Beyond Genre Categories: How Narrative Pattern Coherence and Spanning Distance Shape Film Success

Zhichao Wang
Tohoku University
Faculty of Arts and Letters
wang.zhichao.p2@dc.tohoku.ac.jp

Zeyu Lyu
Tohoku University
Faculty of Arts and Letters
lyu.zeyu.e8@tohoku.ac.jp

Abstract

Prior research on cultural markets has relied on genre labels to distinguish products, overlooking the specific content features that differentiate films within the same genre. We address this gap using tropes as building blocks of narrative structure. From a dataset of 30k tropes across 18k films (TVTropes.org), we identify 29 narrative patterns via community detection and characterize each film by two measures: coherence (how concentrated its tropes are within a few patterns) and spanning distance (how far apart the patterns it combines are). Regression analyses show that coherence improves both audience evaluations and attention, while spanning distance increases evaluations but reduces attention. These findings extend category-spanning theory from genre labels to the internal narrative composition of films, demonstrating how stories are constructed and shape audience responses.

1 Introduction

Artistic creation is shaped by conventions and boundaries (Bourdieu, 1996; DiMaggio, 1987; Lena and Peterson, 2008). In film, as in other cultural industries, producers and audiences share expectations about how stories are told—what kinds of characters appear, how plots unfold, and what narrative devices are appropriate within a given tradition. Previous studies suggest that a work that crosses these boundaries risks confusing audiences who cannot easily place it(Zuckerman, 1999; Hsu, 2006), but may also be rewarded for its originality(Rao et al., 2005; Askin and Mauskapf, 2017). However, these studies have measured conventions and boundaries using genre labels, thereby often overlooking the specific plot and content features of films. For example, *Rush Hour* and *Kill Bill* are both labeled "Action," yet one is a buddy comedy built on slapstick and mismatched partners, while the other is a revenge tragedy driven by violence and betrayal. To capture such differences in narrative structure, we need to look beyond genre labels and into the narrative elements from which films are constructed.

Tropes are the recurring, recognizable narrative elements in cultural products—plot devices, character archetypes, situational setups—that constitute the narrative structure of stories(García-Ortega et al., 2021). A well-known example is *the Chosen One*: an ordinary protagonist marked for an extraordinary destiny, a device found in films as diverse as *The Matrix*, *Kung Fu Panda*, and *Moana*. Tropes do not appear independently; certain tropes tend to co-occur across films, forming what we call *narrative patterns*: groups of tropes that collectively reflect a recurring combination of narrative conventions. *TVTropes*[1], a community-maintained wiki launched in 2004, has documented approximately 33,000 tropes across 127,000 works, providing a comprehensive dataset on how tropes are used across cultural products. This makes it possible to identify the narrative patterns that tropes form and to examine how films' combinations of these patterns relate to market success.

In this paper, we identify narrative patterns from large-scale trope co-occurrence data and characterize each film along two dimensions. *Coherence* captures how concentrated a film's tropes are within a few patterns. *Spanning distance* captures, among films that combine multiple patterns, how far apart those patterns are from one another. We test three hypotheses regarding how these two dimensions relate to audience outcomes. First, we expect that coherence benefits films because tropes concentrated within familiar patterns make the narrative easier for audiences to interpret, leading to more favorable evaluations(Zajonc, 1968). Moreover, coherent films appeal more directly to audiences who already favor particular

[1] https://tvtropes.org/

Proceedings of the 6th International Conference on Natural Language Processing for the Digital Humanities, pages 132–137
July 4, 2026 ©2026 Association for Computational Linguistics

patterns(Keuschnigg, 2015). Therefore: **H1:** Films with higher coherence receive more favorable audience evaluations and attract greater audience attention. Second, we expect that spanning distance affects evaluations and attention differently. Films that combine distant patterns create greater novelty and originality, which may lead to higher evaluations(Uzzi et al., 2013). However, when a film combines patterns that are far apart, it appeals to a narrower overlap of audience tastes, which reduces its potential overall audience(Hannan and Freeman, 1989). Therefore: **H2a:** Films with higher spanning distance are associated with higher audience evaluations. **H2b:** Films with higher spanning distance are associated with lower audience attention.

2 Related Work

The category-spanning literature has identified two mechanisms through which boundary crossing penalizes products: reduced niche fitness and audience confusion (Hsu et al., 2009; Negro and Leung, 2013). In film, Keuschnigg and Wimmer (2017) find that confusion plays a relatively larger role. These penalties are moderated by category contrast and the distance between categories spanned (Kovács and Hannan, 2015). Pontikes (2012) further shows that the same ambiguity penalized by consumers can be rewarded by investors, and distinguishes between variety and atypicality as independent dimensions of boundary spanning. These work suggests that what matters is not simply whether a product crosses boundaries, but which boundaries and how far.

A growing body of research has moved beyond category labels to characterize cultural products by their internal composition. Askin and Mauskapf (2017) use acoustic features to show that optimal differentiation predicts chart success; **?** demonstrate that the popularity of stylistic elements depends on their structural embeddedness among other elements rather than on the categories they belong to; Sgourev et al. (2023) show how visual aesthetic features like color serve as positioning devices relative to peers. These studies share a common insight: product-level features capture dimensions of novelty and conformity that coarse labels miss. Yet in film, the dominant approach continues to rely on genre labels as proxies for narrative structure.

We address this gap using tropes as building blocks of narrative structure. Tropes have been used in a range of computational and cultural studies: tracing how tropes propagate across works (Mellina and Svetlichnaya, 2011), mapping the trope co-occurrence networks (García-Ortega et al., 2021), supporting story generation (Álvarez and Font, 2022; Chou et al., 2023), and identifying social biases embedded in trope usage (Gala et al., 2020). We take a different direction: rather than predicting ratings from individual tropes or generating new narratives, we identify narrative patterns from trope co-occurrence and measure how films' combinations of these patterns—their coherence and spanning distance—relate to audience outcomes, connecting the computational study of tropes to the theoretical framework of category spanning.

3 Data

We collect a large-scale dataset containing approximately 30,924 tropes and 1.2 million trope–film links across 18,164 films from *TVTropes.org*. Then, we matched these films to the IMDb datasets to obtain film metadata (release year, genres) and audience outcomes: *IMDb rating* (audience evaluation) and *number of votes* (audience attention). The dataset covers films released from the 1920s to the 2020s. Table 1 reports descriptive statistics.

Table 1: Descriptive statistics

Variable	**Mean**	**Median [Min; Max]**
Tropes per film	54.66 (64.59)	32 [1; 483]
Films per trope	32.20 (93.33)	10 [1; 3788]
Genres per film	2.37 (1.02)	2 [1; 7]
IMDb rating	6.21 (1.18)	6.40 [1.20; 9.80]
IMDb votes	42967 (114517)	6423 [5; 3152337]

4 Methods

Our research proceeds in three steps. First, we identify *narrative patterns*—groups of tropes that tend to co-occur across films—by detecting trope communities in a co-trope network (Section 4.1). Second, we use each film's distribution of tropes across these patterns to construct two film-level variables: *coherence*, which captures how concentrated a film's tropes are within a few patterns, and *conditional spanning distance*, which captures how narratively distant the bridged patterns are (Section 4.2). Third, we estimate regression models relating these variables to audience outcomes (Section 4.3).

4.1 Narrative Pattern Identification

Certain tropes tend to co-occur in the same films, and these co-occurrence regularities reveal the underlying patterns that organize narrative conventions (Cawelti, 1976). To identify these patterns empirically, we construct a weighted co-trope network in which each node is a trope and each edge weight reflects how often two tropes appear together in the same film. We apply the Leiden algorithm (Traag et al., 2019) to partition this network into communities, yielding 29 groups. Each community is treated as a *narrative pattern*: a data-driven cluster of tropes that collectively reflect a recurring trope combination.

The resulting partition has a distinctive structure. One large community (Community 0, containing 10,573 tropes) appears across nearly all films in the dataset. Its central tropes—such as *Shout Out*, *Foreshadowing*, and *Chekhov's Gun*—are general-purpose narrative devices not specific to any genre or thematic cluster. This community functions as a baseline layer of narrative conventions shared across films. The remaining 28 communities are smaller and more specialized, each organized around a distinctive thematic or narrative domain: for example, one centers on supernatural and horror elements, another on science fiction and futuristic settings, and another on combat and weaponry. Community film coverage among these 28 ranges from 1 to 10,475 films (mean = 4,116; SD = 3,252). Table 3 in the Appendix reports characteristics of five example communities.

4.2 Measuring Coherence and Spanning Distance

Given the 29 narrative patterns identified above, we characterize each film along two dimensions. The first, *coherence*, captures how concentrated a film's tropes are within a few patterns—whether the film commits to a clear narrative identity or disperses its elements across many different conventions. The second, *spanning distance*, captures how narratively distant the patterns are that a film combines—whether it bridges patterns that frequently co-occur (close distance) or patterns that rarely appear together (far distance).

For film m, let n_{mi} denote the number of tropes assigned to community i, and let $N_m = \sum_i n_{mi}$ be the total trope count in film. The share of film m's tropes belonging to community i is $w_{mi} = n_{mi}/N_m$.

4.2.1 Coherence

Some films draw their tropes primarily from one or two narrative patterns; others spread them broadly. We expect this distinction to matter because a film whose tropes cluster within a small number of recognizable patterns presents a clearer narrative backbone, making it easier for audiences to interpret and evaluate. To capture this, we need to measure how concentrated a film's tropes are across patterns. We use the Herfindahl–Hirschman Index (HHI), a standard measure of concentration widely used in economics and organizational research to quantify how a distribution is spread across categories (Rhoades, 1993). In our context, HHI measures the degree to which a film's tropes are concentrated within a few narrative patterns:

$$\mathrm{HHI}_m = \sum_i w_{mi}^2.$$

HHI approaches 1 when a film draws almost all of its tropes from a single pattern, and decreases as tropes spread more evenly.

4.2.2 Spanning Distance

Beyond how concentrated a film's tropes are, we also want to know how far apart the patterns it combines are from one another. Consider two films, each spanning two narrative patterns. One mixes two closely related types of action tropes; the other mixes horror tropes with romantic comedy tropes. Even though both films span the same number of patterns, the second is likely to bridge a larger narrative gap.

To measure how far apart two patterns are, we look at how often they appear together in the same films. Patterns that frequently appear together in films represent familiar combinations that creators and audiences are accustomed to; patterns that rarely co-occur represent more unusual pairings. Formally, let F_i denote the set of films containing at least one trope from community i. The Jaccard distance between trope communities i and j is:

$$d_{ij} = 1 - \frac{|F_i \cap F_j|}{|F_i \cup F_j|}.$$

Given these pairwise distances between patterns, we aggregate them at the film level. Each film's spanning distance is the weighted average distance between the patterns it draws from:

$$\mathrm{SD}_m = \frac{\sum_{i<j} w_{mi}\, w_{mj}\, d_{ij}}{\sum_{i<j} w_{mi}\, w_{mj}}.$$

Spanning distance is high when a film combines narratively distant patterns, and low when the patterns it mixes tend to co-occur in other films as well. 1,056 films draw all of their tropes from a single community, making spanning distance impossible to compute; these films are excluded from the estimation sample.

4.3 Regression Model

We examine how coherence and spanning distance relate to two outcomes: *IMDb rating*, which captures audience evaluation, and *number of IMDb votes*, which captures audience attention. Because vote counts are highly right-skewed, we use $\log(1 + \mathit{Votes})$ as the dependent variable.

We estimate OLS regressions with HC3 robust standard errors:

$$Y_m = \beta_0 + \beta_1 \mathrm{HHI}_m + \beta_2 \mathrm{SD}_m + \beta_3\, \mathit{trope\_count}_m + \gamma_{year} + \delta_{genre} + \varepsilon_m.$$

Both HHI and SD are mean-centered prior to estimation. All models control for the total number of tropes in a film, because films with more tropes tend to span more communities simply because they have more tropes to distribute(García-Ortega et al., 2021). We also include year fixed effects to absorb temporal trends in both film production and TVTropes annotation coverage(García-Ortega et al., 2020), and main-genre fixed effects to ensure that the estimated relationships are not driven by genre differences in ratings or votes(Francemone et al., 2023; Shahid and Islam, 2023).

5 Results

Table 2 reports the main regression estimates. The correlation between HHI and SD in the estimation sample is 0.09, and the variance inflation factors are close to 1, indicating no multicollinearity concern.

HHI is positive and significant in both models ($\beta = 1.089$, $p < 0.01$ for rating; $\beta = 1.207$, $p < 0.01$ for votes). These results support H1, indicating that films whose tropes concentrate within fewer narrative patterns tend to receive higher evaluations and more audiences. Spanning distance shows opposite signs across the two outcomes. For ratings, the coefficient is positive and significant ($\beta = 0.477$, $p < 0.01$); for votes, it is negative and significant ($\beta = -0.557$, $p < 0.01$). These results support H2a and H2b: films that bridge more distant narrative patterns tend to receive higher evaluations but attract smaller audiences.

Table 2: Regression results for trope concentration and spanning distance

	(1) Rating	(2) $\log(1 + \text{Votes})$
HHI_c	1.089^{***}	1.207^{***}
	(0.048)	(0.082)
CTD_c	0.477^{***}	-0.557^{***}
	(0.104)	(0.182)
Trope count	0.005^{***}	0.017^{***}
	(0.0001)	(0.0002)
Year FE	Yes	Yes
Main genre FE	Yes	Yes
Observations	17,108	17,108
R^2	0.258	0.409
Adj. R^2	0.251	0.404

Notes. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6 Discussion

Our results show that narrative structure shapes audience response, and that it does so differently for evaluation and attention. Coherence benefits films on both dimensions. We extended category-conformity arguments (Zuckerman, 1999; Askin and Mauskapf, 2017) from market-level genre labels to the narrative itself: a clear narrative structure helps not only in attracting audience attention but in generating favorable evaluations. Spanning distance reveals that evaluation and attention respond to different mechanisms. Films combining distant patterns may be harder to categorize from external signals, reducing their attention (Hsu, 2006). Yet for audiences who have already watched, narrative novelty may be appreciated rather than penalized (Uzzi et al., 2013). More broadly, this study shows that tropes can be treated not only as descriptive labels but as building blocks of narrative patterns whose combination helps explain success in cultural markets.

Limitations

This study has several limitations. First, TVTropes is a user-edited wiki, so popular films inevitably receive more thorough annotation. This creates a systematic bias between trope coverage and audience attention that our design cannot fully address. Second, the narrative patterns are sensitive to the resolution parameter of the Leiden algorithm, and the presence of a general community may compress spanning distance estimates.

References

Alberto Álvarez and Jose M. Font. 2022. TropeTwist: Trope-based narrative structure generation. In *Proceedings of the International Conference on Computational Creativity*.

Noah Askin and Michael Mauskapf. 2017. What makes popular culture popular? product features and optimal differentiation in music. *American Sociological Review*, 82(5):910–944.

Pierre Bourdieu. 1996. *The rules of art: Genesis and structure of the literary field*. Stanford University Press.

John G. Cawelti. 1976. *Adventure, Mystery, and Romance: Formula Stories as Art and Popular Culture*. University of Chicago Press.

Hsi-Ping Chou, Changxu Huang, and Yun-Nung Chen. 2023. TaleStream: Story generation with tropes. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

Paul DiMaggio. 1987. Classification in art. *American Sociological Review*, 52(4):440–455.

C Joseph Francemone, Kevin Kryston, and Matthew Grizzard. 2023. Examining genre success, co-occurrence, release, and production of 9,068 films over twenty years. *Poetics*, 99:101794.

Dhruvil Gala, Mohammad Omar Khursheed, Hannah Lerner, Brendan O'Connor, and Mohit Iyyer. 2020. Analyzing gender bias within narrative tropes. In *Proceedings of the Workshop on Natural Language Processing and Computational Social Science*.

Rubén Héctor García-Ortega, Pablo García-Sánchez, and Juan Julián Merelo. 2021. Trope community structure in TV series. In *Proceedings of the International Conference on Computational Intelligence in Music, Sound, Art and Design*. Springer.

Rubén Héctor García-Ortega, Pablo García-Sánchez, and Juan Julián Merelo-Guervós. 2020. Startroper, a film trope rating optimizer using machine learning and evolutionary algorithms. *Expert Systems*, 37(6):e12525.

Michael T Hannan and John Freeman. 1989. *Organizational ecology*. Harvard university press.

Greta Hsu. 2006. Jacks of all trades and masters of none: Audiences' reactions to spanning genres in feature film production. *Administrative Science Quarterly*, 51(3):420–450.

Greta Hsu, Michael T. Hannan, and Özgecan Koçak. 2009. Multiple category memberships in markets: An integrative theory and two empirical tests. *American Sociological Review*, 74(1):150–169.

Marc Keuschnigg. 2015. Product success in cultural markets: The mediating role of familiarity, peers, and experts. *Poetics*, 51:17–36.

Marc Keuschnigg and Thomas Wimmer. 2017. Is category spanning truly disadvantageous? new evidence from primary and secondary movie markets. *Social Forces*, 96(1):449–479.

Balázs Kovács and Michael T Hannan. 2015. Conceptual spaces and the consequences of category spanning. *Sociological science*, 2:252–286.

Jennifer C Lena and Richard A Peterson. 2008. Classification as culture: Types and trajectories of music genres. *American sociological review*, 73(5):697–718.

Chris Mellina and Svetlana Svetlichnaya. 2011. Trope propagation across cultural products. *arXiv preprint*.

Giacomo Negro and Ming D Leung. 2013. "actual" and perceptual effects of category spanning. *Organization Science*, 24(3):684–696.

Elizabeth G Pontikes. 2012. Two sides of the same coin: How ambiguous classification affects multiple audiences' evaluations. *Administrative Science Quarterly*, 57(1):81–118.

Hayagreeva Rao, Philippe Monin, and Rodolphe Durand. 2005. Border crossing: Bricolage and the erosion of categorical boundaries in french gastronomy. *American Sociological Review*, 70(6):968–991.

Stephen A Rhoades. 1993. The herfindahl-hirschman index. *Fed. Res. Bull.*, 79:188.

Stoyan V Sgourev, Erik Aadland, and Giovanni Formilan. 2023. Relations in aesthetic space: How color enables market positioning. *Administrative Science Quarterly*, 68(1):146–185.

Muzammil Hussain Shahid and Muhammad Arshad Islam. 2023. Investigation of time series-based genre popularity features for box office success prediction. *PeerJ Computer Science*, 9:e1603.

Vincent A. Traag, Ludo Waltman, and Nees Jan van Eck. 2019. From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports*, 9:5233.

Brian Uzzi, Satyam Mukherjee, Michael Stringer, and Ben Jones. 2013. Atypical combinations and scientific impact. *Science*, 342(6157):468–472.

Robert B Zajonc. 1968. Attitudinal effects of mere exposure. *Journal of personality and social psychology*, 9(2p2):1.

Ezra W. Zuckerman. 1999. The categorical imperative: Securities analysts and the illegitimacy discount. *American Journal of Sociology*, 104(5):1398–1438.

Table 3: Example trope communities identified by the Leiden algorithm for the co-tropes network.

ID	#Tropes	#Films	Dens.	Center trope	Representative tropes
0	10,573	18,164	0.170	Shout-Out	Shout-Out, Oh Crap!, Foreshadowing, Big Bad, Chekhov's Gun
1	3,658	10,475	0.084	Eldritch Abomination	Eldritch Abomination, Our Ghosts Are Different, Demonic Possession, Our Vampires Are Different, Sealed Evil In A Can
2	2,565	9,428	0.083	Doomed Upgrade	Doomed Upgrade, Sci-Fi Flyby, Humans Advance Swiftly, Proof of Commitment, No Warping Zone
3	1,703	7,878	0.090	Converted Into a Weapon	Converted Into a Weapon, Older Than Steam, The Speedster, Single-Power Superheroes, Superhero Capital of the World
4	1,453	7,288	0.129	Face Framed In Shadow	Face Framed In Shadow, Dramatic Unmask, Bank Robbery, No Honor Among Thieves, Evil Wears Black