

Pay Attention to Virality: understanding popularity of social media videos with the attention mechanism

Adam Bielski
Tooploox, Poland

adam.bielski@tooploox.com

Tomasz Trzcinski
Warsaw University of Technology, Tooploox, Poland

tomasz.trzcinski@tooploox.com

Abstract

Predicting popularity of social media videos before they are published is a challenging task, mainly due to the complexity of content distribution network as well as the number of factors that play part in this process. As solving this task provides tremendous help for media content creators, many successful methods were proposed to solve this problem with machine learning. In this work, we change the viewpoint and postulate that it is not only the predicted popularity that matters, but also, maybe even more importantly, understanding of how individual parts influence the final popularity score. To that end, we propose to combine the Grad-CAM visualization method with a soft attention mechanism. Our preliminary results show that this approach allows for more intuitive interpretation of the content impact on video popularity, while achieving competitive results in terms of prediction accuracy.

1. Introduction

Multiple factors make popularity prediction of social media content a challenging task, including propagation patterns, social graph of users and interestingness of content. Current methods for online content popularity analysis focus mostly on determining its future popularity [6, 2, 8, 12]. For instance, [6] and [4] use visual cues to predict the popularity of online images. [13, 14, 11] use machine learning methods, such as Support Vector Regression and recurrent neural networks, applied to visual and textual cues to predict the popularity of social media videos. [2] combines cues from different modalities for the same purpose. Although popularity prediction is an important problem, we believe that it does not address all the challenges faced by online video creators. More precisely, it does not allow to answer the question of many creators: how a given frame or title word contributes to the popularity of the video? Even though correlation does not mean causality, this kind of analysis allows to understand the importance of a given piece of content and prioritize the creation efforts accordingly.

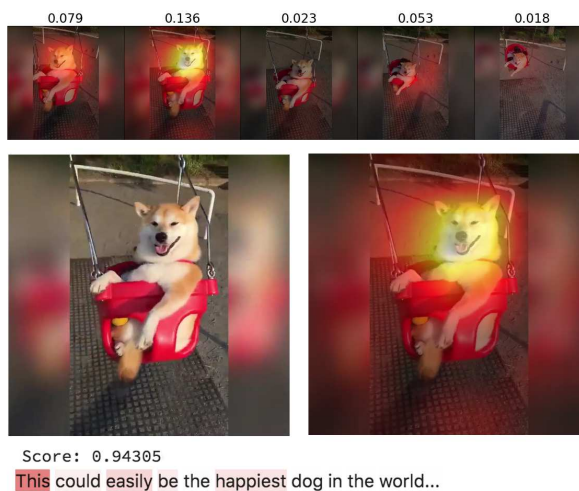


Figure 1. Sample social media video frames and its headline with visualized importance of their parts on predicted video popularity. The **top** row shows 5 consecutive frames of a social media video with their attention weights above. We use those weights to scale the magnitudes of Grad-CAM [9] heatmaps when visualizing the importance of video elements on the popularity score. The **bottom** row shows the frame with the highest attention weight (**left**) with its popularity importance visualization (**right**). For the headline text darker color corresponds to higher importance.

In this paper, we outline a fundamentally different approach to online video popularity analysis that allows social media creators both to predict video popularity as well as to understand the impact of its headline or video frames on the future popularity. To that end, we propose to use an attention-based model and gradient-weighted class activation maps [9], inspired by the recent successes of the attention mechanism in other domains [15, 16]. Although some works focused on understanding the influence of image parts on its popularity [6, 1], our method addresses videos, not images, and exploits the temporal characteristics of video clips through the attention mechanism. By extending the baseline popularity prediction method with the attention mechanism, we enable more intuitive visualization of the impact visual and textual features of the video have on its final popularity, while achieving state-of-the-art results on the popularity prediction task.

2. Popularity prediction with attention

We cast the problem of social media video popularity prediction as a binary classification task, as in [13, 11]. We focus on videos from Facebook and normalize their viewcount by the number of page followers. We assign a popular/unpopular label by splitting our dataset at the median normalized viewcount, following the approach of [10]. We use a cross-entropy loss function to classify a video as popular/unpopular and take a set of video frames and/or headline features as an input.

Video frames. We extract $N = 18$ evenly distributed frames from the first 6 seconds of a video¹. We use 2048-dimensional output of the penultimate layer of ResNet50 [5] pre-trained on ImageNet [3] to get a high-level frame representation as in [13]. For each frame feature vector we apply a learnable linear transformation followed by ReLU, obtaining a sequence of frame embeddings $(q_j)_{j=1}^N$. The final video embedding is a weighted average of these embeddings $v = \sum_{i=1}^N \alpha_i q_i$. Weights α_i are computed with attention mechanism implemented as a two-layer neural network [16]: the first layer produces a hidden representation $u_i = \tanh(W_u q_i + b_u)$ and the second layer outputs unnormalized importance $a_i = W_a u_i + b_a$. W_a can be interpreted as a trainable high level representation of *the most informative vector* in u_i space. Final weights are normalized with softmax: $\alpha_i = \exp(a_i) / \sum_k \exp(a_k)$.

Headline. We represent a headline as a sequence of pre-trained GloVe [7] word vectors $(w_t)_{t=1}^N$. We handle sequences of variable length using a bidirectional LSTM. Similarly to video frames, we use a two-layer attention mechanism on hidden state vectors h_t to let the network learn the importance coefficients β_t for each word. The final text representation d is a weighted average of hidden state vectors $d = \sum_{t=1}^N \beta_t h_t$.

Multimodal prediction. We concatenate previously trained video and text embeddings and train a two-layer neural network for popularity prediction. The intermediate layer output serves as multimodal embedding that captures information from both image and text modality that contribute to image popularity.

Visualizations. We visualize the importance of visual features using Grad-CAM [9]. More precisely, we generate heatmaps pointing to regions contributing to popularity in each frame. To this end, we compute gradients of the popular class score \hat{s} with respect to the output of the last convolutional layer of ResNet50 $A \in R^{K \times K \times F}$. Gradients are then used to compute weights $\gamma_f = \frac{1}{K^2} \sum_{i,j=1}^K \frac{\partial \hat{s}}{\partial A_{i,j}^f}$ that applied to the convolutional output create class activation map $H = \max(0, \sum_{f=1}^F \gamma_f A^f)$. We then normalize the heatmap values to $[0, 1]$ and use attention weights to

¹We use the first seconds of a video as this is how Facebook counts views, but we can extend our method to longer videos through sampling.

Input	Features	Acc [%]	Spearman
Video frames	ResNet50 mean	68.17	0.524
	+ attention	68.87	0.526
Headline	biLSTM [11]	69.47	0.542
	+ attention	68.70	0.525
Multimodal	ResNet + biLSTM	71.94	0.612
	+ attention	72.72	0.607

Table 1. Video popularity prediction results.

scale the heatmap by $\alpha_i / \max(\alpha)$. This way we obtain a sequence-wide normalized heatmap of frame regions influencing the final popularity score.

For visualizations in the text domain, we use attention weights β_t used to compute text representation d . These weights capture relative importance of words in their context to headline popularity, as shown in [16] in the context of sentiment analysis.

3. Experiments

We use a dataset of 37k Facebook videos with 80/10/10 train/validation/test splits. We use validation set to perform randomized search of hyperparameters such as embedding dimensionalities, dropout rates and batch normalization use. For training headline embeddings we use frozen pre-trained GloVe word vectors trained on Wikipedia and Gigaword [7]. As baselines, we use a simple mean of ResNet50 feature vectors as input to two layer neural network (video frames) and concatenation of last states of LSTM (headlines). We use Keras for implementation.

Results. For all methods, we follow the evaluation protocol of [13, 11] and compute the classification accuracy and Spearman correlation between the predicted probability of popular label and normalized view count of a video. Tab. 1 shows the results. Interestingly, almost equal popularity prediction results can be obtained using either video frames or headline features. Combining both modalities leads to noticeable improvement, while adding attention mechanism improves the performance in the multimodal and visual case. For headlines, the performance with attention deteriorates slightly. We speculate that the bi-directional LSTM already learns internal dependencies between hidden states and adding attention cannot help further, while for video frames it enables the network to exploit the temporal dependencies between the frames.

Overall, we see that prediction methods with attention achieve competitive results, while thanks to the attention mechanism we can increase the interpretability of our model. As Fig. 1 shows, extending standard Grad-CAM visualization with the attention mechanism allows us to interpret the influence of each video frame. Although the preliminary results presented in this paper leave place for improvement, they indicate the potential of using attention mechanism to increase the interpretability of popularity prediction methods for social media videos.

References

- [1] X. Alameda-Pineda, A. Pilzer, D. Xu, N. Sebe, and E. Ricci. Viraliency: Pooling local virality. In *CVPR*, pages 484–492. IEEE Computer Society, 2017.
- [2] J. Chen, X. Song, L. Nie, X. Wang, H. Zhang, and T.-S. Chua. Micro tells macro: Predicting the popularity of micro-videos via a transductive model. In *Proc. of the ACM on Multimedia Conference*, 2016.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [4] A. Deza and D. Parikh. Understanding image virality. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [5] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- [6] A. Khosla, A. D. Sarma, and R. Hamid. What makes an image popular? In *Proc. of International World Wide Web Conference (WWW)*, 2014.
- [7] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [8] H. Pinto, J. Almeida, and M. Gonçalves. Using early view patterns to predict the popularity of youtube videos. In *WSDM*, 2013.
- [9] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [10] E. Simo-Serra, S. Fidler, F. Moreno-Noguer, and R. Urta-sun. Neuroaesthetics in Fashion: Modeling the Perception of Fashionability. In *CVPR*, 2015.
- [11] W. Stokowiec, T. Trzcinski, K. Wolk, K. Marasek, and P. Rokita. Shallow reading with deep learning: Predicting popularity of online content using only its title. In *ISMIS*, 2017.
- [12] A. Tatar, M. D. de Amorim, S. Fdida, and P. Antoniadis. A survey on predicting the popularity of web content. *Journal of Internet Services and Applications*, 5, 2014.
- [13] T. Trzcinski, P. Andruszkiewicz, T. Bochenski, and P. Rokita. Recurrent neural networks for online video popularity prediction. In *ISMIS*, 2017.
- [14] T. Trzcinski and P. Rokita. Predicting popularity of online videos using support vector regression. *IEEE Transactions on Multimedia*, 19(11):2561–2570, 2017.
- [15] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In F. R. Bach and D. M. Blei, editors, *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2048–2057. JMLR.org, 2015.
- [16] Z. Yang, D. Yang, C. Dyer, X. He, A. J. Smola, and E. H. Hovy. Hierarchical attention networks for document classification. In *HLT-NAACL*, 2016.