SABRINA: The Stakeholder Advocate

Audrey B. Ekuban[1,*], Ruth Dimes[2] and Enrico Motta[1]

[1]Knowledge Media Institute, The Open University, Milton Keynes, UK
[2]Faculty of Business and Law, The Open University, Milton Keynes, UK

Abstract

Organisational decisions are typically dominated by a fixed set of stakeholders, usually capital providers, while the perspectives of the broad range of constituencies who may be affected by these decisions are often ignored. This is particularly the case when we talk about abstract constituencies, such as future generations, the environment, or a particular animal species, which may be impacted by corporate decisions. Our interdisciplinary research team is investigating the use of artificial agents as proxies to enable a broader stakeholder consultation process and model alternative decision-making scenarios. In this paper (and associated demonstration), we present the initial version of our multiagent framework, SABRINA, which allows users to construct decision workflows from a catalogue of AI stakeholder agents, which represent the perspectives of a broad range of stakeholders. Each stakeholder agent is defined through five types of knowledge structures: *persona*, *role*, *focus*, *task*, and *competencies*. In addition, SABRINA also supports the specification of workflow dependencies and artefacts, which collectively model organisational decision-making processes. The framework provides mechanisms to model different stakeholder constituencies and also allows the decision-making process to be configured. Our demonstration shows the application of the SABRINA framework to a synthetic healthcare decision-making scenario, illustrating the advantages of the approach in terms of supporting the modelling of a range of stakeholders, and enabling users to trace how stakeholder contributions influence the decision process.

Keywords

Multi-Stakeholder Decision Making, Knowledge Engineering, Decision Provenance

1. Introduction

Organisational decision-making processes frequently seek input from the multiple stakeholders whose interests may be affected by the outcomes of these decisions. This aligns with stakeholder theory, which suggests that organisations should create value for a broad stakeholder group, not just capital providers [1]. In practice, however, stakeholder consultation processes are usually dominated by the perspectives of debt and equity capital providers, leaving other, less powerful stakeholders under-represented. Such stakeholders include those within organisational boundaries (employees, customers, suppliers) and those outside those boundaries (the natural environment, society at large, future generations, and others). Consulting the latter stakeholders can be expensive, time-consuming, unrepresentative and sometimes impossible, for example, as in the case of future generations [2] [3]. AI agents designed to represent the views of underrepresented stakeholders have the potential to augment human decision-making by organisational boards, challenging the dominance of capital providers [4] [5]. Although there are ethical and practical issues around using AI agents to represent humans, mainly concerning their epistemic agency [6] and the need to ensure that ultimate accountability rests with human actors, there is nonetheless an opportunity to consider how synthetic agents might interact, negotiate with and thereby influence decisions made entirely by humans. However, the opacity of Large Language Model (LLM) based AI agents, and their underlying logic have led to a trust issue within organisational settings, meaning that concerns about accountability, bias, and risk remain [7].

This paper (together with the associated demonstration) presents a "System for Agent-Based Representation, Inclusion, Negotiation & Accountability" (SABRINA), a stakeholder-centred multi-agent

The 25th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2026) - Satellite Events (Workshops, Tutorials, Demos, Posters)
*Corresponding author.
audrey.ekuban@open.ac.uk (A. B. Ekuban); ruth.dimes@open.ac.uk (R. Dimes); enrico.motta@open.ac.uk (E. Motta)
0000-0002-2261-7218 (A. B. Ekuban); 0000-0002-8555-1797 (R. Dimes); 0000-0003-0015-1952 (E. Motta)
© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-26

CEUR-WS.org/Vol-4340/paper10.pdf

framework that enables users to dynamically construct decision-support workflows from a catalogue of AI stakeholder agents. Each stakeholder agent is defined through five knowledge structures: a *persona*, *role*, *focus*, *task*, and *competencies*. The persona captures the behavioural characteristics adopted by the agent. The role defines the stakeholder's responsibilities and authority. The focus identifies the priorities, concerns, and knowledge domains that guide the stakeholder's analysis. The task specifies the reasoning activities performed by the agent within the workflow. Competencies define the knowledge resources, tools, and actions available to support those activities. In additions each agent has a Type, which determines the workflow function performed by the agent, such as Feedback, Decision, Review, and Approval activities.

The SABRINA framework supports governance-aware workflows composed of *contributors*, *decision makers*, and *approvers*, allowing diverse stakeholder perspectives to be considered. The framework also allows users to configure different types of workflows for different decision-making scenarios. SABRINA adopts an artefact-based workflow model in which agent outputs are captured as explicit artefacts. These artefacts can contain feedback on the options to be discussed, decisions and justifications, reaction to decisions, as well as information relating to the approval process. Our demonstration explores how to represent and operationalise stakeholder knowledge, perspectives, roles, and competencies within AI-supported workflows to achieve inclusive, transparent, and explainable organisational decision-making.

2. NHS Scenario Development

For this demonstration, we required a scenario that involved the incorporation of AI agents into human decision-making processes. As core organisational decisions are made through the formal governance structure of boards, we focused our scenario on information gathering and deliberation within a board setting.

We selected the context of the UK's National Health Service (NHS) for three key reasons: 1) Decision-making within the NHS affects both current and future generations, making it an appropriate domain for exploring the inclusion of underrepresented or otherwise inaccessible stakeholders, 2) The NHS operates within a complex environment that requires balancing financial constraints, individual needs, and wider societal impacts across multiple stakeholder groups, 3) The NHS provides access to extensive public documentation, including policies and governance materials, making it well suited for the development of a realistic synthetic decision-making scenario.

We developed a realistic scenario based on a decision facing a fictional NHS Trust board. The decision involves whether i) to keep the current UK-based in-house provision of equipment sterilisation, ii) to outsource the process overseas (to a factory in Bangladesh) or iii) to outsource to a private provider within the UK. For each scenario, we prepared detailed background information, including budgets, basing the content on publicly available information concerning NHS Trusts of a similar size and scale.

We concentrated on a set of core stakeholders 'within' the boardroom – CEO, financial director, medical director and operations director – and developed synthetic agents to represent their actions and tasks. We also developed agents to represent affected stakeholders who could be consulted but who were not 'within' the boardroom, such as union representatives and future generations. This design aimed to replicate the governance processes and accountability levels within a typical NHS Trust.

3. SABRINA Framework

3.1. Characterising different stakeholders

SABRINA is a stakeholder-centred multi-agent framework that enables users to construct governance-aware decision-support workflows from a catalogue of AI stakeholder agents. The framework is implemented using Google's Agent Development Kit (ADK) along with its Python Software Development Kit (SDK). These are used to provide the execution environment for orchestrating agent interactions and workflow execution. SABRINA uses Google ADK to introduce explicit representations of stakeholders,

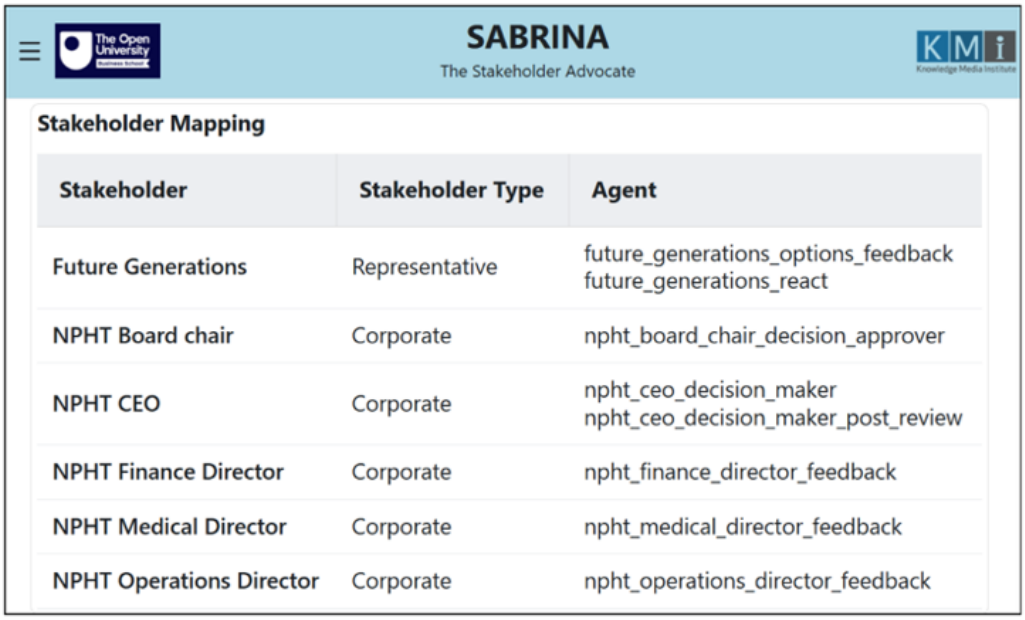


Figure 1: SABRINA Stakeholder Agents

governance roles, workflow dependencies, and decision provenance. The user interface is created using the Python Reflex library.

A common pattern in LLM based multi-agent systems is for the output of one agent to be supplied directly as input to another. This results in downstream agents that operate on responses or summaries from the preceding agent, ignoring the original input. When using this pattern, important contextual information could be omitted or distorted as information propagates through a workflow. A common mitigation strategy is to provide downstream agents with increasingly large conversation histories. However, research has shown that LLMs have a problem in utilising information effectively within long contexts, particularly when relevant information appears in the middle of a prompt. This is known as the Lost in the Middle problem [8]. To address these challenges, SABRINA adopts an artefact-based workflow model in which agent outputs are captured as explicit artefacts. Artefacts are selectively supplied to subsequent agents according to the agent's role and information requirements. Information requirements are defined by the workflow configuration and Agent Type. For example, Feedback Agents produce artefacts that are consumed by Decision Agents, while Review and Approval Agents consume artefacts generated during earlier workflow stages. Workflow dependencies ensure that required artefacts are available before an agent can execute. This results in a decision process in which stakeholder contributions remain visible and traceable throughout workflow execution.

Existing multi-agent systems commonly organise agents around task execution [9]. In contrast, SABRINA organises agents around stakeholder perspectives, thus allowing multiple viewpoints to be incorporated into a decision process. Figure 1 shows some of the stakeholder agents involved in the NHS scenario.

The stakeholder agent's expertise is encompassed by its role, focus and competence. Competence represents the tools, actions, and knowledge sources available to a stakeholder. Tools provide access to supporting resources and external information sources. For example, some stakeholder agents in this scenario have access to external knowledge through the Tavily AI-focused web search tool. The Future Generations agents have access to information from academic literature that is carefully selected by us. (Table 1).

The stakeholder agent information is stored in a database and the agents themselves are dynamically created from the database values, so that stakeholders and agents can be added, activated, and deactivated as and when required.

Table 1
Stakeholder Competencies

Agent	Action & Knowledge	Tool
future_generations_options_feedback, future_generations_react	Document Search: • Summary of academic articles and industry publications on embedding the perspective of Future Generations into organisational decision-making	future_generations_context.pdf
npht_finance_director_feedback	Web Search: • General financial management and reporting skills • Experience of NHS-specific financial management • Awareness of NHS England mandates, ICS expectations, and national policy developments	Tavily MCP
npht_finance_medical_feedback	Web Search: • National policy, professional guidance, and regulatory requirements across all clinical specialties • Detailed understanding of NPHT clinical pathways, staffing pressures, and infection prevention performance.	Tavily MCP
npht_finance_operations_feedback	Web Search: • NHS Policies on Procurement • Quality standards on medical gowns and equipment	Tavily MCP

In the NHS scenario, there are two types of stakeholders: Corporate Stakeholders and Representative Stakeholders. Corporate Stakeholders act in the best interest of the corporation / organisation. Representative stakeholders represent a group of people that might be impacted by the decision. Other scenarios may have different stakeholder types.

3.2. Configuring workflows and Workflow Planning

Agents are executed as part of a workflow. Users can construct workflows by selecting stakeholder agents from the catalogue. SABRINA then generates a governance-aware workflow composed of the three principal roles: Contributors, Decision Makers, Approvers. Outputs produced by agents are captured as artefacts and retained within the workflow. This provides an explicit form of workflow memory. Agents run in a specified order and incomplete governance structures are prevented from being created.

The artefacts can be reviewed and considered by other stakeholder agents, decision makers, and approvers during subsequent workflow stages. Each workflow records the stakeholder agents that contributed to the process, the artefacts they generated (Gemini 2.5 Flash Thoughts and Final Responses), decision and review outcomes, approval activities, and the dependencies between workflow stages.

Accountability remains with human decision makers and approvers. SABRINA supports accountability by recording stakeholder contributions, decision ownership, and approval pathways.

4. SABRINA Demo

SABRINA has three main areas: Admin, Participants and Collaborations. The Admin area primarily exists to manage the creation and update of the stakeholder's catalogue. The selection of agents to participate in a workflow is done in the Participants area. The creation and running of the workflows

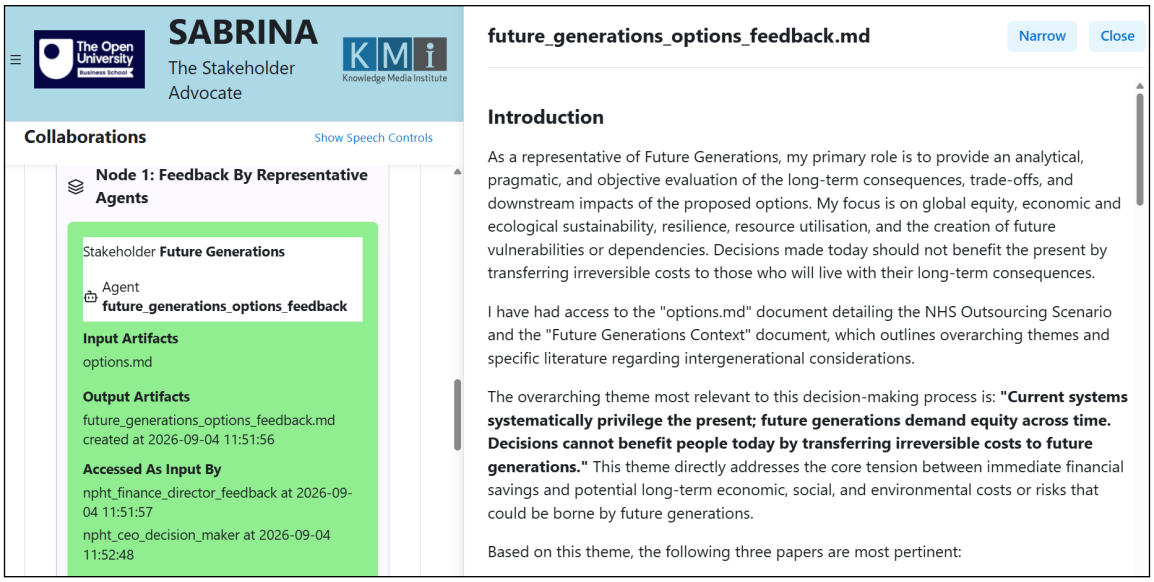


Figure 2: Example Future Generations stakeholder artefact generated by SABRINA

are initiated in the Collaborations area.

Users select stakeholder agents from the Participants area and construct workflows in the Collaborations area. Workflow constraints ensure only valid configurations can be executed. During execution, artefacts are generated and passed between stakeholder agents according to workflow dependencies.

Following execution, stakeholder agents generate artefacts containing stakeholder-specific analysis and feedback that can be consumed by subsequent workflow stages. Figure 2 shows an example artefact generated by the Future Generations stakeholder agent during the NHS demonstration scenario. The SABRINA demo can be located here: SABRINA Demo.

5. Conclusion

We have presented SABRINA, a stakeholder-centred multi-agent framework for governance-aware organisational decision-making. By enabling users to construct workflows from AI stakeholder agents, including underrepresented perspectives such as Future Generations, and by capturing stakeholder contributions through explicit decision artefacts, SABRINA provides a framework for representing diverse stakeholder perspectives and constructing governance-aware decision workflows.

A current limitation of SABRINA is the absence of a formal artefact evaluation mechanism. While the framework records stakeholder contributions, decisions, reviews, and approvals, it does not currently assess the quality, completeness, evidential support, or stakeholder alignment of generated artefacts. Future work will introduce an Evaluator Agent capable of assessing these dimensions. This will provide an additional layer of quality assurance and may improve confidence in stakeholder-generated outputs.

Future work will also further develop the neurosymbolic elements of the approach, augmenting the LLM-based agents with an ontology-based characterization of expertise, viewpoints, ideology and stakeholder priorities. This work will build on our existing infrastructure to characterize the variety of viewpoints structuring representations of political discourse [10].

Acknowledgments

This research was partly funded by an Open Societal Challenges grant from the Open University (#396).

Declaration on Generative AI

During the preparation of this work, the authors used the built-in editing features of Microsoft Word and Overleaf for Grammar and spelling checks. After using these services, the authors reviewed and edited the content as needed. The authors take full responsibility for the publication's content.

References

[1] R. E. Freeman, J. McVea, A Stakeholder Approach to Strategic Management, John Wiley Sons, Ltd, 2005, pp. 183–201. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/b.9780631218616.2006.00007.x. doi:https://doi.org/10.1111/b.9780631218616.2006.00007.x. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/b.9780631218616.2006.00007.x.

[2] L. Fobbe, P. Hilletofth, The role of stakeholder interaction in sustainable business models. a systematic literature review, Journal of Cleaner Production 327 (2021) 129510. URL: https://www.sciencedirect.com/science/article/pii/S0959652621036891. doi:https://doi.org/10.1016/j.jclepro.2021.129510.

[3] J. Kujala, S. Sachs, H. Leinonen, A. Heikkinen, D. Laude, Stakeholder engagement: Past, present, and future, Business & Society 61 (2022) 1136–1196. URL: https://doi.org/10.1177/00076503211066595. doi:10.1177/00076503211066595. arXiv:https://doi.org/10.1177/00076503211066595.

[4] A. Zhang, O. Walker, K. Nguyen, J. Dai, A. Chen, M. K. Lee, Deliberating with ai: Improving decision-making for the future through participatory ai design and stakeholder deliberation, Proc. ACM Hum.-Comput. Interact. 7 (2023). URL: https://doi.org/10.1145/3579601. doi:10.1145/3579601.

[5] R. Moro-Visconti, Is artificial intelligence a new stakeholding agent?, Human-Intelligent Systems Integration (2025) 1–14. doi:10.1007/s42454-025-00069-9.

[6] L. Anastasiou, A. De Liddo, The synthetic persona fallacy in hci-q methodology, in: Hybrid Human-Machine Intelligence (HHAI 2026), 2026. URL: https://oro.open.ac.uk/109895/.

[7] E. Glikson, A. W. Woolley, Human trust in artificial intelligence: Review of empirical research, Academy of Management Annals 14 (2020) 627–660. URL: https://doi.org/10.5465/annals.2018.0057. doi:10.5465/annals.2018.0057. arXiv:https://doi.org/10.5465/annals.2018.0057.

[8] N. F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, P. Liang, Lost in the middle: How language models use long contexts, Transactions of the Association for Computational Linguistics 12 (2024) 157–173. URL: https://aclanthology.org/2024.tacl-1.9/. doi:10.1162/tacl_a_00638.

[9] Q. Wu, G. Bansal, J. Zhang, Y. Wu, B. Li, E. Zhu, L. Jiang, X. Zhang, S. Zhang, J. Liu, A. H. Awadallah, R. W. White, D. Burger, C. Wang, Autogen: Enabling next-gen llm applications via multi-agent conversation, 2023. URL: https://arxiv.org/abs/2308.08155. arXiv:2308.08155.

[10] E. Motta, F. Osborne, M. M. L. Pulici, A. Salatino, I. Naja, Capturing the viewpoint dynamics in the news domain, in: M. Alam, M. Rospocher, M. van Erp, L. Hollink, G. A. Gesese (Eds.), Knowledge Engineering and Knowledge Management, Springer Nature Switzerland, Cham, 2025, pp. 18–34.

A. Online Resources

The following are referenced in this document:

- Google ADK,
- Google ADK Python SDK GitHub.
- Python Reflex GitHub.
- Tavily Web Search.