this large umbrella term. It additionally reveals which countries are championing on what aspects of health equity. The study includes a more fine-grained analysis over AI policy documents published by specific agencies and entities within countries.

Utilizing Latent Dirichlet Allocation (LDA) for text mining and graph analysis, the research identifies and visualizes the presence and variation of health equity concerns across different geopolitical contexts. By providing a comprehensive analysis of policy documents, this research contributes to a deeper understanding of the global landscape of AI in health, offering insights for policymakers and stakeholders aiming to address these critical concerns effectively.

Related Work

AI policies are being studied by researchers seeking to understand values, norms, and priorities in national AI. Work in (Aaronson 2023) leverages the OECD AI policy website and database to understand government efforts to develop AI capabilities and responsible AI. Work in (Schiff et al. 2020; Schiff 2022) explores the role of education and ethics in national AI policies. Ethics is increasingly a concern (Schiff et al. 2021; Van Berkel et al. 2020; Vesnic-Alujevic, Nasci-

[1]https://www.weforum.org/agenda/2023/06/emerging-tech-like-ai-are-poised-to-make-healthcare-more-accurate-accessible-and-sustainable/

mento, and Pólvora 2020). Work in (Van Berkel et al. 2020) subjects national AI policies of 25 countries to LDA and contrasts topic differences between countries but narrows its focus to a set of 11 ethics-related principles.

To the best of our knowledge, all conversations around health equity are largely descriptive. The published AI policy documents from international agencies, nations, and intra-national entities provide us with a unique opportunity to bring an AI-enabled approach to quantitatively expose dimensions of health equity at the international, national, and intra-national level. How does the UN understand health equity, and what do its policy documents speak to? How do the various nations advance the conversation? Do intra-national entities contribute uniquely to this conversation? These are the questions that we answer in this paper, through the methodology described next.

Methods and Materials

We analyze three separate corpora of documents to carry out our analysis in three separate settings, the international, national, and intra-national level. Our methodology pipeline consists of three modules: (1) identification of the themes/-topics present in a corpus of documents through an LDA-based algorithm; (2) identification of topics that articulate various dimensions of health equity; (3) exposure of the varying prioritization of such topics among international agencies, nations, and intra-national entities, as well geopolitical organizations around particular topics.

Datasets

International Corpus: This corpus consists of 38 documents on AI policies and priorities obtained from 6 international organizations (UN, WHO, ITU, WIPO, UNESCO and World Bank).

National AI Infrastructures Corpus: This corpus consists of 54 national AI documents obtained largely from the Organisation for Economic Co-operation and Development (OECD) website[2]. The OECD dataset is not complete, especially when there are multiple national policy documents, which we supplement. A few additional documents were obtained from the Nations' official websites. In such cases, the potentially multiple AI documents for a country are combined into a single document.

Intra-national AI Documents Corpus: This corpus consists of 167 documents obtained from multiple intra-national agencies of 38 countries. It is important to note that a few countries might not have national AI policies, but the agencies operating at the intra-national level might still have their respective AI policy documents.

Distilling Topics from a Corpus of Documents

Each corpus of documents consists of unstructured text data, and so our distillation is unsupervised. As related in section Introduction, we employ LDA due its ability to stay faithful to text (as opposed to the hallucination issue in other generative models (Xu, Jain, and Kankanhalli 2024)), popularity, and good understanding of advantages and shortcomings stress-tested in a variety of applications, including the social sciences. In the interest of space, we direct the reader to the original presentation of the LDA algorithm in (Blei, Ng, and Jordan 2003). In summary, the LDA algorithm learns in an iterative manner probability distributions, representing each document in the corpus as a probability distribution over identified topics (with the user specifying the number of topics T), and each topic itself as a probability distribution over the words in the vocabulary, which is built a priori over the corpus. Two additional, user-specified parameters, α and β control the shape of these distributions.

Due to its practical implementation as an iterative optimization algorithm, the LDA algorithm can provide different results given different initial values, and researchers have developed various methodologies over the years to address this (Hosseiny Marani and Baumer 2023). Building over related approaches, we effectively employ two steps: first, casting a wide net and obtaining various LDA models in a large configuration space of (T, α, β) hyperparameters; and then identifying a few stable topics in this space through a consensus-based approach. The latter requires comparing two sets of topics (corresponding to two different LDA models), and we employ two different similarity measures that are well-suited for comparing sets, the Jaccard Similarity and Rank-Biased Overlap) (Mantyla, Claes, and Farooq 2018; Hosseiny Marani and Baumer 2023). Utilizing these similarity measures, several iterations of the popular leader clustering algorithm identify stable, consensus topics.

Identifying Health Equity Topics

Health equity in AI development and policy brings together two umbrellas: AI-enabled advancement of health and concerns around equity in AI development. We capture the co-occurrence of these two umbrellas as follows. We define SetA keywords(bioethic, biomedical, care, clinic, clinical_trial, clinician, covid, detection, device, diagnosis, disease, doctor, drug, drug_administration, efficacy, health, healthcare, hospital, insurance, medical_device, medicine, patient, physician, prevention, protocol, therapy, treatment, vaccination) to capture various aspects of healthcare, and SetB keywords (accountable, bias, compliance, confidentiality, consent, democratic, discrimination, diversity, equitable, ethics, explainable, fairness, governance, inclusive, informed_consent, just, misinformation, moral, oversight, privacy, regulation, reliable, responsible, robust, safe, secure, sustainable, transparent, trustworthy) to capture aspects of equity in AI development. The reason we employ sets of keywords to identify topics that articulate health equity concerns is because "health equity" has only recently gained traction, and many of the national AI infrastructures precede it in their timeline. Second, what the term captures is evolving based on our collective understanding of the various considerations and so instantiations of ethical AI for healthcare.

We can identify the prominence of a given keyword in a topic, as a topic is itself a probability distribution over words in a corpus. We do so for each of the keywords in SetA and SetB and relate these results visually in section Results (for each set of topics, identified from each of the three separate corpora of documents) via heatmaps, utilizing a light yellow

[2]https://www.oecd.org/

to dark blue color scheme to indicate low to high probability (prominence). We keep the heatmap for SetA separate from that for SetB so we can identify topics that "cover" each of the two umbrellas articulated above so we can expose topics that speak to health equity in various dimensions as related through selected keywords. We refer to these topics as health equity topics.

Health Equity Communities

While we employ this analysis only for the national corpus here, in principle it can be employed at each of the three levels we investigate in this study. However, we seek to expose interesting organizations of countries around specific dimensions/aspects of health equity in AI development and policy. We do so through the nearest-neighbor graph (nnGraph) construction. Let us consider that a health equity topic τ has been identified. Each national document contains this topic with a specific probability, and we can identify alignment between two countries/documents on τ via a simple multiplication of their corresponding probabilities for τ. This is a measure of similarity that can be easily converted into a distance to obtain an nnGraph $G = (V, E)$ as follows. Each vertex $v \in V$ denotes each country/national AI policy plan. Each edge e connects two countries (u, v) if their distance (as defined above) is not above a threshold ϵ. There are different approaches to constructing an nnGraph (Cazals et al. 2014). We favor an ϵ-based approach, as we can employ the kneedle algorithm (Satopaa et al. 2011) over the distribution of the distance of the nearest neighbor of each vertex (over all vertices) to objectively obtain a value for ϵ. The obtained graph allows observing countries that articulate the same aspects around health equity and AI development. As we relate in Section Results, we utilize thickness of edges to relate country proximities. It is worth noting that nnGraphs we obtain are relatively small, so communities and potential hub countries are easily detected visually.

Results

We relate three three separate sets of analysis through which we expose concerns around health equity at three levels, the international, national, and intra-national level. In each setting, our LDA-based algorithm identifies topics that speak to the major themes present in each corpus (at each of the three levels). Co-occurrence of keywords under the two umbrellas of 'health' and 'equity' in each topic allows identifying topics that speak to health equity at each level. Analyzing the importance of each such topic at each of the three levels exposes concerns around equity at the at level of nations, where organizing communities are additionally identified, and at the international and the intra-national level, where we are able to single out specific agencies and entities. Note that multiple documents contribute to each agency at the international level and to each country at the intra-national level, hence singling out specific agencies and documents is feasible. At the national level analysis however, only one national AI policy document is being used per country, and so the nnGraph visualizations are more informative.

Health Equity in International Agencies

Figure 1 relates the topics identified over the UN documents as word clouds, with larger font relating higher probabilities. In the interest of space, we only relate the two topics (Topic 4 and Topic 8) that later analysis in Figure 2 relates as more relevant at the intersection of both health and equity.

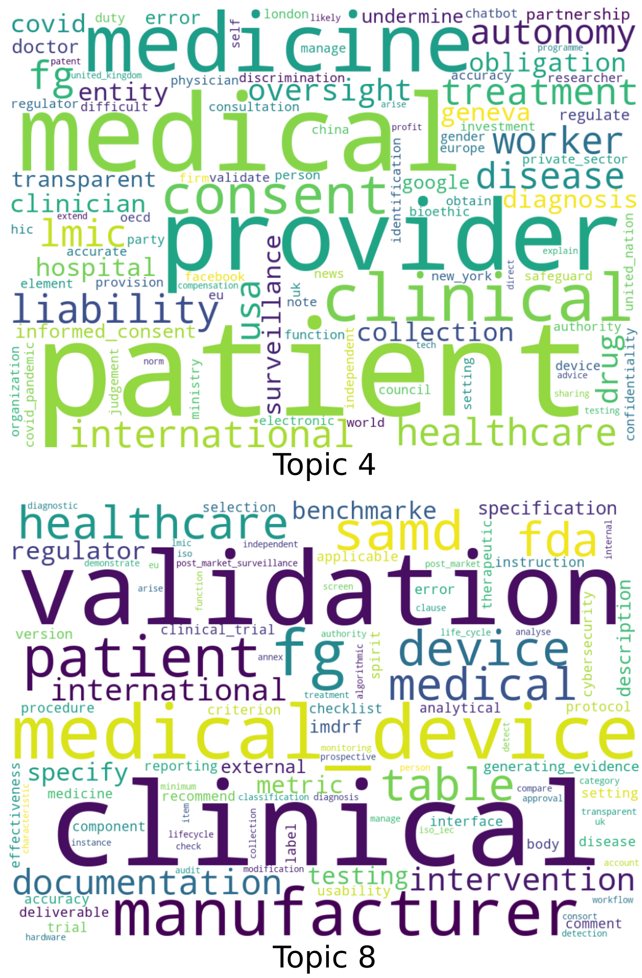

Figure 1: Topics detected over international AI infrastructure documents are visually related as word clouds, with words with higher probability shown in larger font. We restrict our attention here to Topics 4 and 8, which a keyword analysis (as related later) shows to be more relevant at the intersection of health and equity.

In Figure 2 we relate the occurrence of health (left panel) and equity (right panel) keywords in topics 4 and 8. We observe that these topics display concurrent strong signals over instantiations of health equity. Topic 4 deals with the patient-concerning aspect of health equity with words like informed consent, confidentiality, diagnosis, oversight, misinformation, accountable, trustworthy, and inclusive, whereas Topic 8 is more narrowly focused on administration of health treatments and protocols with words like treatment, detection, protocol, drug administration, drug efficacy, transparent, compliance.

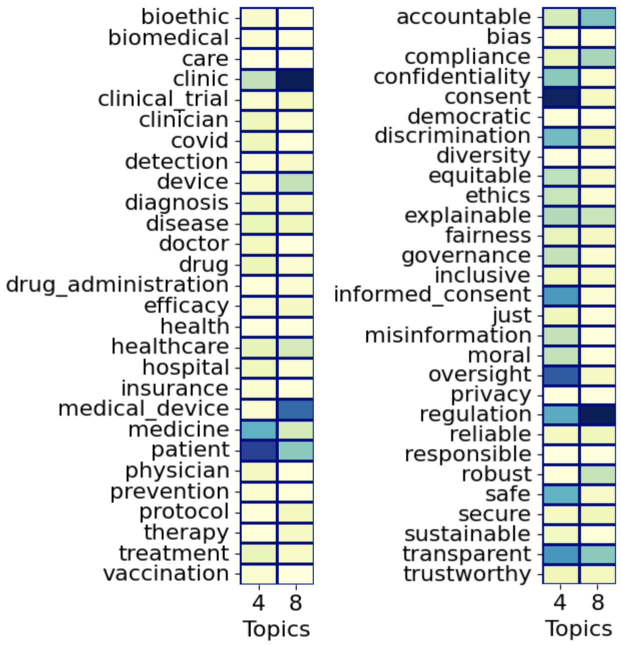

Figure 2: Heatmaps visually relate the probability of each health keyword (left) and equity keyword (right) in identified relevant topics over international AI infrastructures with darker squares implying higher probabilities.

We further analyze which international agencies are more prominent in these 'health equity'-related topics in Table 1 by listing agencies (documents) where a topic occurs with a probability no lower than 0.5. This analysis provides us with more understanding, as it contextualizes the detected topics in specific "champion" international agencies. As we see in Table 1, the World Health Organization (WHO) and the International Telecommunication Union (ITU) lead/champion the health equity topics (Topics 4 and 8), with high probabilities of Topic 4 for their respective documents. This result is not surprising, as WHO's primary objective is health. ITU's prominent presence is attributed to the focus group reports on AI for health. It is worth noting, however, that ITU has made specific contributions advancing health equity, as the results indicate. In particular, the focus on manufacturers and regulators adds a unique perspective on health equity in practical contexts of manufacturing of devices.

Health Equity in National AI Policies

We observe Topics 4 and 5 among national AI policy topics to be most relevant for health and equity, as our keyword analysis indicates. Figure 3 relates these topics as word clouds. Comparison of topics detected over national AI policies with topics detected among international agencies suggests significant overlaps, but one observes a more focused discussion of health equity in international agencies. For instance, while Topic 4 among international agencies covers various aspects of Topics 4 and 5 among national AI documents, Topic 8 in international agencies speaks more concretely and specifically to administration of health treatments, protocols, and devices.

Table 1: Top international agency documents and their probabilities over topics 4 and 8.

Documents with high probability over Topic 4	**Probability**
ITU - Ethics & Governance of AI4H	0.998597
WHO - Ethics & Governance of AI4H - WHO Guidance	0.998555
WHO - Ethics & Governance of AI4H - Guidance on Large Multi-Modal Models	0.974144
WHO - Big Data & AI for Achieving Universal Health Coverage	0.836043
Documents with high probability over Topic 8	
ITU - FGAI4H Good Practices for Health Applications of ML - Considerations for Manufacturers & Regulators	0.998063
WHO - Generating Evidence for AI Based Medical Devices	0.997881
WHO - Regulatory Considerations on AI for Health	0.995811
WHO - Regulatory Considerations on AI4H	0.995598
WHO - AI4H Considerations	0.995228
WHO - AI4H Clinical Evaluation of AI for Health	0.992248
WHO - AI4H AI Training Best Practices Specification	0.916428
ITU - Focus Group on AI for Health Whitepaper	0.697107
WHO - AI4H Ethical Considerations on AI in Dentistry	0.499029

Figure 4 relates the occurrence of health (left panel) and equity (right panel) keywords over topics 4 and 5, which we observe to display concurrent strong signals over instantiations of health equity. While there are other topics with fainter signals that can be presented, we choose to present the topics with higher health equity keyword probabilities in the interest of space. Ranked health and equity keywords in topics show the subtle nuanced differences in health equity aspects of each topic. Topic 4 more subtly relates to disease detection and drugs, whereas Topic 5 is more focused on inclusivity.

Topic-conditioned nnGraphs are now shown for Topics 4 and 5, demonstrating country organizations around such topics in Figure 5. Note that the centrality of a country results from a significantly higher topic probability of that country for a topic as compared to other countries/agencies, suggesting the country as central to the topic, and hence, identified as a champion country for the topic. Topic 4 is led/championed by Thailand, and Topic 5 is led/ championed by India as seen in Figure 5. The thickness of the edges in these nnGraphs denotes the strength of connections between countries, where thicker edges imply thicker connections (where both countries have a higher probability for the topic) and thinner edges suggest weaker connections (where the central country has a higher probability and the connected country has a lower probability for the topic).

Health Equity in Intra-national AI Policies

Figure 6 relates the 'health equity' topics identified over the intra-national documents as word clouds, with larger font relating higher probabilities. Figure 7 relates the occurrence of health (left panel) and equity (right panel) keywords over

Figure 3: Topics detected over national AI policy infrastructure documents as word clouds, restricted to Topics 4 and 5 at the intersection of health and equity.

topics 4 and 9, which we observe to display concurrent signals over instantiations of health equity (although, not as strong when compared to the national and international topics).

We further analyze which intra-national entities are more prominent in these health equity-related topics in Table 2 by listing the documents where a topic occurs with a probability no lower than 0.5. We see that topic 4 mostly has high probabilities for UK documents with contributing documents from agencies like "Center for Data Ethics and Innovation Consultation". Topic 9 is present in documents from agencies of various countries. It is worth noting that Topics 4 and 9 are quire similar in terms of health equity priorities, where Topic 4 speaks to health equity from an accountability perspective, and Topic 9 brings more specificity on ethics and explainability in healthcare.

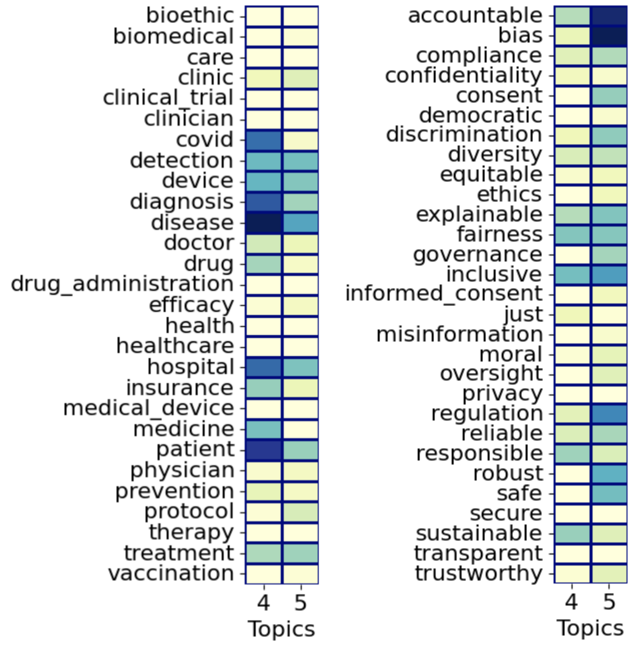

Figure 4: Heatmap visually relates the probability of each health keyword (left) and equity keyword (right) in identified relevant topics over national AI infrastructures.

Table 2: Top intra-national agency documents and their probabilities over topics 4 and 9.

Documents with high probability over Topic 4	**Probability**
UK - Online Harms White Paper	0.994350
UK - Common Regulatory Capacity for AI	0.928059
UK - Addressing Trust in Public Sector Data Use	0.891182
UK - Establishing a Pro Innovation Approach to Regulating AI	0.824286
UK - Scotland AI Strategy Consultation Report	0.816260
UK - Guidelines for AI procurement	0.801965
UK - Review Into Bias in Algorithmic Decision Making	0.774777
UK - AI Ecosystem Survey Informing National AI Strategy Summary Report	0.696890
Australia - Human Rights and Technology Discussion Paper	0.675955
Australia - AI Assurance Framework	0.656405
UK - Center for Data Ethics and Innovation Consultation	0.640238
New Zealand - Algorithm Assessment Report	0.595644
Documents with high probability over Topic 9	
EU - Impact of GDPR on AI - European Parliamentary Research Service	0.893680
US - NIST Principles for Explainable AI	0.876252
Malta - Towards Trustworthy AI	0.834389
Australia - Identity Matching Services Bill 2019	0.789467
Singapore - Model AI Governance Framework Second Edition	0.780308
UK - Data Ethics Framework	0.693026
Thailand - Digital AI Ethics Principle and Guideline	0.615740
EU - Independent High Level Expert Group on AI	0.610825
Singapore - FEAT Principles Updated 7 Feb 19	0.523708
US - AI and Algorithmic Fairness Initiative	0.499315

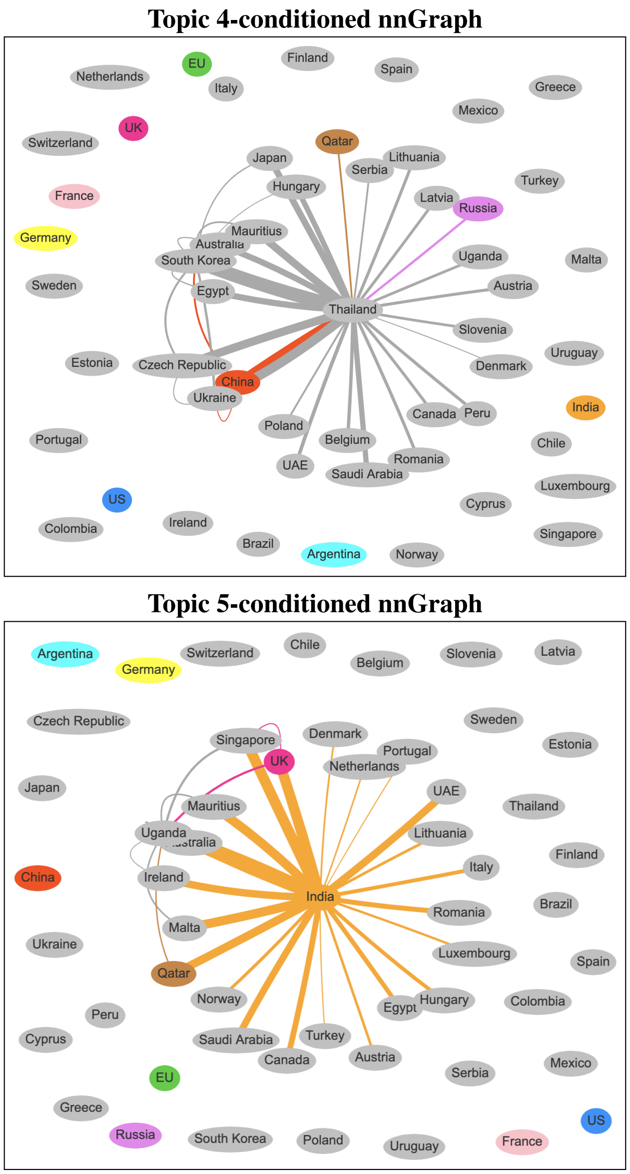

Figure 5: Topic-conditioned nnGraphs over national AI infrastructures

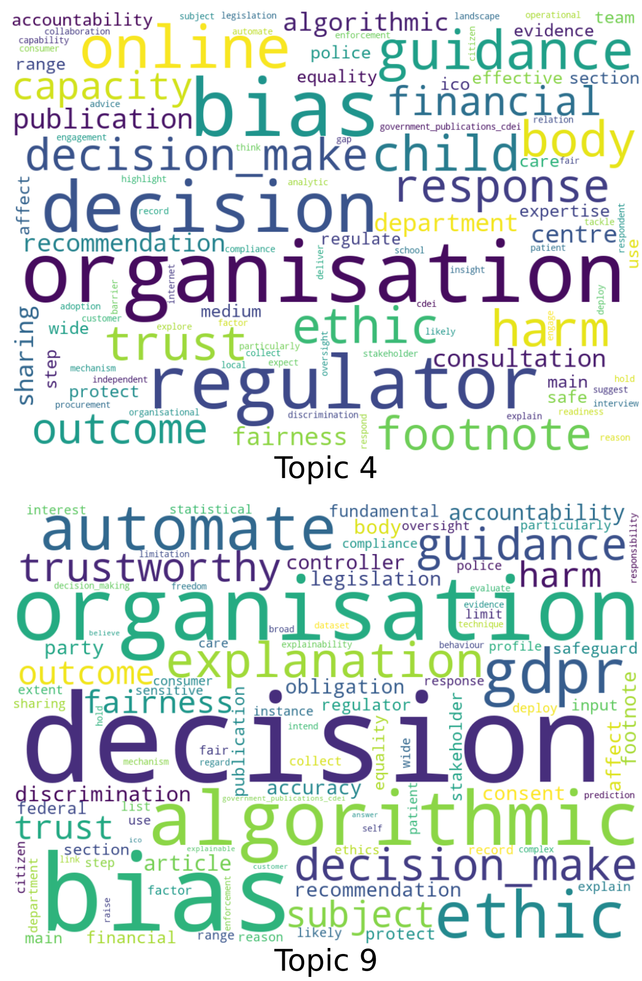

Figure 6: Intra-National AI Infrastructure topics as word clouds, restricted to Topics 4 and 9 at the intersection of health and equity.

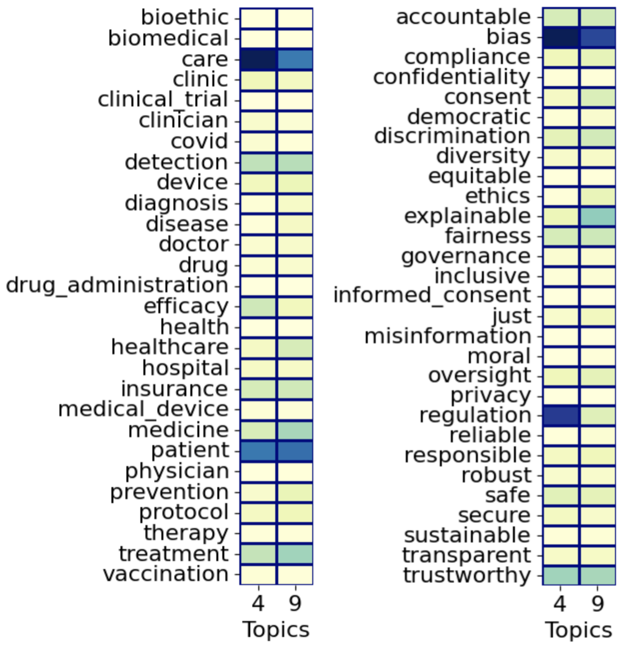

Figure 7: Heatmaps visually relate the probability of each health keyword (left) and equity keyword (right) in identified relevant topics over intra-national AI infrastructures.

Comparing Health Equity Concerns Across International, National and Intra-national AI Policies

The top 500 words of each topic (data not shown) at the three levels (international, national, and intra-national) suggest that while ethical concerns (not specifically related to health) are more prominent at the intra-national level, equitable health concerns are more prominent and have higher specificity at the international level organizations like WHO and ITU but do appear albeit in national plans, albeit with less specificity.

The results indicate that health equity is prominent at the international level, less so at the national level, and even less at the intra-national level. This is expected, as international agencies, such as WHO, specifically focus on health concerns, whereas at the national and the intra-national level, there are multiple concerns involving use of AI (like transportation, data handling, cultivating talent, etc.), which leads to mixed priorities as opposed to a health-driven priority of WHO at the international level.

Figure 8 shows the landscape of health and equity keyword pairs for identified health equity topics over the three levels of analyses. Presence in a level is tracked with a unique color in red (international), green (national) , blue (intra-national), and co-occurrences at levels are related as color summations; e.g., occurrence across all three levels is indicated with a white = red+green+blue; black tracks keywords not present at any levels. The landscape shows that bias, diversity and responsible are primarily dominant at the national and intra-national levels and absent from health-equity topics at the international level. Health-specific equity keyword such as informed consent is not prominent at the intra-national level. Some health keywords (bioethic, clinical trial, drug administration, medical device) are specific to the international level topics (red cells). The keywords privacy and health are absent from all levels (black cells) implying these aspects of health and equity are not addressed in documents at any levels. While chosen topics show both aspects, health and equity, health dominates international level topics, and equity is more prominent than health at the intra-national level.

Conclusion

This study reveals patterns and nuances of equity in AI for health. By providing a comprehensive analysis of policy documents at three levels (international, national, and intra-national), this research contributes to a deeper understanding of the global landscape of AI in health, offering insights for policymakers and stakeholders on global health equity. The study shows significant varying level of sophistication and understanding of the importance of global health equity across international, national, and intra-national AI policies. Country-level policies are more equity-oriented while international organizations - especially due to the influence of WHO and ITU - are more health-focused in their AI plans. This suggests that there is need for increased global collaboration and knowledge sharing to ensure equitable and ethical AI deployment for health.

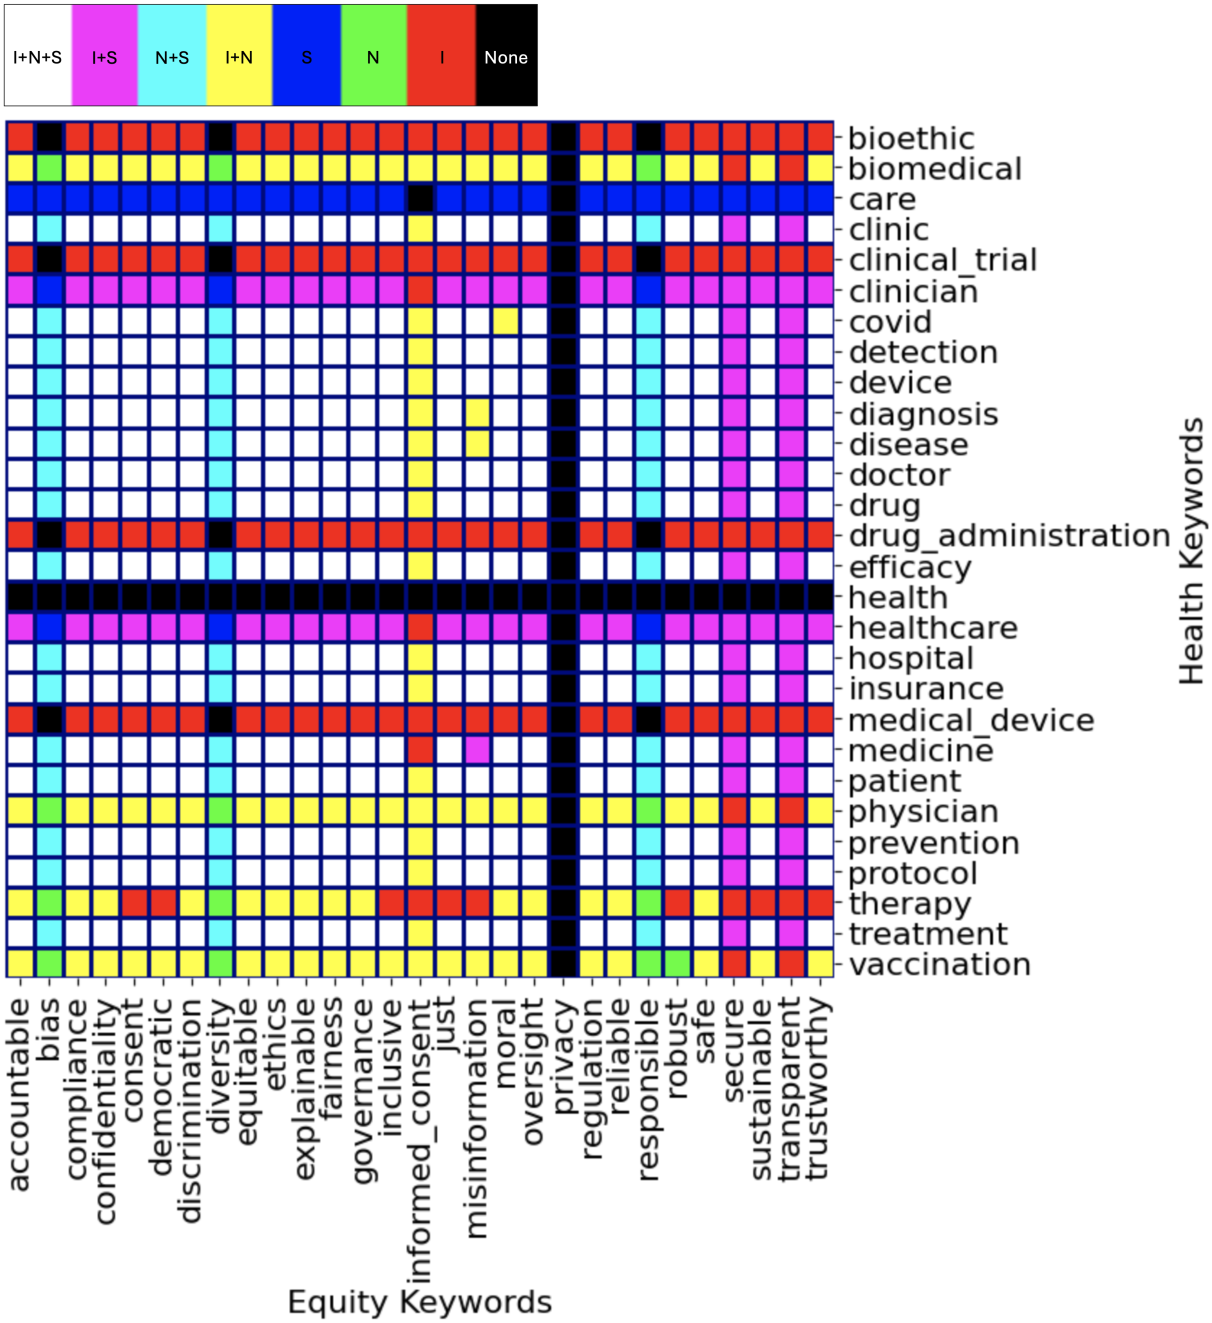

Figure 8: Health Equity Landscape: Top panel relates the colors corresponding to each level (international, national, and intra-national) and their combinations: I for international, N for national, and S for intra-national. The + sign denotes intersection of different levels. Bottom panel presents a colored landscape of each combination of health and equity keywords in health-equity topics identified at the international, national, and intra-national levels. The landscape visually relates the co-occurrence of respective health and equity keywords meeting a threshold probability in identified health-equity topics through an intersection of colors representing each level. Figure shows a dominance of health keywords whereas an absence of equity keywords at the international level. We also see the prominence of equity keywords at the national and intra-national levels.

Acknowledgements

This work is supported in part by a Department of Defense Minerva Program Award to J.P.S and AS. The authors also thank Mason's Minerva AIStrategies team.

References

Aaronson, S. A. 2023. Building trust in AI: a landscape analysis of government AI programs.

Abràmoff, M.; Tarver, M.; and Loyo-Berrios, N. 2023. Considerations for Addressing Bias in AI for Health Equity. *NPJ Digital Medicine*, 6(170).

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.*, 3: 993–1022.

Cazals, F.; Dreyfus, T.; Mazauric, D.; Roth, A.; and Robert, C. 2014. Conformational Ensembles and Sampled Energy Landscapes: Analysis and Comparison. Research Report RR-8610, INRIA.

Hosseiny Marani, A.; and Baumer, E. P. S. 2023. A Review of Stability in Topic Modeling: Metrics for Assessing and Techniques for Improving Stability. *ACM Comput. Surv.*, 56(5).

Lee, D.; and Yoon, S. N. 2021. Application of AI-Based Technologies in the Healthcare Industry: Opportunities and Challenges. *International Journal of Environmental Research and Public Health*, 18(1).

Mantyla, M. V.; Claes, M.; and Farooq, U. 2018. Measuring LDA topic stability from clusters of replicated runs. In *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ESEM '18. New York, NY, USA: Association for Computing Machinery. ISBN 9781450358231.

Satopaa, V.; Albrecht, J.; Irwin, D.; and Raghavan, B. 2011. Finding a "Kneedle" in a Haystack: Detecting Knee Points in System Behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops*, 166–171.

Schiff, D. 2022. Education for AI, not AI for Education: The Role of Education and Ethics in National AI Policy Strategies. *Int J Artif Intell Educ*, 527–563.

Schiff, D.; Biddle, J.; Borenstein, J.; and Laas, K. 2020. What's Next for AI Ethics, Policy, and Governance? A Global Overview. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, 153–158. New York, NY, USA: Association for Computing Machinery. ISBN 9781450371100.

Schiff, D.; Borenstein, J.; Biddle, J.; and Laas, K. 2021. AI Ethics in the Public, Private, and NGO Sectors: A Review of a Global Document Collection. *IEEE Transactions on Technology and Society*, 2(1): 31–42.

Thomasian, N. M.; Eickhoff, C.; and Adashi, E. Y. 2021. Advancing Health Equity with AI. *Journal of Public Health Policy*, 42(4).

Van Berkel, N.; Papachristos, E.; Giachanou, A.; Hosio, S.; and Skov, M. B. 2020. A Systematic Assessment of National AI Policies: Perspectives from the Nordics and Beyond. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, NordiCHI '20. New York, NY, USA: Association for Computing Machinery. ISBN 9781450375795.

Vesnic-Alujevic, L.; Nascimento, S.; and Pólvora, A. 2020. Societal and ethical impacts of AI: Critical notes on European policy frameworks. *Telecommunications Policy*, 44(6): 101961. AI, economy and society.

Xu, Z.; Jain, S.; and Kankanhalli, M. 2024. Hallucination is Inevitable: An Innate Limitation of Large Language Models. arXiv:2401.11817.