

Combining Hierarchical VAEs with LLMs for clinically meaningful timeline summarisation in social media

Jiayu Song^{*1}, Jenny Chim^{*1}, Adam Tsakalidis^{1,3,4}, Julia Ive¹, Dana Atzil-Slonim², Maria Liakata^{1,3}

¹Queen Mary University of London, London, UK

²Bar-Ilan University, Israel

³The Alan Turing Institute, London, UK

⁴European Centre for the Development of Vocational Training, Thessaloniki, Greece

{jiayu.song, c.chim, a.tsakalidis, j.ive, m.liakata}@qmul.ac.uk

dana.slonim@gmail.com

Abstract

We introduce a hybrid abstractive summarisation approach combining hierarchical VAEs with LLMs to produce clinically meaningful summaries from social media user timelines, appropriate for mental health monitoring. The summaries combine two different narrative points of view: (a) clinical insights in third person, generated by feeding into an LLM clinical expert-guided prompts, and importantly, (b) a temporally sensitive abstractive summary of the user’s timeline in first person, generated by a novel hierarchical variational autoencoder, TH-VAE. We assess the generated summaries via automatic evaluation against expert summaries and via human evaluation with clinical experts, showing that timeline summarisation by TH-VAE results in more factual and logically coherent summaries rich in clinical utility and superior to LLM-only approaches in capturing changes over time.

1 Introduction

Social media users discuss different aspects of their lives, providing important clues about their mental health. Previous work (De Choudhury et al., 2013; Coppersmith et al., 2014; Cohan et al., 2018; Chancellor and De Choudhury, 2020) has studied users’ social media posts to help identify depression, bipolar disorder (Yates et al., 2017; Hussein Orabi et al., 2018) or self-harm (Zirikly et al., 2019), with work on multi-task learning to capture user states at a particular moment in time (Benton et al., 2017; Yang et al., 2023). Despite the importance of longitudinal assessments of linguistic and other digital content for mental health clinical outcomes (Velupillai et al., 2018), there is little work on considering the evolution of an individual’s mental health over time through their social media. Tsakalidis et al. (2022b,a) established the task of capturing changes (switches and escalations) in an individual’s mood over time and showed how identifying

these helps predict clinical assessments of suicidal ideation. However, currently clinicians don’t have access to such information to assess individuals’ mental-state and mainly rely on self-reports completed by patients throughout psychotherapy (Crits-Christoph and Gibbons, 2021). Although standardized subjective measures are fundamental to mental health monitoring and research, they have significant limitations, such as the extent of individuals’ self-awareness, their willingness to complete questionnaires, and the limited choice of responses (Kazdin, 2021). Providing concise summaries that can capture fluctuations in individuals’ state-of-mind while emphasizing key clinical concepts, can significantly assist in monitoring, prevention and early detection of mental health issues. Such summaries would augment clinician capacity, provide alternatives to standard questionnaires and compensate for reduced access to mental health services (Schwartz et al., 2023).

To the best of our knowledge we are the first to propose clinically meaningful summaries of social media user ‘timelines’ (sequences of chronologically ordered posts by a user). Driven by the need to concisely summarise time-series language data which can span arbitrary lengths that exceed limits of many contemporary models and render purely extractive methods impractical, we propose a novel hybrid unsupervised abstractive method, Timeline Hierarchical VAE (TH-VAE). Our system makes use of a hierarchical variational autoencoder that compresses timeline information into compact representations and a large language model (LLM), creating a two-layer summary that combines two different narrative points of view. Specifically: (a) a *high-level* summary in third person useful for a clinician, generated by feeding into an LLM expert-guided prompts, and importantly (b) a temporally sensitive abstractive summary of the user’s timeline in first person (*evidence summary*), generated by TH-VAE. The generation of the first person abstrac-

^{*}Equal contribution.

tive evidence summary via TH-VAE is guided by mental health related key-phrases obtained through instruction prompting by an LLM. The final resulting high level summary covers aspects considered to be crucial by clinicians from a wide range of therapeutic approaches, including individuals' diagnosis, intrapersonal and interpersonal patterns and extent of mental state changes over time (Eells, 2022).¹

We make the following contributions:

- We develop a novel abstractive timeline summarisation method (TH-VAE)² based on adapting a hierarchical VAE model (NVAE)(§3.4) to longitudinal social media data (user timelines).
- We provide a new task, the creation of clinically meaningful summaries from social media data. These summaries, generated in a hybrid approach, comprise high-level information in third person consistent with key clinical concepts (diagnosis, inter- and intra- personal aspects, moments of change) and evidence from a user's timeline, generated from the TH-VAE, supporting the assigned high-level insights. (§3)
- We create a dataset of expert-written mental health summaries from longitudinal social media data. A small sample of these is used to help with modeling (§3.5) and the rest is used for evaluation (§4.3).
- We provide a novel detailed evaluation method of the summaries based on preservation of clinical information, summary consistency, and usefulness to clinicians, using semantic similarity based metrics, NLI based inference, as well as expert human evaluation (Zhang et al., 2023b) (§4.3).
- We conduct experiments using different unsupervised summarisation methods based on LLMs and story generation (§4.2), showing superior performance for TH-VAE (§5).

2 Related Work

Timeline summarization aims at concisely summarizing the evolution trajectory of a specific topic along a timeline (Chen et al., 2019, 2023) and has primarily focussed on news datasets. Methodologically it has involved both extractive and abstrac-

tive methods; for example, Allan et al. (2001) define temporal summaries by extracting a sentence per event in a news story while Li et al. (2021) construct a multi-document event graph to capture long distance dependencies between events, weight events and extract an event summary sentence with maximum event coverage. Li and Cardie (2014); Chang et al. (2016); Wang et al. (2021); Hills et al. (2023a) detect important events in an individual's timeline and explore the event trajectory. In Ren et al. (2013) timeline summarisation involves identifying users' interests by defining a social circle from a set of friends and selecting salient tweets to obtain an extractive summary. Chang et al. (2016) also uses extractive summarisation and selects sentences based on different features (e.g., popularity-based, temporal). Work in abstractive timeline summarisation (Martschat and Markert, 2018; Steen and Markert, 2019) involves identifying clusters of news or events to generate abstractive summaries from, or memory-based timeline summarisation to track the trajectory of events (Chen et al., 2019). By contrast we consider a user's timeline, a series of posts shared by an individual over a period of time (Tsakalidis et al., 2022b). Such timelines do not exhibit obvious or consistent topics, contain few events and an explosion of emotions. Our goal in user timeline summarisation is to capture important information and synthesise it.

Summaries in Mental Health. Although summaries are clinically crucial for compiling information about individuals, there is limited literature on the subject, with the primary focus being on expert-generated case study summarization (Eells, 2022). Only recently, researchers have started to use NLP capabilities to automatically generate summaries in the clinical domain. Manas et al. (2021) demonstrated the usefulness of generating summarised diagnoses from a single-session interview. Srivastava et al. (2022) summarised psychotherapy conversations at the level of single counseling sessions proposing that summaries should exploit domain knowledge and conversational elements. On social media, Sotudeh et al. (2022) generated summaries of individual Reddit posts, relying on formatting conventions (i.e. TLDR) to extract short summaries provided by the users themselves without further content constraints. Yang et al. (2023) instruction-tuned LLMs to generate mental health analyses from static social media text. By contrast our work summarises user timelines and combines informa-

¹For a complete list, please see Table 8 in Appendix A.

²<https://github.com/Maria-Liakata-NLP-Group/THVAE-summary/tree/main>

tion from social media posts based on high-level expert domain knowledge, important for evaluating individuals’ progression over time.

Summarising with LLMs. Current work on LLM-based summarisation focuses on news articles or instructional texts (Goyal et al., 2022; Zhang et al., 2023a; Maynez et al., 2023), using simple prompts (e.g. “Summarize the following article:”). Wang et al. (2023) took a multi-step approach, extracting event information from news via curated guiding questions then summarising the prompted outputs. In our work, we summarise longitudinal user generated content and use clinically-informed prompts to generate high-level mental health observations.

Summary Evaluation. Existing mental health summarisation works utilised natural language generation metrics, for example ROUGE (Lin, 2004) being used to measure n-gram overlap against reference documents (Manas et al., 2021; Srivastava et al., 2022; Sotudeh et al., 2022). Srivastava et al. (2022) additionally applied BLEURT (Scialom et al., 2021), a learned metric trained on ratings, QuestEval (Scialom et al., 2021), a metric based on question generation and answering, and MHIC, a metric based on ROUGE that they defined to assess information captured in counselling summaries .

Contrary to prior work, our task involves two-layer mental health summaries combining first-person social media content with high-level clinical concepts in third person, posing unique evaluation challenges. For example, data noisiness makes metrics learned on well-formed texts unsuitable, and evaluation must assess consistency both between summary layers and within the detailed high-level summary itself. To this end, we extend the line of work leveraging natural language inference (NLI) in summary factuality and consistency evaluation (Maynez et al., 2020; Laban et al., 2022).

3 Methodology

Task Given a user’s timeline (a series of posts between two dates (Tsakalidis et al., 2022b)), the goal is to generate an abstractive summary that reflects the user’s mental state and how it changes over time. This summary includes high-level information useful for clinicians in third person, and corresponding evidence from the timeline in first person.

3.1 Architecture Overview

Fig. 1 shows the summary generation process. It consists of two sub-processes:

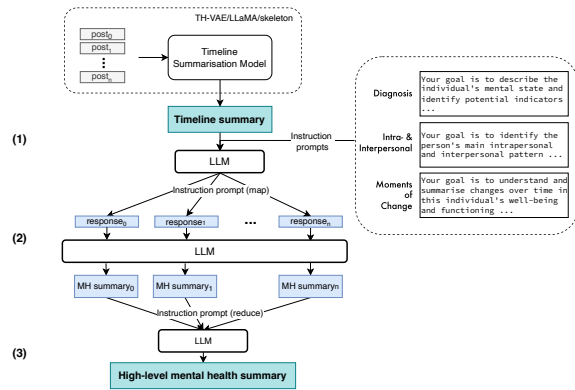


Figure 1: Prompting framework for generating high-level summaries. Taking a first-person summarised timeline as input, we (1) prompt the LLM around different key clinical concepts, (2) summarise extracted inferences into prose per topic, and (3) combine the topic-specific intermediate summaries into a coherent, distilled document.

segment ₁	Back at it again with using sex as self harm ...	IE
	I'm so suicidal right now ...	IEP

segment _i	Can't believe I'm finally starting to find reasons...you'll always have the first...	ISB
	Finally gonna get my sleeping back on track (at least I hope) thanks to staying up all night ...	IS

Figure 2: Each timeline is separated into several segments based on ‘MoC’. We highlight the key phrases.

(I) Abstractive generation of the timeline/evidence summary (§3.4 and Fig. 3). We use three different unsupervised methods for creating the timeline summary in first person: Timeline hierarchical VAE (TH-VAE §3.2 and Fig. 3, right), our key methodological novelty; LLaMA (§4.2); a method from story generation (§4.2).

(II) Generation of the High-level summary (§3.5 and steps (1)-(3) in Fig. 1). We feed the generated timeline/evidence summary into an instruction-tuned LLM (Llama), where prompts originate from a small sample of expert human annotation (§3.2), and generate high-level summaries covering key clinical concepts such as diagnosis, inter- and intra-personal relationships and fluctuations in mood. The following subsections describe our novel timeline summarisation method using an adapted hierarchical VAE (TH-VAE).

3.2 Input to Timeline Summarisation

The input to TH-VAE and the other timeline summarisation methods is a user’s timeline, an-

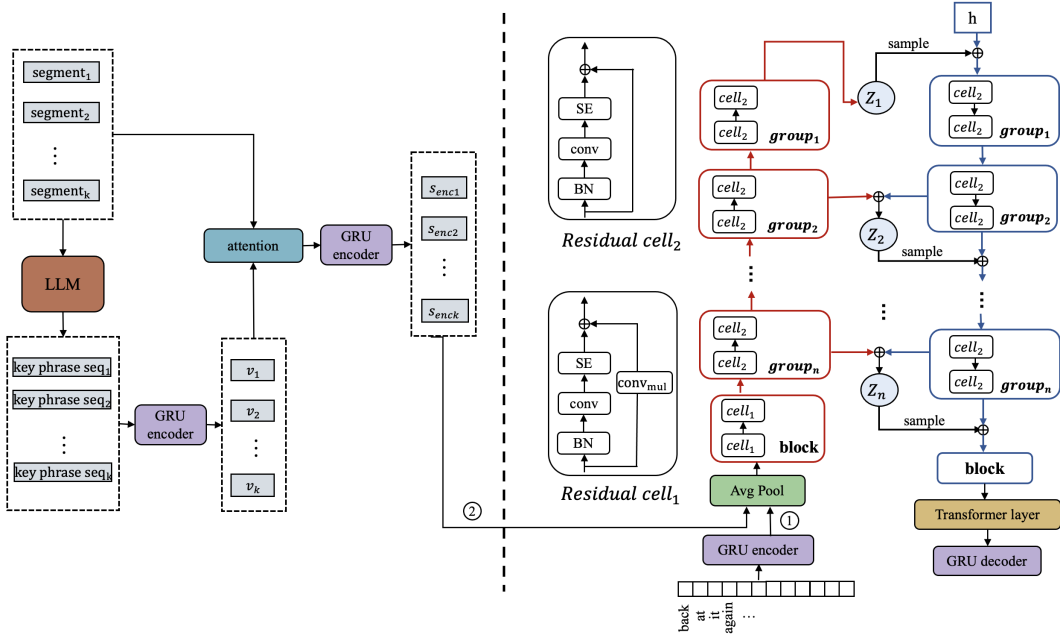


Figure 3: Overview of TH-VAE. The left of the dotted line shows the construction of the k-sentence representation used only during generation, informed by the key-phrases, while the right side shows the hierarchical structure of TH-VAE, and its components. ① and ② represent the input during training and generation respectively.

notated with Moments of Change in mood (MoC)(Tsakalidis et al., 2022b). MoC annotations consist of *Switches* (sudden mood shifts, denoted by ‘IS’–In Switch– and ‘ISB’–In Switch Beginning– tags), and *Escalations* (gradual mood progression, denoted by ‘IE’–In Escalation– and ‘IEP’–In Escalation Peak– tags). We split the whole timeline (see Fig. 2) into several segments (sub-timelines) based on ‘MoC’, so that consecutive posts with the same label (‘IE’ or ‘IEP’), (‘ISB’ or ‘IS’) or ‘0’ are grouped together. This assumes each segment consists of posts of a similar mood type, which facilitates capturing different features and relations between them. This is somewhat similar to news timeline summarisation which clusters around stories or events, with the additional challenge that mood features are more evasive and we hope to model these through latent variables.

Key phrases We asked clinical psychologists to annotate key phrases indicative of users’ mental health in three timelines. These phrases include mood related clues but also information on interpersonal relationships, behaviors or events related to a user’s mental state (see highlights in Fig. 2). We take these annotated timeline/key phrases pairs as examples and prompt LLaMA (Touvron et al., 2023) to annotate the rest of the timelines.

Timeline summary representation For each segment s_i , we input its corresponding key phrase sequence $\{e_1, \dots, e_j, \dots, e_n\}$ into a GRU

encoder (Cho et al., 2014) to get the key phrases encoding $v = GRU([e_1; \dots; e_n])$, which is represented by the last hidden state of the GRU (see left part of Fig3). We calculate the similarity between v and each word embedding w_i in the segment as the weight α_i :

$$\alpha_i = \frac{\cos(v, w_i)}{\sum_{i'=1}^m \cos(v, w_{i'})}$$

Thus s_i can be represented by a series of weighted word embeddings $\{\alpha_1 w_1, \alpha_2 w_2, \dots, \alpha_m w_m\}$, where m is the length of s_i . We encode it with the GRU encoder to get the segment representation $s_{enc_i} = GRU([\alpha_1 w_1; \alpha_2 w_2; \dots; \alpha_m w_m])$. If the timeline is divided into k segments, we can get k segment encodings $s_{enc_1}, \dots, s_{enc_k}$ in this way. We concatenate these encodings in chronological order to get a segment sequence $\{s_{enc_1}, s_{enc_2}, \dots, s_{enc_k}\}$, apply an average pooling operation (Avg Pool)(Lin et al., 2013) over the output of the GRU encoder (See right part of Fig3) and feed it into the hierarchical part of TH-VAE to generate a timeline summary.

3.3 Overview of TH-VAE

Due to the lack of gold summaries for training purposes, we have to construct the summary distribution without any guidance. Thus we need a model that can learn an expressive distribution for

a long timeline (the longest timeline has 124 posts, and the longest of these posts has over 300 words). We also need to construct a mental health related summary distribution that can capture different features and establish the long-range dependencies between these features in the timeline. We propose TH-VAE, an unsupervised abstractive timeline summarization model adapted from NVAE (Vahdat and Kautz, 2020), to construct a more expressive prior for a user timeline.

In the learning process, we split the timeline into several segments (sub-timelines, §3.4), considered to contain consecutive posts with similar mood, and train TH-VAE to learn the distribution of each segment s by reconstructing it.

When generating the evidence summary, we still treat each segment as a unit. To help the model focus on important information during generation we introduce the notion of key phrases (§3.2). We use an automatic method based on an LLM to extract mental health related key phrases from each segment and encode key phrase-segment pairs with an attention mechanism. We concatenate the sequence of segment representations of a timeline in chronological order and input it into the hierarchical structure of TH-VAE to generate the timeline/evidence summary (See Fig. 3, left part).

3.4 Document Reconstruction via TH-VAE

The vanilla VAE assumes a prior $p(\mathbf{z})$ of document \mathbf{x} over latent variables \mathbf{z} to be a Normal Gaussian distribution, and parameterizes an approximate posterior distribution $q_\phi(\mathbf{z}|\mathbf{x})$ given text \mathbf{x} . It uses KL (Kullback–Leibler divergence) to calculate the distance between $p(\mathbf{z})$ and $q_\phi(\mathbf{z}|\mathbf{x})$ and gradually reduces the distance between them in training. Finally, it samples from the hypothesised posterior distribution and generates the document \mathbf{x} . It has been shown that the vanilla VAE can lead to over-regularising the posterior distribution, resulting in latent representations that do not represent well the structure of the data (Klushyn et al., 2019; Alemi et al., 2018; Sønderby et al., 2016; Ranganath et al., 2016; Vahdat and Kautz, 2020). However, for a long document assuming its distribution to be a Gaussian does not provide enough expressive power; we need to be able to consider the structure of different semantic elements and the relationship between them.

The deep Hierarchical VAE (NVAE) (Vahdat and Kautz, 2020), introduced for images, increases expressiveness by introducing several latent variables

to generate large high-quality images, demonstrating the superiority of the hierarchical VAE. Here, we adapt this model for long documents, resulting in Timeline Hierarchical VAE (TH-VAE), and use it as the basis of constructing mental health related timeline representations.

3.4.1 Hierarchical Component

TH-VAE increases the expressiveness of the approximate posterior and prior by partitioning the latent variable \mathbf{z} into l latent variables $\mathbf{z}=\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_l\}$ (Vahdat and Kautz, 2020). The prior is represented by $p(\mathbf{z}) = \prod_l p(\mathbf{z}_l|\mathbf{z}_{<l})$ and it parameterises the approximate posterior distribution $q_\phi(\mathbf{z}|\mathbf{x}) = \prod_l q_\phi(\mathbf{z}_l|\mathbf{z}_{<l}, \mathbf{x})$ which is represented by factorial Normal distributions. This objective is to maximise its lower bound as:

$$L(\theta; \mathbf{x}) = -KL(q_\phi(\mathbf{z}_1|\mathbf{x})||p(\mathbf{z}_1)) \\ \sum_{l=2}^L \mathbb{E}_{q_\phi(\mathbf{z}_{<l}|\mathbf{x})} [-KL(q_\phi(\mathbf{z}_l|\mathbf{x}, \mathbf{z}_{<l})||p(\mathbf{z}_l|\mathbf{z}_{<l}))] \\ + \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})].$$

Before going into the hierarchical architecture, we use a GRU encoder to encode the segment, to reduce the impact from padding. Then we add an Avg Pool (Lin et al., 2013) over the output of the GRU encoder to fix the input length. Both TH-VAE and NVAE use multiple residual cells to construct the hierarchical structure. In TH-VAE we simplify residual cells to work with textual data rather than images, and keep the optimization strategies in NVAE, i.e., BN (batch normalization) with Swish Activation and Squeeze and Excitation (SE). We use two different residual cells: *residual cell*₁ and *residual cell*₂. The input representations first go through a **block** which focuses on capturing the features of a segment and consists of *residual cell*₁. To form *residual cell*₁ we use series BN, *conv* (CNN with one kernel size), SE as well as *conv_{mul}* (CNN with multiple kernel sizes), where the latter helps with capturing the different features. Then, the output of the **block** will go into the layered **groups** (see Fig. 3–right), responsible for learning to capture the relationship between different features in segments and long-range dependencies between them. Each **group** is used to encode the sub-latent variables \mathbf{z}_i and consists of *residual cell*₂s. Since *conv_{mul}* increased parameters without added benefit, we only use *conv* in *residual cell*₂. Finally we add another **block** to integrate information.

During training, the whole hierarchical architecture is used to learn the distribution of each segment, by learning features and long-range dependencies within them via segment reconstruction (as shown in the right part of Fig3). Then during generation a sequence of segments (a whole timeline) is input to TH-VAE to generate similarly structured text. The left part of Fig3 shows the process of encoding the sequence of segments. When decoding, we use the same decoder component as in (Song et al., 2022), comprising a transformer decoder (we load pre-trained parameters from BART) followed by a GRU decoder.

3.5 High-level Mental Health Summarization

We focus on information considered important in summarising individuals’ mental states according to therapeutic approaches (Eells, 2022). Although all users broadly talk about mental health related topics in this dataset, the extent to which clinical concepts appear in each one varies due to natural individual differences. As such, when annotators write gold summaries and when we generate model-written ones, we focus on key clinical concepts that are present, ignoring true negatives.

We prompt an instruction-tuned LLM following a multi-stage framework (Fig. 1) to generate high-level mental health summaries based on timeline summaries. In the *map* stage, we instruct the model to provide inferences based on the timeline summary focusing on key clinical concepts (Appendix A, Table 8), such as presenting issues, inter/intra-personal patterns, and moments of change. Instructions and prompts are in Appendix B. In the *reduce* stage, we iteratively prompt the model to synthesise extracted observations into a concise summary.

4 Experiments

4.1 Evaluation Dataset Creation

We work with three clinical psychology graduate students who are fluent in English to create gold evidence-supported summaries. We use the dataset collected by Tsakalidis et al. (2022b) comprising 500 anonymised user timelines from Talklife. The number of posts in each timeline varies ([12-124]). We sample 30 timelines for annotators to highlight information related to individuals’ mental states and write high-level summaries which include diagnosis, intra- and interpersonal patterns and mental state changes over time. We use these 30 timelines for evaluation and 3 additional held out timelines

for development and in-context learning of key phrase extraction.

4.2 Models & Baselines

We compare our method against existing models for unsupervised abstractive opinion summarisation. For experiment settings, model specifications, and prompts refer to Appendix A and B.

Skeleton-based model is an unsupervised method proposed for story generation which encodes the skeleton (phrases that express the key meaning of sentences) to generate a detailed and polished sentence (Xu et al., 2018). We include it as one of the models to compare against as like TH-VAE it uses key phrases to generate a story/timeline in an unsupervised way. The key phrases provided are the same as for TH-VAE.

LLaMA We prompt a LLM to extract key phrases and then write TLDR-type summaries (Völske et al., 2017) focusing on the key phrases. Resulting summaries are similar to concise user-authored ones commonly found in social media data.

High-level Summary To obtain corresponding mental health summaries, we feed timeline summaries generated via TH-VAE and the above baselines into the LLM prompting framework outlined in §3.5. In addition, to see the benefits of timeline summarisation and specific clinical prompts, we implement a high-level and prompt-only **naive baseline**. It involves splitting timelines into chunks, prompting the LLM to write a mental health summary of each chunk, and rewriting the chunk-level summaries into a single coherent document.

4.3 Evaluation

We use summaries by clinical experts (§4.1) in automatic evaluation. In human evaluation we work with the same experts, and in this case they are asked to rate summaries for factual consistency, salient meaning preservation, and facets of usefulness.³ Details on procedure and metrics are in Appendix A.2-A.3.

Salient information preservation. We adapt MHIC (Srivastava et al., 2022) to assess whether timeline summaries capture clinically relevant information. Given evidence E and timeline sum-

³We merged aspects in human evaluation after a pilot, based on expert feedback. Given the LLM’s ability to output well-formed text, the cognitively taxing nature of the task, and time constraints, we prioritised aspects that demand domain expertise rather than general linguistic quality (e.g. fluency).

Timeline Summary	High-level Summary
I hate school. im so scared.i love the cat so much. i'm gonna go be counter for a while. unknown good night all! i hate my nose. fuck the sat. i got a commitment ring 3 i've got counterbodied under my eyes. i feel alone i feel like everyone hates me. i hate saying this because it sounds bad, but i really want someone to pay attention to me. solid my mom only cares about my brothers learning disabilities and her new boyfriend expressive. my boyfriend goes throug of ptsd like depression symptoms, so he gets distant sometimes levels. i literally have no friends argue. i've always been quiet but sometimes i just want someone never pay attention to me. i argue to know only when't be a body. i have to know. i'll have to be happy.	<p>The individual is likely experiencing a range of mental health issues, including depression, anxiety, and low self-esteem. They report feeling isolated, lonely, and frustrated, with a history of stressful life events. The individual's mood appears to be low, with expressions of sadness, frustration, and anger.</p> <p>It is evident that they have a profound craving for attention and connection with others, as well as a yearning for acceptance and validation. However, their fear of rejection and abandonment hinders them from forming and maintaining healthy relationships. Moreover, their self-criticism and preoccupation with perceived shortcomings indicate a lack of self-compassion and acceptance of their strengths and vulnerabilities.</p> <p>The individual's mood and well-being have been observed to fluctuate over time, with both positive and negative changes experienced. The individual has expressed a range of emotions, including sadness, loneliness, and frustration, as well as moments of happiness and positivity. Noteworthy positive changes include their excitement about having a cat and receiving a commitment ring, which are associated with positive emotions and a sense of joy. However, the individual also struggles with school and experiences anxiety and depression, which are linked to negative emotions such as sadness, fear, and frustration.</p>

Table 1: Example TH-VAE timeline summary and its high-level summary. Examples for all systems in Appendix C.

mary sentences T , we average the maximum recall-oriented BERTScore (Zhang et al., 2020):

$$\text{MHIC}_{sem} = \frac{1}{|E|} \sum_{e \in E} \max_{t \in T} R_{\text{BERT}}(e, t)$$

Factual consistency. To measure whether timeline summaries are consistent with original timelines, we apply the faithfulness score used in traditional summary evaluation with a modified procedure that splits timelines into chunks. Given a chunked timeline D and its timeline summary T , for every sentence t in T , we calculate the maximum probability of a timeline chunk d in D entailing t using a NLI model and average across all summary sentences.

$$\text{FC}_{\text{Timeline}} = \frac{1}{|T|} \sum_{t \in T} \max_{d \in D} \text{NLI}(\text{Entail}|d, t)$$

Next we assess the consistency of high-level model-generated summaries S with human-written ones G , where consistency is the absence of contradiction. We define C to be a function that quantifies the consistency of text B based on text A :

$$C(A, B) = \frac{1}{|A||B|} \sum_{a \in A} \sum_{b \in B} (1 - \text{NLI}(\text{Contradict}|a, b))$$

We calculate the consistency of high-level summaries to gold summaries as $\text{FC}_{\text{Expert}} = C(G, S)$.

Evidence appropriateness. We measure the consistency of high-level summaries S to their accompanying timeline summaries T via $\text{EA} = C(T, S)$.

Coherence. We estimate how easy it is to follow the summary and how effectively the mental health summary integrates information from the timeline summary using BARTScore (Yuan et al., 2021). We evaluate *logical* coherence via intra-summary NLI (IntraNLI), taking the mean consistency of each sentence against all other sentences to assess

Aspect	Metric	LLaMA	TH-VAE	Skeleton	Naive
SMP	MHIC _{sem}	.65	.66	.57	–
	FC _{Timeline}	.63	.63	.21	–
FC	FC _{Expert}	.95	.96	.95	.93
	EA	.97	.97	.95	–
EA	IntraNLI	.95	.96	.95	.93
	BARTScore	-2.96	-3.10	-3.09	–
Coherence	PPL _{Timeline} (↓)	13.80	56.33	31.82	–
	PPL _{High-level} (↓)	9.32	9.30	9.45	11.38

Table 2: Automatic evaluation for salient meaning preservation (SMP), factual consistency (FC), evidence appropriateness (EA), coherence, and fluency. Higher is better, except for PPL. BARTScore uses log likelihood, hence higher (less negative) is better. Best in **bold**, significant improvement over second-best underlined.

the logical interconnection of information within the mental health summary.

Fluency. We separately estimate fluency for timeline and high-level summaries using perplexity (PPL) under GPT-2-XL (Radford et al., 2019).

Usefulness. Summaries should help the clinician understand the client's condition. This is assessed via human evaluation only, with respect to general usefulness and specific categories (diagnosis, intra- and interpersonal patterns and MoC). Details are available in the Appendix in Table 8.

5 Results

5.1 Automatic evaluation

Table 1 shows example summaries. We perform two-tailed permutation tests in our comparisons reporting statistical significance at $\alpha = .05$. TH-VAE and LLaMA generated significantly higher quality summaries compared to other baselines. TH-VAE and LLaMA were comparable on most metrics, preserving mental health information (MHIC_{sem}) while being similarly consistent with the source

(FC_{Timeline}) in timeline summaries and factually consistent with human-written references in high-level mental health summaries (FC_{Expert}).

Two-tailed permutation tests showed that LLaMA timeline summaries were significantly more fluent (PPL_{Timeline}), in line with its tendency to normalise text (see examples, Appendix B). These tests also indicate that high-level summaries were comparably coherent in terms of ease of reading and integrating information from timeline summaries (BARTScore). This is expected since all methods used the same prompting framework to generate high-level summaries. However, TH-VAE achieved significantly higher IntraNLI, suggesting its timeline summaries allow for more logically coherent synthesis of detailed clinical information.

5.2 Human evaluation

We selected three systems for human evaluation: LLaMA, TH-VAE, and the naive LLaMA baseline. This allows us to compare top-performing models and understand how removing timeline summarisation and clinical prompting steps may impact summary quality as perceived by human judges. TH-VAE produced summaries considered the most factually consistent and useful in summarising changes (MoC) among compared models. Human judges found LLaMA summaries generated with clinical prompts to be most useful in other usefulness criteria, whereas LLaMA with a simple summarisation prompt was consistently least useful. Notably, LLaMA summaries *without* clinical prompts were rated as more factually consistent than those *with* clinical prompts, suggesting they adhered to the source timeline, but were impacted by lack of guidance (Table 3). We present a clinician-written qualitative evaluation in §5.4, which shows that LLaMA timelines present more hallucinations than TH-VAE.

Aspect	LLaMA	TH-VAE	Naive
Factual Consistency	3.08	3.35	3.28
Usefulness (General)	3.38	3.28	2.55
(Diagnosis)	3.40	3.25	2.93
(Inter- & Intrapersonal)	3.48	3.33	2.23
(MoC)	3.30	3.35	1.18

Table 3: Human evaluation results based on 5-point Likert scales (1 is worst, 5 is best). Best in **bold**.

5.3 Ablation

We performed ablation studies to investigate the importance of key phrases (§3.2) and elaborate

clinical-expert guided prompts for the final summary generation (§3.5) in TH-VAE and LLaMA. We experimented with (a) removing keyphrases but keeping the expert-guided prompts and (b) keeping the keyphrases, but prompting the LLM to summarise the high-level summary directly without any thematic guidance.

Aspect	Metric	TH-VAE	-keyphrases	-clinical prompts
SMP	MHIC _{sem}	.66	.62	-
FC	FC _{Timeline}	.63	.52	-
	FC _{Expert}	.96	.95	.91
EA	EA	.97	.94	.93
	IntraNLI	.96	.95	.94
Coherence	BARTScore	-3.10	-3.08	-2.74
	Fluency	PPL _{Timeline} (↓)	56.33	81.45
	PPL _{High-level} (↓)	9.30	9.38	13.62

Table 4: Ablation results. Best in **bold**. TH-VAE without clinical prompts uses the same timeline summary as TH-VAE so repeated metrics were removed for brevity.

In both systems, removing keyphrases results in timeline summaries capturing less salient information (MHIC_{sem}), and degradation in logical connectedness (IntraNLI), evidence appropriateness (EA), and factual consistency with gold summaries (FC_{Expert}), showing that *keyphrases help focus generation on mental health related information*. In TH-VAE, removing keyphrases made timeline summaries less consistent with the source (FC_{Timeline}), and we observed the same trend to a greater extent when expert-guided prompts are removed. Thus, *the elaborate prompt does provide an efficient clinical guidance for the LLM to generate summaries*.

In LLaMA, removing keyphrases improves timeline summary faithfulness (FC_{Timeline}) at the expense of clinical informativeness (MHIC_{sem}). This shows the role of keyphrases guided by domain expertise as anchors in summaries of long texts. Consistency with experts (FC_{Expert}) is similar across ablation settings but highest when both are employed, underlining the importance of using these

Aspect	Metric	LLaMA	-keyphrases	-clinical prompts	Naive
SMP	MHIC _{sem}	.65	.59	-	-
FC	FC _{Timeline}	.63	.68	-	-
	FC _{Expert}	.95	.93	.93	.93
EA	EA	.97	.93	.94	-
	IntraNLI	.95	.89	.90	.93
Coherence	BARTScore	-2.96	-2.48	-2.61	-
	Fluency	PPL _{Timeline} (↓)	13.80	11.38	-
	PPL _{High-level} (↓)	9.32	13.78	11.62	11.38

Table 5: Ablation results. Best in **bold**. LLaMA without clinical prompts uses the same timeline summary as LLaMA so repeated metrics were removed for brevity. Naive uses neither keyphrases nor clinical prompts.

components in conjunction.

Aspect	Metric	TH-VAE	-keyphrases	-clinical prompts
SMP	MHIC _{sem}	.66	.62	–
FC	FC _{Timeline}	.63	.52	–
	FC _{Expert}	.96	.95	.91
EA	EA	.97	.94	.93
	IntraNLI	.96	.95	.94
Coherence	BARTScore	-3.10	-3.08	-2.74
	PPL _{Timeline} (↓)	56.33	81.45	–
Fluency	PPL _{High-level} (↓)	9.30	9.38	13.62

Table 6: Ablation results. Best in **bold**. TH-VAE without clinical prompts uses the same timeline summary as TH-VAE so repeated metrics were removed for brevity.

5.4 Qualitative discussion of clinical summaries

The TH-VAE and LLaMA offered moderately insightful details regarding the individual’s diagnosis. Their summaries accurately captured the general aspects of the diagnosis, focusing mainly on evident symptoms while overlooking some critical elements. LLaMA often reached conclusions that were not sufficiently supported by the evidence provided in the timeline. For example, both models noted the individual’s depression, self-harm, and suicidal thoughts but failed to recognize a clear eating disorder. Additionally, LLaMA suggested PTSD without substantial evidence in the provided timeline. However, these models were useful in shedding light on the individual’s self and relational dynamics over time. In contrast, the basic-prompt model presented a very broad summary, missing several vital details and failing to reflect significant clinical concepts. On the other hand, TH-VAE and LLaMA produced more comprehensive summaries, effectively highlighting crucial aspects of the individual’s self-perception, interpersonal relationships, and moments of change. Overall, from a clinical point of view, the quality of the summaries generated by the TH-VAE and LLaMA were quite

Aspect	Metric	LLaMA	-keyphrases	-clinical prompts	Naive
SMP	MHIC _{sem}	.65	.59	–	–
FC	FC _{Timeline}	.63	.68	–	–
	FC _{Expert}	.95	.93	.93	.93
EA	EA	.97	.93	.94	–
	IntraNLI	.95	.89	.90	.93
Coherence	BARTScore	-2.96	-2.48	-2.61	–
	PPL _{Timeline} (↓)	13.80	11.38	–	–
Fluency	PPL _{High-level} (↓)	9.32	13.78	11.62	11.38

Table 7: Ablation results. Best in **bold**. LLaMA without clinical prompts uses the same timeline summary as LLaMA so repeated metrics were removed for brevity. Naive uses neither keyphrases nor clinical prompts.

similar. LLaMA was only slightly lower in factual consistency than the TH-VAE. The TH-VAE and LLaMA were effective in summarizing the intrapersonal and interpersonal patterns and moments of change, but their depiction of diagnostic aspects was only moderately accurate, characterized by some inaccuracies and omissions.

6 Conclusions

We present a novel method for hybrid abstractive summarisation using hierarchical VAE and LLMs and the first approach to creating clinically meaningful mental health summaries from users’ social media timelines. Our approach results in summaries with a dual narrative perspective: high-level third person information useful for clinicians is combined with first person corresponding evidence from users’ timelines. Abstractive timeline summarisation is performed by three different systems (LLM-, TH-VAE- and skeleton-based) whose generation is guided by key-phrases obtained by an LLM through instruction prompting. High-level clinical summaries in third-person are generated by feeding the timeline summaries from all three systems into an LLM. Our proposed timeline summariser, TH-VAE, based on a hierarchical VAE for long texts, can capture long dependencies between sub-timelines and while LLM timeline summaries are the most fluent, they lag behind TH-VAE on logical coherence and factuality. From a clinical psychology viewpoint our work enables clinician access to consented clients’ social media data allowing them to understand changes in their mental state over time. Importantly it enables generation of automated summaries emphasizing essential clinical concepts which can aid mental health professionals to quickly grasp an individual’s psychological condition and progression.

7 Acknowledgements

This work was supported by a UKRI/EPSC Turing AI Fellowship to Maria Liakata (grant no. EP/V030302/1), Keystone grant funding to Liakata from Responsible Ai (grant no. EP/Y009800/1) and the Alan Turing Institute (grant no. EP/N510129/1). We are grateful to our reviewers and thank them for their support. We would finally like to thank our annotation team from Bar Ilan University, Barak Dakar, Dana Rafael, and Hadas Shualy, for their crucial work on highlighting evidence in timelines, writing expert summaries and conducting human evaluation.

Limitations

Our work considers the segmentation of timelines in terms of moments of change as changes in an individual's mood judged on the basis of their self-disclosure of their well-being. This is faced by two limiting factors: (a) users may not be self-disclosing important aspects of their daily lives and (b) while also (Hills et al., 2023b) segment user timelines based on moments of change in mood there may be other appropriate ways to effectively segment timelines into semantically related temporal units. For example timelines could be segmented based on symptoms or life events which could also be evolving over time. Empirically we have not found topics to be an effective way of identifying sub-timelines and segments within a timeline but the best way of segmenting the timelines is an open research direction.

Though our models could be tested in cases of nonself-disclosure (given the appropriate ground truth labels), the analysis and results presented in this work should not be used to infer any conclusion on such cases.

While we believe our methods for clinically meaningful longitudinal summarisation of social media data for mental health monitoring to be applicable to non-social media longitudinal data such as therapy sessions, this remains future work.

In the present study, we have conducted a comparison between timeline summarization using TH-VAE, skeleton-based and LLM-generated summaries. A further qualitative evaluation by a senior clinical therapist found that the summaries generated by Llama often reached conclusions that were not sufficiently supported by the evidence provided in the timeline, and were lower in factual consistency than the TH-VAE. The TH-VAE and Llama were effective in summarizing the intrapersonal and interpersonal patterns and moments of change, but their depiction of diagnostic aspects was only moderately accurate, characterized by some inaccuracies and omissions. These findings will help pinpoint areas where our models can be enhanced and refined.

Ethics Statement

Ethics institutional review board (IRB) approval was obtained from the corresponding ethics board of the lead University prior to engaging in this research study. Our work involves ethical considerations around the analysis of user generated content

shared on a peer support network (TalkLife). A license was obtained to work with the user data from TalkLife and a project proposal was submitted to them in order to embark on the project. The current paper focuses on the summarisation of users' social media timelines for mental health monitoring, by using moments of change (MoC) in mood as the anchors to segment timelines. These changes involve recognising sudden shifts in mood (switches or escalations). Expert clinical annotators were paid fairly in line with University payscales. They were alerted about potentially encountering disturbing content and were advised to take breaks. The annotations are used to provide examples to an in house LLM and evaluate natural language processing models for creating mental health summaries based on users social media timelines. Working with datasets such as TalkLife and data on online platforms where individuals disclose personal information involves ethical considerations (Mao et al., 2011; Keküllüoglu et al., 2020). Such considerations include careful analysis and data sharing policies to protect sensitive personal information. The data has been de-identified both at the time of sharing by TalkLife but also by the research team to make sure that no user handles and names are visible. Any examples used in the paper are paraphrased (generated summaries). Potential risks from the application of our work in being able to summarise the mental health of individuals based on their social media timelines are akin to those in earlier work on personal event identification from social media and the detection of suicidal ideation. Potential mitigation strategies include restricting access to the code base and corpus used for evaluation by requiring an NDA, as with other mental health datasets.

The final high level summaries in all cases are obtained by feeding the timeline summaries into an LLM. Given that LLMs are susceptible to factual inaccuracies, often referred to as 'hallucinations,' and tend to exhibit biases, the clinical summaries they generate may contain errors that could have serious consequences in the realm of mental health decision-making. These inaccuracies can encompass anything from flawed interpretations of the timeline data to incorrect diagnoses and even recommendations for potentially harmful treatments. Mental health professionals must exercise caution when relying on such generated clinical summaries. These summaries should not serve as substitutes

for therapists in making clinical judgments. Instead, well-trained therapists must skillfully incorporate these summaries into their clinical thought processes and practices. Significant efforts are required to establish the scientific validity of the clinical benefits offered by these summaries before they can be integrated into routine clinical practice.

References

- Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A Saurous, and Kevin Murphy. 2018. Fixing a broken elbow. In *International conference on machine learning*, pages 159–168. PMLR.
- James Allan, Rahul Gupta, and Vikas Khandelwal. 2001. [Temporal summaries of news topics](#). In *SIGIR 2001: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, September 9-13, 2001, New Orleans, Louisiana, USA*, pages 10–18. ACM.
- Adrian Benton, Margaret Mitchell, and Dirk Hovy. 2017. [Multitask learning for mental health conditions with limited social media data](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 152–162, Valencia, Spain. Association for Computational Linguistics.
- Stevie Chancellor and Munmun De Choudhury. 2020. Methods in predictive techniques for mental health status on social media: a critical review. *NPJ digital medicine*, 3(1):43.
- Yi Chang, Jiliang Tang, Dawei Yin, Makoto Yamada, and Yan Liu. 2016. Timeline summarization from social media with life cycle models. In *IJCAI*, pages 3698–3704.
- Xiuying Chen, Zhangming Chan, Shen Gao, Meng-Hsuan Yu, Dongyan Zhao, and Rui Yan. 2019. Learning towards abstractive timeline summarization. In *IJCAI*, pages 4939–4945.
- Xiuying Chen, Mingzhe Li, Shen Gao, Zhangming Chan, Dongyan Zhao, Xin Gao, Xiangliang Zhang, and Rui Yan. 2023. Follow the timeline! generating an abstractive and extractive timeline summary in chronological order. *ACM Transactions on Information Systems*, 41(1):1–30.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*.
- Arman Cohan, Bart Desmet, Andrew Yates, Luca Soldaini, Sean MacAvaney, and Nazli Goharian. 2018. [SMHD: a large-scale resource for exploring online language usage for multiple mental health conditions](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1485–1497, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. Quantifying mental health signals in twitter. In *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*, pages 51–60.
- Paul Crits-Christoph and Mary Beth Connolly Gibbons. 2021. Psychotherapy process-outcome research: Advances in understanding causal connections. *Bergin and Garfield’s handbook of psychotherapy and behavior change*, pages 263–296.
- Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. 2013. Predicting depression via social media. In *Proceedings of the international AAAI conference on web and social media*, volume 7, pages 128–137.
- Tracy D Eells. 2022. *Handbook of psychotherapy case formulation*. Guilford Publications.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint*.
- Anthony Hills, Adam Tsakalidis, Federico Nanni, Ioannis Zachos, and Maria Liakata. 2023a. [Creation and evaluation of timelines for longitudinal user posts](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3791–3804, Dubrovnik, Croatia. Association for Computational Linguistics.
- Anthony Hills, Adam Tsakalidis, Federico Nanni, Ioannis Zachos, and Maria Liakata. 2023b. [Creation and evaluation of timelines for longitudinal user posts](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3791–3804, Dubrovnik, Croatia. Association for Computational Linguistics.
- Ahmed Husseini Orabi, Prasadith Buddhitha, Mahmoud Husseini Orabi, and Diana Inkpen. 2018. [Deep learning for depression detection of Twitter users](#). In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, pages 88–97, New Orleans, LA. Association for Computational Linguistics.
- Alan E Kazdin. 2021. Extending the scalability and reach of psychosocial interventions.
- Dilara Keküllüoğlu, Walid Magdy, and Kami Vaniea. 2020. Analysing privacy leakage of life events on twitter. In *Proceedings of the 12th ACM Conference on Web Science*, pages 287–294.

- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Alexej Klushyn, Nutan Chen, Richard Kurle, Botond Cseke, and Patrick van der Smagt. 2019. Learning hierarchical priors in vaes. *Advances in neural information processing systems*, 32.
- Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2022. SummaC: Re-visiting NLI-based models for inconsistency detection in summarization. *Transactions of the Association for Computational Linguistics*, 10:163–177.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Jiwei Li and Claire Cardie. 2014. Timeline generation: Tracking individuals on twitter. In *Proceedings of the 23rd International Conference on World Wide Web, WWW '14*, page 643–652, New York, NY, USA. Association for Computing Machinery.
- Manling Li, Tengfei Ma, Mo Yu, Lingfei Wu, Tian Gao, Heng Ji, and Kathleen McKeown. 2021. Timeline summarization based on event graph compression via time-aware optimal transport. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6443–6456, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.
- Min Lin, Qiang Chen, and Shuicheng Yan. 2013. Network in network. *CoRR*, abs/1312.4400.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Ro{bert}a: A robustly optimized {bert} pretraining approach.
- Gaur Manas, Vamsi Aribandi, Ugur Kursuncu, Amanuel Alambo, Valerie L Shalin, Krishnaprasad Thirunarayan, Jonathan Beich, Meera Narasimhan, and Amit Sheth. 2021. Knowledge-infused abstractive summarization of clinical diagnostic interviews: Framework development study. *JMIR Ment Health*, 8(5):e20865.
- Huina Mao, Xin Shuai, and Apu Kapadia. 2011. Loose tweets: An analysis of privacy leaks on twitter. In *Proceedings of the 10th Annual ACM Workshop on Privacy in the Electronic Society, WPES '11*, page 1–12, New York, NY, USA. Association for Computing Machinery.
- Sebastian Martschat and Katja Markert. 2018. A temporally sensitive submodularity framework for timeline summarization. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 230–240, Brussels, Belgium. Association for Computational Linguistics.
- Joshua Maynez, Priyanka Agrawal, and Sebastian Gehrmann. 2023. Benchmarking large language model capabilities for conditional generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9194–9213, Toronto, Canada. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. Adversarial NLI: A new benchmark for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Rajesh Ranganath, Dustin Tran, and David Blei. 2016. Hierarchical variational models. In *International conference on machine learning*, pages 324–333. PMLR.
- Zhaochun Ren, Shangsong Liang, Edgar Meij, and Maarten de Rijke. 2013. Personalized time-aware tweets summarization. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 513–522.
- Brian Schwartz, Jessica Uhl, and Dana Atzil-Slonim. 2023. Assessments and measures in psychotherapy research: going beyond self-report data. *Frontiers in Psychiatry*, 14:1276222.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano, Alex Wang, and Patrick Gallinari. 2021. QuestEval: Summarization asks for fact-based evaluation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6594–6604, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. 2016. Ladder variational autoencoders. *Advances in neural information processing systems*, 29.
- Jiayu Song, Iman Munire Bilal, Adam Tsakalidis, Rob Procter, and Maria Liakata. 2022. [Unsupervised opinion summarisation in the Wasserstein space](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8592–8607, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Sajad Sotudeh, Nazli Goharian, and Zachary Young. 2022. [MentSum: A resource for exploring summarization of mental health online posts](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2682–2692, Marseille, France. European Language Resources Association.
- Aseem Srivastava, Tharun Suresh, Sarah P. Lord, Md Shad Akhtar, and Tanmoy Chakraborty. 2022. [Counseling summarization using mental health knowledge guided utterance filtering](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*, page 3920–3930, New York, NY, USA. Association for Computing Machinery.
- Julius Steen and Katja Markert. 2019. [Abstractive timeline summarization](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 21–31, Hong Kong, China. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Adam Tsakalidis, Jenny Chim, Iman Munire Bilal, Ayah Zirikly, Dana Atzil-Slonim, Federico Nanni, Philip Resnik, Manas Gaur, Kaushik Roy, Becky Inkster, Jeff Leintz, and Maria Liakata. 2022a. [Overview of the CLPsych 2022 shared task: Capturing moments of change in longitudinal user posts](#). In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 184–198, Seattle, USA. Association for Computational Linguistics.
- Adam Tsakalidis, Federico Nanni, Anthony Hills, Jenny Chim, Jiayu Song, and Maria Liakata. 2022b. [Identifying moments of change from longitudinal user text](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4647–4660, Dublin, Ireland. Association for Computational Linguistics.
- Arash Vahdat and Jan Kautz. 2020. Nvae: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33:19667–19679.
- Sumithra Velupillai, Hanna Suominen, Maria Liakata, Angus Roberts, Anoop D Shah, Katherine Morley, David Osborn, Joseph Hayes, Robert Stewart, Johnny Downs, et al. 2018. Using clinical natural language processing for health outcomes research: overview and actionable suggestions for future advances. *Journal of biomedical informatics*, 88:11–19.
- Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. [TL;DR: Mining Reddit to learn automatic summarization](#). In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.
- Shang Wang, Zhiwei Yang, and Yi Chang. 2021. Bringing order to episodes: Mining timeline in social media. *Neurocomputing*, 450:80–90.
- Yiming Wang, Zhuosheng Zhang, and Rui Wang. 2023. [Element-aware summarization with large language models: Expert-aligned evaluation and chain-of-thought method](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8640–8665, Toronto, Canada. Association for Computational Linguistics.
- Jingjing Xu, Xuancheng Ren, Yi Zhang, Qi Zeng, Xiaoyan Cai, and Xu Sun. 2018. [A skeleton-based model for promoting coherence among sentences in narrative story generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4306–4315, Brussels, Belgium. Association for Computational Linguistics.
- Kailai Yang, Shaoxiong Ji, Tianlin Zhang, Qianqian Xie, Ziyang Kuang, and Sophia Ananiadou. 2023. [Towards interpretable mental health analysis with large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6056–6077, Singapore. Association for Computational Linguistics.
- Andrew Yates, Arman Cohan, and Nazli Goharian. 2017. [Depression and self-harm risk assessment in online forums](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2968–2978, Copenhagen, Denmark. Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [BartScore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 27263–27277. Curran Associates, Inc.
- Tianyi Zhang, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BertScore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2023a. [Benchmarking large language models for news summarization](#).

Youmin Zhang, Fabio Tosi, Stefano Mattocchia, and Matteo Poggi. 2023b. Go-slam: Global optimization for consistent 3d instant reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3727–3737.

Ayah Ziriky, Philip Resnik, Ozlem Uzuner, and Kristy Hollingshead. 2019. Clpsych 2019 shared task: Predicting the degree of suicide risk in reddit posts. In *Proceedings of the sixth workshop on computational linguistics and clinical psychology*, pages 24–33.

A Appendix

A.1 Experimental Settings

TH-VAE We load pre-trained parameters from BART-BASE (Lewis et al., 2020) for pre-trained word embedding and 6 transformer decoder layers in the model. We set the dimensional size of z_i to be the same as the size of word embeddings (768). We set the number of latent variables l as 5, which has the best performance on our dataset. In addition, we set the number of cells in **block** is 3, and the number of cells in each **group** is 1. We use the Adam optimizer (Kingma and Ba, 2015) (learning rate: 5×10^{-4}).

LLM Our experiments use 4bit-quantized LLAMA-2 (Touvron et al., 2023). For keyphrase extraction, we use few-shot prompting on the base pre-trained model LLAMA-2-13B. In zero-shot prompting tasks with detailed instructions (i.e. mental health related inferences), we use LLAMA-2-13B-CHAT to take advantage of its fine-tuning on instruction datasets and human preferences.

Compute We trained TH-VAE with 2 hours on 1 GPU, and spent 20 GPU hours for generating high-level summaries.

A.2 Evaluation Metrics

NLI On metrics that require NLI, we use a ROBERTA model (Liu et al., 2020) fine-tuned on fact verification and NLI (Nie et al., 2020): https://huggingface.co/ynie/roberta-large-snli_mnli_fever_anli_R1_R2_R3-nli. When evaluating evidence appropriateness, we consider text from the timeline summary to be the premise and text from the high-level summary to be the hypothesis. When running the NLI model, we prefix every sentence in the timeline summary with "The individual wrote:". While we did not find statistically significant differences between the selected prefix vs. no prefix and vs. similar alternatives, we decided on prefixing as empirically it seemed to help the NLI model on noisy premises.

Salient Meaning Preservation: MHIC We make the following changes to MHIC (Srivastava et al., 2022). First, instead of ROUGE we measure semantic embedding similarity using BERTScore. Second, instead of computing separate scores based on hard utterance categories, we compute a unified one using the semantic intersection of information highlighted by annotators. We find the intersection of highlighted timeline spans among annotators by (1) directly extracting intersecting substrings, (2) computing pairwise cosine similarity across evidence spans, keeping pairs with similarity $\geq .60$, then selecting the shorter span from each pair, and (3) deduplicating evidences from these steps. We use the sentence-transformers library and MSMARCO-DISTILBERT-BASE-V3 embeddings.

Factual Consistency For FC_{Timeline} , we chunk timeline texts with a cutoff of 60 tokens to match input lengths in the NLI model’s training data.

A.3 Annotation & Human Evaluation

We ran training sessions for both summarisation and evaluation tasks under the supervision of a senior clinical expert to ensure annotators clearly understood task requirements.

Summarisation During training, annotators were introduced to the dataset and task, and provided with guidelines. After reviewing the guidelines and held out examples, we worked on a timeline reserved for annotator training together. The annotators separately worked on another timeline reserved for training. We compared annotations during a second training session. Once we were confident that the team had a shared understanding of task requirements, annotators proceeded to actual timelines used for testing in this paper.

Evaluation We provided the annotators with guidelines and introduced the evaluation task as well as criteria (see Appendix D) in the first training session. We checked agreement on a small set of timelines, then after discussion and clarifications on a second session they were asked to proceed to rating summaries on the remaining test timelines. During evaluation, annotators were presented data on a timeline-by-timeline basis. When rating summaries for a timeline, they would receive the summaries in a randomly shuffled order.

A.4 Clinical Concepts

<p>Diagnosis</p> <p>Presenting issues (what bothers the person and causes distress; triggers).</p> <p>Mental health symptoms, level of functioning, well-being.</p> <p>Physical symptoms.</p> <p>Risk assessment (previous suicidal attempts, intent to suicide, access to lethal means; hopelessness, social isolation, recent loss, impulsivity, dramatic mood swings).</p> <p>Motivation to change.</p> <p>Lifestyle (diet, physical activity, sleep, alcohol/drug/tobacco use, occupation, environment, screen time, healthcare practices).</p> <p>Agency, coping mechanisms, strengths and resources (what helps the person, how they typically cope with stress and difficulties, resilience).</p> <p>Meaning/goals/direction in life.</p> <p>Behaviour (adaptive and maladaptive behavioural patterns).</p> <p>Important events (present and past events in life; traumatic events).</p>
<p>Intrapersonal and Interpersonal patterns</p> <p>Main need/wish/desire.</p> <p>Interpersonal relationships (repetitive interpersonal pattern; conflicts; how others are perceived; social support).</p> <p>Self perception, self esteem.</p>
<p>Moments of change</p> <p>Emotion (sad, happy, etc).</p> <p>Arousal level (high/low).</p> <p>Emotion regulation strategies.</p> <p>Switches (drastic change of one's mood).</p> <p>Escalations (intensification in one's mood).</p> <p>Self understanding (insights about the self and the relationship; ability to reflect and understand repetitive patterns).</p>

Table 8: Clinical concepts important to therapeutic approaches. Our task is to capture and summarize them if such information is present in user timelines.

Appendix B. Instruction Prompts

B.1 Keyphrase Extraction

Task: Choose key phrases in the following posts.

Text: {example post 1}

Keyphrases:{expert key phrases 1}

Text: {example post 2}

Keyphrases:{expert key phrases 2}

Text: {concatenated posts to annotate}

Keyphrases:

B.2 Timeline Summarisation: LLaMA

Write a TLDR as the user (first-person), focusing on the keyphrases.

Keyphrases: {extracted keyphrases}

{concatenated posts to summarise}

TLDR:

B.3 High-level only: LLaMA Naive Baseline

You are a helpful assistant to an expert therapist who reads social media chronological text written by an individual who has mental health concerns.

Summarize the texts below:

{timeline chunk concatenated}

B.4 Map Prompt: Diagnosis

Your goal is to describe the individual's mental state and identify potential indicators that may suggest a mental health diagnosis, considering the following aspects:

1. Presenting Issues: What are the main concerns or stressors evident in the individual's posts? What triggers seem to affect their mental state?
2. Mental Health Symptoms and Functioning: Does the individual exhibit any mental health symptoms? How are their mood, energy levels, and interest in usual activities? Are there noticeable changes in sleep patterns, appetite, concentration, or social interactions? How do they describe their overall well-being and functioning in daily activities?
3. Mental Health Treatment History: Has the individual been in contact with mental health professionals such as psychiatrists or psychotherapists? Are there mentions of current or past outpatient or inpatient mental health treatments? Do they reference taking psychiatric medications?
4. Physical Health: Are there any current or past physical health issues, medical conditions, hospitalizations, or surgeries mentioned?
5. Risk Assessment: Is there evidence of previous suicidal attempts or current suicidal thoughts? Do they have access to lethal means? What level of hopelessness is expressed? Do they discuss social isolation, recent losses, impulsivity, or dramatic mood swings?
6. Lifestyle Factors: What do the individual's posts reveal about their lifestyle habits, such as diet, physical activity, sleep patterns, occupation, environment, screen time, and healthcare practices?
7. Substance Use: Are there any references by the individual to the use of substances like alcohol, drugs, or tobacco? If so, how frequently do they use these substances?
8. Significant Life Events and Family History: Are there references to significant life events like divorce, loss of a close person, experiences of abuse, or neglect? Is there any mention of psychiatric problems or treatments among family members?
9. Motivation and Coping Strategies: What does the individual express about their motivation for change? How do they cope with stress and difficulties? What strengths and resources do they have? What seems to help them? How resilient do they appear? Do they discuss having direction, meaning, or goals in their life?

You must not make anything up. Keep the description concise and only describe observations if they are fully supported by the text.

Here are the texts:

{Timeline summary}

B.5 Map Prompt: Intrapersonal and Interpersonal Patterns

Your goal is to identify the person's main intrapersonal and interpersonal pattern, considering the following aspects:

1. Wish/Need/desire/intention/expectation: What is the person's most dominant need, desire, intention, expectation from others and from themselves? Are there any other needs or wishes that might be indicated in a less obvious way? How well does the individual communicate their needs/ wishes with others?
2. Response of Others: How does this person typically perceive the emotions, behaviors, and thoughts of others? Are there any other perceptions of the other that might be indicated in a less obvious way? Is the individual capable of acknowledging the complex nature of others?
3. Response of self to others: How does this person tend to feel and react to others? Are there any other reactions towards others that might be indicated in a less obvious way?
4. Response of self to self: What is the individual's most dominant emotion, behavior and cognition toward oneself? Are there any other emotions and cognitions towards the self that might be indicated in a less obvious way? What is the level of self-compassion and acceptance of strengths and vulnerabilities?
5. Patterns: What is this individual's predominant dysfunctional intrapersonal and interpersonal pattern? What is this individual's predominant adaptive intrapersonal and interpersonal pattern?

You must not make anything up. Keep the description concise and only describe observations if they are fully supported by the text.

Here are the texts:

{Timeline summary}

B.6 Map Prompt: Moments of Change (MoC)

Your goal is to understand and summarise changes over time in this individual's mood, well-being and functioning (individual/self well being; interpersonal well-being including family and close relationships; social well-being, including work, school and friends). From the following text, identify whether there are changes in the individual's expressed mood, well-being, and functioning.

- a. Overall, does the mood/ well being/ functioning stay consistent or fluctuate over time?
- b. Are there specific positive to negative changes? Which events (personal or interpersonal) are associated with these changes?
- c. Are there specific negative to positive changes? Which events (personal or interpersonal) are associated with these changes?

Answer each concisely and reply None if not available. You must not make anything up. Keep the description concise and only describe observations if they are fully supported by the text.

Here are the texts:

{Timeline summary}

B.7 Reduce Prompt

Rewrite these sentences about the individual so that they form a single clear and coherent document. Combine redundant sentences, but do not remove any clinical information or information about interpersonal relationships. You must not make anything up.

{Previously generated summaries}

Appendix C. Example Summaries

C.1 Timeline Summaries

TH-VAE I hate school. im so scared.i love the cat so much. i'm gonna go be counter for a while. unknown good night all! i hate my nose. fuck the sat. i got a commitment ring 3 i've got counterbodied under my eyes. i feel alone i feel like everyone hates me. i hate saying this because it sounds bad, but i really want someone to pay attention to me. solid my mom only cares about my brothers learning disabilities and her new boyfriend expressive. my boyfriend goes throug of ptsd like depression symptoms, so he gets distant sometimes levels. i literally have no friends argue. i've always been quiet but sometimes i just want someone never pay attention to me. i argue to know only when't be a body. i have to know. i'll have to be happy.

LLaMA Need help passing a breathalyzer test for homecoming, worried about proposal, struggling with school, and feeling insecure about appearance. I love my cat so much, but I have no motivation and can't get out of bed. I really enjoy making others happy, but I struggle with feeling alone and like everyone hates me. What is a healthy weight for a 5'2" 17 year old girl? Is there anything besides vodka that can help me get through homecoming? Someone please pay attention to me. I feel ignored by my family and my chat isn't working. I have no friends and my boyfriend has PTSD-like symptoms. I promise I'm not ignoring you, I just need someone to talk to. My cat snoozes and it's wonderful. "I just wanna be kissed" but I hate my nose and have dark circles under my eyes.

Skeleton don't know what to do anymore i don't want to go to homecoming but i can't help my own time i feel like i'm going to try to share my feelings but i know that i'm feeling alone but i just want someone to pay attention to me. i love my cat so much. i can tell you that i feel good at the same time i'm in love with my partner and i'm glad i could turn up to be my best friend and i know what it happened to me i bet you're going to make me feel better cause i'm sure if i can turn my homecoming if i do help i will be happy but i like a lot of people i want to tell you i appreciate you but i won't help you just want you to go to sleep but i didn't really want to talk to you my way to help my mom so i can help you and make you feel better at the end of my life. i love him so much and i bet it's hard to me but i hope you were a good time but i need to be able to see you but you don't like you but it's okay to see your

C.2 High-level Summaries

TH-VAE The individual is likely experiencing a range of mental health issues, including depression, anxiety, and low self-esteem. They report feeling isolated, lonely, and frustrated, with a history of stressful life events. The individual's mood appears to be low, with expressions of sadness, frustration, and anger.

It is evident that they have a profound craving for attention and connection with others, as well as a yearning for acceptance and validation. However, their fear of rejection and abandonment hinders them from forming and maintaining healthy relationships. Moreover, their self-criticism and preoccupation with perceived shortcomings indicate a lack of self-compassion and acceptance of their strengths and vulnerabilities.

The individual's mood and well-being have been observed to fluctuate over time, with both positive and negative changes experienced. The individual has expressed a range of emotions, including sadness, loneliness, and frustration, as well as moments of happiness and positivity. Noteworthy positive changes include their excitement about having a cat and receiving a commitment ring, which are associated with positive emotions and a sense of joy. However, the individual also struggles with school and experiences anxiety and depression, which are linked to negative emotions such as sadness, fear, and frustration.

LLaMA Based on the information provided, the individual is experiencing a range of mental health symptoms, including depressive and anxiety symptoms, low self-esteem, and suicidal ideation, which are likely interfering with their daily functioning and quality of life. Addressing lifestyle habits such as excessive screen time and lack of physical activity may also be helpful in managing symptoms.

The individual's primary intrapersonal and interpersonal patterns appear to be a mix of dependent and self-centered tendencies, with a lack of self-awareness and self-regulation. These patterns are characterized by negative self-talk, low self-esteem, and a fear of being alone, leading to a reliance on alcohol to cope with stress. The individual also exhibits difficulty regulating their emotions and behaviors, and seeks validation and attention from others while having difficulty setting healthy boundaries.

The individual experiences fluctuating mood and well-being, with feelings of worry, insecurity, and loneliness, as well as moments of joy and happiness when they make others happy. Despite these positive interactions, the individual struggles with negative thoughts and feelings of being ignored by family and friends. They desire positive changes and may benefit from seeking professional help to address symptoms of depression, anxiety, and low self-esteem.

Skeleton Based on the provided social media texts, the individual is struggling with multiple mental health issues, including depression, anxiety, low self-esteem, and possible borderline personality disorder. Their mood appears low, with expressions of feeling alone, lost, and tired, which may indicate depression and anxiety. The individual's desire to "feel better" and "turn up" suggests a willingness to seek help and improve their mental state. The lack of direct references to mental health treatment history and their expressed desire for improvement suggest that they may be open to seeking professional help. Their statements about feeling alone and lost, as well as their desire to "turn up," may indicate a higher risk of suicide.

The individual longs for human connection and yearns for validation from others, yet they feel overwhelmed by their own emotions. They find solace in their attachment to their cat and appreciate their partner's support. However, they still experience a sense of loneliness and disconnection, which they introspectively examine. The individual's primary intrapersonal patterns may be characterized by isolation and disconnection, while their adaptive patterns may involve self-awareness and introspection.

The individual's mood and well-being appear to fluctuate over time, with both positive and negative emotions being experienced. They express feelings of happiness and appreciation for their partner and pets, but also mention feeling alone and isolated, which may suggest a decline in their well-being. The text is written in a stream of consciousness style, providing a glimpse into the individual's inner thoughts and emotions, but it does not provide a complete picture of their mood and well-being.

LLaMA Naive Based on the information provided, it appears that the individual is seeking emotional support and connection. They mention feeling alone and isolated, and express a desire for someone to pay attention to them and listen to their needs. The individual is feeling overwhelmed and frustrated, expressing a desire to leave and pursue their dreams. They are struggling with negative self-talk and body image issues, feeling ignored and alone despite expressing happiness and positivity. They use humor and sarcasm to cope with their emotions, but also mention feeling scared and insecure. The individual is seeking alternative ways to cope with their feelings besides drinking vodka.

Appendix D. Human Evaluation Criteria

Factual consistency	<p>1 - Not at all factually consistent: The summary contains significant inaccuracies or misrepresentations, completely misaligning with the timeline's content.</p> <p>2 - Mostly not factually consistent: The summary contains significant inaccuracies or misrepresentations, poorly reflecting the timeline's content.</p> <p>3 - Somewhat factually consistent: The summary is somewhat accurate, with several inaccuracies or omissions, but retains a basic reflection of the timeline's content.</p> <p>4 - Mostly factually consistent: The summary is largely accurate, with minor inaccuracies or omissions that do not majorly distort overall understanding.</p> <p>5 - Fully factually consistent: The summary is completely accurate, perfectly aligning with the timeline's content without discrepancies.</p>
General usefulness and Salient meaning preservation	<p>1 - Not at all useful: The summary fails to capture any essential information, significantly misrepresenting or omitting critical aspects of the individual's condition.</p> <p>2 - Slightly useful: The summary includes some important details but primarily focuses on irrelevant or less critical information.</p> <p>3 - Moderately useful: The summary captures important information but still includes less relevant details or omits minor key elements.</p> <p>4 - Very useful: The summary highlights most of the crucial information, with only minor irrelevant details.</p> <p>5 - Extremely useful: The summary encapsulates all critical information, providing a comprehensive and clear understanding of the individual's condition, without providing irrelevant information.</p>
Usefulness (diagnosis)	<p>1 - Not at all useful: The summary fails to provide information regarding the individual's diagnosis, or it clearly distorts the individual's diagnosis by incorrectly identifying diagnostic elements.</p> <p>2 - Slightly useful: The summary provides minimal information related to the individual's diagnosis. While the summary includes some correct diagnostic elements, it generally contains irrelevant or incorrect details and omissions.</p> <p>3 - Moderately useful: The summary is generally accurate about the individual's diagnosis but only describes the more obvious aspects, with some information possibly missing or unclear.</p> <p>4 - Very useful: The summary accurately identifies the individual's diagnosis and captures almost all the essential diagnostic information with only minor gaps.</p> <p>5 - Extremely useful: The summary is comprehensive and accurately details all aspects of the individual's diagnosis.</p>

Usefulness (interpersonal and intrapersonal pattern)

The summary provides helpful information about the individuals' main needs and patterns of self and other relationships.

1 - Not at all useful: The summary provides no insight into the individual's interpersonal and intrapersonal patterns.
 2 - Slightly useful: The summary provides a minimal understanding of interpersonal and intrapersonal patterns.
 3 - Moderately useful: The summary covers some key aspects of the individual's interpersonal and intrapersonal patterns but may lack depth or miss important elements.
 4 - Very useful: The summary provides a comprehensive overview of the individual's interpersonal and intrapersonal patterns, with only slight gaps or generalizations.
 5 - Extremely useful: The summary gives a detailed and complete understanding of the individual's interpersonal and intrapersonal patterns.

Usefulness (moments of change)

The summary provides useful information about the individual's changes over time in emotion/cognition and behaviour. Where appropriate, it should help connect information between events and the individual's responses.

1 - Not at all useful: The summary fails to provide any accurate information about whether there are changes in the individual over time.
 2 - Slightly useful: The summary includes information about changes, but they are generally inaccurate and overlook key developments/connections, or they generally contain irrelevant information.
 3 - Moderately useful: The summary accurately describes whether there are changes, although there may be some weaknesses or omissions as well as irrelevant information.
 4 - Very useful: The summary accurately describes whether there are changes and where available offers helpful insights.
 5 - Extremely useful: The summary accurately describes whether there are changes and where available provides clear, well-connected insights about the individual's development over time.
