

Dynamic PageRank Using Evolving Teleportation

Ryan A. Rossi and David F. Gleich

Purdue University
Department of Computer Science
305 N. University St., West Lafayette, IN 47906
{rossi, dgleich}@purdue.edu

Abstract. The importance of nodes in a network constantly fluctuates based on changes in the network structure as well as changes in external interest. We propose an evolving teleportation adaptation of the PageRank method to capture how changes in external interest influence the importance of a node. This framework seamlessly generalizes PageRank because the importance of a node will converge to the PageRank values if the external influence stops changing. We demonstrate the effectiveness of the evolving teleportation on the Wikipedia graph and the Twitter social network. The external interest is given by the number of hourly visitors to each page and the number of monthly tweets for each user.

1 Introduction

Finding important nodes in a graph is a key task in a variety of applications: search engines [24,18], network science [17,8,14], and bioinformatics [27,22], among many others. By and large, these are global measures of node importance and one of the most well-studied measures is PageRank [24,20].

PageRank computes the importance of each node in a directed graph under a random surfer model. When at a node, the random surfer can either:

1. transition to a new node from the set of out-edges, or
2. do something else (e.g., execute a search query, use a bookmark).

The probability that the surfer performs the first action is known as the damping parameter in PageRank. We use α to denote the damping parameter. The second action is called teleporting and is modeled by the surfer picking a node at random according to a distribution called the teleportation distribution vector or personalization vector. These choices only depend on the current node and, consequently, define a Markov chain. This PageRank Markov chain always has a unique stationary distribution for any $0 \leq \alpha < 1$. The importance of a node is proportional to its stationary distribution in this Markov chain. Thus, the computation is governed by the graph, a teleportation parameter α , and a teleportation distribution vector.

The PageRank score is a simple model for the importance of a node in a graph, and there are many variations that may yield more useful scores (for instance [21] models a random walk with a back button). A common complaint

about PageRank models is that they are only defined for static graphs. Motivated by the idea of studying PageRank with dynamic graphs, we formulate a dynamic PageRank model for a static graph with a time-dependent, or evolving, teleportation vector. Intuitively, the teleportation distribution changes based on human dynamics such as recent news and seasonal preferences. For example, in our forthcoming experiments (Section 6), the time-dependent vector is the number of hourly page visits for each page from Wikipedia. We derive the model and algorithms for this dynamic version of PageRank in Section 4. The resulting algorithms scale to large graphs. Moreover, we show that the new model is a generalization of PageRank in the sense that *if the time-dependent vector stops changing then our dynamic score vector converges to the standard PageRank score*.

We make our code and data available in the spirit of reproducible research:

<http://www.cs.purdue.edu/homes/dgleich/codes/dynsyspr-waw>

2 PageRank Notation

In order to place our work in context, we first introduce some notation. Let \mathbf{A} be the adjacency matrix for a graph where $A_{i,j}$ denotes an edges from node i to node j . In order to avoid a proliferation of transposes, we define \mathbf{P} as the transposed transition matrix for a random-walk on a graph:

$$P_{j,i} = \text{probability of transitioning from node } i \text{ to node } j.$$

Hence, the matrix \mathbf{P} is *column-stochastic* instead of row-stochastic, which is the standard in probability theory. Throughout this manuscript, we utilize uniform random-walks on a graph, in which case $\mathbf{P} = \mathbf{A}^T \mathbf{D}^{-1}$ where \mathbf{D} is a diagonal matrix with the degree of each node on the diagonal. However, none of the theory is restricted to this type of random walk and any column-stochastic matrix will do. The PageRank vector \mathbf{x} is the solution of the linear system:

$$(\mathbf{I} - \alpha \mathbf{P})\mathbf{x} = (1 - \alpha)\mathbf{v}$$

for any $0 \leq \alpha < 1$ and any teleportation distribution vector \mathbf{v} such that $v_i \geq 0$ and $\sum v_i = 1$. Table 1 summarizes these notation conventions, and has a few other elements that will be discussed in the forthcoming sections.

3 Dynamic and Evolving Rankings

The PageRank literature is vast, and we now survey some of the other ideas related to incorporating graph dynamics into a PageRank vector, more general models for studying dynamic graphs, and updating PageRank vectors.

Our proposed method is related to changing the teleportation vector in the power method as its being computed. Bianchini et al. [5] noted that the power method would still converge if either the graph or the vector \mathbf{v} changed during the method, albeit to a new solution given by the new vector or graph.

Table 1. Summary of notation. Matrices are bold, upright roman letters; vectors are bold, lowercase roman letters; and scalars are unbolded roman or greek letters.

n	number of nodes in a graph
\mathbf{e}	the vector of all ones
\mathbf{P}	column stochastic matrix
α	damping parameter in PageRank
\mathbf{v}	teleportation distribution vector
\mathbf{x}	solution to the PageRank computation

$\mathbf{v}(t)$	a teleportation distribution vector at time t
$\mathbf{x}(t)$	solution to the Dynamic PageRank computation for time t
θ	decay parameter for time-series smoothing

Our method capitalizes on a closely related idea and we utilize the intermediate quantities explicitly. Another related idea is the Online Page Importance Computation (OPIC) [1], which integrates a PageRank-like computation *with* a crawling process. The method does nothing special if a node has changed when it is crawled again. A more detailed study of how PageRank values evolve during a web-crawl was done by Boldi et al. [7]. Other work has approximated PageRank on graph streams [11].

Outside of the context of web-ranking, O'Madadhain and Smyth propose EventRank [23], a method of ranking nodes in dynamic graphs, that uses the PageRank propagation equations for a sequence of graphs. We utilize the same idea but place it within the context of a dynamical system.

While we described PageRank in terms of a random-surfer model above, another characterization of PageRank is that it is a sum of damped transitions:

$$\mathbf{x} = (1 - \alpha) \sum_{k=0}^{\infty} (\alpha \mathbf{P})^k \mathbf{v}.$$

These transitions are a type of probabilistic walk and Grindrod et al. [16] introduced the related notion of dynamic walks for dynamic graphs.

In the context of popularity dynamics [25], our method captures how changes in external interest influence the popularity of nodes and the nodes linked to these nodes in an implicit fashion. Our work is also related to modeling human dynamics, namely, how humans change their behavior when exposed to rapidly changing or unfamiliar conditions [3]. In one instance, our method shows the important topics and ideas relevant to humans before and after one of the largest Australian Earthquakes.

In closing, we wish to note that our proposed method *does not involve* updating the PageRank vector, a related problem which has received considerable attention [9,19]. Nor is it related to tensor methods for dynamic graph data [26,12].

4 PageRank with Dynamic Teleportation

In order to incorporate dynamics into PageRank, we reformulate a standard PageRank algorithm in terms of changes to the PageRank values for each page. This step allows us to state PageRank as a dynamical system, in which case we can easily incorporate changes into the vector.

The standard PageRank algorithm is the classical Richardson iteration:

$$\mathbf{x}^{(k+1)} = \alpha \mathbf{P} \mathbf{x}^{(k)} + (1 - \alpha) \mathbf{v}.$$

(Note that this iteration is identical to the power method for the PageRank Markov chain.) By rearranging this equation into a difference form, we have

$$\Delta \mathbf{x}^{(k)} = \mathbf{x}^{(k+1)} - \mathbf{x}^{(k)} = \alpha \mathbf{P} \mathbf{x}^{(k)} + (1 - \alpha) \mathbf{v} - \mathbf{x}^{(k)} = (1 - \alpha) \mathbf{v} - (\mathbf{I} - \alpha \mathbf{P}) \mathbf{x}^{(k)}.$$

Thus, changes in the PageRank values at a node *evolve* based on the value $(1 - \alpha) \mathbf{v} - (\mathbf{I} - \alpha \mathbf{P}) \mathbf{x}^{(k)}$. We reinterpret this update as a continuous time dynamical system:

$$\mathbf{x}'(t) = (1 - \alpha) \mathbf{v} - (\mathbf{I} - \alpha \mathbf{P}) \mathbf{x}(t). \tag{1}$$

Other iterative methods also give rise to related dynamical systems, as utilized by [13] for studying eigenvalue solvers.

In the dynamic teleportation model, \mathbf{v} is no longer fixed, but is instead a function of time $\mathbf{v}(t)$:

$$\mathbf{x}'(t) = (1 - \alpha) \mathbf{v}(t) - (\mathbf{I} - \alpha \mathbf{P}) \mathbf{x}(t). \tag{2}$$

Note that this means the PageRank values $\mathbf{x}(t)$ may not “settle” or converge. We see this as a feature of the new model as we plan to utilize information from the evolution and changes in the PageRank values.

Standard texts on dynamical system show that the solution $\mathbf{x}(t)$ is:

$$\mathbf{x}(t) = \exp[-(\mathbf{I} - \alpha \mathbf{P})t] \mathbf{x}(0) + (1 - \alpha) \int_0^t \exp[-(\mathbf{I} - \alpha \mathbf{P})(t - \tau)] \mathbf{v}(\tau) d\tau.$$

If $\mathbf{v}(t) = \mathbf{v}$ is constant with respect to time, then

$$\int_0^t \exp[-(\mathbf{I} - \alpha \mathbf{P})(t - \tau)] \mathbf{v}(\tau) d\tau = (\mathbf{I} - \alpha \mathbf{P})^{-1} \mathbf{v} - \exp[-(\mathbf{I} - \alpha \mathbf{P})t] (\mathbf{I} - \alpha \mathbf{P})^{-1} \mathbf{v}.$$

Hence, for constant $\mathbf{v}(t)$:

$$\mathbf{x}(t) = \exp[-(\mathbf{I} - \alpha \mathbf{P})t] (\mathbf{x}(0) - \mathbf{x}) + \mathbf{x},$$

where \mathbf{x} is the solution to static PageRank: $(\mathbf{I} - \alpha \mathbf{P}) \mathbf{x} = (1 - \alpha) \mathbf{v}$. Because all the eigenvalues of $-(\mathbf{I} - \alpha \mathbf{P}) < 0$, the matrix exponential terms disappear in a sufficiently long time horizon. Thus, when $\mathbf{v}(t) = \mathbf{v}$, nothing has changed. We recover the original PageRank vector \mathbf{x} as the steady-state solution:

$$\lim_{t \rightarrow \infty} \mathbf{x}(t) = \mathbf{x} \text{ the PageRank vector.}$$

This derivation shows that dynamic teleportation PageRank is a generalization of the PageRank vector.

Require:

- a graph $G = (V, E)$ and a procedure to compute $\mathbf{P}\mathbf{x}$ for this graph
- a maximum time t_{\max}
- a function to return $\mathbf{v}(t)$ for any $0 \leq t \leq t_{\max}$
- a damping parameter α
- a time-step h

Ensure: \mathbf{X} where the k th column of \mathbf{X} is $\mathbf{x}(0 + kh)$ for all $1 \leq k \leq t_{\max}/h$ (or any desired subset of these values)

```

t ← 0; k = 1
x(0) ← v(0) (or any other desired initial condition)
while t ≤ tmax - h do
  x(t + h) ← x(t) + h [(1 - α)v(t) - (I - αP)x(t)]
  X(:, k) ← x(t + h)
  t ← t + h; k ← k + 1
end while

```

Fig. 1. In order to compute a sequence of dynamic teleportation PageRank values, we utilize a forward Euler method for the dynamical system: $\mathbf{x}'(t) = (1 - \alpha)\mathbf{v}(t) - (\mathbf{I} - \alpha\mathbf{P})\mathbf{x}(t)$. The resulting procedure looks remarkably similar to the standard Richardson iteration to compute a PageRank vector. A key difference is that there is no notion of convergence.

4.1 Algorithms

In order to compute the time-sequence of PageRank values $\mathbf{x}(t)$, we can evolve the dynamical system (1) using any standard method, for instance a forward Euler or a Runge-Kutta method. At the moment, we only use the forward Euler method for simplicity. This method lacks high accuracy, but is fast and straightforward. Forward Euler approximates the derivative with a first order Taylor approximation:

$$\mathbf{x}'(t) \approx \frac{\mathbf{x}(t + h) - \mathbf{x}(t)}{h},$$

and then uses that approximation to estimate the value at a short time-step in the future:

$$\mathbf{x}(t + h) = \mathbf{x}(t) + h [(1 - \alpha)\mathbf{v}(t) - (\mathbf{I} - \alpha\mathbf{P})\mathbf{x}(t)].$$

Note that if $h = 1$ and $\mathbf{v}(t) = \mathbf{v}$ for all t , then this update becomes the original Richardson iteration. A summary of this derivation as a formal algorithm to compute a dynamic teleportation PageRank time series is given by Figure 1.

4.2 Discussion of the Algorithm & Practical Issues

First, the algorithm we propose easily scales to large networks. This isn't surprising given its close relationship to the Richardson method for PageRank. The major expense is the set of t_{\max}/h matrix-vector products with \mathbf{P} – all of the other work is linear in the number of nodes. It could also be used in a distributed setting if any distributed matrix-vector product is available.

In one sense, the forward Euler method is simply running a power method, but changing the vector \mathbf{v} at every iteration. However, we derived this method based on evolving (2). Thus, by studying the relationship between (2) and the algorithm in Figure 1, we can understand the underlying problem solved by changing the teleportation vector while running the power method. Consequently, we gain additional flexibility in adapting (2) to problems.

Thus far, we also have not discussed how to set $\mathbf{v}(t)$ beyond the brief allusion at the beginning that the dynamic teleportation will be based on Wikipedia pageviews. When we apply the dynamic teleportation PageRank model, we need to pick a relationship between the time-scale of the dynamical system (2) and the time-scale in the underlying application. For instance, does $\mathbf{x}(1)$ correspond to the PageRank values after a second, an hour, a day? There is no “correct” answer and the relationship has implications on the final model.

Suppose that we set $\alpha = 0.85$, $h = 1$, and that $t = 1$ is a minute of time in the application. If we have hourly data on Wikipedia pageviews, then the above algorithm will compute 60 iterations of the power-method between each hour. If we further use the incredibly simple model that $\mathbf{v}(t)$ changes each hour as we get new data, then the forward Euler method is essentially equivalent to running the power-method to convergence after \mathbf{v} changes on the hour. (They are essentially equivalent in the sense that PageRank will have converged to a 1-norm error of 10^{-4} in about 60 iterations.) If, instead, we set $\alpha = 0.85$, $h = 1$, and $t = 1$ to be 20 minutes of time in the application, then we will do 3 iterations of the power method after each hourly change.

In the preceding discussion of the algorithm, we hypothesized that $\mathbf{v}(t)$ changes at fixed intervals based on incoming data. A better idea is to smooth out these “jumps” using an exponentially weighted moving average. We plan to investigate this in the future.

4.3 Ranking from Time-Series

The above equations provide a time-series of dynamic PageRank vectors for the nodes, denoted formally as $\mathbf{x}(t)$, $0 \leq t \leq t_{\max}$. Most applications, however, want a single score, or small set of scores, to characterize the importance of a node. We now discuss a few ways in which these time series give rise to scores. Reference [23] used similar ideas to extract a single score from a time-series.

Transient Rank. We call the instantaneous values of $\mathbf{x}(t)$ a node’s *transient* rank. This score gives the importance of a node at a particular time.

Summary & Cumulative Rank. Any summary function s of the time series, such as the integral, average, minimum, maximum, variance, is a single score that encompasses the entire interval $[0, t_{\max}]$. We utilize the *cumulative rank* in the forthcoming experiments:

$$\mathbf{c} = \int_0^{t_{\max}} \mathbf{x}(t) dt \approx h\mathbf{X}\mathbf{e}.$$

Difference Rank. A node’s difference rank is the difference between its maximum and minimum rank over all time:

$$\mathbf{d} = \max_t[\mathbf{x}(t)] - \min_t[\mathbf{x}(t)].$$

Nodes with high difference rank should reflect important events that occurred within the range $[0, t_{\max}]$. The underlying intuition is that normal nodes are the pages where the Dynamic PageRanks do not change much. While the pages that have large differences in their time-series of PageRanks are topics or news that went viral or becomes popular over time. See Section 6 for more details and Figure 3 for examples such as Rihanna, PricewaterhouseCoopers, Watchmen, and American Idol (season 8).

Having a variety of different scores derived from the same data frequently helps when using these scores as features in a prediction or learning task [4,10].

4.4 Clustering the Time-Series

After applying our forward Euler based algorithm, we have sampled an approximation of this time-series: $\mathbf{X} = \{\mathbf{x}(hk) : k = 1, \dots, t_{\max}/h\}$. By *clustering* these discrete time-series, we can automatically discover patterns such as increasing or decreasing trends, periodic bursts at certain times of the year, and their ilk. Our initial experiments were promising but were omitted due to space.

5 Datasets

In both of the following experiments, we set $h = 1$, and $t = 5$ to represent one period of data – one hour for Wikipedia and one month for Twitter – so that we do 5 iterations of the forward Euler method before incorporating the new data. In each period $\mathbf{v}(t)$ is normalized to sum to 1, but is otherwise unchanged.

Wikipedia Article Graph and Hourly Pageviews. Wikipedia provides access to copies of its database [28]. We downloaded a copy of its database on March 6th, 2009 and extracted an article-by-article link graph, where an article is a page in the main Wikipedia namespace, a category page, or a portal page. All other pages and links were removed. See [15] for more information.

Wikipedia also provides hourly pageviews for each page [29]. These are the number of times a page was viewed for a given hour. These are not unique visits. We downloaded the raw page counts and matched the corresponding page counts to the pages in the Wikipedia graph. We used the page counts starting from March 6, 2009 and moving forward in time.

As an aside, let us note that vertex degrees and cumulated pageviews are uncorrelated with a correlation coefficient of 0.02, indicating that using pageviews will not reinforce any degree bias in the dynamic ranks. In fact, pages with a large number of pageviews may not have high in-degree at all, which provides evidence that pages with large in-degree are not always visited more frequently.

Twitter Social Network and Monthly Tweet Rates. We use a follower graph generated by starting with a few seed users and crawling follows links from 2008. We extract the user tweets over time from 2008 – 2009. A tweet is represented as a tuple $\langle \text{user}, \text{time}, \text{tweet} \rangle$. Using the set of tweets, we construct a sequence of vectors to represents the number of tweets for a given month.

Table 2. Dataset Properties. The pageviews or tweets is denoted as \mathbf{p} .

Dataset	Nodes	Edges	t_{\max} Period	Average p_i	Max p_i
WIKIPEDIA	4,143,840	72,718,664	20 hours	1.3225	334,650
TWITTER	465,022	835,424	6 months	0.5569	1056

6 Empirical Results

In this section, we demonstrate the effectiveness of Dynamic PageRank as a method for automatically adapting page importance based on graph structure and external influence by showing that it provides different insights (§6.1), finds interesting pages (§6.2), and helps predict pageviews (§6.3).

6.1 Ranking from Time-Series

We first use the intersection similarity measure to evaluate the rankings [6]. Given two vectors \mathbf{x} and \mathbf{y} , the intersection similarity metric at k is the average symmetric difference over the top- j sets for each $j \leq k$. If \mathcal{X}_k and \mathcal{Y}_k are the top- k sets for \mathbf{x} and \mathbf{y} , then $\text{isim}_k(\mathbf{x}, \mathbf{y}) = \frac{1}{k} \sum_{j=1}^k \frac{|\mathcal{X}_j \Delta \mathcal{Y}_j|}{2j}$, where Δ is the symmetric set-difference operation. Identical vectors have an intersection similarity of 0.

For the Wikipedia graph, Figure 2 shows the similarity profile comparing \mathbf{d} (from §4.3) to static PageRank, degree, cumulative pageviews \mathbf{p}_c , maximum pageviews difference \mathbf{p}_d , and two other Dynamic PageRank vectors: transient $\mathbf{x}(t_{\max})$ and cumulative \mathbf{c} . The figure suggests that Dynamic PageRank is different from the other measures, even for small values of k . In particular, combining the external influence with the graph appears to produce something new.

6.2 Top Dynamic Ranks

Figure 3 shows the time-series of the top 100 pages by the difference measure. Many of these pages reveal the ability of Dynamic PageRank to mesh the network structure with changes in external interest. This became immediately clear after reviewing significant events from this time period. We find pages related to an Australian earthquake (40, 72, 70), a just released movie “Watchmen” (94, 39, 99), a famous musician that died (2, 95, 68), recent “American Idol” gossip (32, 96, 56), a remembrance of Eve Carson from a contestant on “American Idol” (80, 88, 27), news about the murder of a Harry Potter actor (77), and the




Fig. 2. Intersection similarity between Dynamic PageRank’s difference ranking \mathbf{d} and the other ranking vectors. To more appropriately see the differences, we zoom in on the top 10^5 nodes. See the discussion in the text.

Skittles social media mishap (87). These results demonstrate the effectiveness of the Dynamic PageRank to identify interesting pages that pertain to external interest. The influence of the graph results in the promotion of pages such as Richter magnitude (72). That page was not in the top 200 from pageviews.

In another study, omitted due to space, we performed a clustering of these time-series to identify pages with similar trends. For instance, pages such as Watchmen (37) and Rorschach (94) share strikingly similar patterns. These patterns indicate the page that became important first and the amount of traffic or popularity that diffused over time.

6.3 Predicting Future Pageviews & Tweets

We conclude by studying how well the dynamic PageRank values *predict* future pageviews. Formally, given a lagged time-series [2], the goal is to predict the future value \mathbf{p}_{t+1} (actual pageviews or number of tweets). This type of temporal prediction task has many applications, such as actively adapting caches in large database systems, or dynamically recommending pages.

We performed one-step ahead predictions ($t + 1$) using linear regression. That is, we learn a model of the form:

$$[\bar{\mathbf{f}}(t - 1; \theta) \bar{\mathbf{f}}(t - 2; \theta) \dots \bar{\mathbf{f}}(t - w; \theta)] \mathbf{b} \approx \mathbf{p}(t)$$

where w is the window-size, and $\bar{\mathbf{f}}(\cdot; \theta)$ is an exponentially damped moving average computed from either pageviews, dynamic PageRanks, or both. Using this

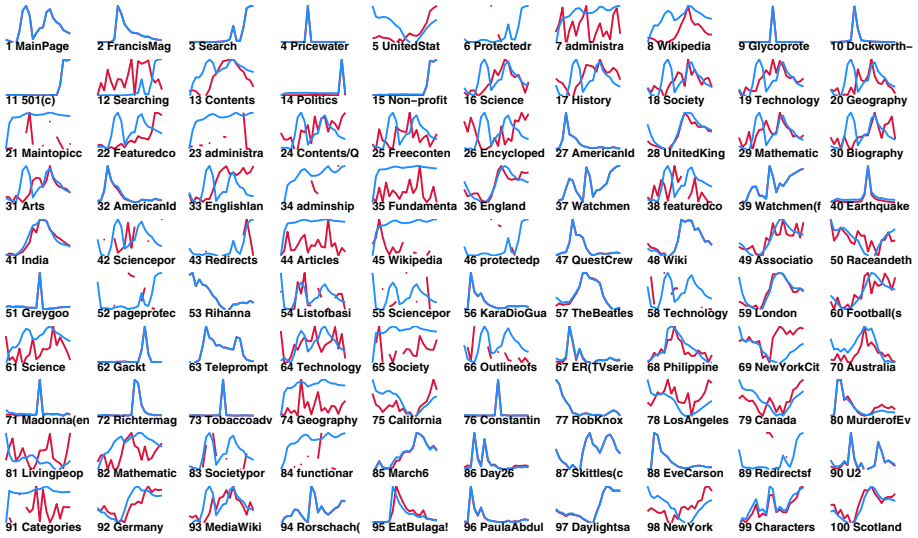


Fig. 3. The top-100 Wikipedia pages that fluctuate the most as determined by the difference ranking from our Dynamic PageRank approach. The x-axis represents time (in hours) while the y-axis represents the Dynamic PageRank value. The blue line represents Dynamic PageRank and the red line represents the hourly pageviews. There exist many interesting time-series patterns such as spikes (40), cyclic/seasonality trends (16-20), and increasing/decreasing trends (39 and 77), among many others. Further analysis and anecdotal evidence was removed due to space.

average is a standard forecasting technique. Specifically, the exponentially damped moving average of a time-series feature $\mathbf{f}(t)$ is:

$$\bar{\mathbf{f}}(t; \theta) = \underbrace{\theta \mathbf{f}(t)}_{\text{new data}} + \underbrace{(1 - \theta) \bar{\mathbf{f}}(t - 1; \theta)}_{\text{old data}}.$$

The exponential factor was $\theta = 0.3$ for Twitter and $\theta = 0.7$ for Wikipedia. Due to the scarcity of the data, we used 0.3 for Twitter since this choice weights past observations more heavily. In the future, we plan to use cross-validation. After fitting, the model predicts $\mathbf{p}(t + 1)$ as $[\mathbf{f}(t; \theta) \bar{\mathbf{f}}(t - 1; \theta) \dots \bar{\mathbf{f}}(t - w + 1; \theta)] \mathbf{b}$. To measure the error, we use symmetric Mean Absolute Percentage Error (or sMAPE) [2].

We study two models.

Base Model. This model uses only the time-series of pageviews or tweet-rates to predict the future pageviews or number of tweets.

Dynamic PageRank Model. This model uses both the Dynamic PageRank time-series and pageviews to predict the future pageviews.

We evaluate these models for prediction on *stationary* and *non-stationary* time-series. Informally, a time-series is weakly stationary if it has properties

(mean and covariance) similar to that of the time-shifted time-series. We consider the top and bottom 1000 nodes from the difference ranking as nodes that are approximately non-stationary (volatile) and stationary (stable), respectively. Table 3 compares the predictions of the models across time for non-stationary and stationary prediction tasks. Our findings indicate that the Dynamic PageRank time-series provides valuable information for forecasting future pageviews.

Table 3. Average SMAPE over all nodes for the two models (lower is better). We also measure the performance of the models for predicting highly volatile nodes (non-stationary) and nodes with relatively stable behavior (stationary). In all cases, the Dynamic PageRank model is more accurate than the base model.

Dataset	Forecasting	Dynamic PageRank	Base Model
WIKIPEDIA	<i>Non-stationary</i>	0.4349	0.5028
	<i>Stationary</i>	0.3672	0.4373
TWITTER	<i>Non-stationary</i>	0.4852	1.2333
	<i>Stationary</i>	0.6690	0.9180

7 Conclusion

We proposed an evolving teleportation adaptation of the PageRank method to capture how changes in external interest influence the importance of a node. This proposal lets us treat PageRank as a dynamical system and seamlessly incorporate changes in the teleportation vector. Furthermore, we demonstrated the utility of using Dynamic PageRank for predicting pageviews. In future work, we hope to include dynamic and evolving graphs into this framework as well.

References

1. Abiteboul, S., Preda, M., Cobena, G.: Adaptive on-line page importance computation. In: WWW, pp. 280–290. ACM (2003)
2. Ahmed, N., Atiya, A., El Gayar, N., El-Shishiny, H.: An empirical comparison of machine learning models for time series forecasting. *Econ. Rev.* 29(5-6), 594–621 (2010)
3. Bagrow, J., Wang, D., Barabási, A.: Collective response of human populations to large-scale emergencies. *PloS one* 6(3), e17680 (2011)
4. Becchetti, L., Castillo, C., Donato, D., Baeza-Yates, R., Leonardi, S.: Link analysis for web spam detection. *ACM Trans. Web* 2(1), 1–42 (2008)
5. Bianchini, M., Gori, M., Scarselli, F.: Inside PageRank. *ACM Transactions on Internet Technologies* 5(1), 92–128 (2005)
6. Boldi, P.: TotalRank: Ranking without damping. In: WWW, pp. 898–899 (2005)
7. Boldi, P., Santini, M., Vigna, S.: Paradoxical effects in PageRank incremental computations. *Internet Mathematics* 2(2), 387–404 (2005)
8. Bonacich, P.: Power and centrality: A family of measures. *American Journal of Sociology*, 1170–1182 (1987)

9. Chien, S., Dwork, C., Kumar, R., Simon, D., Sivakumar, D.: Link evolution: Analysis and algorithms. *Internet Mathematics* 1(3), 277–304 (2004)
10. Constantine, P., Gleich, D.: Random alpha PageRank. *Internet Mathematics* 6(2), 189–236 (2009)
11. Das Sarma, A., Gollapudi, S., Panigrahy, R.: Estimating PageRank on graph streams. In: SIGMOD, pp. 69–78. ACM (2008)
12. Dunlavy, D.M., Kolda, T.G., Acar, E.: Temporal link prediction using matrix and tensor factorizations. *TKDD* 5(2), 10:1–10:27 (2011)
13. Embree, M., Lehoucq, R.B.: Dynamical systems and non-hermitian iterative eigen-solvers. *SIAM Journal on Numerical Analysis* 47(2), 1445–1473 (2009)
14. Freeman, L.: Centrality in social networks conceptual clarification. *Social Networks* 1(3), 215–239 (1979)
15. Gleich, D., Glynn, P., Golub, G., Greif, C.: Three results on the PageRank vector: eigenstructure, sensitivity, and the derivative. In: *Web Information Retrieval and Linear Algebra Algorithms* (2007)
16. Grindrod, P., Parsons, M., Higham, D., Estrada, E.: Communicability across evolving networks. *Physical Review E* 83(4), 046120 (2011)
17. Katz, L.: A new status index derived from sociometric analysis. *Psychometrika* 18(1), 39–43 (1953)
18. Kleinberg, J.: Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)* 46(5), 604–632 (1999)
19. Langville, A.N., Meyer, C.D.: Updating PageRank with iterative aggregation. In: *WWW*, pp. 392–393 (2004)
20. Langville, A.N., Meyer, C.D.: *Google’s PageRank and Beyond: The Science of Search Engine Rankings*. Princeton University Press (2006)
21. Mathieu, F., Bouklit, M.: The effect of the back button in a random walk: application for PageRank. In: *WWW*, pp. 370–371 (2004)
22. Morrison, J.L., Breitling, R., Higham, D.J., Gilbert, D.R.: GeneRank: using search engine technology for the analysis of microarray experiments. *BMC Bioinformatics* 6(1), 233 (2005)
23. O’Madadhain, J., Smyth, P.: Eventrank: A framework for ranking time-varying networks. In: *LinkKDD*, pp. 9–16. ACM (2005)
24. Page, L., Brin, S., Motwani, R., Winograd, T.: *The PageRank citation ranking: Bringing order to the web* (1998)
25. Ratkiewicz, J., Fortunato, S., Flammini, A., Menczer, F., Vespignani, A.: Characterizing and modeling the dynamics of online popularity. *Physical Review Letters* 105(15), 158701 (2010)
26. Sun, J., Tao, D., Faloutsos, C.: Beyond streams and graphs: dynamic tensor analysis. In: *SIGKDD, KDD 2006*, pp. 374–383. ACM, New York (2006)
27. Suzuki, Y., et al.: Identification and characterization of the potential promoter regions of 1031 kinds of human genes. *Genome Research* 11(5), 677–684 (2001)
28. Various. *Wikipedia database dump*, Version from (March 6, 2009), http://en.wikipedia.org/wiki/Wikipedia:Database_download
29. Various. *Wikipedia pageviews* (2011), <http://dumps.wikimedia.org/other/pagecounts-raw/> (accessed in 2011)