

Can we predict the Most Replayed data of video streaming platforms?

Alessandro Duico¹, Ombretta Strafforello^{1,2}, Jan van Gemert¹

¹Delft University of Technology, ²TNO
{a.duico, o.strafforello, j.c.vangemert}@tudelft.nl

Abstract

Predicting which specific parts of a video users will re-play is important for several applications, including targeted advertisement placement on video platforms and assisting video creators. In this work, we explore whether it is possible to predict the Most Replayed (MR) data from YouTube videos. To this end, we curate a large video benchmark, the YTMR500 dataset, which comprises 500 YouTube videos with MR data annotations. We evaluate Deep Learning (DL) models of varying complexity on our dataset and perform an extensive ablation study. In addition, we conduct a user study to estimate the human performance on MR data prediction. Our results show that, although by a narrow margin, all the evaluated DL models outperform random predictions. Additionally, they exceed human-level accuracy. This suggests that predicting the MR data is a difficult task that can be enhanced through the assistance of DL. Finally, we believe that DL performance on MR data prediction can be further improved, for example, by using multi-modal learning. We encourage the research community to use our benchmark dataset to further investigate automatic MR data prediction.

1. Introduction

Video streaming has emerged as a dominant mode of online communication, representing 73% of all internet traffic in 2017 [7], with YouTube leading the way as the most popular platform. Video streaming platforms accumulate, in addition to video data, a substantial amount of metadata, pertaining to users' watching habits and interests. Notably, in May 2022 YouTube released a new feature that shows a line chart of the most frequently replayed moments of each video, the *Most Replayed* data. In addition to aiding YouTube users during video playback, this data can serve various other potential applications, such as optimizing advertisement placement and giving feedback to content creators – for instance, suggesting uninteresting scenes to remove. In both cases, it is desirable to predict the Most Replayed data

before publishing a video. For advertisers, it enables placing advertisements optimally from the very first views, thereby maximizing profits. For content creators, it allows cutting the video appropriately before it reaches the audience, preventing the reputational damage caused by re-uploading a video after collecting the data.

In this work, we investigate whether it is possible to predict the Most Replayed data using Deep Learning (DL). To this end, we collect YTMR500, a dataset of 500 vlog and travel videos with their corresponding Most Replayed data and pre-extracted video features. For a comprehensive description of our dataset collection process, readers are referred to the Supplementary Material (Section 1). We evaluate two DL models on YTMR500, consisting in a fully-connected network and an attention-based architecture, inspired by the PGL-SUM model that Apostolidis *et al.* [3] proposed for video summarization. We compare the results of DL against human performance, which we estimate through a crowdsourced user study.

We make the following contributions: (1) We introduce *YTMR500*, a dataset of 500 videos and the corresponding Most Replayed data, that can foster research on most replayed data predictions in videos; (2) We design a variant of the PGL-SUM [3] architecture, adapted to predict the Most Replayed data for unseen videos; (3) We perform a user study to evaluate human performance on MR data prediction. Our results show that predicting the Most Replayed data is challenging for human annotators and that our model surpasses human performance. The YTMR500 dataset and our code are publicly available¹.

2. Predicting the Most Replayed data in a video

2.1. Problem statement

Given a sequence of segments V in a video, we design a model to learn a function F that maps V into the Most Replayed (MR) data Y . Specifically, the input is a sequence

¹<https://github.com/ombretta/most-replayed-data>

of 1024-dimensional video features, \mathbf{v}_t , and the output is a sequence of scores $y \in [0, 1]$. Formally,

$$\mathbf{V} = \{\mathbf{v}_t\}_0^T, \mathbf{v}_t \in \mathbb{R}^{N=1024} \quad (1)$$

$$Y = \{y_i\}_0^I, I = 100, y_i \in [0, 1] \quad (2)$$

$$F : \mathbf{V} \rightarrow Y. \quad (3)$$

We want the model to determine the relative MR data of video segments when compared to one another, while we do not care about predicting the exact value of the Most Replayed data. Thus, instead of training the model using a Mean Squared Error loss, as in [3], we opt for a ranking loss, namely PyTorch’s `MarginRankingLoss`.

$$\mathcal{L}(\hat{y}_i, \hat{y}_j, s) = \max(0, -s_{i,j} \cdot (\hat{y}_i - \hat{y}_j) + \text{margin}) \quad (4)$$

To match this setup, we construct a ranking of video segments based on the ground-truth (GT) y_i and predicted \hat{y}_i MR data scores. The `MarginRankingLoss` forces the model to predict MR scores that result in a ranking as close as possible to the GT. Concretely, the loss is applied to the video segments in a pairwise fashion. In Equation 4, the targets $s_{i,j}$ must be one of $\{1, -1, 0\}$. Here, 1 indicates $y_i > y_j$, -1 indicates $y_i < y_j$, and 0 indicates $y_i = y_j$. To obtain these targets $s_{i,j}$, during training, we generate a comparison matrix of size $T \times T$ given by $S_{i,j} = \text{sgn}(y_i - y_j)$. Therefore, the ranking of each pair of predictions $\{\hat{y}_i, \hat{y}_j\}$ is encouraged to match the ranking of the GT $\{y_i, y_j\}$.

Interpolation While the Most Replayed data in the YTMR500 dataset have fixed length, the number of video features varies with the duration of the videos. To use a recurrent or attention-based model, the length of the inputs and that of the outputs must match. We overcome this issue in two different ways:

1. interpolating the ground truth to match the size of the frame features;
2. computing a binned average of the frame features, to match the size of the ground truth.

In Case 1, the GT Y is interpolated from its fixed size 100 into a variable size T . As a result, the interpolated GT \tilde{Y} has a different size for each of the videos, depending on the video duration. Considering that the number of frame features T is always greater than 100, this type of interpolation allows us to supply a larger amount of input data to the model, as compared to case 2. Since the number of comparisons in the matrix S grows quadratically with the size of the input, during training we randomly sample a subset of 10,000 comparisons, for each iteration and each video. Furthermore, the random sampling ensures that each video has the same contribution to the loss, independently of T .

In Case 2, we divide the frame features \mathbf{V} in 100 bins of uniform size and compute the average of the features within each bin. Therefore, the binned frame features $\tilde{\mathbf{V}}$ have a constant cardinality of 100, regardless of the video length. Using this strategy, the training becomes computationally easier, as there are fewer elements to be ranked.

2.2. DL models

Fully connected model As a baseline, we use a fully-connected model with two linear layers. We include a Dropout layer to mitigate overfitting on the training set. The final layer uses a Sigmoid activation function to coalesce the outputs into the $[0, 1]$ range, similarly to the GT. It is worth noting that this architecture does not capture any temporal relationships between different segments of the video.

Attention-based model We investigate whether considering the temporal relationships across segments improves the MR prediction accuracy beyond our baseline. We use an attention-based model, following the architecture of PGLSUM (Apostolidis *et al.* [3]) as closely as possible. PGLSUM is based on Multi-Head Attention (MAH) performed globally, *i.e.*, for the whole duration of the video, and locally, on a number of time windows -4 in our setup – obtained from a partitioning of the video. In the global attention module, we use 8 parallel attention heads, similarly to [3] and to the original Transformer [14]. In the local attention modules we use 4 heads. An illustration of the deployed models is provided in the Supplementary Material (Section 2).

2.3. User study

Since there exists no prior work on predicting the most replayed data in a video, there are no indications on the feasibility and difficulty of the task. We perform a user study to gain insights on how difficult predicting the most replayed data is for humans. Our user study contains 30 videos randomly selected from the test set. For the MR data prediction, we use a different setup than for the DL model. We do this because the task performed by the DL model, *i.e.*, constructing a ranking of 100 video segments, is too complex for human annotators. We simplify the task by subdividing each video in 10 shots rather than 100 and averaging the underlying 100 GTs into 10 bins.

We do not ask users to manually fill in the Most Replayed data, to prevent the influence of biases, *e.g.*, a bias towards continuity of the score across segments. Instead, we show a series of side-by-side comparisons of two video segments, to guide the users towards building a ranking. The user study is composed of an introduction about the purpose of the Most Replayed data; the full video, sped up to 30s for convenience; 19 pairwise comparisons, with the addition of an attention check. Each comparison presents participants with two video shots, sped up to 10s, along with the instructions: “*Guess*

which of the two video shots has greater ‘Most replayed’ score.” followed by the mutually exclusive options: “Left”, “Right”, and “CONTROL”. As part of the attention check, in one extra comparison, we place a video with a text overlay asking to choose the “CONTROL” option. Participants who fail this simple check are rejected.

The indices for the binary comparison were derived from the execution of the MergeSort algorithm [10] on 10 elements. The number of comparisons, 19, corresponds to the number of operations that the MergeSort algorithm needs to construct a total ordering. The indices of the segments involved in the comparisons are randomly permuted for each user, so that any imbalance with the MergeSort indices is not reflected in the outcome. Once we have obtained the answers to the pairwise comparisons, we construct a graph of the ordering and perform depth-first graph traversal to reconstruct a unique ranking of the segments.

The user study was crowdsourced to approximately 300 paid workers on Amazon MTurk [1]. Our crowdworkers population corresponds to the average demographics on Amazon MTurk, with a uniform distribution across genders, mainly comprising residents from the US and India, born between 1990 and 2000, as reported by mturk-tracker [5]. The number of users assigned to each video was 10 to 11, which is sufficient to obtain a statistically significant result, according to Carvalho *et al.* [4].

3. Results

Model training We use 5-fold cross validation, with a 80/20 training/test ratio. Therefore, of the 500 videos, 400 are utilized for training, and 100 for testing. To train our models we follow a similar procedure to Apostolidis *et al.* [3] and Fajtl *et al.* [6]. We use the Adam optimizer [9] with learning rate $lr = 5 \times 10^{-5}$ and L2 regularization $\lambda = 1 \times 10^{-5}$. Each batch contains only one sample, which is an entire video. This explains the low learning rate. For the MarginRankingLoss, we always use a margin of 0.01, except for when we are training on 10 video shots, in which case we use a margin of 0.05. We train for 300 epochs, because at that point the training set accuracy reaches a plateau. Even though other research commonly picks the best epoch with respect to the test set [3, 6], we refrain from doing so not to artificially boost our results. Therefore, when reporting the scores we average over the last 50 epochs (250 to 299) and all 5 splits.

Evaluation metrics To evaluate our model, a ranking correlation metric could be used, *e.g.*, Kendall’s τ [8]. However, this would penalize equally errors at the bottom and at the top of the ranking. Furthermore, we prioritize the global ordering rather than the exact position of each element. For instance, permutations of adjacent segments in the predicted ranking should not be heavily penalized.

Hence, we use precision@K, a metric inspired from information retrieval [13]. Precision@K measures how many of the top K results are true positives, divided by K. Given that we do not work with binary labels, we classify the top K video shots in the ground truth as positives and the rest as negatives. Using these labels, the metric corresponds to the proportion of the top K video shots of the predicted ranking that are among the top K video shots in the ground truth. We report precision@K for K in {15, 30, 50}, after interpolating the total number of video shots to 100, when required. The selection of $K = 15$ is inspired by the evaluation practices for video summarization in the literature [2] which typically adhere to a portion of 15% of the total duration. The values of $K = 30$ and $K = 50$ were chosen to assess precision at varying ranking depths. In the context of the user study, given that there are only 10 video shots, we use values of K in {1, 3, 5}. It is worth noting that precision@1 corresponds to top-1 accuracy, a metric commonly used in image classification [12].

3.1. DL models

All models are able to sufficiently fit the training set, obtaining precision@50 above 80% on the test set. However, the performance at test time is only marginally better than random. Results at test time are shown at the top of Table 2. It is surprising that our fully-connected baseline exhibits a satisfactory performance on this task, and the gains of the more complex PGL-SUM architecture are minimal.

We perform an ablation study on the full model, to understand the contribution of each component. At the bottom of Table 2 we report the scores for our model without local attention, without global attention and without the residual connection. We only display Case 1 of the interpolation, since the two cases are almost identical in performance. We discover that the removal of global or local attention does not heavily impact performance, contrary to the removal of the residual connection (shown in Supplementary Material, Figure 2). Therefore, we deduce that the crucial part of the learning is occurring on the input frame features, within the fully-connected layers at the end of the pipeline.

3.2. User study

To measure inter-rater agreement, we compute the Krippendorff’s α [11] among all the users, for each video. Since the average α is -0.017 ± 0.009 , we conclude that the users’ answers are generally not coherent with one another. In Table 1 we compare the precision@K of the users’ rankings and those generated by the DL model, in the simplified scenario with 10 video segments. The top section of the table shows the results computed on the 30 videos included in the user study, while the bottom section shows the results on the complete test set, averaged on 5 splits. Note that for the YTMR500 test set, the standard deviation is computed

Model	prec.@1 (%)	prec.@3 (%)	prec.@5 (%)
Random	10	30	50
User study test set (30 videos)			
Users (avg.)	9.6 ±25.9	31.8 ±20.4	51.3 ±15.3
PGL-SUM trained on 10	18.5 ± 15.0	37.0 ±10.4	53.0 ±6.9
PGL-SUM trained on 100	17.5 ±12.4	39.2 ± 12.6	58.5 ± 9.5
YTMR500 test set (5 splits ×100 videos)			
PGL-SUM trained on 10	14.3 ±2.9	35.0 ±1.5	54.7 ±1.5
PGL-SUM trained on 100	17.5 ± 2.7	38.1 ± 2.2	56.5 ± 1.8

Table 1. Results of our user study compared to our best DL model. Users are not significantly better than random. The DL models are always superior, and perform better when trained on more data, with one exception. N.b. precision@{1, 3, 5} are computed on rankings of 10 segments. The precision of the DL model is averaged over the last 50 epochs. Standard deviation is computed among the 30 videos for the user study test set and among the splits of a 5-fold cross validation for the YTMR500 test set.

between the results of the 5 splits only, not across all the videos in the test set.

We report scores for our attention-based model in two scenarios: firstly, when it is trained on the ground truth interpolated to 10 data points (“trained on 10” in Table 1), to closely match the task given to the users; secondly, when it is trained on the complete ground truth (“trained on 100” in Table 1), and averaged into 10 bins afterward. Naturally, training on more data, in the “trained on 100” case, yields better performance. Users are not able to perform significantly better than random on this task. Note that users do not undergo the same training as our DL models, which means they must base their predictions on their prior knowledge. From the results of our user study, we believe that predicting the Most Replayed data from video segments is a difficult task for humans.

3.3. Discussion

Based on our experiments in Section 3.1, we observe that using a more complex architecture does not induce a significant performance gain over the fully-connected baseline. Contrarily to our speculations, it seems that providing each segment with context about the full video does not improve the accuracy. We hypothesize that the pre-extracted video features alone provide sufficient abstraction to the fully-connected network to generate a satisfactory output from individual segments. As demonstrated by the ablation without a residual connection, hiding the input features from the fully-connected layers hinders performance. As shown

Model	prec.@15 (%)	prec.@30 (%)	prec.@50 (%)
Random	15	30	50
Fully-connected 1	21.5 ±1.8	37.6 ±0.9	56.3 ±0.4
Fully-connected 2	21.3 ±1.6	37.4 ±1.1	56.2 ±0.7
PGL-SUM 1	22.0 ± 1.6	37.9 ± 1.0	56.8 ± 0.6
PGL-SUM 2	22.0 ± 1.9	37.7 ± 0.8	56.4 ± 0.6
PGL-SUM 1 w/o local attention	20.6 ± 2.0	36.4 ± 1.0	55.6 ± 1.0
PGL-SUM 1 w/o global attention	21.5 ± 1.7	37.5 ± 1.1	56.6 ± 0.7
PGL-SUM 1 w/o residual	19.4 ± 1.2	35.6 ± 1.5	55.1 ± 0.9

Table 2. Results on the test set for our models, followed by some ablations. All DL models perform better than random. Surprisingly, the gains of the more complex PGL-SUM architecture are minimal, suggesting that attention between the segments is not fundamental. N.b. precision@{15, 30, 50} are computed on rankings of 100 segments. The numbers 1 and 2 refer to the type of interpolation: Case 1 is when the ground truth is interpolated and Case 2 is when the frame features are averaged.

by the users’ performance in our user study in Section 3.2, predicting the Most Replayed data is an arduous task. One plausible explanation is that the ground truth is noisy and lacks any clear patterns with respect to the input. Upon manual analysis of several videos, we found it hard to justify the location of certain peaks in the Most Replayed data. Another possible explanation is that video-only input does not provide enough information to resolve the problem effectively. Some peaks may be caused by interesting information in the speech, while others could result from dramatic changes in loudness. We defer to further research to incorporate multimodal inputs, particularly from audio channels and text transcripts.

4. Conclusion

The Most Replayed data presents a new source of insight on users’ interests in online video streaming. In this work, we focus on predicting this data using Deep Learning and compare the results against human performance, measured through a user study. For our experiments, we use YTMR500, a novel dataset comprising 500 vlog videos and their corresponding Most Replayed data. All the DL models evaluated on on YTMR500 perform significantly better than random, whereas the human participants are not able to accurately predict the Most Replayed data. We believe future research can further enhance the results, for instance, by incorporating multimodal inputs. We encourage the community to deploy our dataset in follow-up work.

Acknowledgements. This work is part of the research program Efficient Deep Learning (EDL), which is (partly) financed by the Dutch Research Council (NWO).

References

- [1] Amazon mechanical turk. <https://www.mturk.com/>. Accessed: 25 June 2023.
- [2] Evlampios Apostolidis, Eleni Adamantidou, Alexandros I Metsai, Vasileios Mezaris, and Ioannis Patras. Video summarization using deep neural networks: A survey. *Proceedings of the IEEE*, 109(11):1838–1863, 2021.
- [3] Evlampios Apostolidis, Georgios Balaouras, Vasileios Mezaris, and Ioannis Patras. Combining global and local attention with positional encoding for video summarization. In *2021 IEEE International Symposium on Multimedia (ISM)*, pages 226–234, December 2021.
- [4] Arthur Carvalho, Stanko Dimitrov, and Kate Larson. How many crowdsourced workers should a requester hire? *Annals of Mathematics and Artificial Intelligence*, 78(1):45–72, 2016.
- [5] Djelle Difallah, Elena Filatova, and Panos Ipeirotis. Demographics and dynamics of mechanical turk workers. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 135–143, 2018.
- [6] Jiri Fajtl, Hajar Sadeghi Sokeh, Vasileios Argyriou, Dorothy Monekoso, and Paolo Remagnino. Summarizing videos with attention. In Gustavo Carneiro and Shaodi You, editors, *Computer Vision - ACCV 2018 Workshops - 14th Asian Conference on Computer Vision, Perth, Australia, December 2-6, 2018, Revised Selected Papers*, volume 11367 of *Lecture Notes in Computer Science*, pages 39–54. Springer, 2018.
- [7] Cisco VNI Forecast. Cisco predicts more ip traffic in the next five years than in the history of the internet. <https://newsroom.cisco.com/press-release-content?articleId=1955935>, 2017-2022. Accessed: 22 August 2023.
- [8] M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1-2):81–93, 06 1938.
- [9] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [10] Donald Knuth. *Sorting and Searching*, volume 3 of *The Art of Computer Programming*. Addison-Wesley, 2nd edition, 1998.
- [11] Klaus Krippendorff. *Content analysis: An introduction to its methodology*. Sage Publications, 2004.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- [13] Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge, 2008.
- [14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.