

TrendLearner: Early Prediction of Popularity Trends of User Generated Content

Flavio Figueiredo^{1,*}, Jussara M. Almeida^a, Marcos A. Gonçalves^a, Fabricio Benevenuto^a

^aDepartment of Computer Science, Universidade Federal de Minas Gerais
Av. Antônio Carlos 6627, CEP 31270-010, Belo Horizonte - MG, Brazil. Phone: +55 (31) 3409-7541, Fax: +55 (31) 3409-5858

Abstract

Predicting the popularity of user generated content (UGC) is a valuable task to content providers, advertisers, as well as social media researchers. However, it is also a challenging task due to the plethora of factors that affect content popularity in social systems. Here, we focus on the problem of predicting the popularity *trend* of a piece of UGC (object) *as early as possible*. Unlike previous work, we explicitly address the inherent tradeoff between prediction accuracy and remaining interest in the object after prediction, since, to be useful, accurate predictions should be made *before* interest has exhausted. Given the heterogeneity in popularity dynamics across objects, this tradeoff has to be solved on a per-object basis, making the prediction task harder. We tackle this problem with a novel two-step learning approach in which we: (1) extract popularity trends from previously uploaded objects, and then (2) predict trends for newly uploaded content. Our results for YouTube datasets show that our classification effectiveness, captured by F1 scores, is 38% better than the baseline approaches. Moreover, we achieve these results with up to 68% of the views still remaining for 50% or 21% of the videos, depending on the dataset.

Keywords: popularity, trends, classification, social media, ugc, prediction

1. Introduction

The success of Internet applications based on user generated content (UGC)¹ has motivated questions such as: How does content popularity evolve over time? What is the potential popularity a piece of content will achieve after a given time period? How can we predict popularity evolution of a particular piece of UGC? For example, from a system perspective, accurate popularity predictions can be exploited to build more cost-effective content organization and delivery platforms (e.g., caching systems, CDNs). They can also drive the design of better analytic tools, a major segment nowadays [20, 34], while online advertisers may benefit from them to more effectively place contextual advertisements. From a social perspective, understanding issues related to popularity prediction can be used to better understand the human dynamics of consumption. Moreover, being able to predict popularity on an automated way is crucial for marketing campaigns (e.g. created by activists or politicians), which increasingly often use the Web to influence public opinion.

Challenges: However, predicting the popularity of a piece of content, here referred to as an *object*, in a social system is a very challenging task. This is mostly due to the various phenomena affecting the popularity prediction of social media – which were observed on the datasets we use (as well as others) [11, 22, 33] – as well as the diminishing interesting in objects over time, which implies that popularity predictions must

be timely to capture user interest and be useful in real work settings. Both challenges can be summarized as follows:

1. Due to the easiness with which UGC can be created, many factors can affect an object's popularity. Such factors include, for instance, the object's content, the social context in which it is inserted (e.g., social neighborhood or influence zone of the object's creator), the mechanisms used to access the content (e.g., searching, recommendation, top-lists), or even an external factor, such as a hyperlink to the content in a popular blog or website. These factors can cause spikes in the surge of interest in objects, as well as information propagation cascades which affect the popularity trends of objects.
2. To be useful in a real scenario, a popularity prediction approach must identify popularity trends *before the user interest in the object has severely diminished*. To illustrate this point, Figure 1 shows the popularity evolution of two YouTube videos: the video on the left receives more than 80% (shaded region) of all views received during its lifespan in the first 300 days since upload, whereas the other video receives only about half of its total views in the same time frame. If we were to monitor each video for 300 days, most potential views of the first video would be lost. In other words, not all objects require the same monitoring period, as assumed by previous work, to produce accurate predictions: for some objects, the prediction can be made earlier. Thus, the tradeoff should be solved on a *per-object* basis, which implies that determining the duration of the monitoring period that leads to a good solution of the tradeoff for *each object* is part of the problem.

*Corresponding author

Email addresses: flaviouv@dcc.ufmg.br (Flavio Figueiredo), jussara@dcc.ufmg.br (Jussara M. Almeida), mgoncalv@dcc.ufmg.br (Marcos A. Gonçalves), fabricio@dcc.ufmg.br (Fabricio Benevenuto)

¹YouTube, Flickr, Twitter, and so forth

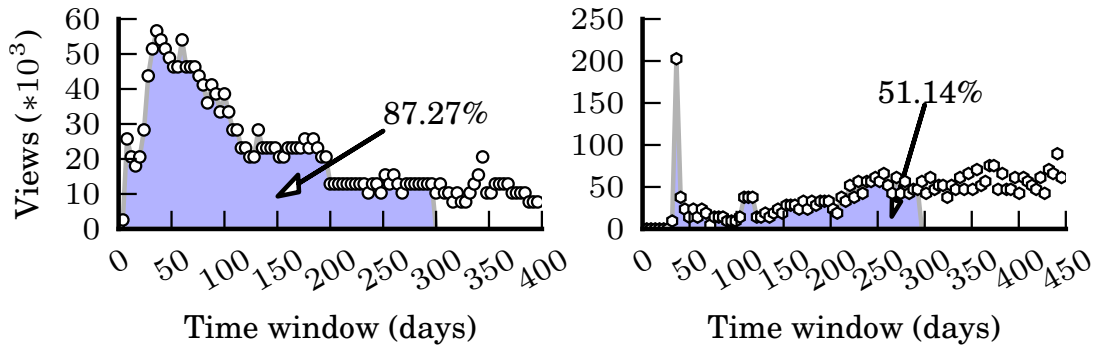


Figure 1: Popularity Evolution of Two YouTube Videos.

These challenges set UGC objects apart from more traditional web content. For instance, news media [3] tends to have clear definitions of monitoring periods, say predicting the popularity of news after one day using information from the first hour after upload. This is mostly due to the timely nature of the content, which is reflected in the popularity trends usually followed by news media [11] – interest is usually concentrated in a peak window (e.g., day) and dies out rather quickly. Thus, mindful of the challenges above, we here tackle the problem of UGC popularity *trend* prediction. That is, we focus on the (hard) task of predicting popularity *trends*. Trend prediction can help determining, for example, if an object will follow a viral pattern (e.g., Internet memes) or will continue to gain attention over time (e.g., music videos for popular artists). Moreover, we shall also show that, by knowing popularity trends *beforehand*, we can improve the accuracy of models for predicting popularity measures (e.g., hits). Thus, by focusing on predicting trends, we fill a gap in current research since no previous efforts has effectively predicted the popularity *trend* of UGC taking into account challenges (1) and (2).

We should stress that one key aspect distinguishes our work from previous efforts to predict popularity [1, 3, 19, 25, 27, 32] – we explicitly address the inherent tradeoff between prediction accuracy and how early the prediction is made, assessed in terms of the remaining interest in the content after prediction. All previous popularity prediction efforts considered fixed monitoring periods for all objects, which is given as input. We refer to this problem as *early prediction*².

In terms of applications, knowing that an object will be popular early on can help advertisers to plan out specific revenue models [13]. Such knowledge can also help out on geographic content sharding [8] for better content delivery. On the other hand, being aware that an object will not be popular at all, as early as possible, allow low access content to be tiered down to lower latency servers/geographic regions, whereas advertisers can use this knowledge to avoid bidding for ads in such content (since they will not generate revenue). Another example are search engines rankings based on predictions [26]. Knowing

that a content is becoming popular can help out in generating better rankings to user queries. However, if we have evidence (based on the trend and remaining interest) that such content is losing popularity (e.g., timely content that users can lose interest over time), such contents may be of less interest to the user. Finally, early prediction is of utmost importance to content producers – knowing whether a piece of content will be follow a certain trend can help in their promotion strategies and in the creation of new content.

TrendLearner: We tackle this problem with a novel two-step combined learning approach. First, we identified popularity trends, expressed by popularity timeseries, from previously uploaded objects. Then, we combine novel time series classification algorithms with object features for predicting the trends of new objects. This approach is motivated by the intuition that it might be easier to identify the popularity trend of an object if one has a set of possible trends as basis for comparison. More important, we propose a new trend classification approach, namely TrendLearner, that tackles the aforementioned tradeoff between prediction accuracy and remaining interest after prediction on a per-object basis. The idea here is to monitor newly uploaded content on an online basis to determine, for each monitored object, the earliest point in time when prediction confidence is deemed to be good enough (defined by input parameters), producing, as output, the probabilities of each object belonging to each class (trend). Moreover, unlike previous work, TrendLearner also combines the results from this classifier (i.e., the probabilities) with a set of object related features [11], such as category and incoming links, building an ensemble learner.

To evaluate our method, we use, in addition to traditional classification metrics (e.g., Micro/Macro F1), two newly defined metrics, specific for the problem: (1) remaining interest (RI), defined as the fraction of all views (up to a certain date) that remain after the prediction, and (2) the correlation between the total views and the remaining interest. While the first metric measures the potential future viewership of the objects, the second one estimates whether there is any bias towards more/less popular objects.

In sum, our main contributions include a novel popularity trend classification method that considers multiple trends, called TrendLearner. The use of TrendLearner can improve the

²We also point out that an earlier, much simpler, variant of our approach, which did not focus on early predictions, was first place on two out of three prediction tasks of the 2014 ECML/PKDD Predictive Analytics Challenge for News Content [10]³, reflecting the quality/effectiveness of our proposal.

prediction of popularity metrics (e.g., number of views). Improvements over state-of-the-art method are significant, being around 33%, at least.

The rest of this article is organized as follows. Next section discusses related work. We state our target problem in Section 3, and present our approach to solve it in Section 4. We introduce the metrics and datasets used to evaluate our approach in Section 5. Our main experimental results are discussed in Section 6. Section 7 offers conclusions and directions for future work.

2. Related Work

Popularity evolution of online content has been the target of several studies. Several previous efforts aimed at developing models to predict the popularity of a piece of content at a given future date. In [19], the authors developed stochastic user behavior models to predict the popularity of Digg’s stories based on early user reactions to new content and aspects of the website design. Such models are very specific to Digg features, and are not general enough for different kinds of UGC. Szabo and Huberman proposed a linear regression method to predict the popularity of YouTube and Digg content from early measures of user accesses [27]. This method has been recently extended and improved with the use of multiple features [25]. Castillo *et al.* [3] used a similar approach as [27] to predict the popularity of news content.

Out of these previous efforts, most authors focused on variations of Linear Regression based methods to predict UGC popularity [3, 18, 25]. In the context of search engines, Radinsky *et al.* proposed Holt-Winters linear models to predict future popularity, seasonality and the bursty behavior of queries [26]. The models capture the behavior of a population of users searching on the Web for a specific query, and are trained for each individual time series. We note that none of these prior efforts focused on the problem of *predicting popularity trends*. In particular, those focused on UGC popularity prediction assumed a fixed monitoring period for all objects, given as input, and did not explore the trade-off between prediction accuracy and remaining views after prediction.

Other methods exploit epidemic modeling of UGC popularity evolution. Focusing on content propagation within an OSN, Li *et al.* addressed video popularity prediction within a single (external) OSN (e.g., Facebook) [21]. Similarly, Matsubara *et al.* [22] created a unifying epidemic model for the trends usually found in UGC. Such a model can be used for tail forecasting, that is, predictions after the peak time window. Again, none of these methods focus neither on trend predictions or on early predictions as we do. Also, tail-part forecasting is very limited when the popularity of an object may exhibit multiple peaks [15, 33]. By focusing on a two step trend identification and prediction approach, combined with a non-parametric distance function, TrendLearn can overcome these challenges.

Chen *et al.* [5] propose to predict whether a tweet will become a trending topic by applying a binary classification model (trending versus non-trending), learned from a set of objects from each class. We here propose a more general approach to

detect *multiple* trends (classes), where trends are first *automatically* learned from a training set. It is also important to note that our solution complements the one by Jiang *et al.* [17], which focused on predicting when a video will peak in popularity. Finally, our solution also exploits the concept of shapelets [31] to reduce the classification time complexity, as we show in Section 4.

We also mention some other efforts to detect trending topics in various domains. Vakali *et al.* proposed a cloud-based framework for detecting trending topics on Twitter and blogging systems [28], focusing particularly on implementing the framework on the cloud, which is complementary to our goal. Golbandi *et al.* [14] tackled trend topic detection for search engines. Despite the similar goal, their solution applies to a very different domain, and thus focuses on different elements (query terms) and uses different techniques (language models) for prediction.

Table 1 summarizes the key functionalities of the aforementioned approaches as well as of our new TrendLearner method. In sum, to our knowledge, we are the first to tackle the inherent challenges of predicting UGC popularity (trends and metrics) as early and accurately as possible, on a per-object basis, recognizing that different objects may require different monitoring periods for accurate predictions. More important, the challenges we approach with TrendLearner (i.e. predicting trends also tackling the tradeoff between prediction accuracy and remaining interest after prediction on a per-object basis) are key to leverage popularity prediction towards practical scenarios and deployment in real systems.

3. Problem Statement

The early popularity trend prediction problem can be defined as follows. Given a training set of previously monitored user generated objects (e.g., YouTube videos or tweets), D^{train} , and a test set of newly uploaded objects D^{test} , do: (1) extract popularity trends from D^{train} ; and (2) predict a trend for each object in D^{test} as early and accurately as possible, particularly before user interest in such content has significantly decayed. User interest can be expressed as the fraction of all potential views a new content will receive until a given point in time (e.g., the day when the object was collected). Thus, by predicting as early as possible the popularity trend of an object, we aim at maximizing the fraction of views that still remain to be received *after prediction*. Determining the earliest point in time when prediction can be made with reasonable accuracy is an inherent challenge of the early popularity prediction problem, given that it must be addressed on a per-object basis. That is, while later predictions can be more accurate, they would imply a reduction of the remaining interest in the content.

In particular, we here treat the above problem as a trend-extraction one combined with as a multi-class classification task. The popularity trends automatically extracted from D^{train} (step 1) represent the classes into which objects in D^{test} should be grouped (step 2). Trend extraction is performed using a time series clustering algorithm [30], whereas prediction is a classi-

Table 1: Comparison of TrendLearner with other approaches

	Trend Identification	Trend Prediction	Views Prediction	Early Prediction
Trending Topics Prediction [5, 28]		✓ (Binary only)		
Linear Regression [25] [27] [3]			✓	
Holt-Winters [26]			✓	
Epidemic Models [21, 22]			✓	
TrendLearner	✓	✓	✓	✓

Table 2: Notation. Vectors (\mathbf{x}) and matrices (\mathbf{X}), in bold, are differentiated by lower and upper cases. Streams ($\hat{\mathbf{x}}$) are differentiated by the hat accent ($\hat{\cdot}$). Sets (D) and variables (d) are shown in regular upper and lower case letters, respectively.

Symbol	Meaning	Example
D	dataset of UGC content	YouTube videos
D^{train}	training set	-
D^{test}	testing set	-
d	a piece of content or object	video
D_i	class/trend i	-
\mathbf{c}_{D_i}	centroid of class i	-
\mathbf{s}_d	time series vector for object d	$\mathbf{s}_d = \langle p_{d,1}, \dots, p_{d,n} \rangle$
$\hat{\mathbf{s}}_d$	time series stream for object d	$\hat{\mathbf{s}}_d = \langle p_{d,1}, \dots \rangle$
$p_{d,i}$	popularity of d at i -th window	number of views
$\mathbf{s}_d^{[i]}$	index operator	$\langle 7, 8, 9 \rangle^{[2]} = 8$
$\mathbf{s}_d^{[i:j]}$	slicing operator	$\langle 7, 8, 9 \rangle^{[2:3]} = \langle 8, 9 \rangle$
\mathbf{S}	matrix with set of time series	all time series

fication task. For the sake of clarity, we shall make use of the term “class” to refer to both clusters and classes.

Table 2 summarizes the notation used throughout the paper. Each object $d \in D^{train}$ is represented by an n -dimensional time series vector $\mathbf{s}_d = \langle p_{d,1}, p_{d,2}, \dots, p_{d,n} \rangle$, where $p_{d,i}$ is the popularity (i.e., number of views) acquired by d during the i^{th} time window after its upload. Intuitively, the duration of a time window w could be a few hours, days, weeks, or even months. Thus, vector \mathbf{s}_d represents a time series of the popularity of a piece of content measured at time intervals of duration w (fixed for each vector). New objects in D^{test} are represented by streams, $\hat{\mathbf{s}}_d$, of potentially infinite length ($\hat{\mathbf{s}}_d = \langle p_{d,1}, p_{d,2}, \dots \rangle$). This captures the fact that our trend prediction/classification method is based on monitoring each test object on an online basis, determining when a prediction with acceptable confidence can be made (see Section 4.2). Note that a vector can be seen as a contiguous subsequence of a stream. Note also that the

complete dataset is referred to as $D = D^{train} \cup D^{test}$.

4. Our Approach

We here present our solution to the early popularity trend prediction problem. We introduce our trend extraction approach (Section 4.1), present our novel trend classification method, TrendLearner (Section 4.2), and discuss practical issues related to the joint use of both techniques (Section 4.3).

4.1. Trend Extraction

To extract temporal patterns of popularity evolution (or trends) from objects in D^{train} , we employ a time series clustering algorithm called K-Spectral Clustering (KSC) [30]⁴, which

⁴We have implemented a parallel version of the KSC algorithm which is available at <http://github.com/flaviovdf/pyksc>. The repository also contains the TrendLearner code

groups time series based on the *shape* of the curve. To group the time series, KSC defines the following distance metric to capture the similarity between two time series \mathbf{s}_d and $\mathbf{s}_{d'}$ with scale and shifting invariants:

$$\text{dist}(\mathbf{s}_d, \mathbf{s}_{d'}) = \min_{\alpha, q} \frac{\|\mathbf{s}_d - \alpha \mathbf{s}_{d'(q)}\|}{\|\mathbf{s}_d\|}, \quad (1)$$

where $\mathbf{s}_{d'(q)}$ is the operation of shifting the time series $\mathbf{s}_{d'}$ by q units and $\|\cdot\|$ is the l_2 norm⁵. For a fixed q , there exists an exact solution for α by computing the minimum of dist , which is: $\alpha = \frac{\mathbf{s}_d^T \mathbf{s}_{d'(q)}}{\|\mathbf{s}_{d'}\|^2}$. In contrast, there is no simple way to compute shifting parameter q . Thus, in our implementation of KSC, whenever we measure the distance between two series, we search for the optimal value of q considering all integers in the range $(-n, n)$ ⁶.

Having defined a distance metric, KSC is mostly a direct translation of the K-Means algorithm [6]. Given a number of trends k to extract and the set of time series, it works as:

1. The time series are uniformly distributed to k random classes;
2. Cluster centroids are computed based on its members. In K-Means based algorithms, the goal is to find centroid \mathbf{c}_{D_i} such that $\mathbf{c}_{D_i} = \arg \min_{\mathbf{c}} \sum_{\mathbf{s}_d \in D_i} \text{dist}(\mathbf{s}_d, \mathbf{c})^2$. We refer the reader to the original KSC paper for more details on how to find \mathbf{c}_{D_i} [30];
3. For each time series vector \mathbf{s}_d , object d is assigned to the nearest centroid based on metric dist ;
4. Return to step 2 until convergence, i.e., until all objects remain within the same class in step 3. Each centroid defines the trend that objects in the class (mostly) follow.

Before introducing our trend classification method, we make the following observation that is key to support the design of the proposed approach: each trend, as defined by a centroid, is conceptually equivalent to the notion of *time series shapelets* [31]. A shapelet is informally defined as a time series subsequence that is in a sense maximally representative of a class. As argued in [31], the distance to the shapelet can be used to classify objects with more accuracy and much faster than state-of-the-art classifiers. Thus, by showing that a centroid is a shapelet, we choose to classify a new object based only on the distances between the object's popularity time series up to a monitored time and each trend.

This is one of the points where our approach differs from the method proposed in [5], which uses the complete D^{train} as reference series, classifying an object based on the distances between its time series and *all* elements of each class. Given $|D^{\text{train}}|$ objects in the training set and k trends (with $k \ll |D^{\text{train}}|$), our approach is faster by a factor of $\frac{|D^{\text{train}}|}{k}$.

Definition: For a given class D_i , a shapelet \mathbf{c}_{D_i} is a time series subsequence such that: (1) $\text{dist}(\mathbf{c}_{D_i}, \mathbf{s}_d) \leq \beta, \forall \mathbf{s}_d \in D_i$; and (2) $\text{dist}(\mathbf{c}_{D_i}, \mathbf{s}_{d'}) > \beta, \forall \mathbf{s}_{d'} \notin D_i$, where β is defined as an optimal

distance for a given class. With this definition, a shapelet can be shown to maximize the information gain of a given class [31], being thus the most representative time series of that class.

We argue that, by construction, a centroid produced by KSC is a shapelet with β being the distance from the centroid to the time series within the class that is furthest away from its centroid. Otherwise, the time series that is furthest away would belong to a different class, which contradicts the KSC algorithm. This is an intuitive observation. Note that a centroid is a shapelet *only* when using K-Means based approaches, such as KSC, to define class labels. In the case of learning from already labeled data a shapelet finding algorithms [31] should be employed.

4.2. Trend Prediction

Let D_i represent class i , previously learned from D^{train} . Our task now is to create a classifier that correctly determines the class of a new object as early as possible. We do so by monitoring the popularity acquired by each object d ($d \in D^{\text{test}}$) since its upload on successive time windows. As soon as we can state that d belongs to a class with *acceptable confidence*, we stop monitoring it and report the prediction. The heart of this approach is in detecting *when* such statement can be made.

4.2.1. Probability of an Object Belonging to a Class

Given a monitoring period defined by t_r time windows, our trend prediction is fundamentally based on the distances between the subsequence of the stream $\hat{\mathbf{s}}_d$ representing d 's popularity curve from its upload until t_r , $\hat{\mathbf{s}}_d^{[1:t_r]}$, and the centroid of each class. To respect shifting invariants, we consider all possible starting windows t_s in each centroid time series when computing distances. That is, given a centroid \mathbf{c}_{D_i} , we consider all values from 1 to $|\mathbf{c}_{D_i}| - t_r$, where $|\mathbf{c}_{D_i}|$ is the number of time windows in \mathbf{c}_{D_i} . Specifically, the probability that a new object d belongs to class D_i , given D_i 's centroid, the monitoring period t_r and a starting window t_s , is:

$$p(\hat{\mathbf{s}}_d \in D_i \mid \mathbf{c}_{D_i}; t_r, t_s) \propto \exp(-\text{dist}(\hat{\mathbf{s}}_d^{[1:t_r]}, \mathbf{c}_{D_i}^{[t_s:t_s+t_r-1]})) \quad (2)$$

where $[x:y]$ ($x \leq y$) is a moving window slicing operator (see Table 2). As in [5, 6, 25], we assume that probabilities are inversely proportional to the exponential function of the distance between both series, given by function dist (Equation 1), normalizing them afterwards to fall in the 0 to 1 range (here omitted for simplicity). Figure 2 shows an illustrative example of how both time series would be aligned for probability computation⁷. That is, for time series of different lengths, we slice a consecutive range of the largest time series so that it has the size of the smallest one. Every slice possible is considered (starting from 1 to $|\hat{\mathbf{s}}_d^{[1:t_r]}|$) and we keep the slice with the smallest distance when computing probabilities.

With Equation 2, we could build a classifier that simply picks the class with highest probability. But this would require t_s and

⁵The l_2 norm of a vector \mathbf{x} is defined as $\|\mathbf{x}\| = \sqrt{\sum_{i=1}^n x_i^2}$.

⁶Shifts are performed in a rolling manner, where elements at the end of the vector return to the beginning. This maintains the symmetric nature of $\text{dist}(\mathbf{s}_d, \mathbf{s}_{d'})$.

⁷In case $|\mathbf{c}_{D_i}| < |\hat{\mathbf{s}}_d^{[1:t_r]}|$, we try all possible alignments of \mathbf{c}_{D_i} with $\hat{\mathbf{s}}_d^{[1:t_r]}$.




Figure 2: Example of alignment of time series (dashed lines) for probability computation.

t_r to be fixed. As shown in Figure 1, different time series may need different monitoring periods (different values of t_s and t_r), depending on the required confidence.

Instead, our approach is to monitor an object for successive time windows (increasing t_r), computing the probability of it belonging to each class at the end of each window. We stop when the class with maximum probability exceeds a class-specific threshold, representing the required minimum confidence on predictions for that class. We detail our approach next, focusing first on a single class (Algorithm 1), and then generalizing it to multiple classes (Algorithm 2).

Algorithm 1 shows how we define when to stop computing the probability for a given class D_i . The algorithm takes as input the object stream \hat{s}_d , the class centroid \mathbf{c}_{D_i} , the minimum confidence θ_i required to state that a new object belongs to D_i , as well as γ_i and γ^{max} , the minimum and maximum thresholds for the monitoring period. The former is used to avoid computing distances with too few windows, which may lead to very high (but unrealistic) probabilities. The latter is used to guarantee that the algorithm ends. We allow different values of γ_i and θ_i for each class as different popularity trends have overall different dynamics, requiring different thresholds⁸. The algorithm outputs the number of monitored windows t_r and the estimated probability p . The loop in line 4 updates the stream with new observations (increases t_r), and function *AlignComputeProb* computes the probability for a given t_r by trying all possible alignments (i.e., all possible values of t_s). For a fixed alignment (i.e., fixed t_r and t_s), *AlignComputeProb* computes the distance between both time series (line 15) and the probability of \hat{s}_d belonging to D_i (line 16). It returns the largest probability representing the best alignment between \hat{s}_d and \mathbf{c}_{D_i} , for the given t_r (lines 17 and 20). Both loops that iterate over t_r (line 4) and t_s (line 15) stop when the probability exceeds the minimum confidence θ_i . The algorithm also stops when the monitoring period t_r exceeds γ^{max} (line 7), returning a probability equal to 0 to indicate that it was not possible to state the \hat{s}_d belongs to D_i within the maximum monitoring period allowed (γ^{max}).

We now extend Algorithm 1 to compute probabilities and monitoring periods for all object streams in D^{test} , considering all classes extracted from D^{train} . Algorithm 2 takes as input the test set D^{test} , a matrix \mathbf{C}_D with the class centroids, vectors θ and γ with per-class parameters, and γ^{max} . It outputs a vector \mathbf{t} with the required monitoring period for each object, and a matrix \mathbf{P} with the probability estimates for each object (row) and class (column), both initialized with 0 in all elements. Given a valid monitoring period t_r (line 6), the algorithm monitors each

object d in D^{test} (line 7) by first computing the probability of d belonging to each class (line 9). It then takes, for each object d , the largest of the computed probabilities (line 11) and the associated class (line 12), and tests whether it is possible to state that d belongs to that class with enough confidence at t_r , i.e., whether: (1) the probability exceeds the minimum confidence for the class, *and* (2) t_r exceeds the per-class minimum threshold (line 13). If the test succeeds, the algorithm stops monitoring the object (line 16), saving the current t_r and the per-class probabilities computed at this window in \mathbf{t} and \mathbf{P} (lines 14-15). After exhausting all possible monitoring periods ($t_r > \gamma^{max}$) or whenever the number of objects being monitored n_{objs} reaches 0, the algorithm returns. At this point, entries with 0 in \mathbf{P} indicate objects for which no prediction was possible within the maximum monitoring period allowed (γ^{max}).

Having \mathbf{P} , a simple classifier can be built by choosing for each object (row) the class (column) with maximum probability. The value in \mathbf{t} determines how early this classification can be done. However, we here employ a different strategy, using matrix \mathbf{P} as input features to another classifier, as discussed below. We compare our proposed approach against the aforementioned simpler strategy in Section 6.

4.2.2. Probabilities as Input Features to a Classifier

Instead of directly extracting classes from \mathbf{P} , we choose to use this matrix as input features to another classification algorithm, motivated by previous results on the effectiveness of using distances as features to learning methods [6]. Specifically, we employ an extremely randomized trees classifier [12], as it has been shown to be effective on different datasets [12], requiring little or no pre-processing, besides producing models that can be more easily interpreted, compared to other techniques like Support Vector Machines⁹. Extremely randomized trees tackle the over fitting problem of more common decision tree algorithms by training a large ensemble of trees. They work as follows: 1) for each node in a tree, the algorithm selects the best features for splitting based on a random subset of all features; 2) split values are chosen at random. The decision of these trees are then averaged out to perform the final classification. Although feature search and split values are based on randomization, tree nodes are still chosen based on the maximization of some measure of discriminative power such as Information Gain, with the goal of improving classification effectiveness.

We extend the set of probability features taken from \mathbf{P} with other features associated with the objects. The set of object features used depends on the type of UGC under study and characteristics of the datasets (D). We here use the features shown in Table 3, which are further discussed in Section 5.2, combining them with the probabilities in \mathbf{P} . We refer to this approach as *TrendLerner*.

Before continuing, we briefly discuss other strategies to combine classifiers as we have done. We experimented with these methods, finding them to be unsuitable to our dataset due to various reasons. For instance, we implemented Co-Training [24], a

⁸Indeed, initial experiments showed that using the same values of γ_i (and θ_i) for all classes produces worse results.

⁹We also used SVM learners, achieving similar results.

Algorithm 1 Define when to stop computing probability of object \hat{s}_d belonging to class D_i , based on minimum confidence θ_i , and minimum and maximum monitoring periods γ_i and γ^{max} .

```

1: function PERCLASSPROB( $\hat{s}_d, \mathbf{c}_{D_i}, \theta_i, \gamma_i, \gamma^{max}$ )
2:    $p \leftarrow 0$ 
3:    $t_r \leftarrow \gamma_i - 1$  ▷ Start at previous window
4:   while  $p < \theta_i$  do ▷ Extend monitoring period
5:      $t_r \leftarrow t_r + 1$  ▷ Move to next current window
6:     if  $t_r > \gamma^{max}$  then ▷ Monitoring period ended
7:       return  $\gamma^{max}, 0$ 
8:     end if
9:      $p \leftarrow \text{AlignComputeProb}(\hat{s}_d, \mathbf{c}_{D_i}, \theta_i, t_r)$ 
10:  end while
11:  return  $t_r, p$  ▷ Return monitoring period and probability
12: end function
13: function ALIGNCOMPUTEPROB( $\hat{s}_d, \mathbf{c}_{D_i}, \theta_i, t_r$ )
14:    $t_s \leftarrow 1; p \leftarrow 0$ 
15:   while  $(t_s \leq |\mathbf{c}_{D_i}| - t_r)$  and  $(p < \theta_i)$  do ▷ Iterate over possible values of  $t_s$ , aligning both series
16:      $p' \propto \exp(-\text{dist}(\hat{s}_d^{[1:t_r]}, \mathbf{c}_{D_i}^{[t_s:t_s+t_r-1]}))$ 
17:      $p \leftarrow \max(p, p')$ 
18:      $t_s \leftarrow t_s + 1$ 
19:   end while
20:   return  $p$ 
21: end function

```

Algorithm 2 Define when to stop computing probabilities for each object in D^{test} , considering the centroids of all classes (\mathbf{C}_D), per-class minimum confidence (θ) and monitoring period (γ), and maximum monitoring period (γ^{max}).

```

1: function MULTICLASSPROBS( $D^{test}, \mathbf{C}_D, \theta, \gamma, \gamma^{max}$ )
2:    $\mathbf{t} = [0]$  ▷ Per-object monitoring period vector
3:    $\mathbf{P} = [[0]]$  ▷ Per-object, per-class probability matrix
4:    $n_{objs} \leftarrow |D^{test}|$  ▷ Number of objects to be monitored
5:    $t_r \leftarrow \min(\gamma)$  ▷ Init  $t_r$  with minimum  $\gamma_i$ 
6:   while  $(t_r \leq \gamma^{max})$  and  $(n_{objs} > 0)$  do
7:     for all  $\hat{s}_d \in D^{test}$  do ▷ Predict class for each object
8:       for all  $\mathbf{c}_{D_i} \in \mathbf{C}_D$  do ▷ Get centroid of each class
9:          $p^{[i]} \leftarrow \text{AlignComputeProb}(\hat{s}_d, \mathbf{c}_{D_i}, \theta_i, t_r)$ 
10:      end for
11:       $\max p \leftarrow \max(\mathbf{p})$  ▷ Get max. probability and corresponding class for  $t_r$ 
12:       $\max c \leftarrow \text{argmax}(\mathbf{p})$ 
13:      if  $(\max p > \theta^{[\max c]})$  and  $(t_r \geq \gamma^{[\max c]})$  then ▷ Stop if maxp and  $t_r$  exceeds per-class thresholds
14:         $\mathbf{t}^{[d]} \leftarrow t_r$  ▷ Save current  $t_r$ 
15:         $\mathbf{P}^{[d]} \leftarrow \mathbf{p}$  ▷ Save current  $\mathbf{p}$  in row  $d$ 
16:         $n_{objs} \leftarrow n_{objs} - 1$ 
17:         $D^{test} \leftarrow D^{test} - \{\hat{s}_d\}$ 
18:      end if
19:    end for
20:     $t_r \leftarrow t_r + 1$ 
21:  end while
22:  return  $\mathbf{t}, \mathbf{P}$  ▷ Return monitoring periods and probabilities
23: end function

```

traditional semi-supervised label propagation approach. However, it failed to achieve better results than just combining the features, most likely because it depends on feature independence, which may not hold in our case. We also experimented with Stacking [9], which yielded similar results as the proposed approach. Nevertheless, either strategy might be more effective on different datasets or types of UGC, an analysis that we leave for future work.

4.3. Putting It All Together

A key point that remains to be discussed is how to define the input parameters of the trend extraction approach, that is, the number of trends k , as well as the parameters of TrendLearner, namely vectors θ and γ, γ^{max} , and the parameters of the adopted classifier.

We choose the number of trends k based primarily on the β_{CV} quality metric [23]. Let the intraclass distance be the dis-

tance between a time series and its centroid (the trend), and the interclass distance be the distance between different trends. The general purpose of the trend extraction is to minimize the variance of the intraclass distances while maximizing the variance of the interclass distances. The β_{CV} is defined as the ratio of the coefficient of variation¹⁰ (CV) of intraclass distances to the CV of the interclass distances. The value of β_{CV} should be computed for increasing values of k . The smallest k after which the β_{CV} remains roughly stable should be chosen [23], as a stable β_{CV} indicates that new splits affect only marginally the variations of intra and interclass distances, implying that a well formed trend has been split.

Regarding the TrendLearner parameters, we here choose to constrain γ^{max} with the maximum number of points in our time series (100 in our case, as discussed in Section 5.2). As for vector parameters θ and γ , a traditional cross-validation in the training set to determine their optimal values would imply in a search over an exponential space of values. Moreover, note that it is fairly simple to achieve best classification results by setting θ to all zeros and γ to large values, but this would lead to very late predictions (and possibly low remaining interest in the content after prediction). Instead, we suggest an alternative approach. Considering each class i separately, we run a one-against-all classification for objects of i in D^{train} for values of γ_i varying from 1 till γ^{max} . We select the smallest value of γ_i for which the performance exceeds a minimum target (e.g., classification above random choice, meaning Micro-F1 greater than 0.5), and set θ_i to the average probability computed for all class i objects for the selected γ_i . We repeat the same process for all classes. Depending on the required tradeoff between prediction accuracy and remaining fraction of views, different performance targets could be used. Finally, we use cross-validation in the training set to choose the parameter values for the extremely randomized trees classifier, as further discussed in Section 6.

We summarize our solution to the early trend prediction problem in Algorithm 3. In particular, TrendLearner works by first learning the best parameter values and the classification model from the training set (*LearnParams* and *TrainERTrees*), and then applying the learned model to classify test objects (*PredictERTrees*), taking the class membership probabilities (*MultiClassProb*) and other object features as inputs. A pictorial representation is shown in Figure 3. Compared to previous efforts [5], our method incorporates multiple classes, uses only centroids to compute class membership probabilities (which reduces time complexity), and combines these probabilities with other object features as inputs to a classifier, which, as shown in Section 6, leads to better results.

5. Evaluation Methodology

This section presents the metrics (Section 5.1) and datasets (Section 5.2) used in our evaluation.

¹⁰The ratio of the standard deviation to the mean.

5.1. Metrics

As discussed in Section 3, an inherent challenge of the early popularity trend prediction problem is to properly address the tradeoff between prediction accuracy and how early the prediction is made. Thus, we evaluate our method with respect to these two aspects.

We estimate prediction accuracy using the *standard* Micro and Macro $F1$ metrics, which are computed from precision and recall. The precision of class c , $P(c)$, is the fraction of correctly classified videos out of those assigned to c by the classifier, whereas the recall of class c , $R(c)$, is the fraction of correctly classified objects out of those that actually belong to that class. The $F1$ of class c is given by: $F1(c) = \frac{2 \cdot P(c) \cdot R(c)}{P(c) + R(c)}$. Macro $F1$ is the average $F1$ across all classes, whereas Micro $F1$ is computed from global precision and recall, calculated for all classes.

To complement the standard metrics above, we propose the use of novel metrics that we define to measure the effectiveness of the early predictions extracted by TrendLearner. These metrics are by no means replacements for standard classification evaluation metrics (such as the $F1$ defined above). That is, given that TrendLearner aims to capture the trade-off between accuracy and early predictions, our proposed novel metrics need to be evaluated together with the traditional ones. Recall that, our objectives are to evaluate both the: (1) accuracy of the classification; and (2) the possible loss of user interest in objects over time.

We evaluate how early our correct predictions are made computing the remaining interest (RI) in the content *after* prediction. The RI for an object s_d is defined as the fraction of all views up to a certain point in time (e.g., the day when the object was collected) that are received *after* the prediction. That is, $RI(s_d, t) = \frac{sum(s_d^{[t^{[d]}+1:n]})}{sum(s_d^{[1:n]})}$ where n is the number of points in d 's time series, $t^{[d]}$ is the prediction time (i.e., monitoring period) produced by our method for d , and function sum adds up the elements of the input vector. In essence, this metric captures the future potential audience of s_d after prediction.

We also assess whether there is any bias in our *correct* predictions towards more (less) popular objects by computing the correlation between the total popularity and the remaining interest after prediction for each object. A low correlation implies no bias, while a strong positive (negative) correlation implies a bias towards earlier predictions for more (less) popular objects. We argue that, if any bias exists, a bias towards more popular objects is preferred, as it implies larger remaining interests for those objects. We use both the Pearson linear correlation coefficient (ρ_p) and the Spearman's rank correlation coefficient (ρ_s) [16], as the latter does not assume linear relationships, taking the logarithm of the total popularity first due to the great skew in their distribution [4, 7, 11].

5.2. Datasets

As case study, we focus on YouTube videos and use two-datasets, analyzed in [11] and publicly available¹¹. The **Top**

¹¹<http://vod.dcc.ufmg.br/traces/youtime/>

Algorithm 3 Our Solution: Trend Extraction and Prediction

```

1: function TRENDExTRACTION( $D^{train}$ )
2:    $k \leftarrow 1$ 
3:   while  $\beta_{CV}$  is not stable do
4:      $k \leftarrow k + 1$ 
5:      $C_D \leftarrow KSC(D^{train}, k)$ 
6:   end while
7:   Store centroids in  $C_D$ 
8: end function
9: function TRENDLEARNER( $C_D, D^{train}, D^{test}$ )
10:   $\theta, \gamma, \mathbf{P}^{train} \leftarrow LearnParams(D^{train}, C_D)$ 
11:   $TrainERTree(D^{train}, \mathbf{P}^{train} \cup \text{obj. feats})$ 
12:   $\mathbf{t}, \mathbf{P} \leftarrow MultiClassProbs(D^{test}, C_D, \theta, \gamma)$ 
13:  return  $\mathbf{t}, PredictERTree(D^{test}, \mathbf{P} \cup \text{obj. feats})$ 
14: end function

```

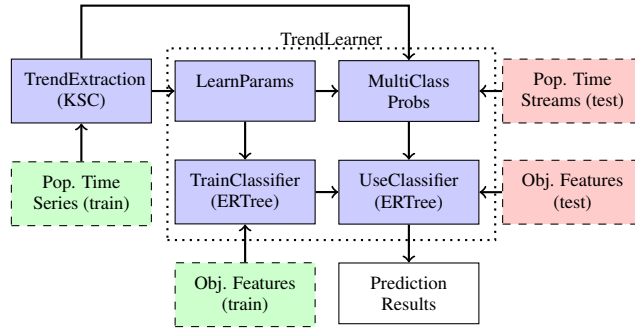


Figure 3: Pictorial Representation of Our Solution

dataset consists of 27,212 videos from the various top lists maintained by YouTube (e.g., most viewed and most commented videos), and the **Random topics** dataset includes 24,482 videos collected as results of random queries submitted to YouTube’s API¹².

For each video, the datasets contain the following features (shown in Table 3): the time series of the numbers of views, comments and favorites, as well as the ten most important referrers (incoming links), along with the date that referrer was first encountered, the video’s upload date and its category. The original datasets contain videos of various ages, ranging from days to years. We choose to study only videos with more than 100 days for two reasons. First, these videos tend to have their long term time series popularity more stable. Second, the KSC algorithm requires that all time series vectors s_d have the same dimension n . Moreover, the popularity time series provided by YouTube contains at most 100 points, independently of the video’s age. Thus, by focusing only on videos with at least 100 days of age, we can use n equal to 100 for all videos. After filtering younger videos out, we were left with 4,527 and 19,562 videos in the Top and Random datasets, respectively.

Table 4 summarizes our two datasets, providing mean μ and standard deviation σ for the number of views, age (in days),

Table 3: Summary of Features

Class	Feature Name	Type
Video	Video category	Categorical
	Upload date	Numerical
	Video age	Numerical
	Time window size (w)	Numerical
Referrer	Referrer first date	Numerical
	Referrer # of views	Numerical
Popularity	# of views	Numerical
	# of comments	Numerical
	# of favorites	Numerical
	change rate of views	Numerical
	change rate of comments	Numerical
	change rate of favorites	Numerical
	Peak fraction	Numerical

and time window duration w ¹³. Note that both average and median window durations are around or below one week. This is important as previous work [2] pointed out that effective popularity growth models can be built based on weekly views.

¹²We do not claim this dataset is a random sample of YouTube videos. Nevertheless, for the sake of simplicity, we use the term Random videos to refer to videos from this dataset.

¹³ w is equal to the video age divided by 99 as the first point in the time series corresponds to the day before the upload day.

Table 4: Summary of analyzed datasets

	Top		Random	
	μ	σ	μ	σ
# of Views	4,022,634	9,305,996	141,413	1,828,887
Video Age (days)	632	402	583	339
Window w (days)	6.38	4.06	5.89	3.42

6. Experimental Results

In this section, we present our results of our trend extraction (Section 6.1) and trend prediction (Section 6.2) approaches. We also show how TrendLearner can be used to improve the accuracy of state-of-the-art popularity prediction models (Section 6.3). These results were computed using 5-fold cross validation, i.e., splitting the dataset D into 5 folds, where 4 are used as training set D^{train} and one as test set D^{test} , and rotating the folds such that each fold is used for testing once. As discussed in Section 4, trends are extracted from D^{train} and predicted for videos in D^{test} .

Since we are dealing with time series, one might argue that a temporal split of the dataset into folds would be preferred to a random split, as we do here. However, we choose a random split because of the following. Regarding the object features used as input to the prediction models, no temporal precedence is violated, as the features are computed only during the monitoring period t_r , before prediction. All remaining features are based on the distances between the popularity curve of the object until t_r and the class centroids (or trends). As we argue below, the same trends/centroids found in our experiments were consistently found in various subsets of each dataset, covering various periods of time. Thus, we expect the results to remain similar if a temporal split is done. However, a temporal split of our dataset would require interpolations in the time series, as all of them have exactly 100 points regardless of video age. Such interpolations, which are not required in a random split, could introduce serious inaccuracies and compromise our analyses.

6.1. Trend Extraction

Recall that we used the β_{CV} metric to determine the number of trends k used by the KSC algorithm. In both datasets, we found k to be stable after 4 trends. We also checked centroids and class members for larger values of k , both visually and using other metrics (as in [30]), finding no reason for choosing a different value¹⁴. Thus, we set $k = 4$. We also analyzed the centroids in all training sets, finding that the same 4 shapes appeared in every set. Thus, we manually aligned classes based on their centroid shapes in different training sets so that class i is the same in every set. We also found that, in 95% of the cases, a video was always assigned to the same class in different sets.

Figure 4 shows the popularity trends discovered in the Random dataset. Similar trends were also extracted from the Top

dataset. Each graph shows the number of views as function of time, omitting scales as centroids are shape and volume invariants. The y-axes are in log scale to highlight the importance of the peak. We note that the KSC algorithm consistently produced the same popularity trends for various randomly selected samples of the data, which are also consistent with similar shapes identified in other datasets [7, 30]. We also note that the 4 identified trends might not perfectly match the popularity curves of all videos, as there might be variations within each class. However, our goal is not to perfectly model the popularity evolution of all videos. Instead, we aim at capturing the most prevalent trends, respecting time shift and volume invariants, and using them to improve popularity prediction. As we show in Section 6.3, the identified trends can greatly improve state-of-the-art prediction models.

Table 5 presents, for each class, the percentage of videos belonging to it, as well as the average number of views, average change rate¹⁵, and average fraction of views at the peak time window of these videos. Note that class D_0 consists of videos that remain popular over time, as indicated by the large positive change rates, shown in Table 5. This behavior is specially strong in the Top dataset, with an average change rate of 1,112 views per window, which corresponds to roughly a week (Table 4). Those videos also have no significant popularity peak, as the average fraction of views in the peak window is very small (Table 5). The other three classes are predominantly defined by a single popularity peak, and are distinguished by the rate of decline after the peak: it is slower in D_1 , faster in D_2 , and very sharp in D_3 . These classes also exhibit very small change rates, indicating stability after the peak.

We also measured the distribution of different types of referers and video categories across classes in each dataset. Under a Chi-square test with significance of .01, we found that the distribution differs from that computed for the aggregation of all classes, implying that these features are somewhat correlated with the class, and motivating their use to improve trend classification.

6.2. Trend Prediction

We now discuss our trend prediction results, which are averages of 5 test sets along with corresponding 95% confidence intervals. We start by showing results that support our approach of computing class membership probabilities using only centroids as opposed to all class members, as in [5] (Section 6.2.1).

¹⁴A possible reason would be the appearance of a new distinct class, which did not happen.

¹⁵Defined by the average of $p_{d,i+1} - p_{d,i}$ for each video d represented by vector $\mathbf{s}_d = \langle p_{d,1}, p_{d,2}, \dots, p_{d,n} \rangle$.

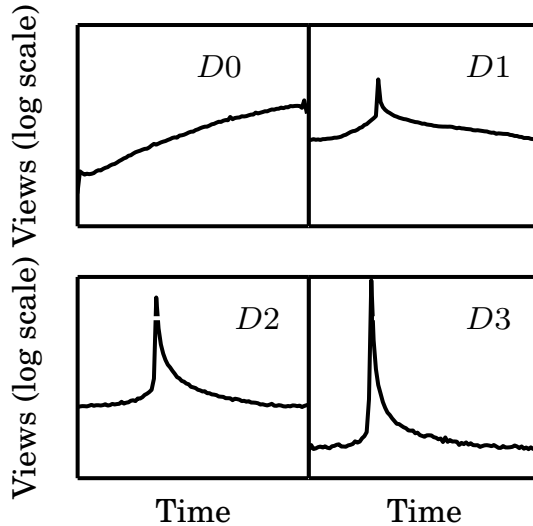


Figure 4: Popularity Trends in YouTube Datasets.

Table 5: Summary of popularity trends (classes)

	D_0	D_1	D_2	D_3
	Top Dataset			
% of Videos	22%	29%	24%	25%
Avg. # of Views	711,868	6,133,348	1,440,469	1,279,506
Avg. Change Rate in # Views	1112	395	51	67
Avg. Peak Fraction	0.03	0.04	0.19	0.40
	Random Dataset			
% of Videos	21%	34%	26%	19%
Avg. # of Views	305,130	108,844	64,274	127,768
Avg. Change Rate in # Views	47	7	4	4
Avg. Peak Fraction	0.03	0.03	0.08	0.28

Table 6: Classification using only centroids vs. using all class members: averages and 95% confidence intervals.

Monitoring period t_r	Centroid		Whole Training Set	
	Micro F1	Macro F1	Micro F1	Macro F1
1 window	.24 ± .01	.09 ± .00	.29 ± .04	.11 ± .01
25 windows	.56 ± .02	.52 ± .01	.53 ± .04	.44 ± .08
50 windows	.67 ± .03	.65 ± .03	.64 ± .05	.57 ± .09
75 windows	.70 ± .02	.68 ± .02	.69 ± .08	.61 ± .12

We then evaluate our TrendLearner method, comparing it with three alternative approaches (Section 6.2.2).

6.2.1. Are shapelets better than a reference dataset?

We here discuss how the use of centroids to compute class membership probabilities (Equation 2) compare to using all class members [5]. For the latter, the probability of an object belonging to a class is proportional to a summation over the

exponential of the (negative) distance between the object and every member of the given class.

An important benefit of our approach is a reduction in running time: for a given object, it requires computing the distances to only k time series, as opposed to the complete training set $|D^{train}|$, leading to a reduction in running time by a factor of $\frac{|D^{train}|}{k}$, as discussed in Section 4.1. We here focus on the classification effectiveness of the probability matrix \mathbf{P} produced by both approaches. To that end, we consider a classifier that assigns the class with largest probability to each object, for both matrices.

Table 6 shows Micro and Macro F1 results for both approaches, computed for fixed monitoring periods t_r (in number of windows) to facilitate comparison. We show results only for the Top dataset, as they are similar for the Random dataset. Note that, unless the monitoring period is very short ($t_r=1$), both strategies produce statistically tied results, with 95% confidence. Thus, given the reduced time complexity, using centroids only is more cost-effective. When using a single window

Table 7: Best values for vector parameters γ and θ (averages and 95% confidence intervals) for the Top dataset

Top Dataset				
	D_0	D_1	D_2	D_3
θ	$.250 \pm .015$	$.257 \pm .001$	$.272 \pm .003$	$.303 \pm .006$
γ	28 ± 16	89 ± 8	5 ± 0.9	3 ± 0.5

Table 8: Best values for vector parameters γ and θ (averages and 95% confidence intervals) for the Random dataset

Random Dataset				
	D_0	D_1	D_2	D_3
θ	$.250 \pm .001$	$.251 \pm .001$	$.269 \pm 0.001$	$.317 \pm 0.001$
γ	33 ± 0.6	74 ± 2	45 ± 9	17 ± 3

both approaches are worse than random guessing (Macro F1 = 0.25), and thus are not interesting.

6.2.2. TrendLearner Results

We now compare our **TrendLearner** method with three other trend prediction methods, namely: (1) **P only**: assigns the class with largest probability in **P** to an object; (2) **P + ERTree**: trains an extremely randomized trees learner using **P** only as features; (3) **ERTree**: trains an extremely randomized trees learner using only the object features in Table 3. Note that TrendLearner combines ERTree and **P + ERTree**. Thus, a comparison of these four methods allows us to assess the benefits of combining both sets of features.

For *all* methods, when classifying a video d , we only consider features of that video available up until $t^{|d|}$, the time window when TrendLearner stopped monitoring d . We also use the same best values for parameters shared by the methods, chosen as discussed in Section 4.3. Both Tables 7 (for the Top dataset) and 8 (for the Random dataset), show the best values of vector parameters γ and θ , selected considering a Macro-F1 of at least 0.5 as performance target (see Section 4.3). These results are averages across all training sets, along with 95% confidence intervals. The variability is low in most cases, particular for θ . Recall that γ^{max} is set to 100. Regarding the extremely randomized trees classifier, we set the size of the ensemble to 20 trees, and the feature selection strength equal to the square root of the total number of features, common choices for this classifier [12]. We then apply cross-validation within the training set to choose the smoothing length parameter (n_{min}), considering values equal to $\{1, 2, 4, 8, 16, 32\}$. We refer to [12] for more details on the parametrization of extremely randomized trees.

Still analyzing Tables 7 and 8, we note that classes with smaller peaks (D_0 and D_1) need longer minimum monitoring periods γ_i , likely because even small fluctuations may be confused as peaks due to the scale invariance of the distance metric used (Equation 1)¹⁶. However, after this period, it is somewhat

¹⁶Indeed, most of these videos are wrongly classified into either D_2 or D_3

Table 9: Comparison of trend prediction methods for both datasets (averages and 95% confidence intervals) for the Top dataset

Top Dataset				
	P only	P+ERTree	ERTree	TrendLearner
Micro F1	$.48 \pm .06$	$.48 \pm .06$	$.58 \pm .01$	$.62 \pm .01$
Macro F1	$.44 \pm .06$	$.44 \pm .06$	$.57 \pm .01$	$.61 \pm .01$

Table 10: Comparison of trend prediction methods for both datasets (averages and 95% confidence intervals) for the Random dataset

Random Dataset				
	P only	P+ERTree	ERTree	TrendLearner
Micro F1	$.67 \pm .02$	$.62 \pm .01$	$.65 \pm .01$	$.71 \pm .01$
Macro F1	$.69 \pm .02$	$.63 \pm .01$	$.63 \pm .01$	$.70 \pm .01$

easier to determine whether the object belongs to one of those classes (smaller values of θ_i). In contrast, classes with higher peaks (D_2 and D_3) usually require shorter monitoring periods, particularly in the Top dataset, where videos have popularity peaks with larger fractions of views (Table 5). Indeed, by cross-checking results in Tables 5, 7 and 8, we find that classes with smaller fractions of videos in the peak window (D_0 and D_1 in Top, and D_0 , D_1 and D_2 in Random) tend to require longer minimum monitoring periods so as to avoid confusing small fluctuations with peaks from the other classes.

We now discuss our classification results, focusing first on the Micro and Macro F1 results, shown in Table 9 and Table 10, for the Top and Random datasets respectively. From both tables we can see that TrendLearner consistently outperforms all other methods in both datasets and on both metrics, except for Macro F1 in the Random dataset, where it is statistically tied with the second best approach (**P** only). In contrast, there is no clear winner among the other three methods across both datasets. Thus, combining probabilities and object features brings clear benefits over using either set of features separately. For example, in the Top dataset, the gains over the alternatives in average Macro F1 vary from 7% to 38%, whereas the average improvements in Micro F1 vary from 7% to 29%. Similarly, in the Random dataset, gains in average Micro and Macro F1 reach up to 14% and 11%, respectively. Note that TrendLearner performs somewhat better in the Random dataset, mostly because videos in that dataset are monitored for longer, on average (larger values of γ_i). However, this superior results comes with a reduction in remaining interest after prediction, as we discuss below.

We note that the joint use of both probabilities and object features renders TrendLearner more robustness to some (hard-to-predict) videos. Recall that, as discussed in Section 4.2.1,

for shorter monitoring periods.

Algorithm 2 may, in some cases, return a probability equal to 0 to indicate that a prediction was not possible within the maximum monitoring period allowed. Indeed, this happened for 1% and 10% of the videos in the Top and Random datasets, respectively, which have popularity curves that do not closely follow any of the extracted trends. The results for the **P** only and **P** + ERTree methods shown in Tables 9 and 10 do *not* include such videos, as these methods are not able to do predictions for them (since they rely only on the probabilities). However, both ERTree and TrendLearner are able to perform predictions for such videos by exploiting the object features, since at least the video category and upload date are readily available as soon as the video is posted. Thus, the results of these two methods in Tables 9 and 10 contemplate the predictions for *all* videos¹⁷.

We now turn to the other side of the tradeoff and discuss how early the predictions are made. These results are the same for all four aforementioned methods as all of them use the prediction time returned by TrendLearner. For all *correctly* classified videos, we report the remaining interest *RI* after prediction, as well as the Pearson (ρ_p) and Spearman (ρ_s) correlation coefficients between remaining interest and (logarithm of) total popularity (i.e., total number of views), as informed in our datasets.

Figure 5(a) shows the complementary cumulative distribution of the fraction of *RI* after prediction for both datasets, while Figures 5(b) and 5(c) (log scale on the y-axis) show the total number of views and the *RI* for each video in the Top and Random datasets, respectively. All three graphs were produced for the union of the videos in all test sets. Note that, for 50% of the videos, our predictions are made before at least 68% and 32% of the views are received, for Top and Random videos, respectively. The same *RI* of at least 68% of views is achieved for 21% of videos in the Random dataset. In general, for a significant number of videos in both datasets, our correct predictions are made before a large fraction of their views are received, particularly in the Top dataset.

We also point out a great variability in the duration of the monitoring periods produced by our solution: while only a few windows are required for some videos, others have to be monitored for a longer period. Indeed, the coefficients of variation of these monitoring periods are 0.54 and 1.57 for the Random and Top datasets, respectively. This result emphasizes the need for choosing a monitoring period on a per-object basis, a novel aspect of our approach, and not use the same fixed value.

Moreover, the scatter plots in Figures 5(b-c) show that some moderately positive correlations exist between the total number of views and *RI*. Indeed, ρ_p and ρ_s are equal to 0.42 and 0.48, respectively, in the Top dataset, while both metrics are equal to 0.39 in the Random dataset. Such results imply that our solution is somewhat biased towards more popular objects, although the bias is not very strong. In other words, for more popular videos, TrendLearner is able to produce accurate predictions by

potentially observing a smaller fraction of their total views, in comparison with less popular videos. This is a nice property, given that such predictions can drive advertisement placement and content replication/organization decisions which are concerned mainly with the most popular objects.

6.3. Applicability to Regression Models

Motivated by results in [25, 29], which showed that knowing popularity trends *beforehand* can improve the accuracy of regression-based popularity prediction models, we here assess whether our trend predictions are good enough for that purpose. To that end, we use the state-of-the-art ML and MRBF regression models proposed in [25]. The former is a multivariate linear regression model that uses the popularity acquired by the object d on each time window up to a reference date t_r (i.e., $p_{d,i}$, $i = 1 \dots t_r$) to predict its popularity at a target date $t_t = t_r + \delta$. The latter extends the former by including features based on Radial Basis Functions (RBFs) to measure the similarity between d and specific examples, previously selected from the training set.

Our goal is to evaluate whether our trend prediction results can improve these models. Thus, as in [25], we use the mean Relative Squared Error (mRSE) to assess the prediction accuracy of the ML and MRBF models in two settings: (1) a general model, trained using the whole dataset (as in [25]); (2) a specialized model, trained for each *predicted* class. For the latter, we first use our solution to predict the trend of a video. We then train ML and MRBF models considering as reference date each value of $\mathbf{t}^{[d]}$ produced by TrendLearner for each video d . Considering a prediction lag δ equal to 1, 7, and 15, we measure the mRSE of the predictions for target date $t_t = \mathbf{t}^{[d]} + \delta$. We also compare our specialized models against the state-space models (SSMs) proposed in [26]. These models are variations of a basic state-space Holt-Winters model that represent query and click frequency in Web search, capturing various aspects of popularity dynamics (e.g., periodicity, bursty behavior, increasing trend). All of them take as input the popularity time series during the monitoring period t_r . Thus, though originally proposed for the Web search domain, they can be directly applied to our context. Both regression and state-space models are parametrized as originally proposed¹⁸.

Table 11 shows average mRSE for each model along with 95% confidence intervals, for all datasets and prediction lags. Comparing our specialized models and the original ones they build upon, we find that using our solution to build trend-specific models greatly improves prediction accuracy, particularly for larger values of δ . The reductions in mRSE vary from 10% to 77% (39%, on average) in the Random dataset, and from 11% to 64% (33%, on average) in the Top dataset¹⁹. The specialized models also greatly outperform the state-space models: the reductions in mRSE over the best state-space model are

¹⁷For the cases with probability equal to 0, the predictions of TrendLearner and ERTree were made with $t_r = \gamma^{max}$, when Algorithm 2 stops. Since we set $\gamma^{max} = 100$, those predictions were made at the last time window, using all available information to compute object features. Nevertheless, note that, in those cases, the remaining interest (*RI*) after prediction is equal to 0.

¹⁸The only exception is the number of examples used to compute similarities in the MRBF model: we used 50 examples, as opposed to the suggested 100 [25], as it led to better results in our datasets.

¹⁹The only exception is the MRBF model for $\delta=1$ in the Top dataset, where general and specialized models produce tied results.

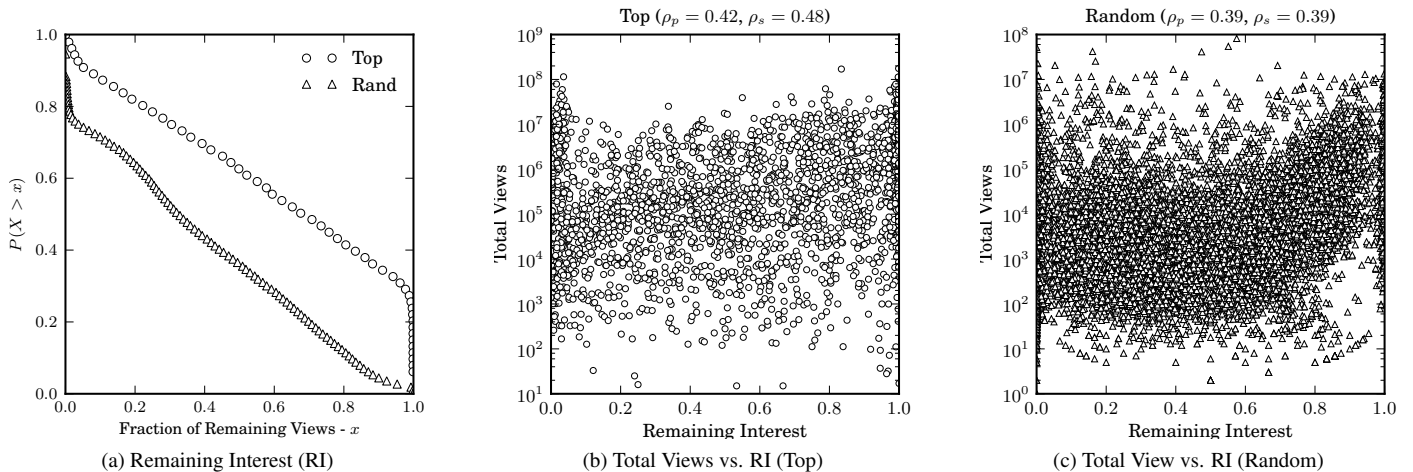


Figure 5: Remaining Interest (RI) and Correlations Between Popularity and RI for Correctly Classified Videos.

Table 11: Mean Relative Squared Error Various Prediction Models and Lags δ (averages and 95% confidence intervals)

Prediction Model	Top Dataset			Random Dataset		
	$\delta = 1$	$\delta = 7$	$\delta = 15$	$\delta = 1$	$\delta = 7$	$\delta = 15$
generalML	.09 \pm .005	.42 \pm .02	.75 \pm .04	.01 \pm .001	.06 \pm .005	.11 \pm .01
generalMRBF	.08 \pm .005	.52 \pm .05	1.29 \pm .17	.01 \pm .001	.1 \pm .01	.26 \pm .03
best SSM	.76 \pm .01	.63 \pm .02	.64 \pm .03	.90 \pm .002	.69 \pm .005	.54 \pm .006
specializedML	.08 \pm .005	.27 \pm .01	.38 \pm .02	.009 \pm .001	.04 \pm .0003	.06 \pm .003
specializedMRBF	.08 \pm .005	.32 \pm .04	.47 \pm .08	.009 \pm .001	.04 \pm .0004	.06 \pm .008

at least 89% and 27% in the Random and Top datasets (94% and 59%, on average). These results offer strong indications of the usefulness of our trend predictions for predicting popularity measures.

Finally, it is important to discuss why the state-space models did not work well in our context. The main reason we found was that Holt-Winters based models can only capture the linear trends in time series, that is, linear growth and decay. By using the KSC distance function, we can identify and group UGC time series with non-linear trends [22, 30], and create specific prediction models for these cases. Also, these models are trained independently for each target object, using early points of the time series. Another possible reason for the low performance in our context might be that, unlike in [26] where the models were trained with hundreds of points of each time series, we here use much less data (only points up to $t^{(d)}$).

7. Conclusions

In this article, we have identified and formalized a new research problem. To the extent of our knowledge, we are the first work to tackle the problem of *early prediction* of popularity trends in UGC. We were motivated in studying this problem based on our previous knowledge on the complex patterns and causes of popularity in UGC [11]. Different from other kinds of content, e.g., news, which have clear definitions of monitoring

periods, target and prediction dates for popularity, the complex nature of UGC calls for a popularity prediction solution which is able to determine these dates automatically. We here provided such a solution – TrendLearner.

We have also proposed a novel two-step learning approach for early prediction of popularity trends of UGC. Moreover, we defined new metrics for measuring the effectiveness of popularity of UGC content, the remaining interest, which is optimized by TrendLearner as to provide not only accurate, but also timely, predictions. Thus, unlike previous work, we address the tradeoff between prediction accuracy and remaining interest in the content after prediction on a per-object basis.

We performed an extensive experimental evaluation of our method, comparing it with state-of-the-art, representative solutions of the literature. Our experimental results on two YouTube datasets showed that our method not only outperforms other approaches for trend prediction (a gain of up to 38%) but also achieves such results before 50% or 21% of videos (depending on the dataset) accumulate more than 32% of their views, with a slight bias towards earlier predictions for more popular videos. Moreover, when applied jointly with recently proposed regression based models to predict the popularity of a video at a future date, our method outperforms state-of-the-art regression and state-space based models, with gains in accuracy of at least 33% and 59%, on average, respectively.

As future work, we plan to further investigate how differ-

ent types of UGC (e.g., blogs and Flickr photos) differ in their popularity evolution as well as which factors (e.g., referrers, content quality) impact this evolution.

Acknowledgments

This research is partially funded by the Brazilian National Institute of Science and Technology for Web Research (MCT/CNPq/INCT Web Grant Number 573871/2008-6), and by the authors' individual grants from Google, CNPq, CAPES and Fapemig.

References

- [1] M. Ahmed, S. Spagna, F. Huici, and S. Niccolini. A Peek Into the Future: Predicting the Evolution of Popularity in User Generated Content. In *Proc. WSDM*, 2013.
- [2] Y. Borghol, S. Mitra, S. Ardon, N. Carlsson, D. Eager, and A. Mahanti. Characterizing and Modeling Popularity of User-Generated Videos. *Performance Evaluation*, 68(11):1037–1055, 2011.
- [3] C. Castillo, M. El-Haddad, J. Pfeffer, and M. Stempeck. Characterizing the life cycle of online news stories using social media reactions. In *Proc. CSCW*, 2014.
- [4] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon. Analyzing the Video Popularity Characteristics of Large-Scale User Generated Content Systems. *IEEE/ACM Transactions on Networking*, 17(5):1357–1370, 2009.
- [5] G. H. Chen, S. Nikolov, and D. Shah. A Latent Source Model for Non-parametric Time Series Classification. In *Proc. NIPS*, 2013.
- [6] A. Coates and A. Ng. Learning Feature Representations with K-Means. *Neural Networks: Tricks of the Trade*, pages 561–580, 2012.
- [7] R. Crane and D. Sornette. Robust Dynamic Classes Revealed by Measuring the Response Function of a Social System. *Proceedings of the National Academy of Sciences*, 105(41):15649–53, 2008.
- [8] Q. Duong, S. Goel, J. Hofman, and S. Vassilvitskii. Sharding social networks. In *Proc. WSDM*, Feb. 2013.
- [9] S. Džeroski and B. Ženko. Is Combining Classifiers with Stacking Better than Selecting the Best One? *Machine Learning*, 54(3):255–273, 2004.
- [10] F. Figueiredo, J. Almeida, and M. Gonçalves. Improving the Effectiveness of Content Popularity Prediction Methods using Time Series Trends. In *Proc. ECML/PKDD Predictive Analytics Challenge Workshop*, 2014.
- [11] F. Figueiredo, F. Benevenuto, M. Gonçalves, and J. Almeida. On the Dynamics of Social Media Popularity: A YouTube Case Study. *ACM Trans. Internet Technol.*, 14(4):24:1–24:23, 2014.
- [12] P. Geurts, D. Ernst, and L. Wehenkel. Extremely Randomized Trees. *Machine Learning*, 63(1):3–42, 2006.
- [13] P. Gill, V. Erramilli, A. Chaintreau, B. Krishnamurthy, D. Papagiannaki, and P. Rodriguez. Follow the Money: Understanding Economics of Online Aggregation and Advertising. In *Proc. IMC*, 2013.
- [14] N. G. Golbandi, L. K. Katzir, Y. K. Koren, and R. L. Lempel. Expediting Search Trend Detection via Prediction of Query Counts. In *Proc. WSDM*, 2013.
- [15] Q. Hu, G. Wang, and P. S. Yu. Deriving Latent Social Impulses to Determine Longevous Videos. In *Proc. WWW*, 2014.
- [16] R. Jain. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley, 1991.
- [17] L. Jiang, Y. Miao, Y. Yang, Z. Lan, and A. G. Hauptmann. Viral video style: A closer look at viral videos on youtube. In *Proc. ICMR*, ICMR '14, 2014.
- [18] J. G. Lee, S. Moon, and K. Salamatian. An Approach to Model and Predict the Popularity of Online Contents with Explanatory Factors. In *Proc. WIC*, volume 1, 2010.
- [19] K. Lerman and T. Hogg. Using a Model of Social Dynamics to Predict Popularity of News. In *Proc. WWW*, 2010.
- [20] J. Leskovec. Social Media Analytics. In *Proc. WWW*, 2011.
- [21] H. Li, X. Ma, F. Wang, J. Liu, and K. Xu. On Popularity Prediction of Videos Shared in Online Social Networks. In *Proc. CIKM*, 2013.
- [22] Y. Matsubara, Y. Sakurai, B. A. Prakash, L. Li, and C. Faloutsos. Rise and Fall Patterns of Information Diffusion. In *Proc. KDD*, 2012.
- [23] D. Menascé and V. Almeida. *Capacity Planning for Web Services: Metrics, Models, and Methods*. Prentice Hall, 2002.
- [24] K. Nigam and R. Ghani. Analyzing the Effectiveness and Applicability of Co-training. In *Proc. CIKM*, 2000.
- [25] H. Pinto, J. Almeida, and M. Gonçalves. Using Early View Patterns to Predict the Popularity of YouTube Videos. In *Proc. WSDM*, 2013.
- [26] K. Radinsky, K. Svore, S. Dumais, J. Teevan, A. Bocharov, and E. Horvitz. Behavioral Dynamics on the Web: Learning, Modeling, and Prediction. *ACM Transactions on Information Systems*, 32(3):1–37, 2013.
- [27] G. Szabo and B. A. Huberman. Predicting the Popularity of Online Content. *Communications of the ACM*, 53(8):80–88, 2010.
- [28] A. Vakali, M. Giatsoglou, and S. Antaris. Social Networking Trends and Dynamics Detection via a Cloud-Based Framework Design. In *Proc. WWW*, 2012.
- [29] J. Yang and J. Leskovec. Modeling Information Diffusion in Implicit Networks. In *Proc. ICDM*, 2010.
- [30] J. Yang and J. Leskovec. Patterns of Temporal Variation in Online Media. In *Proc. WSDM*, 2011.
- [31] L. Ye and E. Keogh. Time Series Shapelets: A Novel Technique that Allows Accurate, Interpretable and Fast Classification. *Data Mining and Knowledge Discovery*, 22(1-2):149–182, 2011.
- [32] P. Yin, P. Luo, M. Wang, and W.-C. Lee. A Straw Shows Which Way the Wind Blows: Ranking Potentially Popular Items from Early Votes. In *Proc. WSDM*, 2012.
- [33] H. Yu, L. Xie, and S. Sanner. Exploring the Popularity Phases of YouTube Videos: Observations, Insights, and Prediction. In *Proc. ICWSM*, 2015.
- [34] D. Zeng, H. Chen, R. Lusch, and S.-H. Li. Social Media Analytics and Intelligence. *IEEE Intelligent Systems*, 25(6):13–16, 2010.