
Dimension-free Concentration Bounds on Hankel Matrices for Spectral Learning

François Denis
Mattias Gybels

Aix Marseille Université, CNRS, LIF, 13288 Marseille Cedex 9, FRANCE

FRANCOIS.DENIS@LIF.UNIV-MRS.FR
MATTIAS.GYBELS@LIF.UNIV-MRS.FR

Amaury Habrard

Université Jean Monnet de Saint-Etienne, CNRS, LaHC, 42000 Saint-Etienne Cedex 2, FRANCE

AMAURY.HABRARD@UNIV-ST-ETIENNE.FR

Abstract

Learning probabilistic models over strings is an important issue for many applications. Spectral methods propose elegant solutions to the problem of inferring weighted automata from finite samples of variable-length strings drawn from an unknown target distribution. These methods rely on a singular value decomposition of a matrix H_S , called the Hankel matrix, that records the frequencies of (some of) the observed strings. The accuracy of the learned distribution depends both on the quantity of information embedded in H_S and on the distance between H_S and its mean H_r . Existing concentration bounds seem to indicate that the concentration over H_r gets looser with its size, suggesting to make a trade-off between the quantity of used information and the size of H_r . We propose new dimension-free concentration bounds for several variants of Hankel matrices. Experiments demonstrate that these bounds are tight and that they significantly improve existing bounds. These results suggest that the concentration rate of the Hankel matrix around its mean does not constitute an argument for limiting its size.

1. Introduction

Many applications in natural language processing, text analysis or computational biology require learning probabilistic models over finite variable-size strings such as probabilistic automata, Hidden Markov Models (HMM), or more generally, weighted automata. Weighted automata

exactly model the class of rational series, and their algebraic properties have been widely studied in that context (Droste et al., 2009). In particular, they admit algebraic representations that can be characterized by a set of finite-dimensional linear operators whose rank corresponds to the minimum number of states needed to define the automaton. From a machine learning perspective, the objective is then to infer good estimates of these linear operators from finite samples. In this paper, we consider the problem of learning the linear representation of a weighted automaton, from a finite sample, composed of variable-size strings i.i.d. from an unknown target distribution.

Recently, the seminal papers of Hsu et al. (2009) for learning HMM and Bailly et al. (2009) for weighted automata, have defined a new category of approaches - the so-called *spectral methods* - for learning distributions over strings represented by finite state models (Siddiqi et al., 2010; Song et al., 2010; Balle et al., 2012; Balle & Mohri, 2012). Extensions to probabilistic models for tree-structured data (Bailly et al., 2010; Parikh et al., 2011; Cohen et al., 2012), transductions (Balle et al., 2011) or other graphical models (Anandkumar et al., 2012c;b;a; Luque et al., 2012) have also attracted a lot of interest.

Spectral methods suppose that the main parameters of a model can be expressed as the spectrum of a linear operator and estimated from the spectral decomposition of a matrix that sums up the observations. Given a rational series r , the values taken by r can be arranged in a matrix H_r whose rows and columns are indexed by strings, such that the linear operators defining r can be recovered directly from the right singular vectors of H_r . This matrix is called the Hankel matrix of r .

In a learning context, given a learning sample S drawn from a target distribution p , an empirical estimate H_S of H_p is built and then, a rational series \hat{p} is inferred from the right singular vectors of H_S . However, the size of H_S increases drastically with the size of S and state of the art approaches

consider smaller matrices $H_S^{U,V}$ indexed by limited subset of strings U and V . It can be shown that the above learning scheme, or slight variants of it, are consistent as soon as the matrix $H_r^{U,V}$ has full rank (Hsu et al., 2009; Bailly, 2011; Balle et al., 2012) and that the accuracy of the inferred series is directly connected to the concentration distance $\|H_S^{U,V} - H_p^{U,V}\|_2$ between the empirical Hankel matrix and its mean (Hsu et al., 2009; Bailly, 2011).

On the one hand, limiting the size of the Hankel matrix avoids prohibitive calculations. Moreover, most existing concentration bounds on sum of random matrices depend on their size and suggest that $\|H_S^{U,V} - H_p^{U,V}\|_2$ may become significantly looser with the size of U and V , compromising the accuracy of the inferred model.

On the other hand, limiting the size of the Hankel matrix implies a drastic loss of information: only the strings of S compatible with U and V will be considered. In order to limit the loss of information when dealing with restricted sets U and V , a general trend is to work with other functions than the target p , such as the *prefix* function $\bar{p}(u) = \sum_{v \in \Sigma^*} p(uv)$ or the *factor* function $\hat{p} = \sum_{v,w \in \Sigma^*} p(vuw)$ (Balle et al., 2013; Luque et al., 2012). These functions are rational, they have the same rank as p , a representation of p can easily be derived from representations of \bar{p} or \hat{p} and they allow a better use of the information contained in the learning sample.

A first contribution is to provide a dimension free concentration inequality for $\|H_S^{U,V} - H_p^{U,V}\|_2$, by using recent results on tail inequalities for sum of random matrices showing that restricting the dimension of H is not mandatory.

However, these results cannot be directly applied to the prefix and factor series, since the norm of the corresponding random matrices are unbounded. A second contribution of the paper is to define two classes of parametrized functions, \bar{p}_η and \hat{p}_η , that constitute continuous intermediates between p and \bar{p} (resp. \hat{p}), and to provide analogous dimension-free concentration bounds for these classes.

These bounds are evaluated on a benchmark made of 11 problems extracted from the PAutomaC challenge (Verwer et al., 2012). These experiments show that the bounds derived from our theoretical results are quite tight - compared to the exact values- and that they significantly improve existing bounds, even on matrices of fixed dimensions.

These results have two practical consequences for spectral learning: (i) the concentration of the empirical Hankel matrix around its mean does not highly depend on its dimension and the only reason not to use all the information contained in the sample should only rely on computing resources limitations. In that perspective, using random techniques to perform singular values decomposition on huge Hankel matrices should be considered (Halko et al., 2011);

(ii) by contrast, the concentration is weaker for the prefix and factor functions, and smoothed variants should be used, with an appropriate parameter.

The paper is organized as follows. Section 2 introduces the main notations, definitions and concepts. Section 3 presents a first dimension free-concentration inequality for the standard Hankel matrices. Then, we introduce the prefix and the factor variants and provide analogous concentration results. Section 4 describes some experiments before the conclusion presented in Section 5.

2. Preliminaries

2.1. Singular Values, Eigenvalues and Matrix Norms

Let $M \in \mathbb{R}^{m \times n}$ be a $m \times n$ real matrix. The singular values of M are the square roots of the eigenvalues of the matrix $M^T M$, where M^T denotes the transpose of M : $\sigma_{max}(M)$ and $\sigma_{min}(M)$ denote the largest and smallest singular value of M , respectively.

In this paper, we mainly use the spectral norms $\|\cdot\|_k$ induced by the corresponding vector norms on \mathbb{R}^n and defined by $\|M\|_k = \max_{x \neq 0} \frac{\|Mx\|_k}{\|x\|_k}$:

- $\|M\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^m |M[i, j]|$,
- $\|M\|_\infty = \max_{1 \leq i \leq m} \sum_{j=1}^n |M[i, j]|$,
- $\|M\|_2 = \sigma_{max}(M)$.

We have: $\|M\|_2 \leq \sqrt{\|M\|_1 \|M\|_\infty}$.

These norms can be extended, under certain conditions, to infinite matrices and the previous inequalities remain true when the corresponding norms are defined.

2.2. Rational stochastic languages and Hankel matrices

Let Σ be a finite alphabet. The set of all finite strings over Σ is denoted by Σ^* , the empty string is denoted by ϵ , the length of string w is denoted by $|w|$ and Σ^n (resp. $\Sigma^{\leq n}$) denotes the set of all strings of length n (resp. $\leq n$). For any string w , let $\text{Pref}(w) = \{u \in \Sigma^* \mid \exists v \in \Sigma^* w = uv\}$.

A *series* is a mapping $r : \Sigma^* \mapsto \mathbb{R}$. A series r is convergent if the sequence $r(\Sigma^{\leq n}) = \sum_{w \in \Sigma^{\leq n}} r(w)$ is convergent; its limit is denoted by $r(\Sigma^*)$. A *stochastic language* p is a probability distribution over Σ^* , i.e. a series taking non negative values and converging to 1.

Let $n \geq 1$ and M be a morphism defined from Σ^* to $\mathcal{M}(n)$, the set of $n \times n$ matrices with real coefficients. For all $u \in \Sigma^*$, let us denote $M(u)$ by M_u and $\sum_{x \in \Sigma} M_x$ by M_Σ . A series r over Σ is *rational* if there exists an integer $n \geq 1$, two vectors $I, T \in \mathbb{R}^n$ and a morphism $M : \Sigma^* \mapsto \mathcal{M}(n)$ such that for all $u \in \Sigma^*$, $r(u) = I^T M_u T$.

The triplet $\langle I, M, T \rangle$ is called an n -dimensional *linear representation* of r . The vector I can be interpreted as a vector of initial weights, T as a vector of terminal weights and the morphism M as a set of matrix parameters associated with the letters of Σ . A *rational stochastic language* is thus a stochastic language admitting a linear representation.

Let $U, V \subseteq \Sigma^*$, the *Hankel matrix* $H_r^{U,V}$, associated with a series r , is the matrix indexed by $U \times V$ and defined by $H_r^{U,V}[u, v] = r(uv)$, for any $(u, v) \in U \times V$. If $U = V = \Sigma^*$, $H_r^{U,V}$, simply denoted by H_r , is a bi-infinite matrix. In the following, we always assume that $\epsilon \in U$ and that U and V are ordered in quasi-lexicographic order: strings are first ordered by increasing length and then, according to the lexicographic order. It can be shown that a series r is rational if and only if the rank of the matrix H_r is finite. The rank of H_r is equal to the minimal dimension of a linear representation of r .

Let r be a non negative convergent rational series and let $\langle I, M, T \rangle$ be a minimal d -dimensional linear representation of r . Then, the sum $I_d + M_\Sigma + \dots + M_\Sigma^n + \dots$ is convergent and $r(\Sigma^*) = I^T(I_d - M_\Sigma)^{-1}T$ where I_d is the identity matrix of size d .

Several convergent rational series can be naturally associated with a stochastic language p :

- \bar{p} , defined by $\bar{p}(u) = \sum_{v \in \Sigma^*} p(uv)$, the series associated with the *prefixes* of the language,
- \hat{p} , defined by $\hat{p}(u) = \sum_{v, w \in \Sigma^*} p(vuw)$, the series associated with the *factors* of the language.

It can be noticed that $\bar{p}(u) = p(u\Sigma^*)$, the probability that a string begins with u , but that in general, $\hat{p}(u) \geq p(\Sigma^*u\Sigma^*)$, the probability that a string contains u as a substring.

If $\langle I, M, T \rangle$ is a minimal d -dimensional linear representation of p , then $\langle I, M, (I_d - M_\Sigma)^{-1}T \rangle$ (resp. $\langle [I^T(I_d - M_\Sigma)^{-1}]^T, M, (I_d - M_\Sigma)^{-1}T \rangle$) is a minimal linear representation of \bar{p} (resp. of \hat{p}). Any linear representation of these variants of p can be reconstructed from the others.

For any integer $k \geq 1$, let

$$S_p^{(k)} = \sum_{u_1 u_2 \dots u_k \in \Sigma^*} p(u_1 u_2 \dots u_k) = I^T(I_d - M_\Sigma)^{-k}T.$$

Clearly, $p(\Sigma^*) = S_p^{(1)} = 1$, $\bar{p}(\Sigma^*) = S_p^{(2)}$ and $\hat{p}(\Sigma^*) = S_p^{(3)}$. Note that $S_p^{(2)} = 1 + \sum_u |u|p(u)$, where $\sum_u |u|p(u)$ is the average length of a string drawn according to p .

Let $U, V \subseteq \Sigma^*$. For any string $w \in \Sigma^*$, let us define the matrices $H_w^{U,V}$, $\bar{H}_w^{U,V}$ and $\hat{H}_w^{U,V}$ by

- $H_w^{U,V}[u, v] = \mathbf{1}_{uv=w}$,

- $\bar{H}_w^{U,V}[u, v] = \mathbf{1}_{uv \in Pref(w)}$ and
- $\hat{H}_w^{U,V}[u, v] = \sum_{x, y \in \Sigma^*} \mathbf{1}_{xuvy=w}$

for any $(u, v) \in U \times V$. For any sample of strings S , let $H_S^{U,V} = \frac{1}{|S|} \sum_{w \in S} H_w^{U,V}$, $\bar{H}_S^{U,V} = \frac{1}{|S|} \sum_{w \in S} \bar{H}_w^{U,V}$ and $\hat{H}_S^{U,V} = \frac{1}{|S|} \sum_{w \in S} \hat{H}_w^{U,V}$.

For example, let $S = \{a, ab\}$, $U = V = \{\epsilon, a, b\}$. We have

$$H_S^{U,V} = \begin{pmatrix} 0 & 0.5 & 0 \\ 0.5 & 0 & 0.5 \\ 0 & 0 & 0 \end{pmatrix}, \quad \bar{H}_S^{U,V} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 0.5 \\ 0 & 0 & 0 \end{pmatrix}, \quad \hat{H}_S^{U,V} = \begin{pmatrix} 2.5 & 1 & 0.5 \\ 1 & 0 & 0.5 \\ 0.5 & 0 & 0 \end{pmatrix}.$$

2.3. Spectral Algorithm for Learning Rational Stochastic Languages

Rational series admit a *canonical linear representation* determined by their Hankel matrix. Let r be a rational series of rank d and $U \subset \Sigma^*$ such that the matrix $H_r^{U \times \Sigma^*}$ (denoted by H in the following) has rank d .

- For any string s , let T_s be the constant matrix whose rows and columns are indexed by Σ^* and defined by $T_s[u, v] = 1$ if $v = us$ and 0 otherwise.
- Let E be a vector indexed by Σ^* whose coordinates are all zero except the first one equals to 1: $E[u] = \mathbf{1}_{u=\epsilon}$ and let P be the vector indexed by Σ^* defined by $P[u] = r(u)$.
- Let $H = LDR^T$ be a reduced singular value decomposition of H : R (resp. L) is a matrix whose columns form a set of orthonormal vectors - the right (resp. left) singular vectors of H - and D is a $d \times d$ diagonal matrix, composed of the singular values of H .

Then, $\langle R^T E, (R^T T_x R)_{x \in \Sigma}, R^T P \rangle$ is a linear representation of r (Baillly et al., 2009; Baillly, 2011; Balle et al., 2012). A quick proof can be found in (Denis et al., 2013).

The basic spectral algorithm for learning rational stochastic languages aims at identifying the canonical linear representation of the target p determined by its Hankel matrix H_p .

Let S be a sample independently drawn according to p :

- Choose sets $U, V \subseteq \Sigma^*$ and build the Hankel matrix $H_S^{U \times V}$,
- choose a rank d and compute a reduced SVD of $H_S^{U \times V}$ truncated at rank d ,
- build the canonical linear representation $\langle R_S^T E, (R_S^T T_x R_S)_{x \in \Sigma}, R_S^T P_S \rangle$ from the right singular vectors R_S and the empirical distribution p_S defined from S .

Alternative learning strategies consist in learning \bar{p} or \hat{p} , using the same algorithm, and then to compute an estimate of p . In all cases, the accuracy of the learned representation mainly depends on the estimation of R . The Stewart formula (Stewart, 1990) bounds the principle angle θ between the spaces spanned by the right singular vectors of R and R_S :

$$|\sin(\theta)| \leq \frac{\|H_S^{U \times V} - H_r^{U \times V}\|_2}{\sigma_{\min}(H_r^{U \times V})}.$$

According to this formula, the concentration of the Hankel matrix around its mean is critical and the question of limiting the sizes of U and V naturally arises. Note that the Stewart inequality does not give any clear indication on the impact or on the interest of limiting these sets. Indeed, Weyl's inequalities can be used to show that both the numerator and the denominator of the right part of the inequality increase with U and V .

3. Concentration Bounds for Hankel Matrices

Let p be a rational stochastic language over Σ^* , let ξ be a random variable distributed according to p , let $U, V \subseteq \Sigma^*$ and let $Z(\xi) \in \mathbb{R}^{|U| \times |V|}$ be a random matrix. For instance, $Z(\xi)$ may be equal to $H_\xi^{U,V}$, $\bar{H}_\xi^{U,V}$ or $\hat{H}_\xi^{U,V}$.

Concentration bounds for sum of random matrices can be used to estimate the spectral distance between the empirical matrix Z_S computed on the sample S and its mean (see (Hsu et al., 2011) for references). However, most of classical inequalities depend on the dimensions of the matrices. For example, it can be proved that with probability at least $1 - \delta$ (Kakade, 2010):

$$\|Z_S - \mathbb{E}Z\|_2 \leq \frac{6M}{\sqrt{N}} \left(\sqrt{\log d} + \sqrt{\log \frac{1}{\delta}} \right) \quad (1)$$

where N is the size of S , d is the minimal dimension of the matrix Z and $\|Z\|_2 \leq M$ almost surely. If $Z = H_\xi^{U,V}$, then $M = 1$; if $Z = \bar{H}_\xi^{U,V}$, $M = \Omega(D^{1/2})$ in the worst case; if $Z = \hat{H}_\xi^{U,V}$, $\|Z\|_2$ is generally unbounded.

These concentration bounds get worse with both sizes of the matrices. Coming back to the discussion at the end of Section 2, they suggest to limit the size of the sets U and V , and therefore, to design strategies to choose optimal sets.

We then use recent results (Tropp, 2012; Hsu et al., 2011) to obtain dimension-free concentration bounds for Hankel matrices. More precisely, we extend a Bernstein bound for unbounded random matrices from (Hsu et al., 2011) to non symmetric random matrices by using the dilation principle (Tropp, 2012).

Let Z be a random matrix, the *dilation* of Z is the symmet-

ric random matrix X defined by

$$X = \begin{bmatrix} 0 & Z \\ Z^T & 0 \end{bmatrix}. \text{ Then } X^2 = \begin{bmatrix} ZZ^T & 0 \\ 0 & Z^T Z \end{bmatrix}$$

and $\|X\|_2 = \|Z\|_2$, $\text{tr}(X^2) = \text{tr}(ZZ^T) + \text{tr}(Z^T Z)$ and $\|X^2\|_2 \leq \text{Max}(\|ZZ^T\|_2, \|Z^T Z\|_2)$.

We can then reformulate the result that we use as follows (Denis et al., 2013).

Theorem 1. *Let ξ_1, \dots, ξ_N be i.i.d. random variables, and for $i = 1, \dots, N$, let $Z_i = Z(\xi_i)$ be i.i.d. matrices and X_i the dilation of Z_i . If there exists $b > 0, \sigma > 0$, and $k > 0$ such that $\mathbb{E}[X_1] = 0$, $\|X_1\|_2 \leq b$, $\|\mathbb{E}(X_1^2)\|_2 \leq \sigma^2$ and $\text{tr}(\mathbb{E}(X_1^2)) \leq \sigma^2 k$ almost surely, then for all $t > 0$,*

$$\Pr \left[\left\| \frac{1}{N} \sum_{i=1}^N X_i \right\|_2 > \sqrt{\frac{2\sigma^2 t}{N}} + \frac{bt}{3N} \right] \leq k \cdot t(e^t - t - 1)^{-1}.$$

We will then make use of this theorem to derive our new concentration bounds. Section 3.1 deals with the standard case, Section 3.2 with the prefix case and Section 3.3 with the factor case.

3.1. Concentration Bound for the Hankel Matrix $H_p^{U,V}$

Let p be a rational stochastic language over Σ^* , let S be a sample independently drawn according to p , and let $U, V \subseteq \Sigma^*$. In this section, we compute a bound on $\|H_S^{U,V} - H_p^{U,V}\|_2$ which is independent from the sizes of U and V and holds in particular when $U = V = \Sigma^*$.

Let ξ be a random variable distributed according to p , let $Z(\xi) = H_\xi^{U,V} - H_p^{U,V}$ be the random matrix defined by $Z_{u,v} = \mathbf{1}_{\xi=uv} - p(uv)$ and let X be the dilation of Z .

Clearly, $\mathbb{E}(X) = 0$. In order to apply Theorem 1, it is necessary to compute the parameters b, σ and k . We first prove a technical lemma that will provide a bound on $\mathbb{E}(X^2)$.

Lemma 1. *For any $u, u' \in U, v, v' \in V$,*

$$|\mathbb{E}(Z_{uv}Z_{u'v})| \leq p(u'v) \text{ and } |\mathbb{E}(Z_{uv}Z_{uv'})| \leq p(uv').$$

Proof.

$$\begin{aligned} \mathbb{E}(Z_{uv}Z_{u'v}) &= \mathbb{E}(\mathbf{1}_{\xi=uv}\mathbf{1}_{\xi=u'v}) - p(uv)p(u'v) \\ &= \sum_{w \in \Sigma^*} p(w)\mathbf{1}_{w=uv}\mathbf{1}_{w=u'v} - p(uv)p(u'v) \\ &= p(u'v)[\mathbf{1}_{u=u'} - p(uv)] \end{aligned}$$

and

$$|\mathbb{E}(Z_{uv}Z_{uv'})| \leq p(u'v).$$

The second inequality is proved in a similar way. \square

Next lemma provides parameters b, σ and k needed to apply Theorem 1.

Lemma 2. $\|X\|_2 \leq 2$, $\mathbb{E}(\text{Tr}(X^2)) \leq 2S_p^{(2)}$ and $\|\mathbb{E}(X^2)\|_2 \leq S_p^{(2)}$.

Proof. 1. $\forall u \in U$, $\sum_{v \in V} |Z_{u,v}| = \sum_{v \in V} |\mathbf{1}_{\xi=uv} - p(uv)| \leq 1 + p(u\Sigma^*) \leq 2$. Therefore, $\|Z\|_\infty \leq 2$. In a similar way, it can be shown that $\|Z\|_1 \leq 2$. Hence,

$$\|X\|_2 = \|Z\|_2 \leq \sqrt{\|Z\|_\infty \|Z\|_1} \leq 2.$$

2. For all $(u, u') \in U^2$: $ZZ^T[u, u'] = \sum_{v \in V} Z_{u,v} Z_{u',v}$.

Therefore,

$$\begin{aligned} \mathbb{E}(\text{Tr}(ZZ^T)) &= \mathbb{E}\left(\sum_{u \in U} ZZ^T[u, u]\right) \\ &= \mathbb{E}\left(\sum_{u \in U, v \in V} Z_{u,v} Z_{u,v}\right) \\ &\leq \sum_{u \in U, v \in V} \mathbb{E}(Z_{u,v} Z_{u,v}) \\ &\leq S_p^{(2)}. \end{aligned}$$

In a similar way, it can be proved that $\mathbb{E}(\text{Tr}(Z^T Z)) \leq S_p^{(2)}$ and therefore, $\mathbb{E}(\text{Tr}(X^2)) \leq 2S_p^{(2)}$.

3. For any $u \in U$,

$$\begin{aligned} \sum_{u' \in U} |\mathbb{E}(ZZ^T[u, u'])| &\leq \sum_{u' \in U, v \in V} |\mathbb{E}(Z_{u,v} Z_{u',v})| \\ &\leq \sum_{u' \in U, v \in V} p(u'v) \\ &\leq S_p^{(2)}. \end{aligned}$$

Hence, $\|ZZ^T\|_\infty \leq S_p^{(2)}$. It can be proved, in a similar way, that $\|Z^T Z\|_\infty \leq S_p^{(2)}$, $\|ZZ^T\|_1 \leq S_p^{(2)}$ and $\|Z^T Z\|_1 \leq S_p^{(2)}$. Therefore, $\|X^2\|_2 \leq S_p^{(2)}$. \square

We can now prove the main theorem of this section:

Theorem 2. *Let p be a rational stochastic language and let S be a sample of N strings drawn i.i.d. from p . For all $t > 0$,*

$$\Pr \left[\|H_S^{U,V} - H_p^{U,V}\|_2 > \sqrt{\frac{2S_p^{(2)}t}{N}} + \frac{2t}{3N} \right] \leq 2t(e^t - t - 1)^{-1}.$$

Proof. Let ξ_1, \dots, ξ_N be N independent copies of ξ , let $Z_i = Z(\xi_i)$ and let X_i be the dilation of Z_i for $i = 1, \dots, N$. Lemma 2 shows that the 4 conditions of Theorem 1 are fulfilled with $b = 2$, $\sigma^2 = S_p^{(2)}$ and $k = 2$. \square

This bound is independent from U and V . It can be noticed that the proof also provides a dimension dependent bound by replacing $S_p^{(2)}$ with $\sum_{(u,v) \in U \times V} p(uv)$, which may result in a significative improvement if U or V are small.

3.2. Bound for the prefix Hankel Matrix $H_{\bar{p}}^{U,V}$

The random matrix $\bar{Z}(\xi) = \bar{H}_\xi^{U,V} - H_{\bar{p}}^{U,V}$ is defined by $\bar{Z}_{u,v} = \mathbf{1}_{uv \in \text{Pref}(\xi)} - \bar{p}(uv)$. It can easily be shown that $\|\bar{Z}\|_2$ may be unbounded if U or V are unbounded: $\|\bar{Z}\|_2 = \Omega(|\xi|^{1/2})$. Hence, Theorem 1 cannot be directly applied, which suggests that the concentration of \bar{Z} around its mean could be far weaker than the concentration of Z .

For any $\eta \in [0, 1]$, we define a smoothed variant of \bar{p} by

$$\bar{p}_\eta(u) = \sum_{x \in \Sigma^*} \eta^{|x|} p(ux) = \sum_{n \geq 0} \eta^n p(u\Sigma^n).$$

Note that $\bar{p}_1 = \bar{p}$, $\bar{p}_0 = p$ and that $p(u) \leq \bar{p}_\eta(u) \leq \bar{p}(u)$ for any string u . Therefore, the functions \bar{p}_η are natural intermediates between p and \bar{p} . Moreover, when p is rational, each \bar{p}_η is also rational.

Proposition 1. *Let p be a rational stochastic language and let $\langle I, (M_x)_{x \in \Sigma}, T \rangle$ be a minimal linear representation of p . Let $\bar{T}_\eta = (I_d - \eta M_\Sigma)^{-1} T$. Then, \bar{p}_η is rational and $\langle I, (M_x)_{x \in \Sigma}, \bar{T}_\eta \rangle$ is a linear representation of \bar{p}_η .*

Proof. For any string u , $\bar{p}_\eta(u) = \sum_{n \geq 0} I^T M_u \eta^n M_\Sigma^n T = I^T M_u (\sum_{n \geq 0} \eta^n M_\Sigma^n) T = I^T M_u \bar{T}_\eta$. \square

Note that T can be computed from \bar{T}_η when η and M_Σ are known and therefore, it is a consistent learning strategy to learn \bar{p}_η from the data, for some η , and next, to derive p .

For any $0 \leq \eta \leq 1$, let $\bar{Z}_\eta(\xi)$ be the random matrix defined by

$$\begin{aligned} \bar{Z}_\eta[u, v] &= \sum_{x \in \Sigma^*} \eta^{|x|} \mathbf{1}_{\xi=uvx} - \bar{p}_\eta(uv) \\ &= \sum_{x \in \Sigma^*} \eta^{|x|} (\mathbf{1}_{\xi=uvx} - p(uvx)). \end{aligned}$$

for any $(u, v) \in U \times V$. It is clear that $\mathbb{E}(\bar{Z}_\eta) = 0$ and we show below that $\|\bar{Z}_\eta\|_2$ is bounded if $\eta < 1$.

The moments $S_{\bar{p}_\eta}^{(k)}$ can naturally be associated with \bar{p}_η . For any $0 \leq \eta \leq 1$ and any $k \geq 1$, let

$$S_{\bar{p}_\eta}^{(k)} = \sum_{u_1 u_2 \dots u_k \in \Sigma^*} \bar{p}_\eta(u_1 u_2 \dots u_k).$$

We have $S_{\bar{p}_\eta}^{(k)} = I^T (I_d - M_\Sigma)^{-k} (I_d - \eta M_\Sigma)^{-1} T$ and it is clear that $S_{\bar{p}_0}^{(k)} = S_p^{(k)}$ and $S_{\bar{p}_1}^{(k)} = S_{\bar{p}}^{(k)}$.

Lemma 3.

$$\|\bar{Z}_\eta\|_2 \leq \frac{1}{1 - \eta} + S_{\bar{p}_\eta}^{(1)}.$$

Proof. Indeed, let $u \in U$.

$$\begin{aligned} \sum_{v \in V} |\bar{Z}_\eta[u, v]| &\leq \sum_{v, x \in \Sigma^*} \eta^{|x|} \mathbf{1}_{\xi=uvx} + \sum_{v, x \in \Sigma^*} \eta^{|x|} p(uvx) \\ &\leq (1 + \eta + \dots + \eta^{|\xi|-|u|}) + S_{\bar{p}_\eta}^{(1)} \\ &\leq \frac{1}{1-\eta} + S_{\bar{p}_\eta}^{(1)}. \end{aligned}$$

Hence, $\|\bar{Z}_\eta\|_\infty \leq \frac{1}{1-\eta} + S_{\bar{p}_\eta}^{(1)}$. Similarly, $\|\bar{Z}_\eta\|_1 \leq \frac{1}{1-\eta} + S_{\bar{p}_\eta}^{(1)}$, which completes the proof. \square

When U and V are bounded, let l be the maximal length of a string in $U \cup V$. It can easily be shown that $\|\bar{Z}_\eta\|_2 \leq l + 1 + S_{\bar{p}_\eta}^{(1)}$ and therefore, in that case,

$$\|\bar{Z}_\eta\|_2 \leq \text{Min}(l + 1, \frac{1}{1-\eta}) + S_{\bar{p}_\eta}^{(1)} \quad (2)$$

which holds even if $\eta = 1$.

Lemma 4. $|\mathbb{E}(\bar{Z}_\eta[u, v] \bar{Z}_\eta[u', v])| \leq \bar{p}_\eta(u'v)$, for any $u, u', v \in \Sigma^*$.

Proof. We have $\mathbb{E}((\mathbf{1}_{\xi=w} - p(w))(\mathbf{1}_{\xi=w'} - p(w'))) = \mathbb{E}(\mathbf{1}_{\xi=w} \mathbf{1}_{\xi=w'} - p(w)p(w'))$. Therefore,

$$\begin{aligned} &\mathbb{E}(\bar{Z}_\eta[u, v] \bar{Z}_\eta[u', v]) \\ &= \sum_{x, x'} \eta^{|xx'|} [\mathbb{E}(\mathbf{1}_{\xi=uvx} \mathbf{1}_{\xi=u'vx'}) - p(u'vx')p(uvx)] \\ &= \sum_{x, x', w} \eta^{|xx'|} p(w) \mathbf{1}_{w=u'vx'} [\mathbf{1}_{w=uvx} - p(uvx)] \\ &= \sum_{x, x'} \eta^{|xx'|} p(u'vx') [\mathbf{1}_{u'vx'=uvx} - p(uvx)] \\ &= \sum_{x'} \eta^{|x'|} p(u'vx') [\sum_x \eta^{|x|} (\mathbf{1}_{u'vx'=uvx} - p(uvx))] \end{aligned}$$

and $|\mathbb{E}(\bar{Z}_\eta[u, v] \bar{Z}_\eta[u', v])| \leq \sum_{x'} \eta^{|x'|} p(u'vx') = \bar{p}_\eta(u'v)$ since $-1 \leq -\bar{p}_\eta(uv) \leq \sum_x \eta^{|x|} (\mathbf{1}_{u'vx'=uvx} - p(uvx)) \leq 1$ i.e. $|\sum_x \eta^{|x|} (\mathbf{1}_{u'vx'=uvx} - p(uvx))| \leq 1$. \square

Lemma 5.

$$\|\mathbb{E}(\bar{Z}_\eta \bar{Z}_\eta^T)\|_2 \leq S_{\bar{p}_\eta}^{(2)} \text{ and } \text{Tr}(\mathbb{E}(\bar{Z}_\eta \bar{Z}_\eta^T)) \leq S_{\bar{p}_\eta}^{(2)}.$$

$$\|\mathbb{E}(\bar{Z}_\eta^T \bar{Z}_\eta)\|_2 \leq S_{\bar{p}_\eta}^{(2)} \text{ and } \text{Tr}(\mathbb{E}(\bar{Z}_\eta^T \bar{Z}_\eta)) \leq S_{\bar{p}_\eta}^{(2)}.$$

Proof. Indeed,

$$\begin{aligned} \|\mathbb{E}(\bar{Z}_\eta \bar{Z}_\eta^T)\|_\infty &\leq \text{Max}_u \sum_{u', v} |\mathbb{E}(\bar{Z}_\eta[u, v] \bar{Z}_\eta[u', v])| \\ &\leq \sum_{u', v, x'} \eta^{|x'|} p(u'vx') \leq S_{\bar{p}_\eta}^{(2)}. \end{aligned}$$

In the same way,

$$\text{Tr}(\mathbb{E}(\bar{Z}_\eta \bar{Z}_\eta^T)) = \sum_{u, v} \mathbb{E}(\bar{Z}_\eta[u, v] \bar{Z}_\eta[u, v]) \leq S_{\bar{p}_\eta}^{(2)}.$$

Similar computations provide all the inequalities. \square

Therefore, we can apply the Theorem 1 with $b = \frac{1}{1-\eta} + S_{\bar{p}_\eta}^{(1)}$, $\sigma^2 = S_{\bar{p}_\eta}^{(2)}$ and $k = 2$.

Theorem 3. Let p be a rational stochastic language, let S be a sample of N strings drawn i.i.d. from p and let $0 \leq \eta < 1$. For all $t > 0$,

$$\text{Pr} \left[\|\bar{H}_{\eta, S}^{U, V} - H_{\bar{p}_\eta}^{U, V}\|_2 > \sqrt{\frac{2S_{\bar{p}_\eta}^{(2)}t}{N}} + \frac{t}{3N} \left[\frac{1}{1-\eta} + S_{\bar{p}_\eta}^{(1)} \right] \right] \leq 2t(e^t - t - 1)^{-1}.$$

Remark that when $\eta = 0$ we find back the concentration bound of Theorem 2, and that Inequality 2 provides a bound when $\eta = 1$.

3.3. Bound for the factor Hankel Matrix $H_{\bar{p}^{U, V}}$

The random matrix $\hat{Z}(\xi) = \hat{H}_\xi^{U, V} - H_{\bar{p}^{U, V}}$ is defined by

$$\hat{Z}_{u, v} = \sum_{x, y \in \Sigma^*} \mathbf{1}_{\xi=xuvy} - \hat{p}(uv).$$

$\|\hat{Z}\|_2$ is generally unbounded. Moreover, unlike the prefix case, $\|\hat{Z}\|_2$ can be unbounded even if U and V are finite. Hence, the Theorem 1 cannot be directly applied either.

We can also define smoothed variants of \hat{p} by

$$\hat{p}_\eta(u) = \sum_{x, y \in \Sigma^*} \eta^{|xy|} p(xuy) = \sum_{m, n \geq 0} \eta^{m+n} p(\Sigma^m u \Sigma^n)$$

which have properties similar to functions \bar{p}_η :

- $p \leq \hat{p}_\eta \leq \hat{p}$, $\hat{p}_1 = \hat{p}$ and $\hat{p}_0 = p$,
- if $\langle I, (M_x)_{x \in \Sigma}, T \rangle$ be a minimal linear representation of p then $\langle \hat{I}_\eta, (M_x)_{x \in \Sigma}, \bar{T}_\eta \rangle$, where $\hat{I}_\eta = (I_d - \eta M_\Sigma^T)^{-1} I$, is a linear representation of \hat{p}_η .

However, proofs of the previous Section cannot be directly extended to \hat{p}_η because \bar{p} is bounded by 1, a property which is often used in the proofs, while \hat{p} is not. Next lemma provides a tool which allows to bypass this difficulty.

Lemma 6. Let $0 < \eta \leq 1$. For any integer n , $(n+1)\eta^n \leq K_\eta$ where

$$K_\eta = \begin{cases} 1 & \text{if } \eta \leq e^{-1} \\ (-e\eta \ln \eta)^{-1} & \text{otherwise.} \end{cases}$$

Proof. Let $f(x) = (x+1)\eta^x$. We have $f'(x) = \eta^x(1 + (x+1)\ln\eta)$ and f takes its maximum for $x_M = -1 - 1/\ln\eta$, which is positive if and only if $\eta > 1/e$. We have $f(x_M) = (-e\eta \ln\eta)^{-1}$. \square

Lemma 7. *Let $w, u \in \Sigma^*$. Then,*

$$\sum_{x,y \in \Sigma^*} \eta^{|xy|} \mathbf{1}_{w=xy} \leq K_\eta \text{ and } \widehat{p}(u) \leq K_\eta p(\Sigma^* u \Sigma^*).$$

Proof. Indeed, if $w = xuy$, then $|xy| = |w| - |u|$ and u appears at most $|w| - |u| + 1$ times as a factor of w .

$$\begin{aligned} \widehat{p}(u) &= \sum_{x,y \in \Sigma^*} \eta^{|xy|} p(xuy) \\ &= \sum_{w \in \Sigma^* u \Sigma^*} p(w) \sum_{x,y \in \Sigma^*} \eta^{|xy|} \mathbf{1}_{w=xuy} \\ &\leq K_\eta p(\Sigma^* u \Sigma^*). \end{aligned}$$

\square

For $\eta \in [0, 1]$, let $\widehat{Z}_\eta(\xi)$ be the random matrix defined by

$$\begin{aligned} \widehat{Z}_\eta[u, v] &= \sum_{x,y \in \Sigma^*} \eta^{|xy|} \mathbf{1}_{\xi=xuy} - \widehat{p}_\eta(uv) \\ &= \sum_{x,y \in \Sigma^*} \eta^{|xy|} (\mathbf{1}_{\xi=xuy} - p(xuy)). \end{aligned}$$

and, for any $k \geq 0$, let

$$S_{\widehat{p}_\eta}^{(k)} = \sum_{u_1 u_2 \dots u_k \in \Sigma^*} \widehat{p}_\eta(u_1 u_2 \dots u_k).$$

It can easily be shown that $\mathbb{E}(\widehat{Z}_\eta) = 0$, $S_{\widehat{p}_\eta}^{(k)} = I^T (I_d - \eta M_\Sigma)^{-1} (I_d - M_\Sigma)^{-k} (I_d - \eta M_\Sigma)^{-1} T$, $S_{\widehat{p}_0}^{(k)} = S_p^{(k)}$ and $S_{\widehat{p}_1}^{(k)} = S_p^{(k+2)}$.

It can be shown that $\|\widehat{Z}_\eta\|_2$ is bounded if $\eta < 1$.

Lemma 8. (Denis et al., 2013)

$$\|\widehat{Z}_\eta\|_2 \leq (1-\eta)^{-2} + S_{\widehat{p}_\eta}^{(1)}.$$

Eventually, we can apply the Theorem 1 with $b = (1-\eta)^{-2} + S_{\widehat{p}_\eta}^{(1)}$, $\sigma^2 = K_\eta S_{\widehat{p}_\eta}^{(2)}$ and $k = 2$ (Denis et al., 2013).

Theorem 4. *Let p be a rational stochastic language, let S be a sample of N strings drawn i.i.d. from p and let $0 \leq \eta < 1$. For all $t > 0$,*

$$\begin{aligned} \Pr \left[\|\widehat{H}_{\eta, S}^{U, V} - H_{\widehat{p}_\eta}^{U, V}\|_2 > \sqrt{\frac{2K_\eta S_{\widehat{p}_\eta}^{(2)} t}{N}} + \frac{t}{3N} \left[\frac{1}{(1-\eta)^2} + S_{\widehat{p}_\eta}^{(1)} \right] \right] \\ \leq 2t(e^t - t - 1)^{-1}. \end{aligned}$$

Remark that when $\eta = 0$ we find back the concentration bound of Theorem 2. We provide experimental evaluation of the proposed bounds in the next Section.

4. Experiments

The proposed bounds are evaluated on the benchmark of PAutomaC (Verwer et al., 2012) which provides samples of strings generated from several probabilistic automata, designed to evaluate probabilistic automata learning. Eleven problems have been selected from that benchmark for which sparsity of the Hankel matrices makes the use of standard SVD algorithms available from NumPy or SciPy possible. The size N of the samples is 20,000 except for the problem 4 where $N = 100,000$. Table 1 shows some target properties of the selected problems: the size of the alphabets and the exact values of $S_p^{(k)}$ computed for the different targets p . Figure 1 shows the typical behavior of $S_{\widehat{p}_\eta}^{(1)}$ and $S_{\widehat{p}_\eta}^{(1)}$, similar for all the problems.




Figure 1. Behavior of $S_{\widehat{p}_\eta}^{(1)}$ and $S_p^{(1)}$ for $\eta \in [0; 1]$.

For each problem, the exact value of $\|H_S^{U, V} - H_p^{U, V}\|_2$ is computed for sets U and V of the form $\Sigma^{\leq l}$, trying to maximize l according to our computing resources. It is compared to the bounds provided by Theorem 2 and Equation (1), with $\delta = 0.05$ (Table 2). The optimized bound ("opt."), refers to the case where σ^2 has been calculated over $U \times V$ rather than $\Sigma^* \times \Sigma^*$ (see the remark at the end of Section 3.1). Tables 3 and 4 show analog comparisons for the prefix and the factor cases with different values of η . Similar results have been obtained for all the problems of PautomaC. We can remark that our dimension-free bounds are significantly more accurate than the one provided by Equation (1). Notice that in the prefix case, the dimension-free bound has a better behavior in the limit case $\eta = 1$ than the bound from Eq. (1). This is due to the fact that in our bound, the term that bounds $\|Z\|_2$ appears in the $\frac{1}{N}$ term while it appears in the $\frac{1}{\sqrt{N}}$ term in the other one.

Additional experiments confirm the implications of these results for spectral learning (see (Denis et al., 2013)).

Table 1. Properties of the problem set.

Problem number	3	4	7	15	25	29	31	38	39	40	42
$ \Sigma $	4	4	13	14	10	6	5	10	14	14	9
$S_p^{(2)}$	8.23	6.25	6.52	13.40	10.65	6.35	6.97	8.09	8.82	9.74	7.39
$S_p^{(3)}$	57.84	31.06	29.61	160.92	93.34	38.11	43.53	65.87	90.81	111.84	62.11

 Table 2. Concentration values from various bounds for $\|H_S^{U,V} - H_p^{U,V}\|_2$ for $U = V = \Sigma^{\leq l}$.

Problem number	3	4	7	15	25	29	31	38	39	40	42
l	8	9	8	5	5	9	7	4	6	4	7
$\ H_S^{U,V} - H_p^{U,V}\ _2$	0.0052	0.0030	0.0064	0.0037	0.0033	0.0045	0.0051	0.0058	0.0049	0.0037	0.0054
Eq. (1)	0.1910	0.0857	0.1917	0.1909	0.1935	0.1908	0.1911	0.1852	0.1925	0.1829	0.1936
Th. 2 (dim. free)	0.0669	0.0260	0.0595	0.0853	0.0761	0.0588	0.0615	0.0663	0.0692	0.0728	0.0634
Th. 2 (opt. U, V)	0.0475	0.0228	0.0527	0.0284	0.0323	0.0472	0.0437	0.0275	0.0325	0.0243	0.0378

 Table 3. Concentration values from various bounds for $\|\overline{H}_S^{U,V} - H_{\overline{p},\eta}^{U,V}\|_2$ (prefix case) for $U = V = \Sigma^{\leq l}$ with $\eta = \frac{1}{2}$ for the first four lines and $\eta = 1$ for the last four lines.

Problem number	3	4	7	15	25	29	31	38	39	40	42
l	8	9	8	5	5	9	7	4	6	4	7
$\ \overline{H}_S^{U,V} - H_{\overline{p},\eta}^{U,V}\ _2$	0.0067	0.0035	0.0085	0.0043	0.0041	0.0055	0.0073	0.0059	0.0061	0.0044	0.0062
Eq. (1)	0.7463	0.3326	0.7545	0.7515	0.7626	0.7250	0.7369	0.7051	0.7068	0.6753	0.7146
Th. 3 (dim. free)	0.0890	0.0339	0.0777	0.1162	0.1026	0.0770	0.0811	0.0884	0.0931	0.0983	0.0844
Th. 3 (opt. U, V)	0.0636	0.0299	0.0697	0.0398	0.0457	0.0621	0.0577	0.0366	0.0432	0.0317	0.0498
$\ \overline{H}_S^{U,V} - H_{\overline{p},\eta}^{U,V}\ _2$	0.0141	0.0059	0.0217	0.0124	0.0145	0.0116	0.0182	0.0132	0.0135	0.0089	0.0127
Eq. (1)	3.1011	1.3079	2.7839	3.5129	3.0283	2.9286	2.6695	2.2395	2.8524	2.5132	2.7863
Th. 3 (dim. free)	0.1784	0.0582	0.1279	0.2967	0.2261	0.1450	0.1547	0.1899	0.2230	0.2472	0.1846
Th. 3 (opt. U, V)	0.1281	0.0518	0.1166	0.1062	0.1057	0.1175	0.1099	0.0778	0.1020	0.0761	0.1077

 Table 4. Concentration values from various bounds for $\|\widehat{H}_S^{U,V} - H_{\widehat{p},\eta}^{U,V}\|_2$ (factor case) for $U = V = \Sigma^{\leq l}$ and $\eta = 1/e$.

Problem number	3	4	7	15	25	29	31	38	39	40	42
l	6	7	5	4	4	6	6	4	4	4	5
$\ \widehat{H}_S^{U,V} - H_{\widehat{p},\eta}^{U,V}\ _2$	0.0065	0.0031	0.0071	0.0042	0.0033	0.0051	0.0072	0.0061	0.0065	0.0047	0.0060
Eq. (1)	0.9134	0.4107	0.9196	0.9466	0.9152	0.9096	0.9219	0.8765	0.8292	0.8796	0.8565
Th. 4 (dim. free)	0.0985	0.0374	0.0858	0.1292	0.1139	0.0849	0.0895	0.0979	0.1033	0.1092	0.0934
Th. 4 (opt. U, V)	0.0601	0.0300	0.0619	0.0364	0.0412	0.0559	0.0589	0.0405	0.0356	0.0349	0.0444

5. Conclusion

We have provided dimension-free concentration inequalities for Hankel matrices in the context of spectral learning of rational stochastic languages. These bounds cover 3 cases, each one corresponding to a specific way to exploit the strings under observation, paying attention to the strings themselves, to their prefixes or to their factors. For the last two cases, we introduced parametrized variants which allow a trade-off between the rate of the concentration and the exploitation of the information contained in data.

A consequence of these results is that there is no a priori good reason, aside from computing resources limitations, to restrict the size of the Hankel matrices. This suggests an immediate future work consisting in investigating recent random techniques (Halko et al., 2011) to compute singular values decomposition on Hankel matrices in order to be

able to deal with huge matrices. Then, a second aspect is to evaluate the impact of these methods on the quality of the models, including an empirical evaluation of the behavior of the standard approach and its prefix and factor extensions, along with the influence of the parameter η .

Another research direction would be to link up the prefix and factor cases to concentration bounds for sum of random tensors and to generalize the results to the case where a fixed number ≥ 1 of factors is considered for each string.

Acknowledgments

This work was supported by the French National Agency for Research (Lampada - ANR-09-EMER-007).

References

- Anandkumar, A., Foster, D.P., Hsu, D., Kakade, S., and Liu, Y.-K. A spectral algorithm for latent dirichlet allocation. In *Proceedings of NIPS*, pp. 926–934, 2012a.
- Anandkumar, A., Hsu, D., Huang, F., and Kakade, S. Learning mixtures of tree graphical models. In *Proceedings of NIPS*, pp. 1061–1069, 2012b.
- Anandkumar, A., Hsu, D., and Kakade, S.M. A method of moments for mixture models and hidden markov models. *Proceedings of COLT - Journal of Machine Learning Research - Proceedings Track*, 23:33.1–33.34, 2012c.
- Bailly, R. *Méthodes spectrales pour l'inférence grammaticale probabiliste de langages stochastiques rationnels*. PhD thesis, Aix-Marseille Université, 2011.
- Bailly, R., Denis, F., and Ralaivola, L. Grammatical inference as a principal component analysis problem. In *Proceedings of ICML*, pp. 5, 2009.
- Bailly, R., Habrard, A., and Denis, F. A spectral approach for probabilistic grammatical inference on trees. In *Proceedings of ALT*, pp. 74–88, 2010.
- Balle, B. and Mohri, M. Spectral learning of general weighted automata via constrained matrix completion. In *Proceedings of NIPS*, pp. 2168–2176, 2012.
- Balle, B., Quattoni, A., and Carreras, X. A spectral learning algorithm for finite state transducers. In *Proceedings of ECML/PKDD (1)*, pp. 156–171, 2011.
- Balle, B., Quattoni, A., and Carreras, X. Local loss optimization in operator models: A new insight into spectral learning. In *Proceedings of ICML*, 2012.
- Balle, B., Carreras, X., Luque, F. M., and Quattoni, A. Spectral learning of weighted automata: A forward-backward perspective. To appear in *Machine Learning*, 2013.
- Cohen, Shay B., Stratos, Karl, Collins, Michael, Foster, Dean P., and Ungar, Lyle H. Spectral learning of Latent-Variable PCFGs. In *ACL (1)*, pp. 223–231. The Association for Computer Linguistics, 2012. ISBN 978-1-937284-24-4.
- Denis, F., Gybels, M., and Habrard, A. Dimension-free Concentration Bounds on Hankel Matrices for Spectral Learning (Long Version). arXiv:1312.6282, 2013.
- Droste, M., Kuich, W., and Vogler, H. (eds.). *Handbook of Weighted Automata*. Springer, 2009.
- Halko, N., Martinsson, P.G., and Tropp, J.A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.
- Hsu, D., Kakade, S.M., and Zhang, T. A spectral algorithm for learning hidden markov models. In *Proceedings of COLT*, 2009.
- Hsu, D., Kakade, S. M., and Zhang, T. Dimension-free tail inequalities for sums of random matrices. *ArXiv e-prints*, 2011.
- Kakade, S. Multivariate analysis, dimensionality reduction, and spectral methods. Lecture Notes (Matrix Concentration Derivations), 2010.
- Luque, F.M., Quattoni, A., Balle, B., and Carreras, X. Spectral learning for non-deterministic dependency parsing. In *Proceedings of EACL*, pp. 409–419, 2012.
- Parikh, A.P., Song, L., and Xing, E.P. A spectral algorithm for latent tree graphical models. In *Proceedings of ICML*, pp. 1065–1072, 2011.
- Siddiqi, S., Boots, B., and Gordon, G.J. Reduced-rank hidden Markov models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS-2010)*, 2010.
- Song, L., Boots, B., Siddiqi, S.M., Gordon, G.J., and Smola, A.J. Hilbert space embeddings of hidden markov models. In *Proceedings of ICML*, pp. 991–998, 2010.
- Stewart, G. W. Perturbation theory for the singular value decomposition. In *SVD and Signal Processing II: Algorithms, Analysis and Applications*, pp. 99–109. Elsevier, 1990.
- Tropp, J.A. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.
- Verwer, S., Eyraud, R., and de la Higuera, C. Results of the PAutomaC probabilistic automaton learning competition. *Journal of Machine Learning Research - Proceedings Track*, 21:243–248, 2012.