Temporal Semantic Analysis of Arabic Mental-Health Discourse using Discrete Wavelet Transform

Rana Salama [1,2], Amal Alqahtani [3], Mona Diab[1] and Abdou Youssef[4]
[1]Carnegie Mellon University, Pittsburgh, Pennsylvania, USA
[2]Faculty of Computers and Artificial Intelligence, Cairo University, Egypt
[3]King Saud University, Riyadh, KSA
[4] George Washington University, Washington DC, USA
rsalama@andrew.cmu.edu, amaalqahtani@ksu.edu.sa
mdiab@andrew.cmu.edu, ayoussef@gwu.edu

Abstract

We present a multiresolution framework for analyzing longitudinal Arabic mental-health discourse drawn from condition-specific Arabic-language X (formerly Twitter) communities. We model each user's posting history as a calendar-time semantic trajectory and apply the Discrete Wavelet Transform (DWT) along the temporal axis to separate semantic variation occurring over finer and broader time scales. To interpret these temporal components linguistically, we introduce *semantic gain*, which measures how strongly corpus-derived phrases become aligned with the reconstructed semantic trajectory when information from a given temporal scale is introduced. We evaluate the framework on longitudinal trajectories from BPD-related and bipolar-related X communities using independent temporal-change measures, chronological permutation tests, cross-scale phrase distinctiveness, and matched high- versus low-gain events. Finer-scale wavelet components align most strongly with local adjacent semantic change, whereas broader-scale components show greater correspondence with distributional change measured by Maximum Mean Discrepancy (MMD). High-gain events also exhibit greater semantic coherence and stronger out-of-window grounding than matched low-gain controls. These findings suggest that longitudinal mental-health discourse on X exhibits multiscale semantic structure, with different forms of linguistic change emerging over different temporal horizons. Combining wavelet decomposition with phrase-based semantic interpretation provides a principled way to characterize both the temporal scale and linguistic content of discourse change. Community labels are treated as discourse contexts rather than indicators of clinical diagnosis[1]

[1]Code for this work is available at https://github.com/engranas/DWT-Temporal-Semantic-Analysis

1 Introduction

Mental health discourse on social media evolves over time. Users may shift between discussions of daily routines, interpersonal relationships, emotional experiences, treatment, support, and culturally situated interpretations of distress. Such variation is difficult to characterize using static post-level representations alone, since the same user may exhibit both persistent semantic themes and short-lived changes over weeks or months. Although computational mental-health research has shown that online language contains measurable signals associated with mental-health communities and self-disclosure, much of this work has emphasized prediction, classification, and dataset construction rather than the temporal organization of discourse itself (Coppersmith et al., 2014a, 2015; Gkotsis et al., 2016; Cohan et al., 2018). We instead study mental-health discourse as a longitudinal semantic signal and ask how its content varies across multiple temporal resolutions. Arabic mental-health discourse is shaped by dialectal variation and by social, cultural, and religious forms of expression, while Arabic mental-health NLP remains comparatively under-resourced. Recent work has begun to address this gap through Arabic mental-health datasets and socioculturally informed analyses (Zaghouani et al., 2026; Alqahtani et al., 2026). In particular, Alqahtani et al. (2026) introduced a longitudinal corpus of Arabic mental-health discourse collected from condition-specific X Communities. In this study, we use that corpus as the basis for our temporal analysis and focus specifically on the borderline personality disorder (BPD) and bipolar-related communities. We treat these community labels as discourse contexts rather than verified clinical diagnoses.

We investigate the use of Discrete Wavelet Transform (DWT) as a multi-resolution framework for this setting. DWT decomposes an ordered signal

Proceedings of the 1st Workshop on Multimodal, Multilingual, and Multicultural Mental Health and Psychotherapy (MultiPsyche), pages 145–158
November 9, 2026 ©2026 Association for Computational Linguistics

into approximation components that capture coarse, slowly varying structure and detail components that capture variation at progressively finer scales. Previous NLP work has applied wavelet and spectral methods to document retrieval, text categorization, feature selection, and semantic representation (Mahajan et al., 2015; Almarwani and Diab, 2021; Salama et al., 2024, 2025). More recently, DWT has been used to compress word and sentence embeddings while preserving much of their semantic information (Salama et al., 2024, 2025). Building on this line of work, we ask whether DWT can also serve as an interpretable temporal representation for longitudinal semantic analysis. To the best of our knowledge, this is the first work to apply DWT to the analysis of longitudinal semantic change in Arabic mental-health discourse, providing a multi-resolution view of how meaning evolves over time.

Our main contributions are:

1. We formulate longitudinal Arabic mental-health discourse from condition-specific X communities as a multi-resolution semantic signal over calendar time and apply DWT to characterize semantic variation at different temporal scales.

2. We develop a reconstruction-based *semantic gain* framework that maps scale-specific DWT variation back to corpus-derived language, enabling linguistic interpretation of finer- and broader-scale temporal changes.

3. We evaluate the framework using independent temporal-change measures, chronological permutation tests, cross-scale semantic distinctiveness, and out-of-window grounding, showing that the identified scale-specific changes are temporally structured and linguistically coherent.

2 Related Work

2.1 Computational Analysis of Mental-Health Discourse

Computational mental-health research has established that social-media language contains signals associated with psychological distress, self-disclosure, and condition-related communities (Coppersmith et al., 2014b). Subsequent studies examined linguistic variation across mental-health communities and developed large-scale datasets for condition-level analysis (Gkotsis et al., 2016; Cohan et al., 2018). While much of this work has emphasized classification and dataset construction, more recent research has highlighted the broader discourse context of mental-health expression, including interpersonal, cultural, religious, and treatment-related themes. Our work builds on this direction by modeling mental-health discourse as a longitudinal semantic signal and analyzing how its linguistic content varies across temporal scales.

2.2 Wavelet and Spectral Methods for Text Representation

Wavelet methods provide a way to analyze signals at multiple resolutions, separating coarse-scale trends from fine-scale variation. In NLP, wavelet-based methods have been explored for document retrieval, text categorization, feature selection, and representation learning. Park et al. (2005) use the Discrete Wavelet Transform for document retrieval, while Darányi et al. (2012) investigate wavelet analysis for text categorization in digital libraries. Mahajan et al. (2015) apply Wavelet Packet Transform to feature selection for short-text classification, showing that wavelet-based representations can be useful for sparse textual inputs.

More recent work connects spectral and wavelet methods to neural text representations. Almarwani and Diab (2021) use the Discrete Cosine Transform to produce efficient sentence embeddings that preserve order-sensitive information. Salama et al. (2024, 2025) apply the Discrete Wavelet Transform to word and sentence embeddings, showing that DWT can compress semantic representations while preserving downstream performance. These studies suggest that transforms developed for signal processing can be useful for representing linguistic structure in neural NLP settings.

3 Discrete Wavelet Transform

DWT provides a multi-resolution representation of a signal by decomposing it into components associated with different temporal scales. Unlike Fourier-based methods, which represent a signal using globally supported frequency components, wavelets are localized in both time and scale, making them well suited for analyzing non-stationary signals whose behavior changes over time (Vetterli and Kovačević, 1995; Daubechies, 1992). Wavelet analysis is defined with respect to a mother wavelet $\psi(t)$, from which a family of scaled and translated

functions is generated:

$$\psi_{a,b}(t) = \frac{1}{\sqrt{a}}\psi\left(\frac{t-b}{a}\right), \quad (1)$$

where a controls the scale of analysis and b determines the wavelet's location along the signal. Smaller scales capture localized, rapidly varying structure, whereas larger scales represent broader, more slowly varying patterns. In the discrete setting, the transform is implemented through a hierarchical filter bank. At each decomposition level j, the signal is convolved with complementary low-pass and high-pass filters and subsequently downsampled, yielding *approximation* and *detail* coefficients:

$$cA_j[k] = \sum_n h[n-2k]\, cA_{j-1}[n], \quad (2)$$

$$cD_j[k] = \sum_n g[n-2k]\, cA_{j-1}[n], \quad (3)$$

where h and g denote the low- and high-pass filters, respectively. The approximation coefficients cA_j retain the lower-frequency structure of the signal, while the detail coefficients cD_j capture variation occurring at the corresponding finer scale (Vetterli and Kovačević, 1995; Daubechies, 1992). Recursively applying the transform to the approximation component produces a hierarchical decomposition,

$$f \longrightarrow \{A_J, D_J, \ldots, D_2, D_1\}, \quad (4)$$

in which A_J represents the coarsest structure and $D_1, \ldots, D_J$ characterize variation across progressively broader scales. This multi-resolution property is particularly useful for longitudinal semantic analysis. When DWT is applied along the temporal dimension of an embedding trajectory, the approximation component represents slowly varying semantic structure, while the detail components capture semantic variation occurring at different temporal resolutions. Importantly, the DWT filters are fixed rather than learned, yielding a deterministic decomposition whose scales are explicitly determined by the sampling interval and decomposition level. Among the available wavelet families, including Haar, Daubechies, Symlets, and Coiflets (Addison, 2017), we use the Haar wavelet. Its compact support and simple pairwise averaging and differencing operations provide a direct correspondence between wavelet levels and calendar-time intervals, which is central to the temporal interpretation developed in this work.[2]

[2]For a broader treatment of wavelet theory and multi-resolution analysis, see Daubechies (1992); Addison (2017); Brunton and Kutz (2019).

4 DWT Temporal Semantic Analysis

To analyze how the semantic content of a user's discourse changes over time and study whether different types of change occur at different temporal resolutions, we model a user's posting history as a semantic time series where each post x_i is preprocessed and encoded using a multilingual sentence encoder. The methodology consists of four stages: (i) Construction of calendar-time semantic trajectories, (ii) DWT multi-resolution temporal decomposition using Haar DWT, (iii) Scale-specific reconstruction and phrase-based semantic interpretation, and (iv) validation using independent temporal-change measures, chronological permutation models, and cross-scale comparisons. Figure 2 summarizes the framework.

Construction of calendar-time semantic trajectories Posts are grouped into fixed 14-day calendar intervals (bins) to preserve a meaningful temporal axis. For user u and interval t, we average the embeddings of all posts within that interval: yielding the longitudinal semantic trajectory. More details are provided in Appendix 10.2.

$$\mathbf{E}_u = [\mathbf{e}_{u,1}, \ldots, \mathbf{e}_{u,T_u}]. \quad (5)$$

DWT Multi-resolution Temporal Decomposition Applying DWT repeatedly to trajectories provides the magnitude and direction of temporal variation in embedding space. DWT compares neighboring portions of a signal. At the first level, it captures relatively local changes between nearby time intervals; at the next level, it compares broader neighboring regions. In this work, we apply a level-2 Haar DWT to each user's trajectory along its temporal axis. The transform is applied independently to each embedding dimension, while the semantic dimensionality of the representation is preserved. The transform generates a set of approximation A_2 and details coefficients for each level of the transform, D_1 and D_2 . With 14-day bins, D_1 captures finer-scale variation over approximately four weeks, while D_2 captures broader variation over approximately eight weeks; A_2 represents the corresponding smoothed temporal structure. These scales refer to the duration of semantic variation, not to different levels of semantic abstraction. Figure 1 illustrates the transform for a single scalar dimension. In the actual experiment, the same oper-

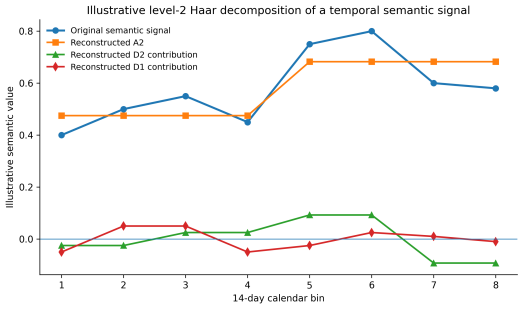


Figure 1: Illustrative level-2 Haar decomposition of a single dimension of a temporal semantic trajectory. The original signal represents eight consecutive 14-day bins. A_2 captures the slowly varying temporal structure, D_2 represents broader variation across approximately four-bin regions, and D_1 captures finer differences between neighboring bins. The actual model applies the same transform independently across all dimensions of each semantic embedding.

ation is applied simultaneously to every dimension of the semantic embedding.

Scale-Specific Reconstruction We then progressively reconstruct the signal by applying inverse DWT to identify the linguistic information contributed at different temporal scales. Where $\mathbf{R}_0$ represents the coarse temporal approximation, $\mathbf{R}_1$ incorporates the broader-scale variation captured by D_2, and $\mathbf{R}_2$ further adds the finer-scale variation captured by D_1, thereby reconstructing the original semantic trajectory up to numerical precision.

$$\mathbf{R}_0 = \widetilde{\mathbf{A}}_2, \tag{6}$$

$$\mathbf{R}_1 = \widetilde{\mathbf{A}}_2 + \widetilde{\mathbf{D}}_2, \tag{7}$$

$$\mathbf{R}_2 = \widetilde{\mathbf{A}}_2 + \widetilde{\mathbf{D}}_2 + \widetilde{\mathbf{D}}_1. \tag{8}$$

The contribution of each detail level can then be expressed as a scale-specific semantic displacement:

$$\mathbf{S}_{D_2}(t) = \mathbf{R}_1(t) - \mathbf{R}_0(t), \tag{9}$$

$$\mathbf{S}_{D_1}(t) = \mathbf{R}_2(t) - \mathbf{R}_1(t). \tag{10}$$

These quantities describe how the semantic representation at a given time point changes when information from a particular temporal scale is introduced. Yet, to provide a linguistic interpretation of the scale-specific displacements, we construct a *phrase vocabulary* (10.3) directly from the cleaned corpus. The vocabulary consists of short lexical expressions observed in the discourse. Each phrase w is assigned an embedding using the same multilingual sentence encoder employed for post representations.

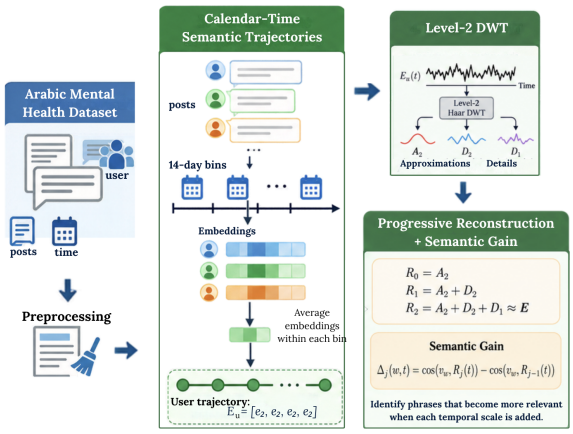


Figure 2: Progressive multi-resolution analysis of longitudinal Arabic mental-health discourse.

4.1 Scale-Specific Semantic Gain

To interpret the linguistic contribution of different coefficients we define the semantic gain introduced by level j as

$$\Delta_j(w,t) = \cos(\mathbf{v}_w, \mathbf{R}_j(t)) - \cos(\mathbf{v}_w, \mathbf{R}_{j-1}(t)). \tag{11}$$

A positive value indicates that phrase w becomes more semantically relevant after the corresponding temporal detail is introduced. We rank phrases by $\Delta_j(w,t)$ to characterize the information contributed by D_2 and D_1. To reduce circularity, candidate phrases used for interpretation must occur outside the temporal window that generated the corresponding DWT component. We similarly ground the coarse approximation A_2 by retrieving semantically related posts and phrases from other periods of the same user's history. Because semantic gain is derived from the wavelet reconstruction, we evaluate whether the detected variation corresponds to independently measurable changes in the semantic trajectory.

5 Experiments

5.1 Dataset

We use the Arabic mental health discourse corpus introduced by Alqahtani et al. (2026), collected from condition-specific Arabic-language X Communities between March 2022 and February 2026. The preprocessed corpus contains 9,582 posts from 1,286 users, covering discourse associated with BPD), bipolar disorder, and ADHD, with 6,324, 2,668, and 590 posts, respectively (Alqahtani et al., 2026). The dataset is predominantly Arabic (9,428 posts), with a small number of English or language-undetermined posts. Each record

contains the post text, timestamp, community label, and anonymized user identifier, enabling the construction of user-level semantic trajectories over calendar time. Community labels are treated as discourse contexts rather than verified clinical diagnoses. Because the present study focuses on longitudinal multi-resolution analysis, we restrict the experiments to the BPD- and bipolar-related communities and remove users appearing in more than one community, leaving 7,944 posts from 1,140 users. Retaining users with at least 20 posts and at least six active 14-day bins, and requiring a contiguous trajectory of at least eight bins (Appendix 10.2), this yields 35 users (22 BPD-related and 13 bipolar-related) with 3,910 posts for our experimental setup. The full filtering pipeline is reported in Appendix 10.1.

5.2 Evaluation Metrics

We evaluate the framework along four dimensions: (i) agreement with alternative embedding-based temporal-change measures, (ii) dependence on chronological ordering, (iii) linguistic distinctiveness across temporal scales, and (iv) semantic quality of high-gain events relative to matched low-gain controls.

Temporal Change Validation. We compare DWT-derived variation with two temporal-change measures that are computationally distinct from the Haar transform but derived from the same embedding trajectories. Agreement with them therefore serves as a consistency check rather than fully external validation.

Maximum Mean Discrepancy. We use RBF-kernel Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) to measure distributional change. At each temporal boundary, the bin-level embeddings of the two preceding bins are compared with those of the two subsequent bins, with the kernel bandwidth set by the median heuristic. Because MMD is not derived from the Haar transform and compares a wider temporal neighborhood than a single pair of bins, it provides a complementary view of broader semantic change. *Adjacent cosine change.* As a local baseline, we compute cosine distance between consecutive bin-level embeddings. Agreement with both MMD and adjacent cosine change is quantified using Spearman rank correlation.

Chronological Permutation. We randomly permute calendar-bin embeddings within each user and recompute DWT detail energy over 1,000 permutations. Embeddings are centered before decomposition, and component energy is computed as the proportion of squared coefficient magnitude assigned to each wavelet component. Comparing observed energy with the permutation distribution tests whether the multi-resolution structure depends on chronological ordering rather than only the marginal distribution of semantic states.

Cross-Scale Semantic Distinctiveness. For each user, we construct phrase profiles from the highest-gain expressions associated with D_1 and D_2. Lexical overlap is measured using Jaccard similarity,

$$J(P_{D_1}, P_{D_2}) = \frac{|P_{D_1} \cap P_{D_2}|}{|P_{D_1} \cup P_{D_2}|}, \qquad (12)$$

with lexical distinctiveness defined as $1 - J(P_{D_1}, P_{D_2})$. Semantic separation is measured using cosine distance between phrase profiles. Significance is assessed using 1,000 within-user permutations that randomly exchange the D_1 and D_2 labels.

Semantic-Gain Quality. High-gain events are compared with low- gain controls matched within the same user and temporal scale. Events are ranked by mean semantic gain, with candidate phrases restricted to discourse outside the source temporal support and its one-bin buffer.

We evaluate phrase coherence, representation shift, and out-of-window grounding. Phrase coherence is the mean pairwise cosine similarity among the highest-gain phrase embeddings. Representation shift is the cosine distance between reconstructed states before and after adding the detail component. Out-of-window grounding is the mean cosine similarity between the centroid of the event's highest-gain phrase embeddings and the three most similar posts by the same user outside the event's temporal support and buffer. Because candidate phrases are drawn from the same out-of-window posts, this measure assesses whether the gained content is consistently retrievable from the user's discourse rather than providing fully independent grounding.

Statistical Testing. To avoid treating multiple events from the same user as independent, event-level measurements are summarized by their median within each user, scale, and event type before testing. Paired user-level comparisons use the Wilcoxon signed-rank test (Wilcoxon, 1945), and

the chronological permutation result is additionally assessed with a sign test. Because each comparison has a pre-specified direction (high-gain > control, $D_1 > D_2$ semantic gain, and observed < permuted detail energy), these tests are one-sided. Per-user permutation p-values are corrected with the Benjamini–Hochberg procedure. Exploratory comparisons between BPD- and bipolar-related trajectories use the two-sided Mann–Whitney U test with Holm correction across energy components.

6 Results

We evaluate the framework on 35 longitudinal trajectories. The trajectories span 8 to 24 consecutive 14-day intervals, representing roughly 16 to 48 weeks of discourse for each user.

6.1 Multi-Resolution Temporal Structure

Figure 3(a) shows the centered DWT energy across temporal scales. Median energy is 0.203 for A_2, 0.303 for D_2, and 0.489 for D_1, with a combined median detail energy of 0.797. The larger D_1 contribution indicates prominent short-range variation, while the non-negligible D_2 energy reflects broader multi-week changes, supporting a multiscale representation of the trajectories. Median user-level semantic gain is 0.053 for D_1 and 0.041 for D_2, with a significant paired difference ($p = .002$). This indicates that finer-scale variation generally produces larger local changes in phrase similarity, without implying that D_1 is inherently more informative. Exploratory condition-level analysis showed a nominal difference in D_2 between BPD-related and bipolar-related trajectories ($p = .046$), while differences in the other energy components were not significant. Given the small sample and multiple comparisons, we treat this result as exploratory.

6.2 Comparison with Embedding-Based Change Measures

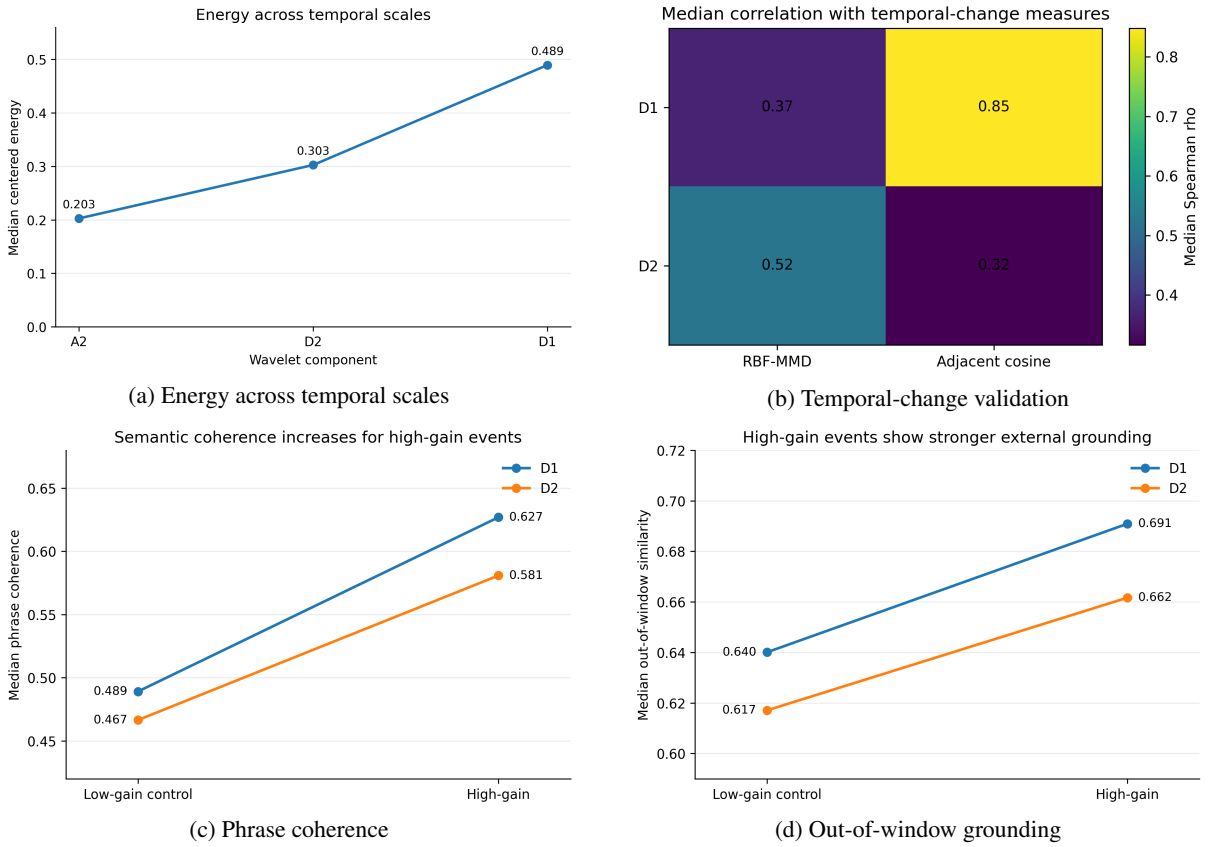
Figure 3*(b)* compares the two detail levels with alternative embedding-based measures of temporal semantic change, computed per user over 5–21 boundaries (median 9). D_1 exhibits a strong association with adjacent-bin cosine change ($\tilde{\rho} = 0.848$), but a weaker association with RBF-MMD ($\tilde{\rho} = 0.367$). In contrast, D_2 correlates more strongly with MMD ($\tilde{\rho} = 0.524$) than with adjacent cosine change ($\tilde{\rho} = 0.316$). This complementary pattern is consistent with the intended multi-resolution interpretation: D_1 primarily reflects localized semantic variation between neighboring states, whereas D_2 is more closely associated with broader changes in the surrounding semantic distribution. The median top-3 overlap between DWT and MMD change locations is 0.667 for both detail levels. Because first-level Haar details are computed from differences between neighboring bins, the strong D_1 association with adjacent cosine change is partly expected by construction and serves mainly as a sanity check. The stronger alignment of D_2 with MMD, which is not derived from the Haar transform, is the more informative result, although both measures share the same embedding representation.

6.3 Chronological Structure

The chronological permutation test further indicates that the observed wavelet structure depends on temporal ordering. Mean detail energy is 0.781 for the observed trajectories, compared with 0.832 after randomizing chronology. The median within-user difference is -0.046, and 28 of 35 users exhibit lower detail energy in their original ordering. The difference is significant under both a paired Wilcoxon test ($p = 4.36 \times 10^{-6}$) and a one-sided sign test ($p = 2.54 \times 10^{-4}$). These results suggest that the observed trajectories exhibit greater temporal coherence than random re-orderings of the same semantic states.

6.4 Cross-Scale Semantic Distinctiveness

The linguistic profiles associated with the two detail levels are also distinct. For each user and scale, the profile is the mean embedding of the 30 phrases with the highest mean semantic gain. The observed mean cosine distance between D_1 and D_2 phrase profiles across the 35 users is 0.078, compared with 0.039 under randomized scale assignments. A within-user permutation test (Section 5.2) yields $p < .001$, the smallest attainable value with 1,000 permutations. Lexical overlap is similarly limited: the median Jaccard similarity between D_1 and D_2 phrase sets is 0.091, corresponding to a median lexical distinctiveness of 0.909. These results indicate that the two temporal scales select largely different phrases rather than simply recovering the same content at different resolutions. Because the two sets are selected by different criteria, this test establishes that they differ, not that they correspond to distinct themes; the absolute profile distance is small, and the thematic interpretation of each scale relies on the qualitative analysis in Section 11.7.



(a) Energy across temporal scales

(b) Temporal-change validation

(c) Phrase coherence

(d) Out-of-window grounding

Figure 3: Summary of the main experimental results. (a) Median centered energy across wavelet components. (b) Median user-level Spearman correlations between D_1 and D_2 and embedding-based temporal-change measures. (c) Median semantic coherence for high-gain events and matched low-gain controls. (d) Median out-of-window grounding (mean similarity to the three most similar out-of-window posts) for high-gain events and controls.

6.5 Semantic-Gain Evaluation

We evaluate 606 events, comprising 303 high-gain events and 303 matched low-gain controls. High-gain events consistently exhibit stronger semantic organization.

As shown in Figure 3c, median phrase coherence for D_1 increases from 0.489 in low-gain controls to 0.627 in high-gain events ($p = 8.9 \times 10^{-9}$). For D_2, coherence increases from 0.467 to 0.581 ($p = 3.9 \times 10^{-6}$).

A similar pattern is observed for independent grounding (Figure 3d). Median out-of-window similarity increases from 0.640 to 0.691 for D_1 and from 0.617 to 0.662 for D_2. These differences are significant for both scales ($p = 4.3 \times 10^{-7}$ for D_1; $p = .003$ for D_2).

Representation shift provides additional evidence at the finer scale. For D_1, high-gain events exhibit a larger median shift than controls ($p = .00087$). The corresponding difference forD_2 is smaller and not significant (0.048 vs. 0.042; $p = .364$), indicating that displacement magnitude alone is insufficient to characterize semantic relevance.

Overall, the results suggest that longitudinal mental-health discourse evolves across complementary temporal scales: D_1 captures more localized semantic fluctuations, whereas D_2 reflects broader shifts in discourse context. Semantic gain further identifies the linguistic themes associated with these changes, offering a way to trace how mental-health-related topics emerge and evolve over time.

7 Discussion

7.1 Multiscale Structure of Mental-Health Discourse

The results indicate that longitudinal Arabic mental-health discourse varies across multiple temporal resolutions. D_1 is most strongly associated with adjacent-bin cosine change, whereas D_2 shows greater correspondence with RBF-MMD. This supports the intended multi-resolution interpretation of the Haar decomposition: D_1 primarily reflects localized semantic variation between neighboring intervals, while D_2 captures broader changes in the

surrounding semantic distribution.

These scales reflect temporal resolution rather than semantic abstraction. A broader component does not necessarily encode more general language; instead, it represents variation developing over a longer time span. This distinction is particularly relevant for mental-health discourse, where themes such as emotions, relationships, treatment, and self-perception may emerge over different temporal horizons.

Chronological permutation further suggests that the observed multi-resolution structure depends on discourse ordering. Most users exhibit lower detail energy in the original trajectory than under random reordering, indicating greater temporal coherence. This supports the interpretation that the wavelet representation captures chronological structure beyond the overall distribution of embeddings, while not implying any specific psychological or clinical process.

Exploratory condition-level comparisons also suggest that BPD-related and bipolar-related discourse may differ more at the broader D_2 scale than in overall wavelet energy. Although the D_2 difference is only nominally significant ($p = .046$), it motivates future work on whether condition-specific discourse differs in the temporal scale over which semantic changes develop.

7.2 Interpretability of Scale-Specific Semantic Change

Semantic gain provides a linguistic interpretation of the temporal decomposition. High-gain events yield more coherent phrase sets than matched low-gain controls at both D_1 and D_2, with stronger out-of-window grounding elsewhere in the same user's history. This suggests that large DWT-derived gains correspond to semantically consistent and independently supported patterns rather than arbitrary movement in embedding space.

The cross-scale analysis further shows low lexical overlap and significant semantic separation between D_1 and D_2, indicating that different temporal scales emphasize distinguishable linguistic content. At the same time, representation-shift magnitude alone is insufficient for interpretation: while high-gain D_1 events show larger shifts than controls, the corresponding D_2 difference is not significant. This suggests that the direction and semantic coherence of change are more informative than magnitude alone.

Accordingly, the framework separates two complementary functions: DWT identifies *when and at what temporal scale* semantic variation occurs, while semantic gain characterizes *what linguistic content* is associated with that variation. For BPD-related and bipolar-related X communities, this enables the study of both short-range fluctuations and broader thematic transitions without treating them as indicators of clinical state.

7.3 Implications for Longitudinal Mental-Health Discourse Analysis

A multi-resolution perspective is useful because mental-health discourse may evolve through both localized fluctuations and broader thematic transitions. Distinguishing these patterns may help examine how themes related to emotions, relationships, treatment, and social support emerge, recur, or persist over time within condition-specific online communities.

These patterns characterize discourse rather than clinical state. With larger cohorts and validated thematic annotations, future work could test whether different mental-health communities exhibit reproducible multiscale discourse signatures and whether particular themes preferentially occur at finer or broader temporal resolutions.

7.4 Future Directions

The analysis is limited to 35 users with sufficiently long and contiguous histories, a selection that favors consistently active posters. Larger cohorts and longer trajectories would support deeper decompositions and better-powered condition comparisons. The findings also depend on the 14-day bins and their alignment, single-bin interpolation, the level-2 Haar decomposition, and a single multilingual encoder on which all measures rely. Future work should test robustness to these choices, including Arabic-specific or alternative encoders. Phrase interpretation would benefit from Arabic orthographic normalization, clitic segmentation, and dialect-aware stopword lists, and from human evaluation in which annotators judge whether high-gain events reflect meaningful shifts and whether the gained phrases describe them accurately. Community labels should be treated as discourse contexts, not clinically verified populations.

8 Conclusion

We presented a proof-of-concept multi-resolution framework that models each user's posting his-

tory as a calendar-time semantic trajectory, separates finer and broader semantic variation with DWT, and interprets each temporal scale through semantic gain. Because both components operate on sentence embeddings, the framework is not tied to Arabic and can be applied to other languages by replacing the condition lexicon, stopword list, and, if needed, the encoder. On BPD- and bipolar-related X trajectories, D_1 aligned with local adjacent change and D_2 with broader distributional change measured by MMD; high-gain events showed more coherent phrases and stronger out-of-window grounding than matched controls; and the two scales selected largely different phrases. These results suggest that online mental-health discourse changes at multiple temporal resolutions, although the current sample does not support strong condition-level comparisons. Community labels are treated solely as discourse contexts, not evidence of clinical diagnosis.

9 Limitations

The data come from condition-specific Arabic-language X communities. Community participation should not be interpreted as evidence of clinical diagnosis, symptom severity, or individual mental-health status. The findings therefore characterize patterns of online discourse rather than clinical trajectories and may not generalize to other languages, platforms, or populations.

Use of AI Assistance. OpenAI ChatGPT was used during manuscript preparation for language editing, organizational and drafting assistance. The tool was also used to assist in exploring alternative presentations of the analyses. Generative AI was not treated as an author and was not used as an authoritative source of scientific evidence.

Ethics Statement

This study uses an existing corpus of public posts from Arabic-language X Communities (Alqahtani et al., 2026). User identifiers were hashed, condition mentions were masked, and only aggregate statistics and short phrases are reported; raw posts are not redistributed, in line with X's terms of service. Community membership is treated solely as a discourse context: the framework does not infer diagnosis, severity, or individual mental-health status and should not be used for clinical decisions. Because methods that track individual discourse over time could be misused to profile or monitor users, we restrict their use to aggregate research under appropriate ethical oversight.

References

Paul S Addison. 2017. *The illustrated wavelet transform handbook: introductory theory and applications in science, engineering, medicine and finance*. CRC press.

Nada Almarwani and Mona Diab. 2021. Discrete cosine transform as universal sentence encoder. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 419–426, Online. Association for Computational Linguistics.

Amal Alqahtani, Rana Aref Salama, and Mona Diab. 2026. Understanding the sociocultural dimensions of mental health discourse in arabic x communities. In *Proceedings of the 11th Social Media Mining for Health Research and Applications (SMM4H-HeaRD 2026) Workshop and Shared Tasks*, pages 286–304.

Steven L Brunton and J Nathan Kutz. 2019. Data-driven science and engineering: machine learning. *Dynamical Systems, and Control*, page 18.

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.

Glen Coppersmith, Mark Dredze, and Craig Harman. 2014a. 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.

Glen Coppersmith, Mark Dredze, and Craig Harman. 2014b. 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, Baltimore, Maryland, USA. Association for Computational Linguistics.

Glen Coppersmith, Mark Dredze, Craig Harman, and Kristy Hollingshead. 2015. From adhd to sad: Analyzing the language of mental health on twitter through self-reported diagnoses. In *Proceedings of the 2nd workshop on computational linguistics and clinical psychology: from linguistic signal to clinical reality*, pages 1–10.

Sándor Darányi, Peter Wittek, and Milena Dobreva. 2012. Using wavelet analysis for text categorization in digital libraries: A first experiment with strathprints. *International Journal on Digital Libraries*, 12(1):3–12.

Ingrid Daubechies. 1992. *Ten lectures on wavelets*. SIAM.

George Gkotsis, Anika Oellrich, Tim Hubbard, Richard Dobson, Maria Liakata, Sumithra Velupillai, and Rina Dutta. 2016. The language of mental health problems in social media. In *Proceedings of the third workshop on computational linguistics and clinical psychology*, pages 63–73.

Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. A kernel two-sample test. *The journal of machine learning research*, 13:723–773.

Anuj Mahajan, Sharmistha Jat, and Shourya Roy. 2015. Feature selection for short text classification using wavelet packet transform. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 321–326, Beijing, China. Association for Computational Linguistics.

Laurence AF Park, Kotagiri Ramamohanarao, and Marimuthu Palaniswami. 2005. A novel document retrieval method using the discrete wavelet transform. *ACM Transactions on Information Systems (TOIS)*, 23(3):267–298.

Rana Salama, Abdou Youssef, and Mona Diab. 2024. Semantic compression for word and sentence embeddings using discrete wavelet transform. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15963–15977, Bangkok, Thailand. Association for Computational Linguistics.

Rana Salama, Abdou Youssef, and Mona Diab. 2025. Combining discrete wavelet and cosine transforms for efficient sentence embedding. *Preprint*, arXiv:2508.00420.

Martin Vetterli and Jelena Kovačević. 1995. *Wavelets and subband coding*. Prentice-hall.

Frank Wilcoxon. 1945. Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6):80–83.

Wajdi Zaghouani, Eman Sedqy Shlkamy, and Mabrouka Bessghaier. 2026. From posts to pressure: An Arabic dataset about stress and mental-health monitoring. In *Proceedings of the 2nd Workshop on NLP for Languages Using Arabic Script*, pages 422–432, Rabat, Morocco. Association for Computational Linguistics.

10 Appendix

10.1 Sample Construction

Table 1 reports the number of posts and users retained after each filtering step. The largest reduction comes from the longitudinal requirements: of the 1,140 eligible users, only 65 have at least 20 posts across at least six active 14-day bins, and 35 of these have a contiguous trajectory of at least eight bins (approximately 16 weeks). Retained trajectories span 8 to 24 bins (16 to 48 weeks).

Filtering step	Posts	Users
Full corpus	9,582	1,289
BPD- and bipolar-related communities	8,992	1,179
Cross-community users removed	7,965	1,145
Non-empty after cleaning	7,944	1,140
$\geq$20 posts and $\geq$6 active bins	4,972	65
Contiguous trajectory $\geq$8 bins	3,910	35

Table 1: Sample attrition from the full corpus to the longitudinal analysis set (22 BPD-related and 13 bipolar-related users).

10.2 Calendar-Time Semantic Trajectories

We first represent each user's posting history as a regularly sampled semantic time series. Posts are preprocessed by removing duplicate content, URLs, user mentions, and explicit condition-name variants that could trivially reveal the source community. Each post x_i is encoded using `paraphrase-multilingual-mpnet-base-v2`, yielding

$$\mathbf{z}_i = f(x_i), \tag{13}$$

where $\mathbf{z}_i \in R^d$ denotes the semantic embedding of post i. Posts are then grouped into non-overlapping 14-day calendar intervals. For user u and temporal bin t, the bin-level representation is defined as

$$\mathbf{e}_{u,t} = \frac{1}{|\mathcal{B}_{u,t}|} \sum_{i \in \mathcal{B}_{u,t}} \mathbf{z}_i, \tag{14}$$

where $\mathcal{B}_{u,t}$ denotes the set of posts authored by user u during interval t. The resulting longitudinal semantic trajectory is

$$\mathbf{E}_u = [\mathbf{e}_{u,1}, \mathbf{e}_{u,2}, \ldots, \mathbf{e}_{u,T_u}]. \tag{15}$$

Calendar-time aggregation is essential to the proposed analysis because it assigns a consistent temporal interpretation to the wavelet scales across users. We interpolate only isolated single-bin gaps; longer gaps define breaks in the trajectory. We retain users with at least 20 posts, at least six active temporal bins, and a contiguous segment of at least eight bins. Trajectory lengths are truncated, when necessary, to a multiple of four to support the level-2 decomposition.

10.3 Phrase Dictionary and Semantic-Gain Interpretation

To interpret the semantic contribution of each DWT component, we construct a phrase dictionary directly from the cleaned discourse corpus. The dictionary contains short lexical units extracted from

ID	Arabic phrase	English gloss
160	أدوية	medications
704	الأفكار السلبية	negative thoughts
875	التعلق العاطفي	emotional attachment
939	الجامعة	university
997	الحب	love
1047	الحمدلله	praise be to God
1119	الدواء	medication
1123	الدوبامين	dopamine
1285	الصراع الداخلي	internal conflict
1408	العلاج	treatment
1413	العلاج السلوكي الجدلي	dialectical behavior therapy
1608	المشاعر	emotions
1702	النوم	sleep
2081	بعلاقات	in relationships
2877	حب الذات	self-love
3243	رمضان	Ramadan
3553	طبيب	doctor

Table 2: Sample entries from the phrase dictionary extracted from the Arabic mental-health discourse corpus. English glosses are provided only for readability and are not used by the model.

the observed posts, including unigrams and short multiword expressions. Each phrase is assigned a unique identifier and encoded using the same multilingual sentence encoder used for the post representations. Thus, every phrase w is associated with an embedding $\mathbf{v}_w$ in the same semantic space as the longitudinal trajectory.

Table 2 shows a small sample of phrases appearing in the dictionary. The dictionary is not manually defined and does not rely on an external mental-health lexicon; rather, it reflects expressions occurring naturally in the corpus.

For a given user and temporal position, the complete dictionary is further restricted to an *out-of-window candidate set*. Specifically, phrases occurring within the temporal support of the DWT component being interpreted are excluded, together with the neighboring buffer interval. Candidate phrases therefore originate from other portions of the observed discourse rather than from the same temporal window that generated the wavelet component. This restriction reduces the possibility of trivially recovering phrases that were already present in the source interval.

Let $\mathcal{V}_{u,t,j}$ denote the candidate phrase set for user u, time t, and detail level j. For each candidate phrase $w \in \mathcal{V}_{u,t,j}$, we compare its embedding with the reconstructed semantic representation before and after adding the corresponding detail component:

$$\Delta_j(w,t) = \cos\left(\mathbf{v}_w, \mathbf{R}_j(t)\right) - \cos\left(\mathbf{v}_w, \mathbf{R}_{j-1}(t)\right). \tag{16}$$

A positive value of $\Delta_j(w,t)$ indicates that the reconstructed semantic state becomes more similar to phrase w after the DWT detail component is introduced. The phrases are then ranked by this increase in similarity.

For example, suppose the out-of-window candidate set for one D_2 event contains expressions such as الحب (love), حب الذات (self-love), التعلق العاطفي (emotional attachment), العلاج (treatment), and النوم (sleep). If the similarity of الحب changes from

$$\cos(\mathbf{v}_{\text{الحب}}, \mathbf{R}_0(t)) = 0.4787 \tag{17}$$

to

$$\cos(\mathbf{v}_{\text{الحب}}, \mathbf{R}_1(t)) = 0.5369, \tag{18}$$

then its semantic gain is

$$\Delta_{D_2}(\text{الحب}, t) = 0.5369 - 0.4787 = 0.0582. \tag{19}$$

The interpretation is therefore not that DWT directly assigns a linguistic label such as "love." Instead, the DWT component moves the reconstructed embedding in a particular semantic direction, and the phrase dictionary is used to identify which naturally occurring expressions become more compatible with that reconstructed state. This provides a linguistic interpretation of the otherwise continuous wavelet representation while grounding the interpretation in expressions that occur in the observed corpus.

11 Interpreting the Haar DWT on Semantic Trajectories

This appendix provides an intuitive description of how the Discrete Wavelet Transform (DWT) is applied to the longitudinal semantic trajectories used in this study and how the resulting coefficients should be interpreted. The explanation focuses on the level-2 Haar transform used in our experiments.

11.1 From Posts to a Temporal Semantic Signal

For each user, posts are first aggregated into consecutive 14-day calendar bins. All posts within a bin are encoded and averaged, producing one semantic vector per interval. A user's trajectory can therefore be written as

$$\mathbf{E}_u = [\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_T], \tag{20}$$

where each $\mathbf{e}_t \in R^d$ is a semantic embedding.

The temporal ordering is essential. The DWT is applied along this *time axis*, independently across the d embedding dimensions. Thus, the transform does not decompose the embedding into more or less abstract meanings. Instead, it decomposes how the semantic representation varies over time.

11.2 First-Level Haar Decomposition

The Haar transform begins by considering neighboring temporal bins in pairs. For two adjacent semantic states $\mathbf{e}_{2k-1}$ and $\mathbf{e}_{2k}$, it computes an approximation coefficient and a detail coefficient:

$$\mathbf{A}_{1,k} = \frac{\mathbf{e}_{2k-1} + \mathbf{e}_{2k}}{\sqrt{2}}, \tag{21}$$

$$\mathbf{D}_{1,k} = \frac{\mathbf{e}_{2k-1} - \mathbf{e}_{2k}}{\sqrt{2}}. \tag{22}$$

The approximation coefficient $\mathbf{A}_{1,k}$ summarizes what is shared across the two neighboring intervals, while $\mathbf{D}_{1,k}$ represents their difference.

Because each temporal bin covers 14 days, a D_1 coefficient is based on approximately four weeks of temporal support. A large magnitude $\|\mathbf{D}_{1,k}\|$ indicates that the semantic representations of the two neighboring bins differ substantially, whereas a small magnitude indicates relative semantic stability.

Importantly, the sign of an individual wavelet coefficient does not represent positive or negative sentiment. It indicates the *direction* of change along a particular embedding dimension. The magnitude of the coefficient, or its energy, is used to quantify the strength of temporal variation.

11.3 Second-Level Haar Decomposition

The first-level approximation sequence is then decomposed again. Two neighboring A_1 coefficients are combined as

$$\mathbf{A}_{2,k} = \frac{\mathbf{A}_{1,2k-1} + \mathbf{A}_{1,2k}}{\sqrt{2}}, \tag{23}$$

$$\mathbf{D}_{2,k} = \frac{\mathbf{A}_{1,2k-1} - \mathbf{A}_{1,2k}}{\sqrt{2}}. \tag{24}$$

The resulting decomposition is therefore

$$\mathbf{E}_u \longrightarrow \{\mathbf{A}_2, \mathbf{D}_2, \mathbf{D}_1\}. \tag{25}$$

A_2 summarizes the slowly varying temporal structure, while D_2 measures differences between broader neighboring portions of the trajectory. With 14-day bins, D_2 has approximately eight weeks of temporal support.

The three components can therefore be interpreted as

Component	Temporal support	Interpretation
A_2	∼8-week aggregation	smoothed level-2 temporal structure
D_2	∼8-week support	broader semantic variation
D_1	∼4-week support	finer semantic variation

These descriptions refer to temporal resolution rather than semantic abstraction. Both A_2 and D_2 are derived at the second decomposition level, but A_2 is the smoothed approximation whereas D_2 captures the corresponding broader-scale detail.

11.4 What Do the Coefficients Represent?

A wavelet coefficient should be interpreted as a *scale-specific temporal difference*. For example, D_1 compares two neighboring 14-day semantic states, whereas D_2 compares two broader neighboring regions after local variation has first been averaged.

Consider an illustrative eight-bin scalar trajectory. A level-2 Haar transform may produce

$$A_2 = [0.950,\ 1.365], \tag{26}$$

$$D_2 = [-0.050,\ 0.185], \tag{27}$$

$$D_1 = [-0.071,\ 0.071,\ -0.035,\ 0.014]. \tag{28}$$

The values themselves are expressed in the orthonormal Haar basis and should therefore not be read as raw semantic scores. Rather, their relative magnitudes indicate how much variation is present at each temporal location and scale. For example, the second D_2 coefficient (0.185) is larger in magnitude than the first (-0.050), indicating a stronger broader-scale transition in the second half of the illustrative trajectory.

In the actual experiment, these coefficients are vectors rather than scalars:

$$\mathbf{D}_{j,k} \in R^d. \tag{29}$$

We therefore summarize their strength using their magnitude or energy. For a detail level D_j, the corresponding energy can be expressed as

$$\mathcal{E}_{D_j} = \sum_k \|\mathbf{D}_{j,k}\|_2^2 . \tag{30}$$

Higher energy indicates greater temporal semantic variation represented at that scale.

11.5 Reconstructing Each Temporal Scale

Although wavelet coefficients quantify temporal variation, they are not themselves ordinary semantic embeddings. To interpret them linguistically, we use the inverse DWT to reconstruct each component at the original temporal resolution.

Let $\widetilde{\mathbf{A}}_2$, $\widetilde{\mathbf{D}}_2$, and $\widetilde{\mathbf{D}}_1$ denote the reconstructed components. We progressively recover the trajectory as

$$\mathbf{R}_0 = \widetilde{\mathbf{A}}_2, \tag{31}$$

$$\mathbf{R}_1 = \widetilde{\mathbf{A}}_2 + \widetilde{\mathbf{D}}_2, \tag{32}$$

$$\mathbf{R}_2 = \widetilde{\mathbf{A}}_2 + \widetilde{\mathbf{D}}_2 + \widetilde{\mathbf{D}}_1. \tag{33}$$

Here, $\mathbf{R}_0$ represents the temporally smoothed semantic trajectory. Adding D_2 produces $\mathbf{R}_1$ and reveals the broader-scale variation that was absent from the approximation. Adding D_1 produces $\mathbf{R}_2$, which restores the finer-scale variation and reconstructs the original trajectory up to numerical precision.

This provides a useful interpretation of each wavelet level as a semantic displacement:

$$\mathbf{S}_{D_2}(t) = \mathbf{R}_1(t) - \mathbf{R}_0(t), \tag{34}$$

$$\mathbf{S}_{D_1}(t) = \mathbf{R}_2(t) - \mathbf{R}_1(t). \tag{35}$$

Thus, rather than interpreting the raw coefficient directly, we ask how the semantic state changes when information from a particular temporal scale is restored.

11.6 From Wavelet Variation to Linguistic Interpretation

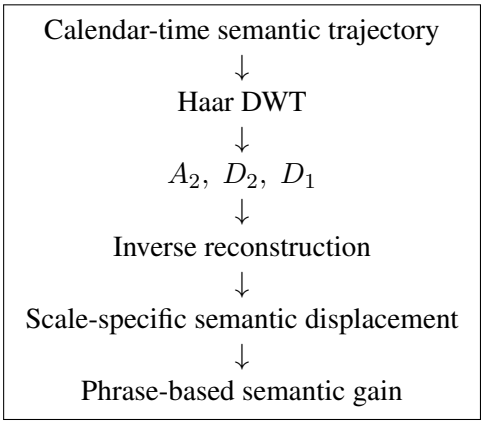
DWT determines the temporal scale and magnitude of change, but it does not directly assign linguistic meaning to that change. We therefore compare the reconstructed states with embeddings of phrases observed in the corpus.

For phrase w with embedding $\mathbf{v}_w$, semantic gain is defined as

$$\Delta_j(w,t) = \cos(\mathbf{v}_w, \mathbf{R}_j(t)) - \cos(\mathbf{v}_w, \mathbf{R}_{j-1}(t)). \tag{36}$$

If $\Delta_j(w,t) > 0$, adding the corresponding wavelet detail component moves the semantic trajectory closer to phrase w. The largest-gain phrases therefore provide a linguistic description of the scale-specific semantic displacement.

The complete interpretation can be summarized as



(37)

Accordingly, the DWT answers *at what temporal scale semantic variation occurs*, while semantic gain provides evidence about *which linguistic content is associated with that variation*.

11.7 Case Study: Scale-Specific Semantic Change

To examine whether the learned wavelet components correspond to observable discourse changes, we inspect an anonymized BPD-related user with a 16-bin trajectory covering approximately 32 weeks. The D_2 component reveals a sequence of broader semantic shifts. Early intervals emphasize pain-related expressions (المعاناة، الألم، جروحك), followed by self-confidence (الثقة بالنفس), perceived worth and invalidation (إنك تستحق، لا يرى قيمتك), and subsequently a highly coherent relationship-related cluster (الحب، للحب، بين الحب، بعلاقات). Later intervals shift toward freedom and control and eventually toward negative thoughts and internal conflict.

The strongest relationship-related D_2 event yields a mean semantic gain of approximately 0.052 and phrase coherence of 0.889. For example, the similarity of الحب to the reconstructed semantic state increases from 0.479 before adding

D_2 to 0.537 afterward ($\Delta = 0.058$). Importantly, this region also corresponds to the largest independent distributional change measured by RBF-MMD ($\mathrm{MMD} = 0.639$). Across the trajectory, D_2 correlates significantly with MMD ($\rho = 0.691$, $p = 0.0089$), supporting the interpretation that the broader wavelet component captures genuine changes in the surrounding semantic distribution.

In contrast, D_1 shows weaker agreement with independent change measures for this user, indicating that fine-scale components are not equally interpretable in every trajectory. This case therefore illustrates both the utility and the limitation of the proposed multi-resolution analysis: D_2 captures a coherent broader thematic transition, whereas finer-scale variation may remain comparatively noisy.

Figure 4 illustrates the temporal behavior of the two detail levels for the same user. The D_2 component shows clearer correspondence with the independent MMD signal, whereas the D_1 component exhibits weaker alignment with adjacent cosine change. This case illustrates that different wavelet scales can capture different forms of temporal semantic variation and that their interpretability may vary across users.

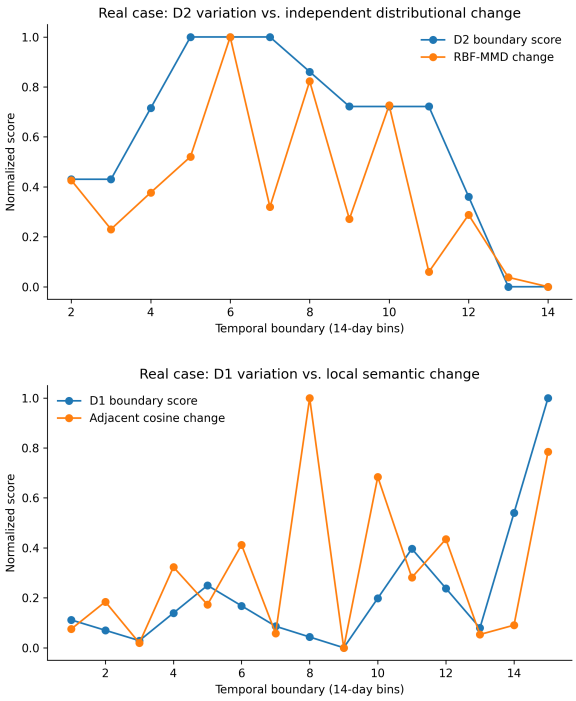


Figure 4: Case-study comparison between DWT-derived temporal variation and independent semantic-change measures for one anonymized user. Top: the D_2 boundary score tracks broader distributional change measured by RBF-MMD. Bottom: the D_1 boundary score is compared with adjacent cosine change, reflecting finer local variation. Scores are min–max normalized within each series for visualization.