

“How’d You Type That So Fast?” A Descriptive Analysis of Counselor Message Text Reuse in Text-Based Crisis Counseling

Stevi N Gligorovic¹ and Jens Kristian R Schou² and Zac E Imel^{1,3}
and Brent Kious^{3,4}

¹Department of Educational Psychology, University of Utah

²Scientific Computing and Imaging Institute, University of Utah

³Department of Psychiatry, University of Utah

⁴Huntsman Mental Health Institute, University of Utah

Abstract

Suicide is a major public health concern, underscoring the importance of understanding communication practices used in crisis intervention. Text-based crisis services are increasingly used, yet little is known about how counselors construct messages across encounters. One understudied feature of this setting is counselor text reuse, or the repeated use of identical or highly similar message content across different clients. Although reuse may support efficiency and consistency, it may raise questions about how personalized responses are across counselors. This study provides a descriptive analysis of counselor text reuse in a large dataset of 4.7 million messages of real-time text-based crisis counseling conversations. Across 136 counselors, mean message similarity was very low, indicating little overall text reuse for most counselors. However, 103 counselors showed at least one instance of detected reuse, and a smaller subset demonstrated more consistent reuse. Reuse was also positively associated with counselor encounter volume across measures of reuse. Frequently reused longer passages primarily involved structured coping-oriented or psychoeducational content, such as coping strategies, grounding exercises, self-care tips, and relaxation techniques. The findings suggest that counselor text reuse is uncommon overall but becomes more likely as counselor encounter volume increases, providing a foundation for future work examining associations with service delivery and client outcomes.

Introduction

Suicide is a serious public health issue, with about 800,000 deaths due to suicide annually (World Health Organization, 2019) and several national declarations concerning children’s mental health (Hua et al., 2023). In the United States, suicide is the second leading cause of death for adolescents and young adults (10–24 years old). In 2022, the American Foundation for Suicide Prevention

reported about 1.6 million suicide attempts and 49,476 deaths by suicide. Notably, adolescent suicide rates have risen (Curtin, 2020). From 1975 to 2016, suicide rates increased by 7.9% among females and 3.5% among males ages 15–19 (Ruch, 2019).

Addressing this crisis requires robust interventions that include restricting access to lethal means, psychotherapy, and crisis hotlines (Goldney, 2005). Crisis hotlines are an accessible, widely utilized public health resource that provide immediate support to at-risk individuals (Gould et al., 2021; Ramchand et al., 2017). With evolving communication preferences, text-based crisis helplines are increasingly integrated into mental health services, particularly for younger populations (Budinger et al., 2015). For example, 988, the national crisis line, has received millions of text messages since its inception in July 2022 (Substance Abuse and Mental Health Services Administration, 2025). Despite the rise of text-based crisis services, the unique affordances and challenges of text-based crisis counseling conversations are relatively under-studied. For example, text-based communication allows counselors to copy-and-paste previously typed messages with important information, facilitating efficiency but potentially diminishing the personalized nature of interactions.

Crisis Lines

Participation in crisis line interactions can lead to reduced suicide risk (Gould et al., 2007; Halford et al., 2021; Mishara and Daigle, 2001). Similarly, text-based interventions are found to be helpful across a multitude of studies (Dwyer et al., 2021). For example, 90% of suicidal texters reported their conversation was helpful and 46.1% reported being less suicidal (Gould et al., 2007); see also (King et al., 2003; Mishara and Daigle, 1997; Pauwels et al., 2025). However, these findings suggest that over half of texters do not experience a reduction

in risk.

Research on crisis counseling outcomes has historically focused on counselor variation in key behaviors such as risk assessment, safety planning, and counselor intervention strategies (Gould et al., 2013; Mishara, 2007; Lake et al., 2022). For example, compliance with suicide-risk assessment guidelines remains inconsistent across crisis hotlines, despite national standards (Gould et al., 2013; Mishara, 2007; Lake et al., 2022). In a crisis hotline study with U.S. veterans, only one-third of eligible callers completed a safety plan (Lauver et al., 2024). In general, more effective crisis counseling emphasizes validation, rapport-building, and collaborative problem-solving (Mishara, 2007). In sum, there is evidence that key behaviors are associated with important outcomes in crisis counseling, but this research has not investigated specific aspects of communication unique to text-based crisis counseling services.

Considerations in Text-Based Counseling

Although text-based counseling shares many attributes with spoken counseling, key differences exist. Text-based counseling remains conversational and allows for the delivery of content in a manner that is similar to spoken counseling, such as risk assessment or psychoeducation. However, text-based counseling differs from spoken counseling in important ways. For example, text-based counselors can manage multiple conversations simultaneously. In such cases, pre-generated messages can improve efficiency by expediting responses, refining presentation, and minimizing repetitive behaviors. However, this practice may result in lengthy, less personalized messages. Additionally, spoken conversations frequently include minimal encouragers to demonstrate active listening, but such responses may seem out of place in text-based formats, potentially necessitating longer messages to convey attentiveness and engagement (Kitchingman et al., 2018; Althoff et al., 2016). Investigating these fundamental elements of text-based crisis counseling is essential to optimizing its effectiveness.

Text Reuse

Text-based crisis counseling requires counselors to manage multiple conversations in real-time, often under high cognitive and emotional demands (Kitchingman et al., 2018). To maintain efficiency, counselors may rely on pre-written templates, frequently reused phrases, or simply develop habits

such that they send messages that are highly similar across their different clients. Additionally, it is notable that some evidence-based therapeutic interventions inherently involve standardized or repetitive forms of communication, such as breathing exercises or grounding techniques. However, excessive reliance on any templated responses could impact client outcomes by reducing the personalization and adaptability of therapeutic interactions. Research suggests that more successful counselors adapt their language dynamically rather than relying on formulaic messaging (Althoff et al., 2016). One way to measure this phenomenon is via text reuse detection (Citron and Ginsparg, 2014), commonly analyzed through textual overlap and, similar to methods used in plagiarism detection, capturing the extent to which language is repeated or varied across messages.

There are many plausible reasons counselors may reuse text, including efforts to provide consistent intervention content, efficiently share structured coping or psychoeducational materials, adhere to any possible evidence-based training materials, and manage the demands of real-time text-based crisis work. The present study is descriptive and focuses on documenting the presence and patterns of text reuse rather than evaluating its causes or associations with client outcomes. Examining whether and how text reuse relates to service delivery or client outcomes remains an important direction for future research.

Text reuse can be assessed in a multitude of ways, namely via message similarity and longest common substrings. Message similarity can be analyzed at the message-level (within a single message) and conversation-level (across multiple messages from the same counselor). For example, one computational method for capturing text reuse, or message similarity, is through the Jaccard similarity index (Citron and Ginsparg, 2014). While there is no prior formal evaluation of text reuse and self-plagiarism in psychotherapy contexts or crisis counseling settings, this method is commonly used in plagiarism detection. For example, Citron and Ginsparg (Citron and Ginsparg, 2014) compared text content of scientific articles between 1991 and 2012 and measured how chronic text reuse is distributed among authors in the dataset. Using an n-gram-based NLP method across more than 750,000 scientific articles, they found that substantial text reuse is not evenly distributed but is concentrated among a small subset of “serial” reusers, while

the majority of authors rarely or never reuse large amounts of text. They also reported a strong negative association between the amount of reused text in an article and its scientific impact ($r \approx -0.74$), suggesting that higher levels of textual reuse are linked to lower influence.

Importantly, psychotherapy and crisis counseling contexts differ substantially from scientific publishing, where text reuse is often interpreted as misconduct or reduced originality. In therapeutic settings, repeated or standardized language may instead reflect evidence-based interventions, safety practices, psychoeducation, or efforts to maintain consistency across client interactions. Accordingly, the present study does not conceptualize text reuse as inherently problematic, but rather examines the extent and patterns of repeated language in text-based crisis counseling.

The longest common substring (LCSS) approach is another avenue to discover copy-pasted messages. LCSS captures the longest contiguous sequence of words shared between two messages, highlighting instances where a counselor has repeated the same block of text verbatim. Unlike Jaccard similarity, which reflects overall similarity in phrasing between messages, LCSS focuses on exact sequential reuse of text, providing a more precise lens on formulaic messaging.

In a study of large language model extraction of psychotherapy dissatisfaction narratives from online forum posts, Steinbrenner used an evaluative metric based on longest common subsequence (LCS) to compare model-identified passages with passages independently identified by a blinded human rater (Lin, 2004). Their LCS metric evaluates the longest sequence of words shared in the same order, without requiring the words to be consecutive. The model achieved an LCS-based similarity score of 0.68, indicating substantial agreement with the human reference and supporting the usefulness of LCS-based methods for identifying similar text despite minor wording differences (Steinbrenner et al., 2025). This method computes overlap in text that allows for skipped words; however, for the purposes of detecting self-plagiarism, we utilize longest common substring rather than subsequence due to its more direct and precise detection of verbatim text reuse. By combining LCSS with broader similarity measures, we can better understand both general patterns of repeated language and specific instances of direct message copying.

For the purpose of our study, we utilize a 10-

gram NLP-based approach using Jaccard similarity indices to detect counselor text reuse across counselor messages. A high similarity score (text reuse) suggests a heavy overlap of text between two messages, indicating substantial reuse of language. Low similarity scores suggest few or no overlapping text between two messages, indicating more tailored or unique language. In addition, we use the longest common substring (LCSS) approach to identify exact, contiguous sequences of repeated words, capturing instances of verbatim message reuse.

1 The Current Study

Despite increased use of text-based crisis counseling, little is known about specific message features. This study examines one key aspect of counselor communication in real-time text-based crisis counseling: text reuse across counselor messages. First, we examined the extent of text reuse across counselors. We examined variability in reuse across counselors and as well as the amount of reuse within counselors. At the message level, we examined the extent that specific messages were reused, both in terms of how many times they were reused, and how many counselors reused the messages. We also clustered reused text to illustrate common patterns and families of repeated content across counselors, highlighting both exact and lightly edited repetitions. By investigating these factors, this study aims to contribute to a more nuanced understanding of how message-level communication looks in crisis counseling, informing best practices for text-based crisis interventions and future research implications that incorporate client outcomes.

2 Methods

2.1 Data

This study utilized de-identified data from a large set of 224,077 crisis counseling encounters (4.7 million messages) collected from a regional text-based crisis and mobile tip line app available 24 hours a day, 7 days per week. The crisis line is implemented in K-12 schools across these regions. The study dataset included crisis encounters from clients of any age located across three different states within the United States who utilized the service between 2017 and 2021. Mobile tips, a system for notifying schools and educators about potentially at-risk student peers, were excluded from the

study sample. Due to the de-identified nature of the dataset, no client-identifying information was available and counselor-level demographic or metadata were limited.

There is a total of 136 licensed or license-eligible counselors who receive additional training in suicide risk assessment and safety planning. Counselor encounters range from 1 to 7,894. Institutional Review Board (IRB) approval was obtained for this study. The crisis line does not systematically collect potentially identifiable information, and text messages were scrubbed of incidental identifying information prior to analysis.

Text reuse analyses were computed on the dataset with additional filtering criteria applied, as described below. Message similarity and LCSS results reflect the subset of messages and encounters that met these criteria.

2.2 Filtering Criteria for Text Reuse

Filters were applied to individual counselor messages, retaining only messages that were longer than 30 words and excluding templated sign-in or sign-off responses using a regular expression pattern. Encounters in which no counselor messages, for example, no messages that were longer than 30 words or non-templated sign-in/sign-off, met these criteria were then excluded from the analytic dataset. This filtering may bias estimates toward longer, more content-rich messages and may underrepresent shorter forms of reuse (e.g., brief templates), which we intentionally exclude to focus on substantive reuse. Message-level filtering resulted in differences in the number of encounters and messages that remained analyzable across counselors. These filtering criteria resulted in a dataset of 332,584 total counselor messages and 92,099 encounters. After message-level filtering, the number of analyzable encounters per counselor ranged from 1 to 5,719. Message similarity score and longest common substring analyses were conducted on these filtered message sets. Estimates for counselors with fewer qualifying messages should be interpreted with caution due to the reduced data density.

To remove scripted sign-in and sign-off responses, all counselor messages were screened for standardized phrases using a predefined pattern-matching regular expression approach. This included common onboarding greetings, for example, greetings or sign-in and sign-off scripts such as “Hi there, I’m glad you reached out. My name

is [name] and I use she/her pronouns. Thanks so much for chatting in today, please let me know how can I best assist you?” and standardized follow-up or “...are you still there?” type messages when a chatter leaves or ends the chat. These messages were identified using regular expressions that matched common variants of these scripts. Messages containing these templated elements were excluded to ensure that only substantive, counselor-generated responses were included in the analysis.

2.3 Measures

First, all counselor and client messages were adjusted for de-identified information to appear as one word (“PLACEHOLDER”). This was done to prevent higher word counts that might arise due to the original de-identified placeholder words that contain two or three words to represent one identifying word.

2.4 Message Similarity

Message similarity was used to assess the degree of text reuse across counselor messages. This was measured using both pairwise Jaccard similarity applied to hashed n-grams and longest common substrings. Jaccard similarity was used as a broader measure of text reuse, capturing overall overlap in phrasing across messages, while the longest common substring approach provided a more specific indicator of exact, sequential reuse. Together, these methods allow for the identification of both general and highly similar forms of repeated language.

Jaccard Similarity Scores. Pairwise Jaccard Similarity scores were calculated for all message pairs across and within each counselor’s encounters. Using the filtered message set described above, each message was split into overlapping 10-word n-grams. The formula for the similarity score is then the total number of shared hashed segments between message A and message B divided by the total number of unique hashes. Hashing converts text segments into a numerical value for ease of computation and comparison over a large set of messages. This formula results in an n-gram Jaccard Similarity score, or a percentage reflecting the extent to which the two messages share the same phrases of length at least 10 words (n-grams). This calculation was repeated on all messages across all encounters for all 136 counselors. The Jaccard similarity formula is defined as follows:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}$$

We compute a mean similarity score for each counselor that aggregates all pairwise similarity scores across all encounters. The degree of overlap between messages is determined by the n-gram length, for example, if $N = 10$, the message similarity is calculated based on every overlapping 10-word sequence shared between messages. The goal is to capture how similar each counselor’s message is relative to other messages by the same counselor across their encounters.

2.5 Longest Common Substring

The longest common substring (LCSS) between two texts is the longest stretch of words that appears contiguously in both texts. In other words, the match between the two texts must be uninterrupted, so the shared material must occur as one continuous block in both texts. Computing LCSS for every pair of messages a counselor has sent then ultimately results in direct text reuse.

Computationally, this can be done efficiently by comparing all prefixes of the two texts and keeping track of how long the current matching run is at each position. Whenever the two current elements match, the run length increases by one, and when they do not, it resets to zero. The longest match discovered in this process is the LCSS. For the purpose of text reuse, we want to discover all LCSS of reasonable length; we chose this hyperparameter to be exact matches of length at least 100. Furthermore, it is possible for the texts to share more than one reused segment, so when an LCSS is discovered we delete it from the two texts and compute the next LCSS, until there are no more matches of length at least 100.

Counting the number of LCSS across all message pairs is not without limitations. A single underlying phrase can produce multiple overlapping LCSS matches across different message pairs. For example, consider the messages A: “Hello, thank you for reaching out today,” B: “thank you for reaching out today,” and C: “thank you for reaching out.” The LCSS between A and B is B, and between A and C is C, however, because C is contained within B, naively counting both overstates the amount of reused text. More generally, overlapping fragments may not nest cleanly: two LCSS spans can partially overlap without either contain-

ing the other, making simple deduplication by substring containment insufficient.

To address this, we cluster LCSS spans that represent variants of the same underlying phrase. Spans shorter than 100 words were discarded to focus on substantive reuse.

2.6 Clustering Descriptives for Reused Messages

We seek to describe and illustrate reuse at two different levels. First, we wanted a message-level estimate of how often counselors reused text content. Second, we wanted to identify recurrent families of reused text, rather than treating every longest common substring as an unrelated artifact. These goals are related, but they are not identical. A single message can contain multiple distinct copied or highly similar submessages, so any attempt to force all reuse into a disjoint set of clusters would blur that structure. We therefore separated prevalence estimation from semantic organization.

To estimate message-level reuse prevalence, we used the extracted LCSS spans together with the associated message ids. For each counselor, we parsed the message IDs and associated every LCSS span with all messages in which it occurred. Next, for each message, we collected all candidate LCSS spans linked to that message. This procedure yields a message-level measure of reuse that is robust to the same message being linked to multiple overlapping substrings.

To identify broader families of reused content, we clustered the LCSS spans semantically. We treated each surviving LCSS as one observation and embedded the span text using MentalBERT (Ji et al., 2021). We chose a semantic embedding approach because clustering based only on surface overlap produced fragmented and often uninformative groups. In particular, lexical methods tend to split together phrases that clearly instantiate the same reused template but differ slightly in wording, while also overemphasizing trivial overlap. By embedding the LCSS spans, we instead aimed to group semantically similar reuse patterns, including lightly edited or paraphrastic variants of a common template.

Each remaining span was tokenized and truncated at 512 tokens, matching the model context limit. This is a limitation of the method because the tail of messages longer than 512 tokens (about 400 English words) will be excluded from the clustering. The resulting span embeddings were clus-

tered with MiniBatchKMeans using $k = 100$. We selected the cluster representative as the LCSS closest to the cluster centroid, since this gives a typical exemplar of the cluster.

For each cluster, we reported the number of distinct message IDs and distinct counselor IDs represented among the member LCSSs. These counts should be interpreted as lower-bound coverage measures for that reuse family. They are not additive across clusters, because the same message can contain multiple distinct reused spans and therefore contribute to more than one cluster and computed clusters are not guaranteed to perfectly separate ground truth clusters. This overlap is a substantive property of the data rather than a defect of the analysis: counselor messages may legitimately combine several reused components in a single response. For this reason, cluster-level counts were used to describe the breadth of a reuse family, while message-level prevalence was estimated separately using the message-collapse procedure above.

2.7 Statistical Analysis

All analyses were conducted in R using “textreus” (Li, 2026), “tidyverse,” and “dplyr” packages (Wickham, 2026). Pairwise Jaccard similarity scores were computed for all messages within each counselors’ encounters. A mean similarity score was also calculated for each counselor to represent the overall degree of text reuse across their set of clients. These scores were summarized descriptively at both the message and encounter level for each counselor. Additionally, LCSS were computed to capture the extent of exact sequential overlap in phrasing between messages. We report overall exploratory and descriptive patterns of message reuse at the encounter and message level in this dataset.

3 Results

3.1 Message Similarity

Similarity was computed at the message level for each counselor. After filtering, all remaining counselor messages were tokenized into 10-gram sequences, and pairwise Jaccard similarity was calculated for every message pair within a counselor. This produced one similarity matrix per counselor, where higher values indicate greater overlap in 10-gram content across messages. In other words, higher scores indicate high text reuse and lower scores indicate low reuse, where 0 indicates no

text overlap between two messages and 1 indicates identical text overlap.

A total of 103 out of 136 (75.7%) counselors demonstrated at least one instance of detected reused text, suggesting that a typical counselor reused message content at some point in their practice. However, across all counselors, mean similarity scores were very low, indicating that reuse accounted for a small proportion of message pairs overall. Counselor-level mean similarity values ranged from 0 to 0.031 ($M = 0.001$, $SD = 0.0034$). For most counselors, the vast majority of message pairs of at least 30 words shared no overlapping 10-word sequences. However, a small subset of counselors demonstrated notably greater reuse than the sample as a whole. Whereas most counselors (83.1%) had less than 0.1% average similarity between their messages, the counselors with the highest levels of reuse had mean similarity scores ranging from 0.85% to 3.1%. This range indicates that, for some counselors, repeated phrasing was a consistent feature of their messaging.

3.2 Longest Common Substring (LCSS)

A total of 72,665 messages (21.86% of 332,584 messages longer than 30 words) included some reuse of at least 10 consecutive words. A total of 13,562 (4%) messages included some reuse of at least 100 consecutive words. Across counselors, the proportion of messages that included reuse of at least 10 and 100 words averaged 0.22 and 0.041, respectively. The corresponding counselor-level counts had medians of 17 and 0, standard deviations of 1,355.61 and 379.88, and ranged from 0 to 8,768 and 0 to 3,110.

Across the 136 counselors, the average number of messages with some reused content was 9.45% among messages of at least 30 words and 1.76% among all messages when reuse was defined as at least 10 consecutive words, with standard deviations of 12.20 and 2.72 percentage points, respectively. For reuse of at least 100 consecutive words, the corresponding averages were 1.29% among messages of at least 30 words and 0.21% among all messages, with standard deviations of 3.15 and 0.61 percentage points, respectively.

The number of counseling encounters varied widely across counselors ($M = 677.20$, $SD = 1,113.68$) for messages of at least 30 words, with a range of 1 to 5,719 and a median of 106. Across all encounters, the average number of encounters

was 1,113.72 ($SD = 1,784.05$), with a range of 1 to 7,894 and a median of 171.

3.3 Counselor Encounter Volume and Message Reuse

Across counselors, the percentage of reused messages was positively correlated with the total number of unique encounters. For reuse of at least 10 consecutive words, the correlation was $r = .51$ among messages of at least 30 words and $r = .43$ across all messages, both $p < .001$. For reuse of at least 100 consecutive words, the corresponding correlations were $r = .50$ and $r = .50$, respectively, both $p < .001$. This suggests that as counselor encounter volume increased, counselors reused content more frequently. The counselor with the largest number of encounters without any observed 10 consecutive word reuse had 163 encounters, corresponding to 1,548 unique messages. Similarly, when reuse was defined as 100 consecutive words, the counselor with the largest number of encounters without detected reuse had 2,724 encounters, corresponding to 25,016 unique messages.

3.4 LCSS Clustering Using MentalBert

Clustering analyses were completed for reused text spans (LCSSs) of length at least 100 words to capture substantive reused messages. For illustration, we report on the top 5 most reused clusters of text discovered by our algorithm. The clustering procedure groups candidate texts that are likely variants of the same underlying phrase, which were reviewed by a human annotator who confirmed cluster membership and identified the representative passage. This semi-automated approach combines the scalability of algorithmic matching with human judgment for disambiguation.

The clusters with the most frequently detected reused LCSS are reported in Table 1. The cluster with the most LCSS occurrences was a list of coping strategies, which was identified as used at least 978 times across 26 counselors. This script typically ranged from 100 to 507 words. A similarly common cluster involved a grounding sensory exercise, observed 790 times across 13 counselors, with messages ranging from 100 to 655 words. Another cluster revealed messages regarding self-care tips, observed 745 times across 11 counselors, with messages ranging from 110 to 314 words.

Other recurring clusters included a different list of coping strategies, observed 578 times across 9 counselors, ranging from 100 to 220 words. A

muscle tension relaxation technique cluster was detected 572 times across 13 counselors, ranging from 101 to 675 words. Overall, the most frequently reused passages involved structured psychoeducational or coping-oriented content, such as breathing exercises, grounding techniques, and self-care guidance, rather than conversational or relational statements.

Discussion

The present study examined text reuse across real-time, text-based crisis counseling encounters. We found that some reuse was detectable for most counselors: 103 of 136 had at least one detected instance of reuse. Repeated wording was not the dominant pattern of counselor communication. Mean pairwise 10-gram Jaccard similarity was extremely low overall ($M = .001$), and most counselors had average similarity values near zero. And finally, counselor text reuse was correlated with encounter volume across analytic definitions of reuse; counselors who handled more encounters showed higher proportions of reused messages. Taken together, these findings suggest that counselor messages were typically composed in unique language at the whole-message level, while selective reuse of particular passages became more common among counselors with greater encounter volume.

Variation across counselor text reuse was a central finding. Reuse was not distributed evenly throughout the sample. Most counselors showed very low average similarity and little repeated overlap across messages, whereas a smaller subset demonstrated more consistent reuse. This pattern suggests that text reuse in this setting is better understood as concentrated among some counselors rather than as a uniform feature of crisis counseling practice. At the same time, reuse was not absent. LCSS analyses showed that approximately one in five counselor messages (30 or more words) contained a repeated segment of at least 10 consecutive words. Much longer repeated passages of 100 or more words were less common, but still appeared in roughly one in 25 qualifying counselor messages. The presence of some counselors with nontrivial encounter volume but no detected reuse at these thresholds further indicates that repeated text was common, but not inevitable.

The Jaccard analyses summarized similarity across all message pairs within counselors, and most pairs did not share any 10-grams, which kept

Cluster / message type	Occurrences	Counselors	Message length (words)
List of coping strategies	978	26	100–507
5-4-3-2-1 grounding exercise	790	13	100–655
Self-care tips	745	11	110–314
Alternative list of coping strategies	578	9	100–220
Muscle tension relaxation technique	572	13	101–675

Table 1: Most frequently detected reused message clusters.

average similarity near zero. In contrast, LCSS identified whether a message contained any repeated segment above a minimum length. A counselor could therefore reuse a limited number of longer passages within otherwise distinct messages and still have very low average pairwise similarity. This distinction is important because it suggests that reuse in this dataset was likely localized to specific components of counselor responses rather than reflecting broadly repetitive communication.

The clustering results help clarify the kinds of content that were most often repeated. Among the most frequently detected long reused passages were lists of coping strategies, grounding exercises, self-care tips, and muscle relaxation instructions. In other words, the most common repeated passages appeared to involve structured, psychoeducational, and coping-oriented material rather than primarily relational or conversational language shared between counselors. These findings suggest that some counselors may draw from preformulated language when delivering common intervention content. However, the present data cannot determine whether these passages came from formal templates, personal notes, training materials, or individual habits that developed over time. Nor can these findings determine whether the use of such language improved or reduced the quality of support.

One of the most important findings was the positive association between counselor encounter volume and text reuse. Across counselors, the proportion of reused messages increased as the number of unique encounters increased, and this pattern was evident across both reuse thresholds examined. For reuse defined as at least 10 consecutive words, the correlation with encounter volume was $r = .51$ among messages longer than 30 words and $r = .43$ across all messages; for reuse defined as at least 100 consecutive words, the corresponding correlations were $r = .50$ and $r = .50$, respectively (all $p < .001$). Because the association was observed across

reuse thresholds and definitions, it does not appear to depend on any one measure of reuse. Although greater encounter volume creates more opportunity to observe repeated passages, the fact that this pattern appeared when reuse was expressed as a proportion of messages suggests that it is unlikely to be explained solely by counselors generating more total text. An interpretation is that counselors with greater encounter volume may develop and repeatedly draw on preferred ways of delivering common intervention content, particularly structured coping or psychoeducational guidance. Rather than viewing text reuse only as an individual stylistic difference, these findings suggest that it also scales with the extent of counselors' practice within the service.

These findings have several implications for understanding communication in digital crisis services. Most importantly, the results do not support the view that counselor messaging is all broadly templated. Instead, they suggest a more selective pattern in which reuse, when it occurs, is concentrated in certain counselors and in certain types of content. This may reflect one way counselors manage recurring tasks in a text-based, high-intensity service environment, particularly when offering concrete coping tools or psychoeducation.

However, this association may also be shaped by unmeasured contextual factors, including counselor tenure, shift timing, fatigue, workload, or differences in training and supervision practices. As a result, encounter volume should be interpreted cautiously and not necessarily assumed to be the sole driver of text reuse patterns.

This study provides a descriptive understanding of text reuse in real-time crisis counseling. Our main finding was that text reuse was positively associated with counselor encounter volume across definitions of reuse. This suggests that reuse increases with counselor encounter volume. Although average message-level similarity was very low for most counselors, text reuse was nevertheless a meaning-

ful and patterned feature of counselor communication. Reuse was concentrated among a subset of counselors and was most often observed in structured coping-oriented or psychoeducational passages. These findings suggest that text reuse in this setting is not pervasive, but neither is it incidental: it becomes more common as counselor encounter volume increases.

The present study was descriptive and focused specifically on extended identical text reuse rather than broader stylistic similarity or effectiveness. The findings establish baseline evidence that communication patterns within text-based crisis counseling are not uniform and that text reuse increases systematically with counselor encounter volume. As digital crisis services continue to expand, future research linking message-level features to client outcomes and examining the full counselor dataset will be essential for understanding whether observed stylistic differences have meaningful implications for service delivery.

Limitations

The number of encounters retained after message-level filtering varied across counselors. This variation reflects differences in how frequently counselors produced longer (> 30-word), non-templated messages within their sampled encounters. As a result, the dataset differed in size across counselors after applying inclusion criteria. Consequently, similarity estimates for counselors with fewer retained messages are based on a smaller pool of observations and should be interpreted with appropriate caution due to reduced data density. Importantly, this variation resulted from analytic filtering decisions rather than unequal sampling or workload differences across counselors. A second limitation concerns the scope of the similarity/text reuse metric. Linguistic similarity was operationalized as identical contiguous 10-word sequences, reflecting the study's focus on extended verbatim text reuse. Accordingly, findings should be interpreted as estimates of longer contiguous segment reuse rather than broader stylistic or thematic similarity.

Finally, although variation in message length and reuse patterns was identified, the study does not evaluate whether particular communication styles are associated with improved client outcomes, perceived helpfulness, or reductions in distress. Therefore, no conclusions can be drawn regarding the effectiveness of longer messages or higher levels of text reuse. Future research linking message-level

features to client outcomes would be necessary to determine whether observed stylistic differences have meaningful clinical implications. Additionally, examining a larger set of counselors and their encounters may provide more stable, accurate estimates of communication patterns and allow for stronger conclusions regarding variability across providers in this setting.

Our LCSS-based analysis has several limitations. First, LCSS is sensitive to minor surface-level discrepancies: a single extra space, a missing line break, or formatting artifacts introduced when copying text from PDF sources can break an otherwise exact match. This contrasts with n-gram Jaccard similarity, which is more robust to such perturbations. Second, when a counselor reuses part of a longer reused passage, the shorter fragment is detected as a separate LCSS match against every message containing the longer passage, inflating its occurrence count. Our clustering procedure using MentalBERT attempts to work around these weaknesses because long messages that differ by a space probably map to the same area and are then clustered together. A clear limitation of using MentalBERT for the clustering is the 512 token limit that cuts off long reused messages. The clustering output requires human review to confirm that grouped spans genuinely represent variants of the same underlying passage, making this a semi-automated rather than fully automated method. Third, the reported counts should be interpreted as lower bounds: at minimum, these counselors used these passages as part of a message at least this many times. Obtaining precise counts is difficult given the overlapping and fragmentary nature of pairwise LCSS output. Fourth, to avoid flagging genuine therapeutic style the minimum LCSS length must be set conservatively high. We chose a threshold of 100 words, accepting that this sacrifices recall of shorter templates in favor of specificity for substantive text reuse.

Ethical Considerations

This study used de-identified data from a text-based crisis counseling service. All messages had personally identifying information removed prior to analysis, and no attempt was made to re-identify individuals. The dataset included sensitive communications from individuals in crisis, and analyses were conducted at the aggregate level to minimize risk of harm or disclosure. Institutional Review

Board (IRB) approval was obtained prior to data access and analysis.

Given the vulnerable nature of the population, care was taken to ensure that findings were reported in a way that avoids stigmatization or misinterpretation of counselor or client behavior. Results are intended to inform understanding of communication patterns at the system level rather than to assess or compare individual providers.

Acknowledgments

AI tools were utilized to assist in the writing of this manuscript and to assist in performing literature search, in accordance with the ACL Policy on AI Writing Assistance cases a and c.

References

- Tim Althoff, Kelsie Clark, and Jure Leskovec. 2016. [Large-scale analysis of counseling conversations: An application of natural language processing to mental health](#). *Transactions of the Association for Computational Linguistics*, 4:463–476.
- M. C. Budinger, M. Cwik, and M. Riddle. 2015. [Awareness, attitudes, and use of crisis hotlines among youth at-risk for suicide](#). *Suicide and Life-Threatening Behavior*, 45(2).
- Daniel T. Citron and Paul Ginsparg. 2014. [Patterns of text reuse in a scientific corpus](#). *Proceedings of the National Academy of Sciences*, 112(1):25–30.
- Sally C. Curtin. 2020. [State suicide rates among adolescents and young adults aged 10–24: United states, 2000–2018](#). Technical report, CDC.
- Angela Dwyer, Antonio de Almeida Neto, Dominique Estival, Wei Li, Catherine Lam-Cassettari, and Mark Antoniou. 2021. [Suitability of text-based communications for the delivery of psychological therapeutic services to rural and remote communities: Scoping review](#). *JMIR Mental Health*, 8(2):e19478.
- Robert D. Goldney. 2005. [Suicide prevention](#). *Crisis*, 26(3):128–140.
- Madelyn S. Gould, S. Chowdhury, A. M. Lake, H. Galfalvy, M. Kleinman, M. Kuchuk, and R. McKeon. 2021. [National suicide prevention lifeline crisis chat interventions: Evaluation of chatters’ perceptions of effectiveness](#). *Suicide and Life-Threatening Behavior*, 51(6):1126–1137.
- Madelyn S. Gould, W. Cross, A. R. Pisani, J. L. Munfakh, and M. Kleinman. 2013. [Impact of applied suicide intervention skills training on the national suicide prevention lifeline](#). *Suicide and Life-Threatening Behavior*, 43(6):676–691.
- Madelyn S. Gould, J. Kalafat, J. L. HarrisMunfakh, and M. Kleinman. 2007. [An evaluation of crisis hotline outcomes part 2: Suicidal callers](#). *Suicide and Life-Threatening Behavior*, 37(3):338–352.
- E. Halford, A. M. Lake, and M. S. Gould. 2021. [Crisis hotlines](#). In *Oxford Textbook of Suicidology and Suicide Prevention*. Oxford University Press.
- L. L. Hua, J. Lee, M. H. Rahmandar, and E. J. Sigel. 2023. [Suicide and suicide risk in adolescents](#). *Pediatrics*, 153(1):e2023064800.
- Shaoxiong Ji, Tianlin Zhang, Luna Ansari, Jie Fu, Prayag Tiwari, and Erik Cambria. 2021. [Mentalbert: Publicly available pretrained language models for mental healthcare](#). *Preprint*, arXiv:2110.15621.
- R. King, B. Nurcombe, L. Bickman, L. Hides, and W. Reid. 2003. [Telephone counselling for adolescent suicide prevention](#). *Suicide and Life-Threatening Behavior*, 33(4):400–411.
- T. A. Kitchingman, C. J. Wilson, P. Caputi, I. Wilson, and A. Woodward. 2018. [Telephone crisis support workers’ psychological distress and impairment](#). *Crisis*, 39(1):13–26.
- A. M. Lake, T. Niederkrotenthaler, R. Aspden, M. Kleinman, A. M. Hoyte-Badu, H. Galfalvy, and M. S. Gould. 2022. [Lifeline crisis chat: Coding form development and findings on chatters’ risk status and counselor behaviors](#). *Suicide and Life-Threatening Behavior*, 52:452–466.
- M. Lauver, M. C. Podlogar, J. J. Herbin, M. Selig, F. K. Muzzey, K. Patelli, J. F. McCarthy, G. J. Hughes, and M. Stacy. 2024. [Implementing and evaluating a veterans crisis line quality improvement initiative](#). *Psychological Services*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- B. L. Mishara and M. S. Daigle. 1997. [Effects of different telephone intervention styles with suicidal callers](#). *American Journal of Community Psychology*, 25(6):861–885.
- B. L. et al. Mishara. 2007. [Which helper behaviors and intervention styles are related to better short-term outcomes](#). *Suicide and Life-Threatening Behavior*, 37(3):308–321.
- Brian L. Mishara and Marc Daigle. 2001. [Helplines and crisis intervention services: Challenges for the future](#). In David Lester, editor, *Suicide Prevention: Resources for the Millennium*, pages 153–171. Brunner-Routledge.
- Kirsten Pauwels, Eva De Jaegere, Patrick Vanderreydt, Silke Aerts, Eva Vande Gaer, and Gwendolyn Portzky. 2025. [Assessing a suicide prevention helpline’s impact on caller crisis level and suicidality](#). *Archives of Suicide Research*, 29(3):621–636.

Rajeev Ramchand, Lisa H. Jaycox, Patricia Ebener, Mary L. Gilbert, Dionne Barnes-Proby, and Pravesh Goutam. 2017. [Characteristics and proximal outcomes of calls made to suicide crisis hotlines in california: Variability across centers](#). *Crisis: The Journal of Crisis Intervention and Suicide Prevention*, 38(1):26–35.

D. A. et al. Ruch. 2019. [Trends in suicide among youth aged 10 to 19 years](#). *JAMA Network Open*, 2(5):e193886.

Tobias Steinbrenner, Christopher Lalk, Alin Kabjesz, Drin Ferizaj, Juan Segundo Pena Loray, Flavio Iovoli, and Julian Rubel. 2025. [Exploring negative experiences in psychotherapy using an nlp approach on online forum data](#). *npj Mental Health Research*, 4(1).

Substance Abuse and Mental Health Services Administration. 2025. [988 lifeline performance metrics](#).

World Health Organization. 2019. [Suicide in the world](#).