Detection of Suicidality in Adolescents with Autism Spectrum Disorders: Developing a Natural Language Processing Approach for Use in Electronic Health Records

Johnny Downs, MRCPsych*[1,2], Sumithra Velupillai*, PhD[1,3], George Gkotsis*, PhD[1], Rachel Holden, MSc[1,4] , Maxim Kikoler, MSc [1,4], Harry Dean, MSc [1], Andrea Fernandes, MSc[1], Rina Dutta, PhD, FRCPsych[1,2]

[1]Department of Psychological Medicine, NIHR Biomedical Research Centre, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK; [2]South London and Maudsley NHS Foundation Trust, London, UK; [3]School of Computer Science and Communication, KTH, Stockholm; [4]University of Canterbury, Southborough, UK.

* authors contributed equally to this work

Abstract

Over 15% of young people with autism spectrum disorders (ASD) will contemplate or attempt suicide during adolescence. Yet, there is limited evidence concerning risk factors for suicidality in childhood ASD. Electronic health records (EHRs) can be used to create retrospective clinical cohort data for large samples of children with ASD. However systems to accurately extract suicidality-related concepts need to be developed so that putative models of suicide risk in ASD can be explored. We present a systematic approach to 1) adapt Natural Language Processing (NLP) solutions to screen with high sensitivity for reference to suicidal constructs in a large clinical ASD EHR corpus (230,465 documents), and 2) evaluate within a screened subset of 500 patients, the performance of an NLP classification tool for positive and negated suicidal mentions within clinical text. When evaluated, the NLP classification tool showed high system performance for positive suicidality with precision, recall, and F1 scores all > 0.85 at a document and patient level. The application therefore provides accurate output for epidemiological research into the factors contributing to the onset and recurrence of suicidality, and potential utility within clinical settings as an automated surveillance or risk prediction tool for specialist ASD services.

Introduction

Over 1 in 6 young people with autism spectrum disorders (ASD) will contemplate or attempt suicide during childhood, making them 30 times more at risk than typically developing children.[1,2] Why children with ASD have higher rates of suicidal behaviours is unclear. It is possible that risk factors for childhood suicidal behaviour found in typically developing children, such as depression or being bullied, are more prevalent or potentially have a greater negative impact in children with ASD.[1] However, very little work has been conducted in ASD cohorts, and findings derived from non-ASD samples cannot be assumed to generalise to children with ASD.[3] A growing number of studies have shown that putative risk factors (both environmental and genetic) for psychiatric outcomes can have different effects in children with neurodevelopmental disorders.[4,5] Therefore individuals with ASD may express and manifest suicidal tendencies and behaviours in ways that differ from those observed in typical development.[6]

Given the widespread adoption of Electronic Health Records (EHRs) in primary and hospital care systems and the rapid growth of health informatics capabilities, longitudinal data from large samples of children with ASD can be used to develop and test new models of suicide risk behaviour. There is considerable potential to adapt EHR research methodologies used in recent epidemiological and risk factor studies[7] and apply these approaches to address the evidence gap in ASD and other vulnerable adolescent groups.[8] Although to capitalise on these developments for suicide research, accurate EHR data extraction systems need to be developed to capture data on those young people with ASD who present to public health services with suicidal thoughts or behaviours.

Information about suicidality in clinical documents is predominantly written in free-text. Haerian et al. showed that using only ICD-9 E-codes to detect patient-level suicide and suicide ideation from clinical text had the lowest positive predictive value (PPV): 0.55, while a combination of codes and Natural Language Processing (NLP) had the highest: 0.97, when applied on EHRs from the New York Presbyterian Hospital/Columbia University Medical Center.[9] They used MedLEE (Medical Language Extraction and Encoding System)[10] to generate Concept Unique Identifiers (CUIs) related to suicidality, and to filter out negated mentions as well as mentions not related to the patient. Anderson et al. applied a rule-based NLP approach to identify positive or negated mentions related to suicidality in the History of Present Illness (HPI) section of EHRs from a distributed health network of primary care organizations in the US, and found that suicidality information was predominantly recorded in free-text.[11]

Because suicidality is routinely assessed in mental health care, the absence or negation of suicidal behaviour is also documented in EHRs. An NLP tool developed specifically for detecting negated mentions of suicide in mental health records using syntactic tree information has been developed in our group[12] with high accuracy (91.9%) when evaluated on 6,000 sentences from mental health EHRs. The tool classifies each target mention (e.g. *suicid**) as negated or positive using a set of negation terms and rules applied on the information provided from a syntactic representation (constituency-based parse tree) of the sentence in which the target term is found, using the Stanford Core NLP toolkit for preprocessing and syntactic representation.[13]

The aim of this study was to extend, further develop and robustly evaluate an NLP approach which could accurately identify suicidality in ASD-patients' clinical records, with the future goal that it may provide data to enable improved risk prediction for related major adverse events, such as suicide attempts. Using EHR documents, such as progress notes, risk assessments and medical correspondence, we examined whether negation detection methods could be used to accurately identify references to suicidality in the EHRs of adolescents with ASD presenting to clinical mental health services. We defined suicidality as either the reporting of the intention to engage in a potentially lethal act towards oneself, or undertaking such acts themselves. To achieve the study aim, we developed coding rules using expert consensus, to define explicit suicidality-related mentions for adolescents with ASD seen in specialist mental health clinics (inpatient and ambulatory). Based on these rules we extended our NLP tool to 1) identify documents containing suicide-related (SR) information (i.e. NLP tool to screen documents) and 2) identify positive and negated references of suicidality on a document and patient level (i.e. NLP to classify SR documents and patients as positive, SR-Pos, or negative, SR-Neg) across a large number of EHRs. We then compared the performance of the NLP tool against expert human-rater case note reviews.

Materials and Methods

Data resources

This study used data extracted from the anonymised, electronic clinical records of a sample of adolescents with ASD referred to South London and Maudsley NHS Foundation Trust (SLaM). This sample and clinical setting has been described elsewhere,[14] but in brief SLaM provides specialist inpatient and outpatient ASD assessment and treatment services for young people from across the UK. Children and adolescents in this study were referred from primary care, child health, and educational and social care services, and typically underwent a multidisciplinary assessment by Child and Adolescent Mental Health Service (CAMHS) clinicians. Primary and secondary psychiatric disorders were diagnosed by CAMHS using the International Classification of Diseases, 10th Revision (ICD-10) multi-axial classification system.

The Clinical Record Interactive Search (CRIS) system[15] was used to produce an anonymised EHR dataset to search on structured data and free text fields for all ASD patients. CRIS was established in 2008 to allow searching and retrieval of full but de-identified clinical information for research purposes with a permission of secondary data analysis, approved by the Oxfordshire Research Ethics Committee C (reference 08/H0606/71+5). The patients were part of an open clinical cohort (entering and leaving the study at different time points) and included children aged 3–17 years with a diagnosis of ASD (ICD-10 F84.0, F84.1, F84.5, F84.9) recorded between 1 January 2008 and 31 December 2013. Free text entries, correspondence and reports were available for this sample from their initial assessment until June 2016. The resulting cohort contained 3,642 unique patients (complete age range). For the purposes of this study, we selected the sample of adolescents who had at least one contact with CAMHS (i.e one free text document in CRIS) between the ages 14 and 18 years, totalling 1,906 patients.

Overall workflow

Figure 1 outlines the overall workflow of our study. There were three main phases. The first phase related to the definition of classification rules to identify suicidality-related information in EHR documents for adolescents with ASD (step 1 below). These rules were then applied in the second phase where a manual review of documents (step 2) was used to inform the development and evaluation of our NLP approach to screen for suicidality related (SR) mentions in documents and filtering out documents with no mentions related to suicidality (NSR) – step 3 below. The NLP approach was then used to extract SR documents for the third phase (step 4). In the third phase, a manual review of documents was performed to annotate mentions of suicidality in SR documents as positive (SR-Pos), negative (SR-Neg) or uncertain (SR-U), step 5. Finally, the NLP approach was evaluated for its ability to correctly classify SR-Pos or SR-Neg in these documents and patients, step 6.

Step 1: Development of a set of classification rules to identify suicidality in adolescents with ASD.

Senior clinicians with expertise in the clinical management of neurodevelopmental disorders (JD) and suicidality assessment (RD) developed a set of rules to classify explicit mentions of suicidality in every document as either positive, negated or unknown. Positive mentions included text that referred to previous attempts, the presence of current or past plans of suicidal acts, command hallucinations related to carrying out a suicide attempt, a desire to be dead, researching methods, having ideas or describing plans of how to end their life or, a clinical opinion of the young person being at an elevated risk of attempting suicide. Negated terms included clinical opinions of the young person not being at elevated risk of suicide, recorded denial of suicidality by the young person (either directly or via third person report). Mentions were classified as uncertain, when aspects of suicidality were referred to, but did not appear to relate to risk of the young person being suicidal, for example references to dreams of being dead, or joking about death, or when references to suicidality were about other people (e.g. family members or friends).

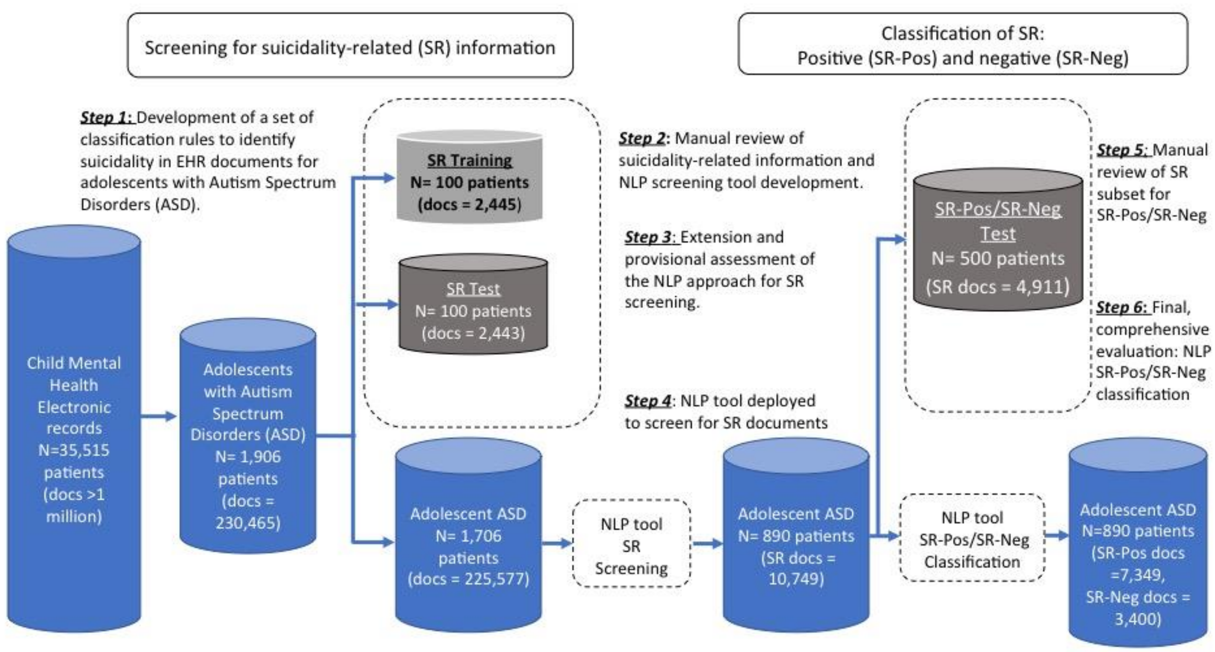

Figure 1. Overall workflow of the study.

Step 2: Manual review of suicidality-related (SR) information and NLP screening tool development.

A randomly extracted subset of 100 patients and their corresponding documents were allocated to a training corpus, and another random selection of 100 patients was allocated to the test corpus. To generate a subset of patients with a reasonable amount of documentation for manual review, the random sample was extracted based on documentation prevalence: each included patient had at least 7 documents (1st quartile) and at most 50 (3rd quartile), yielding a total of 2,445 (training set) and 2,433 (test set) documents in total. One clinically trained annotator (HD) was given these documents grouped for each patient. The annotator reviewed all documents for each patient, marked suicidality-related (SR) expressions, and labelled each SR-expression as either positive, negated or uncertain, according to the rules developed in step 1. However, for this phase, only annotations for SR information (regardless of polarity) were used for analysis.

Step 3: Extension and provisional assessment of the NLP approach for SR screening.

Results from the manual review were used to extend the NLP approach with the addition of new explicit SR expressions. Given the low frequency of the positive or negated SR mentions within the training set, we used the test set to assess precision, recall and F1-score of the tool detecting any SR content (regardless of polarity). Because the end goal is to address overall suicidality risk behaviour, the approach was evaluated on a document and patient level rather than on the mention level.

Step 4: NLP tool deployed to screen for SR documents

The NLP tool was then deployed to filter out documents without any SR mentions (positive or negative) from the original cohort (excluding the already annotated 200 patients). From 1706 patients (225,577 documents), 890 (52.2%) patients had at least one SR document, resulting in a total of 10,749 documents.

Step 5: Manual review of SR subset for identification of positive (SR-Pos) and negative (SR-Neg) suicidality mentions

Two manual coders (RH and MK) were randomly assigned 500 (56.2%) patients from the SR subset. Each annotator was given all documents (for each patient) that were detected by the NLP tool as containing a SR mention. The annotators were not given the NLP system output, but instead were asked to annotate explicit mentions of suicidality (same as above) and label these as positive (SR-Pos), negated (SR-Neg) or uncertain (SR-U). The documents were given to the annotators on a per-patient basis, and each patient was reviewed by one annotator. A subset (n=100) of randomly extracted documents was also used to calculate inter-rater agreement (measured with Cohen's κ and F1-score) on a document-level.

A majority rule was applied when evaluating document-level agreement: all mention-level annotations in each document were first counted, then, if the number of annotations labelled as positive for suicidality outnumbered or equalled the number of annotations labelled as negated, the document-level label was assigned “Positive for suicidality” – SR-Pos, otherwise it was designated “Negated for suicidality” – SR-Neg. To evaluate patient-level performance, priority was given to document-level outcomes: if the patient had at least one document labelled as SR-Pos using the majority rule, the patient-level label was assigned SR-Pos, irrespective of the number of previous or subsequent documents labelled as SR-Neg, i.e. each patient only required a single ‘positive suicidality’ document to be labelled SR-Pos.

Step 6: Final, comprehensive evaluation: NLP SR-Pos/SR-Neg classification

As a final step, the NLP approach was evaluated with precision, recall and F1-score against the manual annotations of the larger, filtered set of documents/patients with SR-Pos and SR-Neg labels, using the same heuristics for document- and patient-level classification assignments as above. Note that the evaluation is only performed on these two labels, i.e. SR-U annotations are not mapped to SR-Pos or SR-Neg. Thus, a false positive or false negative from the NLP approach could be due to an annotation marked as SR-U. A manual error analysis on cases of disagreements between the NLP tool and human annotation labels was also performed to gain a deeper understanding of the results.

Results

Table 1 shows the distribution of SR and non-SR (NSR) documentation and the individual level prevalence amongst the 100 adolescent patients with ASD in the final training set and the 100 patients in the test set. Manual review of both training and test documents revealed that only a small proportion of the corpus contained any SR information: <3% at the document level and around 22% at the patient level, with a similar distribution in the training and test set. Precision, recall, and F1 scores showed high system performance (> 0.8) for both SR and NSR in the test set (table 1).

The lexical markers of suicidality that were added to the NLP tool included *kill himself/herself/themselves/myself*, *end his/her/their life*, *take his/her/their own life*, *want to die*, *were dead*. Note that the NLP tool relies on lemmatised forms in both target expressions and document surface forms in order to achieve a more robust matching, e.g. different verb inflections of *want* will be matched with this approach.

Table 1. Confusion matrix: Screening for suicidality (SR) or non-suicidality (NSR), NLP tool compared to human annotation (A).

		NLP (Training)						NLP (Test)					
		Documents			Patients			Documents			Patients		
		NSR	SR	∑	NSR	SR	∑	NSR	SR	∑	NSR	SR	∑
A	NSR	2374	10	2384	75	2	77	2356	13	2384	73	5	78
	SR	5	56	61	0	23	23	8	56	64	1	21	22
	∑	2379	66	2445	75	25	100	2365	69	2443	74	26	100
	Precision	0.99	0.85		0.99	0.92		0.99	0.81		0.99	0.81	
	Recall	0.99	0.91		0.97	0.99		0.99	0.88		0.94	0.95	
	F1	0.99	0.88		0.99	0.96		0.99	0.84		0.96	0.88	

Table 2 shows the distribution of negated and positive suicidality-related information (SR-Pos/SR-Neg) using our majority rule criteria in 4,911 pre-screened documents derived from 500 patients. Evaluation of the NLP tool (table 2) showed high system performance for SR-Pos with precision, recall, and F1 scores all > 0.83 at a document and patient level. SR-Neg performance measures were lower, especially in recall (0.75 on document level, 0.62 on patient level), but overall good levels of classification were produced (F1 = 0.79 on document level, 0.72 on patient level).

A manual error analysis on a random sample of ten documents where the NLP tool classified a document as positive (SR-Pos) but the human annotator as negated (SR-Neg) was performed to gain a deeper understanding of the reasons behind the lower recall results. The main themes involved:

1) classification of documents with only one suicide-related mention (annotator SR-Neg count = 1, NLP SR-Pos count = 1) due to missing negation term, e.g. '<u>*Nil*</u> ***suicidal***' or error in syntactic parsing due to e.g. badly formatted sentences.

2) cases where the majority heuristic is problematic and negation detection scope is erroneous, e.g. one document annotated with SR-Neg = 2, while the NLP output was: SR-Neg = 1, SR-Pos = 3 contained the following: '*XXX* <u>*denied*</u> *any recent sleep difficulties, excessive fatigue or guilt, changes in appetite or morbid or* ***suicidal*** *ideation*', '*The risk of* ***suicide*** *is low, XXX* <u>*denies*</u> ***suicidal*** *ideation*.'

3) co-reference in combination with majority heuristics (annotator SR-Neg = 3, SR-Pos = 2, NLP SR-Neg = 1, SR-Pos = 2): '*XXX reported that XXX has had* ***suicidal*** *thoughts in the past but has* <u>*no*</u> *current plans on acting on* ***them****$_{co\text{-}reference}$*' (sentence repeated twice in document), '[clinician reporting] *further stated that* ***no*** *evidence of psychosis, self-harming behaviour,* ***suicidal*** *thoughts, sleep or appetite* …'

4) clinically challenging cases and complex information given in the document. Three examples are:

(i) Annotator: SR-U = 4, SR-Neg = 3, SR-Pos = 1; NLP tool output: SR-Neg = 1, SR-Pos = 2. '... *reported fleeting* ***suicidal*** *thoughts… but strongly denied that XXX could act on* ***these***$_{\textbf{co-reference}}$*…'*, '*presenting for serious OD without* ***suicidal****...'*, '*firm denial of* ***suicidal…'***. '*overdose'* marked as uncertain by annotator.

(ii) Annotator: SR-U = 2, SR-Neg = 1, NLP tool output: SR-Neg = 2, SR-Pos = 3 included sentences with information reported by external authorities such as the health care team and the school, references to the past, and includes a conclusive statement towards the end of the document: '*we could not assess*$_{negated}$ ***suicidal*** *ideation as XXX left the room'*, '*unable to assess*$_{negated}$ ***suicidal*** *ideation'*, '*historically*$_{past}$ *has threatened self harm and disclosed* ***suicidal*** *ideation…'*, '*concerns from school about* ***suicidal*** *ideation'*, '*no* ***suicidal*** *ideation expressed'*.

(iii) Annotator: SR-U = 2, SR-Neg = 1, NLP tool output: SR-Pos = 2: 'I tried to assess XXX's **suicidal** risk - XXX does not know if XXX wants to **kill XXXself**', 'XXX does not have any specific plan'

Table 2. Confusion Matrix: Classification of positive and negative SR, document- and patient level assessments. SR-Neg = Suicidality-related (SR) mention is negated (Neg), SR-Pos = Suicidality-related mention is positive (Pos).

		NLP (Test)					
		Documents			Patients		
		SR-Neg	SR-Pos	∑	SR-Neg	SR-Pos	∑
Annotator	SR-Neg	1379	463	1842	81	50	131
	SR-Pos	273	2796	3069	14	355	369
	∑	1652	3259	4911	95	405	500
	Precision	0.83	0.86		0.85	0.87	
	Recall	0.75	0.91		0.62	0.96	
	F1	0.79	0.88		0.72	0.92	

In total, 100 random documents were double annotated (table 3). A document-level assessment using the majority rule yielded an average Cohen's κ of 0.83, F1-scores for SR-Neg and SR-Pos document-level assessment were 0.89 and 0.94 respectively, indicating high agreement.

Table 3. Confusion Matrix: Inter-Rater Agreement on document level. SR-Neg = Suicidality-related (SR) mention is negated (Neg), SR-Pos = Suicidality-related mention is positive (Pos).

		Annotator 1		
		SR-Neg	SR-Pos	∑
Annotator 2	SR-Neg	32	2	34
	SR-Pos	6	60	66
	∑	38	62	100

Discussion and Conclusion

This is the first study to demonstrate that an NLP tool can be used to accurately capture a clinical construct as complex as suicidality within health records of young people with ASD. Our NLP tool identified suicidality-related (SR) mentions with high degrees of precision and recall from clinical free text held within EHRs. This NLP application provides powerful opportunities for surveillance work in adolescent ASD and in other clinical samples, with the potential to improve risk prediction for major adverse events, such as suicide attempts.

The development of this high performance NLP tool was achieved in several steps. First, owing to the potentially distinctive characteristics of the ASD clinical population, and their specialist mental health service provision, we began by building a suicidality terminology from a detailed note review of over 2000 random sets of clinical entries in 100 children with ASD, combined with expert clinical consensus. Because of the limited literature on suicidal terminology in ASD, we used a randomly extracted training and test set from all potential ASD EHR source data, rather than an enriched set filtered by restricted terms (e.g. "suicid" [12] or ICD coding classifications).[16] The rationale for this was to reduce selection bias and loss of sensitivity through the use of training and test data derived using restricted terms or coding classifications.

Random selection from the whole potential corpus also provided us with a better understanding of the overall distribution of suicidality-related information in documents and allowed us to refine and advise on additional terminology. During the training phases, it became clear that there was a low frequency of SR terms (less than 3% of all documents). A much larger corpus was required to conduct an adequate test of the NLP tool's classification performance in discerning positive and negated SR mentions within the documents. We therefore configured the NLP tool to provide a SR screening step and identify a smaller, more feasible volume of suicidality-related documents for human annotation from the whole ASD corpus of c. 220,000 documents. We then tested the positive predictive value (precision) and sensitivity (recall) of the NLP screening tool against the EHR of another 100 children (approximately 2000 human rated documents) which were annotated for any references to suicidality. Finally, we began a comprehensive test of the accuracy of the tool against positive (SR-Pos) or negated (SR-Neg) mentions of suicide within the screened suicide related documents.

The abstraction of mention-level annotations and NLP system predictions to document- and patient-level assessments using simple heuristics (majority rule for document level and SR-Pos priority on patient level) showed that promising results can be obtained even though the NLP tool relies only on a relatively small number of suicidality-related and negation terms. This finding also shows that even though suicidality behaviour is documented with a variety of expressions (e.g. '*took an excessive amount of pills*', '*threw him/herself in front of a train*'), indicative terms (mainly *suicide* in different forms) are typically also used at some point in the documentation and will thus be detected automatically.

Our motivation for applying a majority rule on document level assessments was based on the finding that the main source for false positive errors in our negation detection approach stemmed from cases of question forms (e.g. '*I asked him if he feels suicidal*'), references to the past, etc. Applying this rule was a way of smoothing this error rate. However, the error analysis showed that this approach might be a limitation. In future studies, we aim to compare results with NLP approaches such as ConText[17] where variables relating to the past ('historicity') and subject ('experiencer') are encoded with target terms. We also aim to experiment with other abstraction heuristics, e.g. instead of majority rule, applying a priority hierarchy. In keeping with prior work, another alternative could be to define the annotation task on a document level[9]. Longer term, we aim to compare the predictive validity of different heuristics within our NLP tool, and across other NLP approaches, for later adverse outcomes (i.e. significant suicide attempts or death by suicide), and seek external validity through replication in other EHR systems. Without these further steps, it is difficult to assess the potential clinical impact of differences in precision or recall across NLP tools.

The annotators expressed that it was sometimes challenging to assess suicidality risk based on one document at a time; single documents did not provide sufficient context in all cases. At the same time, given the rare prevalence of suicide-related content in all patient documents, defining a patient-level annotation task using this type of abundant clinical documentation would be very time-consuming. We plan to explore different ways of addressing this issue, one being a nested case-control study design similar to the one presented in Metzger et al.[18]

A strength of this study is that we have not assumed that clinical terms used in more typically developing children or adults generalize to ASD populations. Assessing suicidality in adolescents with ASD often requires a different approach to other patient groups, which in our clinical experience was likely to be reflected in the clinical notes. Young people with ASD presenting to mental health services commonly have severe difficulties with interpersonal interactions, making for a more complex clinical assessment.[19] Clinicians are likely to deliberate within the clinical notes on whether potential behaviours are driven by suicidal ideation. They may have a greater reliance on third person report – i.e. caregivers voicing concerns regarding the young person's suicidality rather than direct accounts from the young person. Also, where a first person account is provided, clinicians will often write verbatim statements (i.e. He told me "I just want to end it", and he "went to the car park to get it done"), providing more atypical clinical terminology for describing suicidality, and increasing the chance of NLP misclassification.

In addition, young people with ASD may not present with suicidality as a principle complaint, but through a behavioural change such as school refusal, with suicidal behaviour emerging through later clinician screening. This may change the emphasis and position within the patient's clinical record relative to other populations where suicidal behaviour is the principle trigger during the first presentation to services. Testing these clinical assumptions empirically using an non-ASD control sample was beyond the scope of the current study, however future work is underway to examine the variability of the NLP tool's accuracy across non-ASD child populations seen in mental health services. NLP applications are commonly validated using randomly extracted documents from EHRs covering a broad range of clinical contexts, seldom rarer clinical populations, such as young people with ASD. As mental health assessment and management needs to be tailored to the developmental needs of the young people in clinic, so should the validation of NLP data extraction tools.

The suicidality outcome data provided by this NLP extraction tool permits analyses of the complex interplay of ASD-specific traits on factors contributing to the onset and recurrence of suicidality. ASD specific mental health services are becoming increasingly available for child and adolescent populations in high-income countries. Although there is more work to be done before clinical application, we believe the NLP tool described provides a step forward in enhancing suicidality surveillance, risk prediction and treatment selection for children with ASD.

Acknowledgements

JD is supported by a Medical Research Council (MRC) Clinical Research Training Fellowship (MR/L017105/1), SV is supported by the Swedish Research Council (2015-00359), Marie Skłodowska Curie Actions, Cofund, Project INCA 600398, HD received salary support from the Foundation of Professional Services to Adolescents, RD is funded by a Clinician Scientist Fellowship (research project e-HOST-IT) from the Health Foundation in partnership with the Academy of Medical Sciences which also funded GG's salary.

References

1 Mayes SD, Calhoun SL, Baweja R, Mahr F. Suicide ideation and attempts in children with psychiatric disorders and typical development. *Crisis J Crisis Interv Suicide Prev* 2015; **36**: 55–60.
2 Segers M, Rawana J. What Do We Know About Suicidality in Autism Spectrum Disorders? A Systematic Review. *Autism Res* 2014; **7**: 507–21.
3 Burack JA, Iarocci G, Flanagan TD, Bowler DM. On mosaics and melting pots: conceptual considerations of comparison and matching strategies. *J Autism Dev Disord* 2004; **34**: 65–73.
4 Caspi A, Langley K, Milne B, *et al.* A replicated molecular genetic basis for subtyping antisocial behavior in children with attention-deficit/hyperactivity disorder. *Arch Gen Psychiatry* 2008; **65**: 203–10.
5 Simonoff E, Pickles A, Charman T, Chandler S, Loucas T, Baird G. Psychiatric Disorders in Children With Autism Spectrum Disorders: Prevalence, Comorbidity, and Associated Factors in a Population-Derived Sample. *J Am Acad Child Adolesc Psychiatry* 2008; **47**: 921–9.
6 Raja M. Suicide risk in adults with Asperger's syndrome. *Lancet Psychiatry* 2014; **1**: 99–101.
7 Barak-Corren Y, Castro VM, Javitt S, *et al.* Predicting Suicidal Behavior From Longitudinal Electronic Health Records. *Am J Psychiatry* 2016; **174**: 154–62.
8 Downs J, Gilbert R, Hayes R, Hotopf M, Ford T. Linking health and education data to plan and evaluate services

for children. *Arch Dis Child* (in press). doi:10.1136/archdischild-2016-311656.
9 Haerian K, Salmasian H, Friedman C. Methods for identifying suicide or suicidal ideation in EHRs. *AMIA Annu Symp Proc AMIA Symp* 2012; **2012**: 1244–53.
10 Friedman C, Alderson PO, Austin JH, Cimino JJ, Johnson SB. A general natural-language text processor for clinical radiology. *J Am Med Inform Assoc* 1994; **1**: 161–74.
11 Anderson HD, Pace WD, Brandt E, *et al.* Monitoring suicidal patients in primary care using electronic health records. *J Am Board Fam Med JABFM* 2015; **28**: 65–71.
12 Gkotsis G, Velupillai S, Oellrich A, Dean H, Liakata M, Dutta R. Don't Let Notes Be Misunderstood: A Negation Detection Method for Assessing Risk of Suicide in Mental Health Records. In: The Third Computational Linguistics and Clinical Psychology Workshop (CLPsych). 2016: 95–105.
13 Klein D, Manning CD. Accurate unlexicalized parsing. In: Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1. Association for Computational Linguistics, 2003: 423–30.
14 Downs J, Hotopf M, Ford T, *et al.* Clinical predictors of antipsychotic use in children and adolescents with autism spectrum disorders: a historical open cohort study using electronic health records. *Eur Child Adolesc Psychiatry* 2016; **25**: 649–58.
15 Perera G, Broadbent M, Callard F, *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* 2016; **6**: e008721.
16 Ting SA, Sullivan AF, Boudreaux ED, Miller I, Camargo CA. Trends in US emergency department visits for attempted suicide and self-inflicted injury, 1993–2008. *Gen Hosp Psychiatry* 2012; **34**: 557–65.
17 Harkema H, Dowling JN, Thornblade T, Chapman WW. Context: An Algorithm for Determining Negation, Experiencer, and Temporal Status from Clinical Reports. *J Biomed Inform* 2009; **42**: 839–51.
18 Metzger M-H, Tvardik N, Gicquel Q, Bouvry C, Poulet E, Potinet-Pagliaroli V. Use of emergency department electronic medical records for automated epidemiological surveillance of suicide attempts: a French pilot study. *Int J Methods Psychiatr Res* 2016. http://onlinelibrary.wiley.com/doi/10.1002/mpr.1522/full (accessed March 7, 2017).
19 Richa S, Fahed M, Khoury E, Mishara B. Suicide in Autism Spectrum Disorders. *Arch Suicide Res* 2014; **18**: 327–39.