

The Push Algorithm for Spectral Ranking

Paolo Boldi Sebastiano Vigna

March 18, 2014

Abstract

The push algorithm was proposed first by Jeh and Widom [6] in the context of personalized PageRank computations (albeit the name “push algorithm” was actually used by Andersen, Chung and Lang in a subsequent paper [1]). In this note we describe the algorithm at a level of generality that make the computation of the spectral ranking of any nonnegative matrix possible. Actually, the main contribution of this note is that the description is very simple (almost trivial), and it requires only a few elementary linear-algebra computations. Along the way, we give new precise ways of estimating the convergence of the algorithm, and describe some of the contribution of the existing literature, which again turn out to be immediate when recast in our framework.

1 Introduction

Let M be a $n \times n$ nonnegative real matrix with entries m_{xy} . Without loss of generality, we assume that $\|M\|_1 = 1$; this means that M is *substochastic* (i.e., its row sums are at most one) and that *at least one row has sum one*.¹

Equivalently, we can think of the arc-weighted graph G underlying M . The graph has n nodes, and an arc $x \rightarrow y$ weighted by m_{xy} if $m_{xy} > 0$. We will frequently switch between the matrix and the graph view, as linear matters are better discussed in terms of M , but the algorithms we are interested in are more easily discussed through G .

As a guiding example, given a directed graph G with n nodes, M can be the transition matrix of its *natural walk*², whose weights are $m_{xy} = 1/d^+(x)$, where $d^+(x)$ is the outdegree of x (the number of arcs going out of x).

We recall that the spectral radius $\rho(M)$ of M coincides with the largest (in modulus) of its eigenvalues, and satisfies³

$$\min_i \|M_i\|_1 \leq \rho(M) \leq \max_i \|M_i\|_1 = \|M\|_1 = 1,$$

¹If this is not the case, just multiply the matrix by the inverse of the maximum row sum. The multiplication does not affect the eigenspaces, but now the matrix satisfies the conditions above. Of course, the values of the damping factor (see further on) have to be adjusted accordingly.

²We make no assumptions on G , so some nodes might be dangling (i.e., without successors). In that case the corresponding rows of M will be zeroed, so M would be neither stochastic, nor a random walk in a strictly technical sense.

³We use row vectors, so the ℓ_1 norm of a matrix is the maximum of the norm of the rows.

where M_i is the i -th row of M (the second inequality is always true; the first one only for nonnegative matrices).

Let \mathbf{v} be a nonnegative vector satisfying $\|\mathbf{v}\|_1 = 1$ (i.e., a distribution) and $\alpha \in [0..1)$. The *spectral ranking*⁴ of M with preference vector \mathbf{v} and damping factor α is defined by

$$\mathbf{r} = (1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1} = (1 - \alpha)\mathbf{v} \sum_{k \geq 0} \alpha^k M^k.$$

Note that the \mathbf{r} needs not be a distribution, unless M is stochastic.⁵ Note also that the linear operator is defined for $\alpha \in [0..1/\rho(M))$, but usually estimating $\rho(M)$ is very difficult. The value $1/\rho(M)$ can actually be attained by a limiting process which essentially makes the damping disappear [8].

We start from the following trivial observation: while it is very difficult to “guess” which is the spectral ranking \mathbf{r} associated to a certain \mathbf{v} , the inverse problem is trivial: given \mathbf{r} ,

$$\mathbf{v} = \frac{1}{1 - \alpha} \mathbf{r} (1 - \alpha M).$$

The resulting preference vector \mathbf{v} might not be, of course, a distribution (otherwise we could obtain *any* spectral ranking using a suitable preference vector), but the equation is always true.

The observation is trivial, but its consequences are not. For instance, consider an indicator vector $\chi_x(z) = [x = z]$. If we want to obtain $(1 - \alpha)\chi_x$ as spectral ranking, the associated preference vector \mathbf{v} has a particularly simple form:

$$\mathbf{v} = \frac{1}{1 - \alpha} (1 - \alpha)\chi_x(1 - \alpha M) = \chi_x - \alpha \sum_{x \rightarrow y} m_{xy} \chi_y. \quad (1)$$

We remark that in the case of a natural random walk, $m_{xy} = d(x)^{-1}$, which does not depend on y and can be taken out of the summation. Of course, since spectral rankings are linear we can obtain $(1 - \alpha)\chi_x$ multiplied by any constant just by multiplying \mathbf{v} by the same constant.

2 The push algorithm

If the preference vector \mathbf{v} is highly concentrated (e.g., an indicator) and α is not too close to one most updates done by linear solvers or iterative methods to compute spectral rankings are useless—either they do not perform any update, or they update nodes whose final value will end up to be below the computational precision.

⁴“Spectral ranking” is an umbrella name for techniques based on eigenvectors and linear maps to rank entities; see [8] for a detailed history of the subject, which was studied already in the late forties.

⁵If M is the natural walk on a graph, \mathbf{r} is not exactly PageRank [7], but rather the *pseudorank* [5] associated with \mathbf{v} and α . The pseudorank is not necessarily a distribution, whereas technically a PageRank vector always is. The distinction is however somehow blurred in the literature, where often pseudoranks are used in place of PageRank vectors. If G has no dangling nodes, the pseudorank is exactly PageRank. Otherwise, there are some differences depending on how dangling nodes are patched [4].

The *push algorithm* uses the concentration of modifications to reduce the computational burden. The fundamental idea appeared first in Jeh and Widom’s widely quoted paper [6], albeit the notation somehow obscures the ideas. Berkhin restated the algorithm in a different and more readable form [2]. Andersen, Chung and Lang [1] applied a specialised version of the algorithm on symmetric graphs. All these references apply the idea to PageRank, but the algorithm is actually an algorithm for the steady state of Markov chains with restart [3], and it works even with substochastic matrices, so it should be thought of as an algorithm for spectral ranking with damping.⁶

The basic idea is that of keeping track of vectors \mathbf{p} (the current approximation) and \mathbf{r} (the *residual*) satisfying

$$\mathbf{p} + (1 - \alpha)\mathbf{r}(1 - \alpha M)^{-1} = (1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1}$$

Initially, $\mathbf{p} = 0$ and $\mathbf{r} = \mathbf{v}$, which makes the statement trivial, but we will incrementally increase \mathbf{p} (and reduce correspondingly \mathbf{r}).

To this purpose, we will be iteratively *pushing*⁷ some node x . A push on x adds $(1 - \alpha)r_x\chi_x$ to \mathbf{p} . Since we must keep the invariant true, we now have to update \mathbf{r} . If we think of \mathbf{r} as a preference vector, we are just trying to solve the inverse problem (1): by linearity, if we subtract from \mathbf{r}

$$r_x \left(\chi_x - \alpha \sum_{x \rightarrow y} m_{xy} \chi_y \right),$$

the value $(1 - \alpha)\mathbf{r}(1 - \alpha M)^{-1}$ will decrease exactly by $(1 - \alpha)r_x\chi_x$, preserving the invariant.

It is not difficult to see why this choice is good: we zero an entry (the x -th) of \mathbf{r} , and we add small positive quantities to a small (if the graph is sparse) set of entries (those associated with the successors of x), increasing the ℓ_1 norm of \mathbf{p} by $(1 - \alpha)r_x$, and decreasing at least by the same amount that of \mathbf{r} (larger decreases happening on strictly substochastic rows—e.g., dangling nodes). Note that since we do not create negative entries, it is always true that

$$\|\mathbf{p}\|_1 + \|\mathbf{r}\|_1 \leq 1.$$

Of course, we can easily keep track of the two norms at each update.

The error in the estimate is

$$\begin{aligned} \|(1 - \alpha)\mathbf{r}(1 - \alpha M)^{-1}\|_1 &= \\ (1 - \alpha) \left\| \mathbf{r} \sum_{k \geq 0} \alpha^k M^k \right\|_1 &\leq (1 - \alpha) \|\mathbf{r}\|_1 \sum_{k \geq 0} \alpha^k \|M^k\|_1 \leq \|\mathbf{r}\|_1. \end{aligned}$$

Thus, we can control exactly the *absolute additive error* of the algorithm by controlling the ℓ_1 norm of the residual.

⁶An implementation of the push algorithm for the computation of PageRank is available as part of the LAW software at <http://law.dsi.unimi.it/>.

⁷The name is taken from [1]—we find it enlightening.

It is important to notice that if M is strictly substochastic it might happen that

$$\|(1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1}\|_1 < 1.$$

If this happens, controlling the ℓ_1 norm of the residual is actually of little help, as even in the case of natural walks the norm above can be as small as $1 - \alpha$. However, since we have the guarantee that \mathbf{p} is a nonnegative vector which approximates the spectral ranking *from below*, we can simply use

$$\frac{\|\mathbf{r}\|_1}{\|\mathbf{p}\|_1} \geq \frac{\|\mathbf{r}\|_1}{\|(1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1}\|_1}$$

as a measure of relative precision, as

$$\frac{\|(1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1} - \mathbf{p}\|_1}{\|(1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1}\|_1} = \frac{\|(1 - \alpha)\mathbf{r}(1 - \alpha M)^{-1}\|_1}{\|(1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1}\|_1} \leq \frac{\|\mathbf{r}\|_1}{\|\mathbf{p}\|_1}.$$

2.1 Handling pushes

The order in which pushes are executed can be established in many different ways. Certainly, to guarantee relative error ε we need only push nodes v such that $r_x > \varepsilon\|\mathbf{p}\|_1/n$, as if all nodes fail to satisfy the inequality then $\|\mathbf{r}\|_1/\|\mathbf{p}\|_1 \leq \varepsilon$.

The obvious approach is that of keeping an indirect priority queue (i.e., a queue in which the priority of every element can be updated at any time) containing the nodes satisfying the criterion above (initially, just the support of \mathbf{v}) and returning them in order of decreasing r_x . Nodes are added to the queue when their residual is larger than $\varepsilon\|\mathbf{p}\|_1/n$. Every time a push is performed, the residual of successors of the pushed node are updated and the queue is notified of the changes.

While this generates potentially an $O(\log n)$ cost per arc visited (to adjust the queue), in intended applications the queue is always very small, and pushing larger values leads to a faster decrease of $\|\mathbf{r}\|_1$.

An alternative approach is to use a FIFO queue (with the proviso that nodes already in the queue are not enqueued again). In this case, pushes are not necessarily executed in the best possible order, but the queue has constant-time access.

Some preliminary experiments show that the two approaches are complementary, in the sense that in situations where the number of nodes in the queue is relatively small, a priority code reduces significantly the number of pushes, resulting in a faster computation. However, if the queue becomes large (e.g., because the damping factor is close to one), the logarithmic burden at each modification becomes tangible, and using a FIFO queue yields a faster computation in spite of the higher number of pushes.

In any case, to reduce the memory footprint for large graphs it is essential to keep track of the bijection between the set of visited nodes and an identifier assigned incrementally in discovery order. In this way, all vectors involved in the computation can be indexed by discovery order, making their size dependent just on the size of the visited neighbourhood, and not on the size of the graph.

2.2 Convergence

There are no published results of convergence for the push algorithm. Andersen, Chung and Lang provide a bound not for convergence to the pseudorank, but rather for convergence to the ratio between the pseudorank and the stationary state of M (which in their case—symmetric graphs—is trivial, as it is proportional to the degree).

In case a priority queue is used to select the nodes to be pushed, when the preference vector is an indicator χ_x the amount of rank going to \mathbf{p} at the first step is exactly $1 - \alpha$. In the follow $d(x)$ steps, we will visit either the successors of x , whose residual is $\alpha/d(x)$, or some node with a larger residual, due to the prioritization in the queue. As a result, the amount of rank going to \mathbf{p} will be at least $\alpha(1 - \alpha)$. In general, if $P_x(t)$ is the *path function* of x (i.e., $P_x(t)$ is the number of paths of length at most t starting from x), after $P_x(t)$ pushes the ℓ_1 norm of \mathbf{r} will be at most $1 - (1 - \alpha) \sum_{0 \leq k \leq t} \alpha^k = \alpha^{t+1}$.

2.3 Some remarks

Precomputing spectral rankings. Another interesting remark⁸ is that if during the computation we have to perform a push on a node x and we happen to know the *spectral ranking of x* (i.e., the spectral ranking with preference vector χ_x) we can simply zero r_x and add the spectral ranking of x multiplied by the current value of r_x to \mathbf{p} . Actually, we could even *never push x* and just add the spectral ranking of x multiplied by r_x to \mathbf{p} at the end of the computation.

Let us try to make this observation more general. Consider a set H of vertices whose spectral ranking is known; in other words, for each $x \in H$ the vector

$$s_x = (1 - \alpha)\chi_x(1 - \alpha M)^{-1}$$

is somehow available. At every step of the algorithm, the invariant equation

$$\mathbf{p} + (1 - \alpha)\mathbf{r}(1 - \alpha M)^{-1} = (1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1}$$

can be rewritten as follows: let \mathbf{r}' be the vector obtained from \mathbf{r} after zeroing all entries outside of H , and let $\mathbf{p}' = \mathbf{p} + \sum_{x \in H} r_x s_x$. Then clearly

$$\mathbf{p}' + (1 - \alpha)\mathbf{r}'(1 - \alpha M)^{-1} = (1 - \alpha)\mathbf{v}(1 - \alpha M)^{-1}.$$

Note that

$$\|\mathbf{r}'\|_1 = \sum_{x \notin H} r_x$$

and

$$\|\mathbf{p}'\|_1 = \|\mathbf{p}\|_1 + \sum_{x \in H} r_x \cdot \|s_x\|_1.$$

So we can actually execute the push algorithm keeping track of \mathbf{p} and \mathbf{r} but considering (virtually) to possess \mathbf{p}' and \mathbf{r}' , instead; to this aim, we proceed as follows:

⁸Actually, a translation of Jeh and Widom's approach based on *partial vectors*, which was restated by Berkhin's under the name *hub decompositions* [2]. Both become immediate in our setting.

- we never add nodes in H to the queue;
- for convergence, we consider the norms of \mathbf{p}' and \mathbf{r}' , as computed above;
- at termination, we adjust \mathbf{p} obtaining \mathbf{p}' explicitly.

Berkhin [2] notes that when computing the spectral ranking of x we can use x as a hub after the first push. That is, after the first push we will never enqueue x again. At the end of the computation, we simply multiply the resulting spectral ranking by $1 + r_x + r_x^2 + \dots = 1/(1 - r_x)$. In this case, $\|\mathbf{p}\|_1$ must be divided by $1 - r_x$ to have a better estimate of the actual norm. Preliminary experiments on web and social graphs show that the reduction of the number of pushes is very marginal, though.

Patching dangling nodes. Suppose that, analogously to what is usually done in power-method computations, we may patch dangling nodes. More precisely, suppose that we start from a matrix M that has some zero rows (e.g., the natural walk of a graph G with dangling nodes), and then we obtain a new matrix P (for “patched”) by *substituting* each zero row with some distribution \mathbf{u} , as yet unspecified.

It is known that avoiding at all the patch is equivalent to using $\mathbf{u} = \mathbf{v}$ [5], modulo a scale factor that is computable starting from the spectral ranking itself. More generally, if \mathbf{u} coincides with the distribution that is being used for preference, no patching is needed provided that the final result is normalized.

For the general case (where \mathbf{u} may not coincide with \mathbf{v}), we can adapt the push method described above as follows: we keep track of vectors \mathbf{p} and \mathbf{r} and of a scalar θ representing the amount of rank that went through dangling nodes. The equation now is

$$\mathbf{p} + (1 - \alpha)(\mathbf{r} + \theta\mathbf{u})(1 - \alpha P)^{-1} = (1 - \alpha)\chi_x(1 - \alpha P)^{-1}$$

When \mathbf{p} is increased by $(1 - \alpha)r_x\chi_x$, we have to modify \mathbf{r} and θ as follows:

- if x is not dangling, we subtract from \mathbf{r} the vector

$$r_x \left(\chi_x - \alpha \sum_{x \rightarrow y} m_{xy} \chi_y \right);$$

- if x is dangling, we subtract just

$$r_x \chi_x$$

and *increase* θ by αr_x .

At every computation step the approximation of the spectral ranking will be given by $\mathbf{p}' = \mathbf{p} + \theta\mathbf{s}$, where \mathbf{s} is the spectral ranking of P with preference vector \mathbf{u} and damping factor α .⁹ As on the case of hubs, we should consider $\|\mathbf{p}'\|_1 = \|\mathbf{p}\|_1 + \theta\|\mathbf{s}\|_1$ when establishing convergence.

⁹Of course, \mathbf{s} must be precomputed using any standard method. If M is the natural walk of a graph G , this is exactly the PageRank vector for G with preference vector \mathbf{u} .

References

- [1] Reid Andersen, Fan R. K. Chung, and Kevin J. Lang. Using PageRank to locally partition a graph. *Internet Math.*, 4(1):35–64, 2007.
- [2] Pavel Berkhin. Bookmark-coloring approach to personalized PageRank computing. *Internet Math.*, 3(1):41–62, 2006.
- [3] Paolo Boldi, Violetta Lonati, Massimo Santini, and Sebastiano Vigna. Graph fibrations, graph isomorphism, and PageRank. *RAIRO Inform. Théor.*, 40:227–253, 2006.
- [4] Paolo Boldi, Roberto Posenato, Massimo Santini, and Sebastiano Vigna. Traps and pitfalls of topic-biased PageRank. In William Aiello, Andrei Broder, Jeannette Janssen, and Evangelos Milios, editors, *WAW 2006. Fourth Workshop on Algorithms and Models for the Web-Graph*, volume 4936 of *Lecture Notes in Computer Science*, pages 107–116. Springer–Verlag, 2008.
- [5] Paolo Boldi, Massimo Santini, and Sebastiano Vigna. PageRank: Functional dependencies. *ACM Trans. Inf. Sys.*, 27(4):1–23, 2009.
- [6] Glen Jeh and Jennifer Widom. Scaling personalized web search. In *Proc. of the Twelfth International World Wide Web Conference*, pages 271–279. ACM Press, 2003.
- [7] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The PageRank citation ranking: Bringing order to the web. Technical Report SIDL-WP-1999-0120, Stanford Digital Library Technologies Project, Stanford University, 1998.
- [8] Sebastiano Vigna. Spectral ranking. *CoRR*, abs/0912.0238, 2009.