

# DYPLODOC: Dynamic Plots for Document Classification

Anastasia MALYSHEVA <sup>a</sup>, Alexey TIKHONOV <sup>b</sup> and Ivan P. YAMSHCHIKOV <sup>c,1</sup>

<sup>a</sup>Open Data Science, Moscow, Russia

<sup>b</sup>Yandex, Berlin, Germany

<sup>c</sup>LEYA Lab, Yandex and Higher School of Economics in St. Petersburg, Russia

**Abstract.** Narrative generation and analysis are still on the fringe of modern natural language processing yet are crucial in a variety of applications. This paper proposes a feature extraction method for plot dynamics. We present a dataset that consists of the plot descriptions for thirteen thousand TV shows alongside meta-information on their genres and dynamic plots extracted from them. We validate the proposed tool for plot dynamics extraction and discuss possible applications of this method to the tasks of narrative analysis and generation.

**Keywords.** narrative analysis, plot dynamics, plot structure, plot extraction

## 1. Introduction

Natural Language Generation (NLG) is one of the areas of Natural Language Processing (NLP). Deep learning enabled various generative applications where generated texts are short and constrained by the context. These could be dialogue responses [1,2], traditional or stylized poetry [3,4,5], end-to-end author or source stylization [6,7,8]. One of the bottle-necks that holds modern Natural Language Processing (NLP) from the generation of longer texts is the concept of *narrative*. There are constant attempts to generate longer blocks of text, such as [9] or [10], yet they succeed under certain stylistic and topical constraints that exclude the problem of narrative generation altogether [11].

Though philosophers and linguists try to conceptualize the notions of plot, narrative arc, action, and actor for almost a century [12,13,14], few of these theoretical concepts could be instrumental for modern NLP. [15] present a machine comprehension corpus for the end-to-end evaluation of script knowledge. The authors demonstrate that though the task is not challenging to humans, existing machine comprehension models fail to perform well, even if they make use of a commonsense knowledge base. Despite these discouraging results, there are various attempts to advance narrative generation within the NLP community, see [16,17]. For a detailed review of approaches to narrative generation, we address the reader to [18], here we just mention several ideas that are relevant within the scope of this paper.

---

<sup>1</sup>Corresponding Author: LEYA Lab, Yandex and Higher School of Economics in St. Petersburg; Kantemirovskaya st. 3, Saint-Petersburg, Russia; E-mail: ivan@yamshchikov.info

One line of thought is centered around some form of a multi-agent system with various constraints that keep generated texts within some proximity of human-written narrative. For example, the paper [19] presents a multi-agent framework for automatic story generation where plots are automatically created by semi-autonomous characters while the resulting plot is converted into text by a narrator agent. [20] experiment with a dialogic form of a story: a deep representation of a story is converted into a dialog that unfolds between two agents with distinctive stylistic characteristics. Such multiagent-based systems show a lot of potential yet do not leverage the existing datasets that include an extensive amount of narratives written by humans.

Another line of research tries to use some NLP methods and statistical procedures to systematize existing narratives using some form of synthetic features that could in some form describe the narrative structure of a given story. [21] represents a story as a cluster of emotional links and tensions between characters that progress over story-time. [22] introduce a concept of narrative arc in a context of large-scale text mining. The author conceptualizes plot structure as a path through a multidimensional space derived from a topic model. This approach allows us to see plot dynamics within any given book or show in detail, yet it is harder to generalize. Topic modeling tends to provide profoundly different sets of topics across genres. When aggregating across all topics within a dataset one gets a high dimensional space that is harder to interpret. For example, similar narrative arcs of a detective story would look different in terms of topic modeling if the same plot takes place in Victorian England or on Mars. One would like to keep the concept of a narrative arc yet to make it more abstract in order to be able to carry comparative narrative research across genres and settings. [23] develop ideas that address this challenge to a certain extent. The authors mathematize a concept of narrative arc that depicts the evolution of an agent's belief over a set of the so-called *universes*. The authors estimate these probabilities with pre-trained probabilistic topic classifiers and combine these estimates with a multi-agent dialogic approach. They demonstrate that the agent that tracks the dynamics of such narrative arcs generates more engaging dialog responses.

This paper shows that the concept of the narrative arc could be applied to a vast variety of narratives in natural languages and does not have to be limited to the setting of a dialogue system. The contribution of this paper is two-fold: (1) we provide a large dataset of longer storylines that could be used for narrative research and (2) use it to demonstrate how the notion of the narrative arc could be applied to longer texts and provide meaningful insights into the structure of the story. This structural insight could be useful both for quantitative research of narrative dynamics and as synthetic features for the generation of narrative structures in the future.

## 2. Data

TVmaze<sup>2</sup> is a huge database with the data on different TV shows that include the storyline as well as certain meta information. Using TVmaze API we have obtained 300+ Mb of data that could be found online as a part of the DYPLODOC package<sup>3</sup>. The dataset includes 13 813 TV shows and 21 962 seasons that consist of more than 300 000 episodes.

---

<sup>2</sup><https://www.tvmaze.com/>

<sup>3</sup><https://github.com/AnastasiaMalysheva/dyplodoc>

Every episode in the dataset is annotated with a condensed description of the main storyline and the narrative structure of a season is defined as a sequence of such descriptions. The data includes main characters, that could be extracted with the NER algorithm, the majority of texts are centered around most reportable events [24], and the narrative is split into clear steps, certain 'quanta' of a story. Almost every season in TVmaze has a tag of a genre, such as (Horror, Food, Drama, etc). We hope that the publication of the dataset could facilitate further quantitative research of narrative structures.

### 3. Plot Dynamics Extraction Tool

Here we leverage the fact that episode descriptions in the TVmaze dataset form seasons and demonstrate how one could build the narrative arc of a season by constructing artificial features that would to a certain degree describe the season's plot dynamics. We call this method Plot Dynamics Extraction (PDE) Tool.

Let us regard a specific TV show  $X$  as a set of  $N$  seasons  $X : \{X_1, X_2, \dots, X_N\}$ . Every season  $X_i$  could be represented as a sequence of episodes  $X_i : \{y_1, y_2, \dots, y_M\}$ . Every  $y_i$  is a description of one particular episode.

We investigated the plot dynamics using two different pre-trained probabilistic topic classifiers. However, unlike [22] we choose fixed topics that form a universal state space of narrative arcs. One could discuss if the provided list of dimensions is complete, yet we believe that two classifiers are sufficient to extract non-trivial narrative arcs out of a variety of texts:

- Classifier of movie genres trained on Wikipedia film synopses provided in [25]. The labels of genres include: Comedy, Drama, Western, Adventure, Animation, Action, Thriller, Family, Romance, Fantasy, Horror, History, Music, Sci-Fi, War, Crime, Musical, Biography, Mystery, Sport.
- DeepMoji: Deep neural network that was trained on a corpus of text that included 64 emojis, see [26]. In their blogpost<sup>4</sup> the authors suggest hierarchic aggregation of 64 emojis. Using this aggregation we limit ourselves to 11 labels in this work. These labels include: Love, Happy, Wink, Deal, Force, Eyes, Fear, Mad, Sad, Music, Misc.

Thus every narrative arc is a trajectory on a union of two simplexes: a 20-simplex of genres and an 11-simplex of emotions. For the vast majority of the cases, the classifier mostly detects one or two most dominant genres and topics respectively, limiting the meaningful development of a given narrative art to one plane on each simplex.

Similarly to [23] we propose to regard a finite set of universes,  $\mathcal{U}$ . One could develop a universe model to assess the compatibility of an episode  $y_i$  with a given universe,  $u \in \mathcal{U}$ . Given such a model, one could update the posterior universe distribution over a sequence of episodes  $y_i$ . For each universe  $u$ , the universe model assigns a likelihood  $p(y_i | y_{1:i-1}, u)$  to an episode  $y_i$ , conditioned on previous episodes  $y_{1:i-1}$  of the season. Introducing  $z(u)$  as the prior distribution over universes, the conditional probability  $z(u | y_i)$  could be defined as follows  $z(u | y_i) = z(u) \times \frac{z(y_i | u)}{z(y_i)}$ .

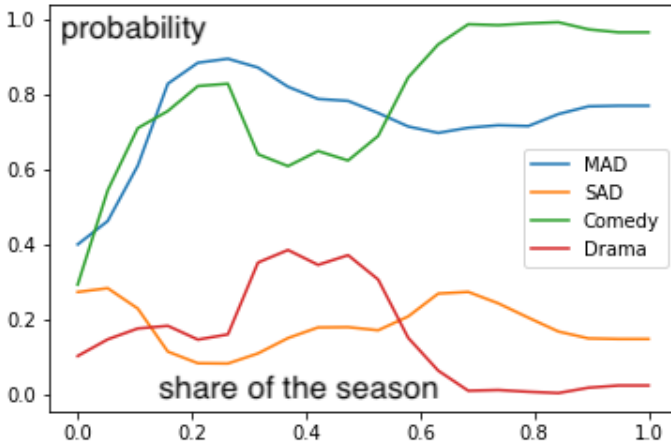
Using this notation [23] obtain the following approximation

<sup>4</sup><https://medium.com/@bjarkefelbo/what-can-we-learn-from-emojis-6beb165a5ea0>

$$p(u|y_{1:i}) = p(u|y_{1:i-1}) \times \frac{z(u|y_i)}{\sum_{u'} p_{i-1}(u'|y_{1:i-1}) z(u'|y_i)} \quad (1)$$

In order to guarantee that the sequence of  $p(u|y_{1:i})$  forms a series of coordinates on the union of 20-simplex of genres and 11-simplex of emoji, we introduce two additional smoothing procedures. First, we control  $z(u)$  distribution update so that it does not degenerate into zero, since in this case Equation (1) becomes obsolete. We also use exponential smoothing average instead of standard multiplicative update.

Figure 1 illustrates the resulting dynamics of one season of 'BoJack Horseman', extracted with a Deepmoji classifier and with a classifier pre-trained on the genres<sup>5</sup>. There are four major non-zero components. On the emoji-simplex the narrative is balancing between MAD and SAD, though the former is prevailing throughout the season, the latter peaks in the episode when BoJack's daughter overdoses and is found in the hospital. On the genre-simplex, the narrative is balancing between Comedy and Drama. BoJack Horseman is a dark comedy animation series so it stands to reason that Comedy dominates the season while Drama peaks in the second quarter of the season when BoJack spends time with his elderly mother in the nursing home.



**Figure 1.** Dynamics of the fourth season of 'BoJack Horseman'. Tags from the simplex of emotions are denoted with capital letters.

#### 4. Experimenting with PDE Tool

Though scores of pre-trained deepmoji and genre classifiers have very intuitive interpretation it is hard to claim that they could be useful for any form of comparative narrative

<sup>5</sup>PDE Tool is available as a part of DYPLODOC package <https://github.com/AnastasiaMalysheva/dyplodoc>

research. The dynamic aspect of these constructed features is clear yet it is not obvious that their dynamics are meaningful in any non-trivial sense. We compare dynamic plots extracted from the season description with static meta-information that is included in the TVmaze data<sup>6</sup>. If different episodes of the season are labeled with a different tag we choose the most frequent tag as the tag for the whole season. If the show was tagged with two values<sup>7</sup>, we treated it like one static tag. For example, 'BoJack Horseman' has Drama | Comedy as a genre, thus Drama, Comedy and Drama | Comedy for three separate static tags.

PDE-tool represents every plot as an array of tags (such as genres or emoticons), in which every tag is represented as a sequence of probability estimations ranging from 0 to 1. We concatenate these representations with extra synthetic features, namely: mean, median, minimum, maximum, and variance for the probability of every tag; area under the curve; the moment of intersection with another tag probability estimation, etc. Thus we represent every text as a multidimensional vector. Using these dynamic based features we clustered the seasons, applying standard K-means clustering over a PCA<sup>8</sup>. We have carried our experiments across all possible combinations of features, a different number of principal components (50, 100, 200, 300), and a various number of clusters (50, 75, 100, 150, 200, 250, 300). Then we calculated mutual information between obtained clusters and static tags for the corresponding seasons. The higher values of mutual information were obtained for the following setup: synthetic features included statistics and area under the curve for every component; the number of principal components equals 200; the number of clusters for K-means equals 200.

With this setup, the mutual information between clusters based on both classifier combined and static tags was 0.57, as shown in Table 1. One could see that both the dynamics based on the genre classifier and the emoji classifier-based dynamics allow inferring static genre tags. The combination of two classifiers yields higher mutual information. Dynamic plots extracted by different classifiers differ significantly and go beyond static genre tags.

## 5. Discussion

Figure 2 shows the resulting t-SNE (t-distributed stochastic neighbor embedding) visualization, see [27] for details, of the clustering based on plot dynamics. The centers of the clusters are labeled with static genre tags. The average size of a cluster in the resulting clustering is 109 and the average share of the dominant tag is 65%. The smallest of the clusters consists of 9 show seasons, the biggest includes 288 seasons. In some clusters the share of the dominant tag reaches 100%, the lowest possible percentage of the dominant tag is 42%.

Using synthetic features based on the PDE tool one could partially infer static human labels for season's genre. Table 1 demonstrates that the dynamic plot structure is non-

<sup>6</sup>7911 seasons out of 21962 have no tags so we exclude them from this experiment.

<sup>7</sup>We filtered the shows that had more than 2 tags for their genre.

<sup>8</sup>Since the length of the seasons might differ and we are interested in comparative research of the obtained arcs we scale every series to the length of ten episodes either uniformly stretching or compressing the resulting timeline.

**Table 1.** Mutual Information (denoted as MI) between k-means clusters based on the features obtained with plot dynamics extraction and the clusters defined by static genre meta-information alongside with the silhouette for the obtained clusterizations.

	MI	Silhouette
Dynamic Genres	0.47	0.23
Dynamic Emojis	0.36	0.16
Combined Dynamic Features	0.57	0.20

**Table 2.** Three biggest clusters in the resulting clusterization that have the highest percentage of a dominant tag.

# of shows	Dominant Static Tag	% of Dominant Static Tag
274	Comedy   Drama"	74
260	Drama   Fantasy"	79
253	Drama   Comedy"	72

trivial and could be useful for quantitative research of narrative. We publish extracted plot dynamics along with the TVmaze data to facilitate further research.

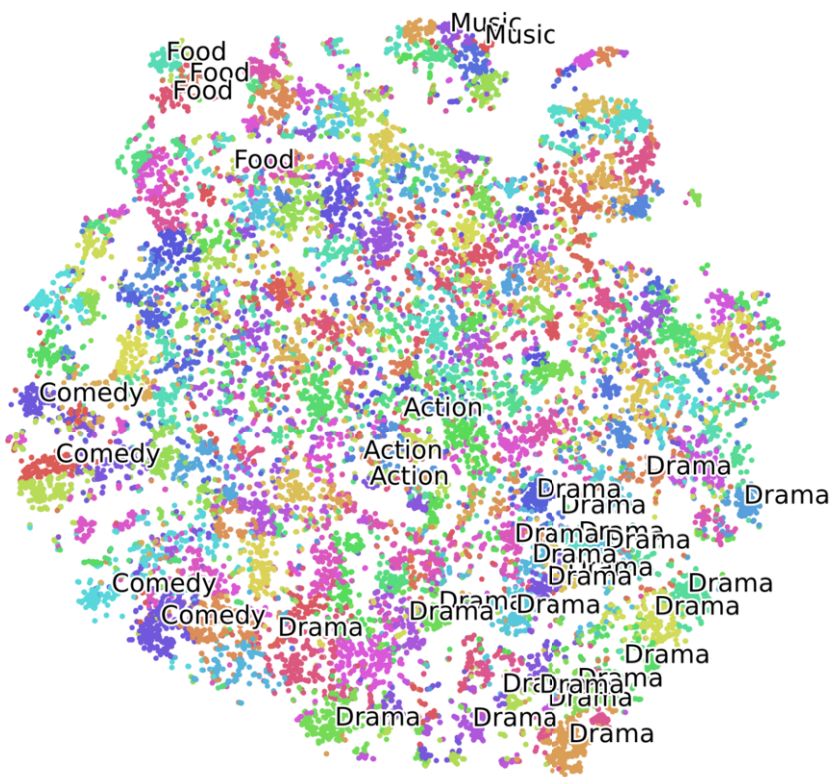
Now let us show how the PDE tool could be used for comparative research of narrative. Let us zoom in on the 200 clusters obtained from dynamics plots extracted from the descriptions of the seasons. Table 2 shows the three biggest resulting clusters with the highest percentage of one static genre tags.

The first cluster could be attributed to darker comedies. Figure 1 shown above is a good representative of this cluster. The second cluster could be attributed to historic dramas. The cluster is characterized by the dominance of Drama and interplay of Sad and Mad emojis. Figure 3 illustrates it with the dynamics of the second season of Victoria. The third cluster is also dominated by Drama but with the certain presence of Comedy as well as with domination of Mad over Sad, see Figure 4 for an example. These are only a few examples that illustrate how rich is the information provided by the PDE tool and its potential for quantitative analysis of narrative.

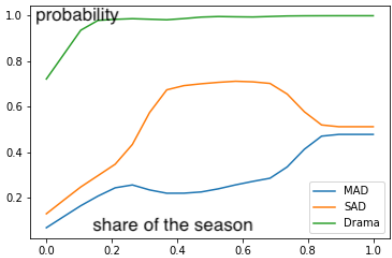
6. Conclusion

This work makes the following contributions: (1) a universal feature extraction method for plot dynamics; (2) a dataset of natural language plot descriptions; (3) finally, the paper shows that plot dynamics features extracted with the proposed method are meaningful and can shed light on the general structure of the narrative.

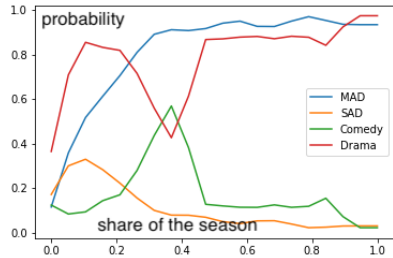
We hope that this paper could help to revive the attention of the community to the problems of narrative analysis and generation at the very least by providing data and tools for further exploration. We also hope to use narrative arch extracted with PDE tool as a form of semi-supervision for narrative-oriented generative models. The extracted dynamic plots are presented alongside the textual data. We hope that the proposed tool would inspire further research on plot dynamics and inform the discussions on narrative structure, particularly in the context of narrative generation.



**Figure 2.** t-SNE visualization of the clustering based on plot dynamics. The labels mark clusters in which the dominant genre tag corresponds to 70% or more shows in the cluster.



**Figure 3.** Dynamics of the second season of 'Victoria'. Tags from the simplex of emotions are denoted with capital letters.



**Figure 4.** Dynamics of the third season of 'Orange is a New Black'. Tags from the simplex of emotions are denoted with capital letters.

References

[1] Li J, Monroe W, Ritter A, Jurafsky D, Galley M, Gao J. Deep Reinforcement Learning for Dialogue Generation. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing; 2016. p. 1192–1202.

- [2] Li Z, Kiseleva J, de Rijke M. Dialogue generation: From imitation learning to inverse reinforcement learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33; 2019. p. 6722–6729.
- [3] Zhang X, Lapata M. Chinese poetry generation with recurrent neural networks. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*; 2014. p. 670–680.
- [4] Wang DZ, He W, Wu H, Wu H, Li W, Wang H, et al. Chinese Poetry Generation with Planning based Neural Network. In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*; 2016. p. 1051–1060.
- [5] Tikhonov A, Yamshchikov IP. Guess who? Multilingual approach for the automated generation of author-stylized poetry. In: *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE; 2018. p. 787–794.
- [6] Hu Z, Yang Z, Liang X, Salakhutdinov R, Xing EP. Toward Controlled Generation of Text. In: *International Conference on Machine Learning*; 2017. p. 1587–1596. Available from: <http://proceedings.mlr.press/v70/hu17e.html>.
- [7] Tikhonov A, Shibaev V, Nagaev A, Nugmanova A, Yamshchikov IP. Style transfer for texts: Retrain, report errors, compare with rewrites. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*; 2019. p. 3927–3936.
- [8] Syed B, Verma G, Srinivasan BV, Natarajan A, Varma V. Adapting Language Models for Non-Parallel Author-Stylized Rewriting. In: *AAAI*; 2020. p. 9008–9015.
- [9] Kedzierski RK. Understanding and Generating Multi-Sentence Texts; 2019.
- [10] Agafonova Y, Tikhonov A, Yamshchikov IP. Paranoid Transformer: Reading Narrative of Madness as Computational Approach to Creativity. *Future Internet*. 2020;12(11):182.
- [11] van Stegeren J, Theune M. Narrative Generation in the Wild: Methods from NaNoGenMo. In: *Proceedings of the Second Workshop on Storytelling*; 2019. p. 65–74.
- [12] Shklovsky V. *Theory of prose* (B. Sher, Trans.). Champaign, IL: Dalkey Archive Press Original work published. 1925.
- [13] Propp V. *Morphology of the Folktale*, trans. Louis Wagner, 2d ed. 1968.
- [14] Van Dijk TA. Philosophy of action and theory of narrative. *Poetics*. 1976;5(4):287–338.
- [15] Ostermann S, Roth M, Pinkal M. MCScript2. 0: A Machine Comprehension Corpus Focused on Script Events and Participants. In: *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (\* SEM 2019)*; 2019. p. 103–117.
- [16] Fan A, Lewis M, Dauphin Y. Strategies for Structuring Story Generation. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*; 2019. p. 2650–2660.
- [17] Ammanabrolu P, Tien E, Cheung W, Luo Z, Ma W, Martin LJ, et al. Story realization: Expanding plot events into sentences. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34; 2020. p. 7375–7382.
- [18] Kybartas B, Bidarra R. A survey on story generation techniques for authoring computational narratives. *IEEE Transactions on Computational Intelligence and AI in Games*. 2016;9(3):239–253.
- [19] Theune M, Rensen S, op den Akker R, Heylen D, Nijholt A. Emotional characters for automatic plot creation. In: *International Conference on Technologies for Interactive Digital Storytelling and Entertainment*. Springer; 2004. p. 95–100.
- [20] Walker MA. M2D: Monolog to Dialog Generation for Conversational Story Telling. In: *Interactive Storytelling: 9th International Conference on Interactive Digital Storytelling, ICIDS 2016, Los Angeles, CA, USA, November 15–18, 2016, Proceedings*. vol. 10045. Springer; 2016. p. 12.
- [21] Pérez y Pérez R. Employing emotions to drive plot generation in a computer-based storyteller. *Cognitive Systems Research*. 2007;8(2):89–109.
- [22] Schmidt BM. Plot archeology: A vector-space model of narrative structure. In: *2015 IEEE International Conference on Big Data (Big Data)*. IEEE; 2015. p. 1667–1672.
- [23] Mathewson KW, Castro PS, Cherry C, Foster G, Bellemare MG. Shaping the Narrative Arc: An Information-Theoretic Approach to Collaborative Dialogue. *arXiv preprint arXiv:190111528*. 2019.
- [24] Ouyang J, McKeown K. Modeling reportable events as turning points in narrative. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*; 2015. p. 2149–2158.
- [25] Hoang Q. Predicting movie genres based on plot summaries. *arXiv preprint arXiv:180104813*. 2018.
- [26] Felbo B, Mislove A, Søgaard A, Rahwan I, Lehmann S. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In: *Proceedings of the 2017*

- Conference on Empirical Methods in Natural Language Processing; 2017. p. 1615–1625.
- [27] Van der Maaten L, Hinton G. Visualizing data using t-SNE. *Journal of machine learning research*. 2008;9(11).