Characterizing Narrative Content in Web-scale LLM Pretraining Data

Teagan Johnson[♣] **Elliott Ash**[♠] **Andrew Piper**[♡] **Maria Antoniak**[♣]

[♣]University of Colorado Boulder [♠]ETH Zürich [♡]McGill University

Abstract

The narrative composition of web-scale LLM pretraining corpora remains largely unexplored, even though narrative is a fundamental mode of human communication. We present the first fine-grained study of narrative features in DOLMA, a 3-trillion-token open pretraining corpus. Drawing on narrative theory, we design a framework spanning three core narrative elements (agency, setting, and events) operationalized as 11 interpretable dimensions. After curating and hand-annotating a diverse set of ∼400 passages, we create an LLM-labeled dataset of 25K passages, and finally, we finetune and validate NARRABERT, two ROBERTA-based models for fine-grained narrative prediction. We apply NARRABERT to ∼13M passages, resulting in a new dataset, NARRADOLMA. We find that narrative structure is measurable at scale across extremely heterogeneous data and narrative qualities are unequally distributed across pretraining sources, topics, and formats in ways that current data curation practices neither measure nor account for. Our framework, dataset, and analyses provide a foundation for understanding how narrative qualities are distributed in LLM pretraining data. We publicly release NARRADOLMA and NARRABERT.

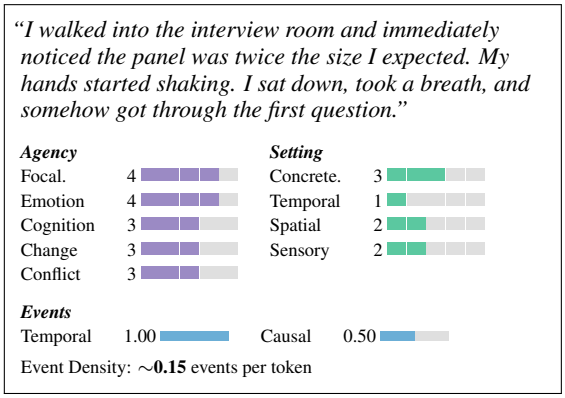

Figure 1: An example web passage scored by NarraBERT across our 12 narrative dimensions. Agency and setting dimensions are rated on a 1–5 Likert scale, temporal sequencing and causal density are passage-level proportions (0–1), and event density is the rate of event triggers per token. This passage scores high on agency and event features but low on setting, a **narrative profile** common to first-person web narratives.

1 Introduction

Narratives are among the most pervasive and cognitively central forms of human communication. From ancient oral traditions to contemporary social media, people use narrative to record experience, build shared knowledge, and reason about cause and effect (Herman, 2011; Boyd, 2009; Gottschall, 2012). Long a central subject of study in NLP (Riedl and Young, 2010), narrative tasks such as story generation have recently become one of the most popular use cases for large language models (LLMs) (Gupta et al., 2026).

The narrative capabilities of LLMs are shaped, in part, by the narrative content present in their pretraining data. However, while recent work has examined pretraining data along dimensions such as quality, toxicity, deduplication, and topic distribution (Lucy et al., 2024; Wettig et al., 2025), the narrative composition of these corpora has received almost no systematic attention. We do not know how much narrative content is present, how it is distributed across sub-corpora and genres, or how narrative-relevant features vary across the heterogeneous web text that dominates these datasets.

This gap matters for several reasons. If narrative content is unevenly concentrated in certain sub-corpora (e.g., books or Reddit), training mixtures that downweight those sources may disproportionately reduce a model's narrative exposure. Conversely, certain genres of narrative may be overrep-

https://huggingface.co/collections/teagrjohnson/narratives-in-llm-pretraining-data

https://github.com/johnsont4/narratives_in_pretraining_data_release

Findings of the Association for Computational Linguistics: EMNLP 2026, pages 48364–48404
October 24-29, 2026 ©2026 Association for Computational Linguistics

resented in ways that skew models toward particular event structures, perspectives, narrative forms, or safety violations. The well-documented creativity deficit in LLM-generated storytelling may be related to training mixtures and not only preference tuning (Tian et al., 2024; Chakrabarty et al., 2024).

In a first thorough investigation of narrative qualities of pretraining data, we map the narrative qualities of DOLMA (Soldaini et al., 2024), a 3-trillion-token open corpus spanning twelve sub-corpora. We build on the theoretical frameworks of Herman (2011) and Piper et al. (2021) and define narrative as the structured sequencing of events through an agent's perspective in a grounded setting as told by a narrator. Given the dataset's extremely heterogeneous nature, **our work treats narrative as a continuous, multi-dimensional structure** (Ochs et al., 2009), not as a binary classification (Pianzola, 2018; Antoniak et al., 2024). Our framework annotates passages along 11 narrative dimensions organized around three core elements from narrative theory: **events** (temporal ordering and causality), **agency** (focalization, emotion, cognition, change of state, and conflict), and **setting** (concreteness, temporal grounding, spatial grounding, and sensory detail). From a validated event detector, we additionally derive **event density** (events per token), giving the 12 features shown in Fig. 1.

Our contributions are as follows:

- an annotation framework grounded in narrative theory that operationalizes 11 dimensions across events, agency, and setting,
- NarraBERT, a pair of efficient classifiers that predict all 11 narrative dimensions, validated against both human and LLM annotations,
- NarraDolma, a dataset of 13M labeled passages sampled from the DOLMA corpus using a principled stratified pipeline, and
- a large-scale analysis showing that narrativity in pretraining data forms a continuous, multidimensional structure and is unequally distributed across pretraining data categories.

A recurring finding across our analysis is that pretraining sources should not be treated as narratively homogeneous. This has direct implications for data curation: decisions made at the source level (e.g., which corpora to include, at what weight) are too coarse to account for the narrative diversity present in pretraining data.

2 Related Work

2.1 Narrative Detection

The focus on predicting the narrativity of texts or passages has gained traction in NLP in recent years, especially for social media content. For example, Ganti et al. (2022) examine narrative detection in online health forums, and Doyle et al. (2024) looks at narrative detection within the context of suicide bereavement forums on Reddit. These studies perform classification within the context of specific topics and online communities.

Two new datasets have been developed to support the task of narrative detection across topics. STORYSEEKER (Antoniak et al., 2024) provides span-level binary annotations on passages from Reddit for their narrative content. NARRADETECT (Piper and Bagga, 2025) provides a large binary-labeled dataset of passages drawn from 18 different genres across books and online texts for narrative content, as well as a small, manually annotated corpus for narrativity using a five-point Likert scale across agency, eventfulness, and world-building.

We extend this work by annotating 11 fine-grained narrative dimensions on a 5-point Likert scale to provide a richer characterization of narrativity that is still computationally tractable.

2.2 Web-Scale Text Corpora and Data Curation

A growing body of work studies the composition of pretraining data (Lucy et al., 2024). Data mixing research has shown that the source proportions of pretraining corpora influence performance on a range of benchmarks (Xie et al., 2023; Rae et al., 2021; Liu et al., 2025; Wettig et al., 2025). RegMix (Liu et al., 2025) learns optimal data mixing weights by training small models on diverse mixtures and using regression to predict performance of unseen mixtures. Wettig et al. (2025) demonstrate the value of fine-grained domain construction by applying RegMix to WEBORGANIZER topic and format labels, optimizing performance on MMLU (Hendrycks et al., 2021) and HellaSwag (Zellers et al., 2019). While these efforts focus on topic, quality, and format as axes of variation, they do not consider narrative structure as a data dimension. Our work is complementary to this line of prior work by introducing *narrative qualities* to the characterization of pretraining data.

3 Our Annotation Framework

The formalization of narrative communication as a multidimensional object of analysis dates back to the work of Propp (1968) and the Russian formalists in the early twentieth century. Such work culminated in Gérard Genette's structuralist model of narrative discourse (Genette, 1983), later synthesized by Herman (2011). The "classic model" of narratology was translated into computational frameworks by Piper et al. (2021) and further refined in Piper (2023); Hamilton et al. (2026), resulting in NARRABENCH, a 50-dimensional benchmark for computational narrative understanding (Hamilton et al., 2026). Fundamental to all of these theoretical models is a tripartite framework that grounds narrative communication in three primary dimensions: agency, eventfulness, and world-building.

The theoretical literature introduces a further distinction between a categorical and scalar understanding of narrative. Classical narratology largely operated on the understanding of narrative as a binary distinction. Later theory, dating back to the work of Giora and Shen (1994), has emphasized instead the concept of "narrativity," that narrative communication exists along a spectrum of degree or intensity (Ochs et al., 2009; Pianzola, 2018; Piper and Bagga, 2022).

By operationalizing these intersecting elements as interpretable rating dimensions for human annotators, we aim to bring narrative-theoretic concepts into contact with large-scale annotation. We provide an overview of our framework here, with full details of the annotation rating scales in App. E.3.

3.1 Agency

Five dimensions capture how characters are represented as agents, rated on a 5-point Likert scale measuring centrality (1 = not central, 5 = extremely central) and drawing on Herman's account of experiential perspective in narrative (Herman, 2009a), which emphasizes that narratives convey "what it's like" to undergo events from a particular consciousness. Fig. 2 shows an example annotation.

Focalization captures the degree to which events are filtered through a specific character's perspective rather than described from an external vantage point, e.g., first-person narration with interiority, close third-person narration, or direct quotation of inner experience, and is inspired by Herman (2009a)'s "what it's like" element of narratives.

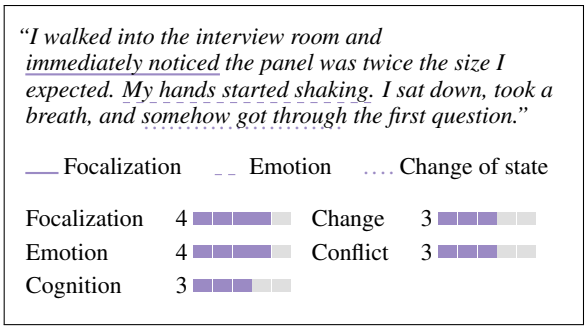

Figure 2: "I [...] immediately noticed" signals focalization, "my hands started shaking" conveys emotion through observable behavior, and "somehow got through" implies a change of state from overwhelmed to managing. Scores reflect the passage as a whole, not individual phrases.

Internal emotion captures the centrality of character's emotional states. Descriptions of emotional behavior can contribute to the score, but internal access raises it further.

Internal cognition captures the centrality of a character's thoughts, reasoning, and motivations. This feature targets cognitive and reasoning interiority rather than perceptual or emotional interiority.

Change of state captures the centrality of a change in a character's condition, encompassing physical, psychological, relational, and existential transformations. In-progress, partially implied, or world-event-entailed changes (e.g., a company going bankrupt that necessarily affects a character) all count toward the score in varying degrees.

Conflict captures the centrality of opposition or tension, encompassing interpersonal tension (character vs. character), internal struggle (character vs. self), opposition to institutions and physical environments (character vs. world).

3.2 Setting

This dimension measures how fully realized the storyworld is along four experiential dimensions, inspired by Herman's conception of narrative world-building (Herman, 2009b). Each dimension is rated on a 5-point Likert scale. Fig. 3 shows an example.

Concreteness captures how tangible and non-abstract the language is. Concreteness is not the same as specificity: a passage can name specific entities and still be abstract if those entities are not rendered perceptually. This maps onto Herman's differentiation between *existents* (things in the storyworld) and their *rendering* (Herman, 2009b), and

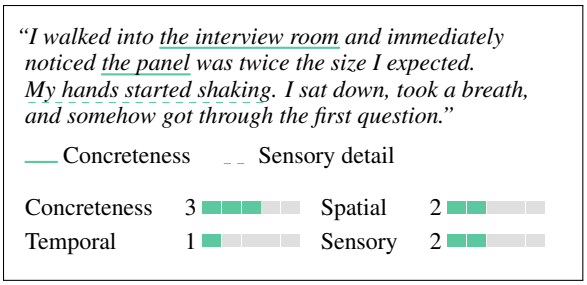

Figure 3: "The interview room" and "the panel" provide concrete referents but are named rather than rendered. "Hands started shaking" is tactile but incidental. Temporal and spatial grounding score low: this could be taking place in any room at any time.

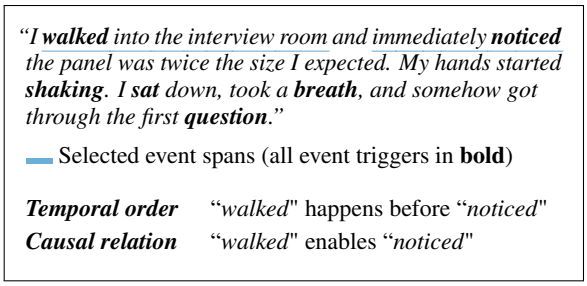

Figure 4: Two randomly selected adjacent event triggers are identified ("walked" and "noticed"). The event "walked into the room" precedes "noticed the panel". The events have a causal relationship: entering the room *enables* the character to notice the panel. Event density counts all triggers in the passage, while temporal order and causal relation describe the selected pair.

is grounded in psycholinguistic evidence linking concreteness to imageability and lexical processing (Richardson, 1975).

Temporal grounding captures how strongly the passage anchors the reader in a particular time, building on Genètte's analysis of narrative time (Genette, 1983). This manifests in two ways: *historical grounding*, which locates the reader in a specific moment (a year, an era), and *cyclical grounding*, which locates the reader within a recurring temporal structure (e.g., season, time of day).

Spatial grounding captures how strongly the passage anchors the reader in a particular place. Spatial grounding manifests in two ways: *geographic grounding* (a country, city, or named landmark) and *proximate grounding* (a room, a building, a street). Geographic grounding is efficient and has the highest ceiling on its own; proximate grounding requires more rendering to score high. The two types work together to produce the highest scores.

Sensory detail captures how central sensory experience is to the text across modalities. This feature measures whether sensory experience is a prominent feature of the text. It is distinct from concreteness: a passage can be concrete without foregrounding any particular sense modality.

3.3 Event Relations

Event relations capture the degree to which a text presents events in temporal sequence and links them through causal relationships. Event trigger detection, a requirement for event relation extraction, has a long history, from TimeBank (Pustejovsky et al., 2006) to neural models trained on ACE (Doddington et al., 2004) and ERE (Wang et al., 2022).

We define an event as a singular, bounded occurrence where something happens at a particular point in time, for which one can identify both *what happened* and *who or what it happened to*. Hypothetical, negated, and future events do not qualify. To measure event relations, we first identify event trigger spans using a RoBERTa-based event detection model fine-tuned on the event trigger data from Sims et al. (2019). We validate the model's performance on our web-scale corpus before proceeding (achieves an F1 of 0.85, see App. C.4).

For **temporal ordering**, annotators indicate which event occurred first in time, selecting *simultaneous* if the start points coincide, or selecting *too hard to tell* if the text does not support a clear ordering. Following Ning et al. (2018), we operationalize temporal order in terms of event start points rather than full event intervals. For **causal relation**, annotators choose among three categories: *direct cause* (Event 1 is sufficient to produce Event 2), *enablement* (Event 1 opens conditions for Event 2 without being sufficient), or *not related* (no causal link is supported by the text or world knowledge). For downstream analysis, we collapse direct cause and enablement into a single *causal* label, yielding a binary causal/not-causal distinction. We also calculate **event density**, the number of event triggers per token in the passage. Fig. 4 shows an example of the event relation annotation with event triggers bolded.

4 Sampling a Pretraining Corpus

Dolma (Soldaini et al., 2024) is an open corpus of over 3 trillion tokens designed to support language model pretraining. We sample from twelve

sources spanning web pages, news, encyclopedic text, books, and social media (Table A6).

Raw DOLMA documents present several challenges for narrative annotation. Full webpage text often includes navigation elements, lists of unrelated items, and fragmented or boilerplate content. Documents also vary enormously in length and register, from three-sentence forum posts to multi-thousand-word articles. Our sampling pipeline addresses these challenges in four steps. Additionally, the DOLMA dataset is comprised of almost entirely English. We acknowledge this as a limitation, and conduct an initial distributional comparison between Arabic, German, French and English in App. D. We find strong correlations across these languages, but more work on multilingual narrative analysis is needed.

Step 1: Initial passage extraction. We draw ∼51M three-sentence passages from ∼15M unique documents across the raw DOLMA v1.7[3] shards, allocating proportionally across sources according to the weights in Table A6. Sentences are segmented using NLTK `sent_tokenize`.

Step 2: Narrative scoring. We score each passage with a ROBERTA-based binary narrative classifier fine-tuned on data from NARRADETECT (Piper and Bagga, 2025) and STORYSEEKER (Antoniak et al., 2024). The classifier returns a continuous confidence score $p \in [0, 1]$. This score is used in Step 4 to focus the majority of final samples on passages with coherent narrative content.

Step 3: Topic classification. We assign topic and format labels to one chunk per document using the WEBORGANIZER topic and format classifier (Wettig et al., 2025), which distinguishes 24 topics and formats. Topic labels are used in Step 4 to balance the final samples across subject domains, preventing high-frequency topics like news from completely dominating the final corpus. Format labels are used later for data analysis.

Step 4: Final samples. We draw two stratified samples from the scored and classified passage pool, both preserving the inter-source proportions in Table A6. The first is the **gold dataset** which consists of 400 passages for human annotation. Of these, 85% are drawn from passages with narrative scores $p > 0.50$ and 15% are drawn without score filtering. For Common Crawl sources, passages are further balanced across topics. We split these 400 passages into two non-overlapping gold sets for the agency and setting dimensions: one for validating LLM classifications (**gold split A**), and one held out for evaluating NARRABERT classifications (**gold split B**). Because event-relation annotations require two confirmed event spans, event pairs are too sparse to split, so we evaluate event relations on the full annotated set at both validation stages.

The second is the full **NARRADOLMA** dataset which consists of ∼13M passages spanning ∼3.7M unique web documents. It follows the same stratified sampling procedure, with the unfiltered proportion shifted from 15% to 25% to increase coverage of non-narrative text. The raw DOLMA source distribution is shown in Fig. A17 and the distribution of NARRADOLMA categories (with topics applied to Common Crawl sources) is shown in Fig. A18.

5 Datasets

5.1 Human-Annotated Gold Dataset

Annotations were collected using a custom POTATO (Pei et al., 2022) annotation service (see App.E.2). One author annotated all three tasks for the full gold dataset ($N = 400$). For verification, a second author and an additional annotator combined for 122 overlapping labels for agency, $N = 120$ for setting, and $N = 275$ for events.

Human inter-annotator agreement is reasonable across all three tasks. For agency and setting, mean $\alpha = 0.75$ (range: 0.69–0.80) and 0.69 (0.65–0.75) respectively, with mean MAE of 0.61 and 0.54. For event relations, mean $\kappa = 0.68$ and mean F1 = 0.90. Per-dimension breakdowns are in App. A.1.

5.2 LLM-Labeled Dataset

Following established validation-first workflows for LLM-assisted annotation (Pangakis et al., 2023), we compare three models against human labels before selecting GEMMA for large-scale labeling: CLAUDE SONNET 4.6, QWEN3-235B-A22B, and GEMMA4-31B.

LLM labeling at scale. Agency and setting labels are produced by a single LLM call per passage. For event relations, we move beyond the single adjacent span pair used in validation to label relations for *all* adjacent event pairs in a passage. We consider two passage-level scores: *temporal sequencing* is the fraction of pairs with temporal relations; *causal density* is the fraction of event pairs

[3] huggingface.co/datasets/allenai/dolma

with causal relations. Both range from 0 to 1. Full prompts are in Fig. A23, A24, and A25. The final LLM-labeled dataset contains 25K passages, stratified by source and topic to preserve their original distribution.

LLM validation. We validate all three models against gold split A ($N = 200$) for agency and setting, and against the full set of human event annotations for event relations. No single model dominates. For agency and setting, mean α across the three models is 0.71 (0.49–0.88) with mean MAE of 0.53 (0.37–0.73). For event relations, mean F1 is 0.78 (0.77–0.79) and mean κ is 0.56 (0.54–0.58). We proceed with GEMMA for large-scale labeling due to its cost effectiveness and sufficient performance. Full per-dimension and per-model agreement breakdowns are in App. A.1.

5.3 NarraDolma and NarraBERT

To scale narrative annotation beyond what LLM inference budgets allow, we use the GEMMA labels to create NARRABERT, a pair of finetuned ROBERTA classifiers trained via knowledge distillation (Pangakis and Wolken, 2024). This converts the LLM's per-passage judgments into lightweight models that can label millions of texts at low cost.

Training. Both classifiers are finetuned on GEMMA pseudo-labels, but on two separate 25K passage draws. For agency and setting, each training example is a passage with nine corresponding Likert-scale labels, drawn unweighted from the corpus. For event relations, each training example is a span-level event pair identified by the Sims et al. (2019) model, each pair with binary labels for temporal ordering and causal relation. Because only passages with at least two spans yield a pair (35% of the passages) the second draw upsamples event-dense passages, giving 61K pairs rather than the 24K the unweighted draw produces at the same annotation cost. Drawing separately keeps the Likert labels on the corpus distribution. We frame the Likert dimensions as regression tasks and the event relation dimensions as binary classification tasks. NARRABERT consists of two ROBERTA-BASE encoders, one of which is a shared encoder with nine regression heads for the agency and setting dimensions and the other a dedicated encoder for event relations. Hyperparameters can be found in App. C.3.

Validation against human annotations. We validate NARRABERT against gold split B ($N = 200$) for agency and setting, and against the full set of human event annotations for event relations. For agency and setting, mean $\alpha = 0.68$ (0.52–0.78) with mean MAE of 0.57 (0.41–0.73). For the two event relation features, mean F1 $= 0.75$ (0.73–0.77) vs. GEMMA's 0.78 (0.77–0.79). Agreement is broadly comparable to the LLMs themselves. Full per-dimension breakdowns for each feature are in App. A.1.

Scaling up labels. We apply NARRABERT to the full NARRADOLMA corpus of ∼13M passages across 12 Dolma sub-corpora. For agency and setting features, the shared encoder produces nine Likert scores per passage in a single forward pass. For event relations, we first run the Sims et al. (2019) event detector to identify event trigger spans in each passage, then apply the event relation encoder to classify every adjacent span pair which yield the temporal sequencing and causal density scores. The event detector itself also yields event density, the number of events per token. This produces a 12-dimensional narrative feature vector for each of the ∼13M passages, which together constitute the NARRADOLMA dataset. For all analyses in §6, we aggregate these passage-level vectors to the document level by averaging across each document's passages, yielding one narrative vector per document (∼3.7M documents), our **analysis corpus** of document-level **narrative profiles**.

6 The Narrative Landscape of Pretraining Data

6.1 The Structure of Narrative Features

Before comparing the narrative features across categories, we examine the distributions of the 12 narrative features across the full NARRADOLMA corpus (Fig. A2). The five agency dimensions cluster near the floor (∼1.0 on a 1–5 Likert scale) while the setting features are more spread out across the full range of scores. High agency features require text to have a perceiving consciousness which seems to be a property the NARRADOLMA corpus lacks, on average, compared to features related to events or setting. Temporal and spatial grounding are features most web text has to some degree since dates and place names are abundant.

To investigate whether the narrative content is recoverable from lexical content, we overlay feature

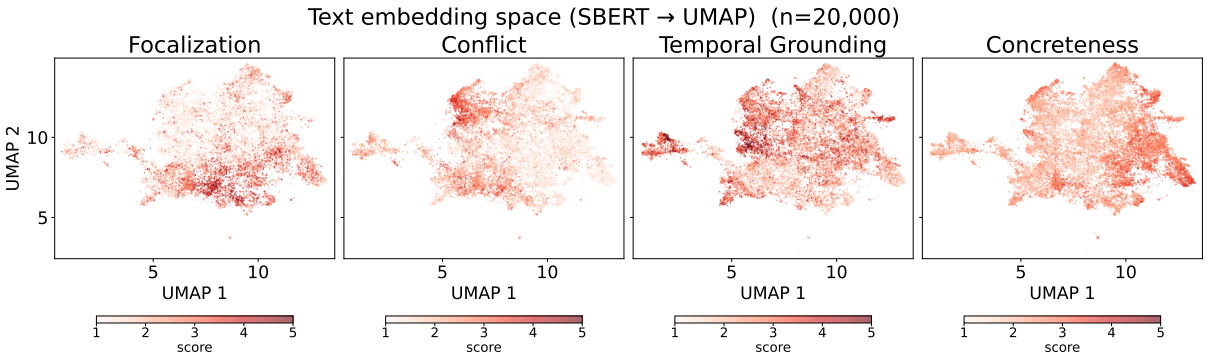

Figure 5: 20K documents randomly sampled from NARRADOLMA, represented as SBERT embeddings and reduced to two UMAP dimensions, overlayed with focalization, conflict, temporal grounding, and concreteness scores. See the SBERT embedding space with annotated topic labels in Fig. A7.

scores onto a UMAP reduction of SBERT embeddings for 20K randomly sampled NARRADOLMA documents. Fig. 5 shows four of the narrative feature scores (focalization, conflict, temporal grounding, and concreteness). Each of the four plots show distinct patterns. If narrative features were a redescription of subject matter, we would expect documents nearby in embedding space to share narrative profiles. However, we see that each feature has a distinct pattern over identical points, with some narrative features more similar than others. Fig. A6 shows all feature scores overlayed upon the reduced SBERT embeddings, and Fig. A7 shows the WebOrganizer topic centroids labeled upon the SBERT embeddings.

6.2 How Narrative Features Relate to Each Other

We expect some narrative features to be correlated, based both on our initial examinations of the data and on the theoretical overlap between subsets of the features. We calculate the Pearson correlation between each of the 12 narrative features shown in Fig. 6. We see "blocks" of high correlation where certain subsets of features are more highly correlated with each other than others. Focalization, emotion, and cognition form an *interiority* block ($r = 0.75, p < 0.005$, peaking at 0.82 between focalization and cognition), and concreteness and sensory detail form a *texture* block ($r = 0.76, p < 0.005$). Three more blocks display moderate correlation: temporal and spatial grounding, temporal sequencing and causal density, and the looser grouping of change of state, conflict, and event density all cohere in the mid-0.4s to low-0.5s. Across blocks, interiority runs mildly against grounding ($r = -0.23, p < 0.005$).

To further explore this correlation we ran principal-component analysis (PCA) on the 10 non-event-relation features (App. B.3). The first three principal components (PCs) explain 71% of the variance and correspond to interpretable axes that are similar to the correlation patterns. These dimensions can be used as summary axes for sampling, in cases where a coarser categorization is helpful.

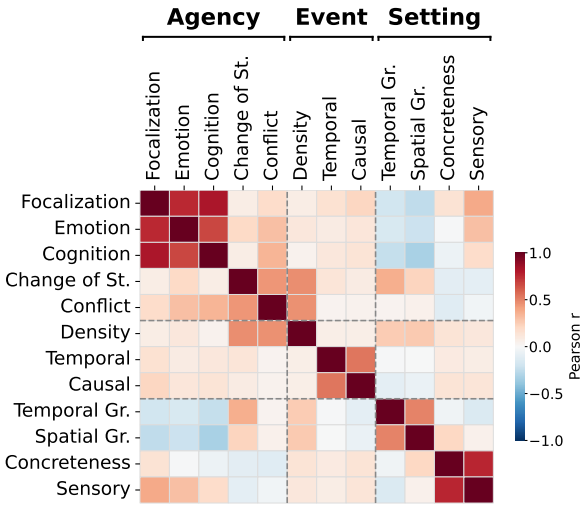

Figure 6: Pearson correlation of each feature pair, across our analysis corpus. Cells involving temporal sequencing or causal density use the 2.2M documents where those features are defined.

6.3 Where Narrativity Concentrates in Pretraining Data

We next examine narrative profiles across DOLMA sources, topics, and formats (§4). The **DOLMA source** is where the text comes from, while WebOrganizer (Wettig et al., 2025) provides the **topic** (what the text is about) and the **format** (how the text is written). These categories have been used

when curating pretraining data mixtures to improve performance on downstream tasks. But can sampling over these categories improve narrative coverage? How is narrative content distributed across these categories?

In the following results, temporal sequencing and causal density are averaged over the 2.2M documents where those features are defined, otherwise results are reported for all 3.7M documents in our analysis corpus.

DOLMA Sources. Fig. 7 shows the mean scores for focalization, conflict, temporal grounding and concreteness by source. Three of the more specifically-curated sources (Reddit, Gutenberg, and Wikipedia) show distinct narrative profiles. Reddit and Gutenberg score high on focalization and conflict, and lower on temporal grounding. Wikipedia shows the exact opposite narrative profile. Fig. A11 shows the full mean z-scored narrative features by DOLMA source.

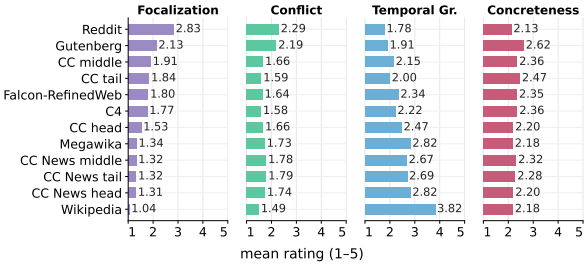

Figure 7: Bar plot of the focalization, conflict, temporal grounding and concreteness feature scores **by source**. All features are shown in App. B.1.

Topics. Fig. 8 shows the mean scores for focalization, conflict, temporal grounding and concreteness by topic. The history topic scores high on temporal grounding but low on focalization, while social life and literature are the exact opposite. Crime & law scores high on conflict but quite low on focalization. Fig. A12 shows the mean z-scored for the full feature set by topic.

Formats. Fig. 9 shows the mean scores for focalization, conflict, temporal grounding and concreteness by format. Narrative is a property of how a text is written (as opposed to what a text is about), and the format labels are the closest of the three categories (source, topic, and format) to describing this. Creative writing, personal blog, user review, and comment section score high on focalization. Tutorial scores high for concreteness but low for

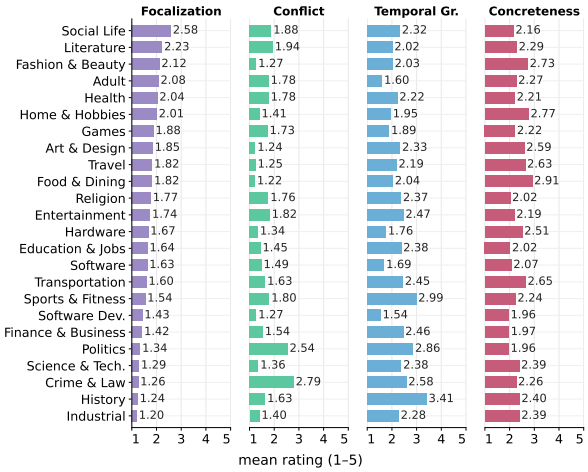

Figure 8: Bar plot of the focalization, conflict, temporal grounding and concreteness feature scores **by topic**. All features are shown in App. B.1.

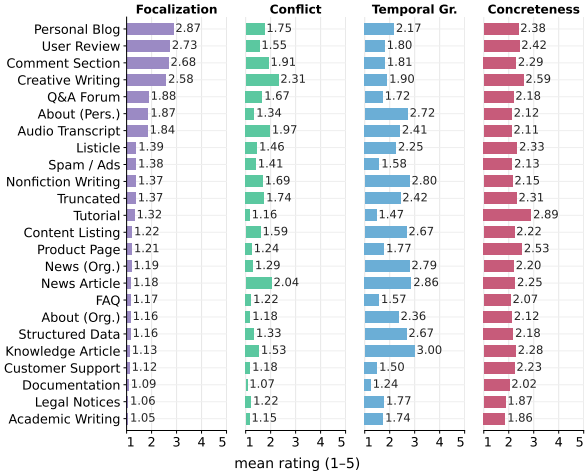

Figure 9: Bar plot of the focalization, conflict, temporal grounding and concreteness feature scores **by format**. All features are shown in App. B.1.

most other features. Fig. A13 shows the full mean z-scored features by format.

Variation within categories. The mean feature scores describe where the categories sit on average, but categories may also display high or lower variation. The within-category standard deviation (SD) in corpus SD units is 0.96 for DOLMA source, 0.92 for topic, and 0.89 for format. A SD of 1 means a category is exactly as internally varied as the whole NARRADOLMA corpus. The more factual and technical subcategories like the Wikipedia **source** (0.68) and the documentation **format** (0.31) have comparably lower SDs. The more creative subcategories like the Reddit **source** (1.09) and the creative writing **format** (1.01) have higher SDs. The full set of SD values across all categories and

features is in App. B.4. This is meaningful because the categories a model practitioner would upweight for more narrative text are precisely the ones with the widest spread of narrative features. Filtering *within* these categories is necessary in order to effectively upweight specific narrative profiles.

7 Discussion

Our analysis reveals three core empirical findings. First, we find that narrative can be measured as a multidimensional, continuous spectrum rather than a binary property. Second, the narrative features are unequally distributed across categories in ways that current curation practices neither measure nor account for. Third, narrative variation within categories is substantial, indicating that source-level, topic-level, and format-level labels are too coarse to capture narrative diversity. Upweighting "narrative" sources would not uniformly increase all narrative qualities. Together these results position narrative structure as a measurable, multidimensional property of web text rather than a single "narrativity" quantity.

These findings have direct implications for data curation. Source-level mixing decisions, as currently practiced, cannot account for the multidimensional structure of narrative in pretraining data. A training mixture heavy on Reddit would amplify interiority while starving the model of storyworld texture and event structure; a mixture heavy on Wikipedia would do the reverse. Narrative-aware curation requires understanding which axes a given source contributes to, not just whether it is "narrative" in aggregate.

Looking forward, this work opens several directions. Controlled data mixing experiments using frameworks such as RegMix (Liu et al., 2025) could establish direct causal links between pretraining narrative composition and downstream narrative capability. Analysis of intermediate training checkpoints could reveal when and how narrative competencies emerge during pretraining. Such audits could also support safety research about the narratives influencing model outputs, whether those narratives are desirable, and how they affect the stories that many people, including children, are already generating via popular chatbots.

Our framework, datasets, and models provide a foundation for these investigations and for treating narrative structure as a primary dimension of pretraining data composition, complementing existing axes of quality, toxicity, and topic distribution.

8 Conclusion

In this work we presented NARRADOLMA, a large-scale corpus characterizing narrative structure in pretraining data, together with NARRABERT, a pair of efficient classifiers for its 11 narrative dimensions. We developed an annotation framework grounded in narrative theory, validated through human annotation and LLM-assisted labeling, and applied it to produce narrative feature vectors for $\sim$13M passages across $\sim$3.7M documents.

9 Limitations

Our study has several limitations. First, NARRADOLMA is a stratified subsample ($\sim$13M passages) of a $>$3-trillion-token corpus, and it deliberately over-represents narrative content. The relative structural patterns we report are therefore informative but absolute prevalence figures do not transfer to raw DOLMA.

Second, human annotation was conducted on 400 passages, a small sample relative to the diversity of web text. This constraint reflects the fine-grained nature of our framework. Each passage requires judgment across 11 dimensions, and the iterative calibration process spanned several months. Primary annotation that was used for model validation was conducted by one author. Some agreement scores between annotators remain modest suggesting some of these dimensions are intrinsically difficult regardless of annotator expertise.

Third, the gold sets for agency and setting validation were non-overlapping, but event annotations were too sparse to split, so event evaluation is not fully independent across the LLM and classifier validation stages.

Fourth, our analysis is limited to English text and to a single theoretical operationalization of narrative. Other languages, narrative traditions, and frameworks may surface different structure. We do conduct an initial distribution comparison between English, Arabic, German, and French detailed in App. D, but leave more robust analysis for future work.

Finally, we do not establish direct causal links between pretraining narrative composition and downstream model behavior. Our focus is on characterizing and mapping narrative structure in pretraining data, which is challenging and required designing new sampling, annotation, and analysis pipelines.

Experiments such as controlled data remixing or checkpoint analysis remain important directions for future work.

Ethical Considerations

This study relies on the open English language pretraining dataset DOLMA, released with the Open Data Commons Attribution License (ODC-By).[4] This dataset includes massive amounts of web scraped data, including toxic, explicit, and personal data posted publicly to the internet. We avoid highlighting such passages in this paper, but we did not avoid them during annotation, as the narrative qualities of this data may have significant implications for safe story generation. We release the human-, LLM-, and NARRABERT-labeled datasets, which include the sampled portion of text along with the Dolma unique ID to allow for future rehydrating. As the data authors could not have an expectation of privacy, we determined that review for this study was not required by our institution's Internal Review Board (IRB).

Potential Risks Because NARRADOLMA includes both passage text and narrative feature annotations, it may make it easier to identify, filter, or prioritize web passages with particular narrative qualities, including emotionally intense, personal, violent, or explicit narratives. This could be misused to construct training mixtures that overrepresent sensational, manipulative, or harmful narrative styles, or to retrieve sensitive personal disclosures from publicly available web text. These risks are heightened because DOLMA contains web-scraped data, including toxic, explicit, and personal content. However, releasing our annotations supports future research into the factors influencing story generation, which is a very popular task for users of modern chatbots and therefore of interest for safety audits. Given the growing prominence of these tools and their effects on society, we believe studies like ours can contribute to better understandings, for researchers and the public.

Use of AI Assistance We used AI assistants for limited support during coding, analysis, and writing. All data decisions, annotations, citations, statistical analyses, interpretations, and final writing decisions were reviewed and made by the authors.

Acknowledgments

We thank our annotators for their careful work annotating the gold dataset, often across difficult and ambiguous passages, and for the many discussions that helped refine the annotation guidelines. We are grateful to our colleagues for feedback on earlier drafts of this work, and to members of the NLP Group at the University of Colorado Boulder for valuable discussion throughout the project. We also thank the anonymous reviewers for their constructive comments. Lastly we thank Rohan Das, Matt Pauk, Advait Deshmukh, and Uma Gunturi for their help with developing the annotation framework.

Computational resources were provided by the University of Colorado Boulder Research Computing group, including the Blanca condo cluster.

[4] https://opendatacommons.org/licenses/by/1-0/

References

Maria Antoniak, Joel Mire, Maarten Sap, Elliott Ash, and Andrew Piper. 2024. Where do people tell stories online? story detection across online communities. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7104–7130, Bangkok, Thailand. Association for Computational Linguistics.

Brian Boyd. 2009. *On the origin of stories: Evolution, cognition, and fiction*. Harvard University Press.

Tuhin Chakrabarty, Philippe Laban, Divyansh Agarwal, Smaranda Muresan, and Chien-Sheng Wu. 2024. Art or artifice? large language models and the false promise of creativity. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–34.

George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. The automatic content extraction (ACE) program – tasks, data, and evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).

Dylan Thomas Doyle, Jay K Ghosh, Reece Suchocki, Brian C Keegan, Stephen Voida, and Jed R Brubaker. 2024. Stories that heal: Characterizing and supporting narrative for suicide bereavement. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 354–366.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Arjun Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Achyutarama Ganti, Steven Wilson, Zexin Ma, Xinyan Zhao, and Rong Ma. 2022. Narrative detection and feature analysis in online health communities. In *Proceedings of the 4th Workshop of Narrative Understanding (WNU2022)*, pages 57–65, Seattle, United States. Association for Computational Linguistics.

Gérard Genette. 1983. *Narrative Discourse: An Essay in Method*. Cornell University Press, Ithaca, NY.

Rachel Giora and Yeshayahu Shen. 1994. Degrees of narrativity and strategies of semantic reduction. *Poetics*, 22(6):447–458.

Jonathan Gottschall. 2012. *The storytelling animal: How stories make us human*. Houghton Mifflin Harcourt.

Neel Gupta, Maria Antoniak, and Melanie Walsh. 2026. Ai fiction in the wild.

Sil Hamilton, Matthew Wilkens, and Andrew Piper. 2026. Narrabench: A comprehensive framework for narrative benchmarking. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3786–3801.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding.

David Herman. 2009a. *The Nexus of Narrative and Mind*, chapter 6. John Wiley & Sons, Ltd.

David Herman. 2009b. *The Third Element; or, How to Build a Storyworld*, chapter 5. John Wiley & Sons, Ltd.

David Herman. 2011. *Basic elements of narrative*. John Wiley & Sons.

Qian Liu, Xiaosen Zheng, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, and Min Lin. 2025. Regmix: Data mixture as regression for language model pre-training.

Li Lucy, Suchin Gururangan, Luca Soldaini, Emma Strubell, David Bamman, Lauren Klein, and Jesse Dodge. 2024. AboutMe: Using self-descriptions in webpages to document the effects of English pretraining data filters. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7393–7420, Bangkok, Thailand. Association for Computational Linguistics.

Qiang Ning, Hao Wu, and Dan Roth. 2018. A multi-axis annotation scheme for event temporal relations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1318–1328, Melbourne, Australia. Association for Computational Linguistics.

Elinor Ochs, Lisa Capps, et al. 2009. *Living narrative: Creating lives in everyday storytelling*. Harvard University Press.

Nicholas Pangakis and Samuel Wolken. 2024. Knowledge distillation in automated annotation: Supervised text classification with llm-generated training labels.

Nicholas Pangakis, Samuel Wolken, and Neil Fasching. 2023. Automated annotation with generative ai requires validation.

Jiaxin Pei, Aparna Ananthasubramaniam, Xingyao Wang, Naitian Zhou, Apostolos Dedeloudis, Jackson Sargent, and David Jurgens. 2022. POTATO: The portable text annotation tool. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 327–337, Abu Dhabi, UAE. Association for Computational Linguistics.

Guilherme Penedo, Hynek Kydlíček, Amir Hossein Kargaran, and Leandro von Werra. 2026. Finetranslations. https://huggingface.co/datasets/HuggingFaceFW/finetranslations.

Federico Pianzola. 2018. Looking at narrative as a complex system: The proteus principle. In *Narrating complexity*, pages 101–122. Springer.

Andrew Piper. 2023. Computational narrative understanding: A big picture analysis. In *Proceedings of the Big Picture Workshop*, pages 28–39.

Andrew Piper and Sunyam Bagga. 2022. Toward a data-driven theory of narrativity. *New Literary History*, 54(1):879–901.

Andrew Piper and Sunyam Bagga. 2025. NarraDetect: An annotated dataset for the task of narrative detection. In *Proceedings of the The 7th Workshop on Narrative Understanding*, pages 1–7, Albuquerque, New Mexico. Association for Computational Linguistics.

Andrew Piper, Richard Jean So, and David Bamman. 2021. Narrative theory for computational narrative understanding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 298–311, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

V. Propp. 1968. *Morphology of the Folktale: Second Edition*. University of Texas Press.

James Pustejovsky, Marc Verhagen, Roser Sauri, Jessica Littman, Robert Gaizauskas, Graham Katz, Inderjeet Mani, Robert Knippen, and Andrea Setzer. 2006. TimeBank 1.2. LDC2006T08.

Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. 2021. Scaling language models: Methods, analysis & insights from training Gopher. *arXiv preprint arXiv:2112.11446*.

John TE Richardson. 1975. Imagery, concreteness, and lexical complexity. *The Quarterly Journal of Experimental Psychology*, 27(2):211–223.

Mark O Riedl and R Michael Young. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268.

Matthew Sims, Jong Ho Park, and David Bamman. 2019. Literary event detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3623–3634, Florence, Italy. Association for Computational Linguistics.

Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Jha, Sachin Kumar, Li Lucy, Xinxi Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma Strubell, Nishant Subramani, Oyvind Tafjord, Pete Walsh, Luke Zettlemoyer, Noah A. Smith, David Wadden, Iz Beltagy, Hannaneh Hajishirzi, Ali Farhadi, and Dirk Groeneveld. 2024. Dolma: An open corpus of three trillion tokens for language model pretraining research. *arXiv preprint arXiv:2402.00159*.

Yufei Tian, Tenghao Huang, Miri Liu, Derek Jiang, Alexander Spangher, Muhao Chen, Jonathan May, and Nanyun Peng. 2024. Are large language models capable of generating human-level narratives? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17659–17681.

Xiaozhi Wang, Yulin Chen, Ning Ding, Hao Peng, Zimu Wang, Yankai Lin, Xu Han, Lei Hou, Juanzi Li, Zhiyuan Liu, Peng Li, and Jie Zhou. 2022. Maven-ere: A unified large-scale dataset for event coreference, temporal, causal, and subevent relation extraction.

Alexander Wettig, Kyle Lo, Sewon Min, Hannaneh Hajishirzi, Danqi Chen, and Luca Soldaini. 2025. Organize the web: Constructing domains enhances pretraining data curation. In *Forty-second International Conference on Machine Learning (COLM)*.

Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V. Le, Tengyu Ma, and Adams Wei Yu. 2023. DoReMi: Optimizing data mixtures speeds up language model pretraining. *arXiv preprint arXiv:2305.10429*.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence?

A Datasets & Validation

This appendix reports the validation of NARRABERT: gold-set label distributions, inter-annotator agreement, and model agreement at each validation stage.

Figure A1 shows the distributions of gold annotated labels across the 11 manually-annotated features. Figure A2 shows the distributions of the full NARRADOLMA corpus labels.

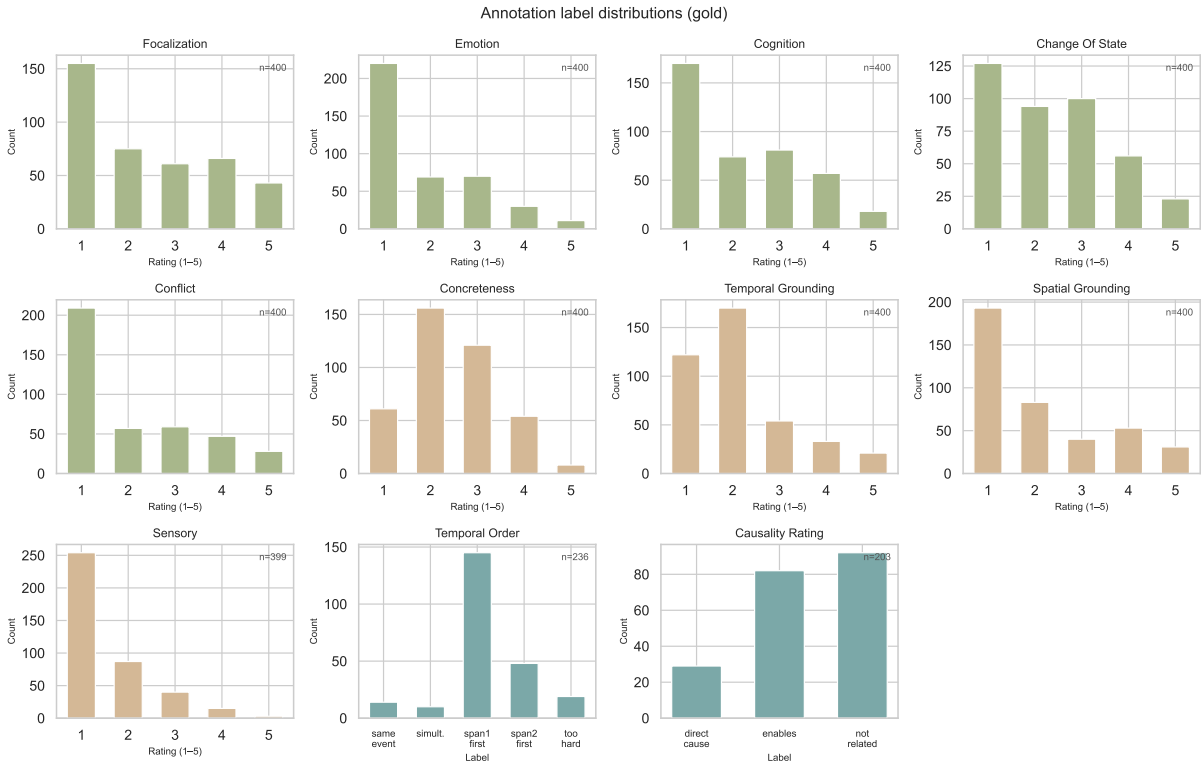

Figure A1: Distribution of gold annotated labels across the 11 manually-annotated features.

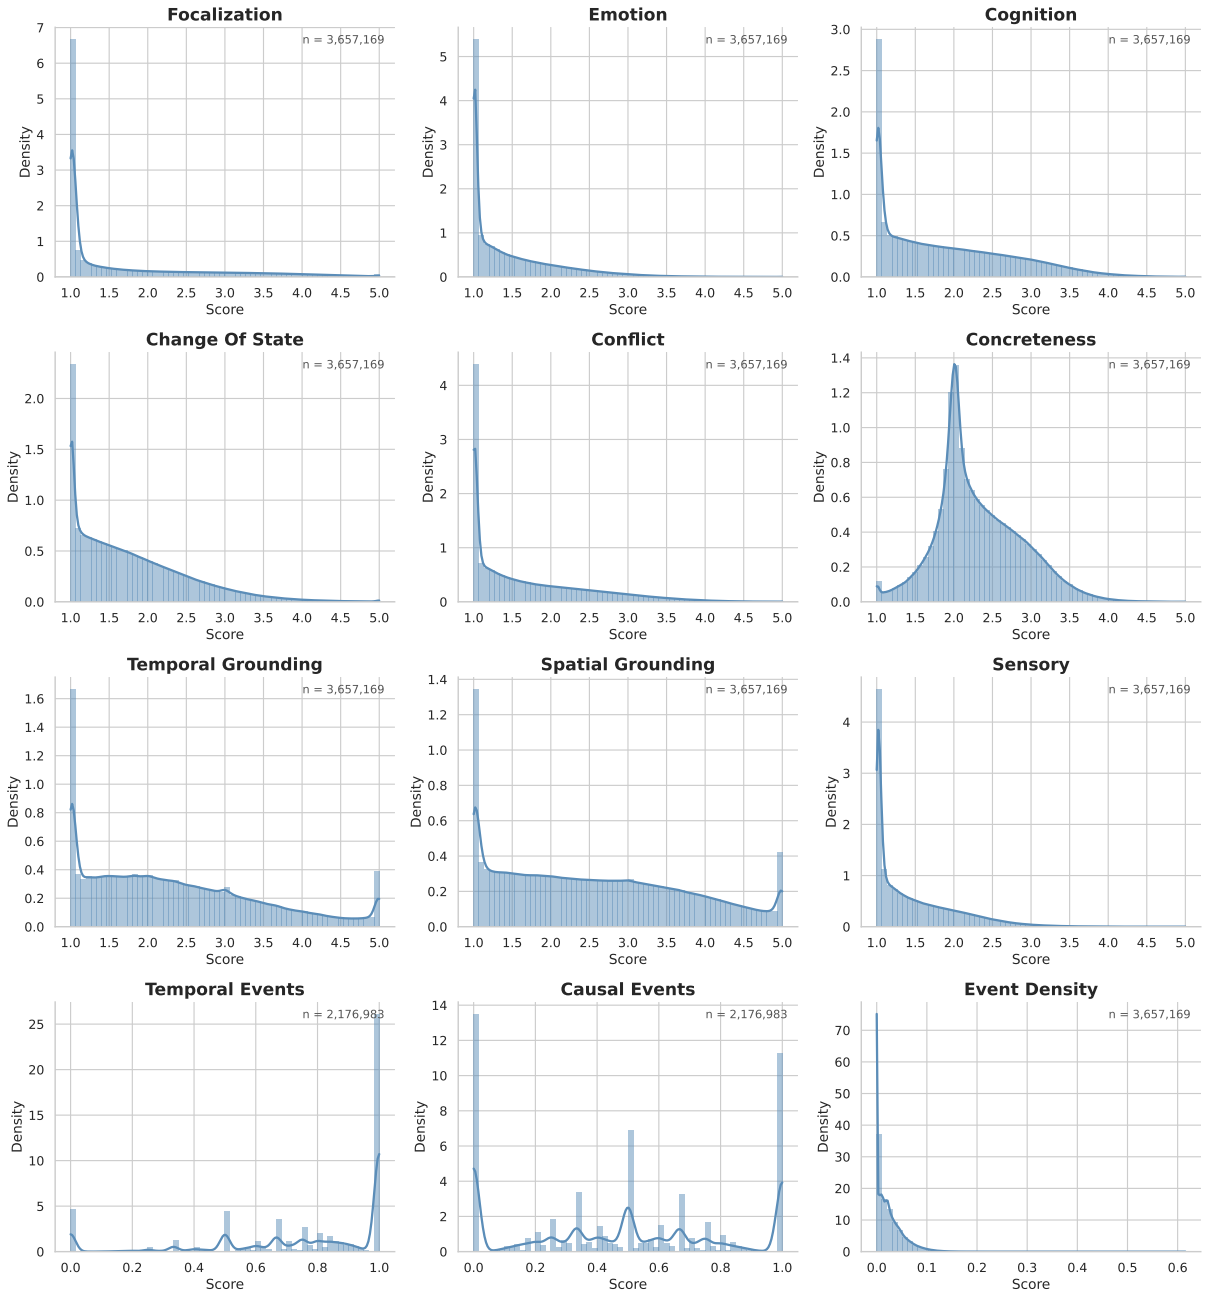

Figure A2: Distributions of each narrative feature in the full NarraDolma corpus. The n in the upper right corner of each plot indicates the number of documents shown in the distribution.

A.1 Agreement Scores

Tables A1, A2, and A3 report full per-dimension agreement scores for each validation stage. Table A1 reports human inter-annotator agreement on the 120 and 122-passage overlap set for agency and setting and on the full event annotation set. Agency and setting dimensions report MAE and Krippendorff's α, while event relations report Cohen's κ and macro F1. Table A2 reports agreement between each of the three LLMs and gold split A ($N = 200$ for agency and setting, full event set for event relations), with the best-performing model per dimension bolded. Table A3 reports NARRABERT agreement against the held-out gold split B ($N = 200$ for agency and setting, full event set for event relations), using annotations not seen during LLM validation.

Dimension	**N**	**MAE**	α	κ	**F1**
Agency					
Focalization	122	0.639	0.795	—	—
Emotion	122	0.541	0.709	—	—
Cognition	122	0.615	0.776	—	—
Change of state	122	0.697	0.691	—	—
Conflict	122	0.549	0.768	—	—
Setting					
Concreteness	120	0.647	0.666	—	—
Temporal gr.	120	0.550	0.713	—	—
Spatial gr.	120	0.583	0.747	—	—
Sensory	120	0.392	0.649	—	—
Event					
Is span event	550	—	—	0.798	0.921
Temporal	123	—	—	0.645	0.891
Causality	82	—	—	0.608	0.901

Table A1: Human inter-annotator agreement. Agency and setting report MAE and Krippendorff's α; event relations report Cohen's κ and F1. For event annotation, annotators are asked to determine whether pre-highlighted event spans (detailed in App. C.4) are events. If both pre-highlighted spans are events, annotators label temporal and causal relations.

		Sonnet			Qwen3			Gemma4		
Dimension	**N**	**MAE**	α/κ	**F1**	**MAE**	α/κ	**F1**	**MAE**	α/κ	**F1**
Agency (α)										
Focalization	200	0.605	0.760	—	0.616	0.762	—	**0.597**	**0.764**	—
Emotion	200	0.505	0.713	—	0.453	0.694	—	**0.436**	**0.749**	—
Cognition	200	**0.545**	0.737	—	0.646	0.678	—	0.559	**0.753**	—
Change of st.	200	**0.565**	**0.739**	—	0.726	0.634	—	0.657	0.703	—
Conflict	200	**0.580**	**0.753**	—	0.637	0.644	—	0.594	0.748	—
Setting (α)										
Concreteness	200	**0.445**	**0.651**	—	0.610	0.602	—	0.516	0.628	—
Temporal gr.	200	**0.425**	**0.765**	—	0.506	0.720	—	0.536	0.728	—
Spatial gr.	200	0.445	0.813	—	**0.378**	**0.883**	—	0.651	0.744	—
Sensory	200	**0.370**	0.574	—	0.425	0.486	—	0.384	**0.639**	—
Event (κ)										
Temporal	218	—	**0.556**	**0.778**	—	—	—	—	0.539	0.769
Causality	180	—	0.550	0.775	—	—	—	—	**0.575**	**0.786**

Table A2: LLM validation against split A of the gold annotations. Agency and setting report MAE and α; event relations report κ and F1. Qwen3 was not evaluated on event relations. Best per dimension in bold (lower is better for MAE, higher for all others).

Dimension	**N**	**MAE**	α	**F1**
Agency				
Focalization	200	0.559	0.782	—
Emotion	200	0.432	0.712	—
Cognition	200	0.615	0.721	—
Change of st.	200	0.734	0.616	—
Conflict	200	0.538	0.719	—
Setting				
Concreteness	200	0.586	0.656	—
Temporal gr.	200	0.570	0.727	—
Spatial gr.	200	0.690	0.659	—
Sensory	200	0.405	0.515	—
Event				
Temporal	236	—	—	0.766
Causality	193	—	—	0.727

Table A3: NarraBERT validation against split B of the gold annotations. Agency and setting report MAE and α; event relations report F1. Validated against held-out human annotations not used during LLM validation.

B Analysis Details & Additional Plots

This appendix expands the corpus analysis.

B.1 Bar Plot Visuals

Figures A3, A4, and A5 show bar plots of the feature scores by DOLMA source, topic, and format respectively.

B.2 UMAP Visuals

Figure A6 shows all 12 feature scores overlayed upon a UMAP reduction of 20,000 randomly sampled passages. Figure A7 shows the topic centroids overlayed upon the SBERT embeddings.

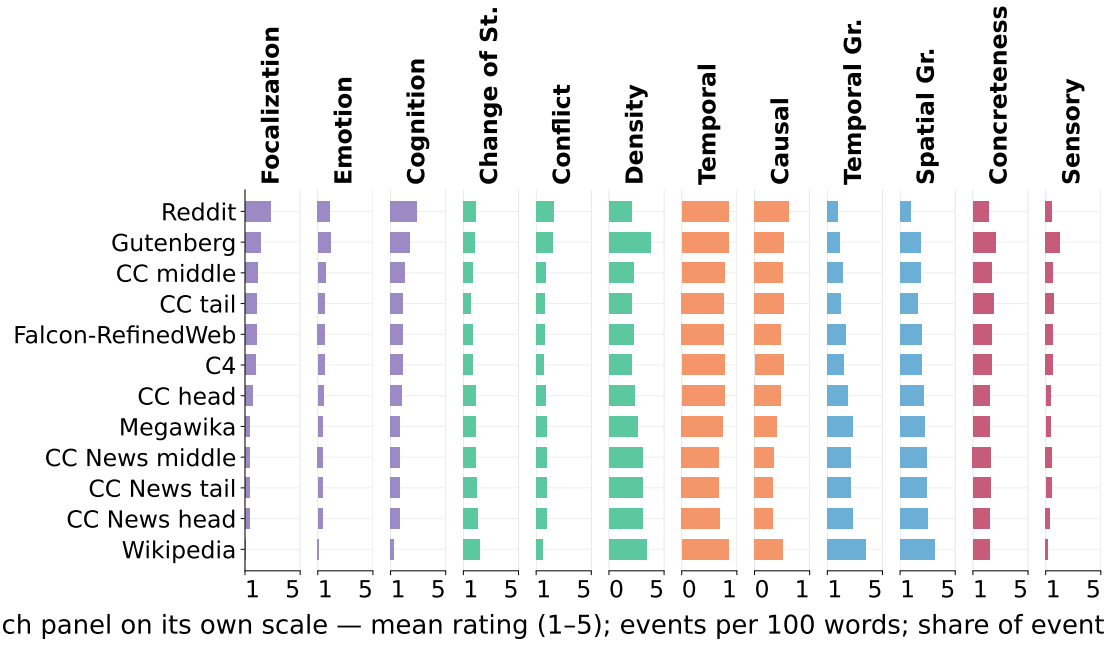

Figure A3: The mean of each feature across the full NARRADOLMA corpus by **source**. Colored by the five covarying "blocks" described in Section 6.2.

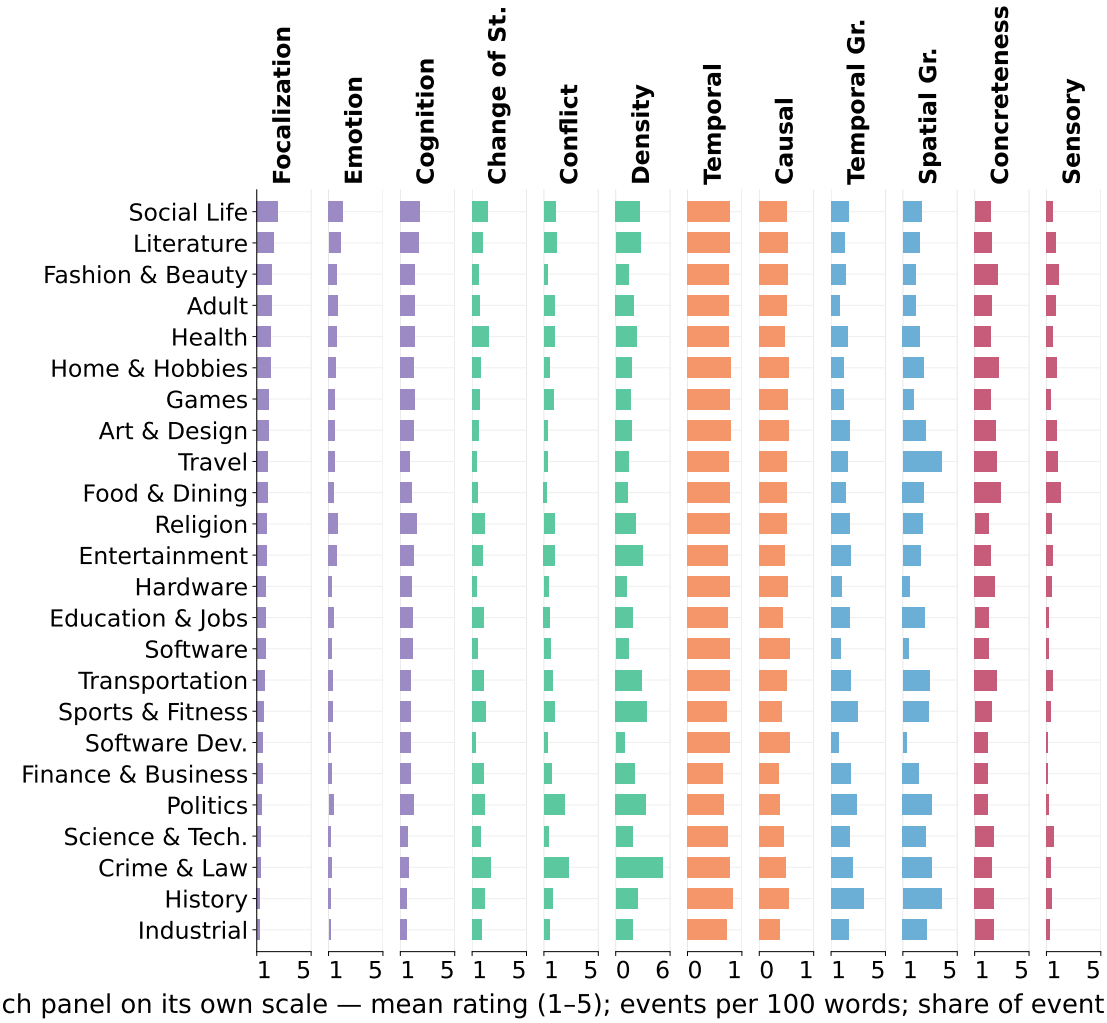

Figure A4: The mean of each feature across the full NARRADOLMA corpus by **topic**. Colored by the five covarying "blocks" described in Section 6.2.

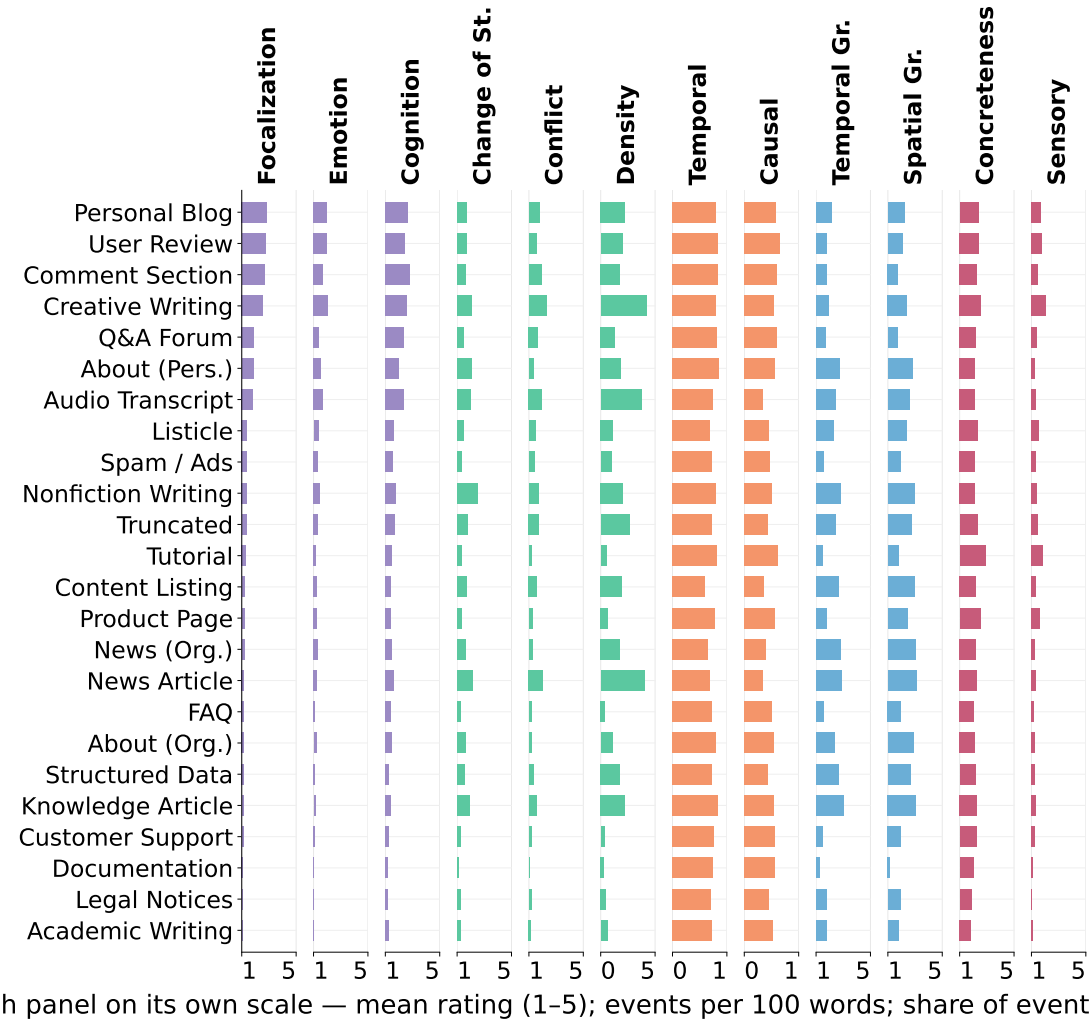

Figure A5: The mean of each feature across the full NARRADOLMA corpus by **format**. Colored by the five covarying "blocks" described in Section 6.2.

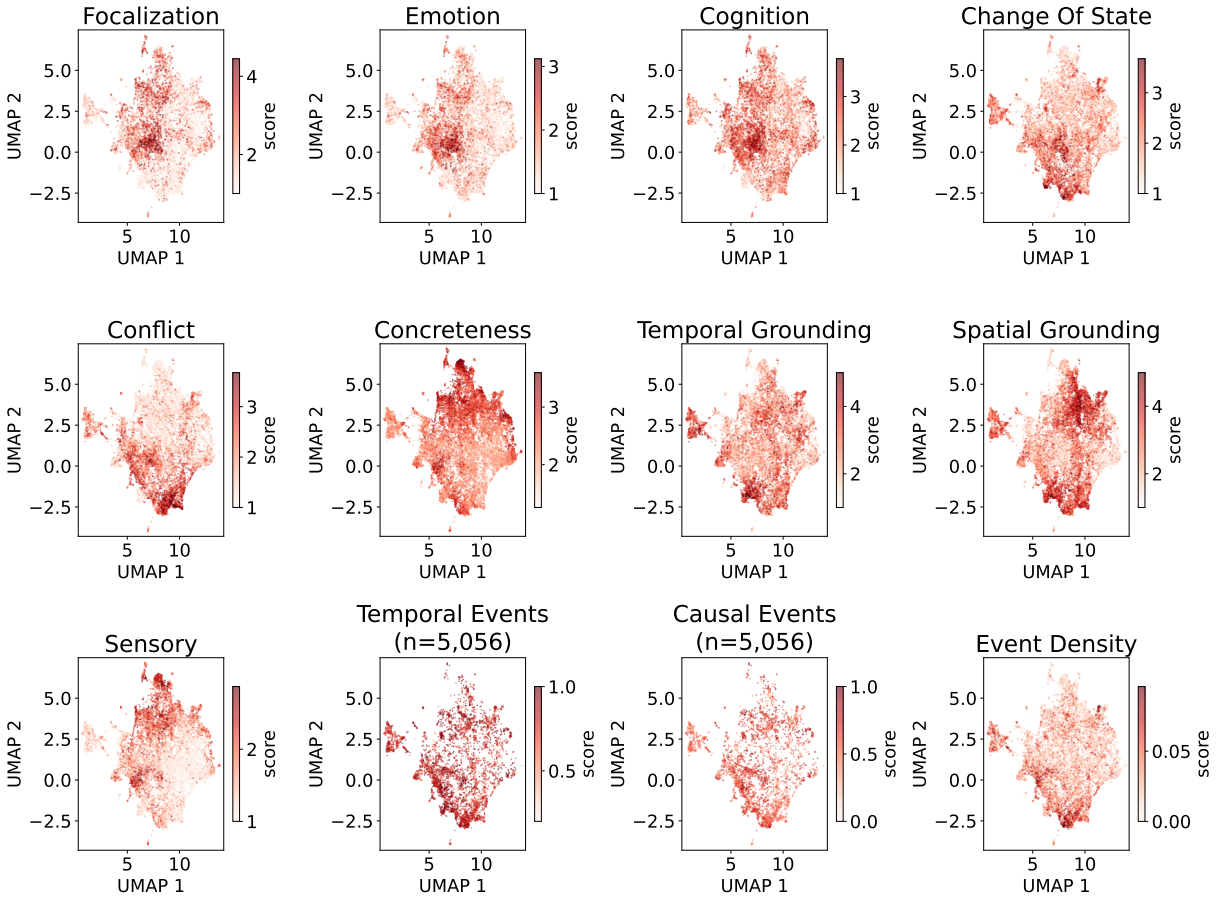

Figure A6: UMAP reduction of SBERT embeddings for 20,000 randomly sampled documents, colored by narrative feature scores.

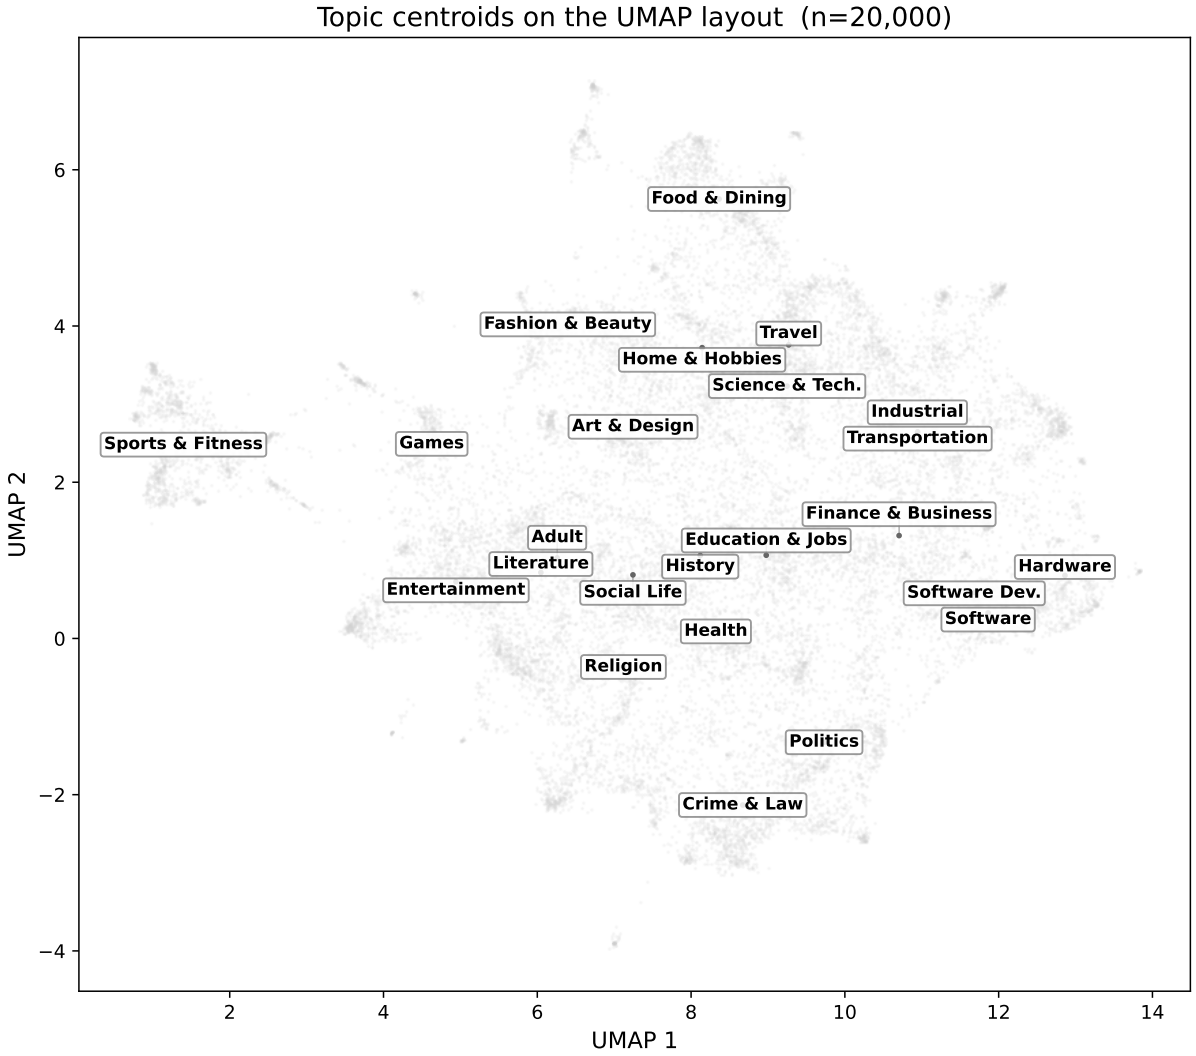

Figure A7: UMAP reduction of SBERT embeddings for 20,000 randomly sampled documents, annotated with the topic centroids from WebOrganizer labels.

B.3 PCA Details

Figure A9 shows the distribution of PC scores over the full NARRADOLMA corpus (with 10 features, removing the two event relation features). Figure A10 shows a bar plot of the proportion of documents within NARRADOLMA that lie within the top quartile of PC scores for each of the first three PCs. Table A5 shows passages that have high PC scores for the first three PCs.

PCA Loadings — top 5 PCs × 10 dimensions

	PC1	PC2	PC3	
Focalization	0.52	-0.04	0.02	Agency
Emotion	0.49	0.05	-0.07	
Cognition	0.51	-0.03	-0.16	
Change Of State	0.09	0.52	-0.12	
Conflict	0.23	0.39	-0.19	
Event Density	0.10	0.49	0.10	Event
Temporal Grounding	-0.19	0.42	0.06	Setting
Spatial Grounding	-0.20	0.38	0.29	
Concreteness	0.09	-0.02	0.68	
Sensory	0.26	-0.06	0.60	

Table A4: Loadings for the first three principal components, explaining ∼71% of variance (see Fig. A8).

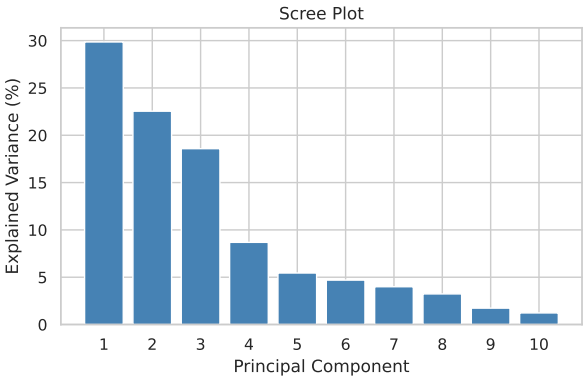

Figure A8: Scree plot of the 10 principal components across the 10 narrative features. We see a significant drop after PC3 to below 10%.

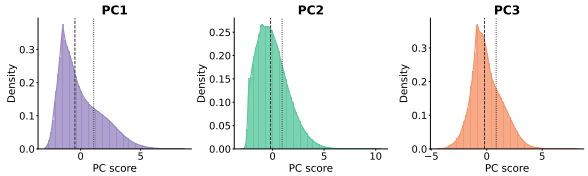

Figure A9: Distributions of the first 3 PCs across all 3.7M documents in NARRADOLMA.

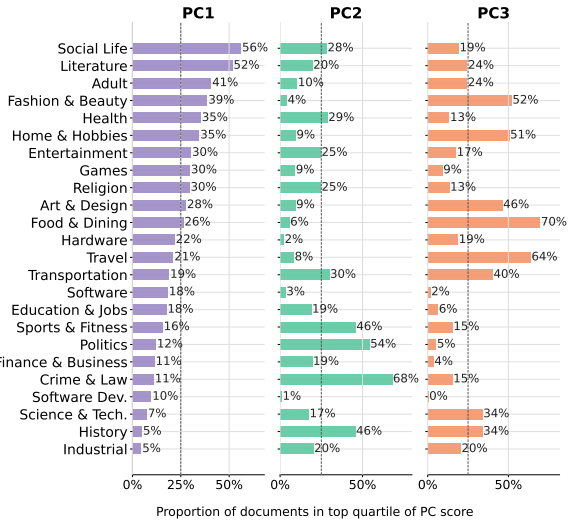

Figure A10: Bar plot of the proportion of documents within NARRADOLMA that lie within the top quartile of PC scores for each of the first three PCs

PC	Score	Topic	Passage
PC1	+8.22	Literature	*In my head everything was a washed-out shade of grey, the color leached away by horror and utter shock. I was still there—in the café, listening to him scream and scream, staring at myself as I stood frozen, the empty cup still clutched in my hand. Wasn't my fault.*
	−3.31	History	*University of Glasgow Research Conference, Glasgow, UK, 16-17 June 2021. Struan, A. and Alexander, M. (2010) Expressions of Civilisation and Colonisation in the Historical Thesaurus of the Oxford English Dictionary. From the Grand Tour to Mass Tourism: The Modern History of the British Abroad, Newcastle, UK, 01-02 Apr 2010.*
PC2	+7.92	Crime & Law	*But at Norwich Magistrates Court on Thursday (August 25) she was found guilty after a doorman who witnessed the attack described seeing her punch the other woman in the face to leave her bloodied. Graham Howton said he had seen her demanding to know who had her handbag before throwing the punch.*
	−3.20	Reddit	*I have a few artists in mind, but I am not sure what the subject should be. I would like to have something meaningful that symbolizes growth; i don't really know of any nonreligious symbols/art/designs that would be good for that.*
PC3	+6.14	Literature	*The sun peeks out long enough to highlight the changing hues, but grayish clouds are skittering across the sky, pushed by a blustery wind that is quickly relieving all the ash trees of their gorgeous golden leaves. As I walk to the post office, my hair blows in my face and the leaves swirl and dance into piles by the sidewalk.*
	−5.00	Reddit	*There was a lot of abuse and violence towards me during my early years and I've had to deal with the ramifications, but in order for me to heal I needed to forgive her for what she did, which ultimately means letting the mother figure I craved in my life die.*

Table A5: Representative passages at the positive and negative extremes of each principal component. High-scoring passages exemplify the narrative quality captured by each axis while low-scoring passages illustrate its absence.

B.4 Narrative Feature Distribution

This section contains additional plots used for the analysis. Figs. A11, A12, and A13 show the mean z-scored narrative features.

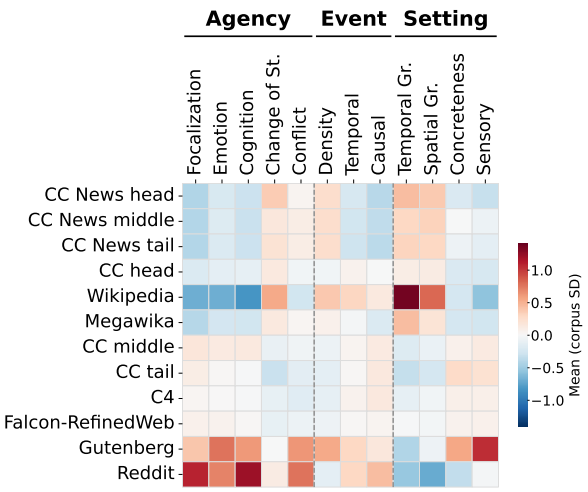

Figure A11: **Mean** z-scored narrative features by **DOLMA source**, ordered by hierarchical clustering on profile similarity.

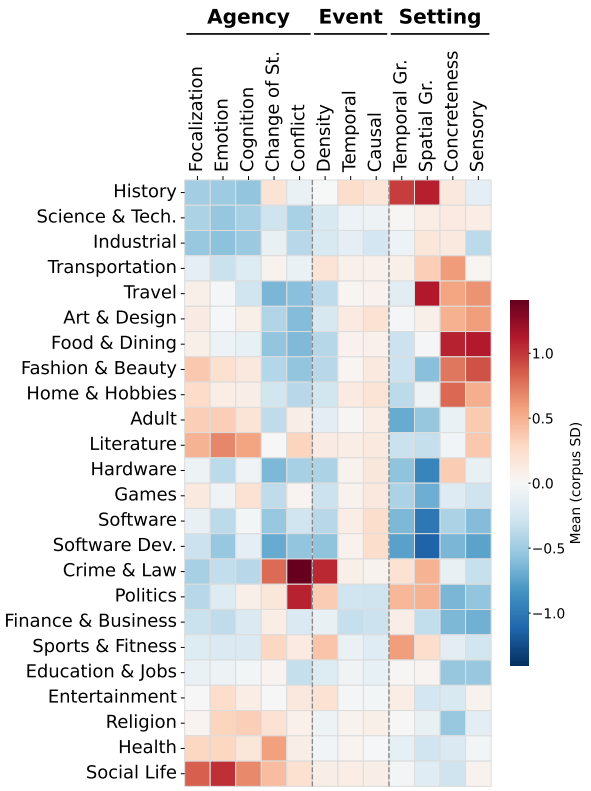

Figure A12: **Mean** z-scored narrative features by **topic**, ordered by hierarchical clustering on profile similarity.

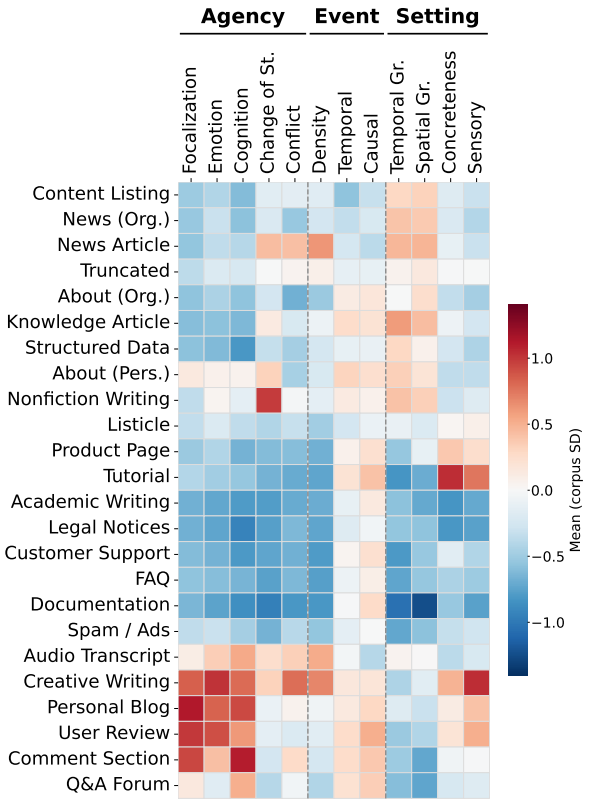

Figure A13: **Mean** z-scored narrative features by **format**, ordered by hierarchical clustering on profile similarity.

Figs. A14, A15, and A16 show the standard deviations (SDs) of the z-scored narrative features.

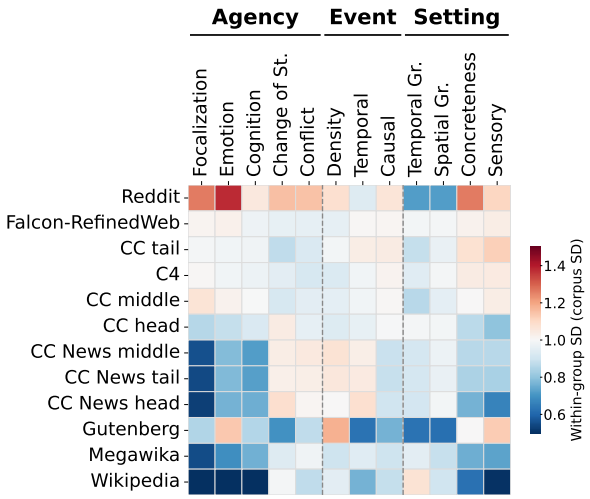

Figure A14: **SD** of z-scored narrative features by **source**, ordered by hierarchical clustering on profile similarity.

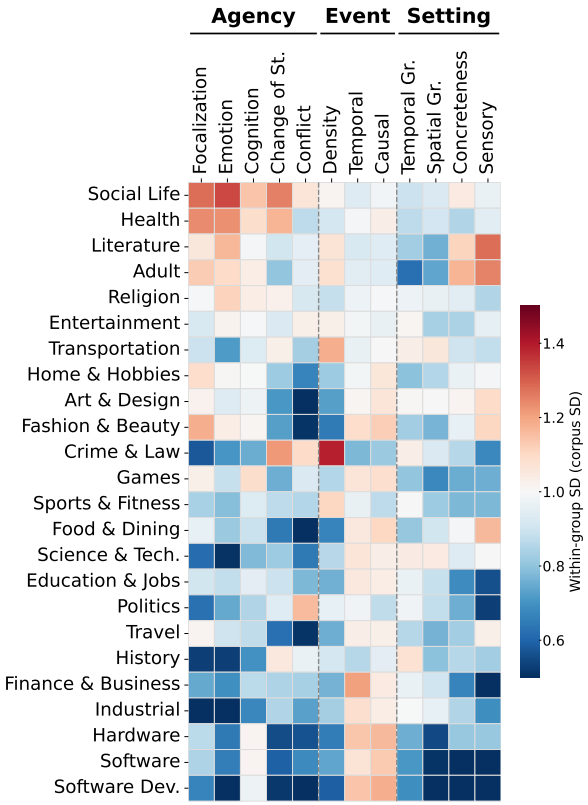

Figure A15: **SD** of z-scored narrative features by **topic**, ordered by hierarchical clustering on profile similarity.

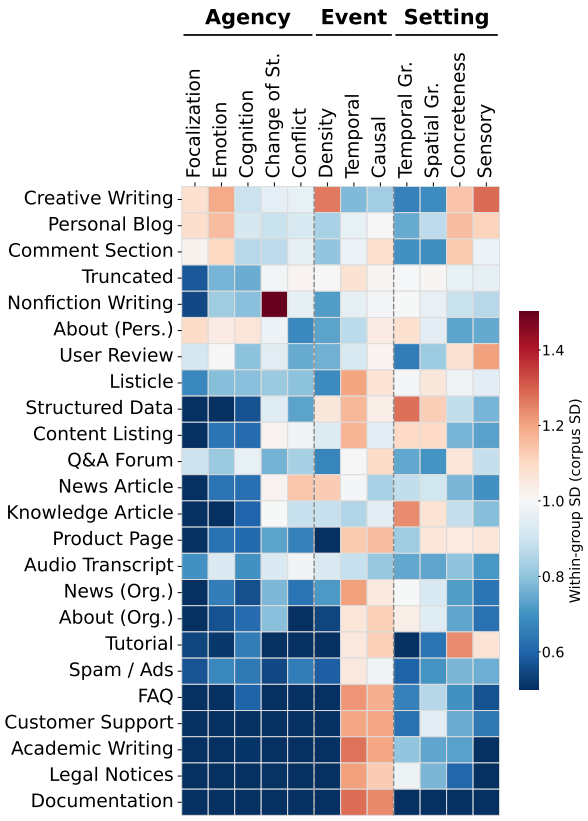

Figure A16: **SD** of z-scored narrative features by **format**, ordered by hierarchical clustering on profile similarity.

C Sampling & Pipeline

This appendix section details how passages were drawn from DOLMA: the source and topic distributions of the sample, the sources we excluded and why, the full sampling configuration, and the validation of the event-span detector that underlies our event features. The final sample makes up the NarraDolma corpus.

C.1 Dolma Source Distributions

Figures A17 and A18 show the distributions of Dolma sources and expanded categories, respectively. Table A6 shows the distribution of the gold set and NARRADOLMA corpus.

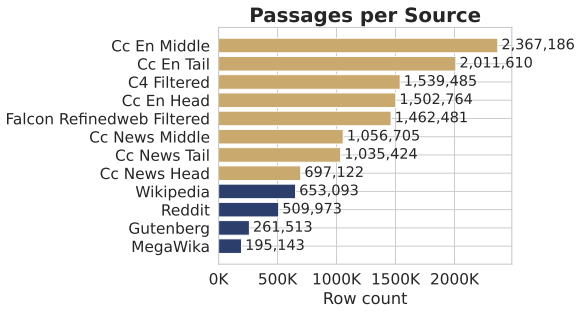

Figure A17: Distribution of Dolma sources that we sampled for NARRADOLMA. These values are at the level of passages.

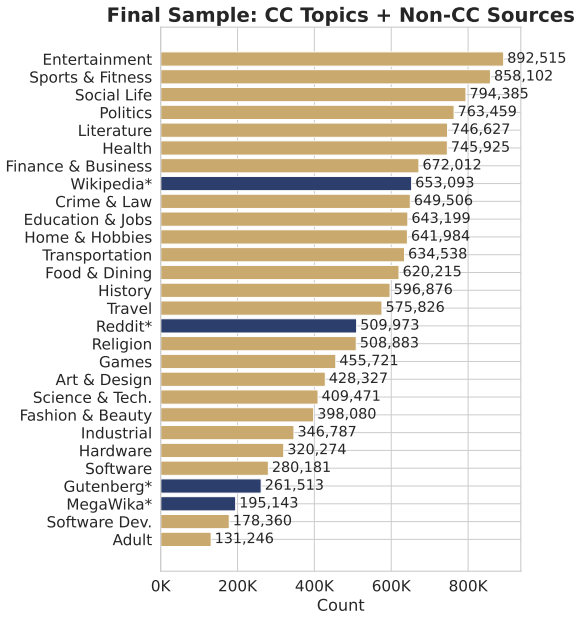

Figure A18: Distribution of Common Crawl topics and non-CC Dolma sources. Tan colors are the aggregated Common Crawl-related Dolma sources partitioned by topic from the WebOrganizer topic classifications. These values are at the level of passages.

C.2 Excluded Sources from Dolma

Table A7 lists Dolma sources that were explicitly excluded (sampling weight set to zero) and the rationale for each.

Source	Target	Gold	NARRADOLMA
CC head	10%	49	1,502,764
CC middle	16%	72	2,367,186
CC tail	16%	66	2,011,610
C4 (filtered)	11%	35	1,539,485
Falcon-RefinedWeb	10%	42	1,462,481
CC News head	5%	14	697,122
CC News middle	8%	27	1,056,705
CC News tail	8%	27	1,035,424
Wikipedia	4%	8	653,093
Megawika	2%	4	195,143
Project Gutenberg	5%	21	261,513
Reddit	5%	35	509,973
Total	100%	400	13,292,499

Table A6: Number of 3-sentence passages sampled across DOLMA sources. *Target* is the normalized sampling proportion, the same proportions used in DOLMA. *Gold* is the per-source allocation for the human-annotated sample. NARRADOLMA is the allocation for the full automatically annotated corpus.

Source	Reason for exclusion
Proof Pile 2 (Algebraic Stack)	Formal mathematics
Proof Pile 2 (OpenWebMath)	Formal mathematics
StarCoder	Source code
RedPajama-arXiv	Academic abstracts
RedPajama-Stack Exchange	Q&A, code snippets
PeS2o	Scientific papers
FLAN-TULU	Instruction templates

Table A7: Dolma sources excluded from sampling.

C.3 Hyperparameters

Table A8 lists all configuration parameters used across the pipeline steps.

Parameter	Value
Initial pool size	51M passages
Unique documents	15M documents
Chunk length	3 sentences
Narrative classifier	RoBERTa
batch size	64
max input length	512 tokens
Narrative confidence threshold	$p \geq 0.50$
Topic classifier	WebOrganizer
batch size	32
max input length	512 tokens
LLM model for annotators	Llama-3.1-8B-Instruct
quantization	4-bit NF4 (bfloat16)
snippet length	1,200 chars $\times$
max new tokens	120
temperature	0.2
top-p	0.9
summary character cap	300 chars

Table A8: Full configuration for the Dolma sampling pipeline.

RoBERTa Training. We fine-tune roberta-base with nine task-specific linear regression heads, one per narrative dimension, appended to the [CLS] token representation. Training uses 5-fold cross-validation on the gold-annotated instances. Within each fold we optimize with AdamW (learning rate 2×10^{-5}, weight decay 0.01) using a linear warmup schedule over the first 10% of training steps, a batch size of 16, and a maximum sequence length of 200 tokens. Training runs for up to 20 epochs with early stopping (patience = 3) monitored by validation MAE on the held-out fold; the final model is then retrained on all gold instances for the median number of best epochs across folds. The loss is masked MSE, which ignores any NaN-valued labels in the multi-task objective. Gradients are clipped to a maximum norm of 1.0. All experiments use random seed 42. Training was performed on an NVIDIA H100 GPU (40 GB MIG partition) on the CU Boulder Blanca research computing cluster. RoBERTa-base is roughly 125M parameters, the 9 regression heads add $\sim$7K.

RoBERTa Inference. Inference over the gold evaluation set uses a batch size of 64 with the same tokenizer settings as training (max length 200, padding to max length, truncation). Raw regression outputs are clipped to the valid Likert range [1, 5] before computing metrics.

LLM Inference. All three LLMs receive identical prompts: the same system prompt and user message are used across models, with scores elicited as integers on a 1–5 scale. No temperature or sampling parameters are explicitly set; each model uses its provider default.

- Claude Sonnet 4.6 is called via the Anthropic Message Batches API with forced tool use (tool_choice: annotate_narrative, max_tokens: 512). The system prompt and tool definition are cached with ephemeral prompt caching to reduce latency and cost.
- Gemma 4 31B and Qwen3-235B-A22B are called via the Doubleword AI async responses API (service_tier: flex, background: true) with 20 concurrent workers. Scores are parsed from the model's plain-text JSON output, with a regex fallback extractor for malformed responses. Gemma 4 31B is roughly 31B parameters and Qwen3-235B-A22B is roughly 235B total parameters (22B active).

C.4 Event Span Detection

Event trigger spans are the lexical anchors around which event meaning is organized (Sims et al., 2019). We detect two complementary types of spans in each passage in order to verify the use of the LitBank event detector for our dataset:

Event spans. Event spans are identified using a RoBERTa-based event detector trained on LitBank (Sims et al., 2019), a corpus of literary text annotated with event mentions following ACE event type definitions. On average, each passage contains 2.4 event spans.

Verb spans. We additionally extract verb spans using the SPACY EN_CORE_WEB_TRF pipeline[5]. On average, each passage contains 6.9 verb spans. Verb spans that overlap with any event span are discarded to avoid redundancy.

Span Pair Selection. From each passage's combined pool of event and verb spans, we select exactly one adjacent pair for manual annotation. Ultimately we find that the LitBank event detection achieves an F1 of 85%, with a precision of 90%. We move forward with the LitBank event detection model as our event detector.

D Multilingual Distribution Comparison

Our work is conducted over the English language. To explore potential multilingual capabilities of our model, we compare the distribution of scores of a sample of Common Crawl web text across English, Arabic, German, and French. Samples are drawn from FineTranslations (Penedo et al., 2026), a release of FineWeb2 machine-translated into English. This dataset contains text that has already been translated from each of these languages to English, so our model is given just English text.

For each language we sample 10,000 documents, which are then sentence-tokenized and cut into 3-sentence chunks. We then keep up to three passages per document which mimicks our Dolma sampling procedure. We apply the topic, narrative classifier, and event detection models to each passage. Fig. A19 shows the distributions of each language across all 12 narrative features. In general, German and French are closest to English across all features. We do notice that the translated Arabic texts have higher spatial grounding.

[5]https://spacy.io/models/en

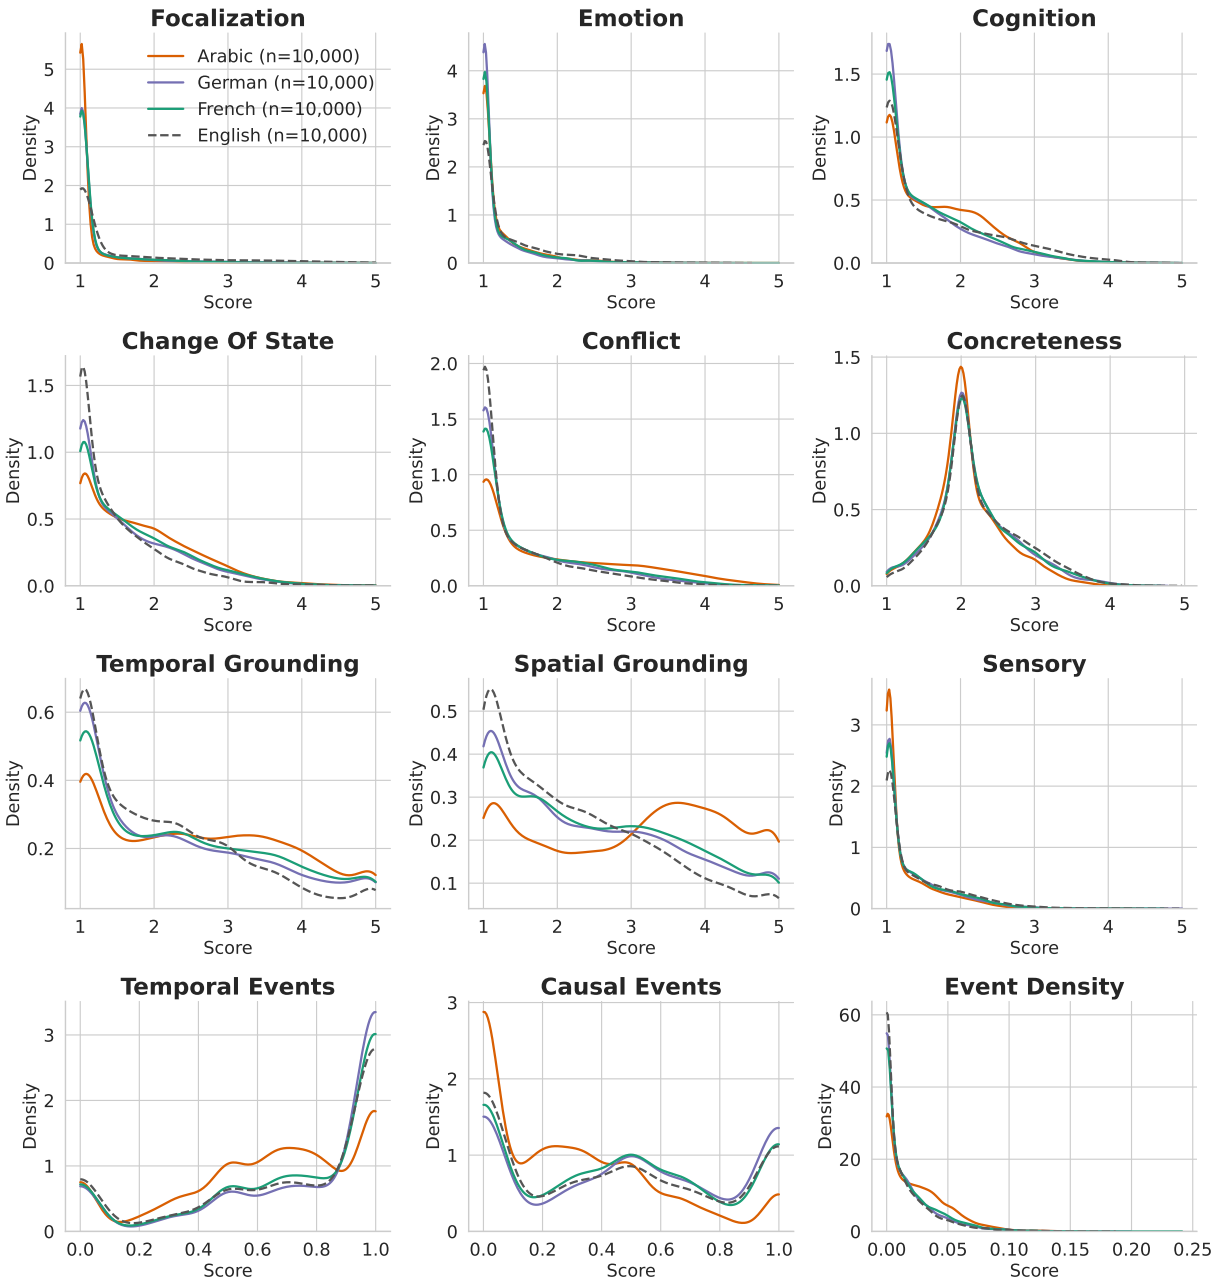

Figure A19: Distribution of the 10K documents' narrative scores across English, Arabic, German, and French.

E Annotation Framework

This appendix documents the full annotation framework: the rating scale for each of the 11 dimensions, the interface annotators used, the document summaries shown to orient them, and the prompts used to elicit the LLM annotations.

E.1 LLM-Generated Document Summaries for Human Annotation

To orient annotators to the genre and topic of each passage without biasing their event, agency, and setting judgments, we pair each chunk with a short machine-generated description of the originating document. We run Llama-3.1-8B-Instruct (Dubey et al., 2024) in 4-bit NF4 quantization over each retrieved document, providing the model with three 1,200-character snippets drawn from the start, middle, and end of the full text.

The system message provided to Llama-3.1-8B-Instruct for summary generation is:

> *"You write short, extremely surface-level summaries of webpages. Do not speculate or interpret. Only describe what is explicitly present."*

E.2 Potato Interface

Figures A20, A21, and A22 show the UI that annotators use to label the sampled texts for agency and setting, respectively.

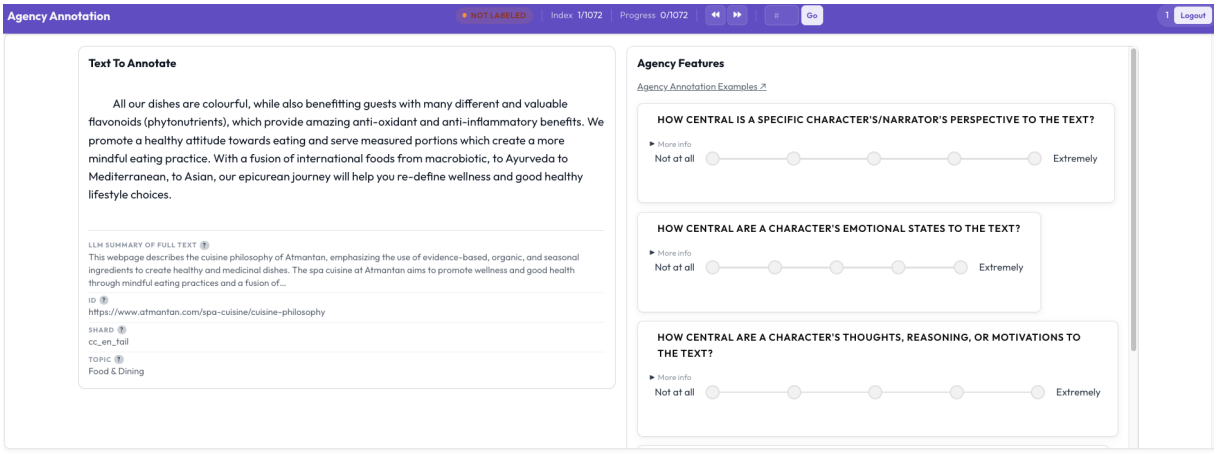

Figure A20: Screenshot of the agency annotation UI built upon the Potato software.

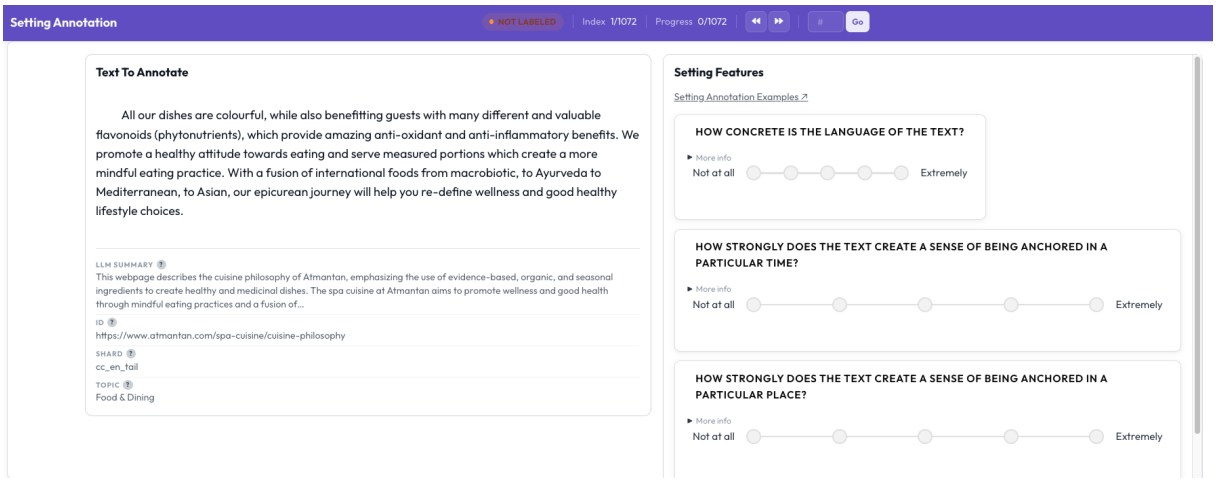

Figure A21: Screenshot of the setting annotation UI built upon the Potato software.

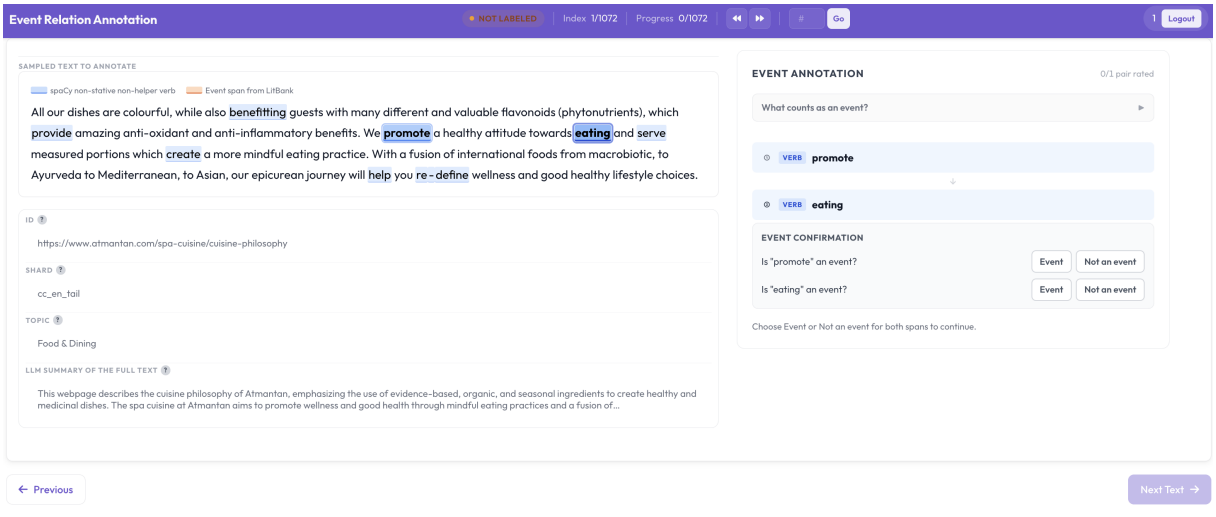

Figure A22: Screenshot of the event relation annotation UI built upon the Potato software.

E.3 Annotator Feature scales

Tables A9-A17 show the feature scales given to annotators for each of the 11 fine-grained narrative features.

Score	Description
1	Not central. The passage describes events and details from a purely external vantage point with no access to any character's inner life.
2	Minimally central. A character's perspective is faintly implied but does not shape how events are presented. The passage remains largely external.
3	Moderately central. A character's perspective is present and shapes some of the passage, but external description and internal access are roughly balanced.
4	Considerably central. Events are predominantly filtered through a character's perceptions and inner life, with only occasional external description.
5	Extremely central. The passage is fully organized around a specific character's or narrator's perspective. Every detail is filtered through their perceptions, feelings, and inner experience.

Table A9: Focalization rating scale.

Score	Description
1	Not central. The passage contains no reference to how a character feels, either through observable behavior or internal access.
2	Minimally central. A character's emotional state is faintly implied through observable behavior (e.g., a gesture or expression), but feelings are not elaborated or sustained.
3	Moderately central. A character's emotional states are present and noticeable, but emotion is not the dominant or organizing feature of the passage.
4	Considerably central. Emotional experience is prominent and shapes how the passage is organized. The reader has meaningful access to how a character feels, either through behavior or internal description.
5	Extremely central. Emotional experience dominates the passage. The reader is given sustained and direct access to a character's feelings, which are foregrounded throughout.

Table A10: Internal emotion rating scale.

Score	Description
1	Not central. The passage contains no reference to what a character thinks, believes, reasons, or wants.
2	Minimally central. A character's mental state is faintly implied but not elaborated. Cognition is incidental to the passage.
3	Moderately central. A character's thoughts or reasoning are present and contribute to the passage, but are not its dominant or organizing feature.
4	Considerably central. A character's beliefs, deliberations, goals, or reasoning are prominent and shape how the passage is organized.
5	Extremely central. The passage is dominated by a character's inner cognitive life including their thoughts, reasoning, beliefs, and desires.

Table A11: Internal cognition rating scale.

Score	Description
1	Not central. Characters remain essentially unchanged across the passage. No change in physical, psychological, relational, or existential condition is present or implied.
2	Minimally central. A minor or incidental change is present or implied, but it is not an organizing feature of the passage.
3	Moderately central. A change in a character's condition is present and contributes meaningfully to the passage, but is not its dominant feature.
4	Considerably central. A change in a character's condition – physical, psychological, relational, or existential – is prominent and shapes how the passage is organized.
5	Extremely central. Change is the dominant and organizing feature of the passage. The transformation (whether completed, in progress, or strongly implied) is what the passage is fundamentally about.

Table A12: Change of state rating scale.

Score	Description
1	Not central. No conflict of any kind is present or implied. The passage describes actions or states without opposition or tension.
2	Minimally central. A hint of tension or opposition is present but remains entirely in the background and does not shape the passage.
3	Moderately central. Conflict is present and noticeable – between characters, within a character, or against external forces – but is not the dominant feature of the passage.
4	Considerably central. Conflict is prominent and shapes how the passage is organized, whether interpersonal, internal, or against external forces or circumstances.
5	Extremely central. Conflict dominates the passage. The opposition or tension is what the passage is fundamentally about, and it is sustained throughout.

Table A13: Conflict rating scale.

Score	Description
1	No physical referents. The passage is composed entirely of abstract ideas, generalizations, categories, or institutional descriptions. Nothing could be explained by pointing.
2	Physical objects or entities are named but not rendered. The text identifies things that exist in the world but provides no perceptual detail. No shape, color, texture, or material quality.
3	Physical referents are present with some rendering. Objects and actions are described with enough specificity to suggest a physical scene, but perceptual detail is sparse or thin.
4	Physical referents are rendered with clear perceptual detail. The passage provides enough specific, tangible description (through surface qualities, physical components, or rendered actions) that the reader can picture the scene.
5	Physical referents are richly and fully rendered. The passage is dense with perceptual particulars. Objects, environments, and actions are described with enough sensory and material detail that the world feels concretely present.

Table A14: Concreteness rating scale.

Score	Description
1	No temporal content. The passage has no markers of time whatsoever; events could be happening at any point in history or none at all.
2	Temporal language present but only sequencing events. Relative markers (e.g., *before*, *after*, *recently*, *a few years ago*) order events without locating the reader in any particular time.
3	Modest temporal location. The passage provides cyclical grounding – a season, a time of day, a day of the week – that gives the reader a felt sense of when without historical specificity. Most cyclical-only passages cap here.
4	Strong temporal location. The passage provides historical grounding – a specific year, era, named event, or cultural period – that anchors the reader in an identifiable moment. In rare cases, rich and sustained cyclical grounding can reach this level.
5	Vivid and sustained temporal immersion. Both historical and cyclical grounding work together, creating a strong and atmospheric sense of a specific moment in time.

Table A15: Temporal grounding rating scale.

Score	Description
1	No spatial content. The passage has no markers of place; events could be happening anywhere or nowhere.
2	Spatial language present but minimal. Either a bare proximate location with no rendering (e.g., *she sat in the kitchen*) or a vague geographic reference (e.g., *somewhere in Europe*). Places the reader no further than knowing a space exists.
3	Modest spatial location. Either a named geographic location that places the reader on the globe without evoking it (e.g., a city name), or a proximate location with some modest rendering that begins to create a felt sense of place.
4	Strong spatial location. Either geographic grounding with atmospheric texture, or a proximate environment rendered vividly enough that the reader feels present in the space.
5	Vivid and sustained spatial immersion. Both geographic and proximate grounding work together with rich rendering, creating a strong and atmospheric sense of a specific place.

Table A16: Spatial grounding rating scale.

Score	Description
1	No sensory content. Events and ideas are conveyed without any appeal to the senses; no modality is invoked.
2	Sensory content present but incidental. One or more senses are mentioned or implied, but the sensory experience is peripheral to the passage's meaning. Removing the sensory language would not substantially change the passage.
3	Sensory content present and noticeable. At least one sense modality recurs or contributes meaningfully to the passage, but sensory experience is not the organizing feature of the text.
4	Sensory experience is central. One or more senses are prominent and clearly organize how the passage is structured, even if the sensory content is not rendered in rich perceptual detail.
5	Sensory experience is dominant and richly rendered. The passage is organized around sensory experience across one or more modalities, with sustained and detailed evocation of what it feels like to perceive the described world.

Table A17: Sensory detail rating scale.

E.4 LLM Annotation Prompt: Agency and Setting

Fig. A23 displays the prompt used to elicit agency annotations and Fig. A24 the prompt for setting annotations from Claude Sonnet 4.6, Qwen3-235B-A22B, and Gemma4-31B. The prompt was held constant across all three models and across LLM validation and large-scale annotation.

You are an expert at identifying narrative qualities in web text. Your task is to read a passage and rate it on nine dimensions across two categories (Agency and Setting) each on a scale from 1 (not at all) to 5 (extremely).

AGENCY DIMENSIONS

These five dimensions capture how characters are represented as agents in the text. Importantly, these dimensions rate the presence and centrality of the feature in the text as written, not what can be inferred about the content.

1. Focalization

How central is a specific character's/narrator's perspective to the text?

A higher score reflects text in which the details and events are primarily presented through a specific character's perspective: their perceptions, thoughts, and feelings shape how the story is told. A lower score reflects text that reports details and events from the outside, describing what is observable without granting access to any character's internal experience.

Focalization is about whether the reader experiences events through a specific perceiving consciousness's inner life, regardless of the formal mechanism. That can be achieved through: direct quotation (if the content renders inner experience), free indirect discourse (*"she couldn't shake the feeling..."*), first person narration with genuine interiority, or close third person narration. When "you" appears in instructional or legal text, it should score low on focalization. However a specific second-person perspective that immerses the reader can score high.

- Score 5: *"she scanned the faces in the crowd, certain that someone was watching her"*
- Score 3: *"He wasn't sure the meeting had gone well. He gathered his things and headed for the door."*
- Score 1: *"she walked through the crowd."*

2. Emotion

How central are a character's emotional states to the text?

A higher score reflects text in which a character's emotional experience is a prominent feature. A lower score reflects text with little or no reference to how a character feels.

- Score 5: *"Completely paralyzed with fear, he was all of a sudden flooded with relief when he heard her voice"*
- Score 3: *"She was nervous about the presentation. She went into the interview room."*
- Score 1: *"he answered the phone."*

3. Cognition

How central are a character's thoughts, reasoning, or motivations to the text?

A higher score reflects text in which a character's thoughts, beliefs, reasoning, goals, or desires are a prominent feature. A lower score reflects text with little or no reference to what a character thinks, intends, or wants.

- Score 5: *"she kept turning the problem over in her mind, convinced there was something she had missed. Was it because they were so quiet in the meeting? Or their body language? She couldn't tell."*
- Score 3: *"She read the email twice. It wasn't entirely clear what they were asking for, but she thought she understood the gist. She drafted a reply."*
- Score 1: *"she sat at her desk"*

4. Change of State

How central is a change in a character's condition or state to the text?

A higher score reflects text in which a change in a character's condition/state is a prominent or organizing feature. A lower score reflects text where characters remain essentially unchanged. The change does not need to be completed within the passage; an in-progress or partially implied change still counts. If a passage describes a world event (a company going bankrupt, an earthquake, a death) that necessarily entails a change in a character's condition, that counts toward the score even if the character's internal response is not elaborated. Change of state should be rated on presence and centrality of the changes themselves, independent of how vividly or dramatically they are rendered.

- Score 5: *"by the time she reached the door, something in her had shifted. She wasn't the same person who had walked in"*
- Score 3: *"By the end of the conversation he felt somewhat better about the situation. It wasn't resolved, but it seemed more manageable than it had before."*
- Score 1: *"she sat at her desk."*

5. Conflict

How central is conflict involving characters to the text?

A higher score reflects text in which conflict is a dominant or organizing feature. A lower score reflects text in which conflict is absent or only incidentally present. Conflict can take many forms: tension between characters, internal psychological struggle, or opposition to institutions, environments, inanimate objects like technology, social forces, or physical events.

- Score 5: *"they had been arguing for hours, neither willing to give ground. The people around them started to stare. They were so caught up in the fight that they didn't even notice."*
- Score 3: *"They disagreed about the route. They took the highway and didn't talk much for the first hour."*
- Score 1: *"they sat across from each other at the table."*

Figure A23: Annotation prompt for agency.

SETTING DIMENSIONS

These four dimensions capture how the text constructs a sense of place, time, and physical presence.

6. Concreteness
How concrete is the language of the text?
A higher score reflects text in which most content could be explained by pointing, demonstrating, or showing (objects, physical states, bodily actions, perceptual qualities). A lower score reflects text in which most content can only be explained using other words (categories, principles, relationships, institutions, ideas). The key distinction is not just whether physical things are present, but how much the text renders them. When a passage enumerates many specific, tangible, pointable things, the cumulative effect creates strong concrete presence.

- Score 5: *"the chipped blue mug sat on the edge of the sink, handle cracked from years of use"*
- Score 3: *"She had been commuting to the same office for six years. The building was on a corner downtown, glass and steel, indistinguishable from the ones on either side of it."*
- Score 1: *"people often form attachments to everyday objects"*

7. Temporal Grounding
How strongly does the text create a sense of being anchored in a particular time?
A higher score reflects text that creates a vivid sense of temporal location. A lower score reflects text that feels as if events could be taking place at any time. Two types contribute to the score: *historical grounding* locates the reader in a specific, unrepeatable moment (a year, an era, a named event); *cyclical grounding* locates the reader within a recurring temporal structure (a season, a time of day). Historical grounding can reach the upper end of the scale even with sparse language; cyclical grounding alone typically reaches a 3.

- Score 5: *"It was the summer of 2015 at 9:02am, before the world changed forever."*
- Score 3: *"This unique summer school gives teenagers the opportunity to work with tutors and musicians."*
- Score 1: *"he walked outside."*

8. Spatial Grounding
How strongly does the text create a sense of being anchored in a particular place?
A higher score reflects text that creates a vivid sense of spatial location. A lower score reflects text that feels as if events could be taking place anywhere. Two types contribute to the score: *geographic grounding* locates the reader on the globe (a country, city, named landmark); *proximate grounding* locates the reader in an immediate physical environment (a room, a building, a street). Geographic grounding can reach a 5 without proximate rendering; vividly rendered proximate grounding without geographic grounding caps at a 4.

- Score 5: *"In the narrow streets of Naples, the apartment's tiled floors kept cool even in August heat."*
- Score 3: *"She worked at an office in London." or "The kitchen was small, cluttered with dishes."*
- Score 1: *"She decided to quit her job."*

9. Sensory
How central are sensory details to the text?
A higher score reflects text in which one or more senses are central to how the passage is organized, where the sensory experience drives, anchors, or sustains the content. A lower score reflects text where events and ideas are conveyed without meaningful appeal to the senses. What matters is both whether a sense modality is central to the passage, and whether it is vividly described.

- Score 5: *"the bread was still warm, its crust crackling under her fingers, the whole kitchen thick with the smell of it"*
- Score 3: *"The bread was still warm when she pulled it from the oven. She set it on the rack and went to wash her hands."*
- Score 1: *"she made breakfast which consisted of bread and eggs."*

Scoring Scale: 1 = Not at all 2 = Slightly 3 = Moderately 4 = Considerably 5 = Extremely

Figure A24: Annotation prompt for setting.

E.5 LLM Annotation Prompt: Event Relations

Fig. A25 displays the prompt used to elicit temporal ordering and causal relation annotations from Claude Sonnet 4.6 and Gemma4-31B at scale. The prompt was held constant across both models and across LLM validation and large-scale annotation.

TEMPORAL ORDERING

Which event started first in the narrative timeline?

- **span1_first** — SPAN1 started before SPAN2
- **span2_first** — SPAN2 started before SPAN1
- **simultaneous** — Both events genuinely share a start point and are distinct events
- **same_event** — The two spans refer to the same event from different angles
- **too_hard_to_tell** — The start times could feasibly be in either order, or the events occur in completely unrelated temporal frames

If either span is not an event, use **not_applicable**.

Inferring order. Use explicit connectives ("before," "after," "then," "when," "while," "as"), backward-looking cues ("preceded," "followed," "prior to"), and logical necessity. A cause precedes its effect, a question is asked before it is answered, a proposal precedes an agreement. When one event logically presupposes the other, infer the ordering.

Reporting and retrospective verbs. Story-world events precede the speech acts or memory acts that describe them, even when the reporting verb appears first in the text. *"The coach expressed postgame that losing the player had impacted the offense"* → span2_first. Memory verbs (*recalls, remembers, reflects*) similarly occur after the events they describe.

Simultaneous vs. sequential. Use simultaneous only when neither event initiates the other and the text marks coincident starts ("just as X, Y"). When "as" or "while" appears, ask whether both events genuinely begin at the same instant (→ simultaneous) or whether one was already underway when the other began (→ span1_first or span2_first). Process/result pairs and physical chains are sequenced even when compressed.
Example: *"While touring in Germany, he released a record."* → span1_first. Touring was already underway when the release happened; "while" marks an ongoing background activity, not a coincident start.

Same event. Use same_event when both spans describe the same occurrence and neither adds new information the other omits. Continuation verbs (*adding, continuing, noting*) extend a prior speech act → use span1_first.

Too hard to tell. Use when: (a) two events describe different aspects of the same episode with no logical ordering between them; (b) the events involve different unrelated actors with no shared context; or (c) the ordering requires world knowledge inference the text does not support. If you find yourself inferring order from narrative position or plausible sequence rather than explicit textual evidence, use too_hard_to_tell. Uncertainty is not a failure.

CAUSAL RELATION

Given a confirmed temporal ordering, how are the two events causally related?

- **direct_cause** — E1 is sufficient to produce E2; given E1, E2 was bound to happen without any intervening decision or action.
  - Physical chain: *"She dropped her phone. The screen cracked."*
  - Involuntary reaction: *"He read the message. His stomach dropped."*
- **enables** — E1 creates a necessary precondition for E2: if E1 had not happened, E2 could not have happened in the same way. The dependency must be traceable in the text; mere co-occurrence is not enough. Enablement can run in either direction.
- **not_related** — No traceable causal dependency. Events that are merely co-present in the narrative without a direct conditional link.

Story-world vs. discourse-world. Story-world events (actions, occurrences, mental states) and discourse events (*said, reported, wrote, announced*) operate at different levels and cannot causally relate to each other. A fire does not cause a spokesperson's statement; the statement is an independent communicative act. When one span is a reporting verb and the other is a story-world event, use **not_related**. Two discourse events can causally relate to each other.

Distinguishing tip. Was E2 bound to happen given E1 alone? → direct_cause. Did E1 create a clear necessary precondition? → enables. Otherwise → not_related.

Figure A25: Annotation prompt for event relations.