

The Shape of Testimony: A Scalable Framework for Oral History Archive Comparison

Itamar Trainin, Renana Keydar, Amit Pinchevski

Hebrew University of Jerusalem
{itamar.trainin, renana.keydar, amitpi}@mail.huji.ac.il

Abstract

Researchers in Holocaust studies have often distinguished between two styles of oral survivor testimony: the USC Shoah Foundation's interviews tend to follow a structured, interviewer-guided format, whereas the Yale Fortunoff Video Archive generally favors a more free-form, open-ended style. This distinction has influenced both scholarly research and the development of later archives. In this study, we critically examine that claim by conducting a large-scale computational analysis of more than 1,600 testimonies from both collections. Leveraging discourse segmentation, topic modeling, and large language model (LLM) based analysis, we quantify the "structuredness" level of testimonies through topic coherence, interviewer-survivor dynamics, and the distribution of question types. Our results generally corroborate the structural differences identified in earlier research, while also revealing significant overlaps between the collections, both within individual interviews and across common narrative patterns. This complicates the simple "structured vs. free-form" dichotomy often applied to these oral histories. Beyond revisiting a foundational claim in Holocaust studies, our work provides a scalable, replicable framework for comparative corpus analysis. As a proof of concept, it suggests broader applications for digital oral history, narrative analysis, and the design of citizen-science annotation platforms.

Keywords: Holocaust Studies, Oral History Language Resources, Computational Archive Comparison, Large language models (LLMs)

1. Introduction

A common claim in the field of Holocaust Studies suggests that the two main oral history video testimony archives, the Yale Fortunoff Video Archive for Holocaust Testimonies (Henceforth: Yale or Fortunoff) and the USC Shoah Foundation (Henceforth: USC or Shoah Foundation), present two opposite styles of interviews. It is said that the interview practices by Yale tend to be loose and open-ended, whereas USC follows a stricter format. Several studies have ascertained this difference by exploring the institutional agenda and policy as well as through close readings of selected testimonies (Wieviorka, 2006; Shenker, 2015; Pollin-Galay, 2018) and recently also by a small-scale empirical study (Presner et al., 2024).

Yet to empirically evaluate these assumptions, a large-scale comparative analysis is required, which is the purpose of this paper. This analysis calls for advanced computational tools, such as those developed recently in the digital humanities and computational history, for examining large narrative corpora. To date, applying these tools to oral testimonies has been limited due to the interpretive challenge and ethical complexities involved.

In what follows, we seek to ascertain the reported difference between the two testimony styles by means of a computational comparison of over 1600 Holocaust survivor testimonies drawn from the USC and Yale archives. Using large language models (LLMs), dialogical segmentation, topic modeling, and question classification, we develop a

reproducible pipeline to measure "structuredness" across interviews. We examine topical coherence, interviewer-survivor dynamics, and the evolution of question types along interviews to assess whether and how these archives actually diverge.

Our findings confirm that USC testimonies tend to follow a more guided and structured format, especially in the early parts of the interview. Yale testimonies, by contrast, display greater topical fluidity, earlier emergence of core themes, and a higher proportion of open-ended questions. However, our analysis also uncovers significant convergences across both collections in later stages of the interviews, as well as similar narrative arcs, which seem to follow the chronology of Holocaust experience.

This article thus makes a dual contribution: (1) it offers a data-driven reassessment of a central historiographical claim within Holocaust testimony scholarship; (2) it introduces a scalable framework for computational comparison of large oral history corpora. By demonstrating how complex, ethically sensitive narratives can be analyzed computationally without compromising their integrity and interpretive richness, we open a pathway for further digital humanities work in trauma studies, oral history, and memory research.

2. Oral history archives of the Holocaust

The Yale Fortunoff Video Archive for Holocaust Testimonies and the USC Shoah Foundation Visual

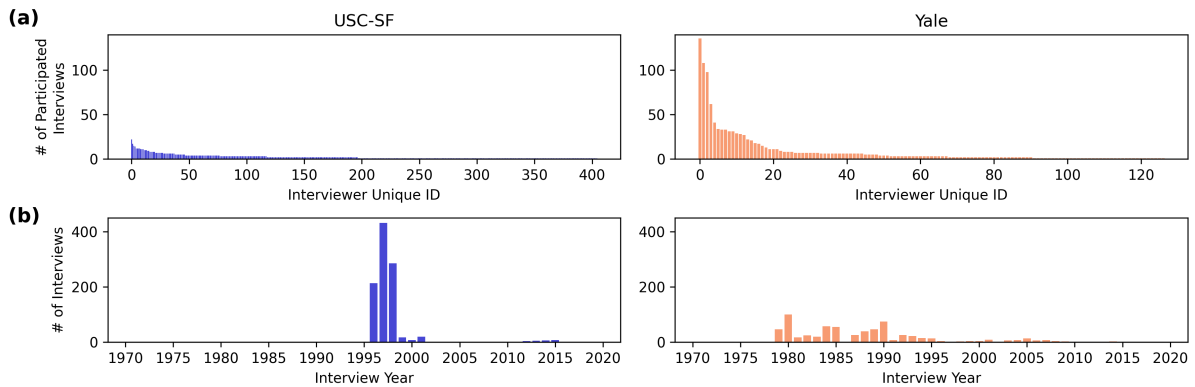


Figure 1: (a) Number of interviews conducted (individually or jointly) by each interviewer in each archive. (b) Annual distribution of testimonies recorded across both archives, computed over the subset of testimonies available for this study.

History Archive are the two main and largest collections of video testimonies of Holocaust survivors. The Fortunoff Archive began in 1979 as a local initiative in New Haven, Connecticut, later relocating to Yale University (Hartman, 1995). Holding more than 4,400 testimonies collected in the US, Europe, and Israel, its mission is to collect, preserve, and share the stories of those who were there. Being the first of its kind, the archive set the standard for projects to come (Felman and Laub, 1992a; Hartman, 1996). The USC archive was founded in 1994 by Steven Spielberg with the goal of recording 50,000 testimonies, for which hundreds of interviewers were trained in multiple countries (Smith, 2022a). It has since surpassed that goal, and in addition to Holocaust testimonies, has expanded to record victim narratives of other genocides (including Armenian, Rwandan, Cambodian) (Assmann, 2011; Smith, 2022a). Both archives were initially based on videotape technology and, over time, came to embrace digital technology, while incorporating and developing new methods to access and engage with testimonies (Assmann, 2011; Shandler, 2022).

At its core, an oral history interview is a process in which an interviewer and an interviewee spend extended time together engaged in storytelling and attentive listening. It is a collaborative act of constructing a narrative. Unlike other forms of qualitative interviewing (such as in-depth interviews), oral history interviews tend to be less narrowly focused in their topical organization (Hesse-Biber and Leavy, 2005). As with other oral history initiatives, each of the Holocaust archives has developed its own interview methodology.

From the outset, interviews recorded at Yale followed a clear code of conduct: putting the witnesses at the center and allowing them to relate their story freely with minimal guidance from interviewers. The interviewers' main role was to provide survivors with a supportive listening companion

during the interviews, all of which took place on the Yale campus (Felman and Laub, 1992a). Conversely, USC interviews take place at the witness's home and are typically preceded by a preparation meeting. USC employs a large number of interviewers in many countries, who all go through formal training and follow a set of elaborate and formulated guidelines (USC Shoah Foundation, 2021) (Smith, 2022b)). The Yale archive has relied on a relatively small group of interviewers, mostly academics, some with psychological training. As will be shown below, despite these disparities, interviews recorded by both archives tend to unfold chronologically, beginning with life before the war, through early signs of change, antisemitism and persecution, experiences during the war (ghetto, captivity, forced labor, extermination), end of war, and finally life after the war. All interviews examined here were transcribed manually by their respective archive, together with speaker identification and speaker-side segmentation.

It is hard to imagine contemporary Holocaust history and memory without the prevalence of survivors' testimonies. The two archives have contributed significantly to the legitimacy of personal narratives as a source for the study of recent history, particularly of traumatic events (Friedländer, 1993; Hartman, 1996; Shandler, 2017). By their nature, testimonies are based on personal memory whose accuracy might vary, but what they nevertheless express is something beyond factual knowledge about past events. It is the relating of the personal experience of how it felt, which expresses the human aspect often missing from grand historical accounts (Felman and Laub, 1992b; Langer, 1993). As such, these narratives of pain and loss demand care and respect while retaining the singularity and integrity of each voice. Over the years, numerous audiovisual interviews with survivors have been recorded. The practical impossibility of watching tens of thousands of such testimonies presents an unprece-

mented challenge for a morally informed study of recent history. It is precisely here that carefully designed digital tools can enable new, previously unattainable modes of engagement with testimony (Keydar, 2020, 2022). Recent advances in natural language processing and large language models have opened new algorithmic ways of engaging with thousands of Holocaust testimonies at scale (Blanke et al., 2019; Naron and Toth, 2020; Ezeani et al., 2024; Presner et al., 2024; Shizgal et al., 2025; Keydar et al., 2026). However, to the best of our knowledge, there has so far been no attempt to harness these tools for large-scale comparative analysis across oral history archives.

3. Data and Methodology

This study analyzes 1,668 Holocaust survivor testimonies: 1,000 from the USC Shoah Foundation archive and 668 from the Yale Fortunoff Archive. Each interview was manually transcribed, annotated with speaker labels, and segmented into question–answer (Q/A) pairs, as outlined below. Due to the sensitive nature of these collections, transcripts are accessible only by direct request to the respective institutions. All testimonies in this study were conducted in English.

3.1. Date of testimonies

The two archives developed and followed highly different collection strategies. The Fortunoff Video Archive, established in 1979, has followed an open-ended acquisition policy, resulting in interviews collected over more than four decades. In contrast, the USC Shoah Foundation conducted a large-scale collection effort primarily between 1994 and 1999. However, within our available dataset, the testimonies from both archives are concentrated towards the initial years of operation. Fig. 1(b) shows the distribution of testimonies per year in the subset of testimonies available for this study.

3.2. Interview Length

USC interviews average 23,396 words¹ ($\sigma = 10,397$), while Fortunoff interviews average 13,622 words ($\sigma = 7,649$). Although both archives use similar methodologies, their protocols differ. One notable outcome is the substantial difference in length: USC interviews are, on average, 1.7 times longer than those from Fortunoff. This results in a greater quantity of content and potentially richer detail in the USC interviews. To mitigate the influence of this disparity, our quantitative metrics normalize for testimony length.

¹A word defined as any whitespace-separated token.

3.3. Distribution of interviewers

Top	USC		Yale	
	Name	#	Name	#
1	Lorrie Fein	22	Vlock Laurel	136
2	Esther Finder	17	Laub Dori	108
3	Joanna Buchan	15	Kline Dana	98
4	Reuben Zylberszpic	14	Rudof Weiner	62
5	Joseph Huttler	12	Millen Susan	41
6	Dina Cohen	12	Frances Cohen	34
7	Ruth Meyer	12	Langer Lawrence	33
8	Zepporah Glass	12	Herz Moss	33
9	Florence Shuster	11	Strochlic Kathy	31
10	Yvonne Walter	11	Katz Helen	31

Table 1: Top 10 participating interviewers and corresponding number of interviews conducted by the interviewer either as an individual or as part of a joint interview.

Fig. 1(a) shows the number of interviews each interviewer conducted in each archive. The Fortunoff corpus engaged a small cohort of interviewers, some of whom conducted over 100 interviews each, reflecting an approach that values a more personal interviewing style. In contrast, no individual interviewer in the USC corpus conducted more than 22 interviews (see top 10 participating interviewers in table 1). This distributed practice aligns with USC archive’s emphasis on standardized interview protocols, which prioritize methodological consistency over personal interviewing style.

3.4. Methodology

Each testimony was segmented into chronological subunits using two complementary strategies, designed to capture both micro- and macro-level structures.

At the micro-level, we employed a topic-oriented segmentation: each Q/A pair was treated as a discrete topical unit, under the assumption (following prior work, (Ifergan et al., 2024)) that individual exchanges typically cover single topics. Short Q/A pairs were merged with adjacent exchanges when appropriate to maintain contextual coherence. This segmentation allowed us to examine the dialogic aspects of interviewer–survivor interaction while retaining the structure of the narrative.

For macro-level analysis, we divided each testimony into $k = 15$ equal-length chronological segments based on prior work ((Ifergan et al., 2024)) and manual validation, enabling a normalized temporal comparison across testimonies of varying lengths and numbers of Q/A pairs. From each segment, we extracted a range of features such as topical diversity, segmental coherence, answer and question lengths, and question type classification using custom pipelines and GPT-based classification.

Thus, in this study, we quantify the differences in the “structuredness” across collections, emphasizing the topical and interviewer–survivor dialogic dynamics. Nonetheless, additional dimensions that may confound a comparative analysis between Yale and USC interviews are discussed in §3.1 - §3.3.

4. Findings

4.1. Topical Sequence Analysis

A central dimension of “structuredness” in oral history interviews is the sequencing of topics across the testimonial timeline. Prior scholarship (e.g., (Piper et al., 2021; Wagner et al., 2022; Ranade et al., 2022; Ifergan et al., 2024; Wagner et al., 2025; Shizgal et al., 2025; Trainin and Abend, 2025)) suggests that highly structured interviews should exhibit a more predictable topical order, narrower chronological coverage per topic, and reduced thematic diversity across segments, whereas free-form interviews are expected to display looser sequencing and greater thematic fluidity.

To evaluate this hypothesis, we apply a computational topic-analysis pipeline modeled on, yet distinct from, earlier manual or semi-supervised approaches. We build upon the framework introduced in (Ifergan et al., 2024), but instead of employing topics generated by LDA (Blei et al., 2003) or BERTopic (Grootendorst, 2022) followed by expansive manual topic labeling and heuristic clustering thresholds, we introduce an alternative LLM-based strategy inspired by (Trainin and Abend, 2025). In doing so, we provide a scalable and fully automatic methodology for identifying emergent topical patterns across thousands of segments.

Our pipeline unfolds in two stages. First, an LLM (ChatGPT) generates a concise descriptive label for every Q/A pair, which serves as the micro-level unit of analysis (see Prompt 7.1). In the second stage (see Prompt 7.2), we prompt the LLM to identify the Top-K recurring topics for each chronological segment. This staged approach captures both the local dynamics of interviewer–survivor exchanges and the broader macro-level narrative structure. Fig. 7 in the appendix presents an illustration of this pipeline.

To validate our methodology, we calculated a Topic Coverage score (Trainin and Abend, 2025) for each inferred topic using a random sample of 50 testimonies. Additionally, we compared the USC topics produced by our pipeline to those derived in (Ifergan et al., 2024) and found strong structural convergence, indicating that the LLM-based approach is robust despite its limited interpretability.

Table 2 visualizes the top three topics per chronological segment. Recurring topics are manually color-coded across segments using a consistent

palette, enabling readers to trace the persistence or transformation of themes over time. Coverage values for each topic are displayed alongside the labels.

Across both corpora, we observe a clear chronological trajectory: testimonies move from prewar life, to the onset of persecution, to deportation and camp experiences, and ultimately to liberation and postwar recovery. One salient structural difference, however, is the USC archive’s recurring concluding segment dedicated to discussing personal photographs, labeled by our method as “Family Memories”. This photo segment appears consistently in USC interviews but is largely absent from the Yale corpus, marking a distinct divergence in institutional interviewing protocols.

Despite the overall similarity of the macro-historical arc, the two corpora diverge in topical behavior and narrative dynamics. USC testimonies display a more predictable and segmented topical progression, with clearer boundaries between phases, an effect aligned with the archive’s structured interview guidelines. Yale testimonies, in contrast, exhibit greater thematic fluidity: topics begin earlier, overlap more frequently, and persist across multiple segments, producing a narrative rhythm that is less tightly scaffolded by interviewer intervention.

Patterns of topical diversity likewise differ. In Yale interviews, the dominant topics in early segments tend to be introspective and affective, with broader diversification emerging only by the third topic. USC interviews, by contrast, show greater topic variation as early as the second and third topics, yielding a more multidimensional narrative distribution. These differences reflect not only interviewing style but also the interactional norms cultivated by each institution.

Taken together, the results complicate the binary distinction between “structured” and “free-form” interviews. Both archives follow a broadly linear historical logic, yet they diverge in how topics arise, persist, and shift across the testimonial timeline. The Yale corpus tends toward emotional continuity and reflective narration, while the USC corpus exhibits a more segmented and guided progression. The combination of LLM-based topic extraction, cross-method validation (LDA, BERTopic), and explicit visualization reveals that “structuredness” in Holocaust testimony is not merely a function of protocol but a dynamic interplay between institutional design, interviewer choices, and survivor narrative agency.

4.2. Question-Answer Dynamics

To measure question and answer dynamics, we performed a simple word count using the NLTK (Bird and Loper, 2004) library. We then plotted changes

USC-SF			
Seg.	Topic 1	Topic 2	Topic 3
1	Childhood Memories (0.82)	Family Heritage (0.87)	Jewish Identity (0.67)
2	Childhood Memories (0.88)	Experiences of Antisemitism (0.48)	Family Dynamics (0.76)
3	Childhood Memories (0.86)	Experiences of Antisemitism (0.61)	Life in the Ghetto (0.36)
4	Life in the Ghetto (0.45)	Family Separation (0.45)	Survival Strategies (0.69)
5	Life in the Ghetto (0.48)	Family Separation (0.52)	Survival Strategies (0.74)
6	Life in the Ghetto (0.51)	Survival Strategies (0.76)	Family Separation (0.53)
7	Family Separation (0.51)	Survival in Hiding (0.56)	Life in the Ghetto (0.52)
8	Survival in Hiding (0.56)	Family Separation (0.51)	Experiences in Auschwitz (0.33)
9	Survival in Hiding (0.58)	Life in Concentration Camps (0.52)	Family Separation (0.47)
10	Survival in Hiding (0.56)	Family Separation (0.46)	Post-War Resilience (0.56)
11	Survival in Hiding (0.52)	Family Loss and Reunion (0.40)	Post-War Identity Struggles (0.54)
12	Survival and Resilience (0.88)	Family Loss and Reunion (0.46)	Post-War Identity Struggles (0.64)
13	Family Loss in the Holocaust (0.44)	Survival and Resilience (0.85)	Post-War Identity Struggles (0.70)
14	Family Memories (0.76)	Survival Reflections (0.79)	Holocaust Legacy (0.69)
15	Family Memories (0.82)	Legacy of Survival (0.78)	Holocaust Remembrance (0.63)

Yale			
Seg.	Topic 1	Topic 2	Topic 3
1	Childhood Memories Before the War (0.76)	Jewish Identity and Anti-Semitism (0.80)	Family Life Before the Holocaust (0.72)
2	Survival in the Ghetto (0.55)	Family Separation (0.50)	Experiences of Anti-Semitism (0.76)
3	Survival in the Ghetto (0.57)	Family Separation (0.53)	Hiding from Persecution (0.54)
4	Survival in the Ghetto (0.59)	Family Separation Trauma (0.60)	Experiences in Auschwitz (0.28)
5	Survival in the Ghetto (0.64)	Family Separation (0.49)	Experiences in Auschwitz (0.32)
6	Survival Strategies in Hiding (0.58)	Family Separation and Loss (0.61)	Experiences in Concentration Camps (0.60)
7	Survival in Hiding (0.54)	Family Separation (0.48)	Life in the Ghetto (0.50)
8	Survival in Hiding (0.53)	Family Separation (0.52)	Life in Concentration Camps (0.58)
9	Survival in Concentration Camps (0.60)	Life in the Ghetto (0.47)	Family Separation and Loss (0.55)
10	Survival in Hiding (0.56)	Family Loss and Trauma (0.68)	Post-War Resilience (0.57)
11	Survival in Hiding (0.56)	Family Loss in the Holocaust (0.42)	Experiences in Concentration Camps (0.67)
12	Survival Strategies (0.77)	Family Loss (0.43)	Post-War Resilience (0.64)
13	Survival in Concentration Camps (0.53)	Family Loss and Trauma (0.71)	Post-War Resilience (0.65)
14	Survival and Trauma (0.89)	Family Loss and Memory (0.68)	Identity After Liberation (0.73)
15	Survival and Memory (0.74)	Family Loss and Resilience (0.55)	Holocaust Education and Remembrance (0.53)

Table 2: Top three topics per chronological segment. Recurring topics are manually color-coded across segments using a consistent palette showing the transformation of themes over time. Coverage values for each topic are displayed alongside the labels.

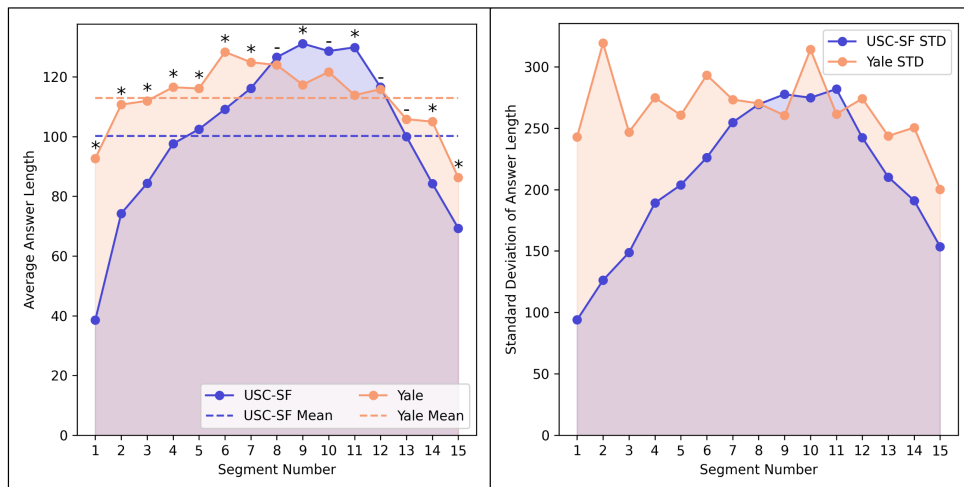


Figure 2: Mean (left) and SD (right) of answer length over time for USC and Yale archives. Statistically significant (*t*-test) differences are marked with an asterisk (*).

in average word count across testimony segments for each corpus, reporting both the mean and variance. A two-sample *t*-test was conducted for each segment to assess whether differences between the corpora were statistically significant (indicated by an asterisk). The overall mean length was plotted as a dashed reference line. See Figures 2, 3.

Answer Length: Yale survivors tended to provide longer answers in the early segments of their

testimonies, whereas responses in the USC corpus gradually lengthened over time. Approximately between segments 1–8, Yale testimonies exhibited longer responses on average. Around segments 8–11, this pattern reversed, with USC answers becoming longer. Toward the end of the testimonies, both collections showed a relative decline in answer length, although Yale participants generally maintained slightly longer responses.

Examining the standard deviation (SD) of an-

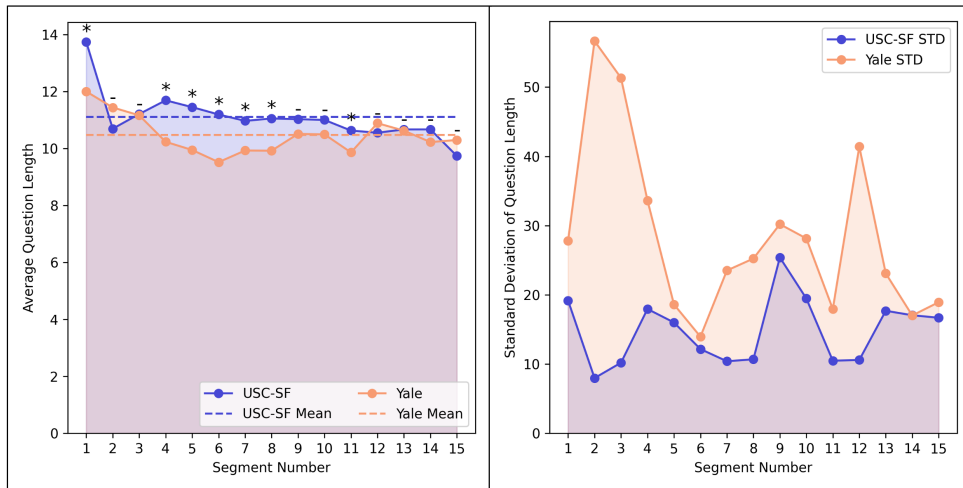


Figure 3: Mean (left) and SD (right) of question length over time for USC and Yale archives. Statistically significant (*t*-test) differences are marked with an asterisk (*).

swer length reveals the degree of variability across testimonies. Higher SD values indicate greater fluctuation in response length, which may reflect differences in survivors' narrative styles, interviewer intervention, or emotional pacing. In this context, Yale testimonies exhibit higher variability, suggesting that survivors alternated between extended reflections and shorter responses, while USC testimonies show more consistent pacing and structure.

Question Length and Frequency: USC interviewers tended to ask longer and more frequent questions, particularly during the early portions of the testimonies. This pattern aligns with the more guided and segmented structure of USC interviews. In contrast, Yale interviewers posed shorter and less predictable questions. The high variability in question length within the Yale corpus (as reflected in SD values) is noteworthy, given the relatively small pool of interviewers, suggesting considerable stylistic diversity and less standardized interviewing practice.

Convergence: Both archives displayed a reduction in answer length toward the final segments, a trend likely associated with standardized closing routines or reflective conclusions. Statistical testing confirmed that differences in answer length between the two corpora were significant in the early interview stages but diminished progressively toward the end.

Taken together, the analysis of question and answer length reveals distinct narrative and methodological dynamics within the two archives. The longer early responses in the Yale corpus suggest a more open and reflective interview style, one that allows survivors to elaborate freely and set the narrative pace. In contrast, the USC interviews begin with shorter, more structured exchanges but expand over time, indicating an interviewer-driven scaffolding that gradually gives way to survivor-led

narration. The higher variability in Yale question and answer lengths reflects more dialogical spontaneity, where interviewers adapt responsively to the survivor's storytelling, whereas the more uniform USC patterns demonstrate adherence to a formalized interview guide. The convergence observed in later segments, marked by shorter responses and reduced differences between the corpora, points to the influence of shared testimonial conventions and closing rituals that shape the end of interviews across institutions. Overall, these patterns highlight that interview structure is not static but evolves through the testimony, balancing institutional design with survivor agency and emotional rhythm.




Figure 4: Intervention Density for USC and Yale archives. Statistically significant (*t*-test) differences are marked with an asterisk (*).

4.3. Intervention Density

To assess the density of interviewer interventions, we measured the duration of uninterrupted survivor speech. Lacking timestamp data, we used word count as a proxy for duration by calculating the proportion of uninterrupted survivor words relative to the total testimony length.

To examine how this measure changes over time, we divided each interview into 15 equal segments based on the cumulative word count. This segmentation, therefore, represents proportional divisions of the testimony length rather than the micro-level segmentation used in the Q/A-based analysis. For each corpus, we computed the average intervention density per segment across all testimonies. This word-based segmentation allows us to capture the relative distribution of interviewer interventions throughout the testimonies, independent of total interview length, and provides a comparable temporal framework across the two archives. The results, presented in Fig. 4, display the average uninterrupted speech density for each segment.

The analysis of intervention density provides further insight into the interactional dynamics that differentiate the two archives. Longer uninterrupted segments of survivor speech indicate greater narrative autonomy, whereas higher intervention density suggests tighter interviewer control or increased dialogical guidance. In this context, the Yale testimonies tend to feature longer uninterrupted stretches, reflecting a more survivor-centered approach that privileges emotional flow and narrative continuity. The USC interviews, by contrast, exhibit a denser pattern of interventions, especially in the earlier portions of the testimonies, which is consistent with a structured interview format emphasizing chronological order. However, in both collections, there is a gradual reduction in intervention density as testimonies progress, suggesting perhaps increased confidence on the survivors' part or greater flexibility on the interviewers' part. This suggests that narrative authority is negotiated throughout the testimonial encounter, resulting in an evolving dynamic between institutional control and individual agency in the making of oral history.

4.4. Question Types

Building on a framework proposed in ((Mittelstadt et al., 2016), Fig. 3.6), we hypothesize that interviewing style affects the distribution of question types—specifically “why,” “what,” “where,” “who,” “when,” “how,” and “other”. Structured interviews are expected to rely more heavily on factual and directive questions (e.g., “what,” “when,” “who”), whereas less structured or conversational interviews are likely to include a broader range of open-ended or ambiguous questions (classified as

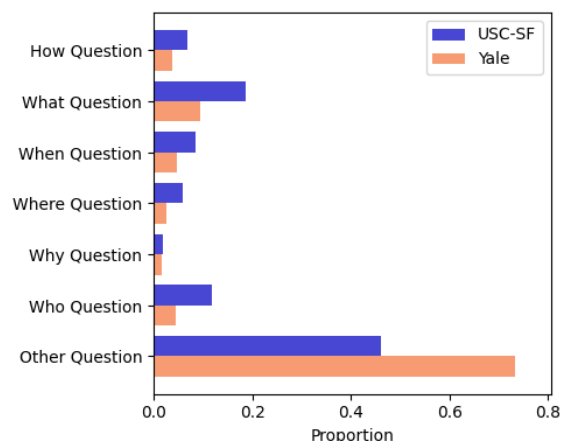


Figure 5: Distribution of question types for USC and Yale archives.

“other”).

To test this hypothesis, we extracted all interviewer questions from both testimony corpora and employed an LLM to classify each into one of the seven categories (see Prompt 7.3 in the Appendix). Model predictions were validated through manual review of a random subset of 50 examples per category (3,500 questions in total). The overall and segment-based distributions of question types are presented in Figures 5 and 6.

The LLM-based classification revealed clear stylistic differences between the two archives. Yale interviewers used a higher proportion of “Other” questions, exhibiting a more flexible and exploratory interviewing style that allows survivors to guide the conversation. In contrast, USC interviewers relied predominantly on “what,” “who,” and “when” questions, particularly in the early segments, consistent with a more structured and documentary-oriented approach. That said, USC interviews show a gradual increase in “Other” question types over time, which might indicate that as rapport developed, the exchanges became more dialogic and open-ended.

Overall, the distribution of question types offers insight into institutional interviewing practices and the dynamics of testimonial co-construction. The USC format prioritizes a more formal interview structure, limiting opportunities for narrative divergence or emotional elaboration. The Yale format fosters greater conversational variability and affective depth, which allows far greater survivor agency in shaping the narrative, but for this reason is more dependent on survivors' compliance and initiative. These findings illustrate how question type mediates the balance between institutional protocol and personal voice, revealing how methodological design and interpersonal rapport together shape the structure and tone of Holocaust testimony.




Figure 6: Distribution of question types over time for USC and Yale archives.

5. Toward a Scalable Framework for Oral History Comparison

Our analysis confirms that USC interviews are, overall, more structured than those in the Yale archive, particularly in question types, topical flow, and interviewer interventions. The USC corpus displays clearer segmental boundaries and a more predictable topical progression, whereas Yale testimonies reveal greater thematic fluidity, variability in questioning style, and more open-ended conversational turns. As the interview unfolds, however, these contrasts diminish as both interview styles move toward similar narrative rhythms shaped by survivor agency, historical chronology, and the affective dynamics of testimony. Ultimately, despite institutional differences, what becomes apparent is that in both collections the dialogical exchange between interviewer and survivor meaningfully shapes the unfolding of testimony alongside the structural template. This finding underscores that “structuredness” in oral history is not an institutional constant but a dynamic property of interaction, emerging from the interplay between methodological design, interpersonal trust, and the survivor’s narrative agency.

Beyond these findings, the study contributes a concrete methodological framework for comparative oral history analysis. The proposed pipeline introduces a sequence of computational strategies designed specifically for dialogic and large-scale testimonial data.

First, it develops a segmentation method that extracts coherent question-answer trajectories from long interviews, enabling alignment across archives that differ in format and length. This segmentation serves as the foundation for cross-corpus comparison, allowing structural and thematic patterns to be analyzed at a shared level of granularity.

Second, the study advances a novel approach to LLM-based topic extraction. Working with oral history presents unique challenges, including contextual overflow, diffuse topic boundaries, and highly

personal narrative structures. To address these, the pipeline employs a staged prompting strategy inspired by map-reduce logic: local prompts generate micro-level topics for each segment, and a secondary synthesis step aggregates these into common themes across the corpus. This approach mitigates the limitations of unsupervised topic modeling and allows for interpretable, reproducible results grounded in qualitative meaning.

Third, the framework integrates iterative evaluation and refinement. Through experimentation with classification strategies, segmentation thresholds, and prompt design, the study develops an ontology of best practices for LLM-assisted oral history analysis. The resulting workflow balances automation and interpretation, enabling scalable comparison while preserving the discursive and emotional complexity of survivor narratives.

Taken together, these contributions establish a reproducible and extensible model for comparing oral history archives at scale. Rather than treating computational methods as an abstraction from humanistic analysis, the pipeline operationalizes interpretive categories such as “structuredness, topical coherence, and interviewer style into measurable and comparable features. It allows distinct institutional collections to be analyzed within a shared framework, revealing how methodological design, interview dynamics, and survivor agency jointly shape the production of testimony.

More broadly, this study demonstrates how digital humanities can move from isolated case studies toward comparative infrastructures that connect archives and traditions of memory work. The pipeline developed here enables not only new forms of analysis but also new questions about the ethics, scale, and dialogical nature of historical testimony. It shows that computational methods, when used critically, can deepen rather than diminish the interpretive work of the humanities, thus illuminating the ways in which stories are told, recorded, and remembered across time and institutional boundaries.

6. Ethics Statement and Limitations

Throughout the study, the OpenAI-GPT suite of models was used, with manual validation of segment samples conducted to confirm model consistency. Furthermore, all data access was obtained with institutional permission and handled strictly in accordance with ethical standards for trauma data. Future work will involve integrating this pipeline into an open-access annotation platform for cross-disciplinary use.

References

- Aleida Assmann. 2011. *Cultural memory and Western civilization: Functions, media, archives*. Cambridge University Press.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Tobias Blanke, Michael Bryant, and Mark Hedges. 2019. [Understanding memories of the holocaust—a new approach to neural networks in the digital humanities](#). *Digital Scholarship in the Humanities*, 35(1):17–33.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Ifeanyi Ezeani, Paul Rayson, Ian N. Gregory, Timothy Cole, Emily Steiner, and Zachary Frank. 2024. The geography of ‘fear’, ‘sadness’, ‘anger’ and ‘joy’: Exploring the emotional landscapes in the holocaust survivors’ testimonies. In *Proceedings of Text2Story@ECIR*, pages 93–103.
- Shoshana Felman and Dori Laub. 1992a. *Testimony: Crises of witnessing in literature, psychoanalysis, and history*. Taylor & Francis.
- Shoshana Felman and Dori Laub. 1992b. *Testimony: Crises of Witnessing in Literature, Psychoanalysis and History*. Routledge, London / New York.
- Saul Friedländer. 1993. *Memory, History, and the Extermination of the Jews of Europe*. Indiana University Press, Bloomington.
- Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.
- Geoffrey H Hartman. 1995. Learning from survivors: The yale testimony project. *Holocaust and Genocide Studies*, 9(2):192–207.
- Geoffrey H Hartman. 1996. *The longest shadow: In the aftermath of the Holocaust*. Indiana University Press.
- Sharlene Nagy Hesse-Biber and Patricia Leavy. 2005. *The Practice of Qualitative Research*. SAGE Publications, Thousand Oaks, CA. Also published in 2006 edition.
- Maxim Ifergan, Omri Abend, Renana Keydar, and Amit Pinchevski. 2024. Identifying narrative patterns and outliers in holocaust testimonies using topic modeling. In *Proceedings of the First Workshop on Holocaust Testimonies as Language Resources (HTRes)@ LREC-COLING 2024*, pages 44–52.
- Renana Keydar. 2020. Listening from afar: An algorithmic analysis of testimonies from the international criminal courts. *University of Illinois Journal of Law, Technology & Policy*, 2020(1):55–83.
- Renana Keydar. 2022. [Changing the lens on survivor testimony: Topic modeling the eichmann trial](#). *Jewish Studies Quarterly*, 29(4):412–435.
- Renana Keydar, Amit Pinchevski, Maxim Ifergan, and Omri Abend. 2026. The testimony of the multitude: Toward a computational model of listening to holocaust testimony. *Holocaust and Genocide Studies*, 40(2).
- Lawrence L Langer. 1993. *Holocaust testimonies: The ruins of memory*. Yale University Press.
- Brent Daniel Mittelstadt, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2):2053951716679679.
- Stephen Naron and Gabor Mihaly Toth. 2020. [Let Them Speak: An Effort to Reconnect Communities of Survivors in a Digital Archive](#), pages 71–94. Springer International Publishing, Cham.
- 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.
- Hannah Pollin-Galay. 2018. *Ecologies of Witnessing: Language, Place, and Holocaust Testimony*. Yale University Press, New Haven, CT.

- Todd Presner, Anna Bonazzi, Rachel Deblinger, Lizhou Fan, Michelle Lee, Kyle Rosen, and Campbell Yamane. 2024. *Ethics of the Algorithm: Digital Humanities and Holocaust Memory*. Princeton University Press, Princeton, NJ.
- Priyanka Ranade, Sanorita Dey, Anupam Joshi, and Tim Finin. 2022. Computational understanding of narratives: A survey. *IEEE Access*, 10:101575–101594.
- Jeffrey Shandler. 2017. *Holocaust Memory in the Digital Age: Survivors' Stories and New Media Practices*. Stanford Studies in Jewish History and Culture. Stanford University Press, Stanford, CA.
- Jeffrey Shandler. 2022. Digitizing holocaust memories. *Jewish studies in the digital age*, 5:25.
- Noah Shenker. 2015. *Reframing Holocaust Testimony*. The Modern Jewish Experience. Indiana University Press, Bloomington, IN.
- Esther Shizgal, Eitan Wagner, Renana Keydar, and Omri Abend. 2025. Computational analysis of character development in holocaust testimonies. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 22721–22745.
- Stephen D Smith. 2022a. *The trajectory of Holocaust memory: The crisis of testimony in theory and practice*. Routledge.
- Stephen D. Smith. 2022b. *The Trajectory of Holocaust Memory: The Crisis of Testimony in Theory and Practice*, 1 edition. Routledge, London.
- Itamar Trainin and Omri Abend. 2025. T5score: A methodology for automatically assessing the quality of llm generated multi-document topic sets. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 26347–26375.
- USC Shoah Foundation. 2021. Interviewer guidelines. https://sfi.usc.edu/sites/default/files/docfiles/uscsf_interviewer_guidelines_03_2021.pdf. PDF document, accessed December 14, 2025.
- Eitan Wagner, Renana Keydar, and Omri Abend. 2025. Unsupervised location mapping for narrative corpora. *arXiv preprint arXiv:2504.05954*.
- Eitan Wagner, Renana Keydar, Amit Pinchevski, and Omri Abend. 2022. Topical segmentation of spoken narratives: A test case on holocaust survivor testimonies. *arXiv preprint arXiv:2210.13783*.
- Annette Wieviorka. 2006. *The Era of the Witness*. Cornell University Press, Ithaca, NY.

Appendix

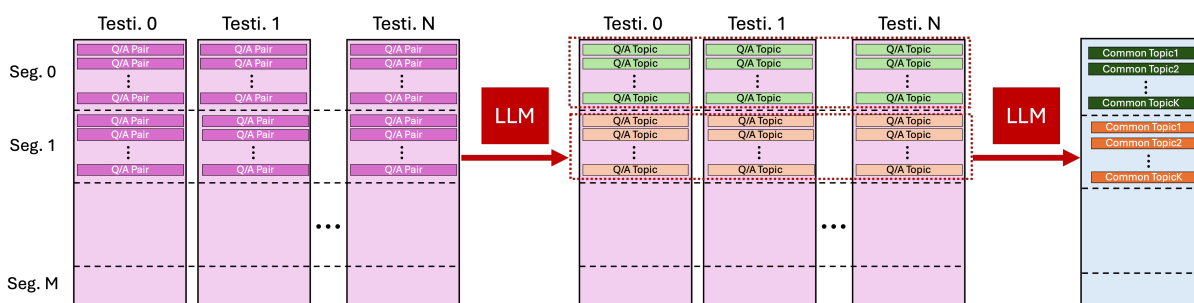


Figure 7: The topical sequence extraction pipeline.

7. LLM Prompts

7.1. Q/A Topic Naming

```
"""
You are a Holocaust researcher.
You will be presented with a short text snippet from a
conversation between an interviewer and a Holocaust survivor.
'INT' represents the interviewer and 'SUBJECT' or
'<survivor_name>' represents the survivor.
Given the text snippet please generate a short title describing
the most prominent topic in the text.
Make sure that the title is short and limited to a few words.
Make sure that the title is comprehensive, specific,
interpretable, and short.
Make sure that the title captures only a single topic.
Output format:
Title: "<title>"
Reason: "<reason>"
Text snippet:
"{text_snippet}"
"""
```

7.2. Common Topics Extraction

```
"""
You are a Holocaust researcher.
You will be presented with a set of titles representing topics
extracted from Holocaust survivor interviews.
```

```
Title Set:
{title_set}
```

Your task is:

- Generate {num_topics} distinct titles that best describe the most common and prominent titles in set.
- Titles must be concise (maximum of a few words), specific, interpretable, and distinct.
- Do NOT combine multiple topics into a single title or use conjunctions like "and".

```
Desired output format:
{output_format}
```

The common titles are:

1.
"""

7.3. Question Type Classification

"""

You are a Holocaust researcher analyzing survivor testimonies. The testimonies are transcripts of an oral interviews. The interviews is composed of speaker sides -- the interviewer questions and the survivor answers. A prefix of "INT" identifies an "interviewer" line. A prefix of "<initials>" identifies a suviror line. You will be given one question asked by an interviewer during a conversation with a Holocaust survivor.

Your task is to classify the question into exactly **one** of the following types, based on its structure and intent:

Question Types:

1. How Question - Asks about a method, process, or manner
2. What Question - Asks for information, descriptions, or clarifications
3. When Question - Asks about time or timing
4. Where Question - Asks about a place or location
5. Why Question - Asks about cause, motivation, or reason
6. Who Question - Asks about a person or group
7. Other Question - Does not match any of the categories above or is ambiguous.

Then, write a brief explanation (no more than one sentence) explaining why you selected that type.

Do not return more than one type.

Output JSON format:

```
{  
  "type": "<one of the 7 types above>",  
  "explanation": "<one-sentence explanation>"  
}
```

Input:

Interviewer Question: "{speaker_line}"
"""