



City Research Online

City St George's, University of London

Citation: Cook, D., Cook, E., Cullen, K., Zachos, K., McManus, S., Bellis, M. A., Feder, G. S. & Maiden, N. (2026). Artificial intelligence in critical synthesis of public health responses to violence: A novel application to UK violence prevention policy. *Public Health*, 255, 106258. doi: 10.1016/j.puhe.2026.106258

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/37390/>

Link to published version: <https://doi.org/10.1016/j.puhe.2026.106258>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).



Short Communication

Artificial intelligence in critical synthesis of public health responses to violence: A novel application to UK violence prevention policy

Darren Cook^{a,*}, Elizabeth Cook^a, Kimberly Cullen^a, Konstantinos Zachos^b, Sally McManus^a, Mark A. Bellis^c, Gene S. Feder^d, Neil Maiden^b

^a Violence and Society Centre, City St George's, University of London, London, UK

^b Institute for Creativity and AI, City St George's, University of London, London, UK

^c Public Health Institute, WHO Collaborating Centre for Violence Prevention, Liverpool John Moores University, Liverpool, UK

^d Centre for Academic Primary Care, University of Bristol, Bristol, UK

ARTICLE INFO

Keywords:

Artificial intelligence
Large language models
Public health policy
Violence and abuse
Policy analysis
Evidence synthesis

ABSTRACT

Objectives: Artificial intelligence (AI) systems are increasingly applied in public health, yet their use for analysing fragmented, multi-sectoral policy landscapes remains underdeveloped. This study aimed to describe the development and preliminary exploration of an AI-enabled tool designed to synthesise evidence from violence-related policy documents in the UK.

Study design: An exploratory, proof-of-concept case study.

Methods: A corpus of publicly available UK policy and strategy documents on violence (N = 343) was compiled through expert review, manual searches of government and third sector organisation websites, and automated web scraping. We used the corpus to train an existing AI framework and deployed it through a question-answer interface. Stakeholders were invited to pose natural-language questions about violence policy and consider the system's utility and the usefulness of its outputs.

Results: Stakeholders reported that the AI-enabled tool facilitated flexible interrogation of violence-related policy documents and supported identification of recurring framings, sectoral differences, and potential policy siloes. Feedback indicated that the system improved the efficiency and transparency of cross-sectoral policy analysis, particularly in the initial stages of an inquiry.

Conclusions: This short communication provides early insight into the potential of AI-enabled tools to support public health policy analysis by structuring and synthesising complex documentary evidence. Such functionality is particularly relevant in areas requiring cross-sectoral collaboration. Further work is required to formally evaluate performance, assess bias, and explore impacts of AI on real-world decision-making prior to wider implementation.

1. Introduction

Artificial intelligence (AI) systems are increasingly applied in public health, offering new ways to synthesise complex information and support decision-making.¹ Recent commentaries highlight the potential of AI to improve the efficiency, transparency, and responsiveness of health systems.² Many applications have focused on service delivery, such as AI-powered chatbots, data surveillance and monitoring, and tracking social media interactions for emerging risks,³ with less attention paid to how AI might support policy analysis.⁴ This is especially true for the

violence prevention sector, where AI is gaining traction as a solution for triaging help-seeking calls, detecting threatening messages, predicting conflict and improving police data, but not for understanding the policy landscape.⁵

Policy responses to violence and abuse present distinct challenges for AI, given they are dispersed across sectors, organisations, and governance levels, resulting in large, heterogeneous collections of documents that are difficult to analyse systematically.⁶ Responsibilities for policy responses to violence span multiple government departments, local authorities, healthcare and education providers, voluntary

This article is part of a special issue entitled: AI and Public Health published in Public Health.

* Corresponding author.

E-mail address: Darren.Cook@citystgeorges.ac.uk (D. Cook).

<https://doi.org/10.1016/j.puhe.2026.106258>

Received 8 January 2026; Received in revised form 4 March 2026; Accepted 17 March 2026

Available online 13 April 2026

0033-3506/© 2026 The Authors. Published by Elsevier Ltd on behalf of The Royal Society for Public Health. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

organisations, unions, employers, and criminal justice agencies. Across these sectors, violence can be defined, framed, and addressed in markedly different ways, making it difficult for analysts to identify common priorities, gaps, or inconsistencies using traditional review methods.

Existing approaches to policy analysis, including manual document review or narrative synthesis, are time-intensive and struggle to capture change over time and cross-sectoral patterns at scale. Consequently, opportunities for greater collaboration across siloed policy domains may be overlooked. There is, therefore, a need for methodological tools that can support the systematic, transparent, efficient and scalable interrogation of policy documents while generating actionable insights for policy analysts.

To address this gap, we developed a novel AI-enabled tool to support exploratory analysis of policy documents through free-form questioning. Building on prior methodological work,^{7,8} this prototype aims to enable users to interrogate a corpus of policy texts to examine how violence is conceptualised, including the identification of victims and perpetrators, the framing of different forms of violence (e.g. bullying, harassment, assault, coercive control), and the types of solutions proposed across sectors. By mapping similarities and differences in these framings, the tool is designed to identify potential policy siloes and areas of misalignment.

This study aims to introduce the development of the tool and report a preliminary, proof-of-concept exploration of its usability and potential utility for public health policy analysis. Focusing on violence-related policy, we examine whether the tool can support the identification of cross-sectoral patterns and generate insights that may be difficult to obtain through conventional methods.

The paper makes three contributions to emerging AI-enabled approaches to public health and policy analysis: we introduce a novel, document-grounded AI tool designed to support exploratory investigation of policy corpora in the violence prevention sector; we demonstrate the potential of the approach to identify cross-sectoral themes that may be challenging to detect through manual review; we provide an initial proof-of-concept evaluation of the tool's usability and potential utility for researchers and analysts working with large, heterogeneous policy documents.

2. Methods

The case study comprised three stages: (1) collating a corpus of eligible policy documents on violence prevention, (2) adapting an existing AI framework using the gathered documents, and (3) stakeholder validation considering the perceived usability and utility of the tool.

2.1. Violence policy corpus

Document selection was based on predefined inclusion criteria:

- publicly accessible,
- published between January 1, 2010 and June 1, 2025,
- UK-based,
- focused on violence or violence prevention,
- contained at least one pre-specified keyword (e.g. 'coercive control', 'child abuse', 'knife crime'; see Supplementary Materials).

We identified an initial set of 119 documents through a manual search of grey literature. Automated web scraping of appropriate web domains (e.g., <https://www.gov.uk>) extended the corpus to 343 documents. The scraper was built in python 3.9 using the requests and BeautifulSoup libraries. Candidate URLs were filtered using a combination of lexical rules, regular expressions, named entity recognition, and sentence embeddings to identify content likely to meet the inclusion criteria. All documents were inspected by two members of the research team to ensure relevance. The final corpus comprised $N = 343$

documents (19% HTML; 81% PDF), with 57% from government and 43% from third-sector organisations. (see Fig. 1).

2.2. Model architecture and domain adaptation

The underlying computational models were initially developed as a framework for identifying innovation opportunities from text corpora. To adapt the system to the present study, the models were retrained on the violence-prevention corpus described above, enabling the identification of domain-specific thematic clusters. The pipeline combines topic modelling, dimensionality reduction, and clustering to transform a collection of semantically related documents into distinct *opportunity spaces* that are subsequently transformed into natural language descriptions using a Large Language Model (LLM). A complete description of the architecture is provided elsewhere.^{7,8}

2.3. System interface and output generation

Users interact with the system through a web-based interface that accepts open-ended queries (see Fig. 2). A question can be asked of government-only documents, third-sector documents, or the full corpus. After submission of a question, the system identifies the most relevant documents and generates structured natural language summaries in response.

The interface provides different response modes based on whether a user requires a constrained or open output. Advanced response styles also include options to adjust the tone and length of a response. Users may also select between a descriptive mode that provides an information synthesis, and an analytic mode that performs critical analysis and interpretation.

2.4. Preliminary usability and utility assessment

An interdisciplinary consortium of 30 violence-prevention researchers (including core members of the research team), alongside eight policymakers and service providers, were invited to participate. Disciplines represented included criminology, demography, epidemiology, economics, psychology, and medicine (see Supplementary Materials for participating institutions).

Participants were provided with access to the tool and invited to submit or test questions relevant to their research or professional practice. Questions reflected diverse disciplinary perspectives on violence prevention and enabled assessment of the tool's capacity for description, retrieval, synthesis, and critical analysis. Engagement ranged from single queries to extended interactive exploration.

Given the exploratory nature of the study, validation focused on qualitative assessment of perceived accuracy, interpretability, and practical usefulness for analysing cross-sectoral policy framings. Feedback was collected via an online form, email, a recorded hybrid workshop, and small group discussions. Responses were collated in Excel and analysed thematically by two members of the research team.

3. Results

Collated stakeholder feedback highlighted several perceived affordances and limitations relevant to policy analysis.

Outputs were generally viewed as well-grounded in the underlying source documents. Syntheses aligned closely with the corpus, and the inclusion of document references and page-level citations supported credibility assessments. Corpus coverage statistics were considered particularly helpful when judging the robustness of responses.

Confidence in the system was closely tied to its bounded and transparent design. Respondents reported greater trust when posing exploratory or sensitive questions compared with general-purpose generative AI systems, attributing this to the tool's closed architecture, which restricted queries to the curated corpus. Explicit identification of the

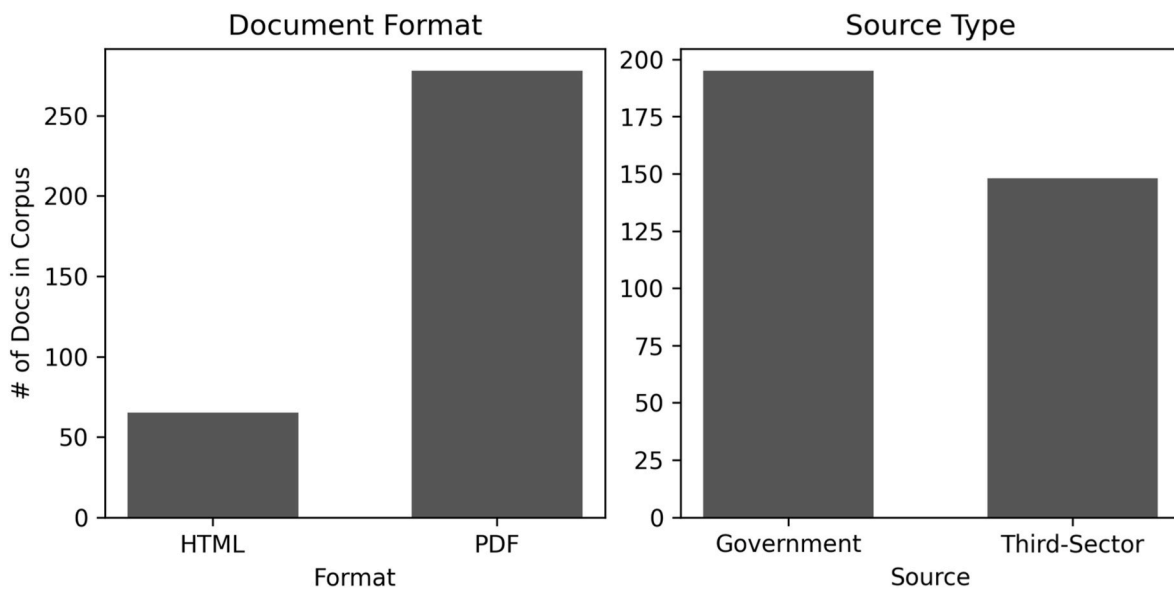


Fig. 1. Distribution of documents in the violence policy corpus by format (HTML, PDF) and source (government, third sector) (N = 343).

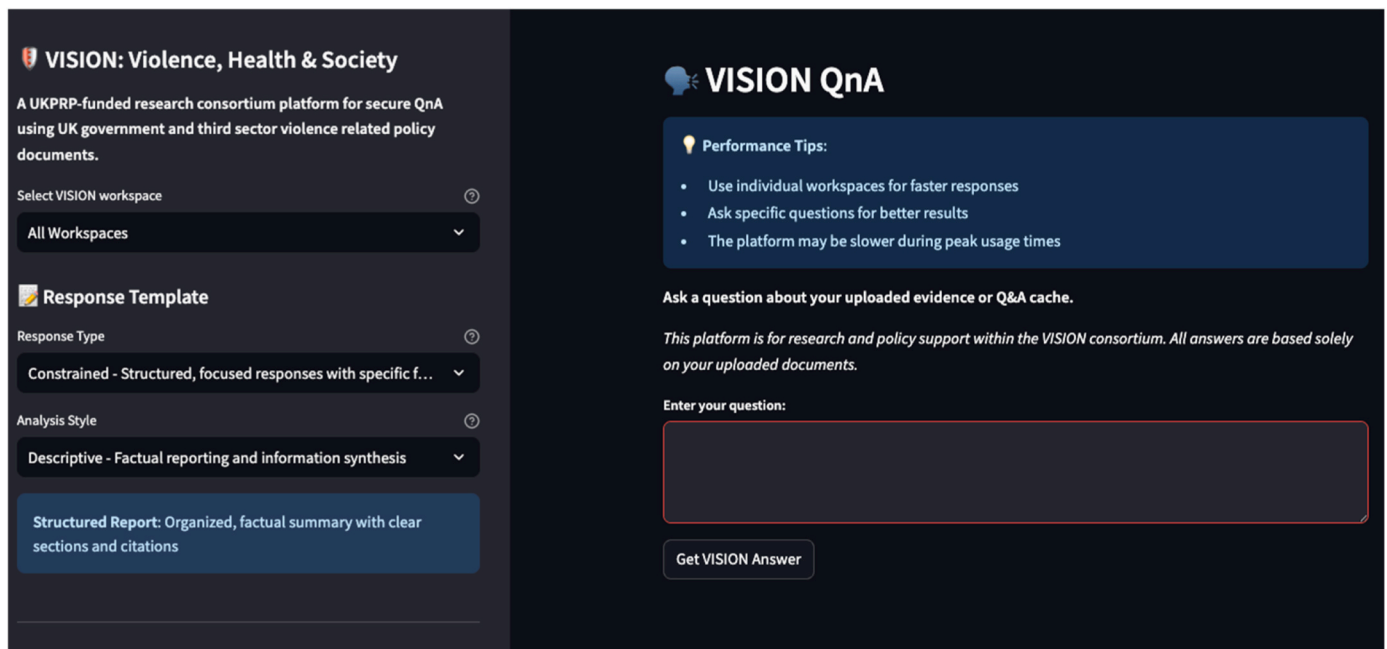


Fig. 2. Demonstration of the user interface.

documents selected for analysis, alongside coverage statistics, was regarded as central to responsible use and appropriate interpretation.

Stakeholders also highlighted the tool's capacity to surface meaningful cross-sectoral differences in policy framings. For example, contrasts between central government and third-sector documents were identified, with the latter placing greater emphasis on survivors and victim support. These distinctions were described as consistent with established understandings in the field.

Clear gains in analytic efficiency were reported, particularly at early stages of inquiry. The system enabled rapid orientation within a fragmented policy landscape and provided structured starting points for more detailed analysis. However, participants emphasised that the tool should be used as an investigative aid rather than a substitute for manual analysis or domain expertise.

Limitations were primarily associated with corpus composition.

Maintaining relevance would require regular updates as the policy landscape evolves. Weaker synthesis in certain areas, particularly local government, was attributed to gaps in available documentation rather than deficiencies in system design. Clear signalling of these constraints was considered important for managing user expectations.

Stakeholders additionally noted a learning curve in question formulation. Variation in outputs according to query phrasing underscored the need for clearer guidance and illustrative examples to support effective and responsible use.

4. Discussion

4.1. Implications for violence-related policy analysis

Policy responses to violence are undergoing scrutiny in the UK,

coinciding with the recent publication of an updated cross-government strategy addressing violence against women and girls.⁹ This renewed focus places increased demands on researchers and policymakers to rapidly synthesise large and fragmented bodies of policy evidence spanning multiple sectors and both local and national government. Traditional approaches to policy review formed around a wholly manual approach may struggle to meet these demands within policy-relevant timeframes.

4.2. Limitations

Our aim in this work was not to test effectiveness or validate the tool, but to illustrate an emerging analytic capability and its potential role within policy-oriented research workflows. Future evaluation work needs to examine gaps, biases and potential implications for inequalities.

4.3. Future work

Future work should include a systematic and formal evaluation, examining the accuracy and relevance of model outputs as well as the practical usability of the tool for both researchers and government-sector analysts and policymakers. Such evaluation would help establish the conditions under which AI-enabled policy analysis can add meaningful value to existing evidence review processes.

While the present study focuses on a single policy domain, applying the approach to larger and more diverse policy corpora represents a further avenue for research, with the potential to test scalability and generalisability across different areas of public health and social policy.

4.4. Conclusion

This short communication contributes by documenting the early application of an AI-enabled tool designed to support exploratory policy analysis. By demonstrating how a document-grounded, closed-domain AI system can be used to interrogate policy framings and identify potential siloes, this work addresses a gap in current public health applications of AI, specifically in the context of violence prevention.

Ethical statement

The project was approved by the committee at City St George's, University of London that considers medium risk applications (ETH21220–299).

Funding

This work was supported by the UK Prevention Research Partnership [Violence, Health and Society; MR-V049879/1], which is funded by the British Heart Foundation, Chief Scientist Office of the Scottish Government Health and Social Care Directorates, Engineering and Physical Sciences Research Council, Economic and Social Research Council, Health and Social Care Research and Development Division (Welsh Government), Medical Research Council, National Institute for Health Research, Natural Environment Research Council, Public Health Agency (Northern Ireland), The Health Foundation, and Wellcome.

Declaration of competing interest

The authors have no conflicts of interest.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.puhe.2026.106258>.

References

1. Olawade DB, Wada OJ, David-Olawade AC, Kunonga E, Abaire O, Ling J. Using artificial intelligence to improve public health: a narrative review. *Front Public Health*. 2023;11. <https://doi.org/10.3389/fpubh.2023.1196397>.
2. Ramezani M, Takian A, Bakhtiari A, Rabiee HR, Ghazanfari S, Mostafavi H. The application of artificial intelligence in health policy: a scoping review. *BMC Health Serv Res*. 2023;23(1):1416. <https://doi.org/10.1186/s12913-023-10462-2>.
3. Panteli D, Adib K, Buttigieg S, et al. Artificial intelligence in public health: promises, challenges, and an agenda for policy makers and public health institutions. *Lancet Public Health*. 2025;10(5):e428–e432. [https://doi.org/10.1016/S2468-2667\(25\)00036-2](https://doi.org/10.1016/S2468-2667(25)00036-2).
4. Misuraca G, van Noordt C. *Overview of the Use and Impact of AI in Public Services in the EU*. Publications Office of the European Union; 2020. <https://data.europa.eu/doi/10.2760/039619>. Accessed December 16, 2025.
5. Pauwels E. Artificial intelligence and data capture technologies in violence and conflict prevention: opportunities and challenges for the international community. *Global Center on Cooperative Security*; 2020:2022. <https://www.jstor.org/stable/resre/p27551>. Accessed December 31, 2025.
6. Clemmow C, Rottweiler B, Pearson E, Gill P. Public or private violence? Understanding the overlap between intimate partner abuse and susceptibility to violent extremism. *Aggress Violent Behav*. 2025;82. <https://doi.org/10.1016/j.avb.2025.102055>.
7. Maiden N, Zachos K, Lockerbie J, Petrianakis K, Brown A. A computational model and tool for generating more novel opportunities in professional innovation processes. *arXiv preprint arXiv:251020402*. 2025 Oct 23.
8. Maiden N, Zachos K, Petrianakis K, et al. Beyond just generative AI for discovering software opportunities. *IEEE Software*. 2025 Feb 13.
9. UK Government. Freedom from violence and abuse: a cross-government strategy to build a safer society for women and girls. https://assets.publishing.service.gov.uk/media/6943d2da9273c48f554cf592/VAWG_01_Strategy_FINAL_171225_WEB.pdf; 2025. Accessed December 31, 2025.