Text Knows What, Tables Know When: Clinical Timeline Reconstruction via Retrieval-Augmented Multimodal Alignment

Sayantan Kumar[1*], Shahriar Noroozizadeh[2,3*], Juyong Kim[2*], Jeremy C. Weiss[1]

[1]*National Library of Medicine, National Institutes of Health, Bethesda, MD, USA*
[2]*Machine Learning Department, Kenneth C. Griffin School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA*
[3]*Heinz College of Information Systems Management and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA*

Clinical language models increasingly operate over electronic health records (EHRs), yet patient records are not stored as temporally grounded trajectories. Clinical notes describe symptoms, assessments, and disease progression, but often compress or narratively reorder events. Structured EHR rows provide timestamps for labs, medications, vitals, and procedures, but capture only part of the clinical story. We formulate clinical timeline reconstruction as retrieval-augmented temporal grounding: constructing a patient trajectory by using narrative text for event semantics and structured rows as partial temporal evidence. We introduce a scaffolded workflow that extracts central narrative events, builds an initial temporal scaffold, attaches non-central events, and calibrates timestamps using retrieved structured EHR rows. We evaluate on 40 discharge summaries, including 15 i2b2-derived and 25 MIMIC-IV summaries, each with manual gold-standard timelines and aligned structured EHR data. Across models, multimodal calibration left event match rates largely unchanged and generally improved temporal performance: mean paired case-level multimodal–unimodal differences were positive in 7 of 12 model-metric comparisons across concordance and AULTC, with none negative. However, uncertainty was substantial given the 40-case sample; paired case-level bootstrap intervals excluded zero only for the DeepSeek V3.2 AULTC improvement. A gap analysis shows that 35.1% of text-derived events have no structured counterpart. These findings support treating structured EHR data as partial temporal evidence for narrative-derived patient trajectories.

Keywords: Clinical timeline reconstruction; electronic health records; clinical narratives; retrieval-augmented generation; multimodal alignment; large language models

1. Introduction

Time is central to clinical reasoning in critical care. Patients often follow rapidly evolving trajectories, requiring symptoms, physiologic measurements, treatments, and outcomes to be interpreted relative to admission, deterioration, intervention, and recovery. Sepsis trajectories exemplify this challenge. Although Sepsis-3 remains important for trial eligibility and computational phenotyping,[1,2] sepsis is better understood as a set of heterogeneous clinical trajectories with a shared dysregulatory response rather than as a single, uniform course. Similar

[*] These authors contributed equally to this work.

© 2026 The Authors. Open Access chapter published by World Scientific Publishing Company and distributed under the terms of the Creative Commons Attribution Non-Commercial (CC BY-NC) 4.0 License.

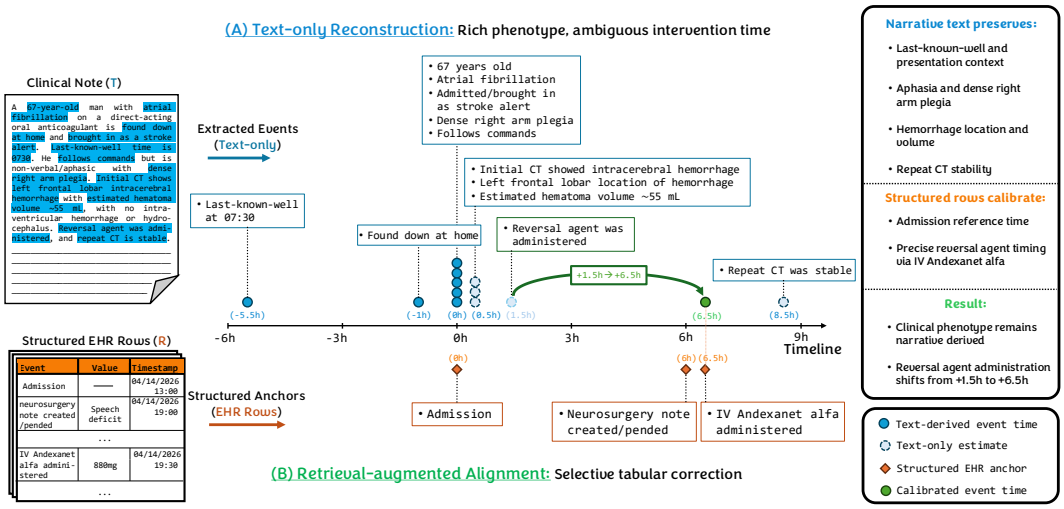


Fig. 1. Motivation for retrieval-augmented temporal grounding using a synthetic stroke-alert case constructed for illustration. (**A**) Text-only reconstruction can recover clinically meaningful narrative events, but intermediate event times may remain ambiguous. (**B**) Retrieved structured EHR rows provide partial temporal evidence that can calibrate text-derived events when clinically relevant, improving temporal placement while preserving narrative event content.

temporal challenges extend across critical illness: risk forecasts depend on when evidence becomes available; treatment-response analyses require distinguishing pre- and post-intervention measurements; and retrospective cohort definitions may change when the timing of symptom onset, diagnostic testing, or treatment initiation is misaligned. This increasing focus on temporal structure is evident in early-warning systems,[3] studies showing that predictive utility varies with timing,[4] and research on the temporal characteristics of sepsis subtypes and progression.[5]

The growing emphasis on time in critical care reveals an important methodological gap: electronic health records (EHRs) are not organized as temporally grounded patient trajectories. Critical-care repositories such as MIMIC-III/IV[6,7] make this challenge explicit by combining structured time-series measurements with narrative clinical documentation. These modalities provide complementary but asymmetric evidence for timeline construction (Figure 1). Structured rows offer relatively precise temporal information from laboratory results, medications, procedures, vital signs, diagnoses, and administrative events. Clinical narratives, by contrast, provide richer accounts of the patient course, including symptoms, clinician assessments, patient-reported histories, functional status, disease progression, and contextual modifiers. For example, a note describing "increasing lethargy and peripheral mottling" may precede an increase in Sequential Organ Failure Assessment (SOFA) scores. Similarly, documented oliguria despite stable creatinine may indicate clinically meaningful states that structured measurements do not fully capture.

Recent studies point to both the promise and the difficulty of this problem. Multimodal annotation of discharge summaries aligned with structured EHR data demonstrated that access to both modalities can improve the precision of inpatient event timelines, highlighting the value of structured rows as temporal evidence when aligned with the narrative record.[8] At the same time, text-only timeline reconstruction remains challenging in realistic EHR-style notes, where dense event structure, abbreviations, and semi-structured formatting make

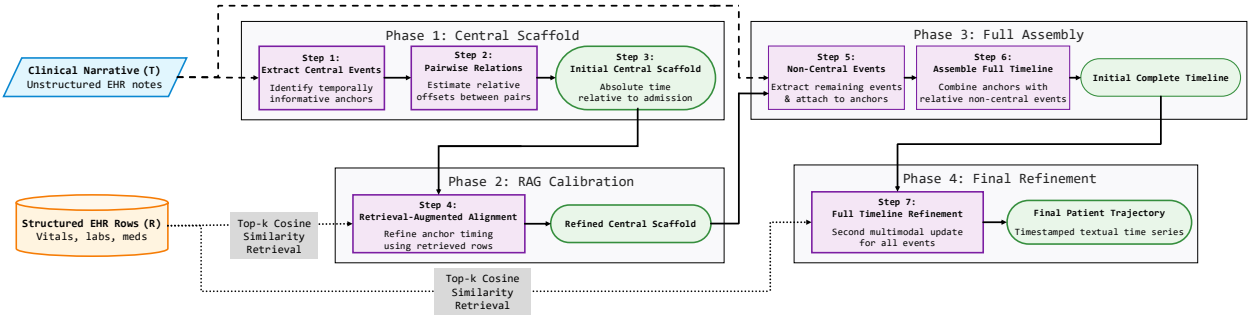


Fig. 2. Overview of the scaffolded retrieval-augmented timeline reconstruction workflow. The method first extracts central events from the clinical narrative, estimates pairwise temporal relations, and builds an initial central-event scaffold. Retrieved structured EHR rows are then used to calibrate this scaffold. Non-central events are extracted and attached relative to central events, after which a final retrieval-augmented refinement step calibrates the complete timeline.

both event recovery and time assignment difficult.[9,10] Together, these findings motivate more explicit alignment between text-based and tabular approaches, with the goal of improving temporal precision while preserving the richer event content available in narrative text.[11,12]

In this work, we investigate whether aligning retrieved structured EHR evidence with text-derived event timelines can improve the temporal precision of reconstructed patient trajectories. Our retrieval-augmented multimodal approach integrates structured rows directly into timeline reconstruction rather than treating textual and tabular data as separate views. The multistep pipeline first identifies temporally informative central events and uses them to create an initial scaffold of the patient trajectory. It then positions additional events relative to this scaffold and refines the resulting timeline using retrieved structured evidence. This decomposition reflects a clinically natural distinction between events that form the backbone of an encounter and those whose timing is best interpreted relative to that backbone. Using this framework, we assess whether and where multimodal alignment improves on text-only timeline reconstruction. Specifically, we examine the effects of introducing structured evidence at different stages of the reconstruction workflow, determine whether improvements are concentrated in temporal metrics rather than event recovery, and analyze which prognostically relevant information appears in text but is missing or delayed in tabular data.

Contributions. In summary, the paper makes the following contributions. **(i) Temporal grounding formulation:** We formulate clinical timeline reconstruction as retrieval-augmented temporal grounding of narrative-derived patient trajectories using structured EHR evidence. **(ii) Multimodal scaffolded reconstruction workflow:** We introduce a workflow that uses anchor events as a temporal backbone and retrieved structured rows as partial evidence for timestamp calibration. **(iii) Multimodal timeline evaluation:** We evaluate multimodal calibration on 40 discharge summaries with aligned structured EHR rows and manual gold-standard timelines, testing whether structured evidence improves temporal localization more than event recovery. **(iv) Information missing from tabular data:** We provide an empirical analysis of clinically relevant events that are absent, delayed, or semantically compressed in structured tabular data, motivating why unstructured text is necessary for constructing clinically meaningful patient trajectories.

2. Related Work

Clinical temporal information extraction. Clinical timeline reconstruction builds on a long tradition of temporal reasoning and temporal information extraction in medicine.[13,14] Early annotation schemes and shared tasks, including i2b2 and THYME, focused on identifying clinical events, time expressions, and temporal relations such as before, after, overlap, and containment.[15–18] Systems such as MedTime and cTAKES-temporal operationalized these ideas for clinical notes by extracting time expressions and temporal relations.[19–21] Pre-LLM work also studied patient timelines and longitudinal record summarization, including timeline generation from clinical notes and HARVEST for summarizing longitudinal patient records.[22,23] These studies established the importance of temporal structure in clinical narratives, but primarily focused on temporal annotation, relation extraction, or summarization rather than grounding narrative-derived absolute timelines in structured EHR evidence.

LLM-based clinical timeline reconstruction. Recent work has moved from relation-centered temporal extraction toward absolute event timelines. Leeuwenberg and Moens studied absolute event timeline extraction from English clinical reports, while Frattallone-Llado et al. showed that aligning discharge summaries with structured EHR data can improve inpatient event timeline precision.[8,24] Textual time-series work has further shown that clinical narratives can be converted into event–time tuples for temporal forecasting, but these approaches are primarily text-only and do not use structured EHR rows as calibration evidence during reconstruction.[9,10,25,26] Recent shared tasks such as ChemoTimelines have renewed interest in patient-level clinical timeline extraction using prompting, fine-tuning, and dictionary-enhanced LLM systems.[27,28] Our work differs from these efforts by treating timeline reconstruction as multimodal temporal grounding: narrative text supplies event semantics, while structured EHR rows are retrieved as partial temporal evidence for timestamp calibration.

Asymmetric evidence in structured and narrative EHR data. Clinical machine learning studies have shown that structured and unstructured EHR data encode complementary information for phenotyping, prediction, and patient representation.[11,12,29] Structured data provide dense timestamped measurements, but often omit symptoms, functional status, patient-reported histories, and clinical reasoning. Narrative notes preserve this context, but are often retrospective and not organized as chronological trajectories. This asymmetry makes clinical timeline reconstruction different from standard multimodal fusion: the goal is not to merge two complete views of the same event sequence, but to use structured rows as partial temporal evidence for a narrative-derived patient trajectory. Our retrieval-augmented scaffolded workflow is designed for this setting, selecting candidate structured rows for each text-derived event and updating timestamps only when the retrieved evidence is clinically relevant.

3. Methods

3.1. *Clinical timeline task and dataset*

Given a clinical narrative T and associated structured EHR rows R, the goal is to reconstruct a textual time series $\hat{S} = \{(e_i, \hat{t}_i)\}_{i=1}^{N}$, where each e_i is a standalone clinical event and each $\hat{t}_i \in \mathbb{R}$ is its predicted timestamp in hours relative to a case-specific reference time. When hospital admission is explicitly described, we use admission as $t = 0$; otherwise, we use the earliest

documented clinical encounter or presentation in the narrative. Events before the reference time receive negative timestamps, and events after it receive positive timestamps.

A clinical event is any patient-specific health-related mention that is semantically interpretable on its own and temporally locatable, including symptoms and signs, diagnoses, tests, procedures, treatments, medication administrations, major clinical states, outcomes, pertinent negatives, and termination or discontinuation events. Our event representation follows prior textual time-series work,[9,25,26,30] but differs from short-span, concept-centric schemes such as i2b2.[15,16] We represent events as standalone clinical findings: spans may include context needed for clinical specificity, conjunctive mentions are split when doing so improves clarity, and modifiers such as negation, uncertainty, severity, laterality, and intent are preserved when they affect event interpretation or temporal placement.

Our evaluation uses **40 discharge summaries** with manual gold-standard timelines and aligned structured EHR rows: **15 i2b2-derived discharge summaries** and **25 MIMIC-IV discharge summaries**. All cases have structured EHR rows aligned to the same hospital encounter. MIMIC-style critical-care records are well suited for this setting because they pair narrative clinical documentation with timestamped structured data, including laboratory values, medication administrations, vital signs, procedures, diagnoses, and administrative events.[6,7] For each case, the narrative note provides the source text for event reconstruction, while the structured record provides external temporal evidence. We convert structured-row timestamps to hours relative to the same case-specific reference time used for the narrative timeline. Each structured row $r_m \in R$ is represented as $r_m = (\nu_m, x_m, s_m)$, where ν_m is the row name or clinical feature, x_m is the observed value when available, and s_m is the timestamp in hours relative to the reference time.

The 15 i2b2-derived cases were randomly sampled from eligible cases; the 25 MIMIC-IV cases were sampled at the patient level, using the first discharge summary for five randomly sampled patients and one randomly selected summary for 20 others, without enrichment for diagnosis, temporal richness, or encounter complexity. Three medical experts annotated disjoint subsets; every case was annotated once from the note and aligned structured record, so no adjudication or inter-annotator agreement was available (Supplementary Material §S1).

We treat structured rows as temporal evidence rather than gold-standard timeline events. Structured EHR tables contain dense repeated measurements, medication records, and machine-generated entries, many of which are not standalone narrative events. Conversely, clinically meaningful narrative events, such as symptoms, functional status, patient-reported histories, and clinical reasoning, may not have direct structured counterparts. Our method therefore uses structured rows to calibrate text-derived event times when clinically relevant, rather than forcing all structured rows into the reconstructed timeline.

3.2. *Scaffolded retrieval-augmented multimodal timeline reconstruction*

Our workflow is designed around three principles. First, it anchors before expanding: the model identifies temporally informative anchor events before reconstructing the full event set. Second, it retrieves candidate structured evidence rather than forcing every narrative event to align to a structured row. Third, it uses preserve-by-default calibration: event times are updated only when retrieved structured evidence is clinically relevant and temporally plausible. The

workflow proceeds in seven clinically interpretable steps (Figure 2).[a]

(1) **Step 1: Extract anchor events.** The model first identifies temporally informative anchor events, also referred to as central events, that define the encounter backbone or serve as temporal referents for other events. Examples include admission, symptom onset, treatment initiation, clinical deterioration, procedures, transfers, and discharge.
(2) **Step 2: Estimate timing among anchors.** The model estimates pairwise time differences (in hours) between selected anchor events.
(3) **Step 3: Build an initial anchor timeline.** The pairwise relations between anchor events are resolved into a provisional timeline backbone in hours relative to the case reference time. When estimates conflict, the model uses higher-confidence relations and the source narrative to construct the initial scaffold.
(4) **Step 4: Calibrate anchors using structured evidence.** For each anchor event, retrieved structured EHR rows are provided to the calibration model together with the narrative and provisional scaffold. The model updates an anchor time only when the structured row clinically supports the change; otherwise, the text-derived estimate is retained.
(5) **Step 5: Extract non-anchor supporting events relative to anchors.** After anchor calibration, the model extracts remaining non-anchor events and attaches each event to a referent anchor with a relative time offset and confidence score. These events include symptoms, exam findings, pertinent negatives, severity descriptors, functional status, medication details, and contextual clinical information.
(6) **Step 6: Assemble the full timeline.** The calibrated anchor times and relative supporting-event offsets are combined into an initial complete timeline, converting all events into absolute times in hours relative to the case reference time.
(7) **Step 7: Refine the complete timeline using structured evidence.** Structured rows are retrieved again for all reconstructed events, not only anchors. The model performs a final conservative calibration pass, updating timestamps when retrieved rows provide clinically consistent temporal evidence and preserving text-derived times otherwise. The output is a temporally grounded textual time series containing clinically meaningful narrative events with calibrated timestamps when structured evidence is relevant.

Retrieval and calibration module. The same retrieval module is used in Step 4 and Step 7. Structured rows are represented by concatenating the row event name and value; timestamps are excluded from the embedding text and passed separately to the calibration model. Rows with the same event name and value are merged, with timestamps represented in hours relative to the case reference time. We embed event strings and structured-row representations using MedTE (`mohammadkhodadad/MedTE-cl15-step-8000`), mean-pool and L2-normalize the embeddings, and rank rows by cosine similarity. For each event, we retrieve the top $k = 5$ structured rows with no fixed similarity threshold. The calibration model receives each row as `name | value | time | similarity`. Calibration follows a preserve-by-default rule: cosine

[a]Step-specific prompts and the full reproducibility pipeline are available in our GitHub repository: `https://github.com/SayantanKumar/multimodal_textual_tabular_alignment`

similarity ranks candidates, irrelevant rows are disregarded, the initial time is retained when no relevant evidence justifies a change, and narrative context resolves repeated timestamps or conflicting rows. The complete decision rule appears in Supplementary Material §S2. We also evaluated $k = 3$, $k = 5$, and $k = 10$, finding negligible performance differences but increased runtime for larger k.

This design treats retrieval as evidence selection rather than forced temporal alignment. Methods such as dynamic time warping assume that two sequences share an approximately comparable event universe and can be aligned through a monotone path. Clinical narratives and structured EHR rows often violate this assumption: many narrative events have no structured counterpart, while structured tables contain dense repeated measurements, medication records, and charted observations that are not standalone narrative events. Our workflow therefore allows structured evidence to calibrate event times when appropriate without requiring every narrative event to have a tabular match.

3.3. *Experimental setup*

Baselines and ablations. We compare the default scaffolded multimodal workflow against text-only, single-step, stage-specific, and forced-alignment baselines. **Text-only reconstruction** uses the same scaffolded workflow but removes both structured-EHR calibration steps, testing whether retrieved structured evidence improves temporal localization beyond narrative-only reconstruction. **Single-step multimodal reconstruction** directly generates a complete timeline from the narrative and retrieved structured rows without constructing an anchor-event scaffold, testing whether the staged anchor/non-anchor decomposition is useful beyond simply adding retrieved evidence to the prompt. To isolate where structured evidence contributes, we compare two stage-specific ablations: **anchor-only calibration**, which calibrates the anchor scaffold but omits final full-timeline calibration, and **final-only calibration**, which reconstructs the full timeline from text before applying structured-EHR calibration. The default workflow applies both anchor-stage and final-stage calibration. Finally, **dynamic time warping (DTW)** serves as a forced temporal-alignment baseline that attempts to align text-derived events with structured rows through sequence alignment rather than retrieval-based evidence selection followed by clinical calibration.

Evaluation metrics. We evaluate predicted timelines along two complementary dimensions: *event recovery* and *temporal localization*, adapted from prior textual time-series work.[9,25] Because predicted and reference timelines may differ in length and event wording, we first align predicted events to manually annotated reference events using a recursive best-match procedure. This produces a one-to-one alignment between predicted and reference events while accommodating unequal timeline lengths. **Event match rate** is the proportion of reference clinical events successfully aligned to a predicted event. Temporal metrics are computed on the matched subset. **Temporal concordance** (c-index) measures whether the predicted timeline preserves the relative ordering of events in the manual reference. **Area Under the Log-Time Cumulative Distribution Function (AULTC)** summarizes absolute timestamp discrepancy on a log-transformed time scale, with higher values indicating smaller timestamp errors across clinically meaningful ranges. Because concordance and AULTC are computed only over matched events, they should be interpreted jointly with event match rate. We estimate un-

Table 1. Performance on 40 discharge summaries at event-matching threshold 0.1 for unimodal and multimodal timeline reconstruction. Values are point estimates with bootstrap 95% confidence intervals in brackets. **Bold** indicates the best point estimate within each metric–modality column.

Model	Event match rate		Concordance		AULTC	
	Unimodal	Multimodal	Unimodal	Multimodal	Unimodal	Multimodal
DeepSeek R1	0.39 [0.32, 0.47]	0.39 [0.32, 0.47]	0.75 [0.69, 0.79]	0.75 [0.70, 0.80]	**0.80 [0.76, 0.82]**	**0.79 [0.77, 0.82]**
KimiK2-Instruct	0.26 [0.21, 0.29]	0.25 [0.22, 0.30]	0.73 [0.69, 0.76]	0.74 [0.68, 0.81]	0.68 [0.61, 0.73]	0.68 [0.62, 0.73]
GPT-OSS-120B	0.32 [0.27, 0.37]	0.32 [0.28, 0.37]	0.74 [0.70, 0.77]	0.75 [0.71, 0.80]	0.76 [0.72, 0.80]	0.77 [0.74, 0.81]
DeepSeek V3.2	**0.49 [0.43, 0.56]**	**0.49 [0.43, 0.56]**	0.74 [0.70, 0.80]	0.77 [0.73, 0.81]	0.71 [0.64, 0.78]	0.77 [0.73, 0.81]
GLM5	0.33 [0.27, 0.39]	0.33 [0.27, 0.40]	**0.80 [0.78, 0.83]**	**0.82 [0.75, 0.84]**	0.78 [0.73, 0.82]	0.78 [0.73, 0.82]
Qwen3.5-397B-A17B	0.38 [0.28, 0.47]	0.40 [0.29, 0.51]	0.78 [0.72, 0.81]	0.79 [0.78, 0.82]	0.74 [0.68, 0.77]	0.77 [0.74, 0.78]

certainty using case-level bootstrap confidence intervals. Each bootstrap sample resamples discharge summaries with replacement and recomputes event matching and timeline metrics, preserving within-case dependence among events. We report 95% confidence intervals for the reconstruction metrics using 200 bootstrap resamples. For direct multimodal–unimodal comparisons, we additionally used 10,000 paired case-level resamples and report confidence intervals for the mean within-case differences in Supplementary Material §S3, Table S1. All implementation and reproducibility details are presented in Supplementary Material §S8.

4. Results

4.1. *Structured EHR evidence calibrates event timing while preserving event recovery*

Table 1 summarizes timeline reconstruction performance for text-only reconstruction and multimodal reconstruction with retrieved structured EHR evidence. Across models, multimodal calibration has little effect on event match rate, indicating that structured rows do not primarily expand the set of events recovered from narrative text. Instead, structured evidence primarily affects temporal localization, with effects that vary by model and metric. These results support the role of structured rows as temporal evidence for text-derived events, rather than as a substitute event source.

Mean paired case-level temporal differences generally favored multimodal calibration: 4 of 6 models showed higher concordance and 3 of 6 showed higher AULTC, while no model showed lower performance on either temporal metric. However, uncertainty was substantial: no paired confidence interval for the multimodal–unimodal concordance difference excluded zero, and only DeepSeek V3.2 showed an AULTC interval entirely above zero (Supplementary Material §S3, Table S1). Performance was generally higher on the i2b2-derived subset, but modality effects were source-dependent (Supplementary Material §S5, Tables S3–S4).

Because event match rate depends on the semantic matching threshold used to align predicted and reference events, we also vary the event-matching threshold and examine the relationship between event recovery and temporal quality. Figure 3 shows AULTC and concordance as functions of event match rate across thresholds. Across operating points, multimodal reconstruction generally occupies a similar event-recovery range as text-only reconstruction while changing temporal performance, especially AULTC, for several models. This supports the same broad conclusion as Table 1: retrieved structured rows primarily affect temporal calibration rather than event discovery.

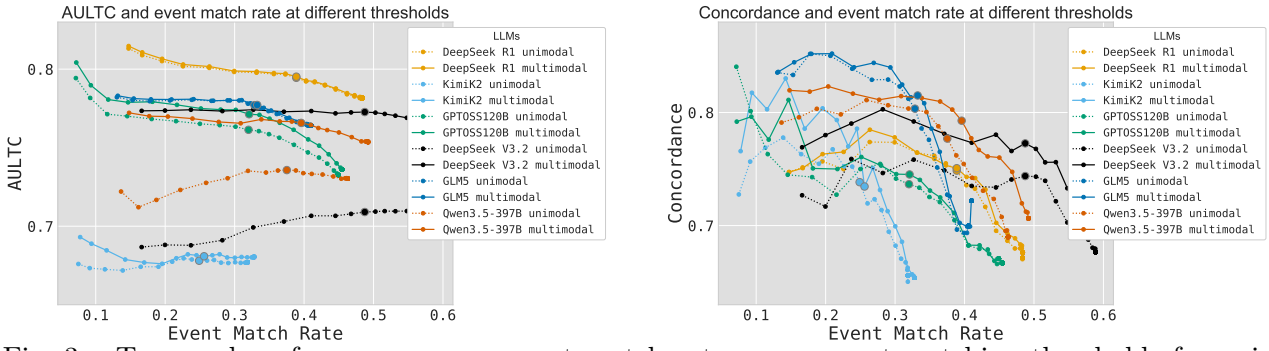


Fig. 3. Temporal performance versus event match rate across event-matching thresholds for unimodal and multimodal variants of each model on the 40-summary evaluation set. (**Left**) AULTC versus event match rate. (**Right**) Temporal concordance versus event match rate.

Scaffold ablations. Table 2A evaluates whether the staged scaffolded workflow is necessary, or whether similar performance can be obtained by generating a complete timeline in one step with retrieved structured rows. The single-step formulation performs worse than the full scaffolded variant, suggesting that clinical timeline reconstruction benefits from first identifying temporally informative anchor events before reconstructing the full event set. The stage at which structured evidence is introduced also matters. Final-stage calibration yields the highest event match rate, but the strongest temporal metrics are obtained when structured evidence is used both to calibrate the anchor scaffold and to refine the final assembled timeline. This supports the design choice of using structured rows first to stabilize the temporal backbone of the encounter and then to recalibrate the expanded trajectory.

Table 2. **(A)** DeepSeek V3.2 scaffold and calibration-stage ablations with bootstrap 95% confidence intervals. **(B)** Retrieval calibration compared with sentence-embedding and lexical Dynamic Time Warping (DTW) alignments.

A. Scaffold and calibration-stage ablations

Ablation	Event match rate		Concordance		AULTC	
	Unimodal	Multimodal	Unimodal	Multimodal	Unimodal	Multimodal
single-step	0.39 [0.35, 0.45]	0.39 [0.35, 0.44]	0.70 [0.65, 0.74]	0.71 [0.65, 0.75]	**0.78 [0.74, 0.83]**	**0.77 [0.73, 0.83]**
central only	–	0.46 [0.42, 0.51]	–	0.72 [0.66, 0.77]	–	0.75 [0.71, 0.80]
final only	**0.51 [0.46, 0.56]**	**0.52 [0.46, 0.58]**	0.68 [0.64, 0.72]	0.69 [0.64, 0.72]	0.72 [0.67, 0.78]	0.73 [0.69, 0.78]
central + final	0.49 [0.43, 0.56]	0.49 [0.43, 0.56]	**0.74 [0.70, 0.80]**	**0.77 [0.73, 0.81]**	0.71 [0.64, 0.78]	**0.77 [0.73, 0.81]**

Note: Central-only, final-only, and central + final correspond to updating the anchor scaffold only, the final timeline only, or both stages, respectively. The central-only condition yields a single final timeline and is undefined for the unimodal setting.

B. Retrieval calibration vs. explicit temporal alignment (Dynamic Time Warping)

Model	Event match rate			Concordance			AULTC		
	Retrieval	DTW(Sent)	DTW(Lex)	Retrieval	DTW(Sent)	DTW(Lex)	Retrieval	DTW(Sent)	DTW(Lex)
DeepSeek R1	0.39	0.39	0.39	0.75	0.75	0.74	0.79	0.80	0.79
KimiK2-Instruct	0.25	0.26	0.26	0.74	0.74	0.73	0.68	0.68	0.68
GPT-OSS-120B	0.32	0.32	0.32	0.75	0.73	0.73	0.77	0.77	0.77
DeepSeek V3.2	0.49	0.49	0.49	0.77	0.76	0.77	0.77	0.77	0.77
GLM5	0.33	0.33	0.33	0.82	0.81	0.81	0.78	0.78	0.77
Qwen3.5-397B-A17B	0.40	0.40	0.40	0.79	0.79	0.79	0.77	0.77	0.77

Explicit temporal alignment. We also compare retrieval-based calibration with dynamic time warping (DTW), an explicit temporal-alignment baseline that aligns text-derived events with structured rows as sequences (Table 2B). The two approaches achieve broadly comparable accuracy: DTW is slightly higher for some model–metric pairs, while retrieval-based calibration is slightly higher for others, and the differences are small relative to the variability reflected by the bootstrap confidence intervals in Table 1. Retrieval and DTW therefore show

Table 3. Structured-data gap analysis. Each text-derived event is compared with structured EHR counterparts from the same encounter.

Category	Definition	Count (%)
Well captured	Semantically adequate structured counterpart within the temporal window	2,827 (33.5%)
Complete absence	No structured counterpart identified	2,963 (35.1%)
Temporal mismatch	Semantically related structured row exists but occurs outside the temporal window	1,463 (17.3%)
Semantic distance	Retrieved row is not clinically equivalent to the narrative event	801 (9.5%)
Detail gap	Structured row omits clinically important attributes such as severity, laterality, duration, progression, or functional status	385 (4.6%)
Total textual events	–	8,439

comparable aggregate accuracy rather than quantitative superiority. Their principal distinction is structural: retrieval presents event-specific candidates and can retain a text-derived estimate when no adequate counterpart exists, whereas DTW imposes monotone sequence alignment across modalities with incomplete and asymmetric event coverage.

4.2. *Complementary coverage of narrative and structured records*

The main reconstruction results show that structured rows can improve event timing when used as calibration evidence. We next ask why structured records cannot simply replace narrative text and then characterize which narrative-derived events are poorly represented in structured EHR data; because this text-to-structure analysis does not assess structured-only events, it supports complementarity rather than completeness of either modality. We compare text-derived timelines against tabular records from the same encounters. The analysis uses the best-performing LLM-generated event timelines as the textual view, MIMIC structured data as the tabular view, and admission or the case reference time as the shared temporal anchor. For each text-derived event, we retrieve a candidate tabular counterpart using embedding-based similarity matching and evaluate it along two dimensions: semantic adequacy and temporal alignment. Semantic adequacy is assessed using retrieval-augmented scoring on a 0–1 scale, while temporal alignment is measured as the absolute difference in hours between the text-derived timestamp and the matched tabular timestamp. Based on these criteria, each event is classified as **well captured**, **complete absence**, **temporal mismatch**, **semantic distance**, or **detail gap**. Table 3 summarizes the distribution of these categories.

Across 8,439 text-derived events, only 33.5% are well captured by structured data, while 35.1% have no structured counterpart at all. When a structured counterpart exists, its timestamp is often useful: the median temporal discrepancy is 4.3 hours, with 58.4% of matched events within 6 hours, 69.5% within 12 hours, and 79.3% within 24 hours. The missing events are not merely low-value narrative detail; they include symptoms, functional status, temporal qualifiers, patient-reported information, severity descriptions, and clinical context. Representative complete-absence examples include **symptom descriptions** such as "substernal chest pain for 1 day" and "developed chest pain"; **denial statements** such as "no seizure-like activity" and "no deformity"; **temporal patterns** such as "pain resolved" and "improving clots on LE"; **severity descriptors** such as "EF 55%" and "non-small cell lung cancer stage IIIA"; and **causal or procedural context** such as "intubated for several days after second surgery" and "repair of leak from pouch." These findings explain why structured rows are useful for temporal calibration but incomplete as representations of patient trajectories.

4.3. *Qualitative case study: text supplies phenotype, tables calibrate time*

Figure 1 illustrates the qualitative behavior of the pipeline using a synthetic stroke-alert case. The narrative supplies clinically decisive phenotype details, including lobar hemorrhage

location, hematoma volume, aphasia, plegia, command following, absence of intraventricular hemorrhage or hydrocephalus, and repeat-CT stability. The structured rows provide sparse but precise temporal anchors, including admission time and IV andexanet alfa administration. The text-only scaffold recovers the reversal event but places it shortly after the initial CT at approximately +1.5h; retrieval-augmented calibration links the event to the medication-administration row and shifts reversal to +6.5h relative to admission.

Pipeline output: calibrated event timeline

Event	Time (h)	Source / calibration note
Last known well at 0730	-5.5	Text clock time relative to admission
Found down at home	-1.0	Text-derived estimate
67 years old	0.0	Text, admission context
Atrial fibrillation	0.0	Text, admission context
Admitted/brought in as stroke alert	0.0	**Table**, anchored by admission row at 1300
Dense right arm plegia	0.0	Text, anchored to admission
Follows commands	0.0	Text, anchored to admission
Initial CT showed intracerebral hemorrhage	0.5	Text-derived estimate
Estimated hematoma volume ∼55 mL	0.5	Text, anchored to initial CT
Left frontal lobar location of hemorrhage	0.5	Text, anchored to initial CT
Reversal agent was administered	6.5	**Table**, corrected by IV andexanet alfa at 1930
Repeat CT was stable	8.5	Text, inferred sequentially post-reversal

This case illustrates the asymmetry motivating the workflow. Text is necessary to recover clinically meaningful phenotype information, while structured rows calibrate selected time-sensitive events when relevant counterparts exist. The resulting trajectory could support human-reviewed cohort definition or eligibility logic by preserving narrative-only clinical details while grounding treatment timing in structured evidence. For example, the 2024 ENRICH trial comparing surgical evacuation versus medical management had eligibility criteria beyond what would be available in structured data alone.[31] Structured data could show admission indication (hemorrhagic stroke), medication events (a confirmation proxy), and functional status scores (neurologic and disability level), but the notes contain key additional ENRICH eligibility features and timing: 55 mL hematoma volume ($30 < x < 80$), left frontal lobar location (lobar/anterior basal ganglionic), aphasia/plegia (new onset of neurological symptoms), no IVH/hydrocephalus and stable repeat imaging (feasibility), as well as localization of the eligibility timing window (within 24 hours of last known well time).

5. Discussion and Conclusion

Our results suggest that the primary benefit of multimodal alignment is temporal calibration rather than event discovery. Across models, structured evidence changed event recovery little, while mean paired case-level temporal differences generally favored multimodal calibration: 7 of 12 model–metric comparisons across concordance and AULTC were positive and none were negative. However, uncertainty in the magnitude of these improvements was substantial given the 40-case sample; paired confidence intervals excluded zero only for the DeepSeek V3.2 AULTC improvement. Source-stratified estimates also differed between i2b2-derived and MIMIC-IV cases (Supplementary Material §S3 and §S5). Within the text-primary formulation, narrative text defines the candidate event set and structured rows provide partial temporal evidence; the results therefore establish the value of narrative content for events represented in the note, not that narrative is the primary source of all clinically meaningful events.

The structured-data gap analysis explains why this distinction matters. Structured EHR data is not merely a noisier version of the narrative record; it offers a different view of the patient trajectory. When relevant counterparts exist, their timestamps can anchor timeline calibration. However, many clinically meaningful narrative events are absent, delayed, or semantically compressed in tabular form, including symptoms, functional status, temporal qualifiers, patient-reported information, and clinical context. These findings characterize asymmetric coverage within the present text-primary analysis; they do not assess structured-only events or establish that either source is complete.

The ablation results also support the scaffolded design. One-shot timeline generation requires a model to identify clinical events, infer relative timings, resolve vague temporal expressions, and decide how to use structured evidence in a single step. By first constructing an anchor-event scaffold, the workflow separates the temporal backbone of the encounter from supporting narrative details. This mirrors clinical reasoning, where events such as admission, diagnostic testing, treatment initiation, deterioration, procedures, and discharge provide temporal reference points for interpreting the rest of the patient course. Staged calibration then allows structured evidence to stabilize the backbone before refining the full trajectory.

Limitations: This study has several limitations. First, event match rates remain modest, so concordance and AULTC measure temporal quality conditional on successful event recovery rather than accuracy over the complete trajectory. Matched events were clinically heterogeneous, but unmatched events more often contained later-course and granular numerical documentation (Supplementary Material §S4, Table S2). Second, the 40-case sample provides only high-uncertainty evidence about stability across data sources and clinical settings; moreover, each case was annotated once by one of three medically trained annotators, without duplicate annotation or inter-annotator agreement. Third, the framework emits one continuous timestamp per event; when timing is vague or absent, it infers a point estimate from context, and the model-generated certainty flag is uncalibrated (Supplementary Material §S6, Table S5). Fourth, the text-primary design may omit clinically meaningful structured-only events. Structured timestamps also have source-specific semantics and may represent collection, ordering, administration, result availability, device recording, or documentation rather than clinical onset. Finally, the seven-call workflow has substantially higher inference cost than the single-step baseline.

Practical use and clinical impact: This framework is not intended for autonomous clinical decision-making, but for temporally grounded research and clinician-reviewed workflows. Reconstructed trajectories could support retrospective cohort construction, treatment-window definition, temporal phenotyping, leakage auditing, and downstream trajectory modeling. Future work should evaluate larger, independently annotated cohorts; support interval-valued timestamps, qualitative relations, and explicit time-unknown outputs; develop bidirectional event discovery that preserves source provenance and timestamp semantics; and test stage-specific cost-quality tradeoffs. Overall, the findings support a retrieval-augmented view in which narrative text supplies candidate-event semantics, structured rows provide partial and source-dependent temporal evidence, and preserve-by-default calibration can improve temporal grounding for some events when relevant evidence is available.

Acknowledgements

This research was supported by the Intramural Research Program of the National Institutes of Health (NIH) and utilized the computational resources of the NIH HPC Biowulf cluster. The contributions of the NIH author(s) are considered Works of the United States Government. S.N. was supported in part by the Carnegie Mellon University TCS Presidential Fellowship, and the Natural Sciences and Engineering Research Council of Canada (NSERC) Canada Graduate Research Scholarship — Doctoral (CGRS D) Fellowship. S.N. was also supported in part by an appointment to the National Library of Medicine Research Program administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy (DOE) and the National Library of Medicine, National Institutes of Health. ORISE is managed by ORAU under DOE contract number DE-SC0014664. The findings and conclusions presented in this paper are those of the authors and do not necessarily reflect the views of the NIH, NLM, DOE, ORAU/ORISE or the U.S. Department of Health and Human Services.

References

1. E. Kyriazopoulou, L. Liaskou-Antoniou, G. Adamis, A. Panagaki, N. Melachroinopoulos, E. Drakou, K. Marousis, G. Chrysos, A. Spyrou, N. Alexiou *et al.*, Procalcitonin to reduce long-term infection-associated adverse events in sepsis. a randomized trial, *American Journal of Respiratory and Critical Care Medicine* **203**, 202 (2021).
2. C. W. Seymour, J. N. Kennedy, S. Wang, C.-C. H. Chang, C. F. Elliott, Z. Xu, S. Berry, G. Clermont, G. Cooper, H. Gomez *et al.*, Derivation, validation, and potential treatment implications of novel clinical phenotypes for sepsis, *JAMA* **321**, 2003 (2019).
3. K. E. Henry, R. Adams, C. Parent, H. Soleimani, A. Sridharan, L. Johnson, D. N. Hager, S. E. Cosgrove, A. Markowski, E. Y. Klein *et al.*, Factors driving provider adoption of the trews machine learning-based early warning system and its effects on sepsis treatment timing, *Nature Medicine* **28**, 1447 (2022).
4. F. Kamran, D. Tjandra, A. Heiler, J. Virzi, K. Singh, J. E. King, T. S. Valley and J. Wiens, Evaluation of sepsis prediction models before onset of treatment, *NEJM AI* **1** (2024).
5. S. Noroozizadeh, J. C. Weiss and G. H. Chen, Temporal supervised contrastive learning for modeling patient risk progression, *Proceedings of Machine Learning Research* **225**, 403 (2023).
6. A. E. Johnson, T. J. Pollard, L. Shen, L.-w. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Anthony Celi and R. G. Mark, MIMIC-III, a freely accessible critical care database, *Scientific data* **3**, 1 (2016).
7. A. Johnson, T. Pollard, S. Horng, L. A. Celi and R. Mark, MIMIC-IV-Note: Deidentified free-text clinical notes (version 2.2) (2023).
8. G. Frattallone-Llado, J. Kim, C. Cheng, D. Salazar, S. Edakalavan and J. C. Weiss, Using multimodal data to improve precision of inpatient event timelines, *Advances in Knowledge Discovery and Data Mining* **14648**, 322 (2024).
9. S. Noroozizadeh and J. C. Weiss, Reconstructing sepsis trajectories from clinical case reports using llms: The textual time series corpus for sepsis, *Proceedings of Machine Learning Research* **333**, 10 (2026).
10. S. Noroozizadeh, S. Kumar and J. C. Weiss, Forecasting clinical risk from textual time series: Structuring narratives for temporal ai in healthcare, *Proceedings of the AAAI Conference on Artificial Intelligence* **40**, 39080 (2026).
11. A. Moldwin, D. Demner-Fushman and T. R. Goodwin, Empirical findings on the role of struc-

tured data, unstructured data, and their combination for automatic clinical phenotyping, *AMIA Joint Summits on Translational Science Proceedings* **2021**, 445 (2021).
12. T. M. Seinen, J. A. Kors, E. M. van Mulligen and P. R. Rijnbeek, Using structured codes and free-text notes to measure information complementarity in electronic health records: Feasibility and validation study, *Journal of Medical Internet Research* **27**, p. e66910 (2025).
13. I. S. Kohane, Temporal reasoning in medical expert systems, in *MEDINFO 86: Proceedings of the Fifth World Congress on Medical Informatics*, eds. R. Salamon, B. Blum and M. Jørgensen (North-Holland, Amsterdam, 1986).
14. L. Zhou and G. Hripcsak, Temporal reasoning with medical data—a review with emphasis on medical natural language processing, *Journal of Biomedical Informatics* **40**, 183 (Apr 2007).
15. Ö. Uzuner, B. R. South, S. Shen and S. L. DuVall, 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text, *Journal of the American Medical Informatics Association* **18**, 552 (2011).
16. W. Sun, A. Rumshisky and O. Uzuner, Evaluating temporal relations in clinical text: 2012 i2b2 challenge, *Journal of the American Medical Informatics Association* **20**, 806 (2013).
17. W. Sun, A. Rumshisky and O. Uzuner, Annotating temporal information in clinical narratives, *Journal of biomedical informatics* **46**, S5 (2013).
18. I. Styler, William F., S. Bethard, S. Finan, M. Palmer, S. Pradhan, P. C. de Groen, B. Erickson, T. Miller, C. Lin, G. Savova and J. Pustejovsky, Temporal annotation in the clinical domain, *Transactions of the Association for Computational Linguistics* **2**, 143 (2014).
19. Y.-K. Lin, H. Chen and R. A. Brown, Medtime: A temporal information extraction system for clinical narratives, *Journal of biomedical informatics* **46**, S20 (2013).
20. G. K. Savova, J. J. Masanz, P. V. Ogren, J. Zheng, S. Sohn, K. C. Kipper-Schuler and C. G. Chute, Mayo clinical text analysis and knowledge extraction system (cTAKES): architecture, component evaluation and applications, *Journal of the American Medical Informatics Association* **17**, 507 (2010).
21. C. Lin, D. Dligach, T. A. Miller, S. Bethard and G. K. Savova, Multilayered temporal modeling for the clinical domain, *Journal of the American Medical Informatics Association* **23**, 387 (2016).
22. A. Nikfarjam, E. Emadzadeh and G. Gonzalez, Towards generating a patient's timeline: extracting temporal relationships from clinical notes, *Journal of biomedical informatics* **46**, S40 (2013).
23. J. S. Hirsch, J. S. Tanenbaum, S. Lipsky Gorman, C. Liu, E. Schmitz, D. Hashorva, A. Ervits, D. Vawdrey, M. Sturm and N. Elhadad, HARVEST, a longitudinal patient record summarizer, *Journal of the American Medical Informatics Association* **22**, 263 (2015).
24. A. Leeuwenberg and M.-F. Moens, Towards extracting absolute event timelines from english clinical reports, *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **28**, 2710 (2020).
25. J. Wang and J. C. Weiss, A large-language model framework for relative timeline extraction from pubmed case reports (2025).
26. S. Kumar and J. C. Weiss, Temporally phenotyping glp-1ra case reports with large language models: A textual time series corpus and risk modeling, *medRxiv* , 2026 (2027).
27. J. Yao, H. Hochheiser, W. Yoon, E. T. Goldner and G. K. Savova, Overview of the 2025 shared task on chemotherapy treatment timeline extraction, in *Proceedings of the 7th Clinical Natural Language Processing Workshop*, eds. A. Ben Abacha, S. Bethard, D. Bitterman, T. Naumann and K. Roberts (Association for Computational Linguistics, Virtual, oct 2025).
28. T. M. Zhang, Z. Sun, S. Zeng, C. Li, N. F. Abernethy, B. D. Lam, F. Xia and M. Yetisgen, UW-BioNLP at ChemoTimelines 2025: Thinking, fine-tuning, and dictionary-enhanced LLM systems for chemotherapy timeline extraction, in *Proceedings of the 7th Clinical Natural Language Processing Workshop*, eds. A. Ben Abacha, S. Bethard, D. Bitterman, T. Naumann and K. Roberts

(Association for Computational Linguistics, Virtual, oct 2025).
29. S. Liu, X. Wang and et al., Multimodal data matters: language model pre-training over structured and unstructured electronic health records, *IEEE Journal of Biomedical and Health Informatics* **27**, 504 (2022).
30. S. Noroozizadeh, S. Kumar, G. H. Chen and J. C. Weiss, Pmoa-tts: Introducing the pubmed open access textual times series corpus (2026).
31. G. Pradilla, J. J. Ratcliff, A. J. Hall, B. R. Saville, J. W. Allen, G. Paulon, A. McGlothlin, R. J. Lewis, M. Fitzgerald, A. F. Caveney *et al.*, Trial of early minimally invasive removal of intracerebral hemorrhage, *New England Journal of Medicine* **390**, 1277 (2024).