

Revealing COVID-19’s Social Dynamics: Diachronic Semantic Analysis of Vaccine and Symptom Discourse on Twitter

Zeqiang Wang¹ Jiageng Wu² Yuqi Wang³ Wei Wang³
Jie Yang⁴ Jon Johnson⁵ Nishanth Sastry¹ Suparna De^{1*}

¹ University of Surrey ² Zhejiang University ³ Xi’an Jiaotong University ⁴ Liverpool University

⁴ Harvard University ⁵ University College London

{zeqiang.wang, s.de, n.sastry}@surrey.ac.uk, jiagengwu@zju.edu.cn, jieynlp@gmail.com,
yuqi.wang17@student.xjtlu.edu.cn, wei.wang03@xjtlu.edu.cn, jon.johnson@ucl.ac.uk

Abstract

Social media is recognized as an important source for deriving insights into public opinion dynamics and social impacts due to the vast textual data generated daily and the ‘unconstrained’ behavior of people interacting on these platforms. However, such analyses prove challenging due to the semantic shift phenomenon, where word meanings evolve over time. This paper proposes an unsupervised dynamic word embedding method to capture longitudinal semantic shifts in social media data without predefined anchor words. The method leverages word co-occurrence statistics and dynamic updating to adapt embeddings over time, addressing the challenges of data sparseness, imbalanced distributions, and synergistic semantic effects. Evaluated on a large COVID-19 Twitter dataset, the method reveals semantic evolution patterns of vaccine- and symptom-related entities across different pandemic stages, and their potential correlations with real-world statistics. Our key contributions include the dynamic embedding technique, empirical analysis of COVID-19 semantic shifts, and discussions on enhancing semantic shift modeling for computational social science research. This study enables capturing longitudinal semantic dynamics on social media to understand public discourse and collective phenomena.

1 Introduction

Social media has become an essential platform for information dissemination and public opinion expression, with over 4.5 billion global users generating vast amounts of textual data daily (Li et al., 2022; Wu et al., 2023a; De et al., 2023). Large-scale, wide-coverage, and real-time social media data provide a new source for computational social science and cultural analysis research, enabling insights into public opinion dynamics, so-

cial event impacts, and collective behavior patterns through large-scale text analysis (Hua et al., 2024). However, social media text analysis faces a unique challenge: the semantic shift phenomenon. Word semantics can change to a certain extent over time due to evolving social events, user interests, cultural trends, and other factors (Harrigan and Dredze, 2022; Wang et al., 2023a). For instance, during the COVID-19 pandemic, “Moderna” initially referred to a biotechnology company but later became associated with its COVID-19 vaccine. Failing to account for such dynamic semantic evolution can affect the accuracy and interpretability of text analysis, especially in long-term longitudinal analyses spanning months or even years.

Existing mainstream semantic shift analysis methods usually require pre-defining a set of anchor words as a reference (Ishihara et al., 2022; Montanelli and Periti, 2023). However, these methods rely excessively on human experts’ prior experiences that are usually highly subjective (Sharifian-Attar et al., 2022), and cannot comprehensively detect semantic shifts in longitudinal corpora. Moreover, social media text presents several challenges for existing semantic shift analysis methods:

1. **Data Sparseness:** Social media text is highly sparse, with only a few hundred core words and categories exhibiting dynamic changes, making it difficult for statistical language models to learn robust semantic representations.
2. **Imbalanced Distribution:** Core words have a highly imbalanced distribution. This long-tail distribution poses challenges for semantic shift modeling, requiring targeted handling of low-frequency and emerging words.
3. **Synergistic Effects:** Topic evolution on social media often manifests as the joint shift of multiple related words. Focusing solely on

*Corresponding Author.

isolated word changes may obscure overall topic-level trends.

To address these challenges, this paper proposes a novel framework aimed at better capturing longitudinal semantic evolution patterns in the context of social media and providing powerful semantic analysis tools for computational social science research. The proposed end-to-end unsupervised dynamic word embedding method is based on word co-occurrence statistics and combines a dynamic update strategy to adaptively capture the trajectory of word semantic evolution over time, without the need for manually defined anchor words. To verify the effectiveness of the method in practical applications, we applied it to a large-scale COVID-19 Twitter dataset spanning more than two years, from February 2020 to April 2022. We focused on the semantic shift phenomenon of vaccine- and symptom-related entity words and observed their semantic shift patterns at different stages of the pandemic, as well as their potential correlations with real epidemic trends.

The main contributions are summarized as follows: 1) We propose a dynamic word embedding method that can adaptively detect semantic shift phenomena in longitudinal corpora without manual annotation; 2) We integrate it with Named Entity Recognition (NER) and Large Language Model (LLM)-based entity normalization to systematically investigate the semantic shift in extensive and longitudinal social media corpora during the COVID-19 pandemic; 3) Our experimental results reveal the semantic evolution patterns of vaccine- and symptom-related topics during the COVID-19 pandemic, offering a new perspective on the study of epidemic evolution and its social impact. This provides a novel paradigm for leveraging computational linguistics in social and medical research.

2 Related Work

Semantic shift analysis is a crucial research problem in the fields of computational linguistics and natural language processing. Over time, the semantics of words evolve, reflecting changes in language usage and socio-cultural dynamics. Researchers have recently proposed various data-driven methods to automatically detect semantic shift phenomena at the word level.

Early studies on semantic shift primarily relied on word frequency statistics, analyzing the temporal variation of word frequencies in books (Michel

et al., 2011) and in web search queries (Choi and Varian, 2012). However, it has been noted that word frequency changes are not always directly correlated with semantic shifts (Kulkarni et al., 2015), with frequency information alone being insufficient to accurately capture subtle semantic variations.

Distributional semantics has been proven to be an effective tool for semantic shift analysis (Turney and Pantel, 2010; Baroni et al., 2014). By leveraging word co-occurrence statistics, distributional representations of words can be learned, directly encoding their semantics as vectors. Jurgens and Stevens (2009) pioneered quantitative analysis of semantic shifts with word vector updates to dynamically model semantic evolution, with subsequent studies employing various word embedding algorithms, such as Latent Semantic Analysis by Sagi et al. (2011) and Local Mutual Information by Gujordava and Baroni (2011). Kim et al. (2014) introduced predictive word embedding models (e.g., word2vec) to the semantic shift detection task, followed by works comparing performance of different types of word embedding models (Hamilton et al., 2016b) which found that predictive models like skip-gram with negative sampling (SGNS) outperform explicit counting methods.

Other studies (Mitra et al., 2015) have looked at characterizing semantic shift types, such as the emergence of novel senses, splitting and merging of word senses, which falls under fine-grained semantic shift analysis. Hamilton et al. (2016a) explored the distinction between cultural and linguistic shifts as the two modes of semantic evolution, with the former closely related to socio-cultural changes and manifesting as global displacement in the word vector space, and the latter reflecting the internal evolution of language and characterized by changes in local neighborhoods. Semantic shift analysis through the topic evolution lens is explored through dynamic topic models to track the temporal variation of topic distributions (Wang and McCallum, 2006) or temporal evolution of Google search topics (Wijaya and Yeniterzi, 2011).

The emergence of contextualized word embeddings capable of capturing diversity and subtle changes in word semantics in pre-trained language models such as ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019) has led to the use of BERT for semantic shift detection, with a novel representation splitting method to improve detection accuracy (Hu et al., 2019). Subsequent studies, such as those by Giulianelli et al. (2020) have con-

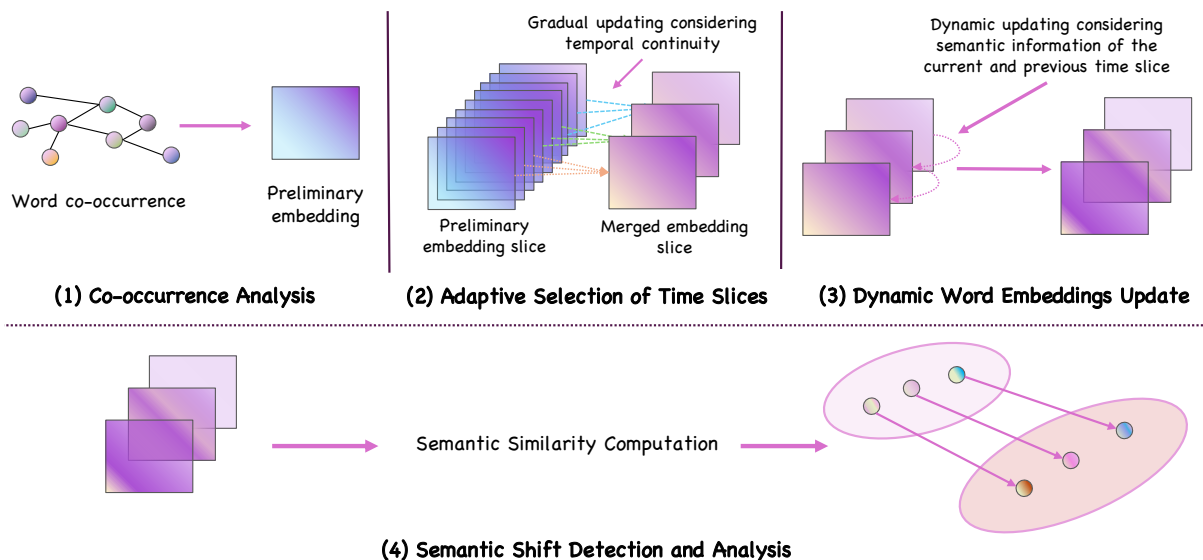


Figure 1: Overall framework of the proposed unsupervised dynamic word embedding method. (1) Co-occurrence Analysis: Word co-occurrence matrices from a diachronic corpus are computed and normalized; (2) Adaptive Selection of Time Slices: Adjacent time slices with high word co-occurrence similarity are merged adaptively; (3) Dynamic Word Embeddings Update: Word embeddings are dynamically updated based on the current and previous time slices. (4) Semantic Shift Detection and Analysis: Word semantic shifts are detected by embedding similarity and tracking the associations of word pairs that change over time.

firmed their superior performance over static ones with systematic evaluations of different types of contextualized word embeddings on this task.

Moreover, some studies focused on the interpretability of semantic shift detection. Rosin et al. (2022) proposed an interpretable semantic shift model based on attention mechanisms, explaining the reasons for word semantic evolution by visualizing attention weights. You (2021) employed causal inference methods to investigate the socio-cultural driving factors behind lexical semantic shifts. These works contribute to a deeper understanding of semantic shift phenomena from both computational and humanistic perspectives. In comparison to the above-mentioned methods, suppressing high-frequency core words and dynamically updating weights ensure that we can detect stable semantic shifts in embeddings.

Regarding the analysis of COVID-19 discussions on social media, several studies have explored this topic from different perspectives. Choudrie et al. (2021) investigated the role of machine learning in identifying misinformation during the pandemic, and explored how older adults interacted with online information, often finding it difficult to verify and preferring traditional media sources. Wicke and Bolognesi (2020) focused on the metaphorical language used in social media, particularly war-related terms, showing how certain

metaphors shaped public conversations around the virus and its effects. Gencoglu and Gruber (2020) analyzed the relationship between pandemic trends, such as infection rates, and shifts in public sentiment on Twitter, using causal modeling to distinguish between events that influenced public attention and those that were merely correlated. These studies collectively offer insights into the dynamic nature of COVID-19 discourse online, shaped by both statistical methods and human perception.

3 Methodology

This section introduces the proposed unsupervised dynamic word embedding method for detecting and analyzing semantic shifts of words in detail. The overall framework is shown in Figure 1. Our method leverages word co-occurrence statistics and adopts adaptive selection and time slice dynamic updating strategies to analyze and capture the evolution of word semantics over time.

3.1 Co-occurrence Analysis

Let $\mathcal{D} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^T\}$ denote the diachronic text corpus, where \mathcal{D}^t represents the text collection of the t -th time slice. We collect and pre-process large-scale text data from a number of time periods, including news articles, books, social media posts, and more, to ensure the data covers a broad temporal span and is representative of each period. We

compute the word co-occurrence frequency matrix $X^t \in \mathbb{R}^{|V| \times |V|}$ on each time slice \mathcal{D}^t , where $|V|$ represents the size of the vocabulary. X_{ij}^t represents the co-occurrence frequency of words w_i and w_j in the t -th time slice, reflecting the contextual information of words in different time periods.

To address the data sparseness, we apply a smoothing method to process co-occurrence matrices. Specifically, add-one smoothing is employed to add a small constant α to all elements of the matrix:

$$\tilde{X}_{ij}^t = X_{ij}^t + \alpha \quad (1)$$

where \tilde{X}^t denotes the smoothed co-occurrence matrix. This smoothing technique helps mitigate the impact of sparse data and improves the robustness of subsequent computation. Furthermore, to handle the unbalanced distribution of core words, we apply frequency-based normalization to the co-occurrence matrices. Specifically, we use logarithmic scaling to balance the importance of core words based on their frequency:

$$\hat{X}_{ij}^t = \log(1 + \tilde{X}_{ij}^t) \quad (2)$$

where \hat{X}^t denotes the normalized co-occurrence matrix. This normalization technique helps reduce the impact of extremely frequent words and gives more weight to less frequent but potentially informative words. We pre-train word vectors using traditional static word embedding methods (such as word2vec (Mikolov et al., 2013) or GloVe (Pennington et al., 2014)) to obtain preliminary embedding for each time slice t , i.e. $W^t \in \mathbb{R}^{|V| \times d}$, where d is the embedding dimension.

3.2 Adaptive Selection of Time Slices

We adaptively adjust the span of time slices based on the temporal granularity of word co-occurrence statistics. First, we compute a similarity measure between word co-occurrence matrices of adjacent time slices. Then, by setting a similarity threshold τ , we merge adjacent time slices with a similarity score higher than the threshold, forming slices with larger time spans. This ensures the continuity of semantic evolution while reducing the number of time slices and improving computational efficiency.

We use cosine similarity between the normalized word co-occurrence matrices \hat{X}^t and \hat{X}^{t+1} of adjacent time slices t and $t + 1$, which is given by:

$$\text{sim}(\hat{X}^t, \hat{X}^{t+1}) = \frac{\text{vec}(\hat{X}^t) \cdot \text{vec}(\hat{X}^{t+1})}{|\text{vec}(\hat{X}^t)|_2 |\text{vec}(\hat{X}^{t+1})|_2} \quad (3)$$

where $\text{vec}(\cdot)$ denotes flattening the matrix into a vector. If the semantic similarity is greater than a predefined threshold τ , the two time slices are merged into one. The corresponding word embeddings are then re-computed based on the merged corpus \mathcal{D}^t and \mathcal{D}^{t+1} . These steps are repeated until there are no similar co-occurrence matrices for any two consecutive time slices.

3.3 Dynamic Word Embeddings Update

We adopt a progressively dynamic update strategy. Given the smoothed word co-occurrence matrix \hat{X}^t of the current time slice t and the word embedding matrix W^{t-1} learned from the previous time slice, we update the word embeddings by minimizing the following objective:

$$W^t = \arg \min_W \left| \hat{X}^t - WW^T \right|_F^2 + \lambda \left| W - W^{t-1} \right|_F^2 \quad (4)$$

where $|\cdot|_F$ denotes the Frobenius norm of a matrix, and λ is a weight coefficient balancing the reconstruction error of the current time slice and the similarity with the embeddings from the previous time slice. This enables word embeddings to adapt to current semantic information while retaining previously learned semantics.

We introduce a forgetting factor $\gamma \in (0, 1]$ to control the speed at which word embeddings forget outdated semantics. Before each update, we decay the word embeddings from the previous time slice:

$$\tilde{W}^{t-1} = \gamma W^{t-1} \quad (5)$$

Then, we substitute the decayed word embeddings \tilde{W}^{t-1} into the optimization objective for updating. A smaller γ value promotes faster forgetting of outdated semantics, while a larger γ value helps maintain the continuity of semantics. To maintain the global semantic stability of word embeddings, we introduce a regularization term in the dynamic update optimization objective:

$$W^t = \arg \min_W \left| \hat{X}^t - WW^T \right|_F^2 + \lambda_1 \left| W - \tilde{W}^{t-1} \right|_F^2 + \lambda_2 \left| WW^T - \bar{X} \right|_F^2 \quad (6)$$

where $\bar{X} = \frac{1}{T} \sum_{t=1}^T \hat{X}^t$, represents the average normalized word co-occurrence matrix across all time slices, and λ_2 controls the weight of global stability. By minimizing the difference between

the inner product of word embeddings WW^T and the average word co-occurrence matrix \bar{X} , this regularization term encourages word embeddings to maintain global semantic consistency across different time and suppresses excessive semantic shift.

3.4 Semantic Shift Detection and Analysis

To analyze semantic shift under complex interactions and capture synergies between words, we introduce higher-order co-occurrence statistics. Specifically, we calculate the cosine similarity between word pairs across different time slices.

By tracking changes in cosine similarity of word pairs across different time slices, we can uncover dynamic patterns of semantic association evolution, such as strengthening, weakening, or shifting of associations. For instance, if the cosine similarity between “*COVID-19*” and “*vaccine*” gradually increases over time, it suggests that their semantic association becomes stronger with the progress of vaccine development and deployment.

4 COVID-19 Case Study

4.1 Dataset construction

We leveraged large-scale Twitter data to support an in-depth analysis of the characteristics and evolution patterns of COVID-19 symptoms. We selected 471,553,966 non-retweet English tweets posted between February 1, 2020, and April 30, 2022, as the raw dataset. This time span covered multiple key stages of the pandemic, including the initial outbreak, global spread, vaccine development, and rollout. The substantial data volume and extended time range provide a comprehensive and longitudinal perspective on the pandemic’s full picture.

To accurately identify self-reported symptoms mentioned in the massive tweet corpus, we designed a rigorous data preprocessing pipeline consisting of three steps:

1. **Strict COVID-19 relevance filtering:** We matched each tweet to a COVID-19 vocabulary consisting of 92 core keywords, derived from an infoveillance study on COVID-19 symptoms (Wu et al., 2023a), to ensure that the filtered tweets were highly relevant to the pandemic. The vocabulary covers multiple semantic categories closely related to the pandemic, such as disease names, virus strains, symptoms, and policies.

2. **Text cleaning and normalization:** To remove noisy and irrelevant information and improve text quality, we performed a series of text cleaning operations, including: a) removing metadata such as URLs, @usernames, and #hashtags; b) converting tweets to lowercase; c) removing stop words, punctuation, and special characters; d) performing lemmatization to merge word variants.
3. **Self-reported symptom matching:** To precisely locate self-reported symptom expressions in the cleaned tweets, we carefully constructed a high-coverage concept hierarchy for COVID-19 symptoms, which contains 10 body systems, 257 core symptoms, and 1,808 synonyms. We adopted a combination of exact matching and fuzzy matching methods to match each tweet against the symptom hierarchy. Exact matching could identify explicitly mentioned standard symptom terms in tweets, while fuzzy matching significantly improved the recognition coverage of colloquial and non-standard symptom expressions through synonym expansion.

Through the three-step preprocessing pipeline, we filtered 948,478 COVID-19-related tweets containing self-reported symptoms from the original over 400 million tweets, accounting for 0.2% of the total tweet volume. Despite the low proportion, these tweets were of rich information and high quality, providing a valuable data foundation for conducting large-scale, real-world COVID-19 symptom analysis.

4.2 Symptom and Vaccine Entity Recognition

To further improve the recognition of symptom and vaccine entities, we used deep learning-based named entity recognition (NER) techniques. We adopted the YATO framework (Wang et al., 2023b), using the pre-trained covid-twitter-bert-v2 model (Müller et al., 2023), and fine-tuned it on the METS-CoV dataset (Zhou et al., 2022) for the fine-grained NER task. The model can automatically identify two types of entities relevant to this study from tweets, including symptoms and vaccines.

We applied the fine-tuned NER model for comprehensive inference on the previously screened, approximately one million tweets on self-reported symptoms and extracted the symptom and vaccine entities. Compared to traditional rule-based or dictionary-matching methods, deep learning-based

NER can better handle the colloquial, non-standard, and complex nature of social media text, thereby significantly improving the precision and recall of entity recognition. As shown in Table 1, it achieved a high F1-score of 0.89 on METS-CoV.

Entity	Precision	Recall	F1-score
Symptom	0.86	0.86	0.87
Vaccine	0.88	0.93	0.90
Overall (macro)	0.87	0.90	0.89

Table 1: NER Performance on METS-CoV.

Through entity recognition on the full set of tweets, we extracted a total number of 93,418 symptom entities and 7,313 vaccine entities. It is worth noting that the extracted entities exhibited diverse forms, including standard symptom terms such as “fever” and “shortness of breath”; non-standard descriptions such as “feeling feverish” and “out of breath”; formal vaccine names like “COVID19 Vaccine”; and colloquial expressions like “pfizer vaccine” and “az jab”. This highlights the flexibility and diversity of semantic expressions in social media text and confirms the necessity of adopting deep learning-based entity recognition methods.

4.3 Entity Normalization via Large Language Model

To enhance entity recognition, we designed a multi-stage entity normalization framework based on the Large Language Model (LLM) (Thirunavukarasu et al., 2023) and Retrieval-Augmented Generation (RAG) (Gao et al., 2023). First, we packaged each identified symptom or vaccine entity with its surrounding tweet context into a query prompt and used it as input to retrieve relevant standardized concepts. We utilized 260 symptom codes from the 11th revision of the International Classification of Diseases (ICD-11¹) as the standardization dictionary for symptoms and 47 vaccine generic names as the standardization dictionary for vaccines.

In the retrieval stage, we employed multiple domain-specific pre-trained language models, including BioLORD (Remy et al., 2024), S-PubMedBert (Deka et al., 2022), and BioSimCSE-BioLinkBERT (Kanakarajan et al., 2022), to retrieve 45 coarse-grained candidate concepts for each query. Next, we used the BGE-Gemmaranker (Li et al., 2023) to re-rank the coarse-grained candidate concepts, and selected the top 15

¹ICD: <https://icd.who.int/en>

as the fine-grained candidate concepts.

In the generation stage, we merged the query prompt with the 15 fine-grained candidate concepts into a new prompt and input it into the Qwen-1.5-32B-chat model (Bai et al., 2023), which selected the most appropriate standardized term from the candidate concepts. To prevent error propagation, we introduced a special “no normalized noun” category in the prompt design. When LLM could not find a suitable standardized term or determined that the entity was not a symptom or vaccine, it would return to this category. After normalization processing, we finally obtained 70,642 symptom entities and 221 corresponding standardized terms, as well as 5,452 vaccine entities and 18 corresponding standardized terms. To evaluate the normalization performance, we randomly sampled 100 entities and their standardized terms for manual inspection, achieving an accuracy of 71%. Table 2 presents the most common physical and psychological symptoms, including their normalized symptoms and raw entities in social media texts. It validates that our method effectively standardizes the informal descriptions found in social media into clinically used ICD codes, facilitating the extraction of useful information for medical research.

4.4 Univariate Semantic Shift Analysis

After obtaining normalized entities for symptoms and vaccines, we further explored their semantic evolution patterns across different stages of the COVID-19 pandemic.

We divided the tweet data into monthly time slices and extracted word co-occurrence statistics for each time slice using the method proposed in Sections 3.1 and 3.2. Then, following the dynamic update strategy described in Section 3.3, we iteratively updated the word embeddings on each time slice. We obtained three word vector embeddings: Feb 2020-Sep 2020, Sep 2020-June 2021, and June 2021-Apr 2022. Based on the number of new COVID-19 cases reported by the WHO² and the related pandemic timeline³, we can roughly correlate these with the three key stages of pandemic development. The first stage corresponds to the initial outbreak of the epidemic, during which the public gradually begins to pay attention to and discuss the epidemic; the second stage corresponds to the first rapid spread, during which people start to imple-

²<https://data.who.int/dashboards/covid19/cases>

³<https://www.cdc.gov/museum/timeline/covid19.html>

Type	Normalized Entity (ICD-11)	Count (freq%)	Raw Entity in Social Media
Physical	Nausea or Vomiting	69670 (6.80%)	nausea head, upset stomach, vomit, covid stomach, vom, nause,...
	Fever	67108 (6.55%)	fever, feverish, cough fever, high temp, heat ...
	Cough	55326 (5.40%)	choking cough, sore throat coughing, chesty cough pain, bad cough, ...
	Disturbances of Smell and taste	53891 (5.26%)	lost tast smell, tasteless, couldnt smell, loss taste, loss smell, ...
	Dizziness and giddness	49178 (4.80%)	dizzy head, dizzyer, brain feels like gone, lightheaded, feeling giddy, ...
Psychological	Anxiety	2,121 (0.20%)	freaking out, nervous wreck, on edge, jittery, stressing out, ...
	Depressed mood	2,013 (0.19%)	feeling down, bummed out, low spirits, in a funk, blue, ...
	Irritability	1,822 (0.17%)	cranky, snappy, grouchy, on a short fuse, annoyed, ...
	Psychomotor agitation	1,672 (0.16%)	can't sit still, fidgety, restless, hyper, pacing, ...
	Aggressive behavior	1,423 (0.13%)	lashing out, losing it, going off, snapping, getting in fights, ...

Table 2: Example of the correspondence between social media symptom entities based on large language model normalization and standardized symptom entities.

ment control measures; the third stage corresponds to the full-blown outbreak of the epidemic, accompanied by the emergence of two variant viruses. More details are in Appendix 5. In classic semantic shift problems, anchor words typically align multiple time series into the same embedding space. However, we used Procrustes analysis to align the first two word vectors to the last one, thus obtaining relatively consistent embeddings.

Procrustes analysis is a linear algebra method used to rotate and translate one matrix onto another while preserving the original structure as much as possible (Grave et al., 2019). Specifically, given two matrices A and B , Procrustes analysis seeks a matrix R such that AR closely approximates B by minimizing the Frobenius norm $\|AR - B\|_F$. As it only involves rotation and translation, it does not alter the length or relative positions of word vectors, thus preserving the intrinsic semantic structure of the word vectors. This method ensures that word embeddings from different time slices are aligned in the same semantic space, which helps maintain the global consistency of word vectors and reduces the impact of semantic shift.

We selected psychological symptoms that underwent relatively large changes during the pandemic for univariate semantic shift analysis. As shown in Figure 2 (a), which introduces temporal dynamics into the univariate semantic shift, and Figure 2 (b), which projects all semantic shifts onto a plane, we found that with the passage of time, some symptoms such as anxiety, depression, and irritability were more scattered in the early stages of the pandemic and gradually clustered as the pandemic progressed (Wu et al., 2023b); some symptoms such as psychomotor agitation and aggressive behavior were initially more clustered and gradually dispersed in the later stages; and some symptoms such as thought withdrawal, thought disorder, delu-

sions, and visual hallucinations remained relatively dispersed throughout (Li et al., 2022).

4.5 Synergistic Effect Analysis

To explore the semantic interactions between symptoms and between symptoms and vaccines, we constructed an evolving network of key symptoms and vaccines based on the dynamically updated word embeddings. We used the normalized entities as nodes and calculated their mutual cosine similarity. As shown in Figure 3, their inter-relationships were depicted at each time slice.

We analyzed the dynamic evolution patterns and compared them with real-world symptom patterns reported in relevant papers (Whitaker et al., 2022; Vihta et al., 2023; Looi, 2023). The results showed that the main symptom combinations were closely related to the pandemic stages. The early stage was dominated by respiratory symptoms such as fever and cough, while taste and smell disorders became increasingly prominent in subsequent stages. During the Delta strain period, the proportion of gastrointestinal symptoms such as diarrhea, nausea, and vomiting increased significantly. We further calculated the semantic similarity changes between the vaccine node and various symptom nodes, and found a significant decrease in the association between vaccine administration and multiple key symptoms. This suggests a potential protective effect of the vaccine on these symptoms.

5 Conclusion and Future work

This study proposes an innovative unsupervised dynamic word embedding method to capture the longitudinal evolution of word semantics in social media texts. We apply this method to symptom and vaccine semantic analysis using large-scale COVID-19 Twitter data, achieving promising results. Our method can effectively identify the dy-

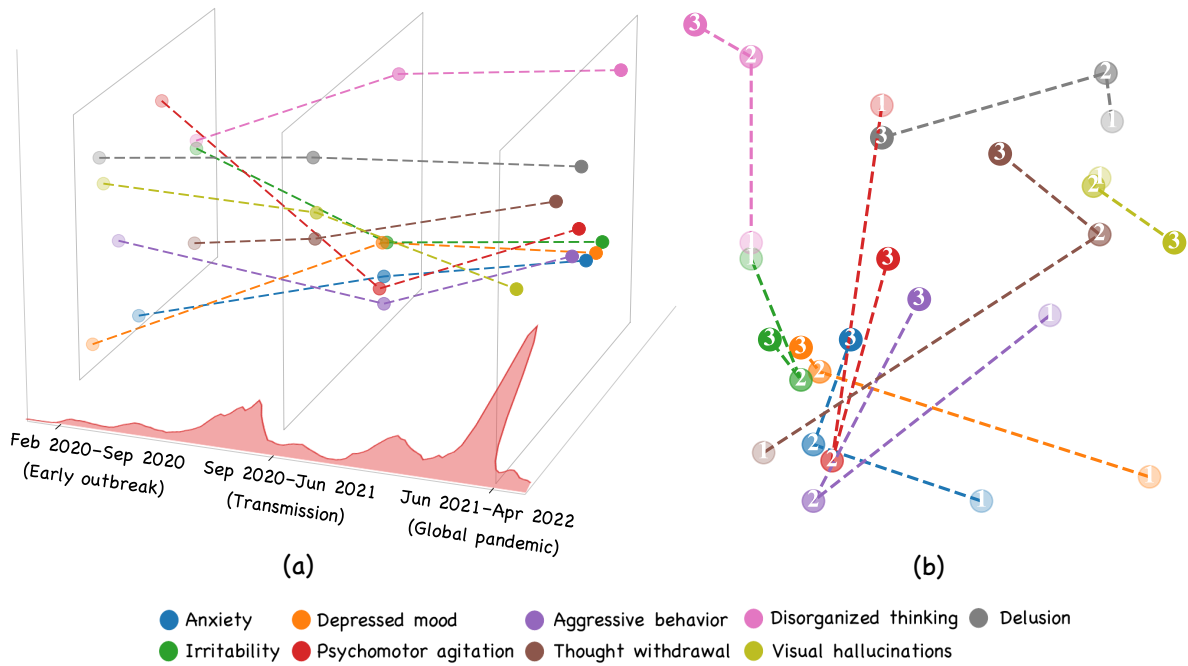


Figure 2: Psychological symptom semantic shift trajectories across different periods. Lines and dots of different colors represent the semantic shift trajectories of various symptoms. Our framework automatically identifies three significant stages of semantic shift, which correspond to the three phases of the pandemic as compared to the number of new COVID-19 cases reported by the WHO. (a) shows the semantic relationships of these symptoms at different stages of the pandemic. (b) visually presents the distribution of symptoms in semantic space across different periods, illustrated using gradients of text and color. Each point in the projection graph represents the semantic position of a symptom at a specific period. By analyzing the trajectories, we observe a converging trend as the global spread of the pandemic progresses, indicating increasing semantic association.

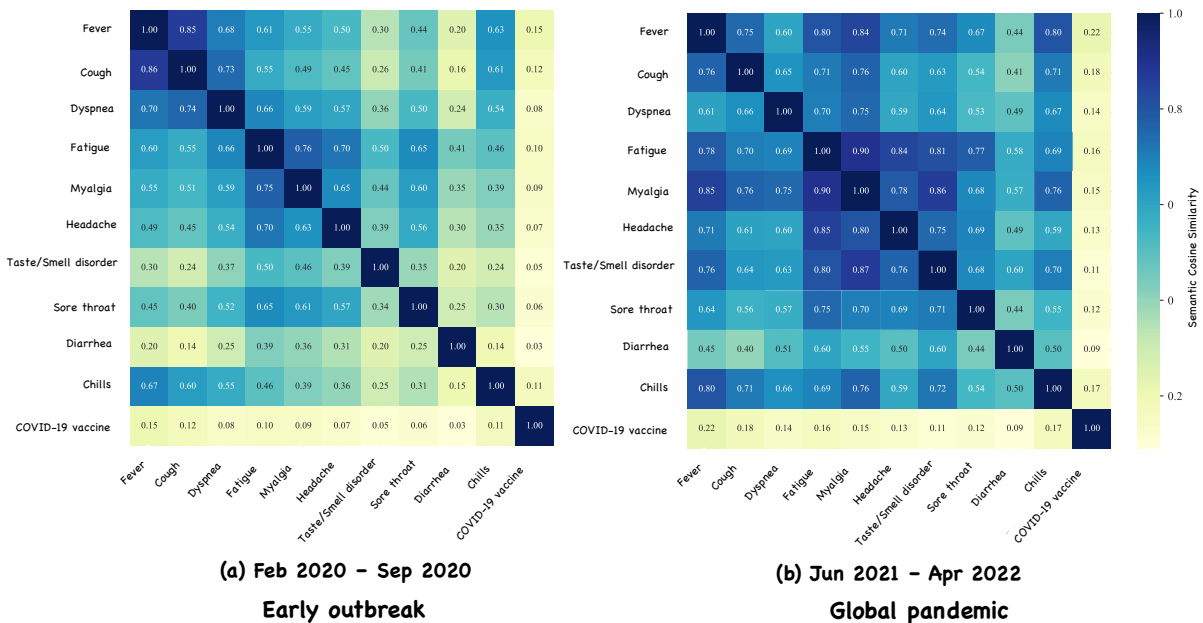


Figure 3: Dynamic longitudinal analysis of COVID-19 symptom-vaccine semantic associations. (a) and (b) show the semantic correlations during the early outbreak (Feb 2020-Sep 2020) and the global pandemic (Jun 2021-Apr 2022), respectively. By analyzing the changes in the semantic associations of symptoms-symptoms and symptoms-vaccine across different periods, we examined the potential combination patterns of symptoms and the sensitivity of symptoms to vaccine protection.

dynamic patterns of symptoms and vaccine semantics at different stages of the epidemic. By introduc-

ing temporal dynamics and iterative update strategies, this method overcomes the limitations of tra-

ditional static word embeddings, achieving longitudinal tracking of semantic content in social media texts. Additionally, we analyze the semantic interactions between different symptoms and between symptoms and vaccines, revealing their synergistic effects over time.

In the future, we plan to expand the application of this method in the field of vertical semantic analysis. On one hand, we will explore more advanced time series modeling techniques to more accurately capture the evolving patterns of word semantics over time. On the other hand, we will integrate semantic shift analysis with technologies such as social network analysis and sentiment analysis to investigate the rich semantic relationships and patterns embedded in social media texts, such as identifying precursor signals of sudden events and revealing the driving factors behind the evolution of public opinion.

Ethics

Our study uses publicly accessible data collected via Twitter’s official API, adhering to the stringent requirements. The research utilized tweets obtained following Twitter’s Privacy Policy, which informs users that their social media content, including profiles and tweets, is public and freely accessible to third parties. To protect privacy, we removed usernames and only analyzed their tweets.

Limitations

There are several limitations in our work: 1) Assessing the accuracy and effectiveness of semantic shift detection methods is inherently challenging due to the lack of normalized evaluation metrics and annotated datasets. The validation of our method relies on case studies and empirical analysis, which may not generalize well enough across different domains or applications; 2) The initial word embeddings are pre-trained using traditional static word embedding methods. The quality and biases inherent in these initial embeddings could influence the subsequent dynamic updates and potentially the detection of semantic shifts.

References

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren,

Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 238–247.

Hyunyoung Choi and Hal Varian. 2012. Predicting the present with google trends. *Economic record*, 88:2–9.

Jyoti Choudrie, Snehasish Banerjee, Ketan Kotecha, Rahee Walambe, Hema Karende, and Juhi Ameta. 2021. Machine learning techniques and older adults processing of online information and misinformation: A covid 19 study. *Computers in human behavior*, 119:106716.

Suparna De, Shalini Jangra, Vibhor Agarwal, Jon Johnson, and Nishanth Sastry. 2023. *Biases and Ethical Considerations for Machine Learning Pipelines in the Computational Social Sciences*, pages 99–113. Springer Nature, Singapore.

Pritam Deka, Anna Jurek-Loughrey, and P Deepak. 2022. Improved methods to aid unsupervised evidence-based fact checking for online health news. *Journal of Data Intelligence*, 3(4):474–504.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.

Oguzhan Gencoglu and Mathias Gruber. 2020. Causal modeling of twitter activity during covid-19. *Computation*, 8(4):85.

Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. Analysing lexical semantic change with contextualised word representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3960–3973.

- Edouard Grave, Armand Joulin, and Quentin Berthet. 2019. Unsupervised alignment of embeddings with wasserstein procrustes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1880–1890. PMLR.
- Kristina Gulordava and Marco Baroni. 2011. A distributional similarity approach to the detection of semantic change in the google books ngram corpus. In *Proceedings of the GEMS 2011 workshop on geometrical models of natural language semantics*, pages 67–71.
- William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016a. Cultural shift or linguistic drift? comparing two computational measures of semantic change. In *Proceedings of the conference on empirical methods in natural language processing. Conference on empirical methods in natural language processing*, volume 2016, page 2116. NIH Public Access.
- William L Hamilton, Jure Leskovec, and Dan Jurafsky. 2016b. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501.
- Keith Harrigian and Mark Dredze. 2022. The problem of semantic shift in longitudinal monitoring of social media: A case study on mental health during the covid-19 pandemic. In *Proceedings of the 14th ACM Web Science Conference 2022*, pages 208–218.
- Renfen Hu, Shen Li, and Shichen Liang. 2019. Diachronic sense modeling with deep contextualized word embeddings: An ecological view. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 3899–3908.
- Yining Hua, Jiageng Wu, Shixu Lin, Minghui Li, Yujie Zhang, Dinah Foer, Siwen Wang, Peilin Zhou, Jie Yang, and Li Zhou. 2024. Streamlining social media information retrieval for public health research with deep learning. *Journal of the American Medical Informatics Association*, page ocae118.
- Shotaro Ishihara, Hiromu Takahashi, and Hono Shirai. 2022. Semantic shift stability: Efficient way to detect performance degradation of word embeddings and pre-trained language models. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 205–216.
- David Jurgens and Keith Stevens. 2009. Event detection in blogs using temporal random indexing. In *Proceedings of the Workshop on Events in Emerging Text Types*, pages 9–16.
- Kamal raj Kanakarajan, Bhuvana Kundumani, Abhijith Abraham, and Malaikannan Sankarasubbu. 2022. [BioSimCSE: BioMedical sentence embeddings using contrastive learning](#). In *Proceedings of the 13th International Workshop on Health Text Mining and Information Analysis (LOUHI)*, pages 81–86, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Yoon Kim, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. Temporal analysis of language through neural language models. In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 61–65.
- Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. Statistically significant detection of linguistic change. In *Proceedings of the 24th international conference on world wide web*, pages 625–635.
- Chaofan Li, Zheng Liu, Shitao Xiao, and Yingxia Shao. 2023. [Making large language models a better foundation for dense retrieval](#). *Preprint*, arXiv:2312.15503.
- Minghui Li, Yining Hua, Yanhui Liao, Li Zhou, Xue Li, Ling Wang, and Jie Yang. 2022. Tracking the impact of covid-19 and lockdown policies on public mental health using social media: Infoveillance study. *Journal of Medical Internet Research*, 24(10):e39676.
- Mun-Keat Looi. 2023. How are covid-19 symptoms changing? *BMJ: British Medical Journal (Online)*, 380:p3.
- Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Google Books Team, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, et al. 2011. Quantitative analysis of culture using millions of digitized books. *science*, 331(6014):176–182.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). *Preprint*, arXiv:1301.3781.
- Sunny Mitra, Ritwik Mitra, Suman Kalyan Maity, Martin Riedl, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2015. An automatic approach to identify word sense changes in text media across timescales. *Natural Language Engineering*, 21(5):773–798.
- Stefano Montanelli and Francesco Periti. 2023. A survey on contextualised semantic shift detection. *arXiv preprint arXiv:2304.01666*.
- Martin Müller, Marcel Salathé, and Per E Kummervold. 2023. Covid-twitter-bert: A natural language processing model to analyse covid-19 content on twitter. *Frontiers in artificial intelligence*, 6:1023281.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.

- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- François Remy, Kris Demuynck, and Thomas De-meester. 2024. [BioLORD-2023: semantic textual representations fusing large language models and clinical knowledge graph insights](#). *Journal of the American Medical Informatics Association*, page ocae029.
- Guy D Rosin, Ido Guy, and Kira Radinsky. 2022. Time masking for temporal language models. In *Proceedings of the fifteenth ACM international conference on Web search and data mining*, pages 833–841.
- Eyal Sagi, Stefan Kaufmann, and Brady Clark. 2011. Tracing semantic change with latent semantic analysis. *Current methods in historical semantics*, 73:161–183.
- Vida Sharifian-Attar, Suparna De, Sanaz Jabbari, Jenny Li, Harry Moss, and Jon Johnson. 2022. [Analysing longitudinal social science questionnaires: Topic modelling with bert-based embeddings](#). In *2022 IEEE International Conference on Big Data (Big Data)*, pages 5558–5567.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature medicine*, 29(8):1930–1940.
- Peter D Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37:141–188.
- Karina-Doris Vihta, Koen B Pouwels, Tim EA Peto, Emma Pritchard, Thomas House, Ruth Studley, Emma Rourke, Duncan Cook, Ian Diamond, Derrick Crook, et al. 2023. Omicron-associated changes in severe acute respiratory syndrome coronavirus 2 (sars-cov-2) symptoms in the united kingdom. *Clinical Infectious Diseases*, 76(3):e133–e141.
- Xuerui Wang and Andrew McCallum. 2006. Topics over time: a non-markov continuous-time model of topical trends. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 424–433.
- Yuqi Wang, Wei Wang, Qi Chen, Kaizhu Huang, Anh Nguyen, Suparna De, and Amir Hussain. 2023a. Fusing external knowledge resources for natural language understanding techniques: A survey. *Information Fusion*, 92:190–204.
- Zeqiang Wang, Yile Wang, Jiageng Wu, Zhiyang Teng, and Jie Yang. 2023b. [YATO: Yet another deep learning based text analysis open toolkit](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 131–139, Singapore. Association for Computational Linguistics.
- Matthew Whitaker, Joshua Elliott, Barbara Bodinier, Wendy Barclay, Helen Ward, Graham Cooke, Christl A Donnelly, Marc Chadeau-Hyam, and Paul Elliott. 2022. Variant-specific symptoms of covid-19 in a study of 1,542,510 adults in england. *Nature Communications*, 13(1):6856.
- Philipp Wicke and Marianna M Bolognesi. 2020. Framing covid-19: How we conceptualize and discuss the pandemic on twitter. *PloS one*, 15(9):e0240010.
- Derry Tanti Wijaya and Reyyan Yeniterzi. 2011. Understanding semantic change of words over centuries. In *Proceedings of the 2011 international workshop on DETecting and Exploiting Cultural diversity on the social web*, pages 35–40.
- Jiageng Wu, Lumin Wang, Yining Hua, Minghui Li, Li Zhou, David W Bates, and Jie Yang. 2023a. Trend and co-occurrence network of covid-19 symptoms from large-scale social media data: Infoveillance study. *Journal of Medical Internet Research*, 25:e45419.
- Jiageng Wu, Xian Wu, Yining Hua, Shixu Lin, Yefeng Zheng, and Jie Yang. 2023b. Exploring social media for early detection of depression in covid-19 patients. In *Proceedings of the ACM Web Conference 2023*, pages 3968–3977.
- Huiling You. 2021. Unsupervised lexical semantic change detection with context-dependent word representations.
- Peilin Zhou, Zeqiang Wang, Dading Chong, Zhijiang Guo, Yining Hua, Zichang Su, Zhiyang Teng, Jiageng Wu, and Jie Yang. 2022. Mets-cov: A dataset of medical entity and targeted sentiment on covid-19 related tweets. *Advances in Neural Information Processing Systems*, 35:21916–21932.

Appendix

Timeline of COVID-19

Figure 4 shows the weekly numbers of self-reported COVID-19 tweets and the total new COVID-19 cases in the four main countries of Twitter users (United States, United Kingdom, Canada, and the Philippines) as reported by the WHO.

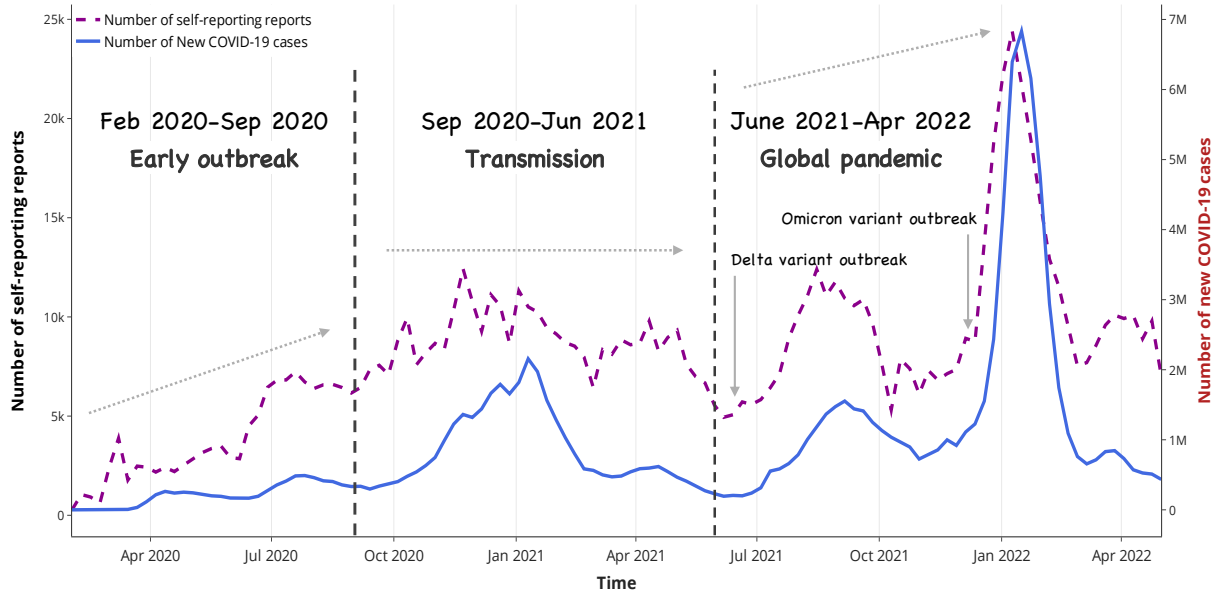


Figure 4: COVID-19 timeline.