INNOVATIONS JOURNALS

Brief Report

Using Natural Language Processing to Detect Suicidal Ideation and Prompt Urgent Interventions: A Retrospective Database Study

Matt Hartz,[1] Doug Hickey,[2] Laura Acosta,[2] Amanda Brooks,[2] Dan Holley,[2] Tom Zaubler[2]

[1]Data Engineering, NeuroFlow, Inc., Philadelphia, PA, USA
[2]Clinical Operations, NeuroFlow, Inc., Philadelphia, PA, USA

Address correspondence to Dr. Dan Holley (dan@neuroflow.com).

Sources of Support: This research was funded by NeuroFlow, Inc.

Conflicts of Interest: All authors are employees of NeuroFlow, Inc., whose product, the NeuroFlow digital behavioral health platform, was evaluated in this study.

Submitted: Sep 11, 2023; First Revision Received: Dec 15, 2023; Accepted: Dec 18, 2023

Hartz M, Hickey D, Acosta L, Brooks A, Holley D, Zaubler T. Using natural language processing to detect suicidal ideation and prompt urgent interventions: a retrospective database study. *Innov Dig Health Diagn Bio*. 2024; 4:6–8. DOI: 10.36401/IDDB-23-10.

Suicide claimed over 48,000 lives in the United States in 2021 and is a leading cause of death among Americans aged 10–54.[1] Suicidal ideation (SI) predicts suicidal behavior[2] and is commonly evaluated in healthcare settings via intake screenings and psychometric assessments.[3] Nevertheless, the timely detection of SI is challenging: Progression from the first instance of SI to the first suicide attempt is often rapid,[4] and roughly half of all adults who die by suicide have contact with a healthcare provider within 1 month before their death.[5] This disconnect underscores the need for asynchronous approaches that complement existing SI-detection frameworks and enable low-latency interventions. Digital behavioral health (dBH) technologies promise to improve SI detection and enable timely interventions by remotely administering prescribed assessments.[6] However, remote assessments alone may fail to detect SI among patients whose assessment compliance has lapsed and those who are compliant but do not indicate SI.

Intriguingly, dBH platforms are well-suited to elicit free-form text entries from patients and remotely screen those entries for SI indicators via natural language processing (NLP)—a field of computer science that enables analysis of ingested text through the application of linguistic, statistical, and machine-learning principles.[7] Here, we used retrospective real-world data from patients who engaged in journaling exercises via the prescribed dBH platform NeuroFlow[6] to show that an asynchronous NLP tokenization[7] protocol can augment the Patient Health Questionnaire-9[2] (PHQ9) by detecting SI in patients who have either fallen out of assessment compliance or who did not indicate SI in their most recent PHQ9. The NLP protocol leveraged keyword detection from a continuously refined SI lexicon, and clinicians reviewed alert-generating text entries to enable further refinement.

We performed a retrospective database study of data collected in the course of routine healthcare operations from 425 patients (316 women, 86 men, 1 transgender, and 22 individuals who did not disclose or indicated "other"; mean [SD] age = 41.67 [15.22] years) across various national insurance payors, healthcare providers, employers, and government/military clients who used the dBH platform to supplement patient care. Patients were selected if they had been flagged via asynchronous NLP screenings as having expressed SI in free-form text responses to platform-administered journaling prompts (see Supplemental Table, available online).

This study was conducted in accordance with the principles of the Declaration of Helsinki. It was exempted from oversight and waived for informed consent following a review by Advarra's Center for IRB Intelligence, which determined that the study constitutes secondary research as defined in 45 CFR 46.104(d)(4).

For each patient's first instance of NLP-detected SI, we evaluated the prior 30 days for PHQ9 compliance. When compliance was met, we evaluated Question 9 (Q9) and interpreted responses greater than 0 as indicative of SI. Question 9 was as follows: "Over the last 2 weeks, how often have you been bothered by... Thoughts that you would be better off dead or of hurting yourself in some way?" The response set included not at all (0), several days (1), more than half the days (2), and nearly every day (3).

Between January 8, 2020, and August 9, 2023, NLP screenings (Fig. 1) detected possible SI in 425 patients, 344 of whom were compliant with PHQ9 assessments. Of those compliant, 177 indicated SI on Q9 of their

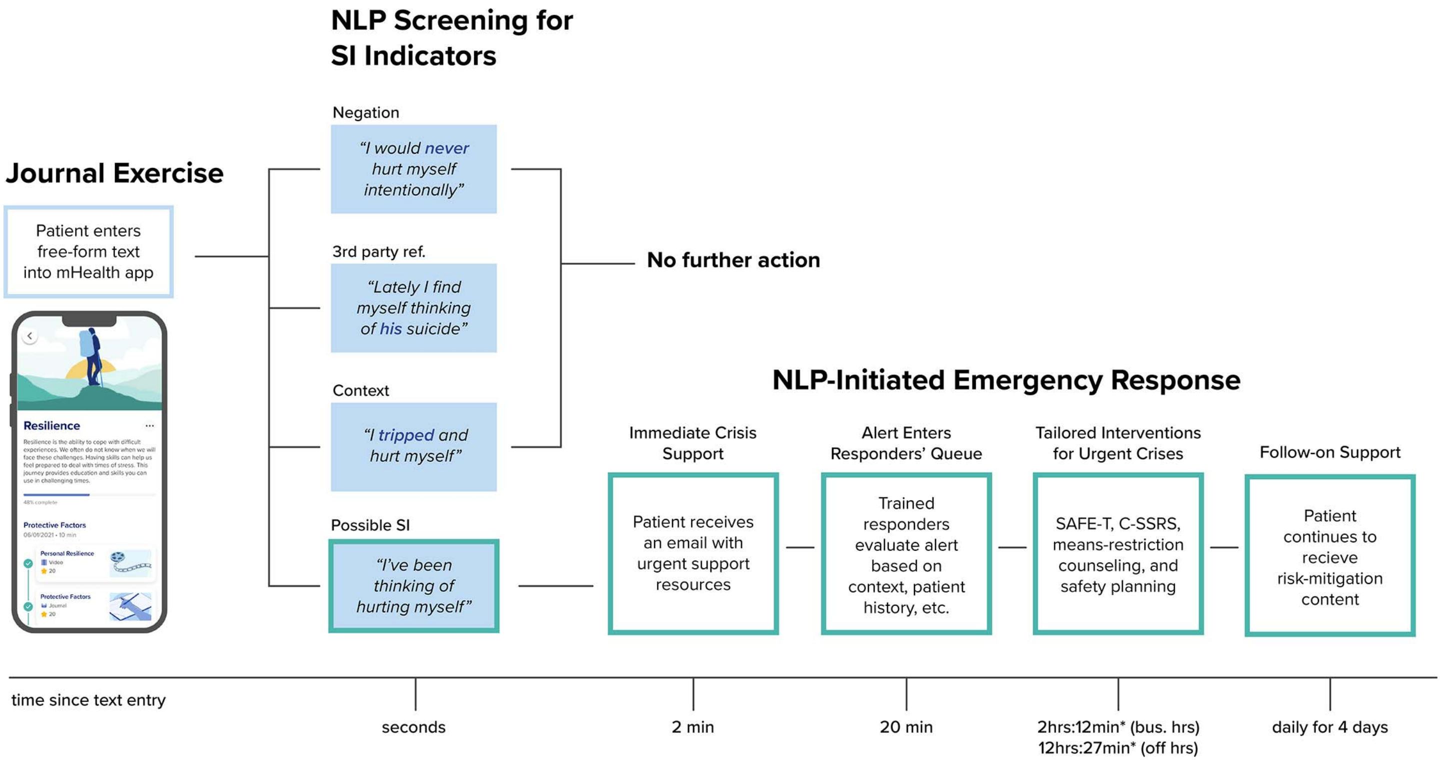

Figure 1. Natural language processing (NLP) for asynchronous suicidal ideation (SI) detection. (*from left to right*) Patients enter free-form text into journaling exercises; for example, via the dBH platform's mobile healthcare (mHealth) app. A proprietary NLP protocol checks for indicators of SI and rules out text containing negations, third-party references, or mitigating contexts. Text that is not covered by any of those checks prompts immediate automated crisis support and routes an alert to trained responders for low-latency evaluation. For patients evaluated as "urgent," responders attempt to make contact to provide personalized intervention. All patients receive follow-on support for 4 days. *Median average. SAFE-T: Suicidal Assessment Five-Step Evaluation and Triage.

most recent assessments (mean PHQ9, SD = 18.45, 5.28), and 167 did not (8.42, 5.92). The remaining 81 of 425 patients were out of assessment compliance.

The dBH platform automatically sent an email containing evidence-based crisis support to each of these 425 patients within 2 minutes of SI detection. These emails included national and local crisis support resources and notification that a check-in would occur in the next 24 hours. The platform simultaneously transmitted an alert for each patient to a response team of mental health professionals trained in crisis intervention (supervised by licensed clinicians) who monitor the alert queue during normal business hours (09:00–20:00 EST). Responders evaluated each alert and attempted outreach to provide tailored interventions to patients determined to be in urgent crises. The median time from SI detection to attempted or completed intervention, including case-note logging, was 2 hours 12 minutes during business hours and 12 hours 27 minutes during off hours. During interventions, responders employed evidence-based suicide prevention strategies[8] by administering the Columbia-Suicide Severity Rating Scale (C-SSRS), initiating the Suicidal Assessment Five-Step Evaluation and Triage (SAFE-T) protocol, conducting safety planning, providing lethal means counseling, and orchestrating further care based on assessed risk. Of note, patients who triggered alerts during off hours received immediate platform-generated support as described above. All patients received platform-generated support focused on suicide risk mitigation and consisting of videos, handouts, and journaling exercises immediately after SI detection and periodically for the next 4 days.

Limited case note availability due to evolving policies and cross-organizational data-sharing strictures is a weakness of our study. Our team performed outreach as described above for 23 of 425 patients who generated alerts, with the remainder serviced by external organizations whose data-sharing policies preclude retrospective case note analysis. Furthermore, even among the patients our team serviced, the dBH platform's data capture and retention policies before March 2022 limit our insights. Case notes were available for only 17 patients serviced by our organization (ie, patients who prompted alerts after March 2022). Our care team evaluated three of these patients as "urgent" and stratified each as "moderate risk" via C-SSRS. Case notes show that these patients received low-latency, personalized care based on their C-SSRS responses, medical history, patient-responder dialogue, and other factors. Each was assessed as "safe" after the intervention. The remaining 14 were either stratified as "low risk" on NLP review, declined completion of the C-SSRS, did not return attempts to contact and were referred to their primary care provider for follow-up, or confirmed safety via text. Future studies that codify a priori cross-organizational data capture and sharing policies will be necessary to enable well-powered retrospective case note analysis.

Our cursory findings evidence the utility of asynchronous NLP to trigger urgent, potentially life-saving interventions for patients whose SI may otherwise go undetected. This approach can augment existing SI-detection frameworks to reach at-risk patients, irrespective of when or where SI occurs. Because this approach is location agnostic and readily scalable, it may be well-suited to support marginalized communities where suicide rates are highest, psychiatric resources are scarce, and social determinants present barriers to care.[9] Our findings pave the way for prospective studies that evaluate outcomes among patients receiving differential support from tiered suicide prevention programs—potentially leading to a causal understanding of the role of dBH technologies in suicide prevention. Finally, future efforts to develop a more comprehensive SI lexicon (eg, using large autoregressive language models[10] promise to complement our NLP approach by making screenings more robust to slang, misspellings, metaphors, and so on.

References

1. Garnett M, Curtin S. Suicide mortality in the United States, 2001–2021. NCHS Data Brief, number 464. U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Health Statistics (NCHS), 2023.
2. Rossom RC, Coleman KJ, Ahmedani BK, et al. Suicidal ideation reported on the PHQ9 and risk of suicidal behavior across age groups. *J Affect Disord*. 2017;215:77–84.
3. Mann JJ, Michel CA, Auerbach RP. Improving suicide prevention through evidenced-based strategies: a systematic review. *Am J Psychiatry*. 2021;178:611–624.
4. Deisenhammer EA, Ing CM, Strauss R, Kemmler G, Hinterhuber H, Weiss EM. The duration of the suicidal process: how much time is left for intervention between consideration and accomplishment of a suicide attempt? *J Clin Psychiatry*. 2009;70:19–24.
5. Ribeiro JD, Gutierrez PM, Joiner TE, et al. Health care contact and suicide risk documentation prior to suicide death: results from the Army Study to Assess Risk and Resilience in Servicemembers (Army STARRS). *J Consult Clin Psychol*. 2017;85:403–408.
6. Pardes A, Lynch W, Miclette M, McGeoch E, Daly BP. Use of a mobile health (mHealth) platform for remote assessment of suicidal ideation, depression, and anxiety: a longitudinal retrospective study. *Innov Digit Health Diagn Bio*. 2021;2:8–15.
7. Nadkarni PM, Ohno-Machado L, Chapman WW. Natural language processing: an introduction. *J Am Med Inform Assoc*. 2011;18:544–551.
8. Brodsky BS, Spruch-Feiner A, Stanley B. The zero suicide model: applying evidence-based suicide prevention practices to clinical care. *Front Psychiatry*. 2018;9:33.
9. Liu S, Morin SB, Bourand NM, et al. Social vulnerability and risk of suicide in US adults, 2016–2020. *JAMA Netw Open*. 2023;6:e239995.
10. Carpenter KA, Altman RB. Using GPT-3 to build a lexicon of drugs of abuse synonyms for social media pharmacovigilance. *Biomolecules*. 2023;13:387.