University of Maryland College Park

Institute for Advanced Computer Studies TR–2003–65
Department of Computer Science TR–4496

On the Method of Pure Truncation for Approximating Stationary Probabilities*

G. W. Stewart†

June, 2003

ABSTRACT

This paper is concerned with approximating the leading components of the stationary vector of a semi-infinite discrete markov chain. The most widely treated method extracts a leading principal submatrix from the matrix of transition probabilities, adjusts its elements so that it becomes stochastic, and takes the stationary vector of the result as the approximation. In this paper, the consequences of taking the normalized Perron vector of the unadjusted matrix as the approximation are explored. Error bounds are derived, and it is shown that the adjusted and unadjusted methods are approximations to one another.

*This report is available by anonymous ftp from `thales.cs.umd.edu` in the directory `pub/reports` or on the web at `http://www.cs.umd.edu/~stewart/`.

†Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (`stewart@cs.umd.edu`). This work was supported in part by the National Science Foundation under Grant CCR0204084.

On the Method of Pure Truncation for Approximating Stationary Probabilities

G. W. Stewart

ABSTRACT

This paper is concerned with approximating the leading components of the stationary vector of a semi-infinite discrete markov chain. The most widely treated method extracts a leading principal submatrix from the matrix of transition probabilities, adjusts its elements so that it becomes stochastic, and takes the stationary vector of the result as the approximation. In this paper, the consequences of taking the normalized Perron vector of the unadjusted matrix as the approximation are explored. Error bounds are derived, and it is shown that the adjusted and unadjusted methods are approximations to one another.

1. Introduction

This paper is concerned with approximation the stationary vector of an infinite Markov chain. Specifically, let P be the transition matrix for a semi-infinite discrete Markov chain having a stationary vector $\boldsymbol{\pi}^{\mathrm{T}}$. Let

$$Q = I - P,$$

so that $\boldsymbol{\pi}^{\mathrm{T}} Q = 0$ and $Q\mathbf{e} = 0$, where $\mathbf{e}$ is the vector consisting of all ones. Unless we can determine the components of $\boldsymbol{\pi}^{\mathrm{T}}$ analytically, we must compute a finite subset of them numerically. Without loss of generality, we can assume that the required components occur at the start of the vector $\boldsymbol{\pi}$. Thus we partition

$$\boldsymbol{\pi}^{\mathrm{T}} = (\boldsymbol{\pi}_n^{\mathrm{T}} \;\; \boldsymbol{\pi}_\bullet^{\mathrm{T}}),$$

where $\boldsymbol{\pi}_n$ has n components, and attempt to approximate some of the components of $\boldsymbol{\pi}_n$. In general, we shall have to take n larger than the number of components we actually require, since the small components in the tail of $\boldsymbol{\pi}_n$ will tend to be less accurate than their larger brethren [see (3.7) and (3.8) below].

Most numerical methods begin by partitioning Q in the form

$$Q = \begin{pmatrix} Q_{nn} & Q_{n\bullet} \\ Q_{\bullet n} & Q_{\bullet\bullet} \end{pmatrix},$$

where Q_{nn} is of order n and attempt to determine an approximation to $\boldsymbol{\pi}_n^{\mathrm{T}}$ from Q_{nn}. The most popular of these is what we will call the method of adjusted truncation.

1. Adjust the elements of Q_{nn} to obtain a matrix $\hat{Q}_{nn}$ such that $I - \hat{Q}_{nn}$ is stochastic and irreducible.
2. Approximate $\boldsymbol{\pi}_n^{\mathrm{T}}$ by the left Perron vector $\hat{\boldsymbol{\pi}}_n^{\mathrm{T}}$ of $I - Q_{nn}$, normalized so that $\hat{\boldsymbol{\pi}}_n^{\mathrm{T}} = 1$.

(1.1)

By the Perron vector we mean the positive eigenvector corresponding to the Perron root of $I - \hat{Q}_{nn}$, in this case 1. (For the theoretical background in nonnegative matrices see [1, 5].)

This method — or rather class of methods — has been treated extensively in the literature, and we will return to it later after we have analyzed the method to be proposed here. For now it is sufficient to note that it has two problems. First, the adjustment in step 2 can be performed in an infinite number of ways, and it is not clear which one will be most effective. Second, if $Q_{n\bullet}$ is substantial, as it often is in practice, the adjustment will not be small, which makes the choice of adjustment all the more critical.

In this paper we will investigate the alternative of omitting step 1 in (1.1). This leads to the following algorithm

1. Approximate $\boldsymbol{\pi}_n^{\mathrm{T}}$ by the left Perron vector $\hat{\boldsymbol{\pi}}_n^{\mathrm{T}}$ of $I - Q_{nn}$, normalized so that $\hat{\boldsymbol{\pi}}_n^{\mathrm{T}}\mathbf{e} = 1$.

(1.2)

To insure that the Perron vector exists and is unique we will assume here and in what follows that

the matrix Q_{nn} is irreducible.

The vector $\hat{\boldsymbol{\pi}}_n^{\mathrm{T}}$ can also be characterized as the normalized left eigenvector corresponding to the smallest eigenvalue of Q_{nn}.

We are goning to establish conditions under which the approximation $\hat{\boldsymbol{\pi}}_n$ is accurate in proportion as $\boldsymbol{\pi}_\bullet$ is small. We will do this by exhibiting a perturbation $\tilde{Q}_{nn}$ of Q_{nn} such that

1. $\boldsymbol{\pi}_n^{\mathrm{T}}\tilde{Q}_{11} = 0$,
2. $\|\tilde{Q}_{nn} - Q_{nn}\|_2 \leq \theta\|\boldsymbol{\pi}_\bullet\|_1$

for some θ that depends on $\boldsymbol{\pi}$ but not on n. When $\boldsymbol{\pi}_\bullet$ is small, Q_{nn} and $\tilde{Q}_{nn}$ are near, and we can use perturbation theory to bound the distance between $\boldsymbol{\pi}_n$ and $\hat{\boldsymbol{\pi}}_n$.

There are three things to note about this approach. First, because the components of $\boldsymbol{\pi}$ sum to one, as n increases $\boldsymbol{\pi}_\bullet$ must ultimately become small. Although it is easy to construct chains in which the decrease in the components of $\boldsymbol{\pi}$ is slow, in many chains the components of $\boldsymbol{\pi}$ rapidly approach zero. Second, the analysis takes place in the finite dimensional space $\mathbb{R}^n$. Only the size of the vector $\boldsymbol{\pi}_\bullet$ must be inferred from the infinite chain. Finally, although our bounds are useful in providing insight into the pure truncation method, they alone are not sufficient to establish the convergence of

the method as $n \to \infty$. We will return to this point after we have established the our main results.

This paper is organized as follows. In the next section we will introduce the perturbation of Q_{nn} and estimate its size. In Section 3 we will assess the accuracy of the approximation $\hat{\boldsymbol{\pi}}_n$ and comment on the meaning of the quantities appearing in in our error bounds. We then give a simple example illustrating the pure truncation method. The paper concludes with a discussion of the relation between the methods of pure and adjusted truncation.

In what follows, $\|\cdot\|_1$ and $\|\cdot\|_2$ will denote the vector and matrix 1- and 2-norms. For more on these norms see [3, 6].

2. The perturbation

In this section we will construct a perturbed matrix

$$\tilde{Q}_{nn} = Q_{nn} + E$$

such that $\boldsymbol{\pi}_n^{\mathrm{T}} \tilde{Q}_{nn} = 0$, and we will give bounds on the perturbing matrix E in terms of $\boldsymbol{\pi}_\bullet^{\mathrm{T}}$. We begin by setting

$$r^{\mathrm{T}} = \boldsymbol{\pi}_n^{\mathrm{T}} Q_{nn}.$$

If we define E by

$$E = -\frac{\boldsymbol{\pi}_n r^{\mathrm{T}}}{\|\boldsymbol{\pi}_n\|_2^2},$$

then it is easily verified that $\boldsymbol{\pi}_n^{\mathrm{T}} \tilde{Q}_{nn} = 0$.

To bound $\|E\|_2$, note that from the relation

$$(\boldsymbol{\pi}_n^{\mathrm{T}} \;\; \boldsymbol{\pi}_\bullet^{\mathrm{T}}) \begin{pmatrix} Q_{nn} & Q_{n\bullet} \\ Q_{\bullet n} & Q_{\bullet\bullet} \end{pmatrix} = (0 \;\; 0),$$

we get

$$r^{\mathrm{T}} = \boldsymbol{\pi}_n^{\mathrm{T}} Q_{nn} = -\boldsymbol{\pi}_\bullet^{\mathrm{T}} Q_{\bullet n}.$$

Hence

$$\|r\|_2 = \|\boldsymbol{\pi}_\bullet^{\mathrm{T}} Q_{\bullet n}\|_2 \leq \|Q_{\bullet n}^{\mathrm{T}} \boldsymbol{\pi}_\bullet\|_1 \leq \|Q_{\bullet n}^{\mathrm{T}}\|_1 \|\boldsymbol{\pi}_\bullet\|_1 \leq \|\boldsymbol{\pi}_\bullet\|_1,$$

the third inequality following from the fact that the absolute row sums of $Q_{\bullet n}$ are not greater than one. If we now set

$$\theta_n = \frac{\|\boldsymbol{\pi}_n\|_1}{\|\boldsymbol{\pi}_n\|_2},$$

then

$$\epsilon_n \stackrel{\text{def}}{=} \|E\|_2 \leq \theta_n \frac{\|\boldsymbol{\pi}_\bullet\|_1}{\|\boldsymbol{\pi}_n\|_1} = \theta_n \frac{\|\boldsymbol{\pi}_\bullet\|_1}{1 - \|\boldsymbol{\pi}_\bullet\|_1} \tag{2.1}$$

Strictly speaking (2.1) does not insure that E becomes small as n increases, since it is possible that θ_n increases as $\|\boldsymbol{\pi}_\bullet\|_1$ decreases. In general, we can only say that $\theta_n \leq \sqrt{n}$, and this bound can be attained (when $\boldsymbol{\pi}_n = \alpha\mathbf{e}$, for some $\alpha > 0$). However, for fixed Q, we have

$$\lim_{n\to\infty} \theta_n = \frac{1}{\|\boldsymbol{\pi}_n\|_2} \equiv \theta_\infty. \tag{2.2}$$

Hence in this case θ_n is uniformly bounded for all n.

We can even say more. We have

$$\begin{aligned} \theta_n &\leq 1/\|\boldsymbol{\pi}_n\|_2 \\ &= 1/\sqrt{\|\boldsymbol{\pi}\|_2^2 - \|\boldsymbol{\pi}_\bullet\|_2^2} \\ &= \theta_\infty/\sqrt{1 - \theta_\infty^2\|\boldsymbol{\pi}_\bullet\|_2^2} \end{aligned} \quad ,$$

so that

$$\theta_n \leq \theta_\infty(1 + O(\|\boldsymbol{\pi}_\bullet\|_2^2)) \leq \theta_\infty(1 + O(\|\boldsymbol{\pi}_\bullet\|_1^2)). \tag{2.3}$$

Thus θ_n is bounded by a quantity that approaches θ_∞ as the *square* of $\|\boldsymbol{\pi}_\bullet\|_1$.

It is worth noting that we have in some sense imitated the adjusted truncation method in perturbing Q_{nn} so that it has a null vector. But in the adjusted truncation method the null vector in question consists of the first n components of the *right* null vector of Q — i.e., a vector whose components are all one — and the perturbation is explicitly computed. In the pure method the null vector consists of the first n components of the *left* null vector of Q. Since these components are unknown, we cannot actually compute $\tilde{Q}_{nn}$. But we can use our bound on $\|E\|_2$ to bound the error in $\hat{\boldsymbol{\pi}}_n$ as an approximation to $\boldsymbol{\pi}_n$.

3. Accuracy

In assessing the accuracy of $\hat{\boldsymbol{\pi}}_1$, it will be convenient to work with the renormalized vectors

$$p_n = \frac{\boldsymbol{\pi}_n}{\|\boldsymbol{\pi}_n\|_2} \quad \text{and} \quad \hat{p}_n = \frac{\hat{\boldsymbol{\pi}}_n}{\|\hat{\boldsymbol{\pi}}_n\|_2}.$$

Let $U = (\hat{p}_n \; V)$ be orthogonal. We are going determine a vector h such that

$$p_n = \hat{p}_n + Vh + O(\epsilon_n^2),$$

so that $\|p_n - \hat{p}_n\|_2 \cong \|h\|_2$.

Remembering that $\hat{p}_n^{\mathrm{T}} Q_{nn} = \nu_n \hat{p}_n^{\mathrm{T}}$, we form the matrix

$$\begin{aligned} T &= U^{\mathrm{T}} Q_{nn} U \\ &= \begin{pmatrix} \hat{p}_n^{\mathrm{T}} \\ V^{\mathrm{T}} \end{pmatrix} Q_{nn} (\hat{p}_n \;\; V) \\ &= \begin{pmatrix} \hat{p}_n^{\mathrm{T}} Q_{nn} \hat{p}_n & \hat{p}_n^{\mathrm{T}} Q_{nn} V \\ V^{\mathrm{T}} Q_{nn} \hat{p}_n & V^{\mathrm{T}} Q_{nn} V \end{pmatrix} \\ &= \begin{pmatrix} \nu_n \hat{p}_n^{\mathrm{T}} \hat{p}_n & \nu_n \hat{p}^{\mathrm{T}} V \\ V^{\mathrm{T}} Q_{nn} \hat{p}_n & V^{\mathrm{T}} Q_{nn} V \end{pmatrix} \\ &\equiv \begin{pmatrix} \nu_n & 0 \\ g & C_n \end{pmatrix}. \end{aligned}$$

The number ν_n is by definition the smallest eigenvalue of Q_{nn}. Because T is block triangular, the matrix C_n contains the remaining eigenvalues of Q_{nn}.

Now let

$$F = U^{\mathrm{T}} E U = \begin{pmatrix} \hat{p}^{\mathrm{T}} E \hat{p} & \hat{p}^{\mathrm{T}} E V \\ V^{\mathrm{T}} E \hat{p} & V^{\mathrm{T}} E V \end{pmatrix} \equiv \begin{pmatrix} \varphi_{11} & f_{12}^{\mathrm{T}} \\ f_{21} & F_{22} \end{pmatrix}.$$

and let

$$\tilde{T} = U^{\mathrm{T}} \tilde{Q}_{nn} U = \begin{pmatrix} \nu_n + \varphi_{11} & f_{12}^{\mathrm{T}} \\ g + f_{21} & C_n + F_{22} \end{pmatrix}$$

Note that because $\|\cdot\|_2$ is invariant under orthogonal transformations we have

$$|\varphi_{11}|, \;\; \|f_{12}\|_2, \;\; \|f_{21}\|_2, \;\; \|F_{22}\|_2 \leq \epsilon_n,$$

where ϵ_n is defined by (2.1).

We are now in a position to prove our main result.

Theorem 3.1. *Assume that C is nonsingular and let*

$$\delta_n = \|C_n^{-1}\|_2^{-1}.$$

If

$$\epsilon_n / \delta_n < 1,$$

then $C_n + F_{22}$ is nonsingular, and

$$p_n = \frac{\hat{p}_n + Vh}{\|\hat{p}_n + Vh\|_2}, \tag{3.1}$$

where

$$h = -(C_n + F_{22})^{-\mathrm{T}} f_{12}. \tag{3.2}$$

Moreover,

$$\nu_n = -h^{\mathrm{T}} g - \varphi_{11} - h^{\mathrm{T}} f_{21}. \tag{3.3}$$

Proof. To show that $C_n + F_{22}$ is nonsingular, we use the well know fact that if $\|X\|_2 < 1$, then $I + X$ is nonsingular and

$$\|I + X\|_2 \leq (1 - \|X\|_2)^{-1}. \tag{3.4}$$

Now write

$$C_n + F_{22} = C_n(I + C_n^{-1} F_{22}).$$

Then $\|C_n^{-1} F_{22}\|_2 \leq \|C_n^{-1}\|_2 \|F_{22}\|_2 \leq \epsilon_n/\delta_n < 1$. Hence $I + C_n^{-1} F_{22}$ is nonsingular, and so is $C_n + F_{22}$.

Now T is similar to $\tilde{Q}_{nn}$ and hence has a left null vector, which we will denote by $(1 \;\; h^{\mathrm{T}})$. Thus we have

$$(1 \;\; h^{\mathrm{T}}) \begin{pmatrix} \nu_n + \varphi_{11} & f_{12}^{\mathrm{T}} \\ g + f_{21} & C_n + F_{22} \end{pmatrix} = (0 \;\; 0). \tag{3.5}$$

Evaluating the second component of this relation, we find

$$f_{12}^{\mathrm{T}} + h^{\mathrm{T}}(C_n + F_{22}) = 0,$$

so that h satisfies (3.2). Transforming back, we find the the null vector of Q_{nn}^{T} is

$$(\hat{p}_n \;\; V) \begin{pmatrix} 1 \\ h \end{pmatrix} = \hat{p}_n + Vh.$$

Hence (3.1) follows from the fact that p_n is the normalized null vector of Q_{nn}^{T}.

The expression (3.3) may be derived by evaluating the first component of (3.5) and solving for ν_n. ■

Turning now to quantitative bounds, we begin by bounding h. From (3.2), we have

$$\|h\|_2 \leq \|(C_n + F_{22})^{-1}\|_2 \|f_{12}\|_2.$$

Now

$$\begin{aligned} \|(C_n + F_{22})^{-1}\|_2 &= \|(I + C_n^{-1} F_{22})^{-1} C_n^{-1}\|_2 \\ &\leq \|(I + C_n^{-1} F_{22})^{-1}\|_2 \|C_n^{-1}\|_2 \\ &\leq (1 - \|C_n^{-1}\|_2 \|F_{22}\|) \|C_n^{-1}\|_2 \quad \text{from (3.4)} \\ &\leq (1 - \epsilon_n/\delta_n)^{-1} \delta_n^{-1} \\ &= 1/(\delta_n - \epsilon_n). \end{aligned}$$

Hence

$$\|h\|_2 \leq \frac{\epsilon_n}{\delta_n - \epsilon_n} = \frac{\epsilon_n}{\delta_n} + O(\epsilon_n^2).$$

The unnormalized approximation to p_n is $\hat{p}_n + Vh$. By the orthonormality of U, we have

$$\|\hat{p}_n + Vh\|_2^2 = (\hat{p}_n + Vh)^{\mathrm{T}}(\hat{p}_n + Vh) = \hat{p}_n^{\mathrm{T}}\hat{p}_n + h^{\mathrm{T}}V^{\mathrm{T}}Vh = 1 + \|h\|_2^2.$$

Hence the normalization of $\hat{p}_n + Vh$ introduces only a second order effect, and we may write

$$\|\hat{p}_n - p_n\|_2 = \|Vh\|_2 + O(\epsilon_n^2) = \|h\|_2 + O(\epsilon_n^2) \leq \frac{\epsilon_n}{\delta_n} + O(\epsilon_n^2)$$

Finally, from (2.1), we have $\epsilon_n \leq \theta_n \|\boldsymbol{\pi}_\bullet\|_1 + O(\|\boldsymbol{\pi}_\bullet\|_1^2)$. Moreover by (2.3), $\theta_n \leq \theta_\infty + O(\|\boldsymbol{\pi}_\bullet\|_1^2)$. Thus we have the following corollary.

Corollary 3.2. *Under the hypotheses of Theorem 3.1,*

$$\|\hat{p}_n - p_n\|_2 \leq \theta_\infty \frac{\|\boldsymbol{\pi}_\bullet\|_1}{\delta_n} + O(\|\boldsymbol{\pi}_\bullet\|_1^2). \tag{3.6}$$

There are several comments to be made about Theorem 3.1 and its corollary.

- The bound (3.6) is a normwise bound and does not guarantee that the individual components of $\hat{p}$ are accurate. However, the bound does provide information about the larger components. In fact, suppose that

$$\|\hat{p}_n - p_n\| = \eta \quad \text{and} \quad |p_i^{(n)}| = \tau\eta \ (\tau > 0), \tag{3.7}$$

where $p_i^{(n)}$ is the ith component of p_n. Since $|\hat{p}_i^{(n)} - p_i^{(n)}| \leq \eta$, we must have

$$\frac{|\hat{p}_i^{(n)} - p_i^{(n)}|}{|p_i^{(n)}|} \leq \frac{1}{\tau}. \tag{3.8}$$

Thus as τ increases the number of significant figures in $\hat{p}_i^{(n)}$ increases proportionally. It is worth noting that this result is true of any norm satisfying $|x_i| \leq \|x\|$.

- The bound is directly proportional to $\|\boldsymbol{\pi}_\bullet\|_1$ — i.e., to the sum of the sizes of the components of $\boldsymbol{\pi}$ not included in the truncation set. In many Markov chains, the components decrease swiftly in magnitude, so that even for small n the size of $\|\boldsymbol{\pi}_\bullet\|_1$ can be made small.

- The bound depends inversely on $\delta_n = \|C_n^{-1}\|_2^{-1}$, which is the smallest singular value of C_n. To get some idea of what this quantity means, let (γ, y) be an eigenpair of C_n with $\|y\|_2 = 1$. Since $C_n y = \gamma y$, we have $\gamma^{-1} y = C_n^{-1} y$. Taking norms, we get

$|\gamma|^{-1} \le \|C_n^{-1}\|_2$, or $\delta_n \le |\gamma|$. Since γ is an arbitrary eigenvalue of C_n and the eigenvalues of C_n are the eigenvalues of Q_{nn} other than ν_n, we have

$$\delta_n \le \min_{\substack{\gamma \in \Lambda(Q_{nn}) \\ \gamma \ne \nu_n}} |\gamma|.$$

• The bounds suggest what must be done to prove the convergence of the method of pure truncation as $n \to \infty$. From (3.6) we see that the accuracy of $\hat{p}$ depends on the ratio of $\|\boldsymbol{\pi}_\bullet\|_1$ to δ_n. Thus we must show that as $n \to \infty$, the quantity $\|\boldsymbol{\pi}_\bullet\|_1$ decreases faster than δ_n. In particular, if δ_n is bounded away from zero as n increases, then the method converges. Whether or not this happens will depend on the properties of the Markov chain in question.

4. An example

For a numerical example of the above analysis, we consider a simple birth-death process with arrival rate $\lambda = 0.4$ and a departure rate of $\mu = 0.45$. The unnormalized stationary vector is given by

$$(1 \;\; \rho \;\; \rho^2 \;\; \cdots),$$

where

$$\rho = \frac{\lambda}{\mu} \cong 0.889.$$

Since we know the vector $\boldsymbol{\pi}$, we can calculate the quantities in our bounds.

All computations were performed in Matlab in IEEE double-precision floating-point arithmetic. The eigenvector $\hat{p}_{nn}$ of Q_{nn} was computed by tow iterations of the inverse power method, starting with a pseudo-random vector whose elements were uniformly distributed in $[0, 1]$. The results are summarized in the table in Figure 4.1.

The first column in this table is the size of the truncated matrix. The second column prints the first component to have a relative error greater than 0.5; i.e.,

$$\text{bad} = \min \left\{ i\colon \frac{|\hat{p}_i^{(n)} - p_i^{(n)}|}{|p_i^{(n)}|} \ge 0.5 \right\}.$$

Figure 4.2 contains a plot or the relative error for $n = 500$. It is seen that the first components are approximated to the highest possible accuracy — about 10^{-16}, which is the rounding unit for IEEE double precision arithmetic. The relative accuracy then decreases, but the method is still doing a fine job of approximating very small components.

In this connection, it is worth noting that in the larger problems it was necessary to do two iterations of the inverse power method to compute the small components of

n	p_{bad}	$\|f_{12}^{\mathrm{T}} C^{-1}\|$	(3.6)	$\nu_2(Q_{11})$	
	δ				
25	3.9e−02	7.1e−02	1.5e+01	1.4e−02	1.3e−02
50	2.3e−03	6.6e−03	3.0e+00	4.2e−03	3.4e−03
75	1.2e−04	4.8e−04	3.2e−01	2.6e−03	1.7e−03
100	6.3e−06	3.1e−05	2.6e−02	2.0e−03	1.1e−03
125	3.3e−07	1.9e−06	1.8e−03	1.8e−03	8.2e−04
150	1.8e−08	1.1e−07	1.2e−04	1.7e−03	6.5e−04
175	9.2e−10	6.3e−09	7.6e−06	1.6e−03	5.4e−04
200	4.9e−11	3.6e−10	4.7e−07	1.6e−03	4.6e−04
225	2.6e−12	2.0e−11	2.8e−08	1.6e−03	4.0e−04
250	1.3e−13	1.1e−12	1.7e−09	1.5e−03	3.5e−04
275	7.1e−15	6.2e−14	9.9e−11	1.5e−03	3.2e−04
300	3.7e−16	3.8e−15	5.7e−12	1.5e−03	2.9e−04
325	2.0e−17	1.2e−15	3.3e−13	1.5e−03	2.6e−04
350	1.0e−18	1.1e−15	1.9e−14	1.5e−03	2.4e−04
375	5.4e−20	1.1e−15	1.1e−15	1.5e−03	2.3e−04
400	2.9e−21	1.1e−15	6.0e−17	1.5e−03	2.1e−04
425	1.5e−22	1.1e−15	3.4e−18	1.5e−03	2.0e−04
450	7.9e−24	1.1e−15	1.9e−19	1.5e−03	1.9e−04
475	4.2e−25	1.1e−15	1.1e−20	1.5e−03	1.8e−04
500	2.2e−26	1.1e−15	5.8e−22	1.5e−03	1.7e−04

Figure 4.1: Birth-death process: $\lambda = .4$, $\mu = .45$

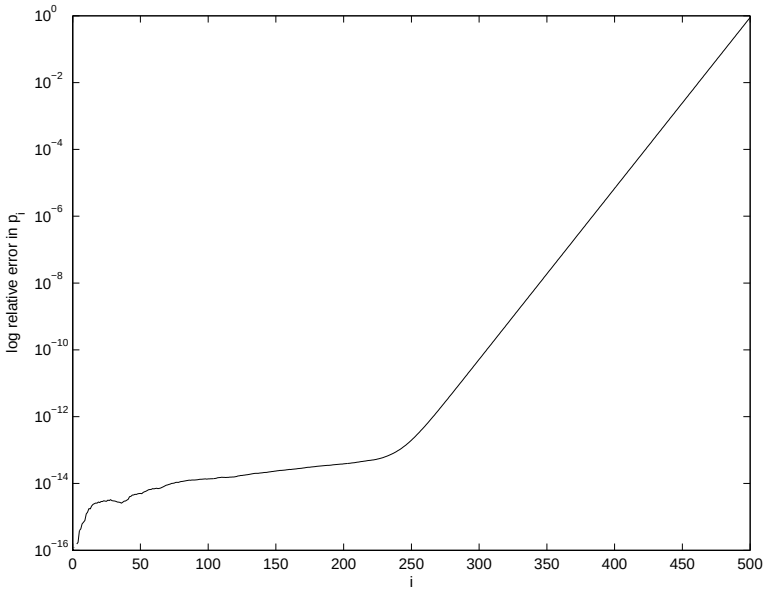

Figure 4.2: $\log\left(\frac{|\hat{p}_i^{(500)} - p_i^{(500)}|}{|p_i^{(500)}|}\right)$

$\hat{p}$ accurately. This fact illustrates an important difference between ordinary eigenvalue problems, in which the components of the desired eigenvector need not vary greatly in size, and Morkov chains where ultimately they must. The analysis of the inverse power method [9, p.619 ff.] [7, §1.1.2] shows that for these problems the inverse power method should converge in a single iteration. However, this analysis is based on norms, and says little about the smaller components. Hence the need for a second iteration.

The third column is a check on our basic approach. In contains the quantity $\|C_n^{-1} f_{12}\|_2$, which up to second order terms in ϵ_n is the norm of $\hat{p}_n - p_n$. By (3.7) and (3.8), these numbers should be larger than $p_{\rm bad}$, but of roughly the same order of magnitude. Note that these numbers stagnate at 10^{-15}, owing to our inability to compute f_{12} accurately enough.

The fourth column contains the bound (3.6). Because, we have used norm inequalities in deriving it, the bound is too large by several orders of magnitude, though it tracks to column one, as it should. As is usual in numerical applications, the bound gives insight but cannot be used practically for determining convergence.

The fifth and sixth columns concern the asymptotic behavior of the second smallest eigenvalue of Q_{nn} and δ_n. The former seems to be approaching a nonzero limit. It is difficult to tell what the latter is doing; however, if it is not bounded away from zero, it seems to be approaching zero far more slowly than $\|\boldsymbol{\pi}_2\|_1$. If this is true, then the truncated method is converging very swiftly to the initial segments of $\boldsymbol{\pi}^{\rm T}$ as n increases.

5. Relation to methods of adjusted truncation

The purpose of this section is to point out an interesting relation between the method of adjusted truncation and the method of pure truncation. Seneta [4, 5] gives a complete treatment of the subject as of about 1980. See Tweedie [8] for a more recent bibliography. Here we will follow Seneta.

The starting point is the following observation. Let $b_n^{(k)}$ be defined by

$$Q_{nn}^{\rm T} b_n^{(k)} = \mathbf{e}_k, \tag{5.1}$$

where $\mathbf{e}_k$ is the vector whose kth component is one and whose other component is zero. For certain classes of chains and certain values of k (generally $k = 1$ or $k = n$) it can be shown that

$$\frac{b_n^{(k)}}{\mathbf{e}^{\rm T} b_n^{(k)}} \to \boldsymbol{\pi}. \tag{5.2}$$

Here $\mathbf{e}$ is the vector consisting of all ones, so that the left hand side of (5.2) is $b_n^{(k)}$ normalized so that it represents a probability distribution.

Unfortunately, the system (3.2) can be shown to become increasingly ill-conditioned with increasing n, and the the vector $b_n^{(k)}$ will be computed inaccurately in finite precision

arithmetic. However, it can be shown that $b_n^{(k)\mathrm{T}}$ is the left Perron vector of the matrix obtained by adjusting the kth column of Q_{nn} so that it becomes stochastic. Under some circumstances the condition of this system is better than that of Q_{nn}, and this was the original rationale for the method of adjusted truncation.

A different point of view emerges if we focus on what is actually being computed by the method of adjusted truncation. The system (5.1) represents one iteration of the inverse power method for the left eigenvector of Q_{nn} corresponding to its smallest eigenvalue, i.e., for the vector $\hat{\boldsymbol{\pi}}_n^{\mathrm{T}}$. When Q_{nn} is sufficiently ill-conditioned, the method typically converges to its limiting accuracy in one iteration and is unaffected by errors introduced by the ill-conditioning [7, §1.1.2] [9, p.619 ff.]. Thus the method of adjusted truncation can be regarded as approximating the method of pure truncation and vice versa.[1] However, to say they approximate one another, is not to say that they are the same method. Although the method based on (5.1) is one step of the inverse power method with a special starting vector, there is no justification for taking a second step, which we saw above is necessary to compute the small components accurately.

Both methods require special justification for individual classes of problems. But the proof techniques are different. Examples for the method of adjusted truncation will be found in the references cited above. For the method of pure truncation, the behavior of the quantity δ_n has plays a central role our analysis. For numerical analysts and perturbation theorists this should come as no surprise; its reciprocal measures the sensitivity of $\hat{\boldsymbol{\pi}}_n^{\mathrm{T}}$ to perturbations in Q_{nn} and also the limiting accuracy of the inverse power method [7, §1.3.2]. Since it has now surfaced in the context of infinite Markov chains, I believe an important problem for further research is to place it in a probabilistic setting.

Acknowledgements

I wish to thank Guy Latouche for his comments on an earlier version of the paper and for a useful list of references. Part of this work was done while I was a faculty appointee at the Mathematical and Computational Sciences Division of the National Institute for Standards and Technology.

References

[1] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences.* Academic Press, New York, 1979. Reprinted by SIAM, Philadelphia, 1994.

[1] Seneta [4] points out the relation of the method of adjusted truncation to the inverse power method, but does not observe that this makes the two methods approximations of one another.

[2] Å. Björck, R. J. Plemmons, and H. Schneider, editors. *Large–Scale Matrix Problems.* North-Holland, New York, 1981. A reprint of Volumn 34, 1980, of *Linear Algebra and Its Applications*.

[3] G. H. Golub and C. F. Van Loan. *Matrix Computations.* Johns Hopkins University Press, Baltimore, MD, second edition, 1989.

[4] E. Seneta. Computing the stationary distribution for infinite Markov chains. *Linear Algebra and Its Applications*, 34:259–267, 1980. Reprinted in [2].

[5] E. Seneta. *Non-Negative Matrices and Markov Chains.* Springer Verlag, New York, 1981.

[6] G. W. Stewart. *Matrix Algorithms I: Basic Decompositions.* SIAM, Philadelphia, 1998.

[7] G. W. Stewart. *Matrix Algorithms II: Eigensystems.* SIAM, Philadelphia, 2001.

[8] R. L. Tweedie. Truncation approximations of invariant measures for Markov Chains. *Journal of Applied Probability*, 35:517–536, 1998.

[9] J. H. Wilkinson. *The Algebraic Eigenvalue Problem.* Clarendon Press, Oxford, 1965.