ORIGINAL RESEARCH
Open access

Contrastive Learning with Prototypical Networks for Few-Shot Detection of Emerging Infectious Disease Outbreaks from Emergency Department Chief Complaints and Triage Notes

Natalie Brown[1*], Emma Taylor[1], James Wilson[2], Oliver Scott[1]

Abstract

Emergency department chief complaints and triage notes are early indicators of health changes during infectious disease outbreaks. These records, made before confirmatory testing, provide a presyndromic view of population health. Traditional syndromic surveillance relies on predefined syndrome categories, which may not align with novel pathogens. Early outbreaks often present as sparse, ambiguous symptom clusters, resulting in few labeled examples for automated detection. This framework suggests using contrastive learning with prototypical networks for few-shot detection of emerging infectious disease syndromes from free-text notes. It leverages historical data to create a robust clinical text embedding space, with a small set of labeled examples defining new syndromes. The system includes a contrastive pre-training encoder, prototypical network, and few-shot classifier. The encoder learns from unlabelled historical notes, and the prototypical network creates syndrome prototypes from a few labeled examples. This framework is designed for situations where public health officials observe early suspect cases but lack mature labeled datasets. It can identify early clusters by comparing incoming notes to emerging syndrome prototypes. Contrastive learning with prototypical networks enables proactive presyndromic surveillance, allowing rapid adaptation during the early phase of an outbreak without relying on large labeled datasets.

Keywords Contrastive learning, Few-shot learning, Prototypical networks, Syndromic surveillance, Emergency department chief complaints, Triage notes

*Correspondence:
Natalie Brown
natalie.brown@gmail.com
[1] Department of Intelligent Healthcare Engineering, University of Bristol, Bristol, United Kingdom
[2] Department of Clinical AI Systems, University of Nottingham, Nottingham, United Kingdom

Introduction

Early outbreak detection remains a central challenge for public health because the interval between first clinical presentation and formal case definition can determine whether an emerging infectious disease is contained or widely disseminated. Existing electronic surveillance systems such as ESSENCE have demonstrated the operational value of monitoring emergency department data streams, but many such systems remain organized around predefined syndromic categories that may not fully represent novel clinical patterns [1]. Artificial intelligence has expanded the surveillance toolkit from local anomaly detection to global epidemic monitoring, yet the capacity to recognize an unknown syndrome from sparse textual evidence remains conceptually underdeveloped [2]. This gap is especially important because emerging outbreaks may initially appear as subtle deviations in ordinary complaint language rather than as clearly labelled disease events [3, 4].

Emergency department chief complaints and triage notes are valuable because they are collected close to the point of care, often before laboratory confirmation, discharge diagnosis, or administrative coding. Studies using emergency department chief complaints, triage narratives, and related clinical text have shown that free-text surveillance can support case identification across febrile, trauma, overdose, and mental health contexts, while also revealing the limitations of relying only on structured discharge information [5-8]. However, new pathogens such as those associated with coronavirus disease, Ebola virus disease, or Zika virus disease may initially present through a small number of heterogeneous complaints, making conventional supervised learning difficult when labelled examples are scarce [4, 9]. Few-shot and zero-shot approaches are therefore conceptually attractive because they address the mismatch between the urgency of detection and the slow accumulation of confirmed labels [10, 11].

Table 1 contrasts the structural assumptions of traditional syndromic surveillance with the proposed few-shot prototype-based framework, highlighting key differences in adaptability, data requirements, and early detection capability.

Table 1. Structural comparison of traditional syndromic surveillance and few-shot prototype-based outbreak detection

Dimension	Traditional Syndromic Surveillance	Proposed Few-Shot Prototype Framework
Syndrome Definition	Predefined categories (e.g., ILI, GI illness)	Dynamically constructed from few-shot examples
Data Dependency	Requires historical labeled categories	Leverages unlabelled text with minimal labeled support
Adaptation to Novel Pathogens	Limited; requires reclassification	High; new prototype defined from few cases
Feature Representation	Keyword rules or supervised features	Contrastively learned semantic embeddings
Early Outbreak Sensitivity	Moderate; depends on category fit	High; detects similarity to emerging patterns
Handling Sparse Data	Weak performance with few examples	Designed explicitly for few-shot conditions
Interpretability	Rule-based or aggregate counts	Prototype-linked examples enable traceability
Scalability	Scales with predefined ontology	Scales via embedding space generalization
False Positive Risk	Elevated under ambiguous categories	Managed via similarity thresholds + unknown class
Role of Human Oversight	Required for validation	Central (human-in-the-loop prototype refinement)

This article proposes a conceptual framework that combines contrastive learning and prototypical networks for few-shot detection of emerging infectious disease outbreaks from emergency department chief complaints and triage notes. Contrastive learning can shape a clinical-text embedding space using unlabelled or weakly labelled historical data, while prototypical networks can classify new notes by their distance to a small number of syndrome prototypes [12-14]. The framework is not presented as an experimental system and does not claim empirical performance; instead, it clarifies a design logic for future implementation, evaluation, and public health integration. The article proceeds from surveillance data and few-shot learning foundations to a framework architecture, contrastive pre-training strategy, and prototypical classification mechanism for emerging outbreak detection [15-17].

Background

ED chief complaints as surveillance data

Emergency department chief complaints are short free-text strings recorded at triage, such as “cough and nasal congestion,” “fever with body aches,” “diffuse rash,” or "vomiting after travel," and triage notes often add contextual detail about duration, exposure, acuity, or associated symptoms. These texts are operationally important because they precede diagnosis and can be aggregated across facilities for near-real-time surveillance, as illustrated by electronic surveillance systems and public health applications that monitor emergency department inputs [1]. Natural language processing studies of triage texts show that these data can support timely case detection, although their utility depends on how well models handle spelling variation, abbreviations, negation, and locally specific clinical language [5]. Because chief complaints may diverge from discharge diagnoses, they should be understood not as definitive labels but as early signals that require probabilistic interpretation and public health review [7, 8].

Few-shot learning in epidemiology

Few-shot learning is relevant to epidemiology because emerging infectious diseases often produce actionable public health concern before large labelled datasets exist. During the early stages of outbreaks associated with coronavirus disease, Ebola virus disease, Zika virus disease, or other rare pathogens, investigators may have only a small number of confirmed or suspected examples from which to infer a broader symptom pattern [4, 9]. Machine learning for emerging infectious disease field responses has therefore emphasized flexible methods that can generalize under data scarcity, heterogeneous data quality, and rapidly changing operational requirements [4]. Reviews of few-shot learning in medical text and biomedical time-series domains similarly frame limited labels as a structural feature of healthcare AI rather than an exceptional inconvenience [10, 18].

Contrastive learning for medical text

Contrastive learning provides a way to learn representations by pulling semantically related examples closer together and pushing unrelated examples farther apart in an embedding space. In clinical natural language processing, contrastive approaches have been used for diagnostic embedding, automated coding, biomedical retrieval, data harmonization, and low-resource domain adaptation, indicating their conceptual relevance to sparse-label healthcare settings [12, 13, 16, 19, 20]. Although methods such as Momentum Contrast, SimCLR, and related contrastive paradigms were not created specifically for emergency department surveillance, their principle of learning invariant representations from augmented views can be adapted to chief complaints and triage notes [12, 13]. In this framework, contrastive learning is used not to generate outbreak labels directly but to produce a clinical-text space in which symptomatically similar complaints become easier to compare under few-shot conditions [16, 19, 21].

Prototypical networks

Prototypical networks are a meta-learning approach in which each class is represented by a prototype, commonly computed as the mean embedding of its labelled support examples [14]. A query example is classified according to its distance from these prototypes, making the method well suited to settings where a new class must be defined from a small number of examples rather than from thousands of labelled cases [14, 22]. In text classification, prototypical and induction-based few-shot methods have shown how learned embeddings can support classification when examples per class are limited, while prompt-based few-shot learning has further broadened the design space for low-label natural language processing [22-24]. For outbreak detection, the prototype concept is especially intuitive because public health officials can define an emerging syndrome through representative complaints and then search for nearby incoming notes in the learned embedding space [10, 11].

Framework Overview

High-level architecture

The proposed architecture begins with a historical corpus of unlabelled emergency department chief complaints and triage notes, which is used to pre-train a clinical-text encoder through a contrastive objective. Once an emerging syndrome is suspected, a small labelled support set of representative complaints is embedded by the encoder and summarized into one or more class prototypes using a prototypical network [14]. Incoming emergency department notes are then embedded as query examples and compared with existing prototypes to estimate whether they resemble the emerging syndrome, background respiratory illness, gastrointestinal illness, rash illness, or an unknown class [1, 5]. The output is not a diagnosis but a surveillance-oriented outbreak probability or similarity score that can support epidemiologic investigation when interpreted alongside time, geography, and clinical context [2, 3].

Figure 1 illustrates the linear architecture of the proposed framework, emphasizing the progression from contrastive representation learning to prototype-based few-shot outbreak detection and human-in-the-loop interpretation

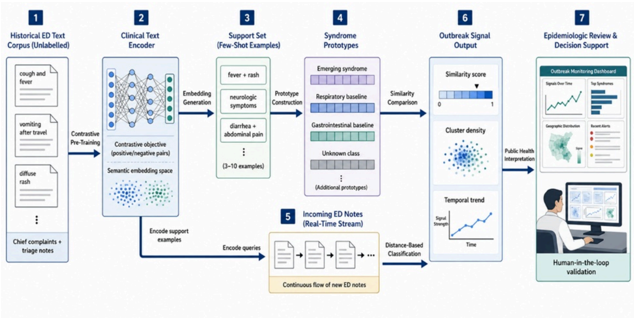

Figure 1. Linear architecture of contrastive learning with prototypical networks for few-shot detection of emerging infectious disease syndromes from emergency department text

Table 2 delineates the functional roles of each architectural component, clarifying how representation learning and prototype-based reasoning jointly enable few-shot outbreak detection.

Table 2. Functional roles of core components in contrastive-prototypical few-shot outbreak detection architecture

Component	**Input**	**Process**	**O**
Contrastive Encoder	Unlabelled ED text	Learns invariant semantic representations via contrastive objective	Der emb
Data Augmentation	Raw complaint text	Generates semantically consistent variants	Positive
Support Set	Few labeled examples	Encoded into embedding vectors	Support
Prototype Constructor	Support embeddings	Computes mean or centroid vectors	Syndrom
Query Encoder	Incoming ED notes	Transforms text into embeddings	Query e
Similarity Module	Query + prototype embeddings	Computes distance (cosine/Euclidean)	Simila
Outbreak Scoring Layer	Similarity outputs over time	Aggregates density and trends	Outbr
Human-in-the-Loop System	Model outputs + raw text	Epidemiologic validation and refinement	Up prototype

Core assumptions

The framework assumes that a health system or public health agency has access to a continuous stream of chief complaints and triage notes from emergency departments, either through local hospital data feeds or syndromic surveillance infrastructure [1]. It also assumes that, during the early phase of a possible outbreak, public health officials can provide a minimal support set, such as three to ten representative complaints from suspected or confirmed cases, even when the syndrome definition remains provisional [4, 10]. A further assumption is that the model can maintain an “unknown” or background class so that not every atypical complaint is forced into a newly defined outbreak category [3]. These assumptions make the framework suitable for conceptualizing early warning workflows, but they also imply that governance, data quality, and epidemiologic oversight are prerequisites for responsible deployment [2, 9].

Design principles

The first design principle is few-shot adaptability: the system should rapidly incorporate a small number of labelled examples without retraining an entire supervised classifier from scratch. The second principle is generalization to unseen symptom patterns, which motivates contrastive pre-training on broad historical complaint language and prototype-based classification for newly defined syndromes [12-14]. The third principle is privacy-conscious learning, because contrastive pre-training can use de-identified or minimally labelled historical notes without requiring manual chart review for every historical encounter [13, 20]. Finally, the framework treats automation as decision support rather than autonomous public health action, aligning machine-generated signals with human epidemiologic review and iterative refinement [1, 2].

Contrastive Pre-Training

Unlabelled data source

The contrastive pre-training stage uses a large historical database of emergency department chief complaints and triage notes collected through local emergency department systems, state-level syndromic surveillance programs, or national surveillance platforms. These data need not contain definitive disease labels, because the encoder learns from relationships among complaint texts rather than from outbreak annotations [12, 13]. Historical corpora can include respiratory, gastrointestinal, dermatologic, neurologic, injury-related, mental health, overdose, and nonspecific complaints, creating a broad semantic foundation for later few-shot adaptation [5-8]. The conceptual advantage is that abundant unlabelled text can be used before an outbreak occurs, while scarce labels are reserved for the moment when a new syndrome must be defined [10, 11].

Contrastive objective

The contrastive objective constructs different views of the same complaint or triage note as positive pairs and treats semantically different complaints as negative or less similar examples. For example, word dropout, synonym replacement, abbreviation normalization, or masking of nonessential tokens can produce alternate views of a complaint while preserving its clinical meaning, supporting invariance to superficial wording variation [12, 13]. The resulting encoder can be optimized with an InfoNCE-style objective in which embeddings of positive pairs are pulled together and embeddings of unrelated complaints are pushed apart, although the framework does not prescribe a single loss implementation [16, 19]. This objective is conceptually useful for syndromic surveillance because early outbreak signals may be expressed through varied triage language that shares latent symptom structure rather than identical wording [3, 5].

Encoder architecture

The encoder may be a transformer-based clinical language model fine-tuned for short emergency department text, a biomedical retrieval model adapted through contrastive learning, or a lightweight sentence encoder selected for public health deployment constraints. Clinical natural language processing research has shown that contrastive domain adaptation and biomedical pre-training can improve representation quality in low-resource settings, while automated coding and retrieval studies demonstrate how contrastive representations can organize medically meaningful text [12, 13, 19]. For resource-constrained surveillance environments, the framework allows smaller encoders if they preserve the core requirement of mapping semantically similar complaints into nearby vectors [20, 21]. The architecture should therefore be chosen not only for representational capacity but also for latency, maintainability, auditability, and compatibility with existing surveillance infrastructure [1, 2].

Prototypical Network for Few-Shot Classification

Support set construction

When an emerging syndrome is suspected, the support set consists of a small number of labelled chief complaints or triage-note snippets judged by public health officials or clinicians to represent the syndrome of concern. In a prototypical network, these examples are passed through the contrastively pre-trained encoder, and their embeddings are averaged to form a prototype for the emerging syndrome class [14]. Additional prototypes may represent comparator syndromes such as seasonal influenza-like illness, nonspecific viral illness, gastrointestinal illness, or rash illness, allowing the system to distinguish the emerging pattern from common background presentations [1, 5, 9]. The support set should be curated carefully because a prototype built from vague or inconsistent examples will encode ambiguity into the classifier, especially when only a handful of examples are available [10, 11].

Query set classification

For each incoming batch of chief complaints or triage notes, the encoder produces query embeddings that are compared with the emerging-syndrome prototype and any comparator prototypes. Distance can be defined through Euclidean distance, cosine distance, or another similarity metric consistent with the embedding geometry learned during contrastive pre-training [14, 22]. A query note is assigned to the closest prototype or given a calibrated similarity score, but the score should be interpreted as surveillance evidence rather than a clinical diagnosis [2]. This distance-based classification is conceptually attractive for emerging disease detection because it can operate when the new class has only a few examples and when the primary task is to identify a cluster of similar presentations for human review [3, 4].

Prototype refinement

Prototype refinement occurs as additional suspected or confirmed cases become available and public health officials update the support set. New embeddings can be incorporated into the prototype incrementally, while older prototypes can be retained to represent previous outbreaks, recurring seasonal syndromes, or ruled-out clusters [14, 23]. This memory-oriented design is important because outbreak detection systems must adapt to concept drift without erasing clinically meaningful representations of prior syndromes [17, 18]. Human-in-the-loop feedback can also correct false positives, remove misleading examples, and preserve interpretability by linking each prototype to the small set of complaints from which it was constructed [1, 2, 25].

Few-Shot Outbreak Detection Workflow

Real-time ingestion

In the proposed workflow, emergency department chief complaints and triage notes are ingested as a continuous stream, with updates occurring hourly, daily, or at another cadence determined by local surveillance capacity. Each incoming text is cleaned, normalized, and encoded through the contrastively pre-trained clinical-text encoder before being compared with active syndrome prototypes [1, 5]. This design aligns with near-real-time syndromic surveillance models in which triage text is available earlier than laboratory confirmation or finalized diagnosis, making it useful for early public health signal generation [5, 26]. The ingestion layer should preserve temporal and geographic metadata at an appropriate level of aggregation so that model outputs can be interpreted as population-level surveillance patterns rather than isolated clinical predictions [1, 2].

Outbreak score

The outbreak score is conceptualized as the proportion or weighted density of recent complaints that fall near the emerging-syndrome prototype above a predefined similarity threshold. This score can be compared with historical baselines for the same facility, region, weekday, season, or demographic stratum, recognizing that emergency department language and visit volume vary across time and setting [1, 17]. A rising score does not prove the existence of an outbreak, but it can indicate that current complaint patterns are becoming more similar to the few-shot syndrome definition than would be expected under baseline conditions [3, 4]. Because early outbreak detection is sensitive to false alarms, the score should be paired with epidemiologic review, external surveillance signals, and confirmatory case investigation before public health action is taken [2, 27].

Handling Emerging Diseases

Rapid response to unknown syndromes

For rapid response to an unknown syndrome, public health officials could provide a small number of early chief complaints or triage-note snippets from suspect cases, such as patients presenting with unusual respiratory, neurologic, gastrointestinal, dermatologic, or exposure-linked symptoms. The system would embed these examples, form a provisional emerging-syndrome prototype, and scan recent and incoming emergency department text for nearby cases within the same catchment area [14, 22]. This approach is consistent with the need for field-responsive machine learning methods that can operate before stable disease definitions and large training sets are available [4, 9]. It also complements broader artificial intelligence surveillance strategies by focusing on the earliest clinical language rather than waiting for administrative codes, laboratory reports, or confirmed diagnoses [2, 27].

Generalisation to unseen symptoms

The framework is most likely to generalize when an emerging syndrome shares latent structure with patterns already represented in historical emergency department text, such as fever with rash, respiratory distress after exposure, or gastrointestinal illness linked to travel. Contrastive learning supports this goal by organizing text according to semantic and clinical similarity, while prototypical networks allow new classes to be defined from a small number of labelled examples [12-14]. However, performance would be expected to degrade when a new disease occupies a truly novel semantic space, when patients describe symptoms in unfamiliar ways, or when triage notes omit critical exposure context [3, 10]. Therefore, the system should include an explicit uncertainty mechanism and an “unknown” region in the embedding space rather than forcing every unusual complaint into an existing prototype [3, 18].

Clinical and Public Health Integration

Alerting dashboard

An alerting dashboard would translate model outputs into interpretable public health information, including a time-series plot of the outbreak score, geographic aggregation of elevated syndrome probability, and facility-level trends suitable for investigation. The dashboard should not present prototype similarity as a diagnosis, because emergency department chief complaints may be vague, incomplete, or discordant with eventual discharge diagnoses [7, 28]. Instead, it should display the evidence trail linking a signal to representative complaints, recent temporal changes, and comparison with historical baseline patterns [1, 17]. Such a dashboard would extend existing syndromic surveillance practice by adding few-shot adaptability while preserving the operational logic of epidemiologic review and situational awareness [2, 27].

Human-in-the-loop

Human-in-the-loop oversight is essential because early outbreak signals require interpretation in light of local epidemiology, care-seeking behavior, testing availability, and clinical documentation practices. Epidemiologists can review high-probability clusters, confirm whether the underlying complaints reflect a plausible syndrome, and remove examples that are misleading, duplicated, miscoded, or driven by noninfectious events [1, 5]. Their feedback can update prototypes, refine thresholds, and distinguish true emerging patterns from transient documentation artifacts or unrelated increases in emergency department visits [14, 25]. This iterative design also helps address privacy and governance concerns by ensuring that automated text analysis remains embedded within accountable public health decision-making rather than functioning as an unsupervised alert generator [13, 20].

Evaluation Strategy

Simulated few-shot benchmarks

A conceptual evaluation strategy can begin with simulated few-shot benchmarks in which one known syndrome or disease-related category is held out during model development and later introduced with only a small support set. The framework could then be assessed by whether incoming notes from the held-out category are ranked closer to the new prototype than to comparator prototypes, using standard detection measures such as discrimination, precision, recall, and calibration without treating these simulations as evidence of real-world outbreak readiness [10, 14]. Such benchmarks should compare contrastively pre-trained encoders with noncontrastive or randomly initialized encoders to isolate the contribution of representation learning [12, 13, 19]. They should also compare prototype-based classification with alternative few-shot text methods, including Siamese, induction-network, and prompt-based approaches, because the most appropriate design may depend on the surveillance corpus and target syndrome [11, 22-24].

Temporal validation

Temporal validation should train or pre-train the encoder on earlier emergency department notes and evaluate signals on later periods that include known outbreaks, seasonal waves, or public health events. This design reflects the practical requirement that surveillance models generalize forward in time despite changes in language, pathogen prevalence, testing practices, and emergency department utilization [1, 17]. Evaluation should examine whether prototype-based scores rise before or near documented increases in relevant cases, while carefully avoiding claims that model alerts represent confirmed outbreak onset without external epidemiologic validation [2, 9]. Temporal validation is especially important for emerging infectious disease surveillance because retrospective random splits can exaggerate apparent generalization when training and test notes share local wording conventions or duplicated documentation patterns [10, 27].

Ablation experiments

Ablation experiments should examine how each conceptual component contributes to the framework, including contrastive pre-training, support-set size, prototype distance metric, comparator-class selection, and incremental prototype refinement. For example, one ablation could compare one, three, five, and ten labelled examples per emerging syndrome, while another could compare Euclidean distance, cosine distance, and learned similarity in the prototype space [14, 22]. Additional ablations could evaluate whether synonym-based augmentation, word dropout, abbreviation normalization, or domain-adaptive contrastive learning changes the structure of the clinical-text embedding space [12, 13, 16]. These analyses would not be intended to manufacture performance claims, but to clarify which design choices are most consequential before prospective deployment in a public health setting [18, 20, 26].

Limitations

Technical limitations

The framework depends on the historical distribution of emergency department notes used for pre-training, which means it may underperform when incoming language, facility mix, patient population, or triage documentation practices shift substantially. Concept drift is particularly relevant for infectious disease surveillance because symptom language can change during public awareness campaigns, testing shortages, media coverage, or the introduction of new clinical screening questions [1, 17]. The prototype mechanism also has an inherent ceiling when a disease produces symptoms or exposure contexts that are absent from the historical embedding space, even if the few-shot support examples are carefully selected [3, 10]. Finally, contrastive objectives can learn shortcuts from formatting, facility-specific wording, or duplicated templates unless preprocessing and evaluation are designed to detect these artifacts [12, 13].

Clinical limitations

Clinical limitations arise because chief complaints and triage notes are noisy, brief, and shaped by patient wording, nurse documentation habits, local triage protocols, and time pressure in emergency care. A complaint such as “fever,” “weakness,” or “flu-like symptoms” may correspond to many infectious and noninfectious conditions, making false positives unavoidable in any presyndromic surveillance framework [5, 7, 28]. Privacy concerns are also substantial because free text can contain names, locations, occupations, travel details, or other potentially identifying information, requiring de-identification, access controls, aggregation, and governance before model development [13, 20]. For these reasons, the framework should be treated as a public health signal-generation tool that requires human review, confirmatory evidence, and ethical oversight rather than as a stand-alone disease detection system [2, 25, 27].

Conclusion

Contrastive learning with prototypical networks provides a coherent conceptual framework for few-shot outbreak detection from emergency department chief complaints and triage notes. The framework uses historical unlabelled clinical text to learn a representation space and then defines emerging syndromes through a small number of representative examples. In this design, early outbreak detection becomes a problem of identifying similarity to provisional syndrome prototypes rather than waiting for large labelled datasets.

The key advantage of the framework is its ability to adapt quickly when a novel disease pattern emerges and only a handful of examples are available. It also leverages abundant unlabelled emergency department text, which is often available before confirmed case labels are established. By combining representation learning with prototype-based classification, the framework offers a practical bridge between presyndromic surveillance and few-shot artificial intelligence.

Future implementation should focus on public syndromic surveillance datasets, local emergency department feeds, and prospective collaboration with public health agencies. Validation should be staged carefully, beginning with retrospective simulations and progressing toward monitored prospective studies under epidemiologic oversight. The ultimate goal is a responsible early warning capability that supports timely investigation of unknown biothreats while preserving clinical caution, privacy, and public health accountability.

Acknowledgements

None

Conflict of interest

None

Ethics statement

None

Financial support

None

Received: 27 Sep 2025 Revised: 05 Dec 2025 Accepted: 15 Jan 2026
Published online: 20 July 2026

References

Burkom H, Loschen W, Wojcik R, Holtry R, Punjabi M, Siwek M, et al. Electronic surveillance system for the early notification of community-based epidemics (ESSENCE): overview, components, and public health applications. JMIR Public Health Surveill. 2021;7(6):e26303. https://doi.org/10.2196/26303.

Zeng D, Cao Z, Neill DB. Artificial intelligence-enabled public health surveillance—from local detection to global epidemic monitoring and control. In: Artificial Intelligence in Medicine. Cambridge (MA): Academic Press; 2021. p. 437-53. https://doi.org/10.1016/B978-0-12-821259-2.00019-3.

Feng R, Hu Q, Jiang Y. Unknown disease outbreaks detection: a pilot study on feature-based knowledge representation and reasoning model. Front Public Health. 2021;9:683855. https://doi.org/10.3389/fpubh.2021.683855.

Chiu HY, Hwang CK, Chen SY, Shih FY, Han HC, King CC, et al. Machine learning for emerging infectious disease field responses. Sci Rep. 2022;12(1):328. https://doi.org/10.1038/s41598-021-04117-3.

Khademi S, Palmer C, Javed M, Dimaguila GL, Clothier H, Buttery J, et al. Near real-time syndromic surveillance of emergency department triage texts using natural language processing: case study in febrile convulsion detection. JMIR AI. 2024;3(1):e54449. https://doi.org/10.2196/54449.

Chenais G, Gil-Jardiné C, Touchais H, Fernandez MA, Contrand B, Tellier E, et al. Deep learning transformer models for building a comprehensive and real-time trauma observatory: development and validation study. JMIR AI. 2023;2(1):e40843. https://doi.org/10.2196/40843.

Goldman-Mellor S, Jia Y, Kwan K, Rutledge J. Syndromic surveillance of mental and substance use disorders: a validation study using emergency department chief complaints. Psychiatr Serv. 2018;69(1):55-60. https://doi.org/10.1176/appi.ps.201700234.

Vivolo-Kantor AM, Smith H IV, Scholl L. Differences and similarities between emergency department syndromic surveillance and hospital discharge data for nonfatal drug overdose. Ann Epidemiol. 2021;62:43-50. https://doi.org/10.1016/j.annepidem.2021.07.006.

Shen Y, Liu Y, Krafft T, Wang Q. Progress and challenges in infectious disease surveillance and early warning. Medicine (Baltimore). 2025;2(1):100071.

Ge Y, Guo Y, Das S, Al-Garadi MA, Sarker A. Few-shot learning for medical text: a review of advances, trends, and opportunities. J Biomed Inform. 2023;144:104458. https://doi.org/10.1016/j.jbi.2023.104458.

Oniani D, Chandrasekar P, Sivarajkumar S, Wang Y. Few-shot learning for clinical natural language processing using siamese neural networks: algorithm development and validation study. JMIR AI. 2023;2:e4429. https://doi.org/10.2196/4429.

Jin Q, Kim W, Chen Q, Comeau DC, Yeganova L, Wilbur WJ, et al. MedCPT: contrastive pre-trained transformers with largescale PubMed search logs for zero-shot biomedical information retrieval. Bioinformatics. 2023;39(11):btad651.

Francis S, Crema Garcia F, Uma K, Mestdagh W, De Moor B, Moens MF. Clinical entity-aware domain adaptation in low resource setting for inflammatory bowel disease. Front Artif Intell. 2025;7:1450477. https://doi.org/10.3389/frai.2024.1450477.

Snell J, Swersky K, Zemel R. Prototypical networks for few-shot learning. Adv Neural Inf Process Syst. 2017;30:4077-87.

Richter-Pechanski P, Wiesenbach P, Schwab DM, Kiriakou C, He M, Geis NA, et al. Few-shot and prompt training for text classification in German doctor’s letters. In: Caring is Sharing – Exploiting the Value in Data for Health and Innovation. Amsterdam: IOS Press; 2023. p. 819-20. https://doi.org/10.3233/SHTI230280.

Tariq A, Goddard K, Elugunti P, Piorkowski K, Staal J, Viramontes A, et al. Contrastive diagnostic embedding (CDE) model for automated coding: a case study using emergency department encounters.

Gumustop S, Gallo-Bernal S, McPeake F, Briggs D, Gee MS, Pianykh OS. Predicting health crises from early warning signs in patient medical records. Sci Rep. 2022;12(1):19267. https://doi.org/10.1038/s41598-022-23957-4.

Li C, Denison T, Zhu T. A survey of few-shot learning for biomedical time series. IEEE Rev Biomed Eng. 2024;18:192-210. https://doi.org/10.1109/RBME.2024.3490503.

Niu K, Wu Y, Li Y, Li M. Retrieve and rerank for automated ICD coding via contrastive learning. J Biomed Inform. 2023;143:104396. https://doi.org/10.1016/j.jbi.2023.104396.

Mallya P, Henao R, Hong C, Wojdyla D, Schibler T, Manchanda V, et al. Automated data harmonization in clinical research: natural language processing approach. JMIR Form Res. 2025;9:e75608. https://doi.org/10.2196/75608.

Zhao Z, Lu W, Peng X, Xing L, Zhang W, Zheng C. Automated ICD coding via contrastive learning with back-reference and synonym knowledge for smart self-diagnosis applications. IEEE Trans Consum Electron. 2024;70(3):6042-53. https://doi.org/10.1109/TCE.2024.3422330.

Sun S, Sun Q, Zhou K, Lv T. Hierarchical attention prototypical networks for few-shot text classification. In: Proc 2019 Conf Empirical Methods Natural Language Processing and 9th Int Joint Conf Natural Language Processing (EMNLP-IJCNLP). 2019. p. 476-85.

Geng R, Li B, Li Y, Zhu X, Jian P, Sun J. Induction networks for few-shot text classification. In: Proc 2019 Conf Empirical Methods Natural Language Processing and 9th Int Joint Conf Natural Language Processing (EMNLP-IJCNLP). 2019. p. 3904-13.

Gao T, Fisch A, Chen D. Making pre-trained language models better few-shot learners. In: Proc 59th Annu Meet Assoc Comput Linguist and 11th Int Joint Conf Natural Language Processing. 2021. p. 3816-30.

Omar M, Brin D, Glicksberg B, Klang E. Utilizing natural language processing and large language models in the diagnosis and prediction of infectious diseases: a systematic review. Am J Infect Control. 2024;52(9):992-1001. https://doi.org/10.1016/j.ajic.2024.03.012.

Sharrab YO, Attar H, Eljinini MA, Al-Omary Y, Al-Momani WA. Advancements in speech recognition: a systematic review of deep learning transformer models, trends, innovations, and future directions. IEEE Access. 2025;13:46925-40. https://doi.org/10.1109/ACCESS.2025.3555305.

Villanueva-Miranda I, Xiao G, Xie Y. Artificial intelligence in early warning systems for infectious disease surveillance: a systematic review. Front Public Health. 2025;13:1609615. https://doi.org/10.3389/fpubh.2025.1609615.

Hsia RY, Niedzwiecki M. Avoidable emergency department visits: a starting point. Int J Qual Health Care. 2017;29(5):642-5.