https://doi.org/10.36315/2026inpact091
Psychological Applications and Trends 2026

PRAXIS: A DOCTORAL PILOT STUDY OF A PRACTICE-BASED INTERVENTION FOR STRESS AND WELLBEING

Sara El Matouk
Universitas Mercatorum (Italy)

Abstract

PRAXIS is an early-stage doctoral protocol. Background: sustained academic pressure, role ambiguity, and limited access to timely support can elevate stress and undermine wellbeing in young adults and early-career researchers. This project aims to develop and pilot PRAXIS: Predictive Risk Assessment with eXplainable & Interpretable Systems, a brief, scalable, practice-based program that translates evidence from applied psychology into daily routines. PRAXIS will integrate psychoeducation about stress, skills for self-regulation and guided reflection to strengthen coping and psychological flexibility. Methods: an iterative mixed-methods approach will be used. In Phase 1, it was designed the protocol, that integrates validated Italian versions of the Columbia Suicide Severity Rating Scale (C-SSRS) and the Orbach–Mikulincer Mental Pain Scale (OMMP-8) with two standardized writing tasks (a distress narrative and a best-possible-self narrative) followed by post-task distress sliders. Annotation will apply Columbia Classification Algorithm of Suicide Assessment (C-CASA) categories alongside a composite severity index (severity_praxis, 0–9 or U). In Phase 2, a feasibility pilot will assess recruitment, retention, adherence, and intervention fidelity, alongside acceptability indicators (satisfaction and perceived usefulness). Expected contribution: although no results are available yet, this submission outlines the theoretical rationale and planned methodology. If feasibility and acceptability are supported, PRAXIS will inform a subsequent randomized controlled trial and provide a transferable toolkit for stress management and wellbeing promotion in academic and community settings.

Keywords: *Intervention development, stress management, wellbeing, psychological flexibility, doctoral research.*

1. Introduction

Suicide is among the leading causes of preventable death globally, and its accurate assessment remains a fundamental challenge for clinical research (Fazel & Runeson, 2020). A critical conceptual difficulty concerns distinguishing suicidal ideation from suicidal behaviour: theories such as the Ideation-to-Action framework and the Three-Step Theory (Klonsky & May, 2015) posit that the two are governed by partially distinct mechanisms, yet most assessment instruments conflate or inadequately separate them. The Columbia Suicide Severity Rating Scale (C-SSRS; Posner et al., 2011) and the Columbia Classification Algorithm of Suicide Assessment (C-CASA; Posner et al., 2007) represent current best practice for standardized, hierarchical taxonomy of suicidal phenomena, but their application demands trained raters and is difficult to scale across institutions.

A compounding problem is the difficulty of integrating data across institutions. Heterogeneous instruments, varying time frames, and incompatible coding schemes make harmonization non-trivial: Campos et al. (2023) found that concordance across 20 instruments and 27 international cohorts is often modest. European data-protection law (GDPR) further constrains sharing of sensitive mental-health data, requiring careful pseudonymization and governance architectures before any multi-site pooling is feasible. Self-report questionnaires, while essential, carry well-documented limits: social desirability effects, under-reporting driven by stigma, and fixed-response formats that constrain idiomatic expression of distress (Khazem et al., 2021; Campos et al., 2023). Standardized short free-text narratives can provide complementary signal: research shows that language features (pronoun patterns, negative affect vocabulary, temporal markers) carry information related to psychological distress and suicidal ideation (Li et al., 2023; Arowosegbe & Oyelade, 2023). However, this evidence base is largely English-language, relies heavily on social media data, and rarely pairs text with validated clinical assessments. A clinically anchored, governed Italian corpus is therefore a meaningful gap to address.

Against this backdrop, PRAXIS pursues three exploratory hypotheses, formulated conservatively given the early stage of this work:

H1 (exploratory): Higher ideation and severity scores may be associated with elevated negative affect markers, increased self-focused pronoun use, and reduced future-oriented language in distress narratives.

H2 (exploratory): Mental pain scores (OMMP-8) may correlate with semantic clusters reflecting entrapment, hopelessness, or social disconnection themes.

H3 (exploratory): Protective (best-possible-self) narratives may show greater future orientation and agency-related language relative to distress narratives.

2. Method

PRAXIS: Predictive Risk Assessment with eXplainable & Interpretable Systems is a cross-sectional, observational study administered entirely online via a secure web-based platform. The platform is selected according to a governance checklist that requires disabled IP logging, encrypted storage, and secure export (see Section 2.4).

2.1. Participants

Participants are undergraduate students aged 19–23, recruited online through institutional channels at Italian universities. Participation is entirely voluntary and unpaid. Inclusion criteria are: Italian as primary language, age 19–23, and provision of informed digital consent. No minimum sample size is pre-specified at this pilot stage, as the primary goal is protocol validation and corpus construction.

2.2. Measures

The battery proceeds in a fixed order. (1) A demographics form collects exact age, region of origin within Italy, and sex. (2) The C-SSRS Baseline in its validated Italian version assesses suicidal ideation on an ordinal 0–5 scale and categorizes suicidal behaviours (none, preparatory, aborted, interrupted, actual, completed, non-suicidal self-injury, intent unknown; Posner et al., 2011). (3) The OMMP-8 in its validated Italian version assesses mental pain as a transdiagnostic proximal risk factor (Landi et al., 2025; Tossani et al., 2021). (4) Brief quantitative distress check items appear between sections to monitor participant state.

2.3. Standardized writing tasks

After the structured measures, participants complete two short writing tasks in fixed order. Task 1 (Distress Narrative) asks participants to write about their feelings and thoughts during a stressful or difficult event in the past six months; they are explicitly instructed not to include names or identifying details. Task 2 (Best Possible Self) asks participants to describe, as realistically as possible, their best possible self one year from now, again without identifying information. Both tasks use identical interface elements: a fixed 6-minute timer, a live word counter, and a minimum threshold of 80 words. These parameters standardize production conditions for cross-participant comparability. Immediately after each task, participants complete two 0–10 sliders rating: distress intensity while writing and perceived controllability or manageability of that distress.

2.4. Data privacy and governance

PRAXIS adopts a privacy-by-design architecture. Each participant is assigned a random alphanumeric identifier at registration; no real names, email addresses, or institutional identifiers are linked to response data. IP addresses and precise geolocation data are not collected. Only non-identifying technical metadata may be retained where operationally necessary (completion time, word count, device type). Participants are instructed before writing not to include names, dates, locations, or other potentially identifying details. Data access is restricted to the principal researcher and designated clinical annotators under a signed data-access agreement, with retention limited to the duration of the active project. Crucially, PRAXIS is a research protocol and is not used for clinical decision-making; this is communicated clearly to participants. A non-clinical escalation screen is shown to participants endorsing active ideation on check items, consistent with a human-in-the-loop safety philosophy (Chandler et al., 2022).

2.5. Data analysis

Each text will receive four annotation fields. First, a primary C-CASA label (Posner et al., 2007) will classify the overall nature of the expressed content. Second, a C-SSRS ideation level (Posner et al., 2011) will be assigned only if ideation content is discernible in the text; otherwise, this field will remain null. Third, a C-SSRS behaviour field will record the highest-order suicidal behaviour detectable (none, preparatory, aborted, interrupted, actual, completed, NSSI, intent_unknown). Fourth, a composite index severity_praxis (0–9 or U) will be derived as max (ideation severity, behaviour severity): behaviour present maps to 6–9; ideation only without behaviour maps to 1–5; no suicidality maps to 0; insufficient information yields U (undetermined). This index provides a single ordinal anchor for downstream modelling while preserving the ideation/behaviour distinction.

A subset of 20–30% of texts, stratified across severity levels, will be independently co-annotated by two clinicians blind to participant demographics. Inter-rater agreement will be computed using weighted Cohen's kappa for ordinal fields and Cohen's kappa or Krippendorff's alpha for categorical fields. Discrepancies will be reviewed collaboratively. The resulting discussion will update the annotation guidelines to version 2. Persistent disagreements will receive the code U.

Analyses will proceed in three steps. First, interpretable psycholinguistic features will be extracted (including first-person singular pronoun frequency, negation density, emotional valence and arousal vocabulary, temporal markers and social reference words) drawing on LIWC-like category logic (Tausczik & Pennebaker, 2010).

These features offer the highest communicability to clinical audiences and will be the primary unit for initial exploratory correlational and group-difference analyses.

Second, word and sentence embeddings will be computed to explore semantic clustering within the Italian corpus, with the aim of identifying families of expressions proximal to suicide-related content in Italian. Third, transformer-based representations will be explored in an explicitly hypothesis-generating frame. Throughout all modelling, the protocol enforces conservative practices: cross-validation and held-out test sets to control overfitting, explicit handling of class imbalance, threshold-sensitivity analysis and comprehensive error analysis foregrounding false-negative cases, which carry the greatest clinical risk. No claim of clinical prediction will be made on the basis of this pilot study (Chandler et al., 2022; Pigoni et al., 2024).

3. Results

All contributions described below are conditional on completion of data collection and annotation, as the study is in an early stage. Methodologically, PRAXIS will produce the first standardized Italian corpus of short texts paired with clinically anchored suicide-related labels, addressing a documented gap in non-English NLP resources for suicidology (Arowosegbe & Oyelade, 2023). The severity_praxis index will offer a reusable composite measure anchored to two internationally validated instruments, serving as a prototype for future cross-site harmonization. The annotation guidelines, platform governance checklist, and privacy-by-design architecture are intended as directly replicable artefacts for other Italian university settings.

Clinically, the psycholinguistic feature set will yield a preliminary map of Italian textual markers associated with levels of suicidal ideation and mental pain, presented in interpretable form to maximize communicability to practitioners. Theoretically, the design allows direct comparison between distress and protective narratives, testing whether future-oriented and agency-related language differs across groups, consistent with theoretical accounts of protective cognition in suicidology (O'Connor & Nock, 2014).

The human-in-the-loop governance model may also serve as a blueprint for responsible deployment of NLP-assisted tools in sensitive clinical research contexts.

4. Discussion

PRAXIS addresses a recognized methodological gap in Italian-language suicidology research: the absence of a standardized, clinically anchored corpus of free-text data collected under reproducible, ethically governed conditions. Existing self-report instruments, while indispensable, are constrained by fixed-response formats and well-documented under-reporting tendencies; standardized elicited narratives may provide complementary, less demand-susceptible signal. By pairing two writing tasks, one distress-oriented, one protective, with validated measures of suicidal ideation, behaviour, and mental pain, PRAXIS grounds textual data in a coherent clinical taxonomy. The severity_praxis index, derived as the maximum of ideation and behaviour severity fields, provides a single ordinal anchor for downstream modelling without obscuring the theoretically important ideation-to-behaviour distinction.

Key methodological strengths include rigid task standardization (fixed timer, word minimum, identical instructions), a structured reliability procedure involving blind clinical co-annotation, an interpretable-first NLP pipeline. Nonetheless, several risks warrant explicit acknowledgment: the study is early-stage and the convenience sample of university students limits generalizability: free-text narratives are not immune to self-presentation effects, annotation at intermediate severity levels may remain contested and transformer-based models trained on a small pilot corpus carry substantial overfitting risk.

Several limitations should be acknowledged. The protocol is at an early stage and all anticipated contributions remain prospective. Participants constitute a convenience sample of young Italian university students, generalizability to older adults, clinical populations, or other cultural contexts is a future target. Free-text narratives, while less constrained than fixed questionnaires, are not immune to self-presentation effects (Khazem et al., 2021). Annotation is inherently labour-intensive and subject to clinician disagreement, particularly at intermediate severity levels; the reliability procedure mitigates but cannot eliminate this. NLP models trained on this pilot corpus will reflect its specific population and task parameters, and any future applied use would require independent clinical validation and regulatory review.

The primary next step will be replication of the same protocol across different student cohorts and, subsequently, additional sample types, to improve external validity. Multi-site replication would require a shared governance framework alongside formal data-sharing agreements aligned with GDPR constraints.

5. Conclusion

PRAXIS (Predictive Risk Assessment with eXplainable & Interpretable Systems) is an early-stage doctoral protocol designed to collect standardized short Italian texts from undergraduate students aged 19–23 and to produce clinically anchored labels of suicide-related risk severity, enabling exploratory extraction of semantic and psycholinguistic indicators within a human-in-the-loop framework. Its primary contribution is not a predictive model but a replicable pipeline: a governed, pseudonymized corpus paired with a coherent annotation scheme, a composite severity index, and an annotation guideline designed to evolve through clinical reliability review. Throughout, the protocol maintains a clear boundary between research signal and clinical decision-making, consistent with responsible NLP practice in sensitive domains. The immediate next steps are the completion of online data collection, finalization of the annotation guidelines to version 2 following blind clinical co-annotation, and execution of the planned exploratory psycholinguistic and embedding analyses. Upon completion of these phases, replication of the protocol across additional student samples and, in subsequent work, across more diverse populations will be pursued, with the longer-term aim of contributing a governance blueprint and open methodological resource for multi-site Italian suicidology research.

References

Arowosegbe, A., & Oyelade, T. (2023). Application of natural language processing (NLP) in detecting and preventing suicide ideation: A systematic review. *International Journal of Environmental Research and Public Health, 20*(2), 1514.

Campos, A. I., Van Velzen, L. S., Veltman, D. J., Pozzi, E., Ambrogi, S., Ballard, E. D., Banaj, N., Başgöze, Z., Bellow, S., Benedetti, F., Bollettini, I., Brosch, K., Canales-Rodríguez, E. J., Clarke-Rubright, E. K., Colic, L., Connolly, C. G., Courtet, P., Cullen, K. R., Dannlowski, U., Dauvermann, M. R., … Rentería, M. E. (2023). Concurrent validity and reliability of suicide risk assessment instruments: A meta-analysis of 20 instruments across 27 international cohorts. *Neuropsychology, 37*(3), 315–329.

Chandler, C., Foltz, P. W., & Elvevåg, B. (2022). Improving the applicability of AI for psychiatric applications through human-in-the-loop methodologies. *Schizophrenia Bulletin, 48*(5), 949–957.

Fazel, S., & Runeson, B. (2020). Suicide. *New England Journal of Medicine, 382*, 266–274.

Khazem, L. R., Anestis, M. D., & Freeman, A. J. (2021). Under-reporting on the MMPI-2-RF extends to extra-test measures of suicide risk. *Psychological Assessment, 33*(8), 789–794.

Klonsky, E. D., & May, A. M. (2015). The Three-Step Theory (3ST): A new theory of suicide rooted in the 'Ideation-to-Action' framework. *International Journal of Cognitive Therapy, 8*(2), 114–129.

Landi, G., Boccolini, G., Chinali, G., et al. (2025). Italian validation of the Orbach and Mikulincer Mental Pain Scale-8 (OMMP-8). *Annals of General Psychiatry, 24*, 46.

Li, T. M. H., Chen, J., Law, F. O. C., et al. (2023). Detection of suicidal ideation in clinical interviews for depression using natural language processing and machine learning: Cross-sectional study. *JMIR Medical Informatics, 11*, e50221.

Nock, M. K., Borges, G., Bromet, E. J., et al. (2008). Suicide and suicidal behavior. *Epidemiologic Reviews, 30*, 133–154.

O'Connor, R. C., & Nock, M. K. (2014). The psychology of suicidal behaviour. *The Lancet Psychiatry, 1*(1), 73–85.

Pigoni, A., Delvecchio, G., Madonna, D., & Brambilla, P. (2024). Machine learning and the prediction of suicide in psychiatric populations: A systematic review. *Translational Psychiatry, 14*, 136.

Posner, K., Brown, G. K., Stanley, B., et al. (2011). The Columbia–Suicide Severity Rating Scale: Initial validity and internal consistency findings from three multisite studies with adolescents and adults. *American Journal of Psychiatry, 168*(12), 1266–1277.

Posner, K., Oquendo, M. A., Gould, M., Stanley, B., & Davies, M. (2007). Columbia Classification Algorithm of Suicide Assessment (C-CASA): Classification of suicidal events in the FDA's pediatric suicidal risk analysis of antidepressants. *American Journal of Psychiatry, 164*(7), 1035–1043.

Taussczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology, 29*(1), 24–54.

Tossani, E., Ricci Garotti, M. G., Mikulincer, M., et al. (2021). Psychometric evaluation of the Italian version of the Orbach & Mikulincer Mental Pain Scale in a non-clinical sample. *Current Psychology, 40*, 1903–1910. https://doi.org/10.1007/s12144-019-0128-4