

# Constructing a mental health analysis system for social media using large language models

*Jinming Zion Liang*

Sendelta International Academy, Shenzhen, Guangdong Province, China

785286205@qq.com

---

**Abstract.** As social media platforms have penetrated every aspect of people's lives, many mental health problems have also arisen along with them. In today's digital age, analyzing mental health trends through these platforms has become critical. In this study, we present a system designed to identify the mental health trends of Weibo users by extracting and analyzing the content posted by users on Weibo, China's leading social media platform. The system is mainly composed of two parts: a data acquisition module and an analysis module. The data collection module uses the Python-based web scraping tool Scrapy to scrape comments from popular topics on Weibo. At the heart of the analysis module is a large language model fine-tuned from a psychological database. The module assesses the topic and specific content of the posts, scoring comments based on criteria such as positivity, alignment with mood disorders, and potential signs of psychoactive substance use. This data is stored and mediated using the relational database MySQL, and then analyzed and visualized using advanced data analysis tools. Through this method, we can timely and comprehensively monitor the mental health status of social media platforms, and provide a solid foundation for further academic research on public mental health.

**Keywords:** ChatGLM, Scrapy, Large Language Models, Artificial Intelligence

---

## 1. Introduction

The incidence of mental health problems in China has risen significantly in recent years, and with it comes an increase in public health issues such as the risk of suicide among individuals diagnosed with mental disorders [1, 2]. Recent research has found a positive correlation between increased social media engagement and a rise in mental health problems as social media use becomes more prevalent, particularly among the teenage population [3]. Excessive use of the internet and mobile phones was also identified as an important factor in self-harm [4]. These trends highlight the urgent need for a new way of analyzing and monitoring mental health issues, particularly through the extensive data available on social media.

Although social media can provide a large amount of potential data for mental health analysis, current methods of analyzing such data have not been very effective. Traditional methods such as questionnaires and keyword screening do not cope well with the large amounts of data generated by social media daily, which limits their application in identifying and tracking long-term mental health trends. Therefore, we need larger scale, more efficient, and more accurate methods to make full use of the potential information in social media data.

Recent advances in Natural Language Processing (NLP) provide new opportunities to better address these problems[5]. The development of pre-trained transformer models and large language models (LLMs) has significantly improved the ability of AI systems to understand and analyze human language [6]. These models are well suited for dealing with dynamic, unstructured, and spontaneous social media texts, making them ideal tools for real-time mental health analysis.

Based on the above premise, we built a new social media mental health analysis system that utilizes a large language model to analyze content posted by users of Weibo, one of the most popular social media platforms in China. By continuously collecting and analyzing text data in real time, the system can generate the latest analysis of mental health trends for different hot topics. It can statistically the distribution of mental health conditions over time and location. This approach has significant advantages over traditional methods, including more timely feedback, less manpower requirements, and greater monitoring scale, enabling better monitoring of mental health changes on social media platforms.

Specifically, the main contributions of this paper are as follows:

1. Propose a mental health monitoring system: We have developed a new system that uses a large language model to analyze social media text data from Weibo in real time to achieve dynamic and continuous monitoring of psychological trends.
2. Improve efficiency of mental health analysis: Our approach reduces the time lag of human observation by automating the data collection and analysis process and significantly improves the efficiency speed, and ease of psychological assessment by analyzing the details of mental health conditions with fine-tuned large language models.
3. Build a visual mental health data statistics platform: Our system's output database is an extremely flexible medium for statistical analysis, with a multitude of ways to interpret, visualize, and study to suit further research.

## 2. Related work

### 2.1. Social media as a predictor for mental health

The use of social media as a tool for predicting mental health issues has been an area of growing interest and research over the past decade. Early studies in this field explored the potential of digital footprints left on platforms like Twitter, Facebook, and Instagram to serve as indicators of mental health status. For example, one of the pioneering studies by De Choudhury et al. [7] demonstrated that social media data could be used to identify individuals with depressive symptoms by analyzing linguistic markers such as negative sentiment, rumination, and reduced social engagement. This research laid the foundation for further exploration into the feasibility of using social media as a predictive tool for mental health conditions.

Over the years, researchers have expanded on these initial findings, developing more sophisticated models and methodologies to capture a wide range of mental health indicators from social media data. Recent research has focused on using machine learning and natural language processing (NLP) techniques to distinguish different mental health problems such as anxiety, depression, and suicidal thoughts. For instance, a study by Reece and Danforth [8] demonstrated the ability to differentiate between healthy and depression-driven content by examining linguistic features and visual cues in Instagram posts, such as color tones, filters, and hashtags. Similarly, Coppersmith et al. [9] showed that Twitter data could be utilized to detect signals associated with post-traumatic stress disorder (PTSD), seasonal affective disorder (SAD), and bipolar disorder.

The current state of research has seen a shift from basic exploratory studies to the application of advanced artificial intelligence (AI) and deep learning techniques, which can process vast amounts of unstructured data to identify subtle patterns associated with mental health issues. For example, Schwartz et al. [10] utilized topic modeling and linguistic inquiry tools to identify psychological stress and mood disorders in Facebook posts. Other studies used deep learning models, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), to analyze social media data and detect mental health conditions with high accuracy [11].

The significance of using social media data for mental health prediction is multifaceted. First, social media platforms offer a wealth of real-time data that can provide insights into users' psychological states, allowing for timely interventions. Unlike traditional methods of mental health assessments, which often rely on self-reports or clinical interviews, social media data can capture continuous and spontaneous expressions of thoughts and emotions, which may offer a more accurate reflection of an individual's mental health status. Second, the accessibility and ubiquity of social media make it a valuable tool for public health monitoring at both the individual and population levels. For instance, large-scale analyses of social media content have been used to monitor mental health trends during major societal events, such as the COVID-19 pandemic, highlighting the impact of external stressors on public mental health [12].

Furthermore, some researchers have shown that analyzing the content of social media posts can be particularly effective in detecting specific mental health conditions [13]. In user-generated text on platforms like Facebook and Twitter, certain linguistic markers, such as frequent use of personal pronouns, negative emotion words, and words associated with cognitive processes, have been linked to psychological conditions like depression, anxiety, and loneliness. These markers can help identify individuals at risk for mental health, providing researchers with valuable data.

### 2.2. LLMs for mental health analysis

Psychological analysis has always been a complex part of Natural Language Processing (NLP), which is closely related to traditional NLP fields such as sentiment classification and text understanding. Traditional natural language processing methods (such as morpheme analysis) are difficult to deeply understand the psychological implications implied in the text, so the accuracy and efficiency of their psychological analysis are relatively low. With the advent of large language models (LLMs), the language understanding ability of computers has been greatly enhanced, which makes it possible to use computers to perform large-scale and detailed psychological analyses. These advances are mainly attributed to the emergence of transformer architectures.

Some well-known models such as BERT (Bidirectional Encoder Representation from Transformer) [14] and GPT (Generate Pre-trained Transformer) [15] leverage the transformer architecture to learn contextual relationships in text. Effectively capturing complex language patterns. Some subsequent models, such as RoBERTa [16] and DistilBERT [17], have further improved the performance and efficiency of these models in various NLP applications by adjusting the structure and details of LLMs.

The application of LLMs to mental health analysis has recently gained momentum. By fine-tuning these models on specific datasets, researchers have identified psychological states and emotional cues from various text data sources, including social media posts and clinical notes. For instance, BERT, when fine-tuned on mental health-specific corpora, has shown a high degree of accuracy in detecting psychological distress and mental health-related discourse on platforms like Reddit [18]. These examples demonstrate LLMs' potential for adapting to different domains and tasks, including mental health monitoring.

One of the primary techniques for optimizing LLMs in mental health analysis is fine-tuning. This process involves adapting a pre-trained model to a smaller, domain-specific dataset, thereby enhancing its ability to recognize linguistic markers and psychological cues relevant to mental health [19-20]. For example, models fine-tuned on datasets related to depression, anxiety, and PTSD have demonstrated improved performance in identifying symptoms associated with these conditions [21]. This adaptability highlights the effectiveness of LLMs in specialized tasks.

Recent advancements in LLMs have led to the development of more sophisticated models, such as GLM-4[22], developed by Zhipu AI. Unlike earlier models like BERT or GPT, GLM-4 demonstrates superior performance in a range of NLP tasks, including multi-lingual understanding, reasoning, and dialogue generation. Furthermore, GLM-4 supports advanced capabilities such as long-context reasoning, web browsing, and custom tool calls, making it especially suitable for complex linguistic analyses. Given these capabilities, GLM-4 holds particular promise for applications in mental health research, especially in analyzing text data from social media to understand psychological trends.

Although LLM has been proven to have great potential for analyzing mental health conditions, there are still many difficulties and challenges in the specific practice process. Examples include the need for a large amount of high-quality training data, privacy protection issues for users, and interpretability of analysis results [23]. Addressing these questions is critical to building a large-scale and efficient mental health monitoring system.

### 2.3. Summary

In summary, the existing research illustrates the value of social media texts in monitoring the mental health of individuals and populations. At the same time, the rapid development of large language models has greatly improved the efficiency and quality of mental health analysis from such textual data. After fine-tuning using psychological datasets, these models demonstrate strong psychological understanding, thus providing a critical foundation for the development of large-scale social media mental health monitoring systems.

## 3. Method

### 3.1. Overview

Our proposed system aims to be able to produce a frequently updated database containing mental health category breakdowns for comments posted on Weibo.

The data from the database can be interpreted, examined, and visualized to allow for fine tracking of mental health trends.

### 3.2. Modules

The system pipeline involves three core modules:

1. Data scraping module: Use the Python Scrapy module to crawl comments from Weibo based on popular topics on the day of crawling and package each comment with its context: post content, topic tags (hot search), and other metadata such as time information and location information.

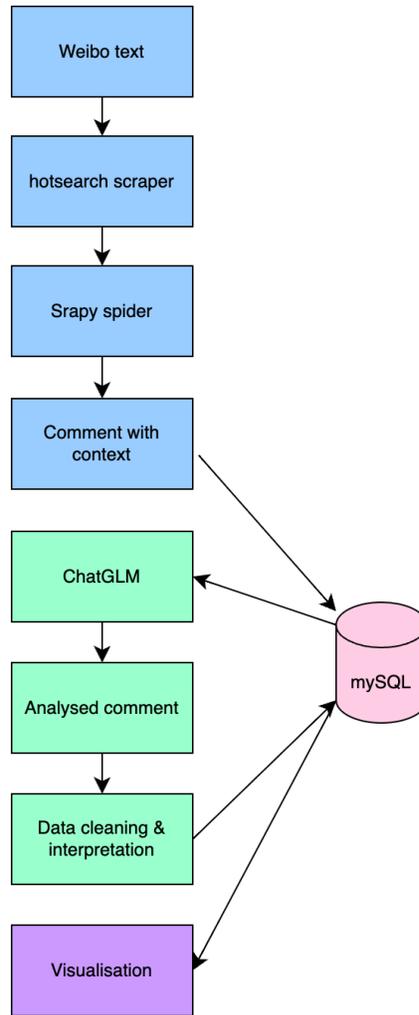
2. Data Analysis module: Using the large language model ChatGLM, pre-trained on various psychological corpora, and prompted to analyze the text according to 11 standards related to mental health, read the packaged comments produced by the data crawling module and output formatted data.

3. Scheduler: A Python script and a MySQL database are used to coordinate the data flow between modules.

Aside from the three more rigid modules, there is also an optional, flexible visualization module, which can be realized in different forms based on the needs of a particular analysis. For demonstration, we have provided an interpretation of the core data, visualized with Tencent's BI online visualization service.

### 3.3. Pipeline

Our system implementation involves two cloud computers from Tencent cloud services running Ubuntu. A MySQL database is used to transfer data between the two machines, as well as the visualization software.



**Figure 1.** Process diagram of the pipeline

## 4. Data scraping module

This module's goal is to be able to periodically start up and scrape user-generated text data, package it with context, and store it in a remote database for the analysis module to read from.

The module runs on one of the two Tencent cloud computers, running on the 64-bit version of Ubuntu Server 24.04 LTS, with a 16-core unspecified CPU, 32 gigabytes of RAM, and no GPU.

The data scraping module starts up according to the scheduler, which is a Python script that runs continually in the background, and also serves as the main body of the execution of the module, linking its separate components together.

### 4.1. Weibo text data overview

Weibo has three levels of context for its text comments:

1. The tags, distinguished by hashtags # in front of and after their name, work as broad identifiers of the topics that will be discussed under posts that contain a particular tag. Tags are ranked by their popularity—how many posts have it in their text body—and the top few as categorized as hot- searches.

2. The posts, which are not dissimilar from posts made on any other social media, contain a tag, text body, and sometimes an image or video. Images or videos in the posts will not be used in the analysis, but because visual media in posts are mostly supplementary, it does not affect the accuracy of the context surrounding comments, and thus the analysis of them.

3. the comments, the subject of analysis. They can contain only text and emotes.

## 4.2. Hotsearch scraper

The first part of the data scraping module is used for acquiring the tag names of Weibo's top hot searches. Using these hot searches as a basis ensures that the module can capture data most relevant to the general population.

It takes advantage of the fact that Weibo stores its raw information publicly, accessible from a URL. For hotsearches, they are included in a URL in the JSON file format. This allows for an optimization that skips the need to navigate through the front-end website to acquire the search data.

Instead, the scraper simply uses Python's requests module to get the raw data, convert it into a Python dictionary, and extract the data by traversing through the list of hot search entries and looking for the key that corresponds with the tag name.

The data is then formatted and written to a file read by the next component, the scrapy spider.

## 4.3. Scrapy

Scrapy is a framework for creating and managing web spiders, interfaced through Python[24]. It is used to gather the bulk of the content of the collection module. Our system uses Scrapy version 2.11.2.

Scrapy projects generally have five core components:

1. The items script, which defines what kind of object is returned by the spider.
2. The middlewares script, is used to modify items that are passed between the server and the spiders.
3. The pipelines script, defines how items re-turned by spiders should be treated and processed.
4. The settings script, which outlines parameters that affect the functionality of the other components.
5. The spiders themselves, which interact with the web and gather data.

Numerous Scrapy projects have been made for various uses and uploaded on GitHub, including ones for scraping Weibo.

To scrape comments from posts based on hot searches, we have employed the Scrapy project database/web search, which we have modified to fit our purpose.

## 4.4. Basic parameters

A legitimate user cookie was set to prevent Weibo from flagging and potentially blocking the spider from scraping data.

Web search filters data by date, which was set to be the date whenever the spider is run to ensure that data scraped on a particular day is up-to-date.

Weibo-search has text filter functionality, where it only scrapes comments from posts containing text specified in its settings script. We leverage this functionality to filter comments by hotsearch.

The raw hotsearch names read from the hotsearch scraper are formatted with hashtags surrounding them to match how they would appear in a post body. This way, by passing in the file with formatted names written by the hot search scraper into the keyword list parameter, the Scrapy spider effectively only scrapes comments if it is from a post that contains a popular topic.

The final data collected is a series of comments from posts that have popular tags. The spider can scrape 19 components of data for each comment entry: id, bid, user ID, screen name, text, article URL, mentioned users, topics, repost count, comment count, attitude count, created at, source pics, video URL, user IP (location), user authentication.

They are then processed into a MySQL database by the Pipeline script, ready for the analysis module to read from.

## 4.5. Scheduler

Using Python's schedule module, the scheduler runs continually in the background and periodically starts the scraping procedure.

The hot search scraping script is run to update the keyword tag filters, after which a Scrapy instance is started as subprocesses by Python's subprocess module, and killed off after a certain time duration.

We have chosen the times 10:00 AM and 10:00 PM for our scheduled times as ballparks for when users are likely to be more active, though we acknowledge can be changed and improved after some statistical analysis.

A duration of 60 minutes was set to be the length of scraping time due to the limitations of the data analysis module, which cannot process data as fast as it can be collected.

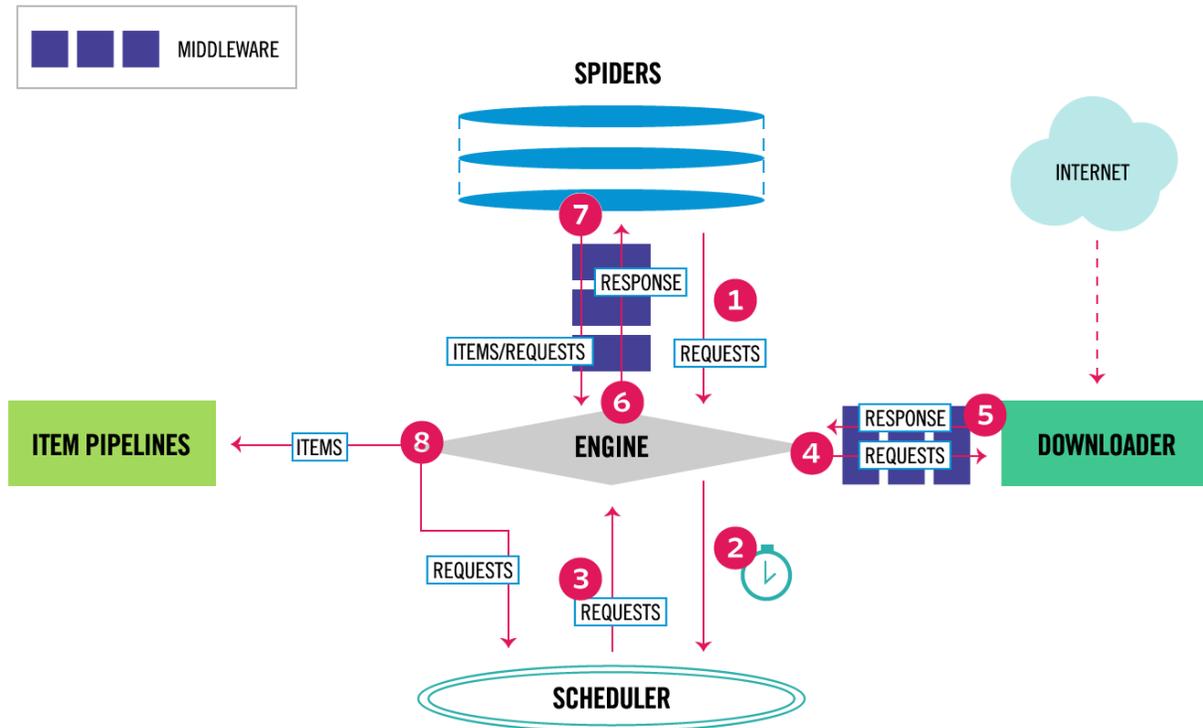
## 5. Data analysis module

The analysis module includes the large language model as its core, feeding comment data from the database filled by the collection module and outputting analysis data into another section of the database for visualization and statistical analyses.

This module runs on the other of the two Ten-cent cloud computers, running on the 64-bit version of Ubuntu Server 22.04 LTS 64, with an 8-core unspecial- fied CPU, 40 gigabytes of RAM, and an NVIDIA V100 GPU to train and run the model.

### 5.1. Model

For the task of interpreting and analyzing Chinese text, we have chosen to use the open-source bilingual chat model project ChatGLM, which includes models that have shown proficiency in text information extraction, formatting, and inference in Chinese [22]. Our system utilizes the GLM-4-9B model.



**Figure 2.** Diagram of the Scrapy architecture

### 5.2. Pre-Training

ChatGLM is a generalized chat model that can process a large variety of text topics. However, since our use of the model does not encompass such a large diversity, it can be tuned to be more especially adept at psychological analysis.

For this purpose, we have elected to use the popular open-source tuning tool LLama-factory [25].

Our pretraining dataset includes various psychology texts such as OpenStax's Psychology 2e textbook, which were originally PDF documents that were first converted into text with Python's PyPDF2 module, then parsed into separate sentences with the NLTK module, and finally underwent more manual cleaning up, removing titles, credits and padding text at the beginning and end of texts. In total, in both English and Chinese, our dataset includes over 20,000 lines of psychology text.

Pretraining for 10 epochs with a learning rate of  $5e-5$ , batch size of 2, gradient accumulation of 8, a cosine LR scheduler, and using the AdamW optimizer, our GLM-4-9b instance was able to achieve a loss of 0.9 from its starting 2.1.

### 5.3. Structure

Due to ChatGLM's nature as a chatbot that is generally used to generate naturalistic language conversationally, the runtime structure template of receiving prompts and generating text for the model is edited for it to be suited to receiving and returning codified data.

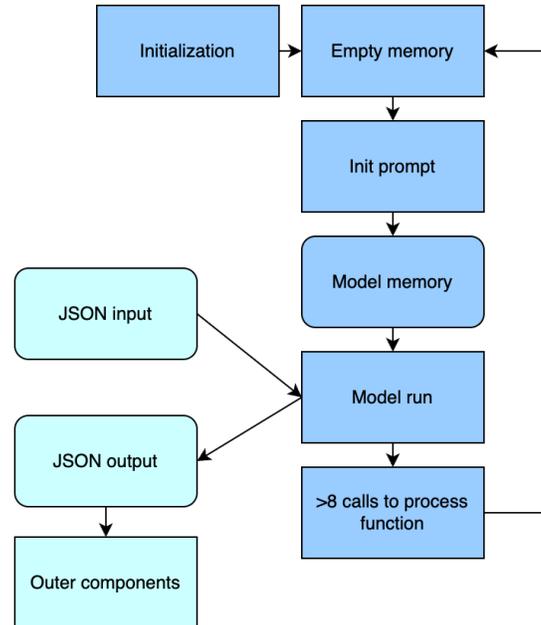
The model receives input as text in JSON format containing 2 keys, post title, and comment text, and outputs the analysis as a Python dictionary containing a key and value for each analysis criterion.

A specially designed initialization prompt is fed into the model on startup that instructs it to expect this JSON input, evaluate the comment text based on the context given by the post title, and output the analysis in a specified JSON format. Moreover, to

prevent hallucinations and the model straying from the intended output, the model memory cache is cleared and the initialization prompt is re-inputted into the model every 8 comments.

The JSON output is then converted into a Python dictionary and returned from the function. To avoid losing analysis data due to the model's mistakes, any conversion error from the JSON parser will make the model re-generate the analysis text.

The model call is then modularized as a Python class, allowing for it to be used outside of its file, making it easier and more concise to use.



**Figure 3.** Diagram of the model usage process

#### 5.4. Analysis criteria

The prompts detail 11 categories of what the model was told to analyze, the first being levels of negativity and the other 10 being different mental health issues, their contents in English is shown in Table 1.

**Table 1.** Classes of mental health issues

English
<b>Organic, including symptomatic mental disorders</b>
<b>Mental and behavioral disorders caused by the use of psychoactive substances</b>
<b>Schizophrenia and paranoia</b>
<b>Anxiety disorders, stress-related physical disorders</b>
<b>Behavioral syndromes associated with physical disorders and physiological factors</b>
<b>Adult personality and behavior disorders</b>
<b>Mental retardation</b>
<b>Mental development disorder</b>
<b>Behavioral and emotional disorders in childhood</b>

These categories are taken from a summary of the classification system of the ICD-10 (International Classification of Diseases), translated into Chinese and included in the prompt.

From the prompt, the model is instructed to give each category a rating from 0 to 5: The sentimental rating, is a rating of the comment from least negative to most negative, and for the psychological categories, it is a rating of very low to very high probability that the author of a comment displays symptoms of that category of mental health issues.

## 5.5. Database

We use a Python wrapper script to communicate with the MySQL database. The script includes a GLM class instance from the model file above that is used for the comment analysis.

Data is fetched from the database in batches. In each batch, some information scraped by the collection module is retrieved, a part of which is processed by the model and the others to provide context for the evaluation provided by the model.

The information needed can vary depending on the needs of a particular statistical analysis; Here we have chosen a baseline out of the 19 pieces of data retrieved for comment for information that can be included as context:

1. post id
2. user IP (user location)
3. repost count
4. comment count on post
5. like count
6. comment time

This information is stored in a temporary table that will hold the final result data. Then, the post title and comment from the row in each batch are formatted into JSON text and fed into the model.

The Python dictionary output from the model is then split into 10 entries, one for each of the psychological categories, each entry containing the name of the category, and the number rating given by the model. This is done to allow for easier analysis, grouping, and labeling of the resultant data. Each of the 10 entries is then added to the temporary table containing the original post content, comment text, and sentimental rating, as well as the context data from the data retrieved from the database. The final resultant data is then stored in another section of the database, ready to be linked to external software for viewing and statistical analysis.

## 6. Data interpretation

Here we demonstrate the analysis capabilities of the system by offering statistical findings and visualizations of the data.

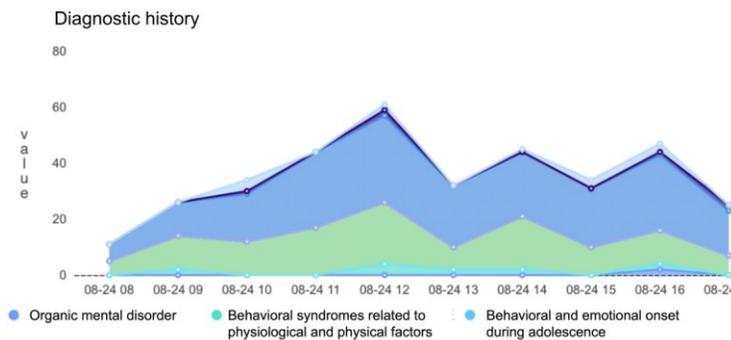
### 6.1. Software

Our visualization and interpretation software of choice was Tencent's BI service. The online service supports the connection of MySQL databases, to let the software read the data to use for analysis and visualization. It also uses a GUI editor, which allows for quick demonstrations of data, our goal for this section.

### 6.2. Data visualisation

Our system was run for 2 days to gather data. Even over a relatively short time interval, the data produced was able to be analyzed and visualized in various ways, serving as a strong proof of concept.

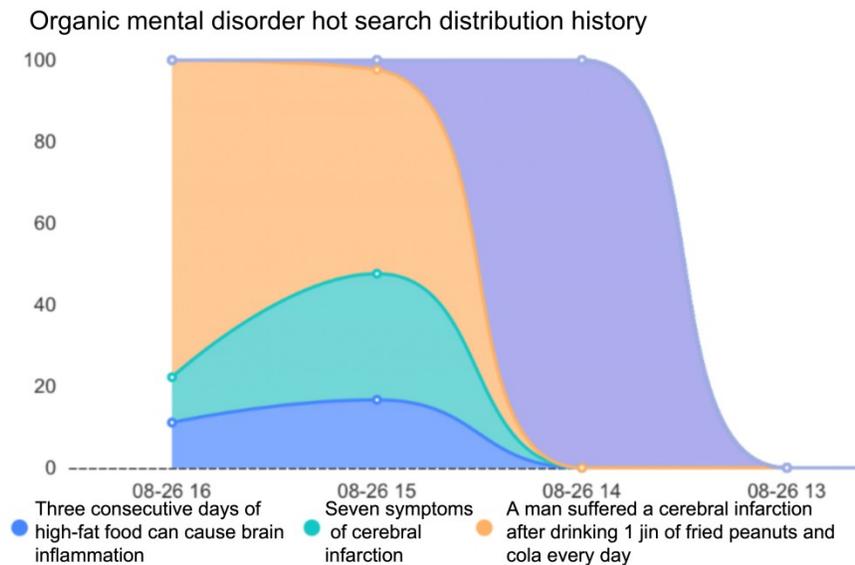
With the category ratings separated by category rather than comment, it allows for various techniques involving segmenting graphs, such as segmenting ratings by category, shown in Figure 4. This data structure also allows for separate analysis of each category, such as visualizing the hot search topics that contribute most to a certain category like Figure 5.



**Figure 4.** Graph of an overview of all comment ratings over time; the y-axis represents the total rating, segmented by each category.

For example, the segment in yellow seen in Figure 4, that translates to "A boy who ate 1 jin(0.1 KG) of fried peanuts and drank cola every day had a cerebral infarction", was the source of the most comments with symptoms of organic mental disorders.

As the IP address is also included in every data entry, map visualizations can also be made, such as viewing every province in China by average comment sentiment, such as in Figure 5, or by overall like count in



**Figure 5.** Graph of the total ratings of organic and symptomatic mental disorders over time, the y-axis segmented by Hotsearch topic.

All in all, these visual representations of the data demonstrate the flexibility of the dataset, and there are many more that can be made depending on the context and needs of any study that uses it.

## 7. Conclusion

We propose a system for analyzing comments from Weibo psychologically. Through Python's scraping technologies and a tuned ChatGLM model, the system produces an extensive database that can be analyzed in a multitude of ways, providing insights into the user base, tracking its various psychological aspects over many variables, and revealing trends that could spur further research.

## Acknowledgments

The research focuses on the psychological analysis of post comments from Weibo using a large language model. By scraping text periodically, and using a fine-tuned large language model, the system proposed by the research can improve on the current state of psychological research of online texts.

This research was done under the advisory of advisor Jack Zhou. The advisory was determined by the school and was not paid. I am the only student participating in this project. The research topic was conceived with a discussion with the advisor based on personal interests in the field of computer science and how it can be used in conjunction with other fields to provide a positive impact.

The data collected for the research is part of the system proposed; It is done through web scraping of publicly available text from Weibo. The data scraped includes post text, comment text, various ID labels of the users and comments, post statistics such as like and repost count, and metadata revolving around the comments such as user IP and time. Then, as detailed in the paper, the data is then processed by a large language model. There were a couple of difficulties the system faced during its stages.

For example, the output data was originally designed to be in full JSON format, but after meeting compatibility problems multiple times, everything outputted was converted into an online MySQL database. There were instances where a GitHub repository project required for the system would refuse to run, such as the large language model and its training software, due to version incompatibility or network errors that were resolved in a trial-and-error fashion by reading the error messages and looking up problems online.

Code debugging was a large part of the development of the system, as with any project that involves a large amount of code. Every file in the final system had gone through countless debugging sessions. Moreover, due to the system running on two cloud

computers, incorrect behavior was relatively difficult to track down, requiring looking at the output values of each module to pinpoint a problem, such as when an end module output would turn up empty when there should be data.

All of this was done under occasional guidance under the advisor.

## References

- [1] Lei, X., Liu, C., & Jiang, H. (2021). Mental health of college students and associated factors in Hubei of China. *PLoS One*, 16(7), e0254183.
- [2] Yu, R., Chen, Y., Li, L., Chen, J., Guo, Y., Bian, Z., Lv, J., Yu, C., Xie, X., Huang, D., Chen, Z., & Fazel, S. (2021). Factors associated with suicide risk among Chinese adults: A prospective cohort study of 0.5 million individuals. *PLoS medicine*, 18(3), e1003545. <https://doi.org/10.1371/journal.pmed.1003545>
- [3] Abi-Jaoude, E., Naylor, K. T., & Pignatiello, A. (2020). Smartphones, social media use, and youth mental health. *CMAJ*, 192(6), E136–E141.
- [4] Wang, L., Liu, X., Liu, Z.-Z., & Jia, C.-X. (2020). Digital media use and subsequent self-harm during a 1-year follow-up of Chinese adolescents. *Journal of Affective Disorders*, 277, 279–286.
- [5] Torfi, A., Shirvani, R. A., Keneshloo, Y., Tavaf, N., & Fox, E. A. (2020). Natural language processing advancements by deep learning: A survey. *arXiv preprint arXiv:2003.01200*.
- [6] Zhou, M., Duan, N., Liu, S., & Shum, H.-Y. (2020). Progress in neural NLP: Modeling, learning, and reasoning. *Engineering*, 6(3), 275–290.
- [7] De Choudhury, M., Gamon, M., Counts, S., & Horvitz, E. (2013). Predicting depression via social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, 7, 128–137.
- [8] Reece, A. G., & Danforth, C. M. (2017). Instagram photos reveal predictive markers of depression. *EPJ Data Science*, 6(1), 15.
- [9] Coppersmith, G., Dredze, M., & Harman, C. (2014). Quantifying mental health signals in Twitter. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, 51–60.
- [10] Schwartz, H. A., Sap, M., Kern, M. L., Eichstaedt, J. C., Kapelner, A., Agrawal, M., Blanco, E., Dziurzynski, L., Park, G., Stillwell, D., Kosinski, M., Seligman, M. E., & Ungar, L. H. (2016). Predicting individual well-being through the language of social media. *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, 21, 516–527.
- [11] Orabi, A. H., Buddhitha, P., Orabi, M. H., & Inkpen, D. (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*, 88–97.
- [12] Chandrasekaran, R., Mehta, V., Valkunde, T., & Moustakas, E. (2020). Topics, trends, and sentiments of tweets about the COVID-19 pandemic: Temporal infoveillance study. *Journal of Medical Internet Research*, 22(10), e22624.
- [13] De Choudhury, M., Counts, S., & Horvitz, E. (2013). Social media as a measurement tool of depression in populations. In *Proceedings of the 5th Annual ACM Web Science Conference*, 47–56.
- [14] Devlin, J. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- [15] Yenduri, G., Ramalingam, M., Selvi, G. C., Supriya, Y., Srivastava, G., Maddikunta, P. K. R., Raj, D. G., Jhaveri, R. H., Prabadevi, B., Wang, W., Vasilakos, A. V., & Gadekallu, T. R. (2023). Generative pre-trained transformer: A comprehensive review on enabling technologies, potential applications, emerging challenges, and future directions.
- [16] Liu, Y. (2019). RoBERTa: A robustly optimized BERT pre-training approach. *arXiv preprint arXiv:1907.11692*.
- [17] Sanh, V. (2019). DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- [18] Inamdar, S., Chapekar, R., Gite, S., & Pradhan, B. (2023). Machine learning driven mental stress detection on Reddit posts using natural language processing. *Human-Centric Intelligent Systems*, 3(2), 80–91.
- [19] Radford, A. (2018). Improving language understanding by generative pre-training.
- [20] Du, J., Grave, E., Gunel, B., Chaudhary, V., Celebi, O., Auli, M., Stoyanov, V., & Conneau, A. (2020). Self-training improves pre-training for natural language understanding. *arXiv preprint arXiv:2010.02194*.
- [21] Xu, X., Yao, B., Dong, Y., Gabriel, S., Yu, H., Hendler, J., Ghassemi, M., Dey, A. K., & Wang, D. (2024). Mental-LLM: Leveraging large language models for mental health prediction via online text data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(1), 1–32.
- [22] Team GLM, Zeng, A., Xu, B., Wang, B., Zhang, C., Yin, D., Feng, G., Zhao, H., Lai, H., Yu, H., Wang, H., Sun, J., Zhang, J., Cheng, J., Gui, J., Tang, J., Zhang, J., Cheng, J., Gui, J., Wu, L., Zhong, L., Liu, M., Huang, M., Zhang, P., Zheng, Q., Lu, R., Duan, S., Zhang, S., Cao, S., Yang, S., Tam, W. L., Zhao, W., Liu, X., Xia, X., Zhang, X., Liu, X., Yang, X., Song, X., Zhang, X., An, Y., Xu, Y., Niu, Y., Yang, Y., Li, Y., Bai, Y., Dong, Y., Qi, Z., Wang, Z., Yang, Z., Du, Z., Hou, Z., Wang, Z. (2024). ChatGLM: A family of large language models from GLM-130B to GLM-4 all tools.
- [23] Ke, L., Tong, S., Chen, P., & Peng, K. (2024). Exploring the frontiers of LLMs in psychological applications: A comprehensive review. *arXiv preprint arXiv:2401.01519*.
- [24] Kouzis-Loukas, D. (2016). Learning Scrapy. Packt Publishing Ltd.
- [25] Zheng, Y., Zhang, R., Zhang, J., Ye, Y., Luo, Z., Feng, Z., Ma, Y. (2024). LLamaFactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.