Enhancing Access to Movie Cultural Heritage through Narrative-Enriched Item Representations

Fabio Ferrero[1,†], Angelo Geninatti Cossatin[1,†], Liliana Ardissono[1,†] and Noemi Mauro[1,†]

[1]*University of Turin, Corso Svizzera 185, Torino, 10149, Italy*

Abstract

Movies are an important form of cultural heritage, encoding social values, narrative traditions, and collective memory over time. However, digital platforms that mediate access to movie collections - including recommender systems, digital archives, and streaming services - typically rely on explicit metadata such as genres and plot summaries, which only partially reflect the semantic richness of cinematic content. In this paper, we propose an approach to augment item representations with implicit narrative information extracted from textual descriptions to provide recommender systems with richer item profiles. Building on an LLM-based method for extracting hidden narrative aspects, their interpretations, and their manifestations from item descriptions, we show how these enriched representations provide more expressive and semantically grounded descriptions of movies. We report the extraction performed on the MovieLens 1M dataset. We further discuss how these narrative-enriched representations extend beyond recommendation performance to support deeper cultural exploration, thematic discovery, and interpretation of movie collections, addressing key challenges in personalization for cultural and natural heritage environments.

Keywords

Cultural Heritage, Generative AI, Personalization

1. Introduction

Movies are a pervasive and influential form of cultural heritage. As cinematic artifacts, films capture the evolution of social norms, artistic practices, emotional vocabularies, and collective imaginaries over decades and cultures. Movies are also strictly related to landscape, narrating its history. For example, Lukinbeal explains that "landscape gives meaning to cinematic events and positions narratives within a particular scale and historical context" [1]. Thus, the preservation and valorization of movie heritage is a concern not only for archivists and historians, but also for anyone interested in how societies construct and transmit meaning over time.

Digital platforms have dramatically expanded access to movie collections. Streaming services, digital archives, and recommender systems [2] now mediate how users discover and engage with films. However, the representations that underpin these systems remain shallow. The dominant paradigm relies on structured metadata, such as genre labels, release years, director names, and cast. Film synopses are available, but can only be searched by keywords or using a limited set of static topics that represent generic themes, as in the British Council Film Archive [3].

This mismatch between the richness of cinematic content and the poverty of its digital representations in metadata (or of the way they are used) has several important consequences. In recommendation systems, this means that user preferences cannot be fully modeled from available item features, leading to suboptimal recommendations. When recommender systems suggest films based on shallow metadata similarity or collaborative filtering signals, as in [4], they tend to surface popular or superficially similar items while failing to draw meaningful connections rooted in narrative or thematic affinity. However,

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden

†These authors contributed equally.

fab.ferrero@unito.it (F. Ferrero); angelo.geninatticossatin@unito.it (A. Geninatti Cossatin); liliana.ardissono@unito.it (L. Ardissono); noemi.mauro@unito.it (N. Mauro)

0009-0007-2941-5030 (F. Ferrero); 0009-0007-5378-7061 (A. Geninatti Cossatin); 0000-0002-1339-4243 (L. Ardissono); 0000-0001-8234-3266 (N. Mauro)

users might appreciate a film not because it shares a genre with something they have seen, but because it explores a similar psychological dynamic, moral conflict, or emotional arc.

In this paper, we propose a framework to extract hidden content [5] within the domain of cultural heritage, with a focus on film production. Rather than treating recommendation performance as the sole ultimate objective, we argue that narrative-enriched item representations serve a broader cultural function: they make the implicit structures of cinematic works legible, enabling richer forms of personalization, discovery, and interpretation. This type of analysis is broadly relevant for cultural heritage, where books, tales, and other types of artistic content are characterized by rich textual descriptions.

Recent advances in Large Language Models (LLMs) offer a promising path forward to build rich movie profiles. LLMs have a remarkable capacity for semantic understanding and commonsense reasoning [6], enabling them to unveil unstated information from textual descriptions [7]. In this work, we explore the adoption of LLMs to extract hidden narrative aspects from movie metadata. Our goal is to augment metadata with an explicit representation of the narrative and emotional characteristics emerging from movie plots. This is a first step towards the development of advanced recommender systems that can exploit both explicit and implicit content to personalize item suggestions.

Specifically, we propose a prompt-based approach to extract narrative hidden content from textual item descriptions using an LLM. The goal is to be agnostic to the selected LLM, an important characteristic given the rapid pace at which Generative Artificial Intelligence is evolving.

We report experimental results on the MovieLens 1M dataset [8], enriched with movie plots from the MPST dataset [9], and discuss how the resulting representations can inform the design of more expressive and culturally aware access systems for film heritage.

The remainder of this paper is organized as follows: Section 2 motivates the extraction of hidden content from movie plots. Section 3 describes the types of hidden content we analyze. Section 4 describes our dataset and the extraction methodology. It also presents a sample extraction. Section 5 discusses some implications of our work. Section 6 discusses the challenges and future work and Section 7 concludes the paper.

2. Movies as Data for Cultural Heritage

Movies function as complex cultural artifacts that encode both explicit and implicit knowledge. At the explicit level, a film can be described through filmographic metadata, such as director, cast, genre, release year, etc., and its narrative content, typically summarized through a plot synopsis. These explicit elements are well-suited to formal representation and have long served as the basis for digital cataloging and retrieval; e.g., see [10, 11, 12, 3, 13].

Movie catalogs enable users to explore their databases using keywords and filters on metadata. However, a film's cultural significance is rarely exhausted by its explicit attributes. Films communicate through narrative structure, visual and auditory aesthetics, character dynamics, symbolic registers, and thematic tensions that operate beneath the surface of the plot. A film about a heist is also, perhaps, a meditation on loyalty and betrayal; a war film may simultaneously constitute an inquiry into the nature of heroism; a romantic comedy may encode specific assumptions about gender and social mobility. These implicit dimensions are precisely what give films their lasting cultural resonance, and should be integrated into information search functions to support more expressive search queries.

Two films classified under the same genre or dealing with the same topic might differ profoundly in their narrative structure, emotional register, moral framework, and thematic complexity. Movie plots convey rich information for revealing such aspects. However, extracting their narrative aspects requires a deep analysis of content and discourse. Thus, these differences remain invisible to systems that employ surface approaches to information retrieval.

The challenge in this work is to develop richer representations that capture something of the implicit semantic content of films. This requires moving beyond keyword extraction and genre classification toward methods capable of identifying and articulating the narrative patterns, emotional dynamics, and thematic structures that characterize cinematic works.

In some studies, film metadata has been enriched with eudaimonic and hedonic scores [14, 15] by analyzing plots, reviews, and audio-video features. However, the narrative aspects of plots have not been explored to augment item-searchable features with a structured representation of this type of information.

3. Extraction of Hidden Narrative Content from Movie Plots

We aim to extend metadata with a structured representation of the main narrative patterns that can be found in movie plots, and of the concrete elements that instantiate such patterns. For this reason, we focus on three types of hidden content:

1. *Hidden Aspect Names*, namely short labels that identify implicit narrative, structural, emotional, or experiential features emerging from textual descriptions, including story mechanics. They represent the main abstract patterns underlying a film, such as narrative configurations, character relationships, emotional developments, or stylistic elements that influence how the item is perceived.

2. For each hidden aspect, we look for a *Hidden Aspect Interpretation* that defines it by contextualizing the pattern in the movie plot.

3. For each hidden aspect, we also look for a *Hidden Aspect Manifestation* that specifies *how* the narrative pattern is instantiated in the movie plot, through events and observations.

4. Dataset Construction

4.1. Source Data

For our work, we used the MovieLens 1M dataset [8], which contains approximately one million ratings from 6,040 users on 3,706 movies, together with genre metadata.

To obtain textual plot descriptions, we extended MovieLens 1M with synopses from MPST [9]. We mapped films across the two datasets via their IMDb identifiers and corresponding entries from MovieLens 25M[1]. The resulting augmented dataset (which we denote as `ML-Plot`) covers 2,209 films with plot synopses, with an average word count of approximately 1,107 words per synopsis.

4.2. Extraction Model

For each movie i in a collection I, let D_i denote its textual descriptions. In principle, an item might have more than one textual description; in our domain, we consider a single one, i.e., the plot synopsis. Moreover, let M_i denote i's associated metadata, e.g., genre labels, considering that a single movie can be classified into multiple genres.

The extraction method we applied prompts a Large Language Model to identify, for each film $i \in I$, a set of triples structured as follows:

$$A_i = \{(name_{i1}, interpr_{i1}, manif_{i1}), \ldots, (name_{ik}, interpr_{ik}, manif_{ik})\} \text{ with } k > 0$$

where $name_{ij}$, $interpr_{ij}$, and $manif_{ij}$ denote, respectively, a Hidden Aspect Name, a Hidden Aspect Interpretation, and a Hidden Aspect Manifestation.

[1]https://grouplens.org/datasets/movielens/25m/

4.3. Prompt Design

We applied structured prompting with few-shot examples to illustrate the expected output format and improve consistency across items. The prompt instructs the LLM to identify implicit narrative patterns, emotional dynamics, thematic structures, and experiential characteristics not explicitly stated in the description of the analyzed item. It is structured to encourage the model to focus on implicit meaning rather than explicit content, identify recurring narrative patterns, provide context-aware descriptions of such patterns, and ground interpretations in specific narrative elements. The prompt is available at the following link, in the online repository: https://bit.ly/3OggyWy.

4.4. Hidden Aspects Extraction

We extracted hidden aspects from all films of the `ML-Plot` dataset by prompting the Qwen3-30B Large Language Model[2]. As a result, on average, each film has 3.14 hidden aspects. Hidden Aspect Names average approximately 3.5 words in length, Interpretations approximately 11 words, and Manifestations approximately 22 words.

Three human annotators independently validated the quality of extracted content on a sample of 100 films, confirming the reliability of the extraction process.

Section 4.5 presents the hidden content extracted for a sample film. At the following link, we report the results of some of the extractions we conducted, with plots and relative hidden content: https://bit.ly/4t1KPI3.

4.5. Case Study: Silence of the Lambs (1991)

Let's consider the famous movie *Silence of the Lambs* (1991), where an FBI trainee consults an imprisoned cannibal (Hannibal Lecter) to track a serial killer who kidnapped a woman named Catherine Martin.[3] For this movie, our extraction procedure surfaced the following hidden content:

- Hidden Aspect Name: *"Trauma-Driven Motivation"*. Hidden Aspect Interpretation: *"Using past psychological wounds to fuel present actions and decisions"*. Hidden Aspect Manifestation: *"The protagonist's nightmares about lambs and her determination to save Catherine Martin to end her own trauma"*.

- Hidden Aspect Name: *"Psychological Predator-Prey"*. Hidden Aspect Interpretation: *"Mutual psychological manipulation exploiting vulnerabilities in a dangerous power dynamic"*. Hidden Aspect Manifestation: *"Lecter's quid pro quo interviews where he gains personal information from the protagonist while giving clues, creating a dangerous power balance"*.

- Hidden Aspect Name: *"Violated Innocence"*. Hidden Aspect Interpretation: *"The horror of physical and psychological violation of the vulnerable"*. Hidden Aspect Manifestation: *"The victims' skinning and objectification, symbolizing the terror of being stripped of identity and autonomy"*.

Note that these representations do not replace the explicit content that is typically encoded by a pre-trained language models (e.g., BERT [16]) from a film's plot. In contrast, they complement that type of information by making its narrative architecture explicit. Specifically, the Hidden Aspect Names provide abstract thematic labels that enable comparison across films at a high level of generality, e.g., matching movies that deal with violated innocence. The Interpretations are slightly more detailed: they supply conceptual content that contextualizes each pattern within the specific film. The Manifestations anchor this conceptual content in concrete narrative events, enabling more precise comparisons among items.

[2] https://ollama.com/library/qwen3:30b-a3b

[3] See https://bit.ly/4cxkKKU for the plot of this movie.

5. Implications for Cultural Heritage Personalization

5.1. Expressive Indexing and Thematic Discovery

Standard movie catalogs are organized primarily along genre lines, supplemented by editorial tags and keyword annotations. While genres provide a useful coarse-grained classification, they are a blunt instrument for cultural exploration. Genre labels cannot distinguish between films that share surface conventions but differ profoundly in their narrative philosophy, emotional register, or thematic depth. A user who has responded strongly to a film's exploration of moral ambiguity will receive poor guidance from a genre-based system that recommends other crime films indiscriminately.

Hidden Aspect Names offer a complementary indexing vocabulary rooted in narrative patterns to extend beyond genre conventions. By organizing films according to recurring thematic structures-such as redemption narratives, such as trauma-driven motivation, psychological predator-prey, and violated innocence, these labels enable a form of thematic discovery that cuts across genre boundaries and surfaces unexpected connections between films from different periods, traditions, and cultures. This is particularly valuable in the context of cultural heritage, where meaningful connections often exist precisely across the conventional boundaries through which collections are organized. The RevIS project explored ontological representations of movie data to model landscape and cultural habits over time [13]. Our work has a similar goal, but grounds on narrative aspects of film plots.

Hidden Aspect Interpretations enrich this thematic vocabulary by providing conceptual descriptions that articulate the emotional and narrative significance of each narrative pattern within the specific context of a film. Rather than simply labeling a theme, they explain what the theme means and why it matters in a particular work. This level of semantic granularity has a potential to support nuanced forms of browsing and recommendation, enabling users to engage with films not as instances of a genre type but as individual cultural artifacts with distinctive narrative identities.

Finally, Hidden Aspect Manifestations ground this conceptual content in the specific events and situations of each film's narrative. By identifying the scenes, character decisions, and plot developments through which a theme is realized, they make it possible to compare movies not only at the level of abstract narrative patterns but at the level of concrete narrative strategies. For example, two films might share the theme of violated innocence while realizing it through very different narrative mechanisms. Manifestations capture this difference in a way that neither genre labels nor plot summaries can.

5.2. Personalization Beyond Popularity

One of the central challenges of cultural heritage recommendation is the dominance of popular works. Recommender systems trained on interaction data tend to reproduce existing consumption patterns, systematically disadvantaging lesser-known films that may be culturally significant but have accumulated fewer ratings. This reinforces a feedback loop in which popular works become more visible while the long tail of cultural heritage remains inaccessible.

The extraction of hidden content, with its capability to enrich item profiles, is aimed at improving recommendation performance by analyzing item descriptions that are available as soon as movies are published. Various works have extracted additional information by exploiting external data sources, such as consumer feedback [7, 17, 18]. However, this is limited to the extent that individual items receive their reviews, and requires a certain level of accumulation to provide useful data for recommendations.

5.3. Support for Narrative Search and Cultural Analysis

Narrative-enriched representations open up new possibilities for search and analysis in cultural heritage contexts. Current search interfaces for movie collections are predominantly metadata-based, supporting queries by title, director, year, genre, or keyword. The hidden aspect framework suggests an alternative: narrative search, in which users can specify the kind of story they are looking for in terms of the patterns, dynamics, and themes they wish to encounter. Queries such as "films exploring the use of inherited objects to anchor personal identity" or "narratives in which absurd situations serve as moral commentary" become tractable when films are represented with the semantic granularity provided by Hidden Aspect Names, Interpretations, and Manifestations.

For cultural researchers and digital humanities scholars, these representations also support new forms of large-scale analysis. The distribution and evolution of narrative patterns across a film collection (which themes are prevalent in which periods, which patterns co-occur, how certain narrative structures migrate across genres and national traditions) becomes an empirical question that can be addressed computationally when films are represented at this level of semantic detail. This opens a rich agenda for the computational study of cultural heritage, complementing the qualitative methods of film studies and cultural history.

The combination of personalization and cultural relevance is particularly challenging in heritage contexts, where the goal is not only to satisfy individual preferences but to facilitate meaningful engagement with the cultural significance of the collection. Narrative-enriched representations support this double objective by providing a semantic vocabulary that is both expressive enough to capture individual thematic interests and grounded enough in the cultural content of films to support substantive engagement with their heritage value.

6. Discussion

The approach presented in this paper raises several important challenges that warrant further investigation.

First, the quality of extracted hidden aspects depends on the reliability of the underlying language model. While LLMs exhibit impressive semantic understanding, their output may be inconsistent across runs or reflect biases present in their training data. The cultural sensitivity of narrative analysis, e.g, the recognition that the same film may be interpreted differently across cultural contexts, is not guaranteed by prompt engineering alone. Future work should investigate the results obtained by exploiting different LLMs to evaluate their consistency, as well as ensemble approaches and bias mitigation strategies for LLM-generated narrative features.

Second, the approach currently relies on the availability of rich textual descriptions. Our extractions demonstrate strong results in the movie domain, where plot synopses provide dense narrative content. However, the applicability of the method to domains with sparser or more functional textual descriptions remains to be established. Cultural heritage collections often include artifacts whose documentation is incomplete or highly technical; adapting the extraction process to these contexts will require additional methodological development.

Finally, the interpretability of hidden aspects as cultural representations deserves attention. The short labels provided by Hidden Aspect Names are necessarily reductive, and the meaning they encode is shaped by the language model's implicit cultural assumptions. A more systematic approach to the cultural validation of extracted representations, drawing on the expertise of film scholars, curators, and heritage professionals, would strengthen the credibility of the approach for cultural heritage applications.

7. Conclusion and Future Work

This paper has argued LLM-based extraction of hidden narrative aspects from film descriptions offers a principled approach to enriching item representations for movie cultural heritage. By surfacing implicit narrative patterns, their contextual interpretations, and their concrete textual manifestations, this approach produces representations that are more semantically expressive than those provided by standard metadata.

Narrative-enriched representations support a broad cultural agenda: enabling thematic discovery across genre boundaries, facilitating personalization that reaches beyond popular works into the long tail of cultural heritage, supporting narrative search in cultural archives, and opening new possibilities for the computational analysis of film collections and for personalized recommendation of movies.

The combination of Large Language Models and cultural heritage research points toward a future in which digital access systems can engage with the semantic depth of cultural artifacts, rather than merely indexing their surface attributes. This paper represents a step in that direction, grounding the approach in empirical results while articulating the broader cultural implications of narrative-enriched item representation.

In our future work, we will experiment with additional LLMs to compare the hidden content they extract. Moreover, we will validate the usefulness of exploiting hidden content in movie recommendation by comparing movie suggestions generated by different recommendation algorithms, e.g., using datasets such as MovieLens, and in a user study.

Acknowledgments

Our University has partially funded this work through the project "Grant for Internationalization".

Declaration on Generative AI

During the preparation of this work, the authors used Claude[4] and Grammarly[5] for the execution of the following tasks: Paraphrase and reword, Grammar and spelling check. After using these tools, the authors reviewed and edited the content as needed. The authors assume complete responsibility for the content of the publication.

References

[1] C. Lukinbeal, Cinematic landscapes, Journal of Cultural Geography 23 (2005) 3–22. doi:10.1080/08873630509478229.

[2] F. Ricci, L. Rokach, B. Shapira, Recommender Systems: Techniques, Applications, and Challenges, Springer US, New York, NY, 2022, pp. 1–35. doi:10.1007/978-1-0716-2197-4_1.

[3] British Council, British Council film archive, 2024. https://film.britishcouncil.org/resources/film-archive.

[4] R. P. Karumur, T. T. Nguyen, J. A. Konstan, Exploring the value of personality in predicting rating behaviors: A study of category preferences on movielens, in: Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '16, Association for Computing Machinery, New York, NY, USA, 2016, p. 139–142. URL: https://doi.org/10.1145/2959100.2959140. doi:10.1145/2959100.2959140.

[5] H. Grice, Logic and conversation, in: P. Cole, J. L. Morgan (Eds.), Syntax and Semantics, volume 3, *Speech Acts*, Academic Press, New York, 1975, pp. 41–58.

[6] J. Lin, X. Dai, Y. Xi, W. Liu, B. Chen, H. Zhang, Y. Liu, C. Wu, X. Li, C. Zhu, H. Guo, Y. Yu, R. Tang, W. Zhang, How can recommender systems benefit from large language models: A survey 43 (2025). URL: https://doi.org/10.1145/3678004. doi:10.1145/3678004.

[7] M. Chaudhary, T. Neal, Implicit aspect extraction: A systematic review, ACM Comput. Surv. 58 (2026). URL: https://doi.org/10.1145/3786590. doi:10.1145/3786590.

[8] F. M. Harper, J. A. Konstan, The movielens datasets: History and context, ACM Trans. Interact. Intell. Syst. 5 (2015). URL: https://doi.org/10.1145/2827872. doi:10.1145/2827872.

[9] S. Kar, S. Maharjan, A. P. López-Monroy, T. Solorio, MPST: A corpus of movie plot synopses with tags, in: N. Calzolari, K. Choukri, C. Cieri, T. Declerck, S. Goggi, K. Hasida, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk, S. Piperidis, T. Tokunaga (Eds.), Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018), European Language Resources Association (ELRA), Miyazaki, Japan, 2018. URL: https://aclanthology.org/L18-1274/.

[4]https://claude.ai/
[5]https://app.grammarly.com/

[10] M. Burch, G. Baulig, T. Boley, A. Mehmeti, D. Kurbanismailova, M. Roswag, O. Streicher, S. Wittig, U. Kloos, IMDb explorer: Visual exploration of a movie database, in: Proceedings of the 11th International Symposium on Visual Information Communication and Interaction, VINCI '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 88–91. URL: https://doi.org/10.1145/3231622.3231633. doi:10.1145/3231622.3231633.

[11] T. Low, C. Hentschel, S. Stober, H. Sack, A. Nürnberger, Exploring large movie collections: Comparing visual berrypicking and traditional browsing, in: MultiMedia Modeling, Springer International Publishing, Cham, 2017, pp. 198–208.

[12] R. Lobato, A. Scarlata, T. Wils, Video-on-demand catalog and interface analysis: The state of research methods, Convergence 30 (2024) 1331–1347. URL: https://doi.org/10.1177/13548565241261992. doi:10.1177/13548565241261992.

[13] S. Alovisio, L. Ardissono, A. Geninatti Cossatin, V. Lombardo, D. Magro, N. Mauro, A. Nasi, L. Scarpinati, From scenes to places: Enhancing film exploration through the revis ontology-based consultation app, J. Comput. Cult. Herit. (2026). URL: https://doi.org/10.1145/3801551. doi:10.1145/3801551, just Accepted.

[14] E. Motamedi, D. Khosh Kholgh, S. Saghari, M. Elahi, F. Barile, M. Tkalcic, Predicting movies' eudaimonic and hedonic scores: A machine learning approach using metadata, audio and visual features, Information Processing & Management 61 (2024) 103610. URL: https://www.sciencedirect.com/science/article/pii/S0306457323003473. doi:https://doi.org/10.1016/j.ipm.2023.103610.

[15] M. Cohen-Kalaf, J. Lanir, P. Bak, O. Mokryn, Movie emotion map: an interactive tool for exploring movies according to their emotional signature, Multimedia Tools and Applications 81 (2022) 14663–14684. doi:10.1007/s11042-021-10803-5.

[16] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: J. Burstein, C. Doran, T. Solorio (Eds.), Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. URL: https://aclanthology.org/N19-1423/. doi:10.18653/v1/N19-1423.

[17] L. Chen, F. Wang, L. Qi, F. Liang, Experiment on sentiment embedded comparison interface, Knowledge-Based Systems 64 (2014) 44–58. URL: https://www.sciencedirect.com/science/article/pii/S0950705114001099. doi:https://doi.org/10.1016/j.knosys.2014.03.020.

[18] L. Chen, G. Chen, F. Wang, Recommender systems based on user reviews: the state of the art, User Modeling and User-Adapted Interaction 25 (2015) 99–154. URL: https://doi.org/10.1007/s11257-015-9155-5. doi:10.1007/s11257-015-9155-5.