The portability of natural language processing methods to detect suicidality from unstructured clinical text in US and UK electronic health records

Marika Cusick (✉ mac2364@med.cornell.edu)

Weill Cornell Medicine

Sumithra Velupillai

King's College London

Johnny Downs

King's College London

Thomas Campion

Weill Cornell Medicine

Rina Dutta

King's College London

Jyotishman Pathak

Weill Cornell Medicine

Research Article

Keywords: suicide, natural language processing (NLP), electronic health records (EHRs), algorithm

Posted Date: February 10th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-155430/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

Abstract

In the global effort to prevent death by suicide, many academic medical institutions are implementing natural language processing (NLP) approaches to detect suicidality from unstructured clinical text in electronic health records (EHRs), with the hope of targeting timely, preventative interventions to individuals most at risk of suicide. Despite the international need, the development of these NLP approaches in EHRs has been largely local and not shared across healthcare systems. In this study, we developed a process to share NLP approaches that were individually developed at King's College London (KCL), UK and Weill Cornell Medicine (WCM), US - two academic medical centers based in different countries with vastly different healthcare systems. After a successful technical porting of the NLP approaches, our quantitative evaluation determined that independently developed NLP approaches can detect suicidality at another healthcare organization with a different EHR system, clinical documentation processes, and culture, yet do not achieve the same level of success as at the institution where the NLP algorithm was developed (KCL approach: F1-score 0.85 vs. 0.68, WCM approach: F1-score 0.87 vs. 0.72). Shared use of these NLP approaches is a critical step forward towards improving data-driven algorithms for early suicide risk identification and timely prevention.

Introduction

Suicide is a major public health concern across the world. The World Health Organization (WHO) reports that nearly 800,000 people die per year by suicide, accounting for nearly 1.4% of all deaths worldwide (1). To prevent suicide, many healthcare institutions have attempted to predict suicidal deaths, but this has not been particularly successful (2). This can be attributed to the fact that within most populations, suicidal phenomena are relatively rare events, especially death from suicide, making it difficult to identify at-risk individuals. Furthermore, current risk detection relies primarily on self-reported questionnaires, which may not always be accurate due to mental health stigma or provider bias (3,4). To improve suicide risk detection, researchers need to be collaborative, by aggregating or linking data from other institutions or when governance limits the pooling of protected health information (PHI), performing meta-analytic evaluations.

Electronic health records (EHRs) are distributed documentation systems that differ from paper records or charts. Making use of EHR data, particularly unstructured clinical notes, offers a novel avenue for suicide risk modeling. EHRs can bring together very large samples for researchers to scrutinize and provide real-world insights into a patient's mental state (5,6). This is particularly true for suicidality, a common precursor to death by suicide (7), as many providers document suicidality in notes rather than as a structured data element in EHR systems (8,9). Based on this insight, several investigators and health systems are now working to develop natural language processing (NLP) approaches to detect suicidality from unstructured clinical notes in the EHRs(10).

As health systems begin to achieve siloed success in detecting clinical conditions from EHR data and clinical notes, a critical next step forward is sharing and implementing these approaches broadly across other organizations, as most do not have the infrastructure and resources to develop and build these approaches de novo. In addition, wide adoption of existing approaches will potentially minimize duplicative effort in the development of NLP algorithms. While there has been a considerable amount of work on the portability of phenotype algorithms to detect physical conditions, such as rheumatoid arthritis (11) and heart disease (12), to our knowledge, there has been little work in the sharing of NLP algorithms to detect more complex clinical phenomena, such as suicidality. The process is further complicated when sharing across institutions located internationally where the practice of clinical psychiatry, coding and documentation varies greatly. As described by Silverman et al. *"There is no internationally agreed-upon set of terms, definitions, or classifications for the range of thoughts, communications, and behaviors that indicate 'suicidality'. The lack of agreement makes it difficult, if not impossible, to compare and contrast different suicide related research studies, clinical reports, or epidemiological surveys, or to make generalizations or extrapolations (13)."*

Psychiatrists in the US and UK document suicide-related issues according to the phenomenology they were taught in medical school and during clinical training, and inevitably “learners are likely to emulate their supervisors’ EHR use” (14). Each clinician seeks to follow best practice national guidelines, such as the American Psychiatric Association (APA) guidelines in USA (15) and Department of Health Best Practice in Managing Risk guidance in the UK (16). Due to the challenge of disparate ways in describing and documenting suicidality (17), the development of algorithms to detect suicidal behavior appear largely up to the interpretation and definition of the healthcare organization. International collaborations using EHRs could offer a parsimonious method to enhance suicide prevention research. However, currently, there is no prior work that can aid international collaborations for sharing and evaluating NLP algorithms to detect suicidality from clinical notes in EHR systems, nor are there empirical findings on the change in detection accuracy when NLP algorithms developed in one institution are implemented in another organization.

To address this, we set out to evaluate how independently developed NLP approaches that detect suicidality translate across differing EHR platforms and classification objectives. In this study, we conducted a portability experiment using NLP approaches and datasets developed independently at two separate academic medical centers in two different countries (US and the UK) with vastly different healthcare systems. Results from our experiment can inform other institutions on how to share NLP algorithms that detect suicidality, improving international collaboration in suicide prevention efforts.

Materials And Methods

Data Source

We used NLP algorithms and EHR data from two large, academic healthcare institutions: South London and Maudsley Foundation NHS Trust, based in South London, UK and Weill Cornell Medicine based in New York City, USA.

The KCL team used data extracted from electronic clinical records from the South London and Maudsley (SLaM) Foundation NHS Trust is one of the UK’s largest and oldest mental health trusts, providing a wide range of inpatient, outpatient and community-based mental health services. The main catchment area is four boroughs in south London (Croydon, Lambeth, Lewisham and Southwark), serving a local population of around 1.3 million people with more than 4,500 employees. Jointly with the Institute of Psychiatry, Psychology & Neuroscience, KCL, SLaM hosts the NIHR Maudsley Biomedical Research Centre (BRC). The Clinical Record Interactive Search (CRIS) database is a resource developed at the Maudsley BRC, making de-identified EHR records available for secondary research use under an extensive governance model (18). The de-identified CRIS database has received ethical approval for secondary analysis: Oxford REC C, reference 18/SC/0372. The data is used in an anonymized and data-secure format under strict governance procedures. All experiments were performed in accordance with guidelines and regulations. The data were used in an entirely anonymized and data-secure format, and patients have the choice to opt-out of their anonymized data being used, and therefore, under UK law, does not require informed consent from patients whose data are represented here.

Weill Cornell Medicine (WCM) is an academic medical center in New York City with 1,600 physicians, over 50 locations throughout the New York City metropolitan region, and 3 million annual patient encounters. WCM has an affiliation with NewYork-Presbyterian Hospital (NYPH), which serves as the primary emergency and inpatient setting for WCM patients. While clinical care is documented in different EHR systems at WCM and NYPH, the Architecture for Research Computing in Health (ARCH) database facilitates the secondary use of EHR data for research by capturing novel research measures and integrating data from multiple EHR systems (19). This study was approved by the WCM Institutional Review Board (IRB). All experiments were performed in accordance with guidelines and regulations. This study was approved for a full waiver of informed consent, as it involves no more than minimal risk to the subjects.

Study Population

The KCL test data consists of 4,911 documents (progress notes, assessments and correspondence notes) from a random sub-cohort of 500 adolescents (13–18 years old) diagnosed with autism spectrum disorders (ICD-10: F84.0, F84.1, F84.5, F84.9) derived from a previously characterized clinical sample. These notes were recorded between January of 2013 and December 2013 and are further described in the Downs et. al paper (20). The clinical documents were annotated for mentions of suicide-related information by trainee clinical psychologists, under senior clinician supervision. As described in Downs et al. (20), suicidality was defined as “as either the reporting of the intention to engage in a potentially lethal act towards oneself, or undertaking such acts themselves”. Each note contained at least one instance of a suicide-related term (e.g. ‘suicid*’, ‘kill him/her/themself’, ‘want to die’), that were then labeled as positive, negated, or uncertain. From the individual annotations, each *document* has then been further labeled as either affirmed/relevant for suicidality (True) or negated (False). There are in total 3,069 documents labeled as True (62.5%) and 1,842 as False (37.5%).

The WCM test data set consists of 837 suicide-related notes for 30 patients selected from a pre-established depression cohort, defined as any patient diagnosed with depression or prescribed an antidepressant. Of the 30 patients, 10 patients had an encounter diagnosis of suicidal ideation (V62.84 (ICD9), R45.85 (ICD10)) in their medical history. The remaining 20 patients were considered to be potentially suicidal, as they had at least ten notes with a key suicidal phrase (“suicidal”, “suicide”, “SI”, or “suicidality”) A large majority of these notes (83%) were documented in the outpatient office setting at psychiatry and internal medicine departments rather than the inpatient hospital setting (13%) between January of 2006 and December of 2019. The dataset was annotated for current suicidality, defined as patients discussing, thinking about or planning for suicide during the documented encounter, by two investigators at WCM with established annotation guidelines. Each note contained at least one instance of a suicidal mention (“suicidal”, “suicide”, “SI”, or “suicidality”), that were then labeled as positive or negative for current suicidality. Based on these annotations, 134 (16.0%) of the documents were classified as either affirmed/relevant for suicidality (True) and 703 (84.0%) documents were classified as negated (False).

NLP Approaches

Two symbolic rule-based NLP approaches were applied: KCL-neg(21) and WCM-si. They were both developed on the basis of the NegEx algorithm(22), an approach to identify negated findings in unstructured clinical text. This algorithm relies on two lexicons: one defining target concepts (e.g. *suicidal*) and the other defining modifiers (e.g. *not*).

The KCL-neg approach was designed to detect any mention of suicidality, regardless of temporality (i.e. current or historical). The target lexicons of the KCL-neg approach included both direct and indirect mentions of suicidality. Direct mentions are any word with the regular expression basis of “suicid”, which includes “suicidal,” “suicidality,” and “suicide.” Indirect mentions include expressions such as “take (his|her|their) life”, “wish to die”, and “life not worth living.” The WCM-si approach was designed for detecting *current* suicidality. The target lexicons for the WCM-si approach were the four key suicidal ideation terms—“suicidal”, “suicide”, “suicidality”, “si”—used to select the EHR note cohort.

The KCL team implemented different sets of modifiers to study the impact on algorithm performance when using previously published lists of modifier terms compared to adapted lists for new use cases. Here, we use two of these modifier sets: anySI-1 and anySI-2. The WCM modifier lexicon set, henceforth called currentSI, included modifiers to negate current suicidality. We categorized both the KCL and WCM modifiers into four different categories: negated, historical, conditional, and unrelated. Examples of each of these modifiers are provided in Table 1. The entire set of modifiers are available on our respective GitHub websites.

Table 1
Modifier categories

Modifier category	AnySI-1	AnySI-2	CurrentSI	Examples
Negated	Yes	Yes	Yes	Denied, negative for, never, not, no
Historical	No	No	Yes	History, previous, attempted, YEAR, DATE
Conditional	No	Yes	Yes	Lifeline, emergency number, even if, return if
Unrelated	No	Yes	Yes	Brother, sister, family history, husband, wife

Portability Experiment

To initiate the portability experiment, investigators from WCM and KCL held several meetings to present and discuss each of the NLP algorithms from both a clinical and technical standpoint. First, it was critical to share the algorithm's main clinical objectives, key details about the study population and site, and guidelines behind the manual annotation of the test set. Next, we shared the technical details of the algorithms and developed a plan of action to transfer compatible code and execution guidelines. Upon realizing that versioning issues may arise, we created a virtual environment with a consistent python installation and package versions. Finally, we confirmed the format of each institution's data set, including the text format. Code for the algorithms are available on our respective GitHub^{1,2} websites.

Evaluation

After executing the two NLP algorithms on the manually annotated datasets, we evaluated the results using both quantitative and qualitative methods. First, to assess portability quantitatively, we compared the algorithms' results using traditional intrinsic evaluation metrics, such as accuracy, precision, recall, and F1-scores. Second, to assess portability with an eye towards the underlying details of each NLP approach, we conducted a thorough qualitative manual error analysis to characterize the most common misclassification scenarios. Based on the specification of each of the approaches, we can deem some classification errors to be expected and better analyze the effectiveness of the approach.

Results

During a 5 month span, the two teams met bi-monthly to outline the technical requirements necessary to port our algorithms to unseen datasets at the other institution. Once all necessary information and details were made available on GitHub, each team successfully executed the ported algorithm on their own test data set with little difficulty. In the event of questions, we kept our communication channels open through email and kept response times within a couple of days.

NLP Approach Results

As demonstrated in Table 2, the ported algorithms did detect suicidality, yet they did not replicate the same level of success in detecting suicidality as at the institution where the algorithm was originally developed. Using the two modifier sets, the KCL-neg approach achieved a maximum macro-average f1-score of 0.85 on its own KCL test dataset, but only resulted in a maximum score of 0.68 on the WCM dataset. Similarly, the WCM-si approach resulted in a macro-average f1-score of 0.87 on their own test dataset, but only achieved a maximum score of 0.72 on the KCL dataset.

Table 2

Results: Precision (P), Recall (R) and F1-score (F1) for affirmed (True) and negated (False) suicide-related instances, applied on two datasets (WCM and KCL) and using two NLP approaches (WCM-si and KCL-neg) with different modifier lexicons.

Data	NLP Approach	Modifier Set	P (T)	R (T)	F1 (T)	P (F)	R (F)	F1 (F)	F1 (avg)
WCM	KCL-neg	AnySI-1	0.39	0.80	0.53	0.95	0.74	0.83	0.68
WCM	KCL-neg	AnySI-2	0.26	0.77	0.38	0.92	0.54	0.68	0.53
WCM	WCM-si	CurrentSI	0.72	0.87	0.79	0.97	0.93	0.95	0.87
KCL	WCM-si	CurrentSI	0.87	0.64	0.74	0.58	0.85	0.69	0.72
KCL	KCL-neg	AnySI-1	0.74	0.86	0.79	0.67	0.5	0.57	0.68
KCL	KCL-neg	AnySI-2	0.87	0.91	0.89	0.84	0.78	0.81	0.85

We observed the same phenomenon on each of the performance metrics (e.g. precision on affirmed instances), as neither of the approaches outperformed the success of the “home algorithms”. While the WCM-si approach was able to achieve higher precision (0.87) than the KCL-neg approach using AnySI-1 modifiers (0.74) for positive instances of suicidality on the KCL data, the KCL-neg approach using AnySI-2 modifiers yielded a similar precision (0.87).

Although the KCL-neg approach using the AnySI-2 modifier set had a better overall performance (macro-average f1-score: 0.85 vs. 0.68) in comparison to the AnySI-1 modifier set on their own KCL data set, the opposite was observed on the WCM data set (macro-average f1-score: 0.53 with AnySI-1 vs. 0.68 with AnySI-2).

Error Analysis

KCL-neg with AnySI-1 modifiers on WCM data

Of the two KCL-neg modifier sets, the AnySI-1 modifier lexicon set proved the most successful (macro-average f1-score of 0.68 vs. 0.53). Thus, using this set, we conducted our qualitative error analysis to determine the reasons for misclassification.

Out of the 206 total errors, there were 177 and 29 false positive and false negative errors, respectively. Of the 29 false negative errors, the majority (69%) can be attributed to the KCL-neg algorithm’s target lexicons not including the term “si.” For this reason, the KCL algorithm was not programmed to detect this type of mention from the clinical notes and automatically classified the note as negative. If we removed these expected instances from the test set, recall for positive mentions of suicidal ideation increased from 0.80 to 0.93.

The 177 false positive errors (Table 3) can be grouped into the following scenarios: missing a negation modifier, non-patient person reference, structured references, conditional mentions, and historical mentions. The table below displays the number of cases in each scenario and several examples. Because the KCL algorithm was not configured to negate historical suicidality, the 45 false positive errors classified to this scenario were to be expected. Taking into account these expected errors, precision for positive mentions increased from 0.39 to 0.46.

Table 3. False Positive Error Scenarios KCL-neg on WCM

Scenario	Cases	Example
Missing a negation modifier	16	“Negative for suicidal ideation”
Non-patient references	23	“Brother (committed suicide at age 27)”
Structured mentions	63	“Suicidal ideation: denies”
Conditional mentions	30	“Children of untreated depressed mothers are also more prone to suicidal behavior”
Historical suicidality	45	“Has another suicide attempt on MM/DD”
Total	177	

WCM-si with currentSI modifiers on KCL data

A similar error analysis was performed on the KCL data with the WCM-si approach. Out of the 1,398 total errors, there were 279 and 1,119 false positive and false negative errors, respectively. Of the 1,119 false negative errors, 529 (47%) were attributed to the WCM-si algorithm not including target terms such as “kill him/herself”, “end his/her life”. An analysis of 250 of the remaining 590 errors revealed that 58 (23%) related to references to the past, which the WCM-si approach was designed to exclude given that its primary focus is on “current” suicidality; 69 (28%) related to complex, long documents where there were several references to suicidal behavior but the algorithm only picked up those related to “suicid*”; 40 (16%) related to missing triggers or erroneous trigger scopes; and 83 (33%) related to other issues, including errors in the gold standard annotations. As the WCM-si approach was not configured to detect indirect mentions and purposefully negated historical mentions, we considered 779 of the false negative errors to be expected, changing the recall for positive mentions of suicidal ideation from 0.64 to 0.83.

Table 4. False Positive Error Scenarios WCM-si on KCL data

Scenario	Cases	Example
Negation and positive mention in same note (often references to the past)	149	“X said they had suicidal thoughts ... although self harm may not be true suicidal intent”
Missing a negation modifier	29	“Nil suicidal ideation”
Structured mentions	35	“Suicide risk: low”
Routine checks, hypotheticals, unrelated	34	“Asked about suicidal thoughts”
Other	32	“Refused to answer questions around suicidal ideation”
Total	279	

For the 279 false positive errors, similar scenarios as for the KCL-neg approach on the WCM dataset were observed (Table 4), with some notable differences: historical mentions were not considered false positives in this case; but missing negation modifiers were observed (e.g. “nil”), as well as structured mentions in forms. Additionally, some false positives were related to documents where there were both negative and positive mentions, where these were classified as negated in the KCL gold standard but the WCM-si approach classified them as positive. Other examples include mentions of routine checks by the clinician, hypotheticals, and mentions related to someone other than the patient. We did not consider any of these errors to be expected.

Discussion

Our study showed that NLP approaches developed to detect complex clinical constructs, such as suicidality, can be successfully ported and shared across institutions, with proper emphasis on clear and effective communication. First, our informative discussions on the technical compatibility of our NLP algorithms (now hosted on GitHub¹) made porting the algorithms a seamless experience. Second, once the approaches were applied to the test data sets, a deeper understanding of the underlying details allowed for a more informative evaluation and identification of potential areas for improvement. While the traditional quantitative measures of portability, such as accuracy and f1-scores, indicated that ported NLP approaches failed to achieve the same level of success as at their home institution, there were inherent reasons for this. Our qualitative error analysis, which took into account each of the approaches slight differences in objectives, indicated that many of the errors were to be expected based on the institution’s guidelines for defining and annotating suicidality. In fact, if expected errors were taken out of consideration, we found the ported algorithm’s results to improve significantly. For the KCL-neg on the WCM data, 65 (32%) of the 206 errors were to be expected, changing the overall F1 score of the algorithm from 0.75 to 0.83. Similarly, of the 1,398 errors that the WCM-si approach made on the KCL data, 587 (42%) of the errors were to be expected based on the algorithm’s configuration and objective, changing the overall F1 score from 0.72 to 0.83.

In addition to a more meaningful evaluation, understanding the underlying details of the NLP approaches may inform how to develop a more generalizable approach. Similar to prior work, we confirmed in this study that suicidality is interpreted and defined differently across institutions and healthcare systems. Among our two institutions, the biggest differences in the definition of suicidality include decisions on temporality and inclusion of indirect suicidality-related terms. While these differences exist, our approaches did also have a number of similarities, including terms in the target lexicons and modifier terms within negation, conditional, and non-experiencer (unrelated categories). This commonality suggests that with further experimentation and collaborations, we can continue to improve on the detection of this complex clinical condition, by developing a portable and generalizable approach.

Both of our institutions have also implemented more novel, state-of-the-art methods for the detection of suicidality, such as a text classification convolutional neural network (CNN) (23), and support vector machines (SVM) (10). However, we decided to experiment with the more basic lexicon and rule-based NLP algorithms for two reasons: ease of portability and human interpretability. With the eventual goal of building generalizable and portable NLP approaches to detect suicidality, we determined that rule-based approaches could be more widely implemented across institutions as they require significantly less technical expertise, computational power, and other resources. In addition, many studies have concluded that simple rule-based approaches achieve similar levels of success to these novel implementations (24–26). A second advantage of rule-based NLP algorithms is their human interpretability. While effective, state of the art machine learning methods are challenging to interpret (27,28) and even more difficult to adjust. In the case of rule-based NLP algorithms, an external institution would have the ability to determine the most common sources of error and make changes to the approach, as it sees fit to the organization’s use case.

There are several limitations to this study. First, because the suicidality NLP approaches and test data sets were developed completely independently of each other, we recognize that they may not be directly comparable. Although a study with a single objective and manual annotation guideline may have been a more robust approach to evaluating portability, we view this to be a real-world experiment and thus replicable by other institutions. Second, our method of data extraction differed amongst the two institutions. While the WCM team only extracted notes with explicit mentions of suicidality, such as “suicidal”, the KCL team extracted notes with a wider search criteria, which included both explicit and “implicit” terms, such as “wish to die.” Third, incorporation of diagnostic codes, such as ICD-9/10 may improve our algorithms’ results. However, the aim of this study was to evaluate and assess portability of our suicidality NLP approaches, rather than prediction results. We hope that results from this study will make NLP algorithms more widely accessible and bolster results of existing suicide prediction algorithms currently reliant on structured EHR data such as diagnosis codes. Finally, because we did not port and evaluate our more advanced machine-learning based algorithms, we cannot say for certain whether these models would have achieved higher success than our lexicon-based NLP approaches. However, based on our goal of developing NLP approaches that can be widely used, we believe this is out of scope for this study and a future area of research.

In an effort to understand patients who are at risk for death by suicide, routinely collected data from healthcare institutions, such as EHRs, can be a valuable resource at scale. Information about suicidal risk behavior is predominantly documented in free text notes in EHRs, leading to an increase in the development of NLP approaches to detect suicidality in unstructured clinical text. However, due to the clinical complexity of suicidality, the lack of consensus on how to define this condition, differences in how clinical assessments are documented in EHRs, and the various ways the task can be modeled for information extraction, the development of relevant NLP approaches have largely been local to each institution’s definitions and interpretations. Thus, these NLP approaches are much less generalizable and portable in comparison to phenotype algorithms for well-defined clinical conditions, such as rheumatoid arthritis(11). However, with a well-defined process to understand the underlying details of an approach, institutions can be well-equipped to make use of an external approach, allowing a larger number of institutions to participate in suicide-related research. This is a critical step forward in learning how to develop more robust, portable, and generalizable NLP methods that can be applied to any clinical text, regardless of the origin EHR system.

Declarations

ACKNOWLEDGEMENTS

Research by MC, TC, and JP was funded in part by NIH grants R01MH105384, R01MH119177 and R01MH121922. JD is supported by NIHR Clinician Science Fellowship award (CS-2018-18-ST2-014) and has received support from a Medical Research Council (MRC) Clinical Research Training Fellowship (MR/L017105/1) and Psychiatry Research Trust Peggy Pollak Research Fellowship in Developmental Psychiatry. This paper by SV, JD, and RD represents independent research part funded by the National Institute for Health Research (NIHR) Biomedical Research Centre at South London and Maudsley NHS Foundation Trust and King’s College London. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health and Social Care. SV and JD are part-funded by the Medical Research Council (MRC) Mental Health Data Pathfinder Award to King’s College London. RD is funded by a Clinician Scientist Fellowship (RE11682: research project e-HOST-IT) from the Health Foundation in partnership with the Academy of Medical Sciences.

AUTHOR CONTRIBUTIONS

SV and MC contributed to the main analysis of this paper. They developed the NLP approaches, conducted the portability experiment, and evaluated the results of the experiment. They contributed heavily to the methods and discussion section. JD and RD provided clinical expertise on suicidality and contributed to the introduction and discussion section. THC

supervised the development of the NLP approach on the WCM side and provided expertise in clinical informatics. JP supervised the entire project and contributed revisions to all sections of the paper.

ADDITIONAL INFORMATION

Competing Interests

RD and SV declare previous research funding received from Janssen. The remainder of the authors (MC, THC, JD, JP) do not declare any competing interests.

Data Availability

The datasets generated during and/or analyzed during the current study are not publicly available because they include personally identifiable data, which each investigator on this study obtained access through our institutions' review boards. CRIS data is made available to researchers with appropriate credentials (provided by the South London and Maudsley NHS Trust) working on approved projects. Projects are approved by a CRIS Oversight Committee, a body set up by and reporting to the South London and Maudsley Caldicott Guardian. On request, and after appropriate credentials have been obtained as well as arrangements with the lead of the respective CRIS project, data presented in this study can be

viewed within the secure system firewall. In a similar fashion, WCM data is made available to investigators who are a part of WCM IRB approved research projects. Researchers are given credentials and access instructions to a subset of EHR data as described in the project's study protocol.

As mentioned in our manuscript, code for our NLP algorithms is available on our GitHub repositories at the following links: https://github.com/KCL-Health-NLP/camhs_pycontext_adaptation, https://github.com/wcmc-research-informatics/SI_Ideation.

Ethics Declarations

This study was approved by the WCM Institutional Review Board (IRB). The de-identified CRIS database has received ethical approval for secondary analysis: Oxford REC C, reference 18/SC/0372. The data is used in an anonymized and data-secure format under strict governance procedures.

References

1. WHO. WHO | Suicide data. *WHO Mental Health* (2019). at <https://www.who.int/mental_health/prevention/suicide/suicideprevent/en/>
2. Kessler, R. C., Bossarte, R. M., Luedtke, A., Zaslavsky, A. M. & Zubizarreta, J. R. Suicide prediction models: a critical review of recent research with recommendations for the way forward. *Mol. Psychiatry* **25**, 168–179 (2020).
3. Thobaben, M. Suicide myths and health care provider bias. *Home Care Provid.* **2**, 109–111 (1997).
4. Blocker, G. M. & Miller, J. A. Unintended consequences: stigma and suicide prevention efforts. *Mil. Med.* **178**, 473 (2013).
5. Walsh, C. G., Ribeiro, J. D. & Franklin, J. C. Predicting Risk of Suicide Attempts Over Time Through Machine Learning. *Clin. Psychol. Sci.* **5**, 457–469 (2017).
6. McCoy, T. H., Jr, Castro, V. M., Roberson, A. M., Snapper, L. A. & Perlis, R. H. Improving Prediction of Suicide and Accidental Death After Discharge From General Hospitals With Natural Language Processing. *JAMA Psychiatry* **73**, 1064–1071 (2016).

7. Brown, G. K., Beck, A. T., Steer, R. A. & Grisham, J. R. Risk factors for suicide in psychiatric outpatients: a 20-year prospective study. *J. Consult. Clin. Psychol.* **68**, 371–377 (2000).
8. Haerian, K., Salmasian, H. & Friedman, C. Methods for identifying suicide or suicidal ideation in EHRs. *AMIA Annu. Symp. Proc.* **2012**, 1244–1253 (2012).
9. Anderson HD, Pace WD, Brandt E, Nielsen RD, Allen RR, Libby AM, West DR, Valuck RJ}. Monitoring Suicidal Patients in Primary Care Using Electronic Health Records. *Journal of the American Board of Family Medicine* **28**, (2015).
10. Fernandes, A. C. *et al.* Identifying Suicide Ideation and Suicidal Attempts in a Psychiatric Clinical Research Database using Natural Language Processing. *Sci. Rep.* **8**, 7426 (2018).
11. Carroll, R. J. *et al.* Portability of an algorithm to identify rheumatoid arthritis in electronic health records. *J. Am. Med. Inform. Assoc.* **19**, e162–9 (2012).
12. Kashyap, M. *et al.* Development and validation of phenotype classifiers across multiple sites in the observational health data sciences and informatics network. *J. Am. Med. Inform. Assoc.* **27**, 877–883 (2020).
13. Silverman, M. M. & De Leo, D. Why There Is a Need for an International Nomenclature and Classification System for Suicide. *Crisis* **37**, 83–87 (2016).
14. Gagliardi, J. P. & Turner, D. A. The Electronic Health Record and Education: Rethinking Optimization. *J. Grad. Med. Educ.* **8**, 325–327 (2016).
15. *Practice Guideline for the Assessment and Treatment of Patients With Suicidal Behaviors.* (American Psychiatric Association, 2010). at https://psychiatryonline.org/pb/assets/raw/sitewide/practice_guidelines/guidelines/suicide.pdf
16. *Best Practice in Managing Risk.* (Department of Health, 2009). at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/478595/best-practice-managing-risk-cover-webtagged.pdf
17. Meyer, R. E. *et al.* Suicidality and risk of suicide—definition, drug safety concerns, and a necessary target for drug development: a consensus statement. *J. Clin. Psychiatry* **71**, e1–e21 (2010).
18. Perera, G. *et al.* Cohort profile of the South London and Maudsley NHS Foundation Trust Biomedical Research Centre (SLaM BRC) Case Register: current status and recent enhancement of an Electronic Mental Health Record-derived data resource. *BMJ Open* **6**, e008721 (2016).
19. Sholle, E. T. *et al.* Secondary Use of Patients' Electronic Records (SUPER): An Approach for Meeting Specific Data Needs of Clinical and Translational Researchers. *AMIA Annu. Symp. Proc.* **2017**, 1581–1588 (2017).
20. Downs, J. *et al.* Detection of Suicidality in Adolescents with Autism Spectrum Disorders: Developing a Natural Language Processing Approach for Use in Electronic Health Records. *AMIA Annu. Symp. Proc.* **2017**, 641–649 (2017).
21. Velupillai, S. *et al.* Identifying Suicidal Adolescents from Mental Health Records Using Natural Language Processing. *Stud. Health Technol. Inform.* **264**, 413–417 (2019).
22. Chapman, W. W., Bridewell, W., Hanbury, P., Cooper, G. F. & Buchanan, B. G. A simple algorithm for identifying negated findings and diseases in discharge summaries. *J. Biomed. Inform.* **34**, 301–310 (2001).
23. Song, X. *et al.* Using Deep Neural Networks with Intra-and Inter-Sentence Context to Classify Suicidal Behaviour. in *Proceedings of The 12th Language Resources and Evaluation Conference* 1303–1310 (2020).
24. Chiticariu, L., Li, Y. & Reiss, F. Rule-based information extraction is dead! long live rule-based information extraction systems! in *Proceedings of the 2013 conference on empirical methods in natural language processing* 827–832 (2013).
25. Topaz, M. *et al.* Mining fall-related information in clinical notes: Comparison of rule-based and novel word embedding-based machine learning approaches. *J. Biomed. Inform.* **90**, 103103 (2019).

26. Tan, W. K. *et al.* Comparison of Natural Language Processing Rules-based and Machine-learning Systems to Identify Lumbar Spine Imaging Findings Related to Low Back Pain. *Acad. Radiol.* **25**, 1422–1432 (2018).
27. Vellido, A., Martín-Guerrero, J. D. & Lisboa, P. J. G. Making machine learning models interpretable. *ESANN* (2012). at <<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.431.5382&rep=rep1&type=pdf>>
28. Martens, D., Vanthienen, J., Verbeke, W. & Baesens, B. Performance of classification models from a user perspective. *Decis. Support Syst.* **51**, 782–793 (2011).