

# Trope Propagation in the Cultural Space

Clayton Mellina  
cmellina@stanford.edu

Stacey Svetlichnaya  
staceys@stanford.edu

December 12, 2011

## Abstract

We investigate the patterns and evolution of cultural ideas and symbols in media using network analysis techniques. We leverage the TV Tropes wiki of media works, cross-referenced with the widely-recognized character/situation types (tropes) that these works contain, to construct a bipartite graph representation of 4,616 films and their associated tropes. We perform community detection on a weighted projection of this graph, explore the correlation between clustering coefficient and external measures of film popularity, and attempt to model trope propagation through works (specifically the process by which film-nodes link to trope-nodes over time). Overall, we find that the TV Tropes dataset is too noisy for the purposes of community detection. However, a trope-based similarity metric proves to be moderately predictive of external measures of film acclaim. Likewise, the change in the frequency of trope occurrences over time indicates a process of cultural drift in which new published works tend to contain the same tropes as more recent works with a higher likelihood than the tropes of older works.

## Introduction

Network analysis techniques have proven useful for understanding relationships between content items: blog posts, news articles, academic papers, tweets, etc. Many earlier studies model the spread and change of ideas by tracking specific words or phrases within a domain or media type (e.g. the blogosphere, news surrounding a certain

event, scientific journals). But how do more abstract and meaning-dense cultural ideas—the unlikely hero, the best friend, the forbidden love, the awkward parents—propagate through media narratives in general? We consider this broader scope of cultural evolution via the universally-recognizable symbols/characters/situations (tropes) of popular media catalogued in the TV Tropes (tvtropes.org) wiki.

TV Tropes is an open wiki that enables users to document common tropes and their appearance in a variety of media, including TV shows, literature, and films. The site defines a trope as a device or convention that a writer can reasonably rely on as being present in the minds and expectations of the audience. Thus, a trope is a unit of literary currency, recurring in works over time and gaining meaning through audience recognition of its connotations and associations. Tropes can be understood as a subset of memes: ideas, behaviours, or styles transmitted among people within a culture. While memes can spread through many forms of social interaction, tropes are spread through their repeated use in entertainment media—the media works themselves are the primary carriers of tropes. The visibility of a work—its cultural prominence, the size of the audience—is therefore a crucial factor in trope propagation. Critical and popular acclaim are reasonable indicators of visibility. Quantitative (if approximate) measures of such acclaim for films include box office earnings and admissions, number of awards, average viewer rating, total budget, and rentals. This external data enables us to validate and contextualize our network-based findings and is easily obtained from IMDb, the Internet Movie Database. Given the difficulty of parsing media type from TV Tropes and of acquiring

standardized popularity measures for books, comics, anime, etc., we focus most of our analysis on the subset of TV Tropes works that are films.

Using the trope-work graph and the aforementioned IMDb film data, we address the following questions: 1) to what extent can trope patterns explain the clustering of similar works by genre?, 2) does the originality of a work, quantified as trope-based similarity to prior works, correlate to critical or popular acclaim?, and 3) how can we model work-node attachment over time—i.e. which tropes will appear in a new work based on the trope/work connectivity in previous time slices?

## Data Preprocessing

DBTropes.org provides a linked-data wrapper over TVTropes.org (Malte, 2010). Our primary data source is the daily data dump of the wiki in RDF N-Triples format available from DBTropes.org. The dump includes the full text content of the wiki and the full hierarchical structure of the trope and work pages. In particular, the RDF encodes all works currently listed in the wiki along with all tropes identified as occurring in each work. Furthermore, both tropes and works are structured within a category hierarchy which encodes, among other things, genre and media type. We see these categorizations as post hoc, with tropes and works providing the primary structure encoded in the wiki, although we utilize category memberships in parsing tropes and works into a network.

After RDF statements are read, we resolve double listings for items and take the subset of works identified as films. In general, works of different media types are not directly comparable without attention to possible confounds. For instance, long-running TV series have a greater opportunity for employing tropes in their narrative and therefore tend to have higher numbers of tropes than do films (e.g., the long-running British sci-fi series *Doctor Who* has 4977 tropes, the most for a single work). Thus, characteristics of work media type influence the documentation of trope inclusion on the wiki. For this reason, and because IMDb provides access to high quality data for movies, we choose to restrict many analyses

to the subset of works in the film category. Preliminary comparison—namely of the degree distributions of films to that of all works—suggests that the connectivity pattern of the film subset is representative of the full set of works (compare the middle and bottom degree distributions on the next page).

We acquire IMDb data using the open-source IMDbpY library. For each movie extracted from DBTropes, we retrieve the basic information: full IMDb title, release year, directors, producers, production companies, and genres. We further parse award features: the number of award nominations, awards won, Oscar nominations, and Oscars won, along with the average IMDb.com user rating and the number of users who rated the film. We also parse box office features: the U.S. gross, worldwide gross, opening weekend earnings, maximum gross earnings over weekends in U.S. theatres, and earnings from rentals, adjusting all these dollar amounts for inflation to 2010 dollars. Note that IMDb does not provide all these values for every movie, hence we use the maximum number of movies for which the relevant feature values are available when testing for correlations to film originality.

## Descriptive Summary

The data dump contains a total of 24,154 work pages, of which 23,687 are unique after double listing information (primarily due to disambiguation pages) is resolved. In total there are 21,538 tropes. A bipartite graph between works and the tropes they contain has nearly 1.6 millions edges, each corresponding to an instance of a trope in a work. Below, we plot the log-log degree distribution for trope nodes, work nodes, and film nodes. The most common trope, plotted as the outlier in the middle figure, is the *Shout Out* trope, which is included in any work that makes an explicit reference to another work. A power-law tail is apparent in all three distributions, although this trend is violated for low degree nodes (a power-law fit to the full distributions is not significant).

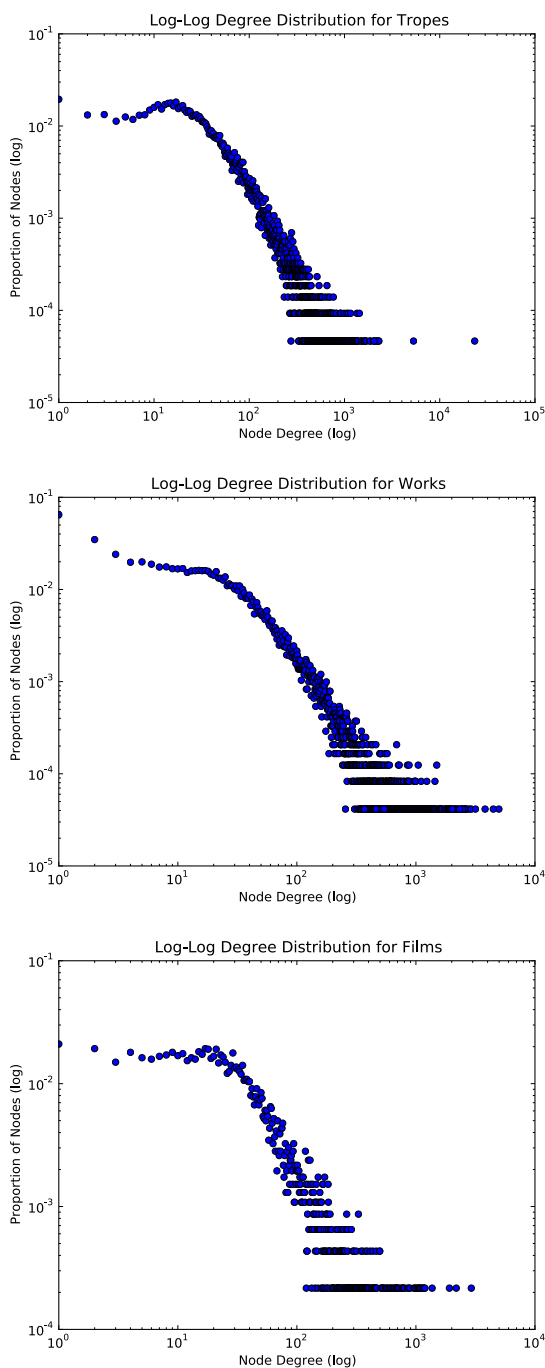
## Network Representations

We use two main network representations in our investigation. The first is the bipartite film graph, the subgraph of the bipartite work graph in which all work nodes are films. The bipartite film graph  $B = (F, T, E)$ , where  $F$  is the set of 4,616 films,  $T$  is the set of 19,142 tropes, and  $E \subseteq (F \times T)$  is the set of 270,980 undirected edges from films to the tropes they feature. A more efficient representation for our community detection analysis and other experiments is a weighted projection of  $B$ ,  $P = (F, E_w)$  in which  $F$  is the set of film nodes and  $E_w$  a set of weighted undirected edges. Each  $(u, v) \in E_w$  is weighted by an intuitive measure of the similarity of the two films  $u$  and  $v$  in terms of tropes ( $N(u), N(v) \subset T$ ) they share. Formally, this is the Jaccard coefficient, defined as the size of the intersection divided by the size of the union of the trope sets:  $w(u, v) = J(u, v) = \frac{N(u) \cap N(v)}{N(u) \cup N(v)}$ . The  $J(u, v)$  weighting scheme aims to normalize the influence of the degree of the  $F$  nodes, assigning higher weights to higher proportions rather than to higher numbers of shared tropes. It has also been found to give high values of precision and good compromises between precision and recall in the specific application of link prediction (Allali et al., 2011).  $E_w$  in the resulting projection  $P$  is about 3.7 million.

## Community Detection

Our first question concerned the dependence, if any, of film genre (drama, comedy, action, etc.) on trope inclusion. Viewing experience suggest that certain tropes are more prevalent in/more characteristic of some genres than others—for example, stormy nights and deserted roads seem to be very popular in horror films. We therefore expect that clustering films by trope-based similarity to other films will yield a partitioning into communities such that genre is more homogeneous within communities and more heterogeneous between communities.

The IMDb data we collected for films in the graph contain multiple genre labels per film, with a total of 27 distinct genres. For simplicity, we analyzed each genre labeling independently. We hypothesized



that films of each genre will occur disproportionately within the same community. This finding would indicate that trope-based similarity is predictive of film inclusion in a genre group. Formally, we consider community detection to be the assignment of films to community bins and expect that this assignment of films to bins will result in the assignment of films of the same genre to the same bins. If bin assignment is independent of genre distribution then the proportion of a genre  $g \in G$  in a community  $c \in C$  is only expected to be the proportion of total film-nodes in the community,  $c/P$ . For each  $c$  and  $g$ , we compare this expected proportion to the observed proportion of films of genre  $g$  in  $c$ :  $\frac{|\{u \in c | \text{genre}(u) = g\}|}{|\{u \in P | \text{genre}(u) = g\}|}$ .

We attempt to extract community structure from  $P$  using the heuristic modularity optimization method presented in Blondel et. al. (2008), shown to outperform other community detection methods on computation time and produce high values for modularity on real networks. This produces 5 communities with a relatively low modularity of 0.137. We tabulated the observed number of films falling in each community for each genre and performed a  $\chi^2$  test for each against the expected number computed from community size as described above. We obtained significance on most of our genres, indicating that community detection does not assign nodes independently of genre. For comparison, we generate a random bipartite graph  $B_r = (F, T, E_r)$  while keeping the degree of each  $u \in F$  and each  $v \in T$  constant. Specifically, we split each edge in the bipartite graph and randomly reconnect the resulting edge spokes. This effectively randomizes the pattern of trope similarity between films without affecting the relative frequency of trope inclusion in films. We then construct a weighted projection  $R$  of  $B_r$  and extract community structure from  $R$  using the same algorithm. Despite the randomization, this also yields 5 communities and a slightly higher modularity of 0.149. A repeated  $\chi^2$  test for this community clustering yielded a similar proportion of significant  $\chi^2$  tests. Thus, it seems that the degree sequence of the network can itself produce seemingly interesting partitions of genre. The significance of these tests on the randomized graph suggest that genre and the size of the trope set of a film are not independent in our data

set. This may reflect genre biases within the wiki community itself. These results suggest that trope similarity is not a sufficiently strong signal to produce a good clustering, much less for drawing genre boundaries.

We believe that a major reason for the insufficiency of trope similarity to produce high-modularity clustering is due to the extremely high proportion of very low edge weights. Since most of the edge weight mass is in low-weighted edges, higher-weighted edges more indicative of similarity do not play a proportionally large role in clustering. The high proportion of edges within a narrow band of edge weight causes the modularity-based clustering algorithm to behave as if the graph were effectively unweighted. Loosely following a divisive algorithm of removing edges in order of increasing weight (Fortunato, 2009), we threshold the edge weight of both the real and randomized graphs, removing all edges with weight lower than 0.05. This resulted in the removal of about 90% of the edges from each graph. This method of community detection obtained higher modularities: 0.46 and 0.53 on the real and randomized graphs respectively. Although modularity improved as expected, the similar modularity values indicate a high likelihood of these communities being due to chance. The edge weight distributions for  $P$  and  $R$  are shown on the next page.

## Work Uniqueness and Acclaim

Our interest in the existence of a correlation between work uniqueness and its popular/critical acclaim was partially inspired by the connection between structural holes and good ideas (Burt, 2004). Burt considers individuals residing near structural holes or acting as local bridges in organizational networks to be advantageously placed, resulting in higher creativity. External metrics for creativity are offered as evidence for this theory. In the absence of a meaningful clustering, we approximate structural holes as films  $u$  that are on average less similar to other films  $v \in F$ , as quantified by the average clustering coefficient for bipartite graphs given in Latapy et al. (2008). In our case, this is the average weight

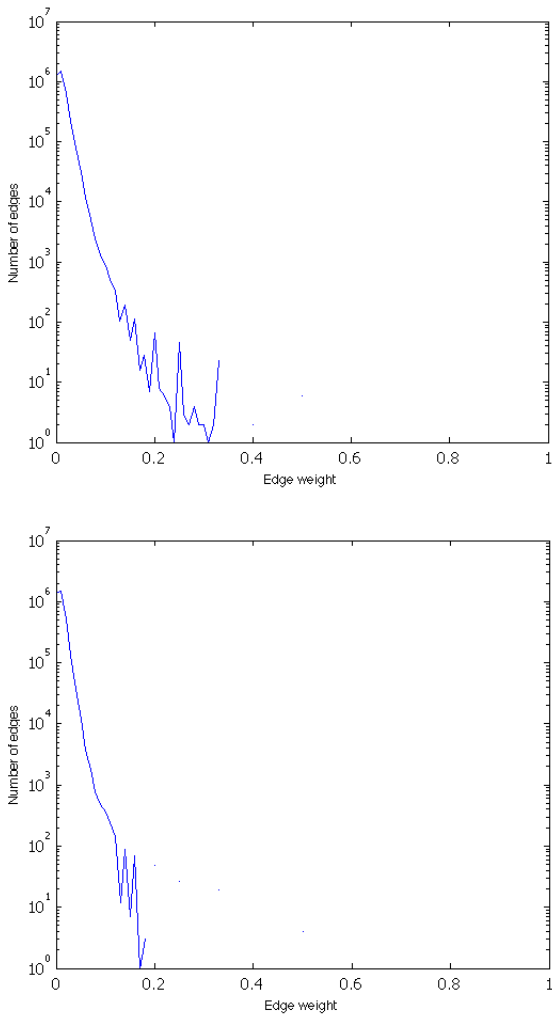


Figure 1: Edge weight distribution for real (top) and randomized (bottom) projected graphs.

of the edges of a given film-node in  $P$ :

$$cc(u) = \frac{\sum_{v \in N(u) | year(v) \leq year(u)} J(u, v)}{|N(u)|} \quad (1)$$

Note the restriction that  $v$  must be a film released in the same year as or earlier than  $u$  to appropriately exclude not-yet-existent films from the context of comparison. The higher the value of  $cc(u)$ , the more similar the film is to other works; the lower this value, the more original the film. We hypothesize that unique combinations of tropes are more creative narratives in the cultural space and that this creativity in films correlates to popular and critical acclaim.

We evaluate this prediction by computing the  $cc(u)$  for all films in  $P$  and testing for a correlation between  $cc(u)$  and the set of numeric feature values parsed via IMDb (award and box office features: the number of award nominations, awards won, Oscar nominations, and Oscars won; the average IMDb.com user rating and the number of users who rated the film, the U.S. gross, worldwide gross, opening weekend earnings, maximum gross earnings over weekends in U.S. theatres; and earnings from rentals, adjusting all these dollar amounts for inflation to 2010 dollars). Only a few of these comparisons resulted in correlation values high enough to be notable. Our overall similarity metric correlated with IMDb rating ( $r(2645) = 0.16$ ,  $p < 0.001$ ), number of IMDb votes ( $r(2645) = 0.27$ ,  $p < 0.001$ ), inflation-adjusted US gross revenue ( $r(2645) = 0.13$ ,  $p < 0.001$ ), and total award nominations ( $r(2645) = 0.10$ ,  $p < 0.001$ ). Although, these correlations are small, they indicate that trope-based similarity of a work to existing works is at least partially predictive of film popularity as operationalized by those network-external variables.

## Evolution of the Bipartite Graph

Taking inspiration from Leskovec et al. (2010), we consider network features that could be used to predict, for a given time-slice of the network, how works created in the next time-slice will connect to already existing nodes. We consider the evolution of the bipartite network from year to year as a process in which new films arrive at each time step and


simultaneously select a set of trope nodes to which to link. For simplicity, we assume that each arriving work-node forms edges to trope-nodes independently. Most previous work has investigated link prediction for a fixed set of nodes, not link prediction for cases in which a new node is added to the network (Kunegis et al., 2010, Allali et al., 2011).

We hypothesized that the inclusion of a given trope in a new work at a given year  $t$  is dependent on the frequency of that trope in the recent years preceding  $t$ . For simplicity, we assume that tropes occur in works independently from one another. Thus, we model work-node arrival as a process in which a film-node flips a weighted coin for each trope to determine whether it will connect to the corresponding trope-node. Let  $T_{tr,t}$  be the set of edges to trope-node  $tr$  added by works arriving in the network at year  $t$ , and let  $W_t$  be the set of film-nodes added to the network at year  $t$ . Thus, the frequency of occurrence of any given trope  $tr$  in films arriving at year  $t$  is computed as  $freq_{tr,t} = |T_{tr,t}|/|W_t|$ . Our hypothesis can be formalized by computing the percentage difference between  $F_{tr,t}$  and  $F_{tr,t-offset}$  for year  $t$  and some previous year  $t - offset$  where  $offset > 1$ . I.e.

$$PD_{tr,t,t-offset} = \frac{|freq_{tr,t} - freq_{tr,t-offset}|}{(freq_{tr,t} + freq_{tr,t-offset})/2} \quad (2)$$

where  $PD_{tr,t,t-offset}$  denotes the percentage difference expressed as a proportion of 1. The average of  $PD_{tr,t,t-offset}$  over all tropes within a given year  $t$ , denoted  $avgPD_{t,t-offset}$ , provides a summary statistic of the extent to which the overall frequency of trope occurrences at year  $t$  differs from the frequency of trope occurrences at year  $t - offset$ . Thus, we hypothesize that there is a positive correlation between  $avgPD_{t,t-offset}$  and  $offset$ . That is, we expect that the frequency of trope occurrences at year  $t$  will differ from the frequency for earlier years more than from that of recent years.

In general, the frequency of most tropes for a given year is extremely low, with most not occurring at all. We expect that frequent tropes for a time period are more indicative of the cultural preference for



tropes in that time period than are the time periods less frequent tropes. For these reasons, we compute  $PD_{tr,t,t-offset}$  for only the top 20 most frequent tropes at year  $t$ . Similarly, we limit our analysis to the period from 1980 to 2010, as earlier years have poor representation in our dataset. Following this procedure, we compute  $avgPD_{t,t-offset}$  for each year from 1980 to 2010 and for each offset from 1 to 10 (this means that the earliest data we use for comparison is from 1970).

We perform two correlation tests on the resulting  $avgPD_{t,t-offset}$  values. First, we find an average  $avgPD_{t,t-offset}$  for each offset value by averaging  $avgPD_{t,t-offset}$  for all years in our target period. This results in 10 numbers, each summarizing the extent to which trope occurrence within a year differs from trope occurrence offset in the past for all years in our target interval. Our prediction is born out, with  $r(8) = 0.92$ ,  $p < 0.001$ . We also compute the correlation of  $avgPD_{t,t-offset}$  to its corresponding offset, resulting in a lower but still significant correlation,  $r(308) = 0.19$ ,  $p < 0.001$ . This smaller correlation suggests that, on the whole, films use recent tropes, but from year to year the extent to which trope frequency depends on recent films varies. This is to be expected, as notable works can influence cultural preferences for future works and serve as cultural touchstones.

## Conclusion

We applied network analysis techniques to user-curated content on the TV Tropes wiki. The special meaning of page types on the wiki allowed us to investigate broad trends evident in the patterns of trope inclusion in published films. While modularity-based community detection methods failed to yield usable community structure in our analysis, the Jaccard weighted film-graph did provide a means of testing the correlation of trope occurrence patterns on network-external popularity metrics. Our findings suggest that popularity, as operationalized by a selection of IMDb metrics, is correlated with the similarity of a film to existing works at the time of its publication, although we cannot conclude that this effect is independent of user-bias. Lastly, trope occurrence patterns in newly published films seems to adhere to a recency effect. It remains to be seen whether this recency effect is driven by the film industry itself or whether trope occurrence patterns track broader patterns of cultural exchange. Future work might attempt to answer specific questions about the history of the film industry and its relationship to broader culture through techniques similar to those explored here.

## References

- Allali, O., Magnien C., Latapy, M. (2011). Link prediction in bipartite graphs using internal links and weighted projection. 2011 IEEE Conference on Computer Communications Workshops, 936-941.
- Blondel, V. et al. (2008). Fast unfolding of communities in large networks. *J. Stat. Mech.*, 10.
- Burt, R. (2004). Structural holes and good ideas. *Am J Soc.*, 110 (2), 349-399.
- Fortunato, S. (2009). Community detection in graphs. *Physics Reports*, 486, 75-174.
- Kunegis, J., De Luca, E., Albayrak, S. (2010). The link prediction problem in bipartite networks. *LNCS*, 6178, 380-389.
- Latapy, M., Magnien, C., Del Vecchio, N. (2008). Basic notions for the analysis of large two-mode networks. *Social Networks*, 30(1), 3148.
- Leskovec, J., Huttenlocher, D., Kleinberg, J.

(2010). Predicting positive and negative links in online social networks. In *Proc. WWW*.

Malte K., Grimnes, G. (2010). DBTropes—a linked data wrapper approach incorporating community feedback. Poster track, *EKAW*.