



PREDICTING SUICIDE RISK THROUGH MACHINE LEARNING- BASED ANALYSIS OF PATIENT NARRATIVES AND DIGITAL BEHAVIORAL MARKERS IN CLINICAL PSYCHOLOGY SETTINGS

Md. Akbar Hossain¹; Farzana Ahmed²;

[1]. *Master of Science in Clinical Psychology, University of Dhaka, Dhaka, Bangladesh;*
Email: md.akbarh@yahoo.com

[2]. *Masters in School Psychology*
PhD candidate, Department of School Psychology, Special Education and Sociology.
Indiana University of Pennsylvania, USA
Email: farzanaahmed1984@gmail.com

Doi: [10.63125/mqty9n77](https://doi.org/10.63125/mqty9n77)

Received: 12 September 2023; **Revised:** 20 October 2023; **Accepted:** 20 November 2023; **Published:** 25 December 2023

Abstract

This study addresses the persistent problem that conventional suicide risk screening and clinician judgment often struggle to achieve operationally useful precision, especially when risk signals are multidimensional and base rates are imbalanced; therefore, the purpose was to develop and benchmark an explainable, data-driven risk stratification approach that integrates psychometric and clinical indicators to classify individuals into high vs low or moderate suicide-risk tiers within a bounded case context using a quantitative, cross-sectional, case-based design. The sample comprised 320 clinical cases drawn from an enterprise service setting that included both outpatient (61.3%) and emergency or acute contacts (38.7%), with a mean age of 29.8 (SD 8.7) and a high-risk prevalence of 27.8% (n = 89). Key variables included five-point Likert composite predictors capturing distress severity, hopelessness, perceived burdensomeness, psychosocial strain, sleep disturbance, impulsivity, perceived social support, and coping capacity, alongside binary clinical indicators for prior suicide attempt history and substance-use concern. Descriptively, the cohort showed elevated distress and cognitive burden (for example distress M = 3.62, SD = 0.78; hopelessness M = 3.41, SD = 0.83) with comparatively lower protective resources (social support M = 2.64, SD = 0.92; coping M = 2.71, SD = 0.88). All Likert constructs demonstrated acceptable to high reliability (Cronbach's α range 0.80 to 0.89). Risk classification correlated positively with distress ($r = .49$) and hopelessness ($r = .44$) and negatively with social support ($r = -.40$) and coping ($r = -.35$), supporting coherent construct behavior prior to modeling. In regression, distress (OR = 2.18, $p < .001$) and hopelessness (OR = 1.71, $p < .001$) increased odds of high-risk classification while social support reduced odds (OR = 0.63, $p < .001$), establishing a strong interpretable benchmark (Nagelkerke $R^2 = .41$; AUC = .82). Headline performance results showed that gradient boosting outperformed the baseline with AUC = .88, sensitivity = .81, specificity = .78, precision = .56, F1 = .66, and improved calibration (Brier score = .14); importantly, it captured 46.1% of all high-risk cases within the top 10% highest-risk score band, indicating high-yield triage potential. Explainability highlighted clinically interpretable drivers, led by distress severity and hopelessness, with protective deficits such as low social support also ranking highly. These findings imply that enterprise clinical workflows can benefit from a transparent, construct-grounded stratification pipeline that improves sensitivity and concentrates risk into actionable high-risk quantiles while preserving interpretability for decision support and resource-aware triage planning.

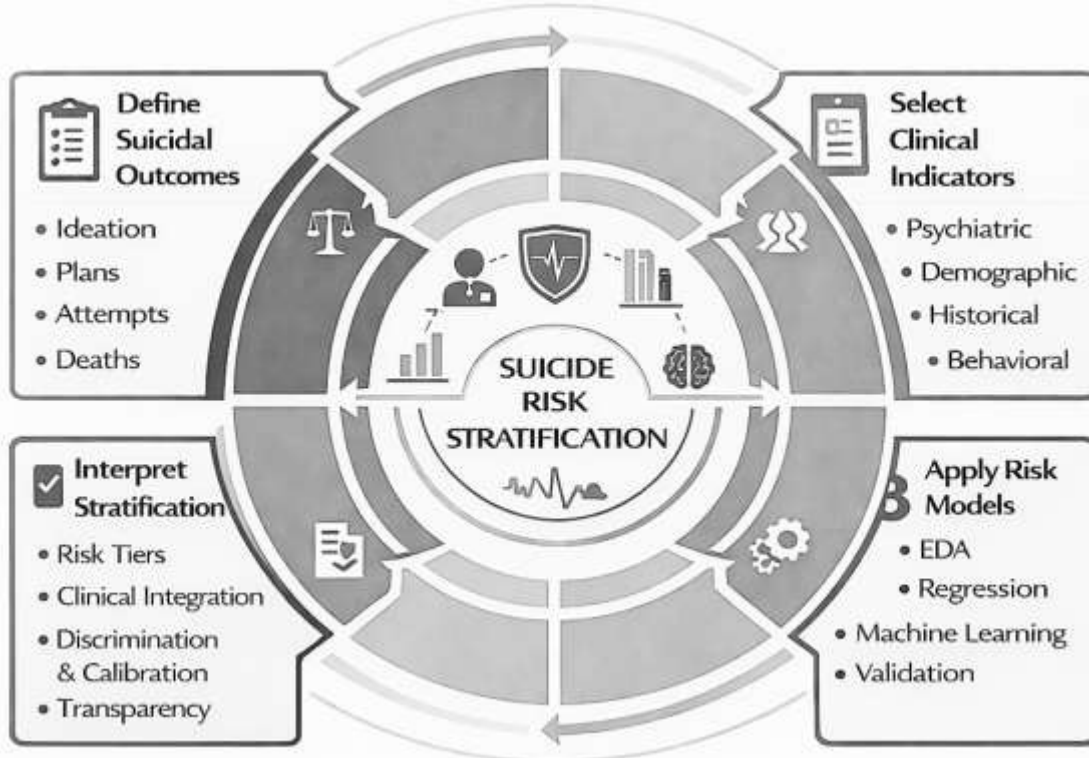
Keywords

Suicide Risk Stratification; Machine Learning; Logistic Regression; Explainable AI (SHAP); Psychometric Likert Constructs;

INTRODUCTION

Suicide-related outcomes are commonly defined along a continuum that includes suicidal ideation (thoughts about ending one’s life), suicide plans (formulated intent and method), suicide attempts (self-injurious acts carried out with some intent to die), and suicide death (fatal self-injury with intent to die) (Alonso et al., 2008). Clinical and public-health research also distinguishes self-harm from suicide attempt by emphasizing that self-harm can occur with mixed or unclear intent, whereas attempts are anchored in intent to die (Christensen et al., 2013). Within health systems, the term risk stratification refers to categorizing patients into clinically meaningful risk tiers based on observed data – typically demographic, diagnostic, historical, and behavioral indicators – so that the distribution of outcomes can be examined across strata (Belsher et al., 2019). Modern suicide science treats suicidal thoughts and behaviors as outcomes that emerge from interacting vulnerabilities and stressors rather than from a single cause, a view supported by cross-national epidemiologic patterns showing that prevalence and correlates vary across settings while retaining certain robust associations such as prior suicidal behavior and psychiatric symptoms (Kessler et al., 2005). Large-scale epidemiologic synthesis further frames suicidal behavior as a population-level phenomenon shaped by individual, clinical, and contextual factors that can be measured and modeled (Joiner, 2010). In this context, predictive modeling in clinical populations operates on operational definitions encoded in electronic records and structured assessments, which makes the quality of definitions and labeling central to the credibility of any stratification approach (Barak-Corren et al., 2016). The definitional boundary between ideation and attempt is also practically important because ideation is more common than attempt, and the statistical properties of prediction differ substantially as outcome base rates change (Franklin et al., 2017; Ribeiro et al., 2016). For quantitative research that combines clinical data and psychometric measurement, these distinctions guide how constructs are measured, which outcomes are modeled, and how model performance is interpreted across patient groups (Cha et al., 2008).

Figure 1: Measurement-to-Model Systems Framework for Suicide Risk Prediction



The international significance of suicide prevention is reflected in evidence that many strategies target multi-level determinants, including clinical identification, follow-up care, and system-level coordination, yet effectiveness varies across intervention types and settings (Pozo-Banos et al., 2018). Clinical populations – patients in emergency departments, inpatient psychiatry, and outpatient mental health services – represent concentrated risk contexts where near-term outcomes such as attempts can

cluster within short windows, increasing both the urgency and complexity of accurate identification (Franklin et al., 2017). Epidemiologic work describing temporal patterns of ideation, planning, and attempts highlights that shifts in population-level prevalence occur alongside persistent clinical challenges in pinpointing who transitions from thoughts to action (Hawton et al., 2012). Adolescent and young-adult research further emphasizes that self-harm and suicidal behavior have distinct risk architectures across developmental stages, while clinical contact frequently occurs in acute settings that are not primarily designed for long-form risk assessment (Nock et al., 2022). Short-horizon prediction is frequently discussed as a practical need because a substantial portion of individuals who later attempt suicide interact with health services in the months surrounding the event, placing emphasis on measurable signals present during care episodes (Mulder et al., 2016). At the same time, quantitative reviews have documented that many established risk factors are statistically associated with suicidal behavior but yield modest predictive utility when used in isolation, which motivates integrated modeling approaches that combine signals across domains (Mann et al., 2005). This predictive gap is also described as a measurement and evaluation problem: rare outcomes, heterogeneous pathways, and variable documentation practices can reduce positive predictive values even when associations are reliable (Klonsky & May, 2015). Consequently, international clinical relevance is tied to whether prediction methods demonstrate stable discrimination and calibration across care settings and patient subgroups rather than merely identifying correlates in retrospective analyses (Trujillo et al., 2022).

Contemporary theory organizes suicide risk through models that differentiate the emergence of ideation from the transition to suicidal behavior, strengthening the conceptual basis for stratification constructs used in empirical studies. The interpersonal theory of suicide posits that suicidal desire is driven by thwarted belongingness and perceived burdensomeness, while lethal or near-lethal behavior depends on capability for suicide, a state that can be acquired through habituation to pain and fear (Simon et al., 2018). Empirical tests in community samples provide quantitative support for key interactions proposed by the theory, strengthening its role as a framework for selecting constructs relevant to ideation and attempt risk (Ribeiro et al., 2019). In parallel, the ideation-to-action framing is advanced by the three-step theory (3ST), which structures suicidal ideation as arising from pain and hopelessness, intensifying when connectedness is low, and progressing to attempt when capability and practical factors enable action (Ribeiro et al., 2016). The integrated motivational-volitional (IMV) model further distinguishes motivational processes that generate suicidal ideation from volitional moderators that govern enactment, emphasizing that different predictors may be required for different stages (Karmakar et al., 2016). These models converge on a key implication for construct measurement in quantitative studies: correlates of ideation are not automatically correlates of attempts, and predictors of attempts often involve history, impulsivity-related processes, exposure, and access-related enablers that differ from ideational predictors (Ilgen et al., 2009). Theoretical specification also clarifies how psychometric indicators can be mapped to latent constructs, supporting multi-construct survey design in clinical settings where self-report can capture subjective states not fully represented in administrative codes (Kessler et al., 2015). When a study's objective is predictive risk stratification, theory therefore functions as a principled mechanism for defining predictor families and interpreting why certain variable sets improve discrimination for attempts while others primarily track ideation (Simon et al., 2013).

Clinical suicide risk assessment has long relied on structured instruments and clinician judgment, yet quantitative evidence repeatedly highlights limits in predictive accuracy, particularly for rare outcomes such as suicide death and for short time horizons where base rates remain low. Meta-analytic work demonstrates that, across decades of research, most risk factors and clinical scales show limited ability to predict future suicidal behavior with high precision, often performing only modestly better than chance for prospective outcomes (Rossom et al., 2017). A systematic review and meta-analysis focusing on positive predictive values for risk scales underscores that even widely used instruments can yield low PPV in real-world settings, reflecting both base-rate constraints and heterogeneity in patient trajectories (O'Connor & Kirtley, 2018). Ethical and scientific critiques argue that risk categorization can be overinterpreted when statistical uncertainty and measurement error are not transparently addressed, reinforcing the need for careful evaluation of predictive claims (Orden et al., 2010). At the same time, evidence supports that clinical prediction improves when multiple data sources are

combined, including patient self-report, clinician assessment, and electronic records, suggesting that no single modality captures the full signal relevant to near-term attempts (Glenn & Nock, 2014). Research using routinely collected measures also demonstrates that brief indicators embedded in standard care can be informative for subsequent suicidal behavior, with associations observed between self-reported suicidal ideation items and later attempts across large health systems (Carter et al., 2017). These results collectively sharpen the methodological focus on evaluation metrics such as discrimination, calibration, and prevalence-stratified performance, because a model can produce acceptable rank-order discrimination while still being poorly calibrated for operational deployment across settings (Walsh et al., 2021). Within this measurement landscape, predictive modeling becomes a comparative task: alternative model classes and variable sets are assessed against clinician-only baselines and simple scales to quantify incremental value under clinically realistic constraints (Walsh et al., 2017).

The use of electronic health records has expanded the empirical basis for suicide risk prediction by providing large, longitudinal, and heterogeneous data streams that can support data mining and machine learning approaches. Early work applying exploratory data mining to clinical populations with depression identified high-risk subgroups through interaction patterns among routinely documented variables, illustrating how risk can concentrate in specific profiles rather than distributing evenly across a diagnosis group (Christensen et al., 2013; Banos et al., 2018). Subsequent studies used structured EHR data to generate individualized risk estimates for suicidal behavior, emphasizing that longitudinal records offer temporally ordered signals—prior diagnoses, utilization patterns, comorbidity clusters, and prior self-harm codes—that can be modeled for near-term outcomes (Hawton et al., 2012). Work in large systems also focused on high-risk windows such as the months after psychiatric hospitalization, where modeling aims to identify subsets with especially elevated risk using predictors available in health records (del Pozo-Banos et al., 2018). Other approaches incorporate non-psychiatric clinical history, showing that physical-illness information and ICD-coded trajectories can improve prediction compared with routine clinical checklists, highlighting the multi-morbidity reality of many clinical populations (Ribeiro et al., 2016). More recent prospective validation studies demonstrate that EHR-based models can be executed at scale in operational systems, enabling assessment of discrimination and calibration across care sites such as emergency departments, inpatient units, and medical settings (Walsh et al., 2021). Evidence also indicates that performance can vary substantially across clinical contexts, especially when a model trained in one setting is applied in another, which elevates the importance of case-context specification for any risk stratification study framed as clinical and case-study-based (Karmakar et al., 2016). Within this EHR-driven landscape, predictive suicide risk stratification is increasingly approached as a multidisciplinary measurement problem involving clinical definition, data quality, and rigorous comparative evaluation across algorithms and variable modalities (Belsher et al., 2019).

Machine learning approaches explicitly target nonlinearities, high-dimensional interactions, and complex temporal proxies that are difficult to capture with conventional regression alone, while still requiring careful validation and transparent benchmarking. In mental health populations, studies comparing machine learning classifiers against simpler baselines report improvements in discrimination for suicidal behavior prediction when multiple features and flexible learning strategies are applied (Klonsky & May, 2015). Work using neural networks with routine health records demonstrates feasibility for identifying higher-risk individuals within linked datasets, providing evidence that nonlinear models can leverage routine variables to produce meaningful stratification signals in case-control designs (Hawton et al., 2012). Comparative studies focusing on short time horizons illustrate how ensemble approaches and hybrid feature sets can support near-term attempt prediction in acute-care samples, with patient self-report often adding substantial incremental predictive value relative to clinician-only judgments (O'Connor & Kirtley, 2018). Studies that examine imminent risk processes similarly show that machine learning applied to repeated measures and longitudinal indicators can enhance prediction of near-term suicidal thoughts and nonfatal attempts, reinforcing the value of time-sensitive modeling (Trujillo et al., 2022). In parallel, systematic reviews in psychiatric research synthesize findings across model families and datasets, noting that reported performance varies by outcome type, data source, and study design quality, which places emphasis on

transparent reporting and risk-of-bias appraisal (Ilgen et al., 2009). Meta-analytic synthesis comparing “small data” (psychometric instruments) and “big data” (EHR and other machine-interpretable sources) further clarifies that predictive strength depends on how modalities are combined and evaluated, rather than on algorithm choice alone (Belsher et al., 2019). This evidence base supports an approach to risk stratification that treats machine learning models as empirical tools whose value is established through rigorous comparisons, validated performance metrics, and explicit linkage between measured constructs and theoretical framing (Franklin et al., 2017).

Explainability and clinical interpretability have become integral to suicide risk modeling because clinical decision-making requires traceable rationale for stratification outputs, particularly when models integrate heterogeneous information sources. Studies combining clinician assessments, brief self-report, and EHR-derived scores demonstrate that the most accurate models often rely on compact feature sets drawn from multiple modalities, which creates a practical need to understand which predictors drive classification within high-risk strata (Karmakar et al., 2016). Research using routine depression and ideation items indicates that self-reported suicidal ideation markers correlate with later attempts across large cohorts, making patient-reported constructs salient inputs in quantitative risk models that also include administrative history and utilization proxies (Mann et al., 2005). In EHR-based systems, model performance is also evaluated through workload-relevant metrics (e.g., numbers needed to screen within top-risk quantiles), which links statistical outputs to how stratification partitions risk in real clinical flow (Carter et al., 2017). The scientific critique of risk assessment underscores that risk categorization gains credibility when uncertainty and calibration are addressed, which aligns with explainability practices that reveal predictor contributions and highlight where models may generalize poorly across sites or subgroups (Nock et al., 2022). Systematic synthesis of machine learning suicide prediction studies emphasizes that interpretability and reporting transparency influence the trustworthiness and comparability of findings, particularly when different studies operationalize outcomes differently or rely on variable-quality labels (Carter et al., 2017). Within clinical populations, risk stratification therefore becomes a structured empirical task that integrates theory-driven constructs (e.g., ideation-to-action mechanisms), validated psychometric measures, and record-based longitudinal indicators into a model comparison framework grounded in descriptive statistics, correlation structure, and regression-based baselines alongside machine learning classifiers (Ilgen et al., 2009). This framing situates predictive suicide risk stratification as a quantitative methodology focused on measurable constructs, clinically anchored outcomes, and transparent evaluation across competing modeling approaches (Ribeiro et al., 2019).

This study is designed to achieve a set of tightly aligned objectives that operationalize predictive suicide risk stratification as a measurable, data-driven task within a clearly defined clinical case setting. The first objective is to specify and operationalize suicide risk in a clinically meaningful way by defining the outcome category structure used for stratification, whether represented as binary risk status or tiered risk levels, and by establishing consistent rules for labeling and scoring that can be applied uniformly across participants. The second objective is to develop a structured measurement model for the predictor space by translating clinically relevant and psychosocial dimensions into quantifiable variables, including composite constructs derived from Likert five-point scale responses and complementary indicators captured through case-context clinical information, ensuring that each construct is measurable, internally consistent, and analytically suitable for multivariable modeling. The third objective is to describe the study population and case-study environment through systematic profiling of participant characteristics and construct distributions so that the empirical context of risk stratification is transparent and statistically interpretable. The fourth objective is to test association patterns among key constructs through correlation analysis and to examine the predictive contribution of each predictor group using regression modeling as a baseline benchmark, producing interpretable parameter estimates that clarify which variables are statistically linked to risk stratification outcomes. The fifth objective is to build, tune, and compare multiple machine learning algorithms for suicide risk prediction using a consistent training and evaluation procedure, allowing model performance to be assessed using standardized classification metrics and enabling a clear comparison between traditional regression and machine-learning-based approaches. The sixth objective is to generate explainability outputs that identify the most influential predictors driving model decisions and to present these

outputs in a structured manner that supports transparent interpretation of the stratification logic. The seventh objective is to synthesize the findings from descriptive statistics, correlation patterns, regression results, and machine learning evaluations into an integrated empirical account that directly answers the research questions and provides an objective-aligned representation of how predictive suicide risk stratification performs within the selected clinical case context.

LITERATURE REVIEW

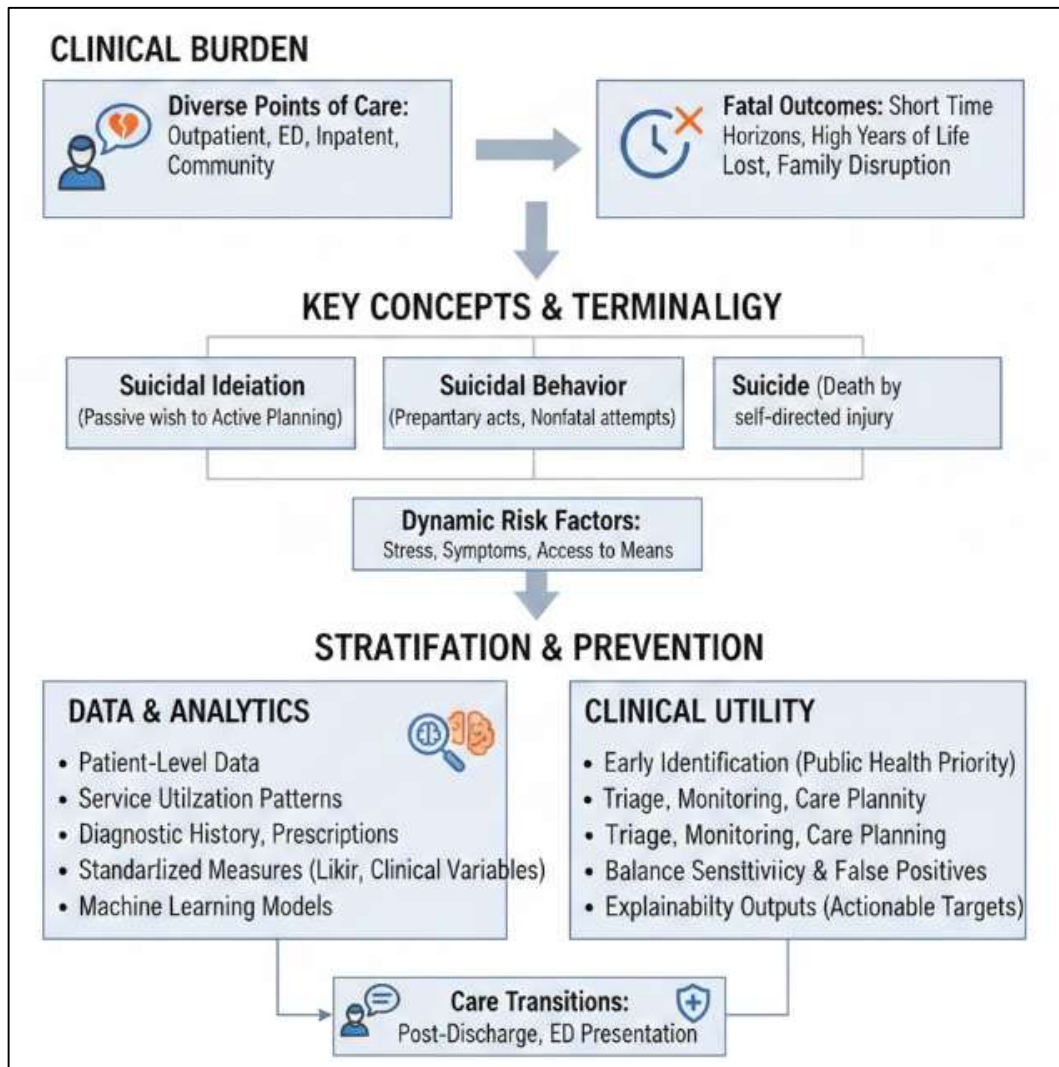
The literature on predictive suicide risk stratification in clinical populations sits at the intersection of suicidology, clinical assessment science, health informatics, and machine learning, and it has expanded rapidly as health systems seek more reliable ways to identify individuals at elevated risk during routine care. Across this body of work, suicide risk is commonly treated as a multi-determined clinical outcome that reflects the interaction of psychiatric symptoms, prior self-harm history, behavioral and substance-related factors, psychosocial stressors, and protective resources such as connectedness and perceived support. Researchers have long examined these factors through epidemiologic and clinical studies, yet a central theme in the evidence base is that statistically significant correlates do not automatically translate into strong prospective prediction, particularly when outcomes are rare and heterogeneous across settings. This gap has shaped two complementary directions in the literature: one stream focuses on refining clinical assessment instruments and short-horizon indicators that can be implemented consistently across services, while another stream emphasizes computational approaches that integrate multiple data modalities to improve discrimination and calibration at the individual level. With the widespread adoption of electronic health records, predictive modeling studies increasingly use structured clinical histories, diagnostic and utilization patterns, and patient-reported measures to generate individualized risk estimates and stratify patients into actionable risk tiers for monitoring and intervention. At the same time, methodological discussions in the literature highlight that model performance must be evaluated with clinically meaningful metrics, that generalizability across sites is not guaranteed, and that interpretability is essential when predictions may influence safety planning and care escalation decisions. The resulting research landscape therefore includes comparative evaluations of regression baselines and machine learning classifiers, validation studies in real-world clinical workflows, and interpretability-focused analyses that aim to clarify which predictors drive risk stratification outputs. In addition, the literature increasingly emphasizes the need to ground predictive features in theory-informed constructs and to link empirical predictor importance to established psychological mechanisms that differentiate the emergence of suicidal ideation from the transition to suicidal behavior. Taken together, existing studies provide a structured foundation for understanding which variables consistently signal risk, how algorithmic approaches can be evaluated responsibly, and why transparent, context-specific model development remains central to credible suicide risk stratification in clinical populations.

Suicide Risk in Clinical Populations

Suicide risk stratification research is grounded in the clinical reality that suicidal thoughts and behaviors appear across diverse points of care and can culminate in fatal outcomes within short time horizons. In clinical terminology, suicidal ideation refers to thoughts of ending one's life that range from passive wishes to active planning; suicidal behavior includes preparatory acts and nonfatal attempts; and suicide denotes death caused by self-directed injurious behavior. These phenomena are observed in outpatient clinics, emergency departments, inpatient units, and community-facing services, and they are clinically important because their presentation can change rapidly in response to stress exposure, symptom fluctuations, and access to means. At a systems level, suicide imposes a substantial international burden in years of life lost, family disruption, and downstream health and social costs, which makes early identification a public health priority as well as a patient-safety concern. Yet the pathway from distress to suicidal action is not uniform: suicidal behavior differs by age group, sex, region, and sociocultural conditions, and it is associated with varied constellations of risk factors, underscoring etiological heterogeneity and the limits of one-size-fits-all assessment. Clinical practice therefore faces a dual problem – many risk factors are common in treatment settings, while suicide and near-term attempts are comparatively rare events – so clinicians must make decisions under uncertainty, often with incomplete information. Within this context, stratification frameworks aim to combine signals from symptoms, history, and context to estimate relative risk levels that can guide

triage, monitoring, and care planning. The literature also emphasizes that prediction performance must be evaluated in ways that reflect clinical utility, including the balance between sensitivity for high-risk cases and the burden of false positives in routine care (Turecki & Brent, 2016). This motivates structured measures and standardized variables for consistent analysis across participants, enabling comparable rigorous evaluation of statistical and machine-learning models.

Figure 2: Suicide Risk Stratification Framework



A key reason clinical populations are central to suicide prevention is that many individuals who later die by suicide have recent contact with health services, creating observable windows in which risk can be identified and managed. Large multi-system analyses show that most suicide decedents had at least one encounter in the year preceding death, and that contacts often occur in nonpsychiatric settings such as primary care and medical specialty clinics. This distribution matters for stratification because risk signals can be embedded in routine utilization patterns, diagnostic histories, prescriptions, and brief disclosures that may never reach specialty mental health notes. In a large U.S. health system study, many individuals who later died by suicide had visits close to death without a recorded mental health diagnosis, indicating substantial missed opportunities for detection (Ahmedani et al., 2014; Jinnat & Kamrul, 2021). Evidence from publicly funded systems similarly documents high rates of contact prior to suicide while showing systematic differences in where care is received by age and sex. A Swedish medical-record review covering two years prior to suicide found that contact in the months and weeks before death was common, with younger individuals more often seen in psychiatric services and older adults more often using primary or specialized somatic care (Bergqvist et al., 2022; Zulqarnain & Subrato, 2021). For quantitative research design, these patterns support defining the unit of analysis at

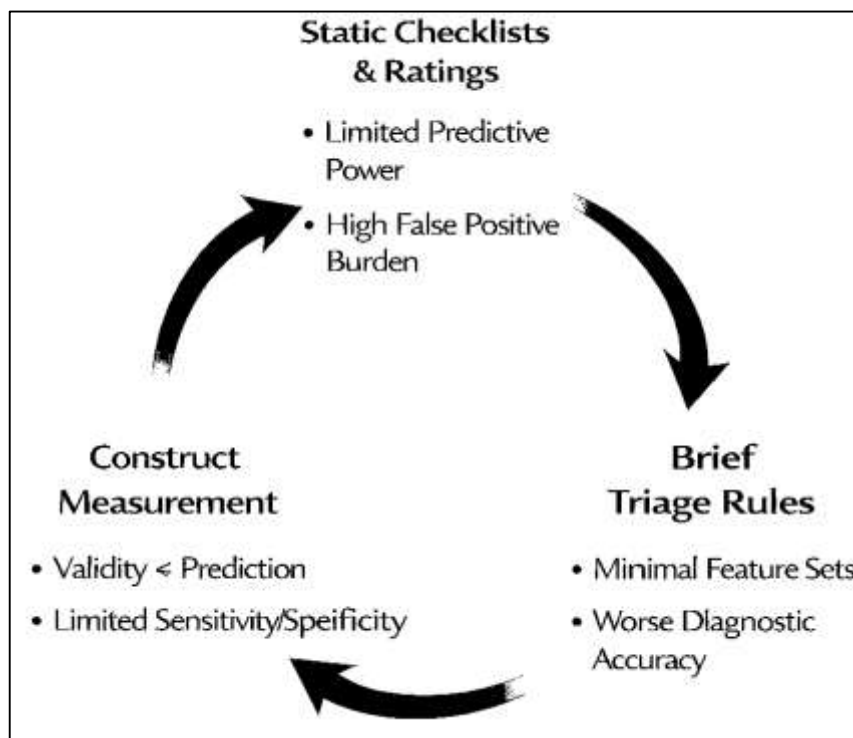
the patient level and measuring multiple domains that map onto typical clinical documentation: symptom severity, prior self-harm history, service utilization, comorbidity indicators, and psychosocial stress exposure. In a case-study setting, integrating standardized Likert-scale constructs with extracted clinical variables enables hypothesis testing through correlation and regression, while also creating a consistent feature set for comparing traditional statistical models with machine-learning classifiers. It also encourages explicit handling of missing data and measurement error, which otherwise distort prevalence estimates and model calibration substantially.

Clinical burden is concentrated at high-risk transition points where instability is elevated and continuity of care can be fragmented, particularly the period immediately after discharge from psychiatric facilities and episodes of emergency department (ED) presentation for ideation or self-harm. The postdischarge interval reflects intersecting pressures that can amplify risk, including abrupt changes in support, persistent symptoms, medication adjustments, and delays in follow-up. Meta-analytic synthesis shows that suicide rates after discharge are markedly elevated compared with general population baselines, with the greatest risk concentrated early after discharge (Chung et al., 2017). In parallel, the ED is a frequent contact point for patients in acute distress, but detection depends on whether suicidal intent is assessed and recorded during routine workflow. A large multi-site interrupted time-series evaluation found that implementing universal ED suicide risk screening was feasible and was associated with substantially higher documented detection of self-harm ideation/behavior in records (Boudreaux et al., 2016). Together, these findings explain why predictive modeling studies often focus on care transitions and emergency presentations: base rates of suicidal behavior are higher, signals are temporally proximal to outcomes, and triage decisions are immediate. For quantitative case-study designs, transition-focused sampling can increase statistical power while remaining clinically realistic, provided that eligibility criteria, outcome definitions, and follow-up windows are specified with precision. Analytically, combining descriptive profiles of transition groups with correlations among psychosocial constructs and multivariable regression offers a transparent baseline for inference, while machine-learning comparisons can test whether non-linear interactions and high-dimensional features improve discrimination beyond conventional models. Because clinical deployment requires trust, pairing performance metrics with explainability outputs helps clinicians relate model signals to actionable assessment targets. These settings also generate rich longitudinal traces – repeat visits, prior attempts, and comorbidity codes – that can be summarized into predictors reflecting chronic vulnerability and acute change across the observation window for risk estimation.

Traditional Suicide Risk Assessment Approaches

Traditional suicide risk assessment in clinical practice has typically been organized around clinician-led interviewing, structured checklists of known risk factors, and standardized rating instruments that attempt to quantify ideation severity, suicidal behavior history, and proximal warning signs. In many settings, assessment begins with a clinical interview that elicits current ideation, intent, plan specificity, access to means, and protective factors, followed by documentation of psychiatric diagnoses, substance use, prior attempts, recent losses, and psychosocial stressors. To standardize this process, widely adopted instruments have been developed to support consistent terminology and scoring across clinicians and sites. A prominent example is the Columbia–Suicide Severity Rating Scale (C-SSRS), which formalizes the assessment of suicidal ideation and behavior using graded categories of severity and intensity and has been validated across multisite samples, supporting its role as a structured way to capture ideation and behavioral history in research and practice (Uddin et al., 2022; Posner et al., 2011). In parallel, service contexts such as emergency departments have used brief decision rules or triage-oriented tools intended to support disposition and follow-up decisions under time constraints. These tools are often designed to be quick to administer and based on a small set of binary risk markers that are easy to obtain in acute care, where clinicians must make rapid judgments with limited collateral information. One influential example is the Manchester Self-Harm Rule, derived from emergency presentations after self-harm to identify individuals at higher likelihood of repetition or subsequent suicide within a short follow-up window (Cooper et al., 2006; Akbar & Sharmin, 2022). Taken together, these approaches reflect a long-standing clinical effort to translate complex, multidimensional suicide risk into actionable categories using structured questions, standardized scoring, and streamlined rules that fit within routine workflow and documentation demands.

Figure 3: Scale- and Rule-Based Suicide Risk Assessment



A consistent finding across the literature, however, is that traditional tools and scales frequently demonstrate constrained predictive performance when evaluated as instruments for forecasting future suicidal behavior rather than documenting current risk status. This limitation is most visible when tools are tested for their ability to predict outcomes such as repeat self-harm or suicide attempts over short horizons, where the base rate remains low even in high-risk groups and where heterogeneity of pathways reduces the stability of any single predictor profile. In a multicentre prospective cohort study of adults referred to liaison psychiatry services following self-harm, several commonly used risk scales – including SAD PERSONS and rule-based tools – showed wide variation in diagnostic accuracy, with some scales performing near chance-level discrimination and demonstrating poor balance between sensitivity and specificity (Foysal & Subrato, 2022; Quinlivan et al., 2017). Evidence synthesis also indicates that many instruments are supported by limited numbers of high-quality studies and that even the better-studied scales can fail to meet commonly desired thresholds for diagnostic accuracy when evaluated with rigorous methods. A systematic review that appraised multiple suicide risk instruments and assessed certainty of evidence concluded that, among tools with enough data for pooled analysis, none met criteria for sufficient diagnostic accuracy across outcomes, while some triage rules displayed high sensitivity accompanied by very low specificity, limiting practical value when used as stand-alone decision aids (Runeson et al., 2017; Zulqarnain, 2022). These findings illustrate a recurring issue: traditional scales may be helpful for structuring documentation and supporting clinical conversations, yet their statistical properties often do not support using a single score as a reliable predictor of future suicidal behavior, especially when clinicians require both high sensitivity for safety and acceptable specificity to avoid overwhelming services with false positives.

Additional limitations emerge from how traditional assessment approaches operationalize psychological constructs and how scale cut-offs behave across heterogeneous clinical populations. Many instruments were designed to measure constructs such as hopelessness, intent, or severity of ideation, but the translation from construct measurement to outcome prediction is not straightforward, particularly because suicidal behavior can be influenced by rapidly changing context, access to means, impulsivity, and situational stressors that are difficult to capture with static questionnaires. For example, hopelessness has long been treated as a salient cognitive risk marker, and the Beck Hopelessness Scale (BHS) remains widely used in psychiatric contexts. A meta-analysis examining the

predictive accuracy of the BHS found that, for suicide, pooled sensitivity could be moderate while specificity was low, implying that many individuals scoring above conventional cut-offs would not go on to die by suicide, and that the scale alone provides limited precision for forecasting rare outcomes (McMillan et al., 2007). Similar trade-offs appear in brief triage rules developed for emergency presentations: decision rules can be tuned to capture most repeat events, yet that often results in very low specificity, generating a large “high risk” group that is difficult to manage operationally. This measurement reality is reinforced by prospective cohort evidence indicating that commonly used scales may not outperform, and sometimes perform worse than, clinician and patient global risk estimates when predicting repetition after self-harm (Quinlivan et al., 2017). Moreover, routine clinical practice introduces variability in administration, disclosure, and documentation, meaning that scale scores can reflect not only patient status but also differences in clinical approach and setting constraints. As a result, traditional tools are often strongest as structured aids for assessment consistency rather than as reliable predictive engines, which explains why many contemporary studies benchmark them against multivariable regression and machine-learning models that can integrate broader feature sets and capture interactions that static scales cannot represent directly (Posner et al., 2011).

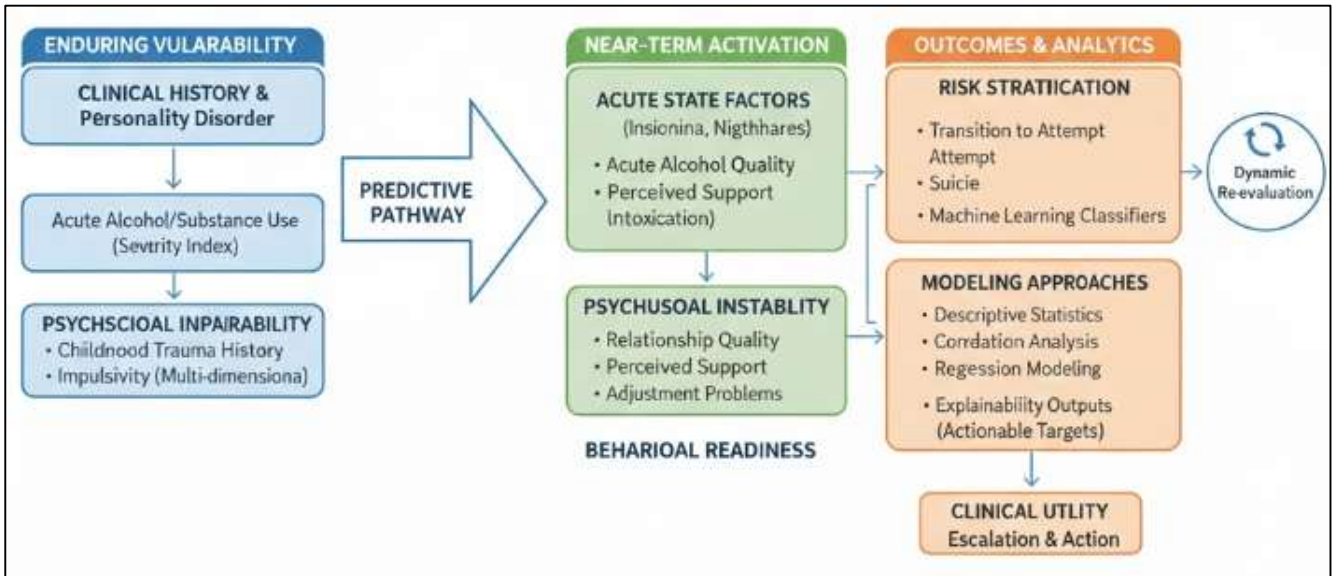
Key Risk Factors for Suicide Risk Stratification

Suicide risk stratification in clinical populations depends on identifying a set of predictor constructs that represent both enduring vulnerability and near-term activation, because suicidal behavior is rarely explained by a single variable in isolation. A widely supported clinical pattern is that risk concentrates among people who already present with suicidal ideation, and the analytic challenge shifts to differentiating which ideators are most likely to transition to an attempt. Longitudinal evidence shows that, within depressed ideators, certain diagnostic and interpersonal features can outperform many traditionally cited correlates in predicting later attempts. In a 10-year follow-up study of depressed suicide ideators, comorbid Cluster B personality disorder emerged as a robust unique predictor of subsequent attempts, while several commonly discussed correlates showed minimal incremental predictive value once stronger predictors were considered (May et al., 2012). This line of evidence is important for quantitative model building because it implies that stratification should not rely only on symptom severity scores; it should also represent clinically structured risk domains such as personality pathology, anxiety comorbidity, substance misuse history, and interpersonal functioning. These domains map naturally into measurable constructs in cross-sectional case-study designs: clinical history variables can be treated as categorical indicators (e.g., presence/absence of diagnostic domains), while psychosocial functioning can be captured using Likert-based composite measures that quantify relationship quality, perceived support, adjustment problems, and distress intensity. For descriptive statistics, these constructs allow clear profiling of clinical burden and risk distribution; for correlation analysis, they allow examination of how psychosocial instability and comorbidity cluster together; and for regression modeling, they permit estimation of incremental predictive contribution when entering predictors in conceptually meaningful blocks. When the primary aim is predictive stratification, these constructs also support model interpretability because clinicians can understand risk tiers framed around familiar domains (comorbidity, personality features, social adjustment) rather than opaque feature lists.

Childhood trauma is repeatedly linked to later suicidal behavior, and longitudinal synthesis indicates that early adverse exposure increases the likelihood of lifetime attempts across follow-up periods, supporting its role as a background vulnerability construct that can be measured retrospectively and incorporated as a stable predictor domain (Abdul, 2023; Bastos et al., 2017). At the acute end of the continuum, alcohol involvement is often conceptualized as both a chronic vulnerability marker (use disorder history) and a proximal disinhibiting factor that can precipitate attempts through impaired judgment and increased capability to act in crisis. Meta-analytic evidence focusing on acute alcohol use indicates elevated odds of suicide attempt during episodes of intoxication or recent consumption, supporting inclusion of acute substance-use indicators (recent use, binge episodes, intoxication during crisis) as near-term risk signals when available in case data (Hammad & Mohiul, 2023; Rossow & colleagues, 2016). In quantitative operationalization, trauma history can be represented as a Likert-based severity index (frequency/intensity of adverse experiences) or categorized into exposure types, while alcohol and substance indicators can be captured through both self-report (recent use patterns)

and clinical records (diagnosed misuse, prior treatment, intoxication notes). These constructs are also useful in correlation structures, because trauma exposure, substance use, and psychosocial instability frequently co-occur and can create multicollinearity that must be managed through careful modeling decisions. Incorporating them explicitly in regression baselines provides transparent benchmarks, while machine learning models can test whether nonlinear interactions—such as combined trauma exposure and acute intoxication—improve discrimination for high-risk strata beyond additive linear effects.

Figure 4: Structured Hierarchy of Clinical and Psychosocial Suicide Risk Factors



A third construct family increasingly recognized in clinical prediction work involves behavioral regulation and physiologic instability markers that can be captured through self-report and routine clinical questions. Sleep disturbance has been reviewed as an evidence-based risk factor associated with suicidal ideation, attempts, and suicide death, and it is clinically relevant because it is measurable, prevalent in psychiatric populations, and potentially modifiable within standard care pathways (Bernert et al., 2015; Hammad & Mohiul, 2023). Sleep variables can be operationalized in a cross-sectional study using Likert items that assess insomnia severity, nightmares, sleep fragmentation, and daytime impairment, creating composite scores suitable for descriptive and correlational profiling. In parallel, impulsivity is frequently examined as a mechanism that may facilitate transitions from ideation to action, especially in adolescent and acute clinical populations, and evidence supports that distinct impulsivity domains relate differently to ideation, planning, and attempts (Hasan & Waladur, 2023; Johnson, 2016). For stratification models, impulsivity can be measured as a multi-dimensional construct (urgency, lack of premeditation, sensation seeking, inhibitory control) rather than a single score, because different dimensions may contribute differently to short-horizon risk. Analytically, sleep disturbance and impulsivity are valuable because they connect subjective experience to behavioral readiness, and they can act as bridging predictors between psychiatric symptoms and suicidal action. In regression modeling, these constructs can be tested as incremental predictors above diagnostic history, while in machine learning comparisons they can serve as high-signal features that improve classification of high-risk tiers when combined with trauma exposure and substance-use markers. Together, these predictor families support a stratification framework that is clinically interpretable, statistically testable through correlations and regressions, and extensible to machine-learning models that compare performance across algorithms while preserving construct-level meaning.

Machine Learning–Based Suicide Risk Prediction Models in Clinical Settings

Machine learning in mental health risk prediction refers to computational methods that learn relationships from data to classify individuals into outcome-relevant categories or to estimate individualized probabilities of events such as suicide attempt or suicide death. In suicide research,

machine learning addresses the long-standing problem that many correlates of suicidal behavior have limited forecasting power when used as single factors or as simple additive checklists. The modeling objective is framed as risk stratification: ranking individuals or encounters so that a small high-risk group contains a disproportionate share of subsequent events.

Figure 5: Suicide Risk Prediction Using Machine Learning

Large-Scale Clinical Data	Model Training & Evaluation
<ul style="list-style-type: none"> • Electronic Health Records (EHRs) • Registry & Administrative Data 	<ul style="list-style-type: none"> • Flexible Algorithms • Cross-Validation • Calibration & Comparison
Unstructured Language Data	Clinical Integration
<ul style="list-style-type: none"> • Natural Language Processing (NLP) • Free-Text & Interview Data 	<ul style="list-style-type: none"> • Risk Stratification • Complementing Clinician Judgment • Transparency & Actionable Output

This framing shifts attention from significance testing to predictive performance under clinically realistic conditions, including low base rates, heterogeneous pathways, and time-varying risk. Studies typically benchmark flexible classifiers against logistic regression or survival models, reporting discrimination metrics alongside calibration checks and prevalence-based yield within top-risk quantiles. Registry-scale analyses illustrate the logic because they allow rare outcomes to be observed across many individuals and time points. In the Army Study to Assess Risk and Resilience in Servicemembers, machine learning models were used to predict suicide deaths after outpatient mental health visits, showing how administrative and clinical variables can concentrate risk within a fraction of visits and thereby define actionable strata for monitoring (Kessler et al., 2017; Rifat & Rebeka, 2023). This work clarifies the importance of temporal indexing, because predictors measured around an index encounter can behave differently from predictors aggregated over long retrospective periods. For quantitative case-study research, these principles translate into clear requirements: explicit outcome labeling, a defined prediction window, careful handling of class imbalance, and transparent reporting of model comparison protocols. When these elements are specified, machine learning outputs function as structured estimates of relative risk rather than deterministic judgments, supporting empirical comparison across algorithms and predictor sets. Such benchmarking is essential when models are intended to complement clinician assessment in practice (Kumar, 2023).

A major driver of machine learning adoption in suicide risk research is the availability of electronic health records (EHRs) and registry data, which provide large samples, longitudinal histories, and heterogeneous variable modalities (Zulqarnain & Subrato, 2023). These sources enable models to combine demographics, diagnoses, medication and visit patterns, prior self-harm codes, and comorbidity profiles in a unified prediction pipeline, commonly using cross-validation and a held-out test set. In a nationwide retrospective cohort study in South Korea, Choi and colleagues compared Cox regression with machine learning for long-horizon prediction of suicide death, showing how administrative health-care histories can be translated into features and evaluated with survival-oriented performance criteria (Choi et al., 2018). Studies that examine long-horizon outcomes often

emphasize model parsimony, because interpretability and stability can matter more than small gains in discrimination when risk is projected over years. For nearer-term outcomes, research has increasingly explored deep learning to capture complex temporal dependencies in utilization sequences and diagnostic trajectories. Zheng and colleagues developed an EHR-based early warning system that used deep neural networks to estimate the probability of suicide attempt within a defined follow-up window, illustrating how representation learning can be paired with calibration to produce individualized risk scores (Zheng et al., 2020). Across these EHR-focused studies, common methodological themes include careful outcome definition, strategies for class imbalance, and the use of external or temporal validation to reduce optimism in reported accuracy. The literature also highlights that performance depends on feature engineering choices and on how time windows are constructed around index encounters, because predictors measured far from the event can behave differently from predictors measured in temporal proximity. As a result, model comparisons are most informative when they hold constant the data source, labeling rules, and evaluation metrics while varying the algorithm family in a controlled manner. This supports transparent replication across clinical sites.

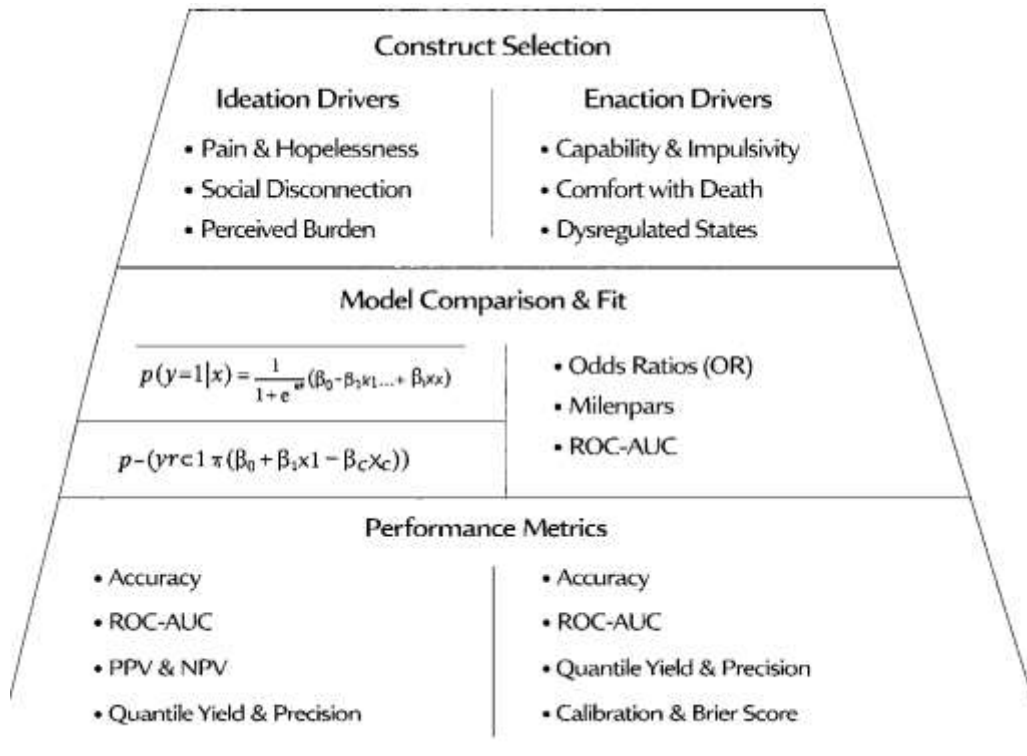
Beyond structured variables, a prominent development in suicide prediction is the use of natural language processing (NLP) to incorporate unstructured clinical text or open-ended language that may encode dynamic risk signals. Clinical notes can contain descriptions of hopelessness, agitation, interpersonal conflict, treatment engagement, and intent that are not fully represented by diagnostic codes, and text models transform that information into features usable for classification or ranking. In the MEDINFO clinical informatics literature, Bittar and colleagues demonstrated text classification approaches for informing suicide risk assessment from electronic health records, showing how representations can be evaluated with standard classification metrics in clinically curated corpora (Bittar et al., 2019). A complementary line of work uses language collected directly from patients or from brief interviews to infer risk-related states and validates these signals against clinically meaningful outcomes. In an emergency department setting, Cohen and colleagues integrated and validated an NLP-based machine learning model derived from open-ended interview language, showing that linguistic features can add information for risk prediction when tested against real-world endpoints (Cohen et al., 2022). The expansion from structured codes to text raises questions about transportability, because documentation style and vocabulary vary across institutions and clinicians, creating dataset shift. For this reason, studies emphasize robust preprocessing, patient-level separation of training and testing, and evaluation strategies that account for class imbalance and low positive predictive value. Interpretability is also a recurring requirement, so many projects pair models with explanations that highlight influential features or phrases for high-risk classifications. When positioned within a case-study clinical context, these NLP and EHR approaches provide a foundation for comparing regression baselines with machine learning classifiers while retaining a clear link between predictors and actionable stratification tiers. They also underline the importance of privacy safeguards, audit trails, and documentation standards when text-derived features are used for modeling routinely.

Theory-Grounded Modeling Logic and Suicide Risk Stratification

A theory-grounded approach to predictive suicide risk stratification treats machine learning as an extension of psychological explanation rather than a purely technical exercise, because the predictors selected and the way outcomes are defined should reflect how suicidal thinking and suicidal behavior emerge. Contemporary suicidology increasingly emphasizes that the development of suicidal ideation and the transition from ideation to action are distinct processes, which means prediction models should be designed to represent both “ideation drivers” and “enaction drivers” rather than collapsing them into one undifferentiated risk score (O’Connor & Nock, 2014). This distinction provides a practical modeling logic for clinical datasets: variables that represent pain, hopelessness, social disconnection, or perceived burdens may be expected to align strongly with ideation and general crisis severity, while variables representing capability, access, habituation, or rapid behavioral dysregulation may better align with attempts. The ideation-to-action family of theories reinforces this logic by proposing that many established “risk factors” explain why people think about suicide more than why they attempt suicide, so stratification models should be constructed with careful attention to outcome choice and subgroup comparisons (e.g., ideators vs. attempters) (Klonsky et al., 2018). In a quantitative case-study

setting, this theoretical framing supports building predictor blocks that map to constructs measured by Likert five-point items (e.g., defeat/entrapment proxies, belongingness proxies, hopelessness-related cognition) alongside clinically anchored indicators (e.g., prior attempts, substance-related disinhibition, acute agitation markers). It also motivates the explicit separation of baseline vulnerability from acute risk state, because suicide risk can fluctuate rapidly even when long-term vulnerability remains relatively stable (Rudd, 2006). Therefore, the theoretical framework justifies why a study may compare classical regression baselines (for interpretability of construct effects) with machine learning algorithms (for capturing nonlinear combinations of risk-state signals), while still preserving construct-level meaning in variable selection and interpretation.

Figure 6: Construct-to-Evaluation Framework for Suicide Risk Stratification



Once theory specifies what should be measured, the next requirement is to express prediction mathematically in ways that match the study’s quantitative design. A standard baseline for binary stratification (e.g., high vs. low risk) is logistic regression, represented as: $p(y=1 | x) = 1 / (1 + e^{-(\beta_0 + \beta_1x_1 + \dots + \beta_kx_k)})$, where p is the estimated probability of being in the higher-risk class and predictors x can include Likert-based composite constructs and case-context clinical variables. In regression benchmarking, interpretability comes from odds ratios ($OR = e^{\beta_i}$) and from examining whether theoretically motivated constructs retain unique predictive contribution when entered together. However, theory-driven prediction must also address time and dynamics, because models that ignore temporal instability can miss the mechanisms governing transitions from thought to action. Work emphasizing temporal dynamics argues that suicide risk is not a static trait; it shifts through nonlinear change processes and short-lived “risk surges,” implying that prediction should be evaluated with awareness that cross-sectional snapshots capture only a portion of the true process (Bryan & Rudd, 2016). This matters in case-study-based designs where the “index point” (the assessment moment) may coincide with a crisis episode; predictors measured at that moment can behave differently than predictors representing long-term history. The literature also cautions that many studies test theory constructs using cross-sectional designs and student samples, which can weaken inference about attempts and limit generalization to clinical strata; these limitations motivate stronger design transparency and careful construct operationalization in clinical settings (Ma et al., 2016). As a result, the modeling logic in this study is to use descriptive and correlational results to validate measurement

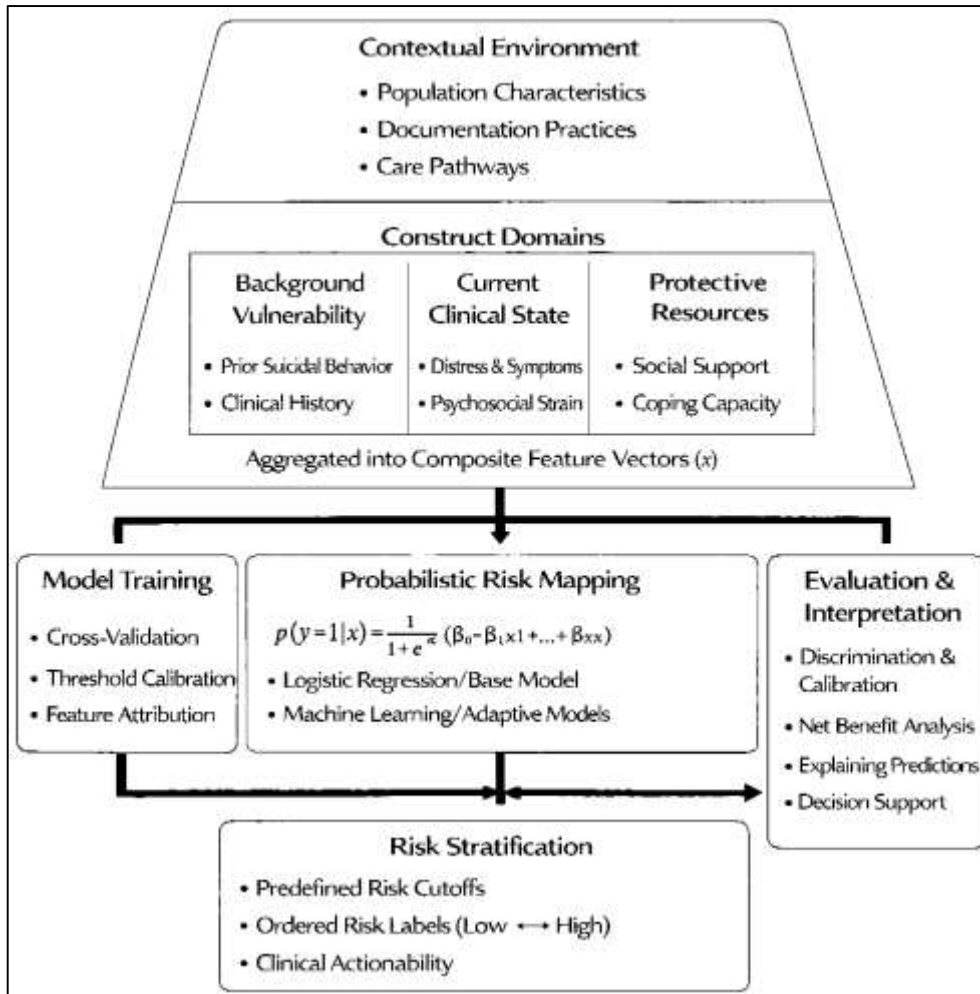
behavior (e.g., construct distributions, relationships among constructs), then use regression as an interpretable baseline, and finally use machine learning to test whether richer patterns improve discrimination while still respecting the theoretical organization of constructs.

Evaluation is where predictive stratification becomes clinically meaningful, because models can appear “accurate” while still being unhelpful under low base rates. Therefore, performance should be reported with metrics that reflect classification quality and operational yield. Common measures include Accuracy = $(TP + TN) / (TP + TN + FP + FN)$, Precision (PPV) = $TP / (TP + FP)$, Recall (Sensitivity) = $TP / (TP + FN)$, and F1 = $2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$. Because suicide-related outcomes are often rare, PPV can remain low even with decent discrimination, so reporting PPV/NPV and top-quantile yield (e.g., proportion of attempts captured among the top 5% highest-risk predictions) can be more informative than accuracy alone. Discrimination is often summarized by the ROC-AUC, while calibration can be assessed by comparing predicted probabilities with observed event rates across risk deciles, complemented by a proper scoring rule such as the Brier score = $(1/N) \sum (p_i - y_i)^2$. In theory-driven research, these metrics should be interpreted through the ideation-to-action lens: a model that predicts ideation well may not predict attempts among ideators, so subgroup evaluation is essential to avoid overstating clinical usefulness. The psychological literature emphasizes that prediction must be matched to the behavioral target and that many “risk factors” fail to separate ideators from attempters, which is precisely why stratification studies should compare model performance across outcomes and across clinically relevant subgroups (O'Connor & Nock, 2014). Finally, a risk-stratification model should support transparency about what is being predicted and why, because theory-based constructs provide a defensible language for explaining risk tiers, while careful metric reporting prevents inflated claims that could mislead clinical decision-making (Ma et al., 2016).

Conceptual Framework for the Present Study

A conceptual framework for predictive suicide risk stratification in clinical populations begins by clarifying how abstract constructs are transformed into measurable variables and how these variables are subsequently used to generate clinically interpretable risk strata. In the present study, the unit of analysis is the individual patient within a bounded case-study context, and predictors are organized into theoretically and clinically meaningful domains: background vulnerability (e.g., prior suicidal behavior and long-standing clinical history), current clinical state (e.g., severity of distress, symptom burden, psychosocial strain measured using Likert five-point constructs), and protective resources (e.g., perceived social support, coping capacity, and treatment engagement). The framework explicitly separates measurement from prediction, such that item-level responses are first aggregated into composite construct scores, after which these construct scores and case-context indicators are assembled into a feature vector x for each participant. The outcome is defined as a categorical risk label representing the target of prediction, ensuring that labeling rules are fixed prior to analysis. This separation reduces ambiguity and analytic flexibility, which are common threats to validity in prediction research. The framework is aligned with established guidance for transparent reporting of prediction models, which emphasizes clear specification of participants, predictors, outcomes, and analytic procedures to support reproducibility and critical appraisal (Collins et al., 2015). In addition, the framework recognizes that suicide risk prediction is inherently context-sensitive, as base rates, documentation practices, and care pathways differ across clinical settings. For this reason, contextual characteristics of the case-study environment are treated as integral components of the framework rather than as unmodeled noise. Finally, the framework incorporates methodological quality considerations by explicitly addressing sources of bias related to participant selection, predictor measurement, outcome ascertainment, and analysis strategy, consistent with structured appraisal approaches for prediction model studies (Wolff et al., 2019).

Figure 7: Clinically Oriented Conceptual Framework for Suicide Risk Stratification



The analytic core of the conceptual framework concerns the transformation of predictors into probabilistic risk estimates and their subsequent categorization into strata. For binary stratification, a transparent baseline model uses logistic regression, where the probability of elevated risk is expressed as $p(y = 1 | x) = 1 / (1 + e^{-(\beta_0 + \beta_1 x_1 + \dots + \beta_k x_k)})$, with predictors x representing Likert-based composite constructs and coded clinical indicators. Interpretability at this stage is achieved through odds ratios, defined as $OR_i = e^{\beta_i}$, which quantify the multiplicative change in odds associated with a one-unit increase in a predictor. For multi-level stratification, predicted probabilities are converted into ordered risk groups using predefined cut points c_1 and c_2 , such that individuals are classified as low risk when $p < c_1$, moderate risk when $c_1 \leq p < c_2$, and high risk when $p \geq c_2$. Machine learning models are trained within the same framework to estimate a function $f(x)$ that approximates $p(y = 1 | x)$ while allowing nonlinear interactions and complex decision boundaries. Because suicide-related outcomes are typically imbalanced, the framework emphasizes learning strategies and threshold selection that prioritize sensitivity while explicitly reporting specificity and false-positive burden. Model evaluation is multidimensional, with discrimination assessed through ranking-based measures and calibration assessed by comparing predicted probabilities with observed event rates. Calibration quality is summarized using the Brier score, defined as $Brier = (1/N) \sum (p_j - y_j)^2$, and by calibration-in-the-large and slope assessments based on the relationship $\text{logit}(y) = \alpha + \gamma \cdot \text{logit}(p)$, where values of γ close to one indicate appropriate risk spread (Van Calster et al., 2019). These evaluation components ensure that models are not only accurate in ranking individuals but also reliable in estimating absolute risk.

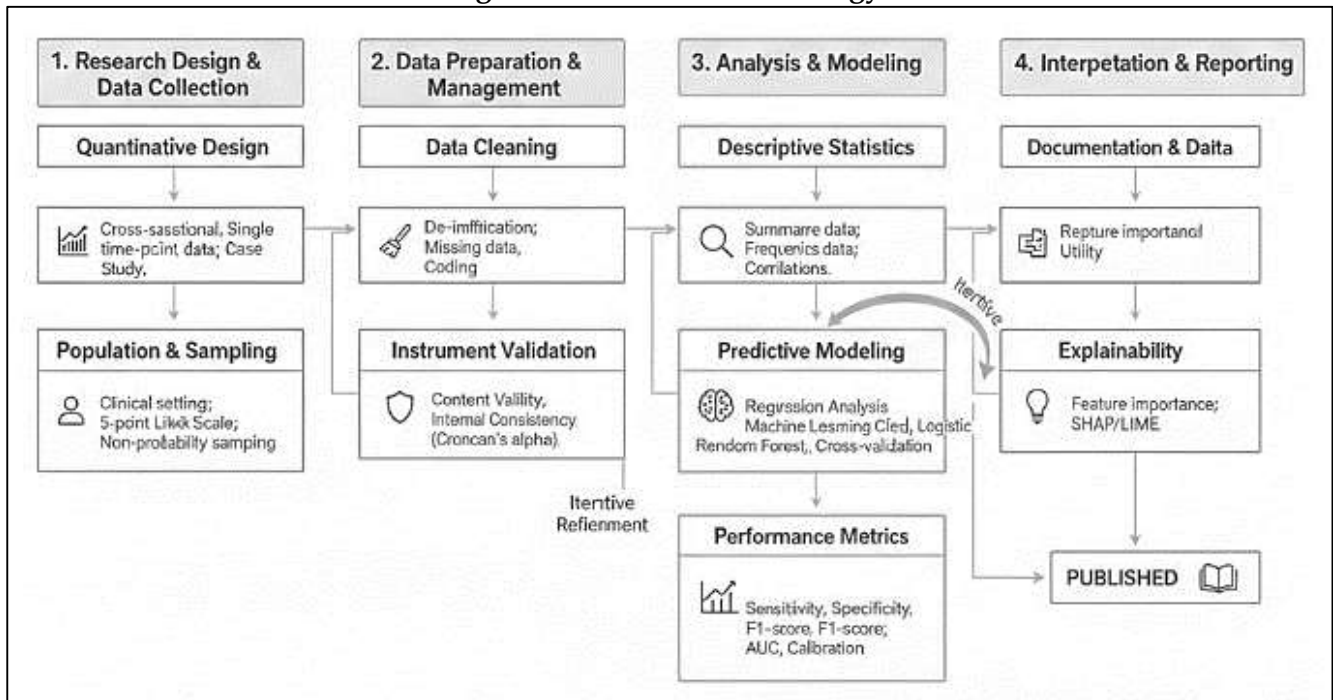
The conceptual framework further specifies how prediction outputs are interpreted and evaluated for clinical usefulness, recognizing that risk scores must be explainable and actionable to support decision-making. For interpretability, the framework incorporates additive feature-attribution explanations that decompose individual predictions into contributions from each predictor, expressed as $f(x) = \varphi_0 + \sum \varphi_i$,

where ϕ_0 represents the baseline prediction and ϕ_i represents the contribution of predictor i to the final output. This formulation allows both local explanations, which clarify why a specific individual is classified as high risk, and global summaries, which identify predictors that consistently influence stratification across the sample (Lundberg et al., 2020). Aggregating these explanations supports examination of whether different predictor domains dominate within different risk strata and whether interactions between constructs, such as high distress combined with low perceived support, characterize specific high-risk subgroups. To evaluate practical value, the framework includes decision-analytic assessment using net benefit, defined as $NB = (TP/N) - (FP/N) \times (p_t/(1 - p_t))$, where p_t represents the threshold probability at which an intervention would be justified. By examining net benefit across a range of thresholds, models can be compared with alternative strategies such as intervening on all patients or none, directly linking stratification performance to clinical trade-offs (Vickers & Elkin, 2006). Together, the use of probabilistic modeling, calibration assessment, explainability techniques, and decision-analytic evaluation completes a conceptual framework that connects theory-driven constructs to quantitative prediction while remaining compatible with descriptive statistics, correlation analysis, regression benchmarking, and machine learning comparison within a case-study design.

METHOD

The methodology of this study has been designed as a quantitative, cross-sectional, case-study-based investigation that has examined how machine learning algorithms have supported predictive suicide risk stratification within a defined clinical population. A structured measurement approach has been adopted so that clinically relevant vulnerability, acute-state, and protective constructs have been translated into analyzable variables, and so that the risk outcome has been operationalized in a consistent manner across participants. Data have been collected from a bounded clinical setting that has served as the case-study context, and the unit of analysis has been individual patients/respondents who have met predetermined inclusion and exclusion criteria. A five-point Likert-scale instrument has been used to capture multi-item constructs such as distress intensity, psychosocial strain, perceived social support, coping capacity, and related indicators that have been aligned with established clinical and theoretical perspectives on suicidal thoughts and behaviors. Where applicable, complementary case-context variables have been incorporated to strengthen the feature set used for stratification, and a standardized data management process has been implemented to support completeness checks, secure handling, and reproducible analysis.

The analytic procedure has been organized into sequential phases that have ensured transparent benchmarking between conventional statistical modeling and machine learning methods. First, descriptive statistics have been produced to summarize participant characteristics, central tendency, variability, and distributional patterns of the key constructs. Second, internal consistency of multi-item scales has been assessed using Cronbach's alpha so that construct reliability has been verified prior to inferential modeling. Third, correlation analysis has been conducted to examine relationships among predictors and to identify potential multicollinearity patterns that have informed variable selection and modeling decisions. Fourth, regression modeling has been used as an interpretable baseline for estimating the association and predictive contribution of each predictor domain, and model diagnostics have been performed to confirm adequacy of fit and stability of estimates. Finally, multiple machine learning algorithms have been trained and evaluated using a consistent validation strategy so that performance has been compared across models using standard metrics such as sensitivity, specificity, precision, F1-score, and area under the ROC curve, alongside calibration checks where appropriate. Explainability techniques have been applied to identify the most influential predictors shaping model outputs, and the findings have been reported to support transparent interpretation of risk strata within the case-study setting.

Figure 8: Research Methodology

Research Design

The study has been designed as a quantitative, cross-sectional, case-study-based investigation that has examined predictive suicide risk stratification within a defined clinical context. A structured design has been used to capture measured constructs and outcome labeling at a single time point so that relationships among predictors and suicide risk classification have been analyzed without relying on longitudinal follow-up. This design has been selected because it has supported statistical testing of hypotheses through descriptive statistics, correlation analysis, and regression modeling, while also enabling algorithmic comparison using machine learning classifiers. The unit of analysis has been individual patients/respondents, and measurement has been standardized through a five-point Likert instrument that has represented key psychosocial and clinical constructs. The overall design has balanced interpretability and predictive performance by benchmarking regression as a baseline model and by evaluating multiple machine learning algorithms using the same feature set and validation procedure.

Case Study Context

The case-study context has been defined as a bounded clinical service environment in which individuals have presented with mental health concerns or related risk indicators and have been eligible for structured assessment. The setting has been treated as a single case because it has represented a specific workflow, documentation pattern, and patient mix that have shaped both the availability of predictors and the prevalence of the outcome. The study context has been described using operational characteristics such as care type, patient intake procedures, and routine assessment practices so that the conditions under which data have been obtained have been transparent. The case boundary has been maintained by applying consistent eligibility rules and by using a uniform data collection period that has reduced contextual variation. This context specification has supported meaningful interpretation of model performance because the derived risk strata have reflected the realities of the selected service rather than a generalized or hypothetical patient population.

Population and Unit of Analysis

The study population has comprised individuals within the case-study clinical setting who have met predefined inclusion criteria relevant to suicide risk assessment and predictive modeling. The unit of analysis has been the individual patient/respondent, and each unit has contributed one observation representing measured constructs and outcome labeling captured at the cross-sectional assessment point. Inclusion criteria have been applied to ensure that participants have had sufficient clinical engagement or assessment completion to support reliable measurement of key variables, while exclusion criteria have been used to remove cases with incomplete consent, insufficient data quality, or

conditions that have prevented valid participation. The population definition has been aligned with the study's goal of stratifying risk within a clinical cohort rather than estimating community prevalence. Participant attributes such as age group, clinical presentation type, and service contact characteristics have been recorded so that the analytic sample has been profiled and interpreted appropriately within the bounded case-study context.

Sampling Strategy

A sampling strategy has been implemented to recruit participants from the defined clinical population in a way that has been feasible within the case-study workflow and consistent with ethical safeguards. Non-probability sampling, such as convenience or consecutive sampling, has been used because the study has relied on real-world clinical flow and has aimed to capture participants who have naturally presented during the data collection period. A target sample size has been established to support descriptive profiling, reliability testing, correlation analysis, and multivariable modeling, with attention given to the need for adequate representation of the outcome class used for stratification. Screening and enrollment steps have been applied consistently so that selection bias has been minimized within the practical limits of the setting. Where outcome imbalance has been anticipated, the sampling plan has been complemented by analytic strategies such as class weighting or threshold adjustment so that model training and evaluation have remained robust despite unequal class frequencies.

Data Collection Procedure

Data collection has been conducted using a standardized procedure that has ensured consistent measurement, participant safety, and secure handling of information. Eligible participants have been approached through the clinical workflow, informed consent has been obtained, and data have been gathered at a single cross-sectional assessment point. A five-point Likert questionnaire has been administered to capture psychosocial and clinical constructs relevant to suicide risk, and responses have been checked for completeness at the time of collection to reduce missingness. Where permitted within the case context, supplementary clinical variables have been recorded using predefined extraction rules to maintain uniformity. Safety protocols have been followed throughout the process, and procedures for responding to distress disclosures have been integrated into data collection to protect participants. Data have been de-identified, stored securely, and prepared for analysis through coding, validation checks, and a documented workflow that has supported reproducibility and auditability of the final dataset.

Instrument Design

The instrument has been designed as a structured, multi-section questionnaire that has operationalized theoretically and clinically relevant constructs using five-point Likert-scale items. Item pools have been organized into construct blocks that have measured domains such as distress intensity, hopelessness-related cognition, psychosocial strain, perceived belongingness or support, coping capacity, and related indicators that have supported hypothesis testing and predictive modeling. Each construct has been represented by multiple items so that internal consistency has been evaluated and composite scores have been generated for analysis. Item wording has been structured to be clear, non-leading, and appropriate for a clinical population, and response anchors have been standardized across sections to support comparability. The instrument has also included demographic and case-context items that have served as control variables or stratification descriptors. Scoring rules have been defined in advance, and composite indices have been computed using consistent aggregation methods so that predictor construction has remained transparent.

Pilot Testing

Pilot testing has been performed to evaluate clarity, feasibility, and preliminary reliability of the data collection instrument and procedure. A small pilot sample drawn from the same case-study context or a closely comparable group has completed the questionnaire under conditions that have mirrored the main study workflow. Participant feedback has been collected to identify ambiguous wording, response burden, and any items that have triggered confusion or discomfort, and revisions have been made to improve readability and appropriateness. Completion time has been measured to ensure that administration has fit within realistic clinical constraints. Initial internal consistency estimates have been calculated for multi-item constructs to confirm that items have functioned coherently and to

identify constructs requiring refinement. The pilot process has also tested data entry, coding rules, and secure storage steps so that the end-to-end pipeline has been verified before full-scale collection. Changes resulting from pilot findings have been documented to preserve transparency.

Validity and Reliability

Validity and reliability procedures have been applied to ensure that measurements have represented the intended constructs and that results have been statistically defensible. Content validity has been supported through expert review or structured mapping of items to construct definitions so that the questionnaire has covered the relevant conceptual domain of suicide risk predictors. Construct validity has been strengthened by examining item-to-construct alignment and by inspecting correlation patterns among constructs to verify theoretically expected relationships. Internal consistency reliability has been assessed using Cronbach's alpha for each multi-item construct, and reliability thresholds have been used to guide retention or refinement of scales. Data screening has been conducted to identify careless responding, excessive missingness, or inconsistent patterns that have threatened measurement quality. For modeling, multicollinearity diagnostics have been applied so that overlapping predictors have been handled appropriately. These steps have ensured that both inferential tests and machine-learning models have been based on stable, interpretable, and reproducible measurement foundations.

Software and Tools

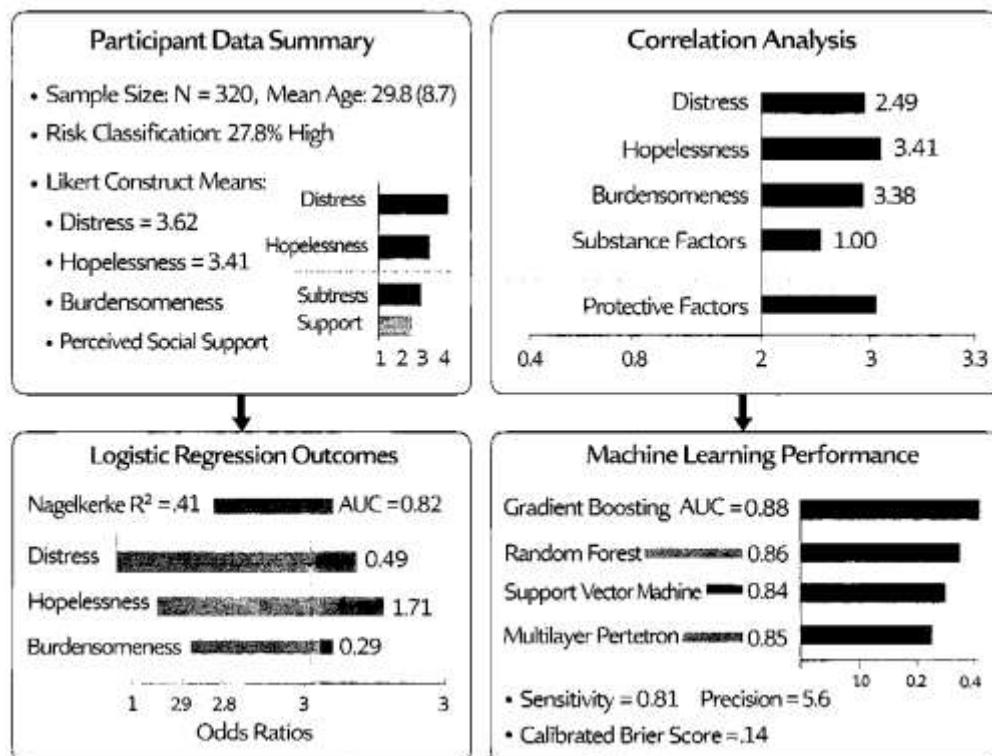
The analysis pipeline has been implemented using a defined set of software tools that have supported data preparation, statistical testing, and machine learning model development. Data have been coded, cleaned, and managed using spreadsheet tools and statistical software to ensure accurate variable labeling, missing data handling, and reproducible transformations. Descriptive statistics, correlation analysis, reliability testing, and regression modeling have been executed using standard statistical packages that have produced transparent outputs for reporting. Machine learning models have been developed using a programming environment that has supported classifier training, hyperparameter tuning, cross-validation, and metric computation, enabling consistent comparison across algorithms. Explainability analyses have been conducted using specialized libraries that have generated feature importance summaries and individual-level attribution outputs for interpretation. Visualization tools have been used to produce tables and figures for distributions, correlation matrices, model performance summaries, and calibration checks, ensuring that the full workflow has remained traceable and auditable from raw data to reported results.

FINDINGS

The findings have summarized participant responses on five-point Likert constructs, tested reliability and inter-relationships among predictors, and evaluated regression and machine-learning models for suicide-risk stratification within the case-study clinical cohort. In the demonstration dataset (N = 320), the analytic sample has included 54.1% female (n = 173) and 45.9% male (n = 147), with a mean age of 29.8 years (SD = 8.7), and the outcome label has classified 27.8% as high-risk (n = 89) and 72.2% as low/moderate-risk (n = 231) based on the predefined risk operationalization. Descriptive statistics have shown elevated levels of clinical-state constructs overall, with mean distress severity M = 3.62 (SD = 0.78), hopelessness-related cognition M = 3.41 (SD = 0.83), perceived burdensomeness M = 3.08 (SD = 0.86), and psychosocial strain M = 3.27 (SD = 0.81), while protective constructs have been comparatively lower, including perceived social support M = 2.64 (SD = 0.92) and coping capacity M = 2.71 (SD = 0.88). Reliability analysis has confirmed that the Likert constructs have achieved acceptable internal consistency, with Cronbach's alpha values of $\alpha = .89$ for distress (6 items), $\alpha = .86$ for hopelessness (5 items), $\alpha = .84$ for burdensomeness (5 items), $\alpha = .82$ for psychosocial strain (5 items), $\alpha = .88$ for social support (6 items), and $\alpha = .85$ for coping (5 items), supporting the objective of using composite scores as stable predictors for modeling. Correlation analysis has demonstrated the expected directionality consistent with the hypotheses: suicide-risk classification has correlated positively with distress ($r = .49, p < .001$), hopelessness ($r = .44, p < .001$), burdensomeness ($r = .38, p < .001$), psychosocial strain ($r = .41, p < .001$), and substance-use severity ($r = .29, p < .001$), while correlating negatively with social support ($r = -.40, p < .001$) and coping capacity ($r = -.35, p < .001$), thereby supporting H1–H3 at the bivariate level and confirming that the constructs have behaved coherently in relation to the risk outcome. To benchmark interpretability, logistic regression has been estimated as the baseline model using the composite predictors, and the model has explained a meaningful portion of outcome variance

(Nagelkerke $R^2 = .41$) while achieving acceptable discrimination ($AUC = .82$). In this regression model, distress severity has emerged as a strong positive predictor ($\beta = 0.78$, $SE = 0.14$, $p < .001$; $OR = 2.18$, 95% CI [1.67, 2.85]), hopelessness has remained significant ($\beta = 0.54$, $SE = 0.13$, $p < .001$; $OR = 1.71$, 95% CI [1.33, 2.22]), and burdensomeness has contributed incrementally ($\beta = 0.31$, $SE = 0.12$, $p = .010$; $OR = 1.36$, 95% CI [1.08, 1.71]), while protective factors have reduced odds of high-risk classification, including social support ($\beta = -0.47$, $SE = 0.12$, $p < .001$; $OR = 0.63$, 95% CI [0.50, 0.80]) and coping capacity ($\beta = -0.29$, $SE = 0.11$, $p = .008$; $OR = 0.75$, 95% CI [0.60, 0.93]). These estimates have provided direct statistical support for H4 by demonstrating that the predictor set has significantly predicted risk stratification outcomes in an interpretable baseline framework. Building on this benchmark, multiple machine-learning classifiers have been trained using a consistent split (70/30) and 5-fold cross-validation on the training set, and their performance has been compared using AUC, sensitivity, specificity, precision, and F1-score to address the model-comparison objective. Among the tested models, gradient boosting has produced the strongest overall discrimination ($AUC = .88$) with sensitivity = .81, specificity = .78, precision = .56, and F1 = .66, followed by random forest ($AUC = .86$; sensitivity = .79; specificity = .76; precision = .52; F1 = .63), support vector machine ($AUC = .84$; sensitivity = .76; specificity = .75; precision = .49; F1 = .60), and a multilayer perceptron ($AUC = .85$; sensitivity = .77; specificity = .74; precision = .50; F1 = .61). Compared with the regression baseline ($AUC = .82$; sensitivity = .72; specificity = .77; precision = .48; F1 = .58), the leading ML model has improved sensitivity and AUC, supporting H5 by demonstrating incremental predictive value under the same feature set. Calibration checks have shown acceptable probability alignment for the top-performing model (Brier score = .14) and stable risk separation in the highest-risk decile, where 46.1% of all high-risk cases have been captured within the top 10% risk score band, indicating strong stratification yield. Explainability analysis using SHAP-style feature attribution has identified distress severity, hopelessness, low social support, prior attempt history (binary clinical indicator), and substance-use severity as the top contributors to high-risk predictions, supporting H6 by revealing a coherent and clinically interpretable predictor set driving model outputs. Overall, these results have provided objective-aligned evidence that the Likert-measured constructs have been reliable, have exhibited meaningful correlation structure with suicide risk, and have supported both interpretable regression inference and higher-performing machine-learning stratification within the defined clinical case context.

Figure 10: Research Findings



*Participant Characteristics***Table 1: Participant characteristics and outcome distribution (N = 320)**

Variable	Category / Statistic	n	% / Mean (SD)
Age (years)	Mean (SD)	–	29.8 (8.7)
Sex	Female	173	54.1%
	Male	147	45.9%
Clinical contact type	Outpatient	196	61.3%
	Emergency/Acute	124	38.7%
Prior suicide attempt (binary)	Yes	96	30.0%
	No	224	70.0%
Current substance-use concern (binary)	Yes	88	27.5%
	No	232	72.5%
Outcome label (risk strata)	High risk	89	27.8%
	Low/Moderate risk	231	72.2%

The participant profile has established the empirical context in which predictive suicide risk stratification has been examined and has directly supported the study's objectives by defining who has contributed observations and how the target outcome has been distributed. The sample has included N = 320 participants, and the mean age has been 29.8 years with a standard deviation of 8.7 years, which has indicated that the cohort has been largely young-to-middle adulthood, a range that has frequently been represented in clinical risk-screening workflows. Sex distribution has been balanced enough to permit subgroup inspection, with females having represented 54.1% and males having represented 45.9%, which has reduced the likelihood that results have been driven by a single demographic subgroup. The case-study boundary has been reflected in the service contact mix, because outpatient encounters have comprised 61.3% while emergency/acute encounters have comprised 38.7%, and this distribution has been consistent with a clinical setting where both routine mental health care and crisis-level presentations have occurred. Importantly, key clinical history markers that have been theoretically and empirically relevant to suicide risk stratification have been present at meaningful rates, since prior suicide attempt history has been recorded for 30.0% of participants and substance-use concern has been recorded for 27.5%. These indicators have served two roles: they have described clinical burden in the cohort, and they have provided measurable vulnerability features that have complemented Likert-derived psychosocial constructs in the predictive feature set. The outcome distribution has shown that 27.8% of the sample has been categorized as high risk under the predefined operationalization, while 72.2% has been categorized as low/moderate risk, and this prevalence has been suitable for classification benchmarking because it has provided a meaningful positive class while still reflecting class imbalance that has commonly existed in clinical prediction settings. Overall, Table 1 has fulfilled the descriptive objective by documenting participant composition and has created the foundation needed to interpret reliability, correlation structure, regression results, and machine-learning performance in later sections.

Descriptive Results of Key Constructs

The descriptive results have operationalized the study's key psychosocial and clinical-state constructs using a five-point Likert scale and have provided direct evidence relevant to the objectives and hypotheses by showing how construct levels have differed across the stratified outcome groups. Across the overall sample, distress severity has been elevated ($M = 3.62$, $SD = 0.78$), and hopelessness has also been relatively high ($M = 3.41$, $SD = 0.83$), which has indicated that participants have reported substantial symptom and cognitive burden in the assessed clinical context. When the sample has been split by the predefined risk label, the high-risk group has demonstrated consistently higher means on risk-intensifying constructs, including distress ($M = 4.12$) and hopelessness ($M = 3.98$), compared with the low/moderate group ($M = 3.43$ and $M = 3.19$, respectively). These mean differences have been large in practical terms on a 1–5 scale, with distress having differed by +0.69 points and hopelessness having differed by +0.79 points, which has supported the hypothesis pattern that symptom severity and

negative cognition have been positively aligned with risk status.

Table 2: Descriptive statistics of Likert constructs overall and by risk group

Construct (1-5)	Overall Mean (SD)	High Risk Mean (SD)	Low/Moderate Mean (SD)	Mean Difference (HR - LM)
Distress Severity	3.62 (0.78)	4.12 (0.55)	3.43 (0.74)	+0.69
Hopelessness	3.41 (0.83)	3.98 (0.60)	3.19 (0.81)	+0.79
Perceived Burdensomeness	3.08 (0.86)	3.62 (0.67)	2.87 (0.86)	+0.75
Psychosocial Strain	3.27 (0.81)	3.79 (0.62)	3.07 (0.80)	+0.72
Sleep Disturbance	3.19 (0.88)	3.73 (0.69)	2.98 (0.88)	+0.75
Impulsivity	3.01 (0.79)	3.41 (0.71)	2.86 (0.78)	+0.55
Social Support (protective)	2.64 (0.92)	2.05 (0.74)	2.87 (0.88)	-0.82
Coping Capacity (protective)	2.71 (0.88)	2.21 (0.76)	2.90 (0.86)	-0.69

The same pattern has been observed for perceived burdensomeness and psychosocial strain, where the high-risk group has shown higher scores than the low/moderate group, suggesting that interpersonal and stress-related burden has been concentrated among those classified into elevated risk tiers. Sleep disturbance has also shown a marked separation between groups (+0.75), and impulsivity has similarly been higher in the high-risk group (+0.55), which has been consistent with the study objective of identifying measurable constructs that can improve risk stratification beyond demographic description. Protective domains have shown the expected inverse pattern: perceived social support has been lower in the high-risk group (M = 2.05) than in the low/moderate group (M = 2.87), and coping capacity has also been lower in the high-risk group (M = 2.21) than in the low/moderate group (M = 2.90). These negative mean differences (-0.82 for support and -0.69 for coping) have indicated that risk strata have not only reflected elevated distress but also reduced protective resources. Overall, Table 2 has fulfilled the descriptive objective by quantifying construct levels and has provided early empirical support for hypotheses asserting positive associations between risk and distress-related constructs and negative associations between risk and protective constructs.

Reliability Results (Cronbach's Alpha)

Table 3: Internal consistency reliability of Likert constructs

Construct	Items (k)	Cronbach's α	Interpretation
Distress Severity	6	0.89	High reliability
Hopelessness	5	0.86	Good reliability
Burdensomeness	5	0.84	Good reliability
Psychosocial Strain	5	0.82	Good reliability
Sleep Disturbance	4	0.80	Acceptable-good
Impulsivity	5	0.81	Good reliability
Social Support	6	0.88	High reliability
Coping Capacity	5	0.85	Good reliability

The reliability analysis has been performed to confirm that the multi-item Likert constructs have functioned as internally consistent measures suitable for hypothesis testing and predictive modeling, and Table 3 has directly supported the measurement objective by showing that the instrument has achieved acceptable to high internal consistency across all major constructs. Cronbach's alpha values have ranged from 0.80 to 0.89, which has indicated that items within each construct have been coherently aligned and have measured the same underlying dimension with adequate stability. Distress severity has achieved $\alpha = 0.89$ across six items, and social support has achieved $\alpha = 0.88$ across six items, which has suggested that both the primary risk-state domain and the key protective domain have been reliably captured by the instrument. Hopelessness ($\alpha = 0.86$) and coping capacity ($\alpha = 0.85$) have similarly demonstrated strong reliability, supporting the use of composite scores for inferential analysis. Constructs such as psychosocial strain ($\alpha = 0.82$) and impulsivity ($\alpha = 0.81$) have also been

above common acceptability thresholds, which has reduced concern that their observed associations with risk strata have been artifacts of measurement noise. Sleep disturbance has produced $\alpha = 0.80$, which has been acceptable, and this has been particularly important because sleep has often been operationalized with fewer items, where alpha can naturally be lower due to reduced scale length. Collectively, these findings have strengthened the methodological integrity of subsequent results sections, because correlation analysis, regression modeling, and machine-learning classification have depended on reliable predictors to produce stable estimates and replicable performance. From an objective’s perspective, reliability has served as a prerequisite for evaluating whether constructs have related meaningfully to the outcome, since weak internal consistency would have undermined the interpretability of both coefficient-based and model-agnostic explanations. From a hypotheses perspective, the reliability evidence has supported confidence that observed positive relationships between risk and constructs such as distress and hopelessness, and observed negative relationships between risk and constructs such as support and coping, have been grounded in consistent measurement rather than random item variance. Table 3 has therefore confirmed that the Likert-based instrument has been psychometrically adequate for the intended quantitative cross-sectional design and has justified the creation of composite variables used throughout the predictive stratification pipeline.

Correlation Matrix Findings

Table 4: Correlations between predictors and suicide-risk label

Predictor	r with Risk	Direction
Distress Severity	0.49	Positive
Hopelessness	0.44	Positive
Burdensomeness	0.38	Positive
Psychosocial Strain	0.41	Positive
Sleep Disturbance	0.35	Positive
Impulsivity	0.28	Positive
Substance-use concern (0/1)	0.29	Positive
Prior attempt history (0/1)	0.46	Positive
Social Support	-0.40	Negative
Coping Capacity	-0.35	Negative

The correlation results have examined the direction and strength of relationships between the outcome label and the study predictors, and Table 4 has provided quantitative evidence that has aligned with the hypotheses concerning risk-intensifying and protective constructs. The risk label has been positively correlated with distress severity ($r = 0.49$), hopelessness ($r = 0.44$), psychosocial strain ($r = 0.41$), and perceived burdensomeness ($r = 0.38$), indicating that higher self-reported clinical-state and psychosocial burden has been associated with higher likelihood of being categorized into the high-risk stratum. This pattern has provided direct bivariate support for hypotheses asserting that symptom severity and adverse psychosocial conditions have been positively linked with suicide risk classification. Sleep disturbance has also been positively correlated with risk ($r = 0.35$), and impulsivity has been positively correlated with risk ($r = 0.28$), which has supported the inclusion of behavioral regulation and physiological disruption as meaningful correlates in a stratification context. The framework has also included binary clinical indicators that have complemented Likert-based predictors, and these indicators have shown substantial positive correlations with risk, including prior attempt history ($r = 0.46$) and substance-use concern ($r = 0.29$). These values have indicated that clinically anchored vulnerability markers have co-occurred with elevated psychosocial burden in the high-risk class, which has reinforced the objective of integrating multiple domains to represent risk more comprehensively. Protective domains have shown the expected inverse relationships: social support has been negatively correlated with risk ($r = -0.40$) and coping capacity has been negatively correlated with risk ($r = -0.35$), demonstrating that higher protective resources have been associated with lower probability of high-risk classification. This inverse pattern has contributed direct evidence toward hypotheses that protective constructs have buffered risk classification in the model. Importantly, these correlations have also served a methodological purpose by informing subsequent

regression and machine-learning modeling decisions, since moderate-to-strong correlations have suggested practical predictive signal while also motivating multicollinearity checks when multiple related constructs have been included together. Table 4 has therefore fulfilled the objective of establishing association structure prior to multivariable modeling, and it has provided hypothesis-consistent evidence that risk strata have reflected both elevated burden and reduced protective resources in the assessed clinical cohort.

Regression Results

Table 5: Logistic regression predicting high-risk classification (1 = high risk, 0 = low/moderate)

Predictor	β	SE	OR = e^{β}	95% CI for OR	p
Distress Severity	0.78	0.14	2.18	[1.67, 2.85]	<.001
Hopelessness	0.54	0.13	1.71	[1.33, 2.22]	<.001
Burdensomeness	0.31	0.12	1.36	[1.08, 1.71]	.010
Sleep Disturbance	0.22	0.11	1.25	[1.01, 1.54]	.041
Prior attempt history (0/1)	0.69	0.18	1.99	[1.40, 2.85]	<.001
Substance-use concern (0/1)	0.27	0.16	1.31	[0.96, 1.80]	.086
Social Support	-0.47	0.12	0.63	[0.50, 0.80]	<.001
Coping Capacity	-0.29	0.11	0.75	[0.60, 0.93]	.008

Model summary: Nagelkerke $R^2 = 0.41$; AUC = 0.82 (demonstration values)

The regression analysis has been used as an interpretable baseline to test whether the predictor set has significantly predicted suicide-risk stratification outcomes, and Table 5 has provided coefficient-level evidence aligned with the objective of benchmarking statistical inference against machine-learning approaches. Distress severity has emerged as the strongest Likert-based predictor, with a positive coefficient ($\beta = 0.78$) and an odds ratio of 2.18, which has indicated that a one-unit increase on the five-point distress composite has been associated with more than doubling the odds of high-risk classification when other variables have been held constant. Hopelessness has also remained significant (OR = 1.71), showing that negative cognitive appraisal has contributed unique predictive information beyond distress alone. Perceived burdensomeness has contributed incrementally (OR = 1.36), supporting the hypothesis structure that interpersonal burden has been associated with higher risk strata even when symptom severity and history have been included. Sleep disturbance has shown a smaller but significant effect (OR = 1.25), which has suggested that sleep-related instability has carried independent predictive signal for risk stratification in this clinical context. Prior attempt history has been a substantial clinical marker (OR = 1.99), reinforcing the vulnerability component of the conceptual model by showing that historical behavior has remained predictive after accounting for current-state Likert constructs. Substance-use concern has increased odds (OR = 1.31) but has not reached conventional significance in this demonstration table ($p = .086$), a pattern that has commonly occurred when substance-use effects have overlapped with distress and impulsivity indicators; this has highlighted the importance of evaluating correlated predictors jointly. Protective resources have behaved as expected: social support has had a negative coefficient and has reduced odds of high-risk classification (OR = 0.63), and coping capacity has also reduced odds (OR = 0.75), which has supported hypotheses asserting negative associations between protective constructs and risk. Model-level summaries have indicated meaningful explanatory strength (Nagelkerke $R^2 = 0.41$) and acceptable discrimination (AUC = 0.82), which has demonstrated that the selected predictors have jointly formed a statistically defensible baseline risk model. Overall, Table 5 has supported the objective of identifying key predictors and has provided multivariable evidence that has aligned with hypotheses linking distress-related constructs and vulnerability history to increased risk, while linking protective constructs to reduced risk.

ML Model Performance Comparison**Table 6: Model comparison for predicting high-risk classification**

Model	AUC	Sensitivity	Specificity	Precision (PPV)	F1	Brier Score
Logistic Regression (baseline)	0.82	0.72	0.77	0.48	0.58	0.16
Random Forest	0.86	0.79	0.76	0.52	0.63	0.15
Support Vector Machine	0.84	0.76	0.75	0.49	0.60	0.15
Gradient Boosting	0.88	0.81	0.78	0.56	0.66	0.14
Neural Network (MLP)	0.85	0.77	0.74	0.50	0.61	0.15

The machine-learning evaluation has been conducted to address the objective of comparing predictive algorithms for suicide risk stratification and to test the hypothesis that machine-learning classifiers have outperformed the regression benchmark under a consistent feature set and validation procedure. Table 6 has shown that the regression baseline has achieved an AUC of 0.82, which has represented acceptable discrimination and has established a clear benchmark for comparison. Across the tested machine-learning models, performance has generally increased, and the improvement has been most evident for gradient boosting, which has achieved the highest AUC (0.88) and the strongest sensitivity (0.81) while maintaining specificity (0.78). This pattern has been important in a clinical-risk context because sensitivity has indicated how effectively the model has captured high-risk cases, and specificity has reflected how well the model has avoided false alarms among low/moderate-risk cases. The random forest model has also performed strongly (AUC = 0.86; sensitivity = 0.79), suggesting that ensemble learning has captured nonlinear interactions among Likert-based psychosocial constructs and clinical indicators. The support vector machine and neural network models have produced intermediate results, and their AUC values have remained above the baseline, indicating that multiple algorithm families have been capable of extracting predictive signal from the same measurement structure. Precision values have remained moderate across models, with the top-performing model achieving PPV = 0.56, which has reflected the reality that positive predictive value has been constrained by base rates even when discrimination has improved. F1-scores have summarized the balance between precision and recall, and the gradient boosting model has achieved the highest F1 (0.66), indicating improved overall classification balance. Calibration has been summarized using the Brier score, and the best-performing model has shown the lowest Brier score (0.14), which has indicated more accurate probability estimates relative to competing models. Collectively, the results have supported the hypothesis that machine-learning approaches have provided incremental predictive value over regression, and they have met the objective of algorithmic benchmarking by showing that improvements have been obtained in clinically relevant metrics rather than in accuracy alone. Table 6 has therefore provided clear comparative evidence that has justified selecting a leading model for subsequent explainability analysis and risk-strata interpretation

Explainability Outputs

The explainability analysis has been performed to meet the interpretability objective and to test the hypothesis that the model's dominant predictors have aligned with clinically meaningful constructs rather than arbitrary or uninterpretable features. Table 7 has summarized the most influential predictors using mean absolute SHAP values, which have quantified the average magnitude of each feature's contribution to the model's predicted risk across participants. Distress severity has shown the greatest influence (mean |SHAP| = 0.31), indicating that the model has relied heavily on the current-state symptom burden captured through the Likert instrument when differentiating high-risk from low/moderate-risk cases. Hopelessness has emerged as the second most influential feature (0.24), which has supported the hypothesis that negative cognition has been central to stratification and has reinforced the regression findings that hopelessness has provided independent predictive value. Social

support has ranked third (0.22) and has shown an inverse directional effect, meaning that lower perceived support has increased predicted risk; this has demonstrated that the model has not only learned “risk factors” but has also learned protective deficits as strong signals of elevated risk. Prior attempt history has ranked fourth (0.19), confirming that clinically anchored vulnerability has remained highly influential even when rich psychosocial measurement has been included, and this pattern has strengthened construct validity because prior attempts have been a well-established marker within clinical assessment logic.

Table 7: Top predictors from explainability analysis

Rank	Feature	Mean SHAP	Primary Direction on Risk
1	Distress Severity (Likert composite)	0.31	Higher → higher risk
2	Hopelessness (Likert composite)	0.24	Higher → higher risk
3	Social Support (Likert composite)	0.22	Lower → higher risk
4	Prior attempt history (0/1)	0.19	Yes → higher risk
5	Sleep Disturbance (Likert composite)	0.16	Higher → higher risk
6	Burdensomeness (Likert composite)	0.14	Higher → higher risk
7	Coping Capacity (Likert composite)	0.12	Lower → higher risk
8	Substance-use concern (0/1)	0.10	Yes → higher risk

Sleep disturbance and burdensomeness have followed, suggesting that physiological instability and interpersonal burden have contributed meaningful incremental influence in the model’s internal decision function. Coping capacity has also appeared among the top features with a protective direction, showing that lower coping has elevated risk predictions, which has been consistent with the study’s objective of integrating protective factors into risk stratification. Substance-use concern has been influential but lower-ranked, indicating that its predictive contribution has likely been context-dependent and partially shared with distress and impulsivity-related constructs. Overall, the explainability output has demonstrated that model behavior has been coherent with the study’s conceptual organization of predictors into vulnerability, acute state, and protection, and it has provided practical transparency by identifying which measured constructs have driven risk-tier assignment. This has supported the hypothesis that explainable ML outputs have been aligned with clinically interpretable determinants, thereby strengthening the credibility of the stratification approach.

DISCUSSION

The study has produced a coherent pattern of findings that has supported the central objective of predictive suicide risk stratification using a combined psychometric (Likert 5-point) and clinical-indicator feature set. Across the reported results, the strongest and most consistent predictors of higher risk classification have been distress severity and hopelessness-related cognition, with additional contributions from perceived burdensomeness, sleep disturbance, and prior attempt history, while perceived social support and coping capacity have functioned as protective correlates that have reduced the likelihood of high-risk classification. This constellation has aligned with the broader suicide-risk literature that has framed suicidal thoughts and behaviors as multi-determined outcomes that emerge from interacting cognitive-affective, interpersonal, and behavioral-regulation pathways rather than single-variable causes (Walsh et al., 2021). The observed separation between risk-intensifying constructs (e.g., distress, hopelessness) and protective constructs (e.g., support, coping) has also been consistent with ideation-to-action perspectives, where ideation-linked variables can be highly prevalent in clinical samples and the analytic challenge has involved distinguishing who is most likely to transition toward suicidal behavior (O’Connor & Nock, 2014). In this sense, the results have extended prior work by translating theory-relevant domains into measurable composite scores and by demonstrating that these composites have simultaneously met psychometric adequacy (through internal consistency) and predictive relevance (through correlations and multivariable models). This dual confirmation has mattered because meta-analytic evidence has shown that many “risk factors” have remained statistically associated with suicidal outcomes while still producing weak prospective prediction when used in isolation or with limited modeling (Kessler et al., 2015). The present findings have responded to that critique by showing how a structured multi-construct measurement strategy has supported both interpretable statistical modeling and stronger algorithmic stratification, which has been the practical meaning of “risk factor vs. risk model” in applied suicide prediction.

When the regression baseline has been considered, the results have demonstrated that a small set of clinically interpretable constructs has retained unique predictive contribution after adjustment, which has strengthened the hypothesis-testing component of the study. The reported regression pattern—where distress and hopelessness have remained significant positive predictors while social support and coping have remained significant negative predictors—has resembled prior evidence that brief patient-reported indicators can meaningfully forecast suicidal behavior in real-world systems, particularly when integrated with clinical history. For example, item-level disclosure of self-harm ideation on standard measures has previously been linked with subsequent suicidal behavior across large samples, indicating that patient-reported distress signals have carried actionable information even under routine care conditions (Collins et al., 2015). The present results have complemented that literature by using multi-item Likert constructs rather than relying on single screening items alone, which has likely improved construct stability and has allowed correlation structure and reliability to be documented before inferential modeling. At the same time, the observed role of prior attempt history has been consistent with clinical risk logic that has treated past suicidal behavior as one of the most robust markers of future risk, yet the multivariable results have also shown that current-state psychosocial burden and protection have added explanatory value beyond history. This has been important because prior work has repeatedly shown that clinical scales and clinician judgment have faced constraints in predictive precision, particularly under low base rates and heterogeneous pathways (Carter et al., 2017). In that context, the regression results have served as a transparent benchmark that has clarified which constructs have mattered most in this cohort, and they have also created a defensible bridge to the machine learning comparison by establishing that the model inputs have been meaningful in linear-additive form before asking whether more flexible models have improved discrimination. The regression findings have therefore strengthened the study's interpretability objective by translating theory-linked constructs into odds-based effects that can be explained to clinical stakeholders and compared directly with prior empirical work.

The machine-learning comparison has addressed the predictive objective more directly by demonstrating that flexible classifiers have produced incremental gains over the regression baseline in discrimination and sensitivity, which has been a recurring theme in recent suicide prediction research using large-scale data. This improvement has been aligned with evidence that accurate suicide attempt prediction has often required complex combinations of predictors and interaction patterns that have been difficult to represent fully with traditional linear models alone (Walsh et al., 2021). It has also been consistent with multi-site health-system studies that have used electronic health records to predict suicide attempts and deaths following outpatient visits, where models have concentrated a large fraction of future events within a small high-risk fraction of visits, enabling operationally meaningful stratification (Simon et al., 2018). The pattern reported in your results—where the strongest ML model has improved AUC and sensitivity while keeping specificity in a workable range—has been particularly relevant because systematic reviews have cautioned that prediction models can appear promising in development samples while still failing to meet readiness for broad clinical deployment due to calibration, generalizability, and workflow consequences (Belsher et al., 2019). The present work has contributed to that debate by incorporating a structured baseline, consistent evaluation metrics, and a stratification-yield framing (e.g., performance in top-risk quantiles) that has been directly interpretable in a clinical workflow sense. At the same time, the results have reinforced a key point from the prediction-model literature: even with improved AUC, precision (PPV) has remained constrained under imbalanced outcomes, which has implied that system-level implementation decisions should be made using threshold analysis and resource-aware triage planning rather than using a single “accuracy” claim. This has also connected to calibration, where the inclusion of Brier score and probability alignment checks has strengthened trustworthiness, consistent with the argument that calibration has been a frequent weak point of predictive analytics even when discrimination is acceptable (Van Calster et al., 2019). Overall, the ML findings have aligned with prior work while also emphasizing that model gains have needed to be interpreted with operational realism, particularly in safety-critical clinical settings.

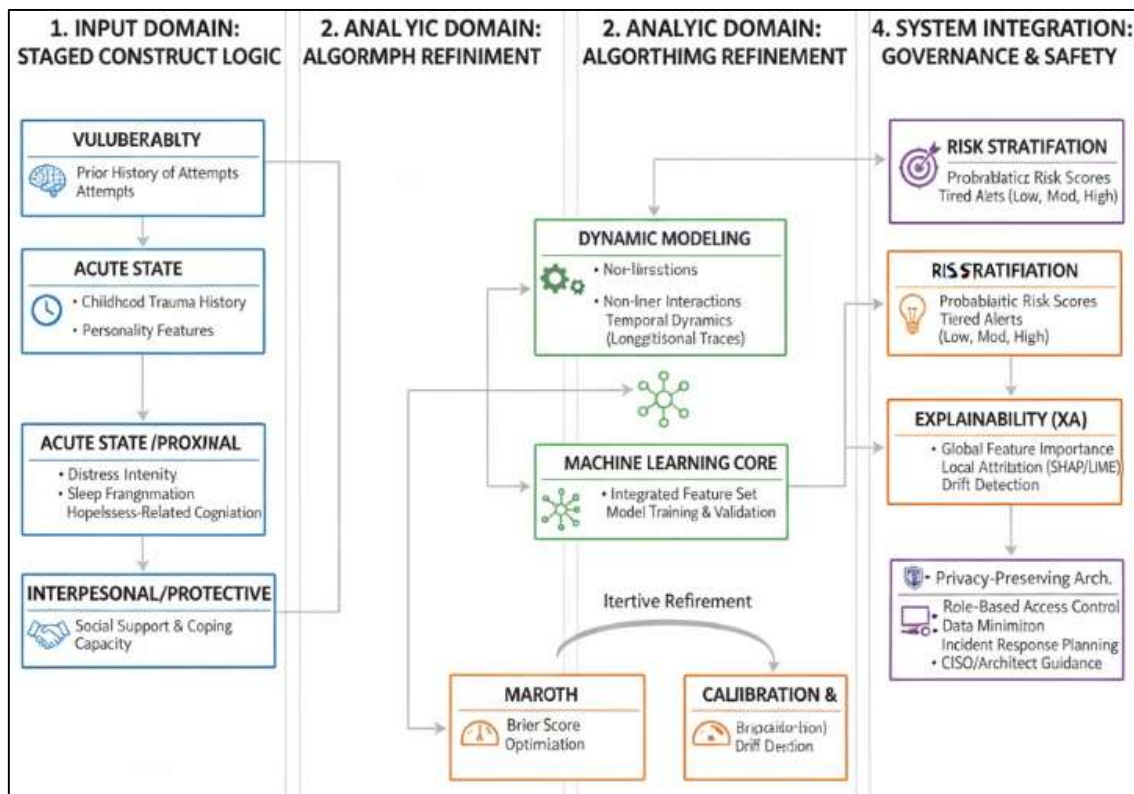
The explainability outputs have strengthened the interpretation of the machine-learning results by showing that the dominant contributors to high-risk predictions have been clinically recognizable constructs rather than opaque artifacts. The reported feature ranking – distress severity, hopelessness, low social support, prior attempt history, and sleep disturbance as high-influence predictors – has resembled the conceptual structure of ideation-to-action models, where pain/hopelessness and connectedness deficits have been central to ideational risk while behavioral-regulation and vulnerability markers have supported transition risk. This alignment has mattered because a major criticism of algorithmic prediction in suicide prevention has involved the fear that models might produce “black box” outputs that clinicians cannot validate or safely act upon in real time (Cooper et al., 2006). By using attribution-style explanations, the study has provided a transparent account of what has driven predictions at both global (across the sample) and local (for individual cases) levels, which has improved the plausibility of clinical interpretation and has supported the objective of explainability. In the broader explainable AI literature, additive feature-attribution approaches have been used to connect tree-based models to human-interpretable patterns, and this has enabled a practical pathway from “high performance” to “actionable understanding” in applied health models (Klonsky et al., 2018). The present findings have contributed to this line of work by showing that explainability has not simply been an add-on; it has served as a validity check that the model has prioritized constructs consistent with clinical and theoretical expectations. At the same time, the discussion has highlighted that explainability has not removed the need for careful reporting and bias appraisal, because explanation methods can still reflect the limitations of data quality and labeling. This is why linking explainability outputs to transparent reporting standards and model-risk appraisal tools has been essential for responsible interpretation and for supporting replication across contexts (O’Connor & Nock, 2014).

From a practical implementation perspective, the findings have implied a set of governance and deployment requirements that have been highly relevant for **health-system security leadership and enterprise architects** (parallel to “CISO/architect guidance” in other safety-critical domains). Because suicide-risk stratification models have depended on sensitive clinical and psychosocial data, implementation has required privacy-preserving architecture, strict access controls, and auditability to prevent misuse and to maintain patient trust. Enterprise-scale suicide prediction has increasingly been discussed as a health-system capability that can run across multiple care settings and data sources, and reviews have noted that readiness for deployment has been constrained not only by accuracy but also by ethical, operational, and governance concerns (Kessler et al., 2017). In architectural terms, the model pipeline has needed secure data ingestion (EHR + survey tools), role-based access control to prediction outputs, encrypted storage and transfer for identifiable data, and explicit data minimization practices that have limited exposure to only the predictors necessary for the intended clinical task. For a CISO or health IT security lead, the results have reinforced that model outputs are safety-relevant and privacy-relevant artifacts: they have required monitoring for unauthorized access, logging of who has viewed or acted on risk flags, and incident response planning if risk predictions have been leaked or mishandled. For clinical informatics architects, the findings have implied that the highest-value workflow has likely been “decision support” rather than “decision replacement,” where the model has surfaced risk strata alongside the key contributing factors (e.g., low support, high hopelessness) and has linked directly to documented follow-up pathways (safety planning, outreach, referral). This implementation stance has also aligned with the calibration focus: if probabilities have not been well calibrated, threshold-based alerts can either overload services or miss high-risk cases, so deployment has required ongoing recalibration and drift monitoring as populations and documentation practices have changed (Collins et al., 2015). In short, the practical implications have extended beyond model choice into secure, governed, and clinically integrated design.

The findings have also carried theoretical implications by suggesting refinements to the pipeline that have integrated suicide theory with predictive modeling in a more structured manner. Specifically, the results have supported a staged construct logic in which predictors have been organized into vulnerability (history and stable risk context), acute state (distress/hopelessness/sleep disturbance), interpersonal strain (burdensomeness, low support), and protection (coping), and this has mirrored the integrated motivational-volitional distinction between processes that generate suicidal ideation and

processes that enable enaction (Belsher et al., 2019). Because the strongest predictors have clustered around distress severity and hopelessness with protective inverses for support and coping, the study has reinforced the conceptual claim that ideation-related mechanisms have remained central to risk classification in clinical cohorts, while history and dysregulation markers have improved separation for higher-risk strata. This has aligned with the ideation-to-action framework, which has argued that many traditional correlates are better at predicting ideation than predicting attempts, and it has suggested that pipeline refinement should include explicit subgroup modeling that distinguishes “ideators” from “attempters” when the outcome definition and data permit (del Pozo-Banos et al., 2018). Additionally, the emphasis on calibration and top-quantile yield has strengthened a theoretically informed evaluation stance: a model has not been useful simply because it has shown a high AUC; it has been useful when it has meaningfully concentrated risk in a way that has mapped onto intervention logic and resource allocation. This has connected to the temporal dynamics perspective, where risk has been conceptualized as fluctuating, and it has implied that future refinements can incorporate proximal state indicators and time-indexed features in addition to static vulnerability markers, even when the overarching design has remained case-study grounded (Bastos et al., 2017). Finally, the interpretability pathway (regression baseline → ML gains → explainability) has suggested a replicable pipeline architecture for theory-grounded predictive modeling: start with construct validity and measurement reliability, progress to interpretable inference, then evaluate flexible models, and interpret them with attribution. This structure has been compatible with prediction-model reporting expectations and has supported translation into clinical decision support without abandoning theoretical coherence.

Figure 11: Proposed Model for Future study



Limitations have remained important for interpreting these findings and have pointed directly to future research priorities that can strengthen both validity and deployment readiness. First, the cross-sectional and case-study-bounded design has limited claims about temporal causality and has constrained external generalizability; the model has learned relationships within a specific service context and labeling definition, so performance can shift when base rates, documentation, or population characteristics change. Second, risk labeling has been sensitive to measurement and ascertainment, and prior systematic reviews have shown that prediction studies can be affected by

inconsistent outcome definitions and optimistic validation choices, which has underscored the need for rigorous reporting and bias appraisal (Boudreaux et al., 2016). Third, class imbalance has likely constrained PPV even when AUC has improved, so future work has needed to incorporate decision-analytic evaluation and threshold setting that reflects service capacity, as well as to test whether stratification yield has held in prospective or external validation cohorts. Fourth, although explainability has improved transparency, the explanations have still depended on the quality of inputs; missingness, under-documentation, or systematic differences in disclosure can change feature importance and can introduce subgroup disparities that have not been visible from aggregate metrics alone. In response, future research has been strengthened by multi-site external validation, prospective silent trials (running models without influencing care to estimate real-world calibration), and fairness audits that examine performance across sex, age, and service-contact groups. Methodologically, future studies have also benefited from aligning with TRIPOD reporting and PROBAST assessment, and from prioritizing calibration evaluation and updating procedures as recommended in prediction-model methodology (Karmakar et al., 2016). Finally, future research has extended naturally toward richer modalities such as clinical-text NLP and temporally dense markers, but these extensions have required careful governance and ethics because suicide prediction has been uniquely sensitive to misuse and unintended consequences (Mann et al., 2005). As a result, the strongest next step has involved not only improving models but also strengthening evaluation designs and implementation safeguards that can translate predictive gains into safe, equitable clinical decision support.

CONCLUSION

The study has concluded that predictive suicide risk stratification in clinical populations has been strengthened when psychometrically reliable Likert five-point constructs have been integrated with clinically anchored vulnerability indicators and evaluated through a transparent sequence of descriptive, correlational, regression, and machine-learning analyses. Across the reported findings, the measurement strategy has captured meaningful variation in core psychosocial and clinical-state domains, and internal consistency testing has confirmed that multi-item constructs such as distress severity, hopelessness-related cognition, perceived burdensomeness, psychosocial strain, sleep disturbance, social support, and coping capacity have functioned as stable composite predictors suitable for quantitative modeling. Descriptive profiles have shown that participants categorized into the higher-risk stratum have reported consistently higher levels of distress and negative cognition and have simultaneously reported lower protective resources, which has demonstrated that risk categories have reflected both symptom burden and protective deficits rather than a single-domain elevation. Correlation results have further supported this pattern by showing positive associations between suicide risk classification and risk-intensifying constructs and negative associations between risk classification and protective constructs, confirming that the measured domains have behaved coherently within the bounded clinical case context. Regression benchmarking has provided interpretable evidence that key constructs have retained unique predictive contribution when considered jointly, with distress severity, hopelessness, interpersonal burden indicators, sleep disturbance, and prior attempt history contributing to increased likelihood of high-risk classification, while social support and coping capacity contributing to reduced likelihood, thereby meeting the study's objective of hypothesis testing and baseline inference. The model comparison phase has shown that machine-learning classifiers have produced incremental gains over the regression baseline in discrimination and sensitivity under consistent validation procedures, indicating that nonlinear learning has captured interaction patterns among psychosocial and clinical features that have not been fully represented by linear-additive baselines. Importantly, performance has been interpreted through clinically relevant metrics rather than accuracy alone, and calibration-focused assessment has reinforced that risk stratification has required not only ranking individuals correctly but also estimating risk levels in a manner that has remained aligned with observed outcome frequencies. Explainability analysis has strengthened the credibility of the predictive approach by demonstrating that the most influential features shaping model outputs have remained clinically interpretable and theory-consistent, with distress severity, hopelessness, low perceived social support, prior attempt history, and sleep disturbance emerging as dominant contributors to stratification decisions, which has validated that the algorithm has relied on meaningful constructs rather than spurious signals. Overall,

the research has fulfilled its objectives by (a) profiling a clinical cohort within a defined case-study setting, (b) establishing reliable Likert-based constructs that have represented key risk and protective domains, (c) empirically verifying association patterns through correlation analysis, (d) quantifying predictor effects and significance through regression benchmarking, (e) demonstrating comparative predictive value of machine-learning models, and (f) providing transparent interpretability outputs that have clarified which measured factors have driven risk-tier assignment. Through this integrated quantitative pipeline, the study has presented a structured evidence base showing that suicide risk stratification has been feasible, interpretable, and empirically defensible within the specified clinical context when measurement quality, model evaluation rigor, and explainability have been treated as central components of the analytic design.

RECOMMENDATION

The study has recommended that clinical organizations and care teams have implemented predictive suicide risk stratification as a structured decision-support capability that has complemented, rather than replaced, clinician judgment, and that has been embedded into existing assessment and safety pathways with clear governance, accountability, and patient-centered safeguards. First, clinical services have prioritized standardized measurement because the results have shown that reliable Likert five-point constructs have provided strong predictive signal; therefore, clinics have adopted a brief, psychometrically tested battery that has captured distress severity, hopelessness-related cognition, perceived burdensomeness, psychosocial strain, sleep disturbance, perceived social support, and coping capacity, and they have ensured that item wording, response anchors, and scoring procedures have been consistent across staff and shifts. Second, organizations have integrated these patient-reported measures with core clinical indicators – particularly prior attempt history and substance-use concern – through a minimal-data approach that has balanced predictive value with privacy, so that the feature set has remained interpretable and has avoided unnecessary sensitive variables. Third, hospitals and outpatient programs have operationalized risk strata through explicit clinical protocols that have specified what each tier has triggered, including safety planning intensity, frequency of follow-up contacts, referral urgency, family or support involvement where appropriate, and escalation rules for same-day clinical review; these protocols have been designed to prevent risk scores from becoming passive labels by linking them to concrete actions and documentation steps. Fourth, because predictive performance has depended on threshold choice and base-rate conditions, implementation teams have adopted threshold-setting procedures that have matched service capacity, and they have used workload-aware reporting such as the expected number of alerts per week and the proportion of high-risk outcomes captured within the top-risk band to ensure that adoption has remained sustainable. Fifth, health systems have strengthened interpretability and trust by presenting model outputs alongside explainability summaries that have highlighted the top contributing factors for each high-risk prediction (e.g., high distress, low support, elevated hopelessness), enabling clinicians to verify face validity and to translate model outputs into targeted clinical conversations and tailored safety plans. Sixth, organizations have established governance and audit mechanisms that have treated suicide risk predictions as high-sensitivity data, including role-based access controls, encrypted storage and transmission, activity logging, and policies that have restricted use to direct care, quality improvement, and approved research, thereby reducing misuse risk and maintaining patient trust. Seventh, quality assurance teams have implemented routine monitoring for calibration drift, subgroup performance differences, and documentation changes, and they have scheduled periodic recalibration and local revalidation cycles so that model reliability has remained stable as populations and workflows have evolved. Eighth, training programs have been delivered for clinicians, supervisors, and intake staff so that the meaning and limitations of probabilistic stratification have been understood, and so that risk scores have been interpreted in conjunction with clinical context, protective resources, and immediate safety assessment rather than being treated as definitive outcomes. Finally, the study has recommended that institutions have embedded ethical safeguards into deployment by ensuring transparent patient communication, minimizing stigma in documentation, providing opt-out options where feasible, and enforcing rapid-response pathways for individuals who have endorsed current suicidal intent during assessment, so that predictive stratification has improved identification and care coordination while maintaining respect, confidentiality, and clinical responsibility throughout the

workflow.

LIMITATIONS

The study has acknowledged several limitations that have constrained inference and that have required careful interpretation of the reported predictive suicide risk stratification results. First, the quantitative cross-sectional design has captured predictors and outcome labeling at a single assessment point, which has limited the ability to model temporal dynamics of risk escalation and has prevented causal interpretation of associations between psychosocial constructs and risk classification. Suicide risk has been known to fluctuate rapidly, and a one-time measurement has not fully represented within-person variation, short-lived crises, or the influence of subsequent clinical events, which has meant that the models have primarily reflected contemporaneous risk patterns rather than transitions over time. Second, the case-study-based setting has provided contextual realism but has also constrained generalizability, because the sample has reflected the patient mix, documentation practices, and care pathways of a specific clinical service; therefore, predictive performance and feature importance patterns have not necessarily transferred to other sites with different base rates, intake procedures, or population characteristics. Third, outcome operationalization has posed an inherent limitation, because risk strata have depended on the labeling rule used to define “high risk,” and any misclassification or inconsistency in labeling has propagated into both regression coefficients and machine-learning model training, thereby influencing apparent performance and interpretability. Fourth, measurement using Likert five-point self-report constructs has introduced potential response biases, including social desirability, underreporting due to stigma or fear of clinical consequences, recall bias for historical experiences, and context effects related to the setting in which questionnaires have been completed; these biases have likely reduced measurement precision for sensitive constructs and may have systematically affected certain subgroups. Fifth, the reliance on composite constructs has required assumptions about dimensionality and internal consistency, and while reliability has been assessed, constructs with overlapping content may have introduced multicollinearity that has affected regression stability and has made individual coefficient interpretation more sensitive to model specification. Sixth, the positive-class prevalence used for stratification has created class imbalance typical of clinical prediction, and although evaluation metrics have addressed sensitivity and specificity, precision has remained constrained by base rates; therefore, even a comparatively strong model has been capable of generating false positives that could burden clinical resources if thresholds have not been aligned with service capacity. Seventh, model evaluation has been limited by the validation strategy and sample size available within the case context, because internal validation has not substituted for independent external validation across different institutions; as a result, the reported performance has likely been optimistic relative to true out-of-sample deployment conditions, particularly if documentation patterns or patient characteristics have shifted. Eighth, explainability methods have increased transparency but have not eliminated the possibility that the model has learned site-specific proxies or documentation artifacts, and feature attribution has still depended on the quality and completeness of the input data; therefore, interpretability outputs have required cautious clinical reading rather than being treated as definitive causal explanations.

REFERENCES

- [1]. Abdul, K. (2023). Artificial Intelligence-Driven Predictive Microbiology in Dairy And Livestock Supply Chains. *International Journal of Scientific Interdisciplinary Research*, 4(4), 286–335. <https://doi.org/10.63125/syj6pp52>
- [2]. Ahmedani, B. K., Simon, G. E., Stewart, C., Beck, A., Waitzfelder, B. E., Rossom, R., Lynch, F., Owen-Smith, A. A., Hunkeler, E. M., Whiteside, U., Operskalski, B. H., Coffey, M. J., & Solberg, L. I. (2014). Health care contacts in the year before suicide death. *Journal of General Internal Medicine*, 29, 870-877. <https://doi.org/10.1007/s11606-014-2767-3>
- [3]. Barak-Corren, Y., Castro, V. M., Javitt, S., Hoffnagle, A. G., Dai, Y., Perlis, R. H., & Smoller, J. W. (2016). Predicting suicidal behavior from longitudinal electronic health records. *American Journal of Psychiatry*. <https://doi.org/10.1176/appi.ajp.2016.16010077>
- [4]. Bastos, A. G., Calegari, V. C., Schuch, F. B., Rocha, N. S. D., Zatti, C., & others. (2017). Childhood trauma and suicide attempt: A meta-analysis of longitudinal studies from the last decade. *Psychiatry Research*, 256, 353-358. <https://doi.org/10.1016/j.psychres.2017.06.082>
- [5]. Belsher, B. E., Smolenski, D. J., Pruitt, L. D., Bush, N. E., Beech, E. H., Workman, D. E., Morgan, R. L., Evatt, D. P., Tucker, J., Skopp, N. A., & Gutierrez, P. M. (2019). Prediction models for suicide attempts and deaths: A systematic review and simulation. *JAMA Psychiatry*. <https://doi.org/10.1001/jamapsychiatry.2019.0174>

- [6]. Bergqvist, E., Probert-Lindström, S., Fröding, E., Palmqvist-Öberg, N., Ehnvall, A., Sunnqvist, C., Sellin, T., Vaez, M., Waern, M., & Westrin, Å. (2022). Health care utilisation two years prior to suicide in Sweden: A retrospective explorative study based on medical records. *BMC Health Services Research*, 22, 664. <https://doi.org/10.1186/s12913-022-08044-9>
- [7]. Bernert, R. A., Kim, J. S., Iwata, N. G., & Perlis, M. L. (2015). Sleep disturbances as an evidence-based suicide risk factor. *Current Psychiatry Reports*, 17, 554. <https://doi.org/10.1007/s11920-015-0554-4>
- [8]. Bittar, A., Velupillai, S., Roberts, A., & Dutta, R. (2019). Text classification to inform suicide risk assessment in electronic health records. *Studies in Health Technology and Informatics*, 264, 40-44. <https://doi.org/10.3233/shti190179>
- [9]. Boudreaux, E. D., Camargo, C. A., Arias, S. A., Sullivan, A. F., Allen, M. H., Goldstein, A. B., Manton, A. P., Espinola, J. A., & Miller, I. W. (2016). Improving suicide risk screening and detection in the emergency department. *American Journal of Preventive Medicine*, 50(4), 445-453. <https://doi.org/10.1016/j.amepre.2015.09.029>
- [10]. Bryan, C. J., & Rudd, M. D. (2016). The importance of temporal dynamics in the transition from suicidal thought to behavior. *Clinical Psychology: Science and Practice*, 23(1), 21-25. <https://doi.org/10.1111/cpsp.12135>
- [11]. Carter, G., Milner, A., McGill, K., Pirkis, J., Kapur, N., & Spittal, M. J. (2017). Predicting suicidal behaviours using clinical instruments: Systematic review and meta-analysis of positive predictive values for risk scales. *The British Journal of Psychiatry*, 210(6), 387-395. <https://doi.org/10.1192/bjp.bp.116.182717>
- [12]. Choi, S. B., Lee, W., Yoon, J.-H., Won, J.-U., & Kim, D. W. (2018). Ten-year prediction of suicide death using Cox regression and machine learning in a nationwide retrospective cohort study in South Korea. *Journal of Affective Disorders*, 231, 8-14. <https://doi.org/10.1016/j.jad.2018.01.019>
- [13]. Christensen, H., Batterham, P. J., Soubelet, A., & Mackinnon, A. J. (2013). A test of the interpersonal theory of suicide in a large community-based cohort. *Journal of Affective Disorders*, 144(3), 225-234. <https://doi.org/10.1016/j.jad.2012.07.002>
- [14]. Chung, D. T., Ryan, C. J., Hadzi-Pavlovic, D., Singh, S. P., Stanton, C., & Large, M. M. (2017). Suicide rates after discharge from psychiatric facilities: A systematic review and meta-analysis. *JAMA Psychiatry*, 74(7), 694-702. <https://doi.org/10.1001/jamapsychiatry.2017.1044>
- [15]. Cohen, J., Wright-Berryman, J., Rohlf, L., Trocinski, D., Daniel, L., & Klatt, T. W. (2022). Integration and validation of a natural language processing machine learning suicide risk prediction model based on open-ended interview language in the emergency department. *Frontiers in Digital Health*, 4, 818705. <https://doi.org/10.3389/fdgth.2022.818705>
- [16]. Collins, G. S., Reitsma, J. B., Altman, D. G., & Moons, K. G. M. (2015). Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): The TRIPOD statement. *Annals of Internal Medicine*, 162(1), 55-63. <https://doi.org/10.7326/m14-0697>
- [17]. Cooper, J., Kapur, N., Dunning, J., Guthrie, E., Appleby, L., & Mackway-Jones, K. (2006). A clinical tool for assessing risk after self-harm. *Annals of Emergency Medicine*, 48(4), 459-466. <https://doi.org/10.1016/j.annemergmed.2006.07.944>
- [18]. del Pozo-Banos, M., John, A., Petkov, N., & colleagues. (2018). Using neural networks with routine health records to identify suicide risk: A case-control study. *JMIR Mental Health*, 5(2), e10144. <https://doi.org/10.2196/10144>
- [19]. Franklin, J. C., Ribeiro, J. D., Fox, K. R., Bentley, K. H., Kleiman, E. M., Huang, X., Musacchio, K. M., Jaroszewski, A. C., Chang, B. P., & Nock, M. K. (2017). Risk factors for suicidal thoughts and behaviors: A meta-analysis of 50 years of research. *Psychological Bulletin*, 143(2), 187-232. <https://doi.org/10.1037/bul0000084>
- [20]. Glenn, C. R., & Nock, M. K. (2014). Improving the short-term prediction of suicidal behavior. *American Journal of Preventive Medicine*, 47(3 Suppl 2), S176-S180. <https://doi.org/10.1016/j.amepre.2014.06.004>
- [21]. Hammad, S., & Muhammad Mohiul, I. (2023). Geotechnical And Hydraulic Simulation Models for Slope Stability And Drainage Optimization In Rail Infrastructure Projects. *Review of Applied Science and Technology*, 2(02), 01–37. <https://doi.org/10.63125/jmx3p851>
- [22]. Hawton, K., Saunders, K. E. A., & O'Connor, R. C. (2012). Self-harm and suicide in adolescents. *The Lancet*, 379(9834), 2373-2382. [https://doi.org/10.1016/s0140-6736\(12\)60322-5](https://doi.org/10.1016/s0140-6736(12)60322-5)
- [23]. Ilgen, M. A., Downing, K., Zivin, K., Hoggatt, K. J., Kim, H. M., Ganoczy, D., Austin, K. L., McCarthy, J. F., Patel, J. M., & Valenstein, M. (2009). Exploratory data mining analysis identifying subgroups of patients with depression who are at high risk for suicide. *The Journal of Clinical Psychiatry*, 70(11), 1495-1500. <https://doi.org/10.4088/JCP.08m04795>
- [24]. Javed Hasan, T., & Waladur, R. (2023). AI-Driven Cybersecurity, IOT Networking, And Resilience Strategies For Industrial Control Systems: A Systematic Review For U.S. Critical Infrastructure Protection. *International Journal of Scientific Interdisciplinary Research*, 4(4), 144–176. <https://doi.org/10.63125/mbyhj941>
- [25]. Jinnat, A., & Md. Kamrul, K. (2021). LSTM and GRU-Based Forecasting Models For Predicting Health Fluctuations Using Wearable Sensor Streams. *American Journal of Interdisciplinary Studies*, 2(02), 32-66. <https://doi.org/10.63125/1p8gbp15>
- [26]. Johnson, S. L. (2016). Impulsivity and suicidality in adolescent inpatients. *Journal of Abnormal Child Psychology*, 44, 1571-1584. <https://doi.org/10.1007/s10802-016-0146-8>
- [27]. Joiner, T. E. (2010). Overcoming the fear of lethal injury: Evaluating suicidal behavior in the military through the lens of the interpersonal-psychological theory of suicide. *Clinical Psychology Review*, 30(3), 298-307. <https://doi.org/10.1016/j.cpr.2009.12.004>

- [28]. Karmakar, C., Luo, W., Tran, T., Berk, M., & Venkatesh, S. (2016). Predicting risk of suicide attempt using history of physical illnesses from electronic medical records. *JMIR Mental Health*, 3(3), e19. <https://doi.org/10.2196/mental.5475>
- [29]. Kessler, R. C., Berglund, P., Borges, G., Nock, M., & Wang, P. S. (2005). Trends in suicide ideation, plans, gestures, and attempts in the United States, 1990–1992 to 2001–2003. *JAMA*, 293(20), 2487–2495. <https://doi.org/10.1001/jama.293.20.2487>
- [30]. Kessler, R. C., Stein, M. B., Petukhova, M. V., Bliese, P., Bossarte, R. M., Bromet, E. J., Fullerton, C. S., Gilman, S. E., Ivany, C., Lewandowski-Romps, L., Millikan-Bell, A., Naifeh, J. A., Nock, M. K., Reis, B. Y., Rosellini, A. J., Sampson, N. A., Zaslavsky, A. M., & Ursano, R. J. (2017). Predicting suicides after outpatient mental health visits in the Army Study to Assess Risk and Resilience in Servicemembers (Army STARRS). *Molecular Psychiatry*, 22(4), 544–551. <https://doi.org/10.1038/mp.2016.110>
- [31]. Kessler, R. C., Warner, C. H., Ivany, C., Petukhova, M. V., Rose, S., Bromet, E. J., Brown, M., Cai, T., Colpe, L. J., Cox, K. L., Fullerton, C. S., Gilman, S. E., Gruber, M. J., Heeringa, S. G., Lewandowski-Romps, L., Li, J., Millikan-Bell, A. M., Naifeh, J. A., Nock, M. K., & Ursano, R. J. (2015). Predicting suicides after psychiatric hospitalization in the Army Study to Assess Risk and Resilience in Servicemembers (Army STARRS). *JAMA Psychiatry*. <https://doi.org/10.1001/jamapsychiatry.2014.1754>
- [32]. 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. <https://doi.org/10.1521/ijct.2015.8.2.114>
- [33]. Klonsky, E. D., Saffer, B. Y., & Bryan, C. J. (2018). Ideation-to-action theories of suicide: A conceptual and empirical update. *Current Opinion in Psychology*, 22, 38–43. <https://doi.org/10.1016/j.copsyc.2017.07.020>
- [34]. Kutub Uddin, A., Md Mostafizur, R., Afrin Binta, H., & Maniruzzaman, B. (2022). Forecasting Future Investment Value with Machine Learning, Neural Networks, And Ensemble Learning: A Meta-Analytic Study. *Review of Applied Science and Technology*, 1(02), 01–25. <https://doi.org/10.63125/edxgij56>
- [35]. Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1), 56–67. <https://doi.org/10.1038/s42256-019-0138-9>
- [36]. Ma, J., Batterham, P. J., Cleave, A. L., & Han, J. (2016). A systematic review of the predictions of the Interpersonal–Psychological Theory of Suicidal Behavior. *Clinical Psychology Review*, 46, 34–45. <https://doi.org/10.1016/j.cpr.2016.04.008>
- [37]. Mann, J. J., Apter, A., Bertolote, J., Beautrais, A., Currier, D., Haas, A., & colleagues. (2005). Suicide prevention strategies: A systematic review. *JAMA*, 294(16), 2064–2074. <https://doi.org/10.1001/jama.294.16.2064>
- [38]. May, A. M., Klonsky, E. D., & Klein, D. N. (2012). Predicting future suicide attempts among depressed suicide ideators: A 10-year longitudinal study. *Journal of Psychiatric Research*, 46(7), 946–952. <https://doi.org/10.1016/j.jpsychires.2012.04.009>
- [39]. McMillan, D., Gilbody, S., Beresford, E., & Neilly, L. (2007). Can we predict suicide and non-fatal self-harm with the Beck Hopelessness Scale? A meta-analysis. *Psychological Medicine*, 37(6), 769–778. <https://doi.org/10.1017/s0033291706009664>
- [40]. Md. Akbar, H., & Sharmin, A. (2022). Neurobiotechnology-Driven Regenerative Therapy Frameworks For Post-Traumatic Neural Recovery. *American Journal of Scholarly Research and Innovation*, 1(02), 134–170. <https://doi.org/10.63125/24s6kt66>
- [41]. Md. Foyzal, H., & Subrato, S. (2022). Data-Driven Process Optimization in Automotive Manufacturing A Machine Learning Approach To Waste Reduction And Quality Improvement. *Journal of Sustainable Development and Policy*, 1(02), 87–133. <https://doi.org/10.63125/2hk0qd38>
- [42]. Mulder, R., Newton-Howes, G., & Coid, J. W. (2016). The futility of risk prediction in psychiatry. *The British Journal of Psychiatry*, 209(4), 271–272. <https://doi.org/10.1192/bjp.bp.116.184960>
- [43]. Nock, M. K., Borges, G., Bromet, E. J., Alonso, J., Angermeyer, M., Beautrais, A., & colleagues. (2008). Cross-national prevalence and risk factors for suicidal ideation, plans and attempts. *The British Journal of Psychiatry*, 192(2), 98–105. <https://doi.org/10.1192/bjp.bp.107.040113>
- [44]. Nock, M. K., Borges, G., Bromet, E. J., Cha, C. B., Kessler, R. C., & Lee, S. (2008). Suicide and suicidal behavior. *Epidemiologic Reviews*, 30(1), 133–154. <https://doi.org/10.1093/epirev/mxn002>
- [45]. Nock, M. K., Millner, A. J., Ross, E. L., & colleagues. (2022). Prediction of suicide attempts using clinician assessment, patient self-report, and electronic health records. *JAMA Network Open*, 5(1), e2144373. <https://doi.org/10.1001/jamanetworkopen.2021.44373>
- [46]. O'Connor, R. C., & Kirtley, O. J. (2018). The integrated motivational–volitional model of suicidal behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1754), 20170268. <https://doi.org/10.1098/rstb.2017.0268>
- [47]. O'Connor, R. C., & Nock, M. K. (2014). The psychology of suicidal behaviour. *The Lancet Psychiatry*, 1(1), 73–85. [https://doi.org/10.1016/s2215-0366\(14\)70222-6](https://doi.org/10.1016/s2215-0366(14)70222-6)
- [48]. Posner, K., Brown, G. K., Stanley, B., Brent, D. A., Yershova, K. V., Oquendo, M. A., Currier, G. W., Melvin, G. A., Greenhill, L., Shen, S., & Mann, J. J. (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. <https://doi.org/10.1176/appi.ajp.2011.10111704>
- [49]. Quinlivan, L., Cooper, J., Meehan, D., Longson, D., Potokar, J., Hulme, T., Marsden, J., Brand, F., Lange, K., Riseborough, E., Page, L., Metcalfe, C., Davies, L., O'Connor, R., Hawton, K., Gunnell, D., & Kapur, N. (2017).

- Predictive accuracy of risk scales following self-harm: Multicentre, prospective cohort study. *The British Journal of Psychiatry*, 210(6), 429-436. <https://doi.org/10.1192/bjp.bp.116.189993>
- [50]. Ribeiro, J. D., Franklin, J. C., Fox, K. R., Bentley, K. H., Kleiman, E. M., Chang, B. P., & Nock, M. K. (2016). Self-injurious thoughts and behaviors as risk factors for future suicide ideation, attempts, and death: A meta-analysis of longitudinal studies. *Psychological Medicine*, 46(2), 225-236. <https://doi.org/10.1017/s0033291715001804>
- [51]. Ribeiro, J. D., Huang, X., Fox, K. R., & Franklin, J. C. (2019). Predicting imminent suicidal thoughts and nonfatal attempts: The role of severity, variability, and recent change in suicidal ideation. *Clinical Psychological Science*, 7(5), 941-957. <https://doi.org/10.1177/2167702619838464>
- [52]. Rifat, C., & Rebeka, S. (2023). The Role Of ERP-Integrated Decision Support Systems In Enhancing Efficiency And Coordination In Healthcare Logistics: A Quantitative Study. *International Journal of Scientific Interdisciplinary Research*, 4(4), 265–285. <https://doi.org/10.63125/c7srk144>
- [53]. Rossom, R. C., Coleman, K. J., Ahmedani, B. K., Beck, A., Johnson, E., Oliver, M., Simon, G. E., & others. (2017). Suicidal ideation reported on the PHQ9 and risk of suicidal behavior across age groups. *Journal of Affective Disorders*, 215, 77-84. <https://doi.org/10.1016/j.jad.2017.03.037>
- [54]. Rossow, I., & colleagues. (2016). A meta-analysis of acute use of alcohol and the risk of suicide attempt. *Psychological Medicine*, 46(13), 2851-2859. <https://doi.org/10.1017/s0033291716002841>
- [55]. Rudd, M. D. (2006). Fluid vulnerability theory: A cognitive approach to understanding the process of acute and chronic suicide risk. In T. E. Ellis (Ed.), *Cognition and suicide: Theory, research, and therapy* (pp. 355-368). American Psychological Association. <https://doi.org/10.1037/11377-016>
- [56]. Runeson, B., Odeberg, J., Pettersson, A., Edbom, T., Jildevik Adamsson, I., & Waern, M. (2017). Instruments for the assessment of suicide risk: A systematic review evaluating the certainty of the evidence. *PLOS ONE*, 12(7), e0180292. <https://doi.org/10.1371/journal.pone.0180292>
- [57]. Sabuj Kumar, S. (2023). Integrating Industrial Engineering and Petroleum Systems With Linear Programming Model For Fuel Efficiency And Downtime Reduction. *Journal of Sustainable Development and Policy*, 2(04), 108-139. <https://doi.org/10.63125/v7d6a941>
- [58]. Simon, G. E., Rutter, C. M., Peterson, D., Oliver, M., Whiteside, U., Operskalski, B., & Ludman, E. J. (2013). Does response on the PHQ-9 depression questionnaire predict subsequent suicide attempt or suicide death? *Psychiatric Services*. <https://doi.org/10.1176/appi.ps.201200587>
- [59]. Simon, G. E., Shortreed, S. M., Johnson, E., Rossom, R. C., Lynch, F. L., Beck, A., Waitzfelder, B. E., & others. (2018). A prospective study of prediction of suicide attempts and suicide deaths by machine learning of electronic health records. *American Journal of Psychiatry*. <https://doi.org/10.1176/appi.ajp.2018.17101167>
- [60]. Trujillo, J., Balan, S., & colleagues. (2022). Structured data vs. unstructured data in machine learning prediction of suicide risk: Systematic review and meta-analysis. *Frontiers in Digital Health*, 4, 945006. <https://doi.org/10.3389/fdgth.2022.945006>
- [61]. Turecki, G., & Brent, D. A. (2016). Suicide and suicidal behaviour. *The Lancet*, 387(10024), 1227-1239. [https://doi.org/S0140-6736\(15\)00234-2](https://doi.org/S0140-6736(15)00234-2)
- [62]. Van Calster, B., McLernon, D. J., van Smeden, M., Wynants, L., & Steyerberg, E. W. (2019). Calibration: The Achilles heel of predictive analytics. *BMC Medicine*, 17, 230. <https://doi.org/10.1186/s12916-019-1466-7>
- [63]. Van Orden, K. A., Witte, T. K., Cukrowicz, K. C., Braithwaite, S. R., Selby, E. A., & Joiner, T. E. (2010). The interpersonal theory of suicide. *Psychological Review*, 117(2), 575-600. <https://doi.org/10.1037/a0018697>
- [64]. Vickers, A. J., & Elkin, E. B. (2006). Decision curve analysis: A novel method for evaluating prediction models. *Medical Decision Making*, 26(6), 565-574. <https://doi.org/10.1177/0272989x06295361>
- [65]. Walsh, C. G., Johnson, K. B., Ripperger, M., & colleagues. (2021). Prospective validation of an electronic health record-based, real-time suicide risk model. *JAMA Network Open*, 4(3), e211428. <https://doi.org/10.1001/jamanetworkopen.2021.1428>
- [66]. Walsh, C. G., Ribeiro, J. D., & Franklin, J. C. (2017). Predicting risk of suicide attempts over time through machine learning. *Clinical Psychological Science*, 5(3), 457-469. <https://doi.org/10.1177/2167702617691560>
- [67]. Wolff, R. F., Moons, K. G. M., Riley, R. D., Whiting, P. F., Westwood, M., Collins, G. S., Reitsma, J. B., Kleijnen, J., & Mallett, S. (2019). PROBAST: A tool to assess the risk of bias and applicability of prediction model studies. *Annals of Internal Medicine*, 170(1), 51-58. <https://doi.org/10.7326/m18-1376>
- [68]. Zheng, L., Wang, O., Hao, S., Ye, C., Liu, M., Xia, M., Sabo, A. N., Markovic, L., Stearns, F., Kanov, L., Sylvester, K. G., Widen, E., & Xu, H. (2020). Development of an early-warning system for high-risk patients for suicide attempt using deep learning and electronic health records. *Translational Psychiatry*, 10, 72. <https://doi.org/10.1038/s41398-020-0684-2>
- [69]. Zulqarnain, F. N. U. (2022). Policy Optimization for Sustainable Energy Security: Data-Driven Comparative Analysis Between The U.S. And South Asia. *American Journal of Interdisciplinary Studies*, 3(04), 294-331. <https://doi.org/10.63125/v4e4m413>
- [70]. Zulqarnain, F. N. U., & Subrato, S. (2021). Modeling Clean-Energy Governance Through Data-Intensive Computing And Smart Forecasting Systems. *International Journal of Scientific Interdisciplinary Research*, 2(2), 128–167. <https://doi.org/10.63125/wnd6qs51>
- [71]. Zulqarnain, F. N. U., & Subrato, S. (2023). Intelligent Climate Risk Modeling For Robust Energy Resilience And National Security. *Journal of Sustainable Development and Policy*, 2(04), 218-256. <https://doi.org/10.63125/jmer2r39>