Combatting Veteran PTSD with Deep Learning on Longitudinal UK Health Data: A Comprehensive Review

Gbenga Adeniyi Adediran [1], **Andrew Ayemere Okhueigbe** [2], **Ruth Ese Otaigboria** [3], **Chiamaka Pamela Agu** [4], **Chizube Obinna Chikezie** [5]

[1] *Leeds Beckett University*
City Campus, Leeds, LS1 3HE, United Kingdom

[2] *University of Denver*
2199 S. University Blvd., Denver, CO 80208, USA

[3] *Northern Illinois University*
1425 W. Lincoln Hwy., DeKalb, IL 60115-2828, USA

[4] *University of New Haven*
300 Boston Post Rd, West Haven, CT 06516, USA

[5] *Ulster University*
Saint James Buildings, 61-95 Oxford St, Manchester, M1 6FQ, UK

DOI: 10.22178/pos.122-8

LCC Subject Category: BF1-990

Received 27.08.2025
Accepted 27.09.2025
Published online 30.09.2025

Corresponding Author:
Gbenga Adeniyi Adediran

Abstract. PTSD in UK veterans often emerges long after service, complicating detection. The objective of this review is to assess whether deep learning applied to longitudinal NHS and Ministry of Defence data can facilitate a shift from late recognition to anticipatory intervention, and to define the requirements for its safe deployment. Evidence on sequence models (LSTM, transformers), multimodal integration of structured records and clinical notes, and explanation layers that render outputs legible to clinicians is synthesised. UK constraints, fragmented records, inconsistent veteran identifiers, and uneven digital maturity limit scale; feasible mitigations include secure data environments and federated training. Priority actions include standardising veteran coding, establishing a Veteran Health Analytics Hub, conducting prospective trials with health-economic endpoints, and integrating risk scores with concise, validated explanations into existing workflows. Properly implemented, longitudinal deep learning can reduce missed cases, accelerate access to effective support, and enable services to learn from their data while protecting privacy and trust.

Keywords: veteran PTSD; deep learning; longitudinal EHR; NHS data; model interpretability.

INTRODUCTION

PTSD is a complex psychiatric disorder arising after traumatic experiences, characterised by symptoms such as intrusive memories, avoidance, negative mood shifts, and hyperarousal. Military veterans are a high-risk group for PTSD due to combat and service-related traumas [1]. In the UK Armed Forces, recent estimates put the rate of probable PTSD around 9–10%, with higher rates among combat-exposed subgroups [2]. This prevalence is roughly double that of the general UK population (approximately 4–5%) [3], reflecting the unique burdens of military service. PTSD in veterans often co-occurs with depression, anxiety, and substance misuse, compounding its impact on well-being and long-term functioning [4]. Moreover, PTSD can have a delayed onset and protracted course, sometimes emerging months or years after service separation [5]. Such temporal complexity challenges traditional screening and suggests that longitudinal monitoring may be necessary to identify at-risk individuals who initially appear resilient [6, 7].

Early intervention is critical; untreated PTSD is linked to chronic physical comorbidities (e.g., pain, cardiometabolic illness) and elevated suicide risk [8]. However, timely diagnosis is often hampered by stigma, heterogeneous symptom profiles, and reliance on self-report or delayed help-seeking [9]. These challenges have spurred interest in data-driven approaches to augment PTSD detection and treatment. In particular, the rise of electronic health records (EHRs) and "big data" has enabled the leveraging of routine healthcare data to identify emerging mental health issues. The UK's NHS, which provides universal coverage including for veterans, houses rich longitudinal records of diagnoses, prescriptions, and hospital admissions [10, 11]. Integrating these records with military service data (e.g., deployment history or combat injuries) could enable more proactive identification of veterans who may develop PTSD, allowing for preventative care or early therapy.

Machine learning (ML), particularly deep learning, offers powerful methods for detecting latent patterns in large and complex datasets. ML algorithms can "learn" predictive signatures from high-dimensional health data that might elude conventional analytic techniques. Early studies demonstrated the feasibility of ML models for PTSD risk stratification: for example, researchers have used pre-deployment survey data to predict which soldiers will screen positive for PTSD post-deployment [12]. Deep learning, a subset of ML using multi-layered neural networks, is particularly adept at modelling sequential data and uncovering non-linear relationships in multimodal inputs. These capabilities align well with the needs of PTSD prediction, where time sequences of medical encounters and subtle changes in behaviour or physiology might presage symptom onset. In parallel, the NHS and UK research bodies have begun prioritising artificial intelligence (AI) innovations in healthcare [13]. A national NHS AI Lab has been established to accelerate safe deployment of AI, and pilots such as "Foresight" are training models on population-scale NHS data [14]. Such initiatives create a timely opportunity to apply deep learning for veteran mental health.

In this context, our review examines how deep learning on longitudinal UK health data can help combat PTSD in veterans. We outline the current state-of-the-art in ML for PTSD, highlight UK-specific data considerations, and discuss key issues of model interpretability, clinical integration, and system readiness. By drawing on both foundational studies and the latest literature (2015–2025), we aim to provide a comprehensive overview and a strategic roadmap for researchers and practitioners interested in harnessing deep learning to improve outcomes for veterans with PTSD.

METHODS

We conducted a narrative literature review focusing on peer-reviewed studies and authoritative reports relevant to PTSD, veterans, and deep learning applications in health data. No formal PRISMA systematic search or meta-analysis was performed, given our broad exploratory scope. Instead, we iteratively searched academic databases (e.g., PubMed, IEEE Xplore) and reference lists for English-language publications from approximately 2000 to 2025, with a particular emphasis on the period from 2015 to 2025, to capture contemporary research on deep learning. Key search terms included combinations of "PTSD," "veteran," "machine learning," "deep learning," "longitudinal," "electronic health records," "UK," "NHS," "interpretability," and "healthcare AI." We also reviewed UK government and NHS digital strategy documents, as well as seminal older works, to provide background and contrast. Studies spanning a range of methodologies (e.g., retrospective EHR analyses, prospective cohort ML studies, and systematic reviews) were included to ensure a holistic understanding.

Sources were selected for relevance to one or more of the following topics: (a) deep learning techniques for sequential or longitudinal medical data; (b) machine learning applications to PTSD or mental health; (c) UK veteran mental health epidemiology or health services; and (d) ethical, interpretability, or implementation aspects of healthcare AI. Given the narrative (rather than systematic) approach, we did not calculate inter-rater agreement or perform risk-of-bias assessments. Rather, we synthesised findings qualitatively and organised the discussion around thematic headings. The lack of formal meta-analytic pooling is a limitation of our method; however, our goal was to integrate diverse strands of evidence into an analytical perspective, identifying consensus, gaps, and implications for future work. The following sections (Literature Review, Results, and Discussion) reflect this integrative analysis. All source materials are cited as live hyperlinks for transparency.

Literature Review

Early efforts to predict or identify PTSD using computational models date back at least two decades, initially employing simpler classifiers and regression models on questionnaire data or limited clinical features [15]. For example, researchers in the 2000s explored logistic regression models using demographics and trauma exposure scores to estimate PTSD risk, laying the groundwork for later data-driven approaches [16]. By the mid-2010s, interest shifted toward machine learning algorithms capable of handling more complex inputs and interactions. A 2015 study by Karstoft et al., for instance, used an ensemble of decision trees to differentiate soldiers with chronic vs. recovering PTSD trajectories, highlighting that non-linear ML methods could outperform traditional statistical predictions in this domain [17]. Around the same time, authors applied unsupervised learning to longitudinal PTSD symptom scores, uncovering distinct trajectory subtypes (resilient, delayed onset, etc.) [18], which underscored the heterogeneous and temporal nature of PTSD courses. These foundational works suggested that capturing the time dimension and multi-factorial nature of PTSD could improve predictive accuracy, a hypothesis that deep learning would soon test on a larger scale.

The advent of deep neural networks brought significant advances. Deep learning models can automatically learn complex feature representations from raw data, reducing the need for manual feature engineering. In healthcare, a landmark model was the "Deep Patient," which, in 2016, demonstrated that unsupervised deep feature learning on EHR data could improve predictions of various diseases [19]. Authors trained stacked denoising autoencoders on a hospital's entire EHR warehouse, deriving a latent patient representation that consistently outperformed both original raw data and shallow models in forecasting future diagnoses [19]. This result was foundational in showing the promise of deep learning on longitudinal health records. Subsequent studies extended deep learning to mental health. For example, in a *Translational Psychiatry* meta-review, authors noted an emerging trend of applying deep multilayer perceptrons and convolutional neural nets to psychiatric data, albeit often with small sample sizes. By 2019–2020, mental health-focused deep learning research had accelerated: transformer networks, which excel at sequence modelling, were introduced to EHR data (e.g., TransformEHR using encoder–decoder transformers for disease outcome prediction) [20], and recurrent neural networks (RNNs) like LSTMs were being routinely used to model patient timelines in depression and psychosis studies [21].

Specific to PTSD, the literature of the past decade includes both systematic reviews and innovative modelling studies. A 2023 systematic review by [22] identified 41 studies using ML for PTSD diagnosis, concluding that AI-based methods show strong potential but require careful model selection, feature curation, and validation to be clinically useful. They stressed that integrating diverse data types often improved performance. Similarly, a 2024 meta-analysis by [23] found that machine learning could classify PTSD with ~89% overall accuracy, and that models using multi*dimensional data (e.g., combining clinical, psychological, and physiological inputs) achieved up to 96% accuracy*, outperforming those using any single data type. This echoes findings in other domains that multimodal learning can provide a more comprehensive understanding of a condition. In terms of concrete model applications, several recent studies have leveraged large clinical datasets: authors [12] developed and validated an ML model on 4,711 U.S. Army soldiers' data to predict post-deployment PTSD, successfully stratifying high-risk individuals before symptoms emerged. Authors [24] focused on EHR data, using Canadian primary care records with both structured fields and unstructured text; their best model (a hybrid deep neural network) could identify PTSD patients with an area under the curve (AUC) of 0.88, and importantly, the unstructured **clinical notes over multiple visits provided a stronger predictive signal than structured data alone*. The study underlines the value of longitudinal narrative data (progress notes, etc.) to detect the often subtle indicators of PTSD.

The literature reveals a convergence of evidence that deep learning applied to longitudinal health data is a promising avenue for PTSD risk detection and diagnosis. Earlier foundational research established feasibility, and more recent works have demonstrated high accuracy when modern architectures and rich data are used. At the same time, across these studies, several common themes emerge: PTSD should be approached as a temporal phenomenon (not a static diagnosis), combining data sources (medical history, text notes, possibly even wearables or social data)

improves prediction, and achieving real-world impact will require addressing interpretability and integration challenges. Building on these insights, we next examine the current state of the art in this field and analyse key discussion areas, particularly relevant to UK veterans and the NHS context.

RESULTS AND DISCUSSION

Current Capabilities of Deep Learning for PTSD. Thanks to advances in algorithms and the growing availability of health datasets, state-of-the-art models can sift through years of patient data to flag those at risk of PTSD or related disorders with impressive accuracy [25]. Recent machine learning models have moved beyond simple classification toward more nuanced predictions, for instance, forecasting the *future trajectory* of PTSD symptoms or the likelihood of relapse at a given time point [26]. Cutting-edge architectures, such as transformers (originally developed for natural language processing), are now being applied to sequences of healthcare events. These models employ self-attention mechanisms to capture long-range dependencies in a patient's record. Early experiments suggest that transformers outperform traditional RNNs in handling the irregular timing and varied types of medical events found in EHRs [27]. For example, a 2023 study reported that a transformer-based model could predict a range of 12-month outcomes (including new mental health diagnoses) from prior EHR data more accurately than logistic regression or RNN benchmarks [27]. Figure 1 provides a conceptual framework for how such deep learning models ingest longitudinal veteran health data and produce predictions of PTSD risk.

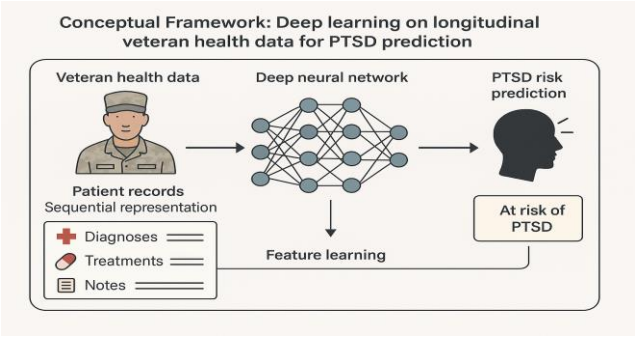

Figure 1 – Conceptual Framework: Deep learning on longitudinal veteran health data for PTSD prediction

Patient records are extracted and processed into sequential representations (e.g. timelines of diagnoses, treatments, and notes). A deep neural network (such as a multilayer transformer or recurrent model) learns abstract features from these sequences, capturing temporal patterns. The learned representation can then be used for predictive tasks, here, identifying individuals at risk of developing PTSD, which in turn can inform early interventions [19].

As Figure 1 suggests, a hallmark of modern approaches is the integration of multimodal data. PTSD risk factors are diverse, encompassing medical history, psychological assessments, service records, and social determinants; therefore, models that only consider one aspect of the data may miss the broader picture. Today's best-performing systems often combine structured data (such as diagnoses or prescription codes) with unstructured text (clinician notes, mental health questionnaires), and sometimes even biomarker or sensor data. The author's work [24] is illustrative: their model's sensitivity improved markedly when free-text notes (documenting patients' moods, sleep problems, trauma narratives, etc.) were included alongside coded data. In general, natural language processing (NLP) techniques, including domain-specific BERT-based transformers, have enabled the extraction of PTSD-related signals from clinical text at scale (for instance, flagging phrases indicative of flashbacks or hypervigilance in psychotherapy notes). This multimodal trend aligns with the meta-analysis above, which finds that multidimensional models achieve ~96% accuracy, compared to ~86–90% for single-source models [23].

In the UK, an exciting development is the national-scale modelling of health outcomes using deep learning techniques. In 2024, researchers from University College London and King's College London piloted an AI model named "Foresight" on de-identified NHS data encompassing 57 million individuals [14]. Leveraging techniques akin to large language models, Foresight was trained on sequences of routine health events (hospital admissions, accident & emergency visits, medications, etc.) to predict subsequent outcomes. Initially focused on forecasting COVID-19, this pilot demonstrated that it is feasible to train a single model using an entire country's longitudinal health records. The model showed accuracy in predicting various medical events (new diagnoses, complications) and illustrated

the potential of large-scale data to capture patterns even for minority conditions or subpopulations [14]. For veteran PTSD, such an approach could be transformative: by representing the full "health trajectory" of millions, including those with military service, a model might learn subtle precursors of PTSD that occur in the years following discharge (e.g., increasing primary care visits for sleep issues or chronic pain).

Another cutting-edge feature is the integration of model interpretability techniques into AI focused on PTSD. Researchers recognise that clinical uptake of deep learning will require more than just accuracy; models must also explain their reasoning. To that end, recent studies often accompany their deep models with interpretability methods such as attention weight visualisation (highlighting which past events most influenced a PTSD prediction) or feature importance rankings derived via SHAP (SHapley Additive exPlanations) values [28, 29]. For example, in a deep learning study on EHR data, the authors computed the contribution of each feature. They found that prior diagnoses of sleep disorders and frequent pain complaints were among the top predictors of later PTSD diagnoses, consistent with clinical intuition and thus helping to validate the model [30, 31]. Such tools not only build trust but can yield novel insights; a model might reveal, say, that a sequence of minor physical health issues in a veteran (headaches, hypertension) strongly preceded their PTSD onset, suggesting a stress mechanism worth further examination. The literature indicates a growing synergy between predictive performance and interpretability, a positive development given that black-box models alone are often viewed skeptically in medicine.

Finally, it's worth noting that the technical barriers to deployment are receiving more attention in state-of-the-art discussions. The real-world implementation of these AI models in the NHS presents challenges, including data security, integration with electronic health record systems, and obtaining regulatory approval. However, progress is evident: secure data environments (SDEs) have been established to enable the training and validation of models on sensitive NHS data without compromising privacy. Regulators, such as the UK Medicines and Healthcare products Regulatory Agency (MHRA), are piloting sandbox initiatives for AI-as-a-medical-device oversight [32]. And the NHS AI Transformation Directorate emphasises principles of explainability and human oversight for any AI tools deployed in patient care [32]. These parallel efforts mean that the technical state-of-the-art is increasingly about not just model accuracy, but also how well models can be integrated into clinical workflows, audited, and scaled across the health system. In summary, the field now has powerful deep learning models capable of processing longitudinal UK health data for the prediction of PTSD. The next steps are to refine these models for the veteran context and ensure they can be responsibly deployed.

Modelling PTSD as a Temporal Disorder: The Case for Longitudinal Learning. Unlike many medical conditions that have a clear onset or diagnosis date, PTSD often follows a dynamic temporal course [33]. In delayed-onset PTSD, symptoms may fluctuate, recur after periods of remission, or appear for the first time long after the triggering trauma. PTSD is "characterised by delayed onset and prolonged duration" in a significant subset of cases [5]. Some military studies estimate that more than one-third of PTSD cases might be classified as delayed-onset, only manifesting clinically months or years post-trauma or service. This temporality means that a single cross-sectional snapshot of a patient's status is often insufficient to gauge their true risk or illness trajectory [34, 35]. For example, two veterans might both screen negative for PTSD one year after discharge, yet one could remain resilient while the other begins showing subtle prodromal signs, leading to full PTSD three years later [34]. Traditional diagnostic approaches, which rely on point-in-time assessments (e.g., a questionnaire during a clinic visit), risk misclassifying the second veteran as a false negative.

*Longitudinal data and temporal modelling address this gap***.** By continuously or periodically monitoring health indicators over time, one can detect patterns suggestive of emerging PTSD. Deep learning models, especially sequence-based architectures, are well-suited to exploit such longitudinal patterns. Recurrent neural networks (RNNs), such as LSTMs and gated recurrent units, can maintain a memory of past events when scanning through a patient's timeline. In contrast, transformer models can attend to events regardless of when they occurred [36]. These models effectively treat a patient's history as a sequence to be learned, rather than a set of independent observations. In doing so, they can

capture temporal dependencies (e.g., the worsening of insomnia and hyperarousal symptoms every winter) and evolving risks (e.g., a veteran's danger might spike following a divorce or retirement from work; transitions are often observable in health records through indirect signals). Studies have shown that incorporating prior time points improves the prediction of PTSD and related outcomes [37]. For instance, a machine learning analysis of trauma survivors found that trajectories of symptom change in the first few months post-trauma were highly predictive of chronic PTSD; patients who showed even small residual symptoms at 1 month tended to either recover or exacerbate along specific paths, which a temporal model could learn to forecast [37].

Another benefit of modelling PTSD as a temporal disorder is the ability to perform early warning surveillance. By the time PTSD is formally diagnosed (often many months after symptom onset, due to delays in help-seeking), the disorder may have already impaired the person's life significantly. However, if a deep learning model monitoring longitudinal data could flag "PTSD likely" earlier, perhaps when only mild symptoms are documented, clinicians could intervene proactively. Encouraging evidence comes from the authors [38], who reviewed longitudinal EHR-based models and found that ML/DL algorithms can detect diseases *earlier than clinical diagnosis* in various conditions. Applied to PTSD, the findings could mean identifying patients for further evaluation before they hit crisis points. For example, imagine a veteran who, over two years, has escalating complaints of chronic pain, interpersonal problems noted in therapy sessions, and a couple of emergency visits for panic attacks; a longitudinal model might recognise this pattern as similar to prior patients who eventually were diagnosed with PTSD, thus prompting an alert for a PTSD screening. Preliminary research indicates that this kind of temporal risk flagging is feasible; one study achieved the successful prediction of PTSD onset 14 months in advance by analysing subtle changes in functional MRI and psychophysiological data over time (though MRI is not routine in practice, analogous patterns in EHRs could exist) [39].

In summary, treating PTSD as a moving target over time, rather than a binary state, aligns with clinical reality and improves predictive modelling. Longitudinal deep learning leverages the rich temporal context available in veterans' health records. It enables the capture of delayed effects (e.g., the latent impact of combat exposure that surfaces years later) and cumulative burdens (e.g., multiple minor stressors that compound until a threshold is crossed). For UK veterans, whose PTSD may surface after they've left military service and lost the daily support structure of unit cohesion, such modelling is especially pertinent [40]. The case for longitudinal learning is thus strong: it mirrors the nature of PTSD and yields. This approach provides earlier warnings and potentially differentiates between transient distress and the onset of a chronic disorder. Future research should continue to refine temporal models, incorporating irregular time intervals and life events, to exploit the longitudinal dimension of veteran data fully.

Bridging Data Gaps in UK Veteran Health Records. A critical challenge in applying deep learning to UK veterans' PTSD is the fragmentation and gaps in their health data. Unlike the United States, which has the Veterans Health Administration providing an integrated medical record for many former service members, the UK does not have a single, unified veterans' health system [41, 42]. During active duty, servicemembers receive healthcare via the Ministry of Defence's Defence Medical Services (DMS) [43]. Upon leaving the forces, their care transitions to the NHS and civilian providers [44]. In principle, a summary of their military medical record is supposed to be transferred to their NHS general practitioner (GP). In practice, this handover is inconsistent, and historically, no routine flag or indicator existed in NHS health records to identify individuals who are veterans [45]. Authors [46] note that England and Wales lacked a universally applied coding system to identify veteran status in secondary mental health care systems, making it difficult to even count the number of veterans accessing those services. Although initiatives such as the Armed Forces Covenant and GP accreditation schemes now encourage GPs to code patients as veterans, under-recording remains widespread; veterans often do not disclose their military history, or providers fail to update records accordingly [47, 48].

Bridging this gap requires both data linkage and data sharing. On the linkage front, efforts have begun to connect disparate sources. For example, researchers funded by the Forces in Mind Trust

conducted a feasibility study linking a mental health trust's patient records with a list of known veterans to see if service history could be inferred. They developed a "Military Service Identification Tool" to search NHS records for clues (like specific injuries or terms) that suggest a patient is a veteran [49]. Such work demonstrated that it is possible to identify many veterans retrospectively in EHR data, but the process is labour-intensive and imperfect. A more systematic solution would be *to embed veteran status into NHS data at the point of registration,* for instance, when someone registers with a GP after leaving the military, using the standard Read code or SNOMED code for "Military veteran" (codes Xa8Da or 13J [48, 50]). This is now NHS policy in England; veterans are advised to tell their GP and have their records coded, but compliance is uneven [51].

Even assuming we can identify veterans in NHS data, the next gap is the content of their military health records. Service-related information, like deployment history, combat exposures, any mental health evaluations during service, or physical injuries, often resides only in MoD archives and is not automatically merged with NHS records. Yet these data are relevant for PTSD risk modelling; for example, a history of traumatic brain injury or participation in intense combat operations is a strong risk factor for subsequent PTSD [52]. Currently, a researcher or AI model trying to predict PTSD from NHS records alone might lack that context. Bridging the gaps requires secure data-sharing agreements between the MoD and NHS databases to enable the linking of individuals' records. Figure 2 illustrates the envisioned data flow: MoD service records (containing dates of service, roles, and medical events from DMS) would be joined with post-service NHS health records, creating a *combined longitudinal dataset* for each veteran, which can then be used to train the model.

Military service data (e.g., deployment history, DMS medical records) and civilian NHS data (GP and hospital records, mental health services) are linked, with appropriate consent and de-identification, to produce a unified longitudinal health profile of veterans. This combined dataset is used to train deep learning models (such as sequential neural networks), which then output predictions (e.g. PTSD risk scores). Secure data environments and governance frameworks ensure privacy during this process.

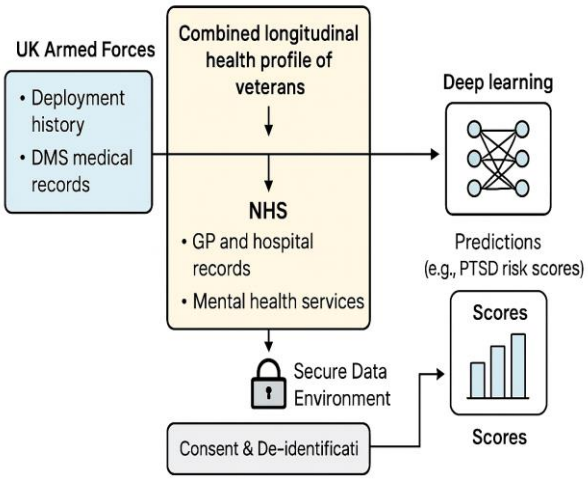

Figure 2 – Data flow linking UK Armed Forces service records with NHS health records for deep learning

Table 1 – Themes, example deep-learning methods, and UK-specific implications for linking MoD and NHS records to model veteran PTSD risk

Theme	Example Deep-Learning Methods	UK-Specific Implication
Temporal Modelling	LSTM / GRU sequences; transformer EHR models	Leverages longitudinal NHS records to flag delayed-onset or relapsing PTSD [22]
Interpretability	Attention heat-maps; SHAP/LIME attribution	Meets MHRA & NICE expectations for clinician-explainable AI in mental-health care [53]
Data Integration	Federated learning; synthetic-data augmentation	Links MoD service data with NHS EHRs while preserving GDPR privacy [54]
Infrastructure Readiness	Model compression; edge-inference deployment	Eases rollout across NHS trusts with uneven GPU/IT capacity [55]

Achieving this linkage in practice faces technical and ethical hurdles. Data security is paramount; veteran data is sensitive and must be handled in strict compliance with GDPR and other regulations. The NHS England Secure Data Environment (SDE) approach, as used in the Foresight project, offers one model: data from different sources can be linked within a secure enclave that researchers query without extracting identifiable records [56]. For example, the MoD and NHS Digital could collaborate to place relevant data into an SDE for veteran health research, where an AI model could be trained within the environment, without exposing individual identities. This federated or enclave-based training preserves privacy while allowing large-scale analysis [57]. Another approach is federated learning across institutions: rather than pooling data centrally, algorithms are sent to train separately on MoD data and NHS data and then aggregate the learned parameters [58]. This way, raw data remains in its original systems. Federated deep learning is still an emerging technique in healthcare, but it may be suitable for scenarios where direct MoD–NHS data sharing is restricted.

Beyond identity linkage, we must also address data quality and completeness. Many UK veteran health records suffer from missing data, e.g., an Afghanistan deployment might not be noted in NHS files, or mental health symptoms discussed with a charity counsellor might never enter NHS systems. Real-world NHS data itself is often incomplete or inconsistently coded; one report observes that NHS data is frequently "incomplete, mislabelled or unsuitable" for AI use, limiting model performance [59, 60]. For veterans, this limitation is compounded by potential care outside the NHS (through charities like Combat Stress or private providers) [61]. Bridging gaps thus also means pulling in non-NHS data where possible (with permission), such as charity treatment records or research cohort data (like the King's Centre for Military Health Research cohort studies).

In summary, bridging data gaps for UK veterans entails:

(1) systematically identifying who the veterans are in health datasets (e.g. enforcing the use of veteran status codes nationwide),

(2) linking MoD and NHS records to get a full picture of service and post-service health,

(3) improving data accessibility and quality.

Progress is being made: the Armed Forces Covenant has heightened awareness of veterans' needs in the NHS, and projects like "Op Courage" (the Veterans Mental Health and Wellbeing Service) create dedicated pathways that might generate more consistent data [44, 62]. Nonetheless, significant work remains to be done to integrate these data streams. For deep learning models to truly capitalise on longitudinal data, they must have *longitudinal data to begin with*. Bridging the current gaps will require policy support, technical infrastructure (as depicted in Figure 2), and trust-building among stakeholders (veterans must trust that data sharing will be used to help them, not penalise them). The payoff, however, could be substantial, enabling models that understand a veteran's journey from enlistment through civilian life and using that understanding to predict better and prevent PTSD.

Interpretability and Trust in Deep Learning Applications. In applying deep learning to any medical domain, one confronts the classic tension between model complexity and interpretability. This is especially pronounced in mental health, where diagnoses like PTSD are nuanced and sensitive, and clinical decisions (e.g. whether to intervene based on a risk score proactively) carry substantial ethical implications. Clinicians, patients, and healthcare leaders alike have expressed concerns about "black-box" AI: How did the model arrive at this prediction? Can we trust it, or is it picking up spurious correlations? In the context of veteran PTSD, these questions are vital; a false prediction could lead to unwarranted anxiety or stigma for the veteran, while a missed prediction could mean a lost opportunity to prevent suffering [63, 64]. Therefore, interpretability and trustworthiness are not optional features; they are requirements for any deep learning solution in this space.

One key aspect is the need for transparent algorithms. Traditional PTSD risk assessment relies on human-understandable factors (e.g., the presence of nightmares, hypervigilance, startle responses, etc.) assessed through clinical interviews or questionnaires. If an AI model suggests a veteran is at high risk of PTSD "because of patterns in data," it must provide a rationale that clinicians can comprehend and evaluate. Various strategies exist: interpretable model design (using simpler models or constrained

models), *post-hoc* explanation tools, and visualisations of what the model "attended" to. For deep neural networks, which are inherently complex, *post-hoc interpretability* is the usual approach. For example, an attention-based model might highlight specific past medical events that most contributed to the PTSD risk score, perhaps a sequence of physical injury, relationship breakdown, and job loss notes that parallels known risk narratives. Alternatively, one can use feature importance analysis to rank predictors: a model might implicitly use features like "number of prior mental health visits" or "sleep medication refills" strongly in predicting PTSD; surfacing those can reassure clinicians that the model aligns with known risk factors. In study [24], the authors found that information from narrative notes (such as mentions of nightmares or flashbacks) had high predictive value, which both made intuitive sense and lent credence to the model's operation.

Beyond technical explanations, trust is also fostered by rigorous validation and user involvement. Healthcare professionals are more likely to trust AI if they know it has been stress-tested in conditions similar to their own practice and if they have a say in its development. A Nature review on trust in AI noted that in healthcare, *"transparency and explainability are the most important factors of trust"* among clinicians [65]. Physicians want to understand how an AI arrives at its conclusions and to verify that these align with clinical logic [66]. The same review highlighted other factors eroding trust, including the paucity of high-quality clinical trials for AI tools and concerns about bias [65]. In the PTSD prediction realm, this means AI models should be prospectively evaluated, for instance, tested on a fresh cohort of veterans to see if those flagged at high risk show higher PTSD incidence over time.

Additionally, involving mental health professionals in the design phase (e.g., selecting the outputs the model provides, such as a risk category with an explanation versus a raw score) can enhance acceptance. There is also the veteran perspective: veterans may worry that an algorithm labelling them with "PTSD risk" could affect how they're treated or even have implications for things like insurance or employment if not kept confidential [67]. Thus, building trust requires communicating to veteran patients how the AI is intended as a supportive tool for better care, not a diagnostic stamp or a replacement for human judgment.

A notable development is the pursuit of explainable AI (XAI) techniques specifically tailored to healthcare [68]. Some research is looking at hybrid models that incorporate causal reasoning or more structured knowledge (like risk factor graphs) into deep learning so that the outputs are inherently more interpretable. For example, a model might provide output such as "High PTSD risk, primarily due to the combination of prior depression + chronic pain + recent bereavement," a format that a clinician can readily grasp. While not many PTSD-specific models have achieved this level of semantic explanation yet, it's a direction worth exploring. It's also worth noting that simpler predictive tools (like the PTSD Checklist questionnaire) are already in use and are fully transparent; any AI system needs to offer a *net gain* significant enough to warrant introducing a black-box element, and part of that gain should be mitigated by making the black box as glass-like as possible [69].

Finally, trust is bolstered by robust ethical guardrails. For deep learning in veteran PTSD, this means ensuring the model does not inadvertently reinforce bias (for example, if training data were mostly male veterans, does it perform equally well for female veterans?) [32]. It also means having clear protocols: an AI might flag someone, but a human clinician should verify and discuss the matter with the patient before any action is taken, employing the so-called "human-in-the-loop" approach that NHS guidance endorses [32, 70]. The NHS AI Ethics Initiative and others emphasise that algorithms in care should be *assistive*, not autonomous, and this principle must be adhered to to maintain user trust [59].

In conclusion, interpretability and trust are about making the deep learning model not only accurate but also legible and reliable to those it serves. A powerful PTSD prediction algorithm is of little use if clinicians distrust it or if veterans feel alienated by it. By incorporating explainability (through technical means such as attention weights or SHAP values) and embedding the model's use within a framework of clinician-patient interaction (rather than as a sole decision-maker), we enhance acceptance [71]. The maxim "with great power comes enormous responsibility" applies: deep learning affords us powerful predictive tools, and it is our responsibility to wield them in a transparent, fair, and empathetic manner, particularly in the domain of mental health, where trust and understanding are of paramount importance.

*Deployment Barriers and Systemic Readiness in the NHS***.** Deploying a deep learning solution for PTSD risk detection in the NHS involves more than just having a successful model. The surrounding health system must be *ready* to integrate, support, and sustain this innovation. The UK's NHS is a large, complex mosaic of organisations with varying digital maturity levels. While some hospitals and GP practices are equipped with modern electronic health record systems and data science teams, others continue to struggle with outdated technology and limited IT resources. Basic infrastructure challenges are holding back the potential of AI: a recent analysis noted that many NHS staff are enthusiastic about AI's promise but are hindered by "crumbling" IT infrastructure, outdated computers that crash, slow internet, and systems not designed to support AI workloads [72]. In fact, it's estimated that between *10% and 50% of NHS technology systems need modernisation* to be AI-ready [72]. This uneven digital terrain means any PTSD-AI tool might work smoothly in one trust but face technical obstacles in another.

One barrier is simply ensuring access to data and computing where it's needed. Some NHS trusts have already partnered with academia to gain access to high-performance computing for AI, while others lack even the personnel to extract data from their EHRs. There are reported instances of EHR vendors not facilitating simple data export, with some requiring additional fees or proprietary tools to extract patient data [72]. This poses a challenge: our deep learning model for PTSD will only be as effective as the data pipeline feeding it. If an NHS trust cannot readily compile a veteran's longitudinal record (due to technical or contractual barriers), the model can't function. To address this, national bodies are pushing for data standards and interoperability. The NHS Long Term Plan and recent Data Strategy call for open APIs and data portability in clinical systems [73]. Overcoming vendor lock-in and fragmentation is critical; one Health Foundation report cited the lack of standardised data access across the NHS as a particular barrier to scaling AI innovations.

Another deployment barrier is workflow integration. Even if a PTSD risk model exists in the background, how does it deliver its output to clinicians in a useful way? Busy GPs or mental health clinicians will not want to log into a separate AI dashboard; they would need the alert or score surfaced in their existing workflow (e.g., a flag on the patient's electronic chart). Integration with electronic record interfaces (EMIS, SystmOne, etc. in UK primary care, or hospital EHR systems, is thus a practical hurdle. Such integration often requires collaboration with the EHR software companies and assurance that the AI's suggestions align with clinical pathways. For example, if a veteran patient is flagged as high-risk for PTSD, there should be a clear protocol: perhaps a prompt for the GP to ask certain questions or refer to a veteran-specific clinic (like Op Courage). Developing these protocols and ensuring the AI output triggers appropriate action (without causing alert fatigue) is a part of deployment that goes beyond the AI model itself.

Workforce skills and training also play a role. Many NHS clinicians have limited training in AI. A survey by the British Medical Association found that clinicians are often unsure about how to evaluate or work alongside AI tools. Deploying a deep learning model for PTSD would necessitate educating care teams about the model's capabilities, its limitations, and how to interpret its outputs. Without this, there's a risk of misuse (either blind over-reliance or dismissive underuse). Additionally, data science expertise within the NHS is growing but still limited; trusts may rely on external partners to maintain AI systems. If a model encounters issues (e.g., starts drifting in accuracy), internal capacity to detect and correct that is essential for safe ongoing deployment.

A further systemic consideration is regulatory and legal clearance. In the UK, an AI that influences clinical decisions can be considered a medical device, subject to regulation. The MHRA's pilot of an AI sandbox and the forthcoming AI-specific regulatory framework mean that our PTSD model would likely require assessment for safety, efficacy, and bias before widespread use [32]. Compliance with GDPR is also fundamental: models must justify the data they use and how they use it. Fortunately, there is momentum on this front; NHSX (now part of NHS England) published guidelines, "Artificial Intelligence: How to Get It Right," which outline best practices for validation and the ethics of AI in healthcare [74]. One key recommendation is conducting prospective trials. The cost is a barrier in itself: evaluating an AI tool in real-world NHS settings can be time-consuming and expensive, requiring approvals, data sharing agreements, and outcome tracking. Many innovators struggle with this "last

mile" of evidence generation [32], which can delay deployment.

Lastly, we must consider organisational culture and readiness for change. Introducing AI into mental health services for veterans could face resistance if staff fear it might replace human roles or if patients worry about privacy. Transparent communication and involving stakeholders (clinicians, veterans, and IT staff) from the outset can mitigate this. For example, co-design workshops with clinicians to determine how the AI should present information will result in a tool that better suits practice [32]. The King's Fund in 2025 emphasised building both the technical and human capability, training the workforce, engaging leadership, and establishing governance as essential to realising AI benefits in care.

In summary, the NHS is making strides, albeit unevenly: some trusts operate on cutting-edge systems that are on par with those of *leading global institutions, while others still rely on paper notes and fax machines* [32, 75]. This discrepancy hinders the large-scale implementation of AI, as any solution must be compatible with the most basic infrastructure. Investment in core IT (up-to-date hardware, network, and interoperable software) is needed. Encouragingly, national initiatives (like the NHS AI Lab and targeted funding for digital upgrades in hospitals) are addressing this [13]. As these barriers slowly fall, the system will become more ready to absorb innovations like our PTSD deep learning model. In planning deployment, one should adopt a phased approach: pilot in digitally mature settings (to prove value and iron out kinks), then incrementally expand while providing support to less digitally mature sites. The NHS's sheer size means scaling a solution is non-trivial, but also that the *impact of success is immense*. Overcoming the deployment barriers is the price to pay for that impact. If we can surmount these challenges — modernising infrastructure, integrating workflows, training users, and ensuring robust evaluation — we can truly leverage deep learning to support veteran mental health across the UK.

Implications for Veteran Mental Health Services in the UK

Integrating deep learning PTSD risk models into veteran mental health services could significantly enhance the proactive and personalised care veterans receive. Currently, UK veterans access mental health support through both mainstream NHS services and veteran-specific programs, such as Op COURAGE, the Veterans' Mental Health and Wellbeing Service [44, 62]. These services, launched in recent years, are designed as entry points where veterans can be assessed and routed to appropriate care (therapy, medication, social support) [76, 77]. An accurate PTSD prediction model could act as a force multiplier for such services by enabling early identification and outreach. For instance, a veteran who does not actively seek help might still show signs in their health records (perhaps repeated attendances for somatic complaints or notes about marital strain) that the model flags as high PTSD risk. The service could then reach out to this individual, even inviting them for a screening or offering resources, rather than waiting for self-presentation, thereby aligning with the NHS's prevention agenda [41].

Moreover, deep learning could help stratify and tailor interventions. Not all veterans with PTSD have the same needs; some may respond well to trauma-focused cognitive-behavioural therapy, others might need medication or social interventions (housing, employment support), and some have complex comorbidities requiring intensive care. By analysing patterns in the data, AI models might identify subgroups or *profiles* of PTSD among veterans. For example, one cluster might be veterans whose PTSD is highly associated with chronic pain and physical injury, suggesting they benefit from integrated pain management and mental health care. Another cluster might involve those with primarily moral injury or guilt-related PTSD, indicating a different therapeutic approach. This type of data-driven clustering (sometimes achieved through unsupervised deep learning) can inform service development, ensuring that UK veteran services segment their care pathways intelligently [78]. It complements the traditional "assessment then refer" model by providing an evidence-based second opinion on what trajectory a patient might follow or what resources they are likely to need.

A further implication is in **resource planning and policy**. The UK's commitment to the Armed Forces Covenant has brought veteran mental health to the forefront of policy [79]. Policymakers could use insights from longitudinal data analyses to allocate funding or design programs [80]. For instance, if the deep learning analysis reveals that a significant proportion of

high-risk veterans are not engaging with existing services, it could prompt targeted outreach initiatives or adjustments in how services are advertised and delivered. It might highlight gaps; perhaps certain regions or cohorts (like early service leavers) have worse outcomes, prompting targeted interventions [81]. Over time, deploying predictive analytics could also allow the NHS to evaluate the impact of interventions. If high-risk individuals are identified and offered preventive care, the data should show reduced progression to full PTSD or improved recovery rates, thereby justifying the approach [82]. This ties into a learning health system model, where data continuously inform improvements in care.

For frontline clinicians and support workers, the integration of AI predictions could change practice by introducing a system of alerts or reminders for veteran patients. GPs, for example, if alerted that "this patient is a veteran with elevated PTSD risk," might be prompted to ask about nightmares or hypervigilance during a routine visit for something else. It effectively increases awareness at the point of care. Similarly, therapists in veteran charities could potentially benefit from having access to a fuller NHS history (if data sharing is established), ensuring continuity and a more holistic understanding. The AI can serve as a *safety net*, catching those who might otherwise fall through the cracks (such as veterans who frequently miss appointments or move between regions, patterns that an algorithm can notice and flag for follow-up).

It is important, however, to manage expectations and maintain a human-centred approach. Veteran mental health services are fundamentally built on trust, confidentiality, and rapport. An AI should augment but never replace the human touch. For veterans, knowing that the health system is proactively monitoring their wellbeing (in a non-intrusive, supportive manner) could reinforce the message that *"we care and we're keeping an eye on your wellbeing as part of our covenant commitment."* Conversely, any sense of surveillance or loss of control must be avoided. This underscores the need for clear communication about data use; veterans should ideally consent to and be aware of the use of predictive models in their care, framed positively as an aid for their clinicians [79].

In practice, if properly implemented, veteran PTSD prediction tools could enable earlier, more precise care, which might reduce the severity and chronicity of PTSD in the community. Better outcomes for individuals (e.g., maintaining employment and family stability due to timely support) translate into societal benefits, lowering healthcare costs, and fulfilling the moral obligation to those who serve [83]. We might also see an indirect destigmatisation: if the system is actively identifying and helping veterans with PTSD, it normalises seeking help and treating PTSD as a manageable health condition, not a personal failing. Over time, successes in the veteran sector could even be transferred to civilian trauma care or other mental health arenas.

In summary, the implications are promising; deep learning could become an invaluable tool in the *proactive armamentarium* of veteran mental health services in the UK. It aligns well with current policy directions of personalised, preventive care. To realise these benefits, however, careful integration, ethical oversight, and continuous evaluation will be necessary, ensuring that the use of AI truly enhances the support structure around each veteran rather than becoming a mere tech add-on.

CONCLUSIONS

Building on this review, we propose several recommendations for research and practice at the intersection of deep learning, UK health data, and veteran PTSD:

1. Establish a Unified Veteran Data Repository: To facilitate advanced modelling, stakeholders should create a secure, integrated repository of veteran health data. This could be a virtual federated database linking MoD service records (with key variables such as deployment dates, roles, and in-service medical episodes) and NHS records (including primary care, secondary care, and mental health services) for veterans. Initiatives like the NHS England Secure Data Environment provide a blueprint [84]. A dedicated "Veteran Health Analytics Hub" could dramatically accelerate research by providing pre-linked, de-identified longitudinal datasets to approved studies [85]. We recommend the MoD, NHS, and academic partners pilot such an initiative for a cohort (e.g., all veterans of a recent deployment) to demonstrate feasibility and value.

2. Advance Longitudinal Modelling Techniques: Researchers should continue to refine deep learning models that handle irregular, long-term

sequences typical of health data [27]. This includes experimenting with transformer architectures on EHR sequences (beyond initial studies) and exploring hybrid models that combine data-driven learning with clinical knowledge (e.g., models that enforce certain temporal logic, such as the principle that PTSD cannot occur before trauma) [27, 86]. Special attention should be given to calibrating models over long follow-up periods, as many veteran outcomes unfold over decades, necessitating models with mechanisms to refresh or update with new data. Continuous-learning systems that update with new patient data (while preventing catastrophic forgetting) would be valuable for keeping predictions current [87].

3. Focus on Multimodal and Contextual Inputs: PTSD risk is influenced by context beyond clinical variables [88]. Future models should integrate social and environmental data where available, for example, incorporating data on social care, employment status, or housing (some of which might be recorded in GP notes or social services databases). Researchers may also utilise natural language processing on veterans' self-reported narratives, such as those from online forums or research interviews (if ethical), to enrich models with psychosocial context. The aim is to construct a holistic picture of each veteran's life. Additionally, leveraging sensor data (like sleep trackers or heart rate variability from wearable devices) could provide early physiological signs of PTSD exacerbation [89]. Pilot programs that provide veterans with wearable technology and feed that data into prediction models could be trialled to see if it boosts accuracy.

4. Strengthen Explainability and Clinician Interface: We recommend that any predictive model for PTSD be paired with an explainability module before deployment. Research should test different explanation approaches with clinicians and veterans to see which builds the most trust and understanding. For example, does a simple rule-based explanation suffice ("Key factors: increasing anxiety medication usage and recent bereavement") or are visual timelines better? User-centred design studies can optimise how model outputs are communicated [90]. Moreover, the interface in clinical systems should be refined: perhaps risk scores are embedded as a small icon in the GP's screen that expands to show details when clicked. Co-designing this interface with end users (GPs, psychiatrists, and veteran support workers) is recommended to ensure it enhances rather than hinders workflow [32].

5. Rigorous Prospective Trials and Health Economic Evaluation: While retrospective accuracy studies are valuable, the true test is prospective impact. We call for trials of the deep learning PTSD risk system in a real-world setting. For example, a cluster-randomised trial could be conducted in which some NHS regions implement the AI-assisted veteran care pathway. In contrast, others continue with standard care, measuring outcomes such as PTSD incidence, time to treatment, symptom severity reduction, and patient satisfaction. Alongside, perform a health economics analysis: does the early intervention prompted by AI lead to cost savings (e.g. fewer hospitalisations, quicker return to work)? [91] Evidence of cost-effectiveness will be crucial for the NHS to adopt it broadly. Additionally, such trials should include monitoring for any unintended consequences (like over-referral or increased anxiety from risk labelling) to ensure patient safety and benefit.

6. Education and Training Modules: To prepare the system, develop training programs about AI in mental health for both clinicians and patients [92]. For clinicians, short courses or e-learning can demystify how the PTSD risk model works, its limitations, and how to communicate its results to patients [93]. For veterans, informational materials can explain that an AI might assist their care ("the NHS is using advanced computer models that look at your health data to support you better; here's what that means…"). Transparent education will preempt misunderstandings and promote collaborative use of the tool. We also suggest involving veteran advocacy groups in crafting the messaging to ensure it resonates and alleviates concerns about privacy or stigma.

7. Policy Frameworks for Ethical AI Use: On a policy level, we recommend that the Department of Health and Social Care and MoD develop joint guidelines for AI in veteran healthcare [94]. This could cover data governance specifics for veteran records, consent processes for data use (perhaps an "opt-out" model with safeguards), and clarity on liability (e.g., if an AI misses a PTSD case, how is accountability handled?). Additionally, ensure alignment with emerging AI ethics frameworks by implementing routine bias audits of the model (to evaluate performance across subgroups, such as female versus male veterans or different service

branches) and instituting an oversight committee that includes veterans, ethicists, and clinicians to review the AI's operation [95] continuously.

8. Extend to Related Mental Health Needs: Finally, while PTSD is a focal point, veterans often face a suite of mental health challenges (depression, anxiety, substance use, traumatic brain injury effects). Future research can enhance models to predict or detect comorbid conditions and transitions between them, such as the risk of alcohol misuse relapse in veterans with PTSD. A comprehensive deep learning system might eventually track overall "veteran wellbeing" using multiple outputs. This aligns with holistic care: addressing not just PTSD in isolation but the full spectrum of a veteran's mental health. Multi-task learning models, which can output predictions for several related conditions, could be explored [22].

By pursuing these recommendations, we can move from concept to reality in using deep learning to combat PTSD among UK veterans. It will require a multidisciplinary, collaborative effort, bringing together data scientists, clinicians, NHS IT personnel, MoD officials, ethicists, and the veteran community. The momentum exists and the need is clear; with thoughtful action, the next decade could see AI-driven mental health support become an integral part of how the UK honours its commitment to those who have served.

The risk of posttraumatic stress in UK veterans fluctuates over time, not due to clinic visits. Deep learning applied to longitudinal records, linking Ministry of Defence service histories with NHS data, offers a practical route from reactive care to anticipatory medicine. Sequence models can detect subtle precursors, integrate multimodal signals, and support earlier, more tailored intervention without displacing clinical judgment.

Real impact, however, depends on system choices. Veteran identifiers and interoperable data flows must be enforced; secure data environments or federated approaches should govern linkage; and outputs must be legible at the point of care, pairing risk scores with concise, clinician-tested explanations. Deployment requires routine bias audits, prospective evaluation, and workforce preparation so that models augment decisions rather than automate them.

The prize is clear: fewer missed cases, faster access to effective support, and services that learn from their data. With a focused program, establishing a Veteran Health Analytics Hub, conducting prospective trials with health-economic endpoints, and co-designing interfaces with clinicians and veterans, the UK can leverage longitudinal data and modern AI to create a reliable early-warning system for those who have served.

REFERENCES

1. Morgan, L., & Aldington, D. (2019). Comorbid chronic pain and posttraumatic stress disorder in UK veterans: a lot of theory but not enough evidence. *British Journal of Pain, 14*(4), 256–262. doi: 10.1177/2049463719878753

2. Jones, S., Smith, J., Green, N., Doe, D., & Foster, N. (2025). Current perspectives on the mental health of UK military personnel and veterans. *British Medical Bulletin, 154*(1).

3. White, H., Williams, C., Miller, S., Foster, N., Harris, S., & Moore, S. (2025, June 26). *Chapter 3: Posttraumatic stress disorder*. Retrieved from https://digital.nhs.uk/data-and-information/publications/statistical/adult-psychiatric-morbidity-survey/survey-of-mental-health-and-wellbeing-england-2023-24/posttraumatic-stress-disorder

4. Rytwinski, N. K., Scur, M. D., Feeny, N. C., & Youngstrom, E. A. (2013). The Co-Occurrence of Major Depressive Disorder Among Individuals With Posttraumatic Stress Disorder: A Meta-Analysis. *Journal of Traumatic Stress, 26*(3), 299–309. doi: 10.1002/jts.21814

5. Andrews, B., Brewin, C. R., Philpott, R., & Stewart, L. (2007). Delayed-Onset Posttraumatic Stress Disorder: A Systematic Review of the Evidence. *American Journal of Psychiatry, 164*(9), 1319–1326. doi: 10.1176/appi.ajp.2007.06091491

6. Bonanno, G. A., Mancini, A. D., Horton, J. L., Powell, T. M., Leardmann, C. A., Boyko, E. J., et al. (2012). Trajectories of trauma symptoms and resilience in deployed U.S. military service members: Prospective cohort study. *British Journal of Psychiatry, 200*(4), 317–323. doi: 10.1192/bjp.bp.111.096552

7. Galatzer-Levy, I. R. (2014). Empirical characterisation of heterogeneous posttraumatic stress responses is necessary to improve the science of posttraumatic stress: Commentary. *Journal of Clinical Psychiatry, 75*(9), e950–e952. doi: 10.4088/JCP.14com09368

8. Sareen, J. (2014). Posttraumatic stress disorder in adults: Impact, comorbidity, risk factors, and treatment. *Canadian Journal of Psychiatry, 59*(9), 460–467. doi: 10.1177/070674371405900902

9. Smith, J. R., Workneh, A., & Yaya, S. (2020). Barriers and facilitators to help-seeking for individuals with posttraumatic stress disorder: A systematic review. *Journal of Traumatic Stress, 33*(2), 137–150. doi: 10.1002/jts.22469

10. NHS England Digital. (2025). Research powered by data. Retrieved from https://digital.nhs.uk/data-and-information/research-powered-by-data

11. PTSD UK. (2025). NICE guidelines in the UK for PTSD. Retrieved from https://www.ptsduk.org/what-is-ptsd/nice-guidelines-for-ptsd/

12. Papini, S., Norman, S. B., Campbell-Sills, L., Sun, X., He, F., Kessler, R. C., et al. (2023). Development and validation of a machine learning prediction model of posttraumatic stress disorder after military deployment. *JAMA Network Open, 6*(6), e2321273. doi: 10.1001/jamanetworkopen.2023.21273

13. NHS England Digital. (2025). AI knowledge repository. Retrieved from https://digital.nhs.uk/services/ai-knowledge-repository

14. Hughes, O. (2025). Researchers pilot AI model to predict future health outcomes. Retrieved from https://www.digitalhealth.net/2025/05/researchers-pilot-ai-model-to-predict-future-health-outcomes/

15. Blekic, W., D'Hondt, F., Shalev, A. Y., & Schultebraucks, K. (2025). A systematic review of machine learning findings in PTSD and their relationships with theoretical models. *Nature Mental Health, 3*(1), 139–158.

16. Karstoft, K. I., Galatzer-Levy, I. R., Statnikov, A., Li, Z., & Shalev, A. Y. (2015). Bridging a translational gap: Using machine learning to improve the prediction of PTSD. *BMC Psychiatry, 15*(1), 30.

17. Karstoft, K. I., Statnikov, A., Andersen, S. B., Madsen, T., & Galatzer-Levy, I. R. (2015). Early identification of posttraumatic stress following military deployment: Application of machine learning methods to a prospective study of Danish soldiers. *Journal of Affective Disorders, 184*, 170–175. doi: 10.1016/j.jad.2015.02.043

18. Galatzer-Levy, I. R., Ankri, Y., Freedman, S., Israeli-Shalev, Y., Roitman, P., Gilad, M., et al. (2013). Early PTSD symptom trajectories: Persistence, recovery, and response to treatment: Results from the Jerusalem Trauma Outreach and Prevention Study (J-TOPS). *PLOS ONE, 8*(8), e70084.

19. Miotto, R., Li, L., Kidd, B. A., & Dudley, J. T. (2016). Deep Patient: An unsupervised representation to predict the future of patients from the electronic health records. *Scientific Reports, 6*(1), 26094.

20. Cascarano, A., Mur-Petit, J., Hernández-González, J., Camacho, M., de Toro Eadie, N., Gkontra, P., et al. (2023). Machine and deep learning for longitudinal biomedical data: A review of methods and applications. *Artificial Intelligence Review, 56*(2), 1711–1771. doi: 10.1007/s10462-023-10561-w

21. Razavi, M., Ziyadidegan, S., Mahmoudzadeh, A., Kazeminasab, S., Baharlouei, E., Janfaza, V., et al. (2024). Machine learning, deep learning, and data preprocessing techniques for detecting, predicting, and monitoring stress and stress-related mental disorders: Scoping review. JMIR Mental Health, 11(1), e53714. doi: 10.2196/53714

22. Wu, Y., Mao, K., Dennett, L., Zhang, Y., & Chen, J. (2023). Systematic review of machine learning in PTSD studies for automated diagnosis evaluation. *NPJ Mental Health Research, 2*(1), 16.

23. Wang, J., Ouyang, H., Jiao, R., Cheng, S., Zhang, H., Shang, Z., et al. (2024). The application of machine learning techniques in posttraumatic stress disorder: A systematic review and meta-analysis. *NPJ Digital Medicine, 7*(1), 121.

24. Zafari, H., Kosowan, L., Zulkernine, F., & Signer, A. (2021). Diagnosing posttraumatic stress disorder using electronic medical record data. *Health Informatics Journal, 27*(4), 14604582211053259. doi: 10.1177/14604582211053259

25. Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., et al. (2018). Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine, 1*(1), 18.

26. Islam, M. M., Hassan, S., Akter, S., Jibon, F. A., & Sahidullah, M. (2024). A comprehensive review of predictive analytics models for mental illness using machine learning algorithms. *Healthcare Analytics, 6*, 100350. doi: 10.1016/j.health.2024.100350

27. Yang, Z., Mitra, A., Liu, W., Berlowitz, D., & Yu, H. (2023). TransformEHR: Transformer-based encoder-decoder generative model to enhance prediction of disease outcomes using electronic health records. *Nature Communications, 14*(1), 7857. doi: 10.1038/s41467-023-43491-7

28. Jin, D., Sergeeva, E., Weng, W. H., Chauhan, G., & Szolovits, P. (2021). *Explainable deep learning in healthcare: A methodological survey from an attribution view*. Retrieved from http://arxiv.org/abs/2112.02625

29. Kohan, A., Zahedi, A., Alizadehsani, R., Tan, R. S., & Acharya, U. R. (2025). Application of explainable artificial intelligence (XAI) techniques in patients with intracranial hemorrhage: A systematic review. WIREs Data Mining and Knowledge Discovery, 15(3), e70031. 10.1002/widm.70031

30. Løkhammer, S., Koller, D., Wendt, F. R., Choi, K. W., He, J., Friligkou, E., et al. (2024). Distinguishing vulnerability and resilience to posttraumatic stress disorder evaluating traumatic experiences, genetic risk and electronic health records. *Psychiatry Research, 337*, 115950. doi: 10.1016/j.psychres.2024.115950

31. Vali, M., Nezhad, H. M., Kovacs, L., & Gandomi, A. H. (2025). Machine learning algorithms for predicting PTSD: a systematic review and meta-analysis. *BMC Medical Informatics and Decision Making, 25*(1), 34. doi: 10.1186/s12911-024-02243-2

32. Hall, P. D. (2025). Barriers to AI innovation in healthcare. BCS, The Chartered Institute for IT. Retrieved from https://www.bcs.org/articles-opinion-and-research/barriers-to-ai-innovation-in-healthcare/

33. Committee on the Assessment of Ongoing Efforts in the Treatment of PTSD, Institute of Medicine. (2012). History, diagnostic criteria, and epidemiology. In Treatment for Posttraumatic Stress Disorder in Military and Veteran Populations: Initial Assessment. National Academies Press (US). Retrieved from https://www.ncbi.nlm.nih.gov/books/NBK201095/

34. Galatzer-Levy, I. R., & Bryant, R. A. (2013). 636,120 ways to have posttraumatic stress disorder. *Perspectives on Psychological Science, 8*(6), 651–662. doi: 10.1177/1745691613504116

35. Porter, B., Bonanno, G. A., Frasco, M. A., Dursa, E. K., & Boyko, E. J. (2017). Prospective posttraumatic stress disorder symptom trajectories in active duty and separated military personnel. *Journal of Psychiatric Research, 89*, 55–64. doi: 10.1016/j.jpsychires.2017.01.014

36. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2023). *Attention is all you need*. Retrieved from http://arxiv.org/abs/1706.03762

37. Tomas, C. W., Fitzgerald, J. M., Bergner, C., Hillard, C. J., Larson, C. L., & deRoon-Cassini, T. A. (2022). Machine learning prediction of posttraumatic stress disorder trajectories following traumatic injury: Identification and validation in two independent samples. *Journal of Traumatic Stress, 35*(6), 1656–1671. doi: 10.1002/jts.22908

38. Swinckels, L., Bennis, F. C., Ziesemer, K. A., Scheerman, J. F. M., Bijwaard, H., de Keijzer, A., et al. (2024). The use of deep learning and machine learning on longitudinal electronic health records

for the early detection and prevention of diseases: Scoping review. *Journal of Medical Internet Research, 26*(1), e48320. doi: 10.2196/48320

39. Ben-Zion, Z., Simon, A. J., Rosenblatt, M., Korem, N., Duek, O., Liberzon, I., et al. (2025). Connectome-based predictive modeling of PTSD development among recent trauma survivors. *JAMA Network Open, 8*(3), e250331. doi: 10.1001/jamanetworkopen.2025.0331
40. Williamson, V., Stevelink, S. A. M., & Greenberg, N. (2018). Occupational moral injury and mental health: Systematic review and meta-analysis. *British Journal of Psychiatry, 212*(6), 339–346. doi: 10.1192/bjp.bp.117.202339
41. Burnett, H., Garratt, K., & Powell, T. (2025). Veterans: Access to health services. Retrieved from https://commonslibrary.parliament.uk/research-briefings/cbp-10126/
42. Charles, A. (2022). Integrated care systems explained. The King's Fund. Retrieved from https://www.kingsfund.org.uk/insight-and-analysis/long-reads/integrated-care-systems-explained
43. Gov.uk. (2025). Health and wellbeing services for armed forces personnel and their families. Retrieved from https://www.gov.uk/guidance/looking-after-our-armed-forces
44. Juškaitė, S., Stone, J., Greenberg, N., Dyball, D., & Fear, N. T. (2025). Current perspectives on the mental health of UK military personnel and veterans. *British Medical Bulletin, 154*(1), ldaf003. doi: 10.1093/bmb/ldaf003
45. Williamson, C., Palmer, L., Leightley, D., Pernet, D., Chandran, D., Leal, R., et al. (2023). Military veterans and civilians' mental health diagnoses: An analysis of secondary mental health services. *Social Psychiatry and Psychiatric Epidemiology, 58*(7), 1029–1037. doi: 10.1007/s00127-023-02492-7
46. Mark, K. M., Leightley, D., Pernet, D., Murphy, D., Stevelink, S. A. M., & Fear, N. T. (2019). Identifying veterans using electronic health records in the United Kingdom: A feasibility study. *Healthcare, 8*(1), 1. doi: 10.3390/healthcare8010001
47. Mellotte, H., Murphy, D., Rafferty, L., & Greenberg, N. (2017). Pathways into mental health care for UK veterans: A qualitative study. *European Journal of Psychotraumatology, 8*(1), 1389207. doi: 10.1080/20008198.2017.1389207
48. Neath, T. (2025). Support for service leavers, veterans and their families. Warwickshire County Council. Retrieved from https://www.warwickshire.gov.uk/support-armed-forces/support-service-leavers-veterans-families/
49. Forces in Mind Trust. (2022). Researchers develop new tool to identify veterans in secondary healthcare services. Retrieved from https://www.fim-trust.org/news-policy-item/researchers-develop-new-tool-to-identify-veterans-in-secondary-healthcare-services/
50. Wessex LMCs. (2025). Registering and care of military personnel & veterans. Wessex LMCs. Retrieved from https://www.wessexlmcs.com/guidance/registering-and-care-of-military-personnel-veterans/
51. NHS. (N. d.). Healthcare for the Armed Forces community. Retrieved from https://www.nhs.uk/nhs-services/armed-forces-community/
52. Royal College of General Practitioners. (2025). Veterans' Health Hub: Guidance for GPs. Veterans' Retrieved from https://elearning.rcgp.org.uk/mod/page/view.php?id=14970
53. Joyce, D. W., Kormilitzin, A., Smith, K. A., & Cipriani, A. (2023). Explainable artificial intelligence for mental health through transparency and interpretability for understandability. *NPJ Digital Medicine, 6*, 6. doi: 10.1038/s41746-023-00742-4

54. University of Oxford. (2024). New system protects patient data through federated learning. Retrieved from https://www.ndm.ox.ac.uk/news/new-system-protects-patient-data-through-federated-learning

55. Petrova, A. (2025). Energy-efficient machine learning inference in edge computing for healthcare IoT. Retrieved from https://nhsjcs.com/2025/02/25/energy-efficient-machine-learning-inference-in-edge-computing-for-healthcare-iot/

56. NHS England. Foresight AI case study. *NHS England Digital*. Retrieved from https://digital.nhs.uk/data-and-information/research-powered-by-data/case-studies/foresight-ai

57. King's College London. (2021). Validating the Military Service Identification Tool. Retrieved from https://clinicaltrials.gov/study/NCT04299789

58. NHS England. (2025). Artificial Intelligence. Retrieved from https://transform.england.nhs.uk/information-governance/guidance/artificial-intelligence/

59. England NHS. (2025). Artificial intelligence (AI) and machine learning. Retrieved from https://www.england.nhs.uk/long-read/artificial-intelligence-ai-and-machine-learning/

60. Thompson, C. L., & Morgan, H. M. (2020). Ethical barriers to artificial intelligence in the national health service, United Kingdom. *Bulletin of the World Health Organization, 98*(4), 293–295. doi: 10.2471/BLT.19.237289

61. Combat Stress. (N. d.). Mental health services for veterans. Retrieved from https://combatstress.org.uk

62. Finnegan, A., Salem, K., Green, N., Ainsworth-Moore, L., & Ghomi, M. (2023). Evaluation of the NHS England' Op COURAGE' High Intensity Service. *BMJ Military Health, 171*(1), e002385. doi: 10.1136/bmjmilitary-2022-002385

63. Rosenbacke, R., Melhus, Å., McKee, M., & Stuckler, D. (2024). How explainable artificial intelligence can affect clinicians' trust. *JMIR AI, 3*(1), e53207. doi: 10.2196/53207

64. Rudin, C. (2019). Stop explaining black box machine learning models. *Nature Machine Intelligence, 1*(5), 206–215. doi: 10.1038/s42256-019-0048-x

65. Afroogh, S., Akbari, A., Malone, E., Kargar, M., & Alambeigi, H. (2024). Trust in AI: Progress, challenges, and future directions. *Humanities and Social Sciences Communications, 11*(1), 1568. doi: 10.1057/s41599-024-02668-7

66. Amann, J., Blasimme, A., Vayena, E., Frey, D., & Madai, V. I. (2020). Explainability for artificial intelligence in healthcare. *BMC Medical Informatics and Decision Making, 20*(1), 310. doi: 10.1186/s12911-020-01332-6

67. Tucci, V., Saary, J., & Doyle, T. E. (2022). Factors influencing trust in medical artificial intelligence. *Journal of Medical Artificial Intelligence, 5*.

68. Sadeghi, Z., Alizadehsani, R., Cifci, M. A., Kausar, S., Rehman, R., & Mahanta, P. (2024). A review of explainable artificial intelligence in healthcare. *Computers & Electrical Engineering, 118*, 109370. doi: 10.1016/j.compeleceng.2023.109370

69. Weathers, F. W., Litz, B. T., Keane, T. M., Palmieri, P. A., Marx, B. P., & Schnurr, P. P. (2013). PTSD Checklist for DSM-5 (PCL-5). *National Center for PTSD*. Retrieved from https://www.ptsd.va.gov/professional/assessment/adult-sr/ptsd-checklist.asp

70. GOV.UK. (2023). Using AI to improve patient bed allocation decisions in hospital. Retrieved from https://www.gov.uk/government/case-studies/using-ai-to-improve-patient-bed-allocation-decisions-in-hospital

71. Alam, M. N., Kaur, M., & Kabir, S. (2023). Explainable AI in Healthcare: Enhancing Transparency and Trust upon Legal and Ethical Consideration. *International Research Journal of Engineering and Technology, 10*(06).

72. Mistry, P. (2025). Infrastructure for innovation: getting the NHS and social care ready for AI. Retrieved from https://www.kingsfund.org.uk/insight-and-analysis/long-reads/infrastructure-nhs-social-care-ai

73. Refsum, C., Bradshaw, A., Large, O., & Ward, D. (2024). Preparing the NHS for the AI Era: A Digital Health Record for Every Citizen. Retrieved from https://institute.global/insights/public-services/preparing-the-nhs-for-the-ai-era-a-digital-health-record-for-every-citizen

74. Joshi, I., & Morley, J. (2019). Artificial Intelligence: How to get it right. Retrieved from https://law.queensu.ca/sites/lawwww/files/files/Research/MachineMD/NHSX_AI_report.pdf

75. Horton, T., & Gerhold, M. (2024). Can technology and AI 'save the NHS'? The Health Foundation. Retrieved from https://www.health.org.uk/features-and-opinion/blogs/can-technology-and-ai-save-the-nhs-a-look-at-the-main-party-manifestos

76. GOV.UK. (2024). Access Op COURAGE: The Veterans Mental Health and Wellbeing Service. Retrieved from https://www.gov.uk/support-for-veterans/op-courage-the-veterans-mental-health-and-wellbeing-service

77. NHS. (2021). Mental health support for veterans, service leavers and reservists. Retrieved from https://www.nhs.uk/nhs-services/armed-forces-community/mental-health/veterans-reservists/

78. Benedict, T. M., Nitz, A. J., Gambrel, M. K., & Louw, A. (2023). Pain neuroscience education improves posttraumatic stress disorder, disability, and pain self-efficacy in veterans and service members with chronic low back pain: Preliminary results from a randomised controlled trial with 12-month follow-up. *Military Psychology, 36*(4), 376–392. doi: 10.1080/08995605.2023.2188046

79. Armed Forces Covenant. (2025). Covenant support for Service leavers and Veterans: Healthcare. Retrieved from https://www.armedforcescovenant.gov.uk/armed-forces-community/service-leavers-and-veterans/healthcare/

80. GOV.UK. (2022). Veterans' Strategy Action Plan: 2022 to 2024. Retrieved from https://www.gov.uk/government/publications/veterans-strategy-action-plan-2022-to-2024/veterans-strategy-action-plan-2022-to-2024-html

81. Buckman, J. E. J., Forbes, H. J., Clayton, T., Jones, M., Jones, N., Greenberg, N., Sundin, J., Hull, L., Wessely, S., & Fear, N. T. (2012). Early Service leavers: a study of the factors associated with premature separation from the UK Armed Forces and the mental health of those that leave early. *The European Journal of Public Health, 23*(3), 410–415. doi: 10.1093/eurpub/cks042

82. Kilbourne, A. M., Evans, E., & Atkins, D. (2021). Learning health systems: Driving real-world impact in mental health and substance use disorder research. FASEB BioAdvances, 3(8), 626–638. doi: 10.1096/fba.2020-00124

83. RAND Health Care. (2019). Improving the Quality of Mental Health Care for Veterans. Retrieved from https://www.rand.org/pubs/research_briefs/RB10087.html

84. NHS England. (2025). Secure Data Environments (SDEs). Retrieved from https://transform.england.nhs.uk/key-tools-and-info/data-saves-lives/secure-data-environments/

85. UCL. (2025). AI model trained on de-identified data from 57 million people. Retrieved from https://www.ucl.ac.uk/news/2025/may/ai-model-trained-de-identified-data-57-million-people

86. Li, Y., Rao, S., Solares, J. R. A., Hassaine, A., Ramakrishnan, R., Canoy, D., Zhu, Y., Rahimi, K., & Salimi-Khorshidi, G. (2020). BEHRT: Transformer for Electronic Health Records. *Scientific Reports, 10*(1). doi: 10.1038/s41598-020-62922-y

87. Armstrong, J., & Clifton, D. (2022). Continual learning of longitudinal health records. Retrieved from http://arxiv.org/abs/2112.11944

88. Worthington, M. A., Mandavia, A., & Richardson-Vejlgaard, R. (2020). Prospective prediction of PTSD diagnosis in a nationally representative sample using machine learning. *BMC Psychiatry, 20*(1). doi: 10.1186/s12888-020-02933-1

89. Guichard, L., An, X., Neylan, T. C., Clifford, G. D., Li, Q., Ji, Y., Macchio, L., Baker, J., Beaudoin, F. L., Jovanovic, T., Linnstaedt, S. D., Germine, L. T., Bollen, K. A., Rauch, S. L., Haran, J. P., Storrow, A. B., Lewandowski, C., Musey, P. I., Hendry, P. L., … McLean, S. A. (2024). Heart rate variability wrist-wearable biomarkers identify adverse posttraumatic neuropsychiatric sequelae after traumatic stress exposure. *Psychiatry Research, 342*, 116260. doi: 10.1016/j.psychres.2024.116260

90. Park, A. H., Patel, H., Mirabelli, J., Eder, S. J., Steyrl, D., Lueger-Schuster, B., et al. (2025). Machine learning models predict PTSD severity and functional impairment: A personalised medicine approach for uncovering complex associations among heterogeneous symptom profiles. *Psychology Trauma: Theory, Research, Practice, and Policy, 17*(2), 372–386. doi: 10.1037/tra0001472

91. Kastrup, N., Holst-Kristensen, A. W., & Valentin, J. B. (2024). Landscape and challenges in economic evaluations of artificial intelligence in healthcare: A systematic review of methodology. *BMC Digital Health, 2*(1), 39. doi: 10.1186/s44247-024-00039-3

92. Charow, R., Jeyakumar, T., Younus, S., Dolatabadi, E., Salhia, M., Al-Mouaswas, D., et al. (2021). Artificial Intelligence Education Programs for Health Care Professionals: Scoping Review. *JMIR Medical Education, 7*(4), e31043. doi: 10.2196/31043

93. Schubert, T., Oosterlinck, T., Stevens, R. D., Maxwell, P. H., & van der Schaar, M. (2025). AI education for clinicians. *EClinicalMedicine, 79*, 102968. doi: 10.1016/j.eclinm.2024.102968

94. Gerhold, M., Nell, T., Hardie, T., & Horton, T. (2024). Priorities for an AI in health care strategy. The Health Foundation. Retrieved from https://www.health.org.uk/reports-and-analysis/briefings/priorities-for-an-ai-in-health-care-strategy

95. GOV.UK. (2024). MHRA's AI regulatory strategy ensures patient safety and industry innovation into 2030. Retrieved from https://www.gov.uk/government/news/mhras-ai-regulatory-strategy-ensures-patient-safety-and-industry-innovation-into-2030