Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

ForcesAtlas: construction of a multi-source, AI-assisted evidence discovery database of Armed Forces research

Daniel Leightley

daniel.leightley@kcl.ac.uk

King's College London

Rhea Varghese

King's College London

Toslima Khatun

King's College London

Iain Marshall

King's College London

Simon Dymond

Swansea University

Amos Simms

King's College London

Alan Finnegan

Northumbria University

Charlotte Williamson

King's College London

Research Article

Keywords: Database, Armed Forces, Military Resource, Repository

Posted Date: October 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11376106/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare potential competing interests as follows: DL is a member of the UK Army Reserve Force. This work has been undertaken as part of his civilian employment. AS is a

full-time member of the UK Armed Forces.

Abstract

Background

Research concerning serving Armed Forces personnel, veterans and their families spans multiple disciplines and is distributed across databases with differing coverage, terminology and indexing methods. This fragmentation makes evidence discovery challenging, particularly where military-connected populations are described inconsistently. ForcesAtlas was developed as an accessible, continuously updated evidence-discovery platform to bring relevant research together across scholarly sources.

Methods

ForcesAtlas identifies candidate articles from PubMed, Europe PMC, OpenAlex and Semantic Scholar, with Crossref used for enrichment. Matching and deduplication are undertaken at article ingestion. A large language model is used to assess relevance from title, publication year and abstract, along with full-text processing (where are available). The platform is implemented using Node.js, Express and PostgreSQL. Classification reliability was assessed using a balanced sample of 2,000 records independently reviewed by at least two reviewers. Known-item coverage was assessed against included publications from five systematic reviews spanning different areas of Armed Forces research.

Findings:

As of September 2026, ForcesAtlas contained more than 2.4 million article records, of which approximately 86,000 were publicly discoverable. Pairwise reviewer agreement on inclusion ranged from 92.5% to 95.2% (Cohen's k = 0.849 to 0.904), with 92.9% three-way agreement among records assessed the reviewers (Fleiss' k = 0.906). Reviewers confirmed 1,843 of 2,000 ForcesAtlas classifications (92.2%) in a sample comprising equal numbers of included and rejected records, with all 140 unresolved records counted as failures. Across five systematic reviews, mean known-item coverage was 87.8% (range 70.0% to 100.0%). In total, 135 of 149 included articles were identified in ForcesAtlas.

Conclusion

ForcesAtlas provides a domain-specific evidence layer across multiple scholarly sources, combining multi-source acquisition, article matching across sources and deduplication. Human validation demonstrated high reproducibility in relevance assessment, while known-item benchmarking indicated high coverage across a range of Armed Forces research domains. ForcesAtlas is intended to complement rather than replace existing databases and specialist evidence resources.

Background

Evidence-informed policy, clinical practice, service provision, research and evaluations, depend on researchers (including students) and decision makers identifying relevant research quickly and efficiently. The increasing volume and breadth of academic output make this difficult (1, 2). Publication databases differ in the journals, disciplines and document types that they index, while their retrieval mechanisms differ in functionality, field searching, export capability and reproducibility (3). Methodological studies have shown that relevant records may be uniquely available through individual databases and that search systems vary substantially in their suitability for systematic evidence retrieval (4, 5).

This challenge is important for research concerning the Armed Forces community. The population is heterogeneous, encompassing serving Regular and Reserve personnel, veterans, military families and other military-connected populations (6). Research questions extend beyond a conventional definition of military medicine and can include topics such as mental and physical health, occupational exposure, rehabilitation, transition from service, employment, housing, family functioning, substance use, education, criminal justice, digital inclusion and access to services (7–9). This may be further defined in relation to demographic variables such as age, sex, ethnic minorities and families (10, 11).

Terminology introduces a second discovery problem. Military-connected populations can be described using a wide range of terms, including military personnel, Service personnel, Armed Forces, active duty, serving personnel, reservists, veterans, ex-Service personnel, former Service members, military spouses, military children and military families, alongside country-specific occupational and organisational terminology (12–14). Relevant articles may therefore use different or non-standardised descriptors for substantively similar military-connected populations. Conversely, the presence of military-related terminology does not necessarily indicate that an article is related to the Armed Forces community, creating the potential for false-positive identification when simple keyword-based approaches are used (15).

No general index is designed to resolve this domain-specific problem. For example, PubMed provides broad biomedical coverage, Europe PMC provides life-science literature infrastructure (16), OpenAlex provides a broad open scholarly knowledge graph (17), and Semantic Scholar provides a large-scale literature graph and discovery environment (18). Their breadth is a major strength, but none is intended to act as a curated Armed Forces database spanning health and non-health domains.

Specialist resources provide important complementary functions. In the US, PTSDpubs provides international, cross-disciplinary literature concerning post-traumatic stress disorder (PTSD) and traumatic stress and applies specialist trauma-focused indexing (19). Military REACH identifies, evaluates and translates research concerning military and veteran families (20). In the UK, the Centre for Evidence for the Armed Forces Community (formerly the Forces in Mind Trust Research Centre) provides a repository of research and policy material and research-community resources (21). These initiatives

demonstrate the value of domain-specific evidence infrastructure, but they have different population, document, curation and translation objectives.

Here, we present ForcesAtlas (www.forcesatlas.com) which was developed in response to the specific methodological problem of discovering Armed Forces-relevant evidence across a large, multidisciplinary and internationally distributed scholarly literature (22). The objective was to create a freely accessible evidence-discovery platform in which likely military relevance is established upstream, allowing users subsequently to search the relevant databases by research topic without reconstructing extensive military-population filters for every question. This article describes the development of ForcesAtlas, including its multi-source evidence collection, indexing, provenance, AI-assisted relevance assessment and metadata processing, human quality assurance, search functionality and software-development approach.

Construction and content

Design principles and scope

ForcesAtlas was designed around six principles: 1) broad disciplinary coverage, 2) domain relevance, 3) source traceability, 4) low barriers to use, 5) continuous updating; and 6) human oversight of processing. It was developed as a Research Viable Product (RVP) before being developed into a Minimum Viable Product (MVP) using agile development methodology (23).

The platform is freely accessible without a user account and is intended for researchers, clinicians, policy teams, charities, service providers, students and the wider Armed Forces community. The primary record represents a scholarly article, with available source metadata including title, authors, year, journal, abstract and persistent identifiers, alongside platform-derived enrichment. ForcesAtlas focuses on peer-reviewed research concerning the health, wellbeing and occupational functioning of serving Armed Forces personnel, reservists, veterans and their immediate family members, alongside operational military medicine and combat-casualty care. Its scope spans multiple disciplines within these areas. Relevance is defined at the level of the Armed Forces community rather than a single health condition or discipline. Because the use of military terminology in the paper may be incidental, inclusion is treated as a classification problem rather than simple keyword matching.

Source selection and evidence acquisition

ForcesAtlas uses multiple sources as no single source provides sufficient breadth for the intended domain. A wide search strategy is used to be as inclusive as possible to article inclusion in the extraction from each source. The sources contributing to evidence, cross-checking or metadata enrichment are summarised in Table 1.

Table 1
Data sources used by ForcesAtlas.

Resource	Principal role	Contribution
PubMed	Evidence identification and metadata	Biomedical and health literature, PMID and structured bibliographic metadata
Europe PMC	Evidence identification and enrichment	Life-science literature, abstracts, identifiers and open literature metadata
OpenAlex	Evidence identification and enrichment	Broad multidisciplinary scholarly coverage and linked scholarly entities
Semantic Scholar	Evidence identification and cross-checking	Multidisciplinary scholarly metadata and literature-graph information
Crossref	Identifier checking and metadata enrichment	DOI verification and bibliographic metadata
Unpaywall (part of OpenAlex)	Access enrichment	Identification of legitimate open-access article locations
Manual identification	Quality assurance and gap recovery	Investigation or addition of known relevant publications not satisfactorily resolved through automated sources

Overlapping sources permit records and identifiers to be cross-checked, while source-specific records extend coverage beyond any single index. Candidate articles enter a processing workflow before inclusion. Source and processing metadata are retained through persistent identifiers, batch information, source filenames, enrichment timestamps and selected source labels, although complete source membership is not represented for every record.

Technical architecture and implementation

ForcesAtlas is implemented as a Node.js application. PostgreSQL is the primary datastore and search layer. Records are stored in a central articles table containing bibliographic, current workflow stage, relevance and metadata; authors, tags and countries are represented in separate relational tables linked through many-to-many relationships (see Fig. 1).

Searching is executed in PostgreSQL using a weighted English text-search vector derived from title, abstract and generated summary and indexed with GIN; *pg_trgm* indexes support partial and fuzzy matching, while application logic handles Boolean and fielded queries. PostgreSQL-backed queues support LLM and batch processing through a dedicated worker. Complex grouped searches are persisted as database jobs so that processing can advance in bounded batches and resume after interruption.

Index checking, matching and deduplication

We perform weekly checks on each data sources. PubMed checks search existing records by exact PMID then DOI, whereas Europe PMC, OpenAlex and Semantic Scholar use DOI then PMID. If no identifier match is found, all four use PostgreSQL trigram title similarity greater than 0.85 together with an exact publication year. For all records, we use Crossref, for DOI enrichment which fills any missing data fields prior to analysis. The database restricts exact DOI and PMID; meaning that no record can have the same identifiers.

AI-assisted article relevance classification

Candidate records identified through scheduled source checks, or through batch imports, enter an approval-stage workflow and are placed on a PostgreSQL-backed LLM queue. Relevance screening uses title, publication year when available and abstract. Normal batch queueing requires a non-empty abstract; records without abstracts are skipped by the routine path, although rejected articles can subsequently be rescued through open-access PDF retrieval and if present, then be re processed.

ForcesAtlas assesses article relevance using GPT-4o-mini, providing the article's title, available abstract and publication year. Full-text analyses additionally use available PDF text with references removed. The model returns a relevance category (high, medium, low or none), a confidence score from 0 to 100 and an explanation. High or medium ratings proceed to metadata enrichment. High-rated articles passing additional checks are eligible for human-initiated batch inclusion, while other articles remain available for human review.

The model returns a relevance rating of high, medium, low or none, alongside a confidence score (0 to 100) and reasoning. High relevance indicates that at least one inclusion route is explicit and central to the article's aims, methods or conclusions. Medium relevance indicates a direct and substantive connection within a mixed civilian and military scope, or a less specific operational application. Low relevance indicates an explicit connection that goes beyond a passing mention but is secondary or limited. None indicates that neither inclusion route is satisfied. Confidence separately represents the model's certainty in assigning the rating.

Classification used a zero-shot prompting approach, without worked classification examples. The complete prompt templates, generation settings and rules governing subsequent processing are provided in Supplementary A.

Quality assurance and human review

Human review operates through the review pipeline. Once LLM processing is complete, approval-stage and eligible under-review records are available to reviewers, who can approve or reject them; administrators can also correct bibliographic metadata, relevance fields, tags, countries and review notes. Partial duplicate matches, relevance-policy recoveries, records rescued through full-text processing, public article flags and selected published-record rescans can be routed for review.

AI-assisted software development

Generative AI is used both within evidence processing and to build and modify the platform itself; these uses were treated separately because they have different failure modes and assurances. The platform was developed iteratively using large language model-based coding assistants (Codex, OpenAI). The emerging term “vibe coding” describes a development approach in which a developer or domain specialist expresses desired behaviour in natural language and an AI system generates or modifies software in response (24).

In biomedical research, this approach has been proposed as a means of shortening the distance between scientific intent and working software, while increasing the need for validation, traceability and domain oversight. In ForcesAtlas, AI-assisted programming translated research requirements into candidate implementations for search behaviour, processing workflows, metadata reconciliation, review interfaces and provenance-related functions.

Generated code was treated as provisional and checked by a computer scientist against expected outputs and real bibliographic examples. Unexpected behaviour prompted inspection, modification, rollback or manual correction. Given that seemingly functioning code can still implement scientifically incorrect search, matching or processing behaviour, acceptance was based on whether the implementation reproduced the intended scientific behaviour rather than whether it merely compiled or appeared plausible. ForcesAtlas therefore used vibe coding as AI-assisted iterative software rather than autonomously.

Search architecture and user interface

ForcesAtlas provides a PubMed-style evidence-discovery interface for simple and structured retrieval. Public search executes in PostgreSQL using the weighted full-text index and trigram matching described above, with application logic supporting Boolean operators, field constraints and relational searches over authors, tags and countries. Complex grouped searches persisted in database job tables and processed in bounded batches, allowing results to remain resumable for 24 hours. Search results can be exported into common bibliographic formats for reference-management and evidence-synthesis workflows.

Data and evaluation

Current database

At a September 2026 snapshot, ForcesAtlas contained more than 2.4 million article records, of which approximately 86,000 were publicly discoverable. Rejected, approval-stage and under-review records were retained in the database but excluded from public search to allow for deduplication detection and reduce future AI processing.

The ten most frequently assigned tags encompassed population descriptors, health topics and research methods. These were Veterans (48,796 records; 56.4%), Military Personnel (28,632; 33.1%), Active Duty (13,161; 15.2%), Service Members (11,529; 13.3%), Post Traumatic Stress Disorder (11,388; 13.2%), Mental Health (9,743; 11.3%), Quantitative (7,854; 9.1%), Veterans Health (5,449; 6.3%), Risk Factors (4,959; 5.7%) and Cohort Study (4,903; 5.7%). This distribution represents a concentration of veteran-related research, with post-traumatic stress disorder and mental health prominent among the leading health-related tags. ForcesAtlas tag counts.

The ten most frequently recorded countries of research focus were the United States (65,822 records; 76.0%), United Kingdom (2,791; 3.2%), Afghanistan (2,294; 2.6%), Iraq (2,164; 2.5%), Ukraine (1,445; 1.7%), Israel (1,181; 1.4%), Canada (1,165; 1.3%), Australia (1,036; 1.2%), Vietnam (845; 1.0%) and Iran (505; 0.6%). Country labels represent the geographical focus of the research rather than the location of its authors.

Validation of classification performance

To evaluate ForcesAtlas relevance classification, a validation sample of 2,000 articles was selected, stratified by existing ForcesAtlas status: 1,000 publicly included articles and 1,000 rejected articles. Reviewers were blinded to the ForcesAtlas decision and independently checked the title and abstract. Each publication was classified as Include, Exclude or Insufficient information.

Three reviewers participated (DL, RV, TK). Every record was reviewed by at least two reviewers. A common set of 100 articles was reviewed by all three reviewers; of the remaining records, 634 were allocated to reviewers A and B, 633 to reviewers A and C, and 633 to reviewers B and C. Reviewer A screened 1,367 records, Reviewer B screened 1,367 and Reviewer C screened 1,366 (see Table 2).

Table 2
Reviewer screening outcomes.

Reviewer	Records screened	Include	Exclude	Insufficient information
A	1,367	650 (47.5%)	709 (51.9%)	8 (0.6%)
B	1,367	650 (47.5%)	715 (52.3%)	2 (0.1%)
C	1,366	653 (47.8%)	708 (51.8%)	5 (0.4%)

Pairwise agreement among records with determinate Include/Exclude decisions ranged from 92.5% to 95.2%, with Cohen's *k* values from 0.849 to 0.904. Among the 100 articles reviewed by all three reviewers, 99 received had agreeable decisions from all three reviewers; 92 of these 99 records (92.9%) had complete three-way agreement, with Fleiss' $k = 0.906$ (see Table 3).

Table 3
Inter-rater agreements.

Comparison	Records analysed	Raw agreement	*k*
Reviewer A vs B	730	92.5%	0.849
Reviewer A vs C	725	94.1%	0.881
Reviewer B vs C	730	95.2%	0.904
All three reviewers	**99**	**92.9%**	**0.906**

In total, 885 articles were classified as Include and 975 articles as Exclude. The remaining 140 articles (7.0%) were unresolved because of disagreement between determinate decisions (*n* = 126) or at least one Insufficient information response (n = 14). All 140 were retained and counted as failures in the analysis. They comprised 104 records originally included and 36 originally rejected by ForcesAtlas. Reviewers confirmed 1,843 of 2,000 ForcesAtlas inclusion or rejection decisions (92.2%; 95% CI 90.9% to 93.2%), with all 140 unresolved records counted as failures. This estimate describes performance in the balanced validation sample, comprising equal numbers of included and rejected records.

Known-item coverage across systematic reviews

To assess known-item coverage, five recent systematic reviews were selected to represent different areas of Armed Forces research (8, 25–28). For each review, the publications reported as included were extracted and checked assessed if they were present and visible in ForcesAtlas. The five reviews covered veterans with criminal-legal involvement, suicide among UK Armed Forces veterans, military-to-civilian transition, mental health interventions for military partners, and disease and non-battle injury among deployed personnel.

Across the five reviews, 149 articles were included, of which 135 were present in ForcesAtlas (see Table 4). Coverage ranged from 70.0% to 100.0%. All 29 articles included in Lee *et al*. (2025) (25) were identified in ForcesAtlas. Coverage was 94.7% for Edwards *et al*. (2025) (28), 89.7% for Alcover *et al*. (2024) (27), 84.6% for Smith *et al*. (2025) (8), and 70.0% for Randles *et al*. (2025) (26). The unweighted mean known-item coverage across the five reviews was 87.8%. Of the 14 articles not in ForcesAtlas, eight were judged during verification to be records that would otherwise be expected to be included, two were non-peer-reviewed records, and four had other citation, metadata or identification issues.

Table 4
Verification of known-item coverage across five systematic reviews.

Review	Included Articles	Present in ForcesAtlas	Expected to be present but missing	Coverage (%)
Alcover *et al*. (2024) (27)	78	70	5	89.7
Randles et al (2025) (26)	10	7	0	70.0
Edwards *et al*. (2025) (28)	19	18	1	94.7
Smith *et al*. (2025) (8)	13	11	2	84.6
Lee *et al*. (2025) (25)	29	29	0	100.0
Total / Mean (%)	**149**	**135**	**8**	**87.8**

Utility and discussion

Evidence discovery and synthesis

A central design choice in ForcesAtlas is to move identification of military-connected populations upstream from each individual search. Conventional searches often require researchers to construct extensive military-population terminology before adding the substantive topic of interest. As likely domain relevance is established during database construction, users can instead search within a pre-classified Armed Forces evidence environment. This does not make a search automatically systematic; rather, it changes the point in the process at which domain relevance is established. For example, a user interested in gambling can begin with gambling concepts rather than first building a comprehensive military-population filter.

The approach supports exploratory searches, scoping, policy questions, horizon scanning and early evidence-synthesis work, including identification of terminology, key authors and candidate search concepts. Search results can be exported into reference-management and review workflows, while the curated corpus can support candidate evidence identification, question development and broad evidence mapping.

A multi-index domain-specific database may also have value as a supplementary information source in systematic and scoping reviews. Methodological studies show that reliance on a single bibliographic database can miss relevant evidence (29). Upstream sources can omit literature, automated classification can produce false negatives, indexing can be delayed and some records may have insufficient metadata for confident classification. ForcesAtlas should be treated as an additional

evidence-discovery layer, not a substitute for methodological judgement, critical appraisal or specialist information-retrieval expertise.

Aggregated metadata can be used to examine publication activity across years, topics, countries, journals and authors, supporting evidence mapping and identification of apparent areas of concentration or under-representation. Apparent evidence gaps require validation because low counts may also reflect terminology, source coverage, classification performance or missing metadata.

Comparison with existing resources

ForcesAtlas complements rather than replaces the current infrastructure provided by PubMed, Europe PMC, OpenAlex and Semantic Scholar by adding domain-specific classification, source-aware multi-source acquisition and evidence discovery. The relationship with specialist resources is similarly complementary rather than substitutive. PTSDpubs provides deeper trauma-specific indexing, including a specialised thesaurus and information on tests and measures (19). Military REACH has a particular focus on military and veteran families, evaluating research and translating findings into practical resources (20). The Centre for Evidence for the Armed Forces Community provides a wider UK focused repository, policy and evidence summaries and research-community infrastructure (21). ForcesAtlas focuses on automated multi-source scholarly discovery, route-specific cross-source matching, processing traceability and high-volume domain curation.

Index coverage and validation

Cross-source matching provides operational information about overlap and source-specific discovery during ingestion, but the current schema does not retain a canonical article-to-source membership history for every record. Consequently, source contribution and index coverage cannot yet be inferred reliably from stored article provenance alone. Planned validation will therefore combine known-item reference sets with direct checks against the underlying source indexes. Strengthening article-level source membership and versioned provenance is also a future development priority.

Methodological implications of AI-assisted development

AI-assisted programming can lower the technical barriers to developing research infrastructure, but it can also make it easier to produce software whose limitations are difficult to recognise (24). These approaches can be understood as a spectrum, from accepting generated code with little inspection to incorporating coding assistants within development processes that retain human responsibility for requirements, code review and verification. The extent of AI involvement alone therefore provides limited information about the reliability of the resulting software.

In ForcesAtlas, generated code was treated as provisional and checked by a computer scientist against intended behaviour and real bibliographic examples. This was relevant to search, record matching and relevance-processing workflows, where an implementation could run successfully while producing scientifically inappropriate results. Unexpected behaviour prompted investigation and revision, rollback or manual correction. These did occur frequently.

Limitations

ForcesAtlas cannot guarantee comprehensive coverage because it depends on upstream data sources and heterogeneous metadata. However, having multiple databases does mitigate this risk, and enables wider coverage. Missing abstracts can prevent routine relevance screening. This limitation may depend on acquisition order because scheduled-source matches are currently skipped rather than used to reconcile existing metadata. An article first acquired without an abstract may therefore remain unprocessed even if an abstract is available from another source. To resolve this, is to check alternative sources, including PubMed and Europe PMC, for missing abstracts before treating a record as unavailable for routine screening. Cross-source matching is route-specific, identifier normalisation is not uniform across all ingestion paths, and the current schema does not retain complete article-to-source membership.

Human oversight does not mean that every record has been individually validated, because ForcesAtlas also uses automated and bulk decision pathways. These pathways support acquisition and curation across millions of candidate records, a scale at which exhaustive manual screening would require substantial reviewer resources. Automation can therefore support broader coverage and more frequent updating, while human review provides a means of investigating uncertain decisions and correcting identified errors. Comparative staff time and costs remain to be evaluated. Inclusion in ForcesAtlas indicates domain relevance rather than methodological quality or endorsement of an article's conclusions.

Future development

Priorities are formal validation of database coverage and relevance performance, harmonisation of identifier matching and entity resolution, explicit article-to-source provenance and versioning, expansion to grey literature, and additional support for human-led evidence synthesis. Quality-assurance development should also strengthen reviewer attribution and auditable correction pathways. Future evaluation should examine search efficiency, reproducibility and accessibility for users without specialist information-retrieval training.

Conclusions

ForcesAtlas provides a specialist evidence-discovery layer across public scholarly sources for research concerning serving Armed Forces personnel, veterans and their families. Its architecture combines multi-source acquisition, ingestion-specific matching and deduplication, PostgreSQL-backed search, AI-assisted relevance assessment and operational human review.

AI-assisted software development, including vibe coding, reduced the distance between domain requirements and implementation but did not remove the need for verification. ForcesAtlas complements bibliographic databases, specialist resources and formal evidence-synthesis methods;

empirical evaluation of coverage, retrieval and classification performance, together with stronger source provenance, is the next step in establishing its methodological contribution.

Declarations

Ethics approval and consent to participate

Not required.

Consent for publication

Not applicable.

Competing interests

DL is a member of the UK Army Reserve Force. This work has been undertaken as part of his civilian employment. AS is a full-time member of the UK Armed Forces.

Funding

No funding.

Availability of data and materials

ForcesAtlas is freely accessible at www.forcesatlas.com and can be searched without registration. Bibliographic records are derived from public scholarly metadata services and remain subject to the terms and licences of their originating sources.

References

1. Bornmann L, Haunschild R, Mutz R (2021) Growth rates of modern science: a latent piecewise growth curve approach to model publication numbers from established and new literature databases. Humanit Soc Sci Commun 8(1):224. 10.1057/s41599-021-00903-w
2. Marshall IJ, L'Esperance V, Marshall R, Thomas J, Noel-Storr A, Soboczenski F et al (2021) State of the evidence: a survey of global disparities in clinical trials. BMJ Glob Health 6(1):e004145. 10.1136/bmjgh-2020-004145
3. Bramer WM, Rethlefsen ML, Kleijnen J, Franco OH (2017) Optimal database combinations for literature searches in systematic reviews: a prospective exploratory study. Syst Rev 6(1):245. 10.1186/s13643-017-0644-y

4. Gusenbauer M, Haddaway NR (2020) Which academic search systems are suitable for systematic reviews or meta-analyses? Evaluating retrieval qualities of Google Scholar, PubMed, and 26 other resources. Res Synth Methods 11(2):181–217. 10.1002/jrsm.1378
5. Marshall IJ, Marshall R, Wallace BC, Brassey J, Thomas J (2019) Rapid reviews may produce different results to systematic reviews: a meta-epidemiological study. J Clin Epidemiol 109:30–41. 10.1016/j.jclinepi.2018.12.015
6. Peat D, Perrmann-Graham J, Stone CB (2026) Veterans and military-connected individuals in the civilian workforce: an integrative review and research agenda. Organ Manag J 23(1):69–96. 10.1108/OMJ-10-2024-2323
7. Walker J, Misca G (2025) Reflections on a journey from research to impact: Influencing policy meeting the needs of UK Armed Forces families. Fam Court Rev 63(4):728–741. 10.1111/fcre.70019
8. Smith A, Rafferty L, Croak B, Greenberg N, Khan R, Langston V et al (2025) A systematic review of military-to-civilian transition, The role of gender. Singer DE, editor. PLOS ONE. ;20(2):e0316448. 10.1371/journal.pone.0316448
9. Coventry T, Coventry M, Russell K (2025) Veteran Pathways to Homelessness and Housing Solutions: A Scoping Review. J Veterans Stud 11(2):170–194. 10.21061/jvs.v11i2.810
10. Salem K, Randles R, Sapre B, Finnegan A (2023) Experiences of ethnic minority personnel in the armed forces: A systematic review. J Mil Veteran Fam Health 9(1):5–14. 10.3138/jmvfh-2022-0019
11. Gribble R, Goodwin L, Fear NT (2019) Mental health outcomes and alcohol consumption among UK military spouses/partners: a comparison with women in the general population. Eur J Psychotraumatology 10(1):1654781. 10.1080/20008198.2019.1654781
12. Gribble R, Mahar AL, Keeling M, Sullivan K, McKeown S, Burchill S et al (2020) Are we family? A scoping review of how military families are defined in mental health and substance use research. J Mil Veteran Fam Health 6(2):85–119. 10.3138/jmvfh-2019-0054
13. Dodge J, Wortham W, Kale C, Williamson V, Ross A, Maher S et al (2025) Programs to Address Violence for Military Families: a Systematic Review. J Fam Violence 40(2):383–399. 10.1007/s10896-023-00586-8
14. Townsend ML, Green H, Fabrianesi B, Braunack-Mayer A (2026) Military personnel perspectives on participating in health research: A scoping review. Cioffi A. editor PLOS One 21(4):e0346884. 10.1371/journal.pone.0346884
15. Leightley D, Pernet D, Velupillai S, Stewart RJ, Mark KM, Opie E et al (2020) The Development of the Military Service Identification Tool: Identifying Military Veterans in a Clinical Research Database Using Natural Language Processing and Machine Learning. JMIR Med Inf 8(5):e15852. 10.2196/15852
16. The Europe PMC Consortium (2015) Europe PMC: a full-text literature database for the life sciences and platform for innovation. Nucleic Acids Res 43(D1):D1042–D1048. 10.1093/nar/gku1061

17. Priem J, Piwowar H, Orr R, OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts [Internet]. arXiv; 2022 [cited 2026 Sep 3]. Available from: http://arxiv.org/abs/2205.0183310.48550/arXiv.2205.01833
18. Ammar W, Groeneveld D, Bhagavatula C, Beltagy I, Crawford M, Downey D et al Construction of the Literature Graph in Semantic Scholar. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers) [Internet]. New Orleans - Louisiana: Association for Computational Linguistics; 2018 [cited 2026 Sep 3]. pp. 84–91. Available from: http://aclweb.org/anthology/N18-301110.18653/v1/N18-3011
19. U.S. Department of Veterans Affairs, National Center for PTSD PTSDpubs [Internet]. 2026 [cited 2026 Sep 3]. Available from: https://www.ptsd.va.gov/ptsdpubs/index.asp
20. Military REACH, Military REACH [Internet]. Auburn University; 2026 [cited 2026 Sep 3]. Available from: https://militaryreach.auburn.edu/
21. Centre for Evidence for the Armed Forces Community About the Centre for Evidence for the Armed Forces Community [Internet]. 2026 [cited 2026 Sep 3]. Available from: https://www.centreforevidence.org/about
22. ForcesAtlas, ForcesAtlas Armed Forces & Veterans Health Research Database [Internet]. 2026 [cited 2026 Sep 3]. Available from: https://forcesatlas.com/
23. Leightley D Navigating the software development landscape: Introducing the term Research Viable Product (RVP) [Internet]. London: King’s College London; 2024 Jun [cited 2024 Oct 10]. Available from: http://doi.org/10.17605/OSF.IO/5N37J
24. Moore JH, Tatonetti N (2025) Vibe coding: a new paradigm for biomedical software development. BioData Min. ;18(1):46, s13040-025-00462–9. 10.1186/s13040-025-00462-9
25. Lee M, Hwang S, Lee J (2025) Non-pharmacological interventions for mental health among partners of military members and Veterans: A systematic review. J Mil Veteran Fam Health 11(3):139–152. 10.3138/jmvfh-2023-0111
26. Randles R, Burroughs H, Green N, Finnegan A (2025) Prevalence and risk factors of suicide and suicidal ideation in veterans who served in the British Armed Forces: a systematic review. BMJ Mil Health 171(2):166–172. 10.1136/military-2023-002413
27. Alcover KC, Howard K, Poltavskiy E, Derminassian AD, Nickel MS, Allard RJ et al (2024) Disease and Non-Battle Injury in Deployed Military: A Systematic Review and Meta-analysis. Mil Med 189(Supplement3):21–30. 10.1093/milmed/usae033
28. Edwards ER, Fortuna A, Holliday R, Addison H, Tsai J (2025) Prevalence of mental health conditions, substance use disorders, suicidal ideation and attempts, and experiences of homelessness among Veterans with criminal-legal involvement: A meta-analysis. Clin Psychol Rev 115:102533. 10.1016/j.cpr.2024.102533
29. Goossen K, Hess S, Lunny C, Pieper D (2020) Database combinations to retrieve systematic reviews in overviews of reviews: a methodological study. BMC Med Res Methodol. ;20(1):138.

10.1186/s12874-020-00983-3

Figures

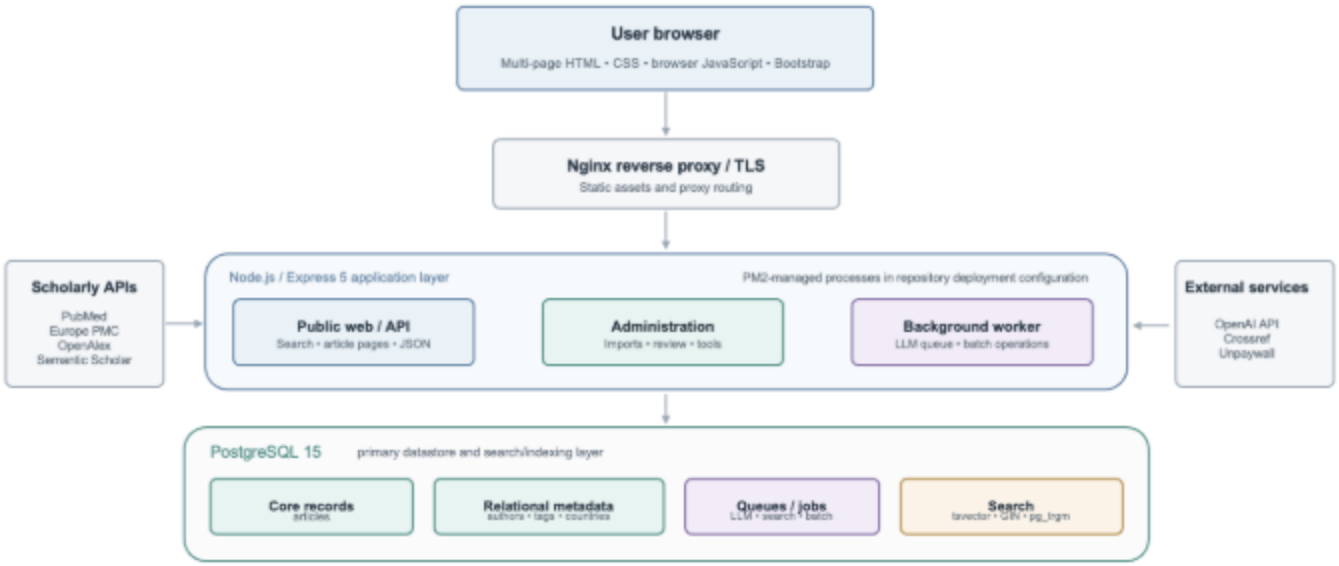


Figure 1

Forces Atlas overall architecture.

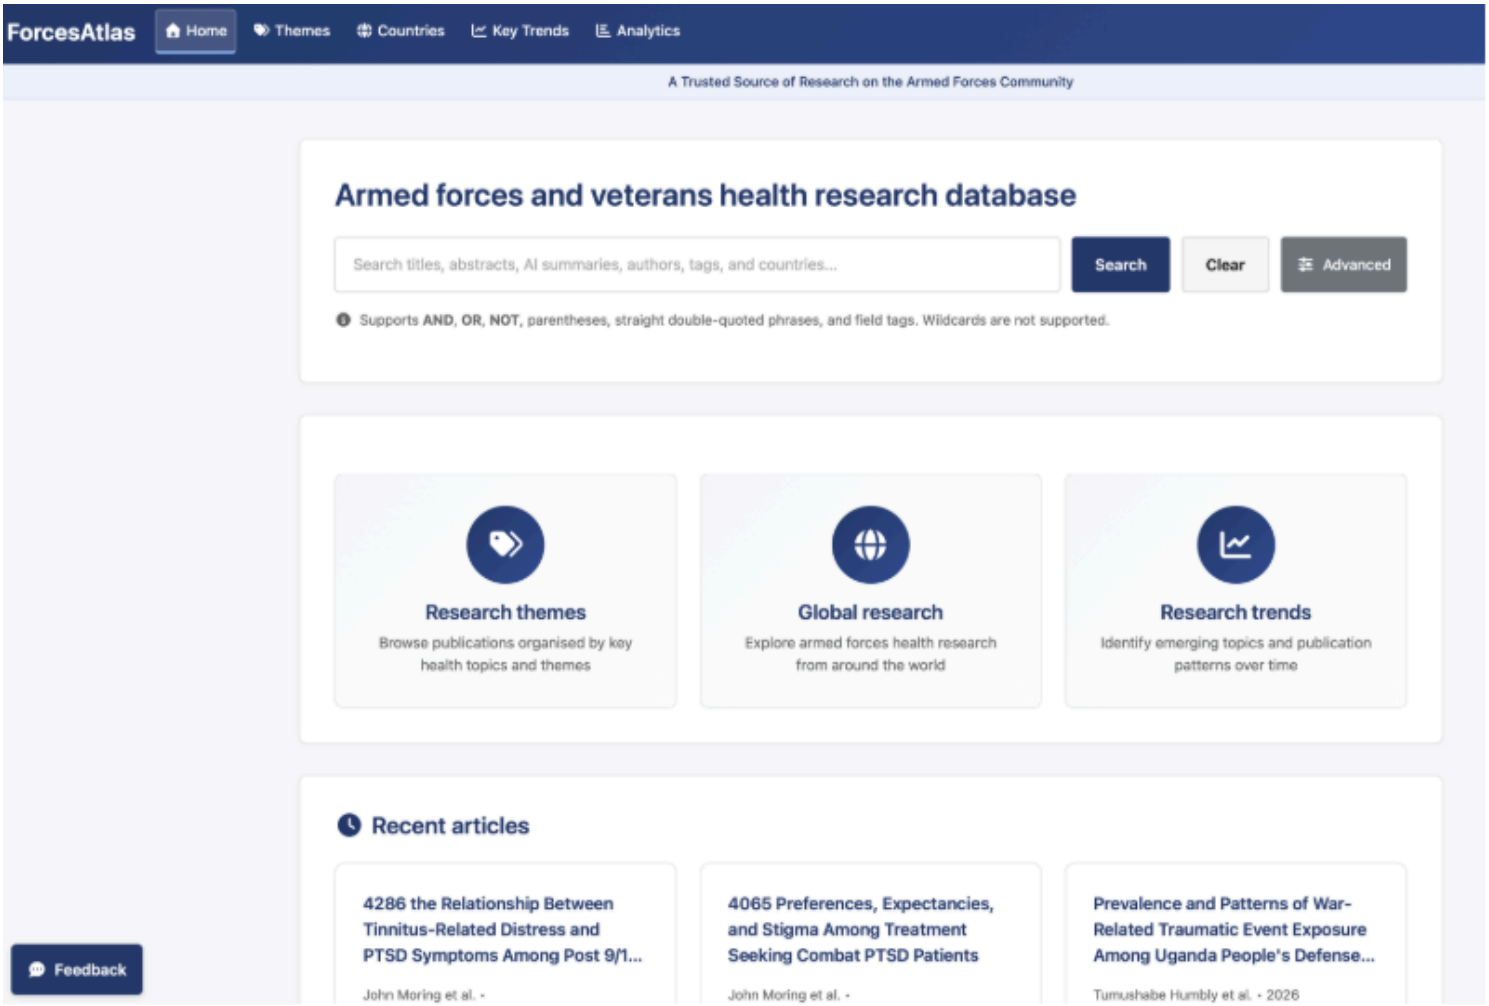


Figure 2

Example of the ForcesAtlas landing page.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementA.docx