

Heavy-Tailed Analogues of the Covariance Matrix for ICA

Joseph Anderson
The Ohio State University
andejose@cse.ohio-state.edu

Navin Goyal
Microsoft Research
navingo@microsoft.com

Anupama Nandi
The Ohio State University
nandi.10@osu.edu

Luis Rademacher
University of California, Davis
lrademac@ucdavis.edu

Abstract

Independent Component Analysis (ICA) is the problem of learning a square matrix A , given samples of $X = AS$, where S is a random vector with independent coordinates. Most existing algorithms are provably efficient only when each S_i has finite and moderately valued fourth moment. However, there are practical applications where this assumption need not be true, such as speech and finance. Algorithms have been proposed for heavy-tailed ICA, but they are not practical, using random walks and the full power of the ellipsoid algorithm multiple times. The main contributions of this paper are:

- (1) A practical algorithm for heavy-tailed ICA that we call HTICA. We provide theoretical guarantees and show that it outperforms other algorithms in some heavy-tailed regimes, both on real and synthetic data. Like the current state-of-the-art, the new algorithm is based on the centroid body (a first moment analogue of the covariance matrix). Unlike the state-of-the-art, our algorithm is practically efficient. To achieve this, we use explicit analytic representations of the centroid body, which bypasses the use of the ellipsoid method and random walks.
- (2) We study how heavy tails affect different ICA algorithms, including HTICA. Somewhat surprisingly, we show that some algorithms that use the covariance matrix or higher moments can successfully solve a range of ICA instances with infinite second moment. We study this theoretically and experimentally, with both synthetic and real-world heavy-tailed data.

1 Introduction

Independent component analysis (ICA) is a computational and statistical technique with applications in areas ranging from signal processing to machine learning and more. Formally, if S is an n -dimensional random vector with independent coordinates and $A \in \mathbb{R}^{n \times n}$ is invertible, then the ICA problem is to estimate A given access to i.i.d. samples of the mixed signals $X = AS$. We say that X is generated by an *ICA model* $X = AS$. The recovery of A (the *mixing matrix*) is possible only up to scaling and permutation of the columns. Moreover, for the recovery to be possible, the distributions of the random variables S_i must not be Gaussian (except possibly one of them). Since its inception in the eighties (see (Comon and Jutten 2010) for historical

remarks), ICA has been thoroughly studied and a vast literature exists (e.g. (Hyvarinen, Karhunen, and Oja 2001; Comon and Jutten 2010)). The theory is well-developed and *practical* algorithms—e.g., FastICA (Hyvarinen 1999), JADE (Cardoso and Souloumiac 1993)—are now available along with implementations, e.g. (Cichocki et al.). However, to our knowledge, rigorous complexity analyses of these assume that the fourth moment of each component is finite: $\mathbb{E}(S_i^4) < \infty$. If at least one of the independent components does not satisfy this assumption we will say that the input is in the *heavy-tailed regime*. Many ICA algorithms first pre-process the data to convert the given ICA model into another one where the mixing matrix A has orthogonal columns; this step is often called *whitening*. We will instead call it *orthogonalization*, as this describes more precisely the desired outcome. Traditional whitening is a second order method that may not make sense in the heavy-tailed regime. In this regime, it is not clear how the existing algorithms would perform, because they depend on empirical estimation of various statistics of the data such as the covariance matrix or the fourth cumulant tensor, which diverge in general for heavy-tailed data. For example, for the covariance matrix in the mean-0 case this is done by taking the empirical average $(1/N) \sum_{i=1}^N x(i)x(i)^T$ where the $\{x(i)\}$ are i.i.d. samples of X . ICA in the heavy-tailed regime is of considerable interest, directly (e.g., (Kidmose 2001b; 2001a; Shereshevski, Yeredor, and Messer 2001; Chen and Bickel 2004; 2005; Sahmoudi et al. 2005; Wang, Kuruoglu, and Zhang 2009; J. Eriksson and Koivunen 2001; Clemencon and Slim 2007; Bermond and Cardoso 1999)) and indirectly (e.g., (Bickson and Guestrin 2010; Gael, Teh, and Ghahramani 2009; Welling, Osindero, and Hinton 2002)) and has applications in speech and finance. We also mention an informal connection with robust statistics: Algorithms solving heavy-tailed ICA might work by focusing on samples in a small (but high probability) region to get reliable statistics about the data and avoid the instability of the tail. Thus, if the data has outliers, the outliers are less likely to affect such an algorithm.

Recent theoretical work (Anderson et al. 2015) proposed a polynomial time algorithm for ICA that works in the regime where each component S_i has finite $(1 + \gamma)$ -moment for $\gamma > 0$. This algorithm follows the two phases of several ICA algorithms: (i) Orthogonalize the independent components. The purpose of this step is to apply an affine transformation

to the samples from X so that the resulting samples correspond to an ICA model where the unknown matrix A has orthogonal columns. (ii) Learn the matrix with orthogonal columns. Each of these two phases required new techniques: (1) *Orthogonalization via uniform distribution in the centroid body*. The input is assumed to be samples from an ICA model $X = AS$ where each S_i is symmetrically distributed (w.l.o.g, see Sec. 2) and has at least $(1 + \gamma)$ -moments. The goal is to construct an *orthogonalization* matrix B so that BA has orthogonal columns. In (Anderson et al. 2015), the inverse of the square root of the covariance matrix of the uniform distribution in the centroid body is one such matrix. (2) *Gaussian damping*. The previous step allows one to assume that the mixing matrix A is orthogonal. The modified second step is: If X has density $\rho_X(t)$ for $t \in \mathbb{R}^n$, then the algorithm constructs another ICA model $X_R = AS_R$ where X_R has pdf proportional to $\rho_X(t) \exp(-\|t\|_2^2/R^2)$, where $R > 0$ is a parameter chosen by the algorithm. This explains the term Gaussian damping. This achieves two goals: (1) All moments of X_R and S_R are finite. (2) The product structure of is retained. This follows from two facts: A has orthogonal columns, and the Gaussian has independent components in any orthonormal basis. Because of these properties, the model can be solved by traditional ICA algorithms.

The algorithm in (Anderson et al. 2015) is theoretically efficient but impractical. Their orthogonalization uses the ellipsoid algorithm for linear programming, which is not practical. It is not clear how to replace their use of the ellipsoid algorithm by practical linear programming tools, as their algorithm only has oracle access to a sort of dual and not an explicit linear program. Moreover, their orthogonalization technique uses samples uniformly distributed in the centroid body, generated by a random walk. This is computationally efficient in theory but, to the best of our knowledge, only efficient in practice for moderately low dimension.

Our contributions. Our contributions are experimental and theoretical. We provide a new and practical ICA algorithm, HTICA, building upon the previous theoretical work in (Anderson et al. 2015). HTICA works as follows: (1) Compute an orthogonalization matrix B . (2) Pre-multiply samples by B to get an orthogonal model. (3) Damp the data, run an existing ICA algorithm. For step (1), we propose two theoretically sound and practically efficient ways below, *orthogonalization via centroid body scaling* and *orthogonalization via covariance*. Our algorithm is simpler and more efficient, but needs a more technical justification than the method in (Anderson et al. 2015). We demonstrate the effectiveness of HTICA on both synthetic and real-world data.

Orthogonalization via centroid body scaling. We propose a more practical orthogonalization matrix than the one from (Anderson et al. 2015) (orthogonalization via the uniform distribution in the centroid body, mentioned before). First, consider the centroid body of random vector X , denoted ΓX (this is really a function of the *distribution* of X ; formal definition in Sec. 2). For intuition, it is helpful to think of the centroid body as an ellipsoid whose axes are aligned with the independent components of X . The centroid body is in general not an ellipsoid, but it has certain symmetries aligned with the independent components. Let random vector Y be

a scaling of X along every ray so that points at infinity are mapped to the boundary of ΓX , the origin is mapped to itself and the scaling interpolates smoothly. One such scaling is obtained in the following way: It is helpful to consider how far a point is in its ray with respect to the boundary of ΓX . This is given by the *Minkowski functional* of ΓX , denoted $p : \mathbb{R}^n \rightarrow \mathbb{R}$, which maps the boundary of ΓX to 1 and interpolates linearly along every ray. We can then achieve the desired scaling by first mapping a given point to the boundary point on its ray (the mapping $x \mapsto x/p(x)$) and then using the function \tanh , which maps $[0, \infty)$ to $[0, 1]$ with $\tanh(0) = 0$ and $\lim_{x \rightarrow \infty} \tanh(x) = 1$ to determine the final scale along the ray, namely, $\tanh p(x)$. More formally, our scaling is the following: Let Y be $\frac{\tanh p(X)}{p(X)} X$. We show in Sec. 4.1 that $B = \text{Cov}(Y)^{-1/2}$ is an orthogonalization matrix when $\text{Cov}(Y)$ is invertible. In order to make this practical, one needs a practical estimator of the Minkowski functional of ΓX from a sample of X . In Sec. 4.1 and 5, we describe such an algorithm and provide a theoretical justification, including finite sample estimates. The proposed algorithm is much simpler and practical than the one described in (Anderson et al. 2015). In particular, it avoids the use of the ellipsoid algorithm by the use of a closed-form linear programming representation of the centroid body (Prop. 4, Lemma 5) and new approximation guarantees between the empirical (sample estimate) and true centroid body of a heavy-tailed distribution. In Sec. 4.1, we discuss our practical implementation and show results where orthogonalization via centroid body scaling produces results with smaller error.

Orthogonalization via covariance. Previously, (e.g., in (Chen and Bickel 2004)), the empirical covariance matrix was used for whitening in the heavy-tailed regime and, surprisingly, worked well in some situations. Unfortunately, the understanding of this was quite limited. We give a theoretical explanation for this phenomenon in a fairly general heavy-tailed regime: Covariance-based orthogonalization works well when each component S_i has finite $(1 + \gamma)$ -moment, where $\gamma > 0$. We also study this algorithm in experimental settings. As we will see, while orthogonalization via covariance improves over previous algorithms, in general orthogonalization via centroid body has better performance because it has better numerical stability; but there are some situations where orthogonalization via covariance matrix is better.

Empirical Study. We perform experiments on both synthetic and real data to see the effect of heavy-tails on ICA.

In the synthetic data setting, we generate samples from a fixed heavy-tailed distribution and study how well the algorithm can recover a random mixing matrix (Sec. 3).

To study the algorithm with real data, we use recordings of human speech provided by (Donohue 2009). This involves a room with different arrangements of microphones, and six humans speaking independently. The speakers are recorded individually, so we can artificially mix them and have access to a ground truth. We study the statistical properties of the data, observing that it does indeed behave as if the underlying processes are heavy-tailed. The performance of our algorithm shows improvement over using FastICA on its own.

See supplementary material for proofs and extra detail.

2 Preliminaries

Heavy-tailed distributions arise in a wide range of applications (e.g., (Nolan 2015)). They are characterized by the slow decay of their tails. Examples of heavy-tailed distributions include the Pareto and log-normal distributions.

We denote the pdf of random variable Z by ρ_Z . We will assume that our distributions are symmetric, that is $\rho(x) = \rho(-x)$ for $x \in \mathbb{R}$. As observed in (Anderson et al. 2015), this is without loss of generality for our purposes. This follows from the fact that if $X = AS$ is an ICA model, and if we let $X' = AS'$ be an i.i.d. copy of the same model, then $X - X' = A(S - S')$ is an ICA model with components of $S - S'$ having symmetric pdfs. One further needs to check that if the components of S are away from Gaussians then the same holds for $S - S'$; see (Anderson et al. 2015). We formulate our algorithms for the symmetric case; the general case immediately reduces to the symmetric case.

For $K \subseteq \mathbb{R}^n$, K_ϵ denotes the set of points that are at distance at most ϵ from K . The set $K_{-\epsilon}$ is all points for which an ϵ -ball around them is still contained in K .

An important related family of distributions is that of stable distributions (e.g., (Nolan 2015)). In general, the density of a stable distribution has no closed form, but is fully defined by four real-valued parameters. Some stable distributions do admit a closed form, such as the Cauchy and Gaussian distributions. For us the most important parameter is $\alpha \in (0, 2]$, known as the stability parameter; we will think of the other three parameters as being fixed to constants.

We use the notation $\text{poly}(\cdot)$ to indicate a function which is asymptotically upper bounded by a polynomial expression of the given variables.

Definition 1 (Centroid body). *Let $X \in \mathbb{R}^n$ be a random vector with finite first moment, that is, for all $u \in \mathbb{R}^n$ we have $\mathbb{E}(|\langle u, X \rangle|) < \infty$. The centroid body of X is the compact convex set, denoted ΓX , whose support function is $h_{\Gamma X}(u) = \mathbb{E}(|\langle u, X \rangle|)$. For a probability measure \mathbb{P} , we define $\Gamma \mathbb{P}$, the centroid body of \mathbb{P} , as the centroid body of any random vector distributed according to \mathbb{P} .*

Note that for the centroid body to be well-defined, the mean of the data must be finite. This excludes, for instance, the Cauchy distribution from consideration in the present work.

3 HTICA and experiments

In this section, we show experimentally that heavy-tailed data poses a significant challenge for current ICA algorithms, and compare them with HTICA in different settings. We observe some clear situations where heavy-tails seriously affect the standard ICA algorithms, and that these problems are frequently avoided by using the heavy-tailed ICA framework. In some cases, HTICA does not help much, but maintains the same performance of plain FastICA.

To generate the synthetic data, we create a simple heavy-tailed density function $f_\eta(x)$ proportional to $(|x| + 1.5)^{-\eta}$, which is symmetric, and for $\eta > 1$, f_η is the density of a distribution which has finite $k < \eta - 1$ moment. The signal S is generated with each S_i independently distributed from f_{η_i} . The mixing matrix $A \in \mathbb{R}^{n \times n}$ is generated with

each coordinate i.i.d. $\mathcal{N}(0, 1)$, columns normalized to unit length. To compare the quality of recovery, the columns of the estimated mixing matrix, \hat{A} are permuted to align with the closest matching column of A , via the Hungarian algorithm. We use the Frobenius norm to measure the error, but all experiments were also performed using the well-known Amari index (Amari et al. 1996); the results have similar behavior and are not presented here.

3.1 Heavy-tailed ICA when A is orthogonal: Gaussian damping and experiments

Focusing on the third step above, where the mixing matrix already has orthogonal columns, ICA algorithms already suffer dramatically from the presence of heavy-tailed data. As proposed in (Anderson et al. 2015), Gaussian damping is a preprocessing technique that converts data from an ICA model $X = AS$, where A is unitary (columns are orthogonal with unit l_2 -norm) to data from a related ICA model $X_R = AS_R$, where $R > 0$ is a parameter to be chosen. The independent components of S_R have finite moments of all orders and so the existing algorithms can estimate A .

Using samples of X , we construct the damped random variable X_R , with pdf $\rho_{X_R}(x) \propto \rho_X(x) \exp(-\|x\|^2/R^2)$. To normalize the right hand side, we can estimate

$$K_{X_R} = \mathbb{E} \exp(-\|X\|^2/R^2)$$

so that

$$\rho_{X_R}(x) = \rho_X(x) \exp(-\|x\|^2/R^2)/K_{X_R}.$$

If x is a realization of X_R , then $s = A^{-1}x$ is a realization of the random variable S_R and we have that S_R has pdf $\rho_{S_R}(s) = \rho_{X_R}(x)$. To generate samples from this distribution, we use rejection sampling on samples from ρ_X . When performing the damping, we binary search over R so that about 25% of the samples are rejected. For more details about the technical requirements for choosing R , see (Anderson et al. 2015).

Figure 1 shows that, when A is already a perfectly orthogonal matrix, but where S may have heavy-tailed coordinates, several standard ICA algorithms perform better after damping the data. In fact, without damping, some do not appear to converge to a correct solution. We compare ICA with and without damping in this case: (1) FastICA using the fourth cumulant (“FastICA - pow3”), (2) FastICA using log cosh (“FastICA - tanh”), (3) JADE, and (4) Second Order Joint Diagonalization as in, e.g., (Cardoso 1989).

3.2 Experiments on synthetic heavy-tailed data

We now present the results of HTICA using different orthogonalization techniques: (1) Orthogonalization via *covariance* (Section 4.2) (2) Orthogonalization via the *centroid* body (Section 4.1) (3) the ground truth, directly inverting the mixing matrix (*oracle*), and (4) No orthogonalization, and also no damping (for comparison with plain FastICA) (*identity*).

The “mixed” regime in the left and middle of Figure 2 (where some signals are *not* heavy-tailed) demonstrates a very dramatic contrast between different orthogonalization methods, even when only two heavy-tailed signals are present.




Figure 1: (Left, middle): The error of ICA with and without damping (solid lines and dashed lines, resp.), with unitary mixing matrix. The error is averaged over ten trials, in 3 and 10 dimensions where $\eta = (6, 6, 2.1)$ and $\eta = (6, \dots, 6, 2.1, 2.1)$, resp. (Right): The difference between the errors of FastICA with and without damping in 2 dimensions, averaged over 40 trials. For a single cell, the parameters are given by the coordinates, $\eta = (i, j)$. Red indicates that FastICA without damping does better than FastICA with damping, white indicates that the error difference is 0 and the blue indicates that FastICA with damping performs better than without damping. Black indicates that FastICA without damping failed (did not return two independent components).




Figure 2: The Frobenius error of the recovered mixing matrix with the ‘pow3’ and ‘tanh’ contrast functions, on 10-dimensional data, averaged over ten trials. The mixing matrix A is random with unit norm columns, not orthogonal. In the left and middle figures, the distribution has $\eta = (6, \dots, 6, 2.1, 2.1)$ while in the right figure, $\eta = (2.1, \dots, 2.1)$ (see Section 3.2 for a discussion).

In the experiment with different methods of orthogonalization it was observed that when all exponents are the same or very close, orthogonalization via covariance performs better than orthogonalization via centroid and the true mixing matrix as seen in Figure 2. A partial explanation is that, given the results in Figure 1, we know that equal exponents favor FastICA without damping and orthogonalization (*identity* in Figure 2). The line showing the performance with no orthogonalization and no damping (“identity”) behaves somewhat erratically, most likely due the presence of the heavy-tailed samples. Additionally, damping and the choice of parameter R is sensitive to scaling. A scaled-up distribution will be somewhat hurt because fewer samples will survive damping.

3.3 ICA on speech data


While the above study on synthetic data provides interesting situations where heavy-tails can cause problems for ICA, we provide some results here which use real-world data, specifically human speech. To study the performance of HTICA on voice data, we first examine whether the data is heavy-tailed. The motivation to use speech data comes from observations by the signal processing community (e.g. (Kidmose 2000))

that speech data can be modeled by α -stable distributions. For an α -stable distribution, with $\alpha \in (0, 2)$, only the moments of order less than α will be finite. We present here some results on a data set of human speech according to the standard cocktail party model, from (Donohue 2009).

To estimate whether the data is heavy-tailed, as in (Kidmose 2000), we estimate parameter α of a best-fit α -stable distribution. This estimate is in Figure 3 for one of the data sets collected. We can see that the estimated α is clearly in the heavy-tailed regime for some signals.

Using data from (Donohue 2009), we perform the same experiment as in Section 3.2: generate a random mixing matrix with unit length columns, mix the data, and try to recover the mixing matrix. Although the mixing is synthetic, the setting makes the resulting mixed signals same as real. Specifically, the experiment was conducted in a room with chairs, carpet, plasterboard walls, and windows on one side. There was natural noise including vents, computers, florescent lights, and traffic noise through the windows.

Figure 3 demonstrates that HTICA (orthogonalizing with centroid body scaling, Section 4.1) applied to speech data yields some noticeable improvement in the recovery of the



Signal	$\hat{\alpha}$
1	1.91
2	1.93
3	1.61
4	1.60
5	1.92
6	2.00

		σ_{\min}		$\sigma_{\max}/\sigma_{\min}$	
Orthogonalizer		Centroid	Covariance	Centroid	Covariance
Samples	1000	0.9302	0.9263	27.95	579.95
	3000	0.9603	0.9567	20.44	410.11
	5000	0.9694	0.9673	19.25	490.11
	7000	0.9739	0.9715	18.90	347.68
	9000	0.9790	0.9708	20.12	573.18
	11000	0.9793	0.9771	18.27	286.34

Figure 3: (Left): Error of estimated mixing matrix on the “perimeterreg” data, averaged over ten trials. HTICA is more robust than FastICA. (Middle): Stability parameter α estimates of each component in the “perimeterreg” data. Values below 2 are in the heavy-tailed regime. (Right): Smallest singular value and condition number of the orthogonalization matrix BA computed via the centroid body and the covariance. The data was sampled with parameter $\eta = (6, 6, 6, 6, 6, 6, 6, 6, 2.1, 2.1)$.

mixing matrix, primarily in that it is less susceptible to data that causes FastICA to have large error “spikes.” Moreover, in many cases, running only FastICA on the mixed data failed to even recover all of the speech signals, while HTICA succeeded. In these cases, we had to re-start FastICA until it recovered all the signals.

4 New approach to orthogonalization and a new analysis of empirical covariance

As noted above, the technique in (Anderson et al. 2015), while being provably efficient and correct, suffers from practical implementation issues. Here we discuss two alternatives: orthogonalization by *centroid body scaling* and orthogonalization by using the empirical covariance. The former, orthogonalization via centroid body scaling, uses the samples already present in the algorithm rather than relying on a random walk to draw samples which are approximately uniform in the algorithm’s approximation of the centroid body (as is done in (Anderson et al. 2015)). This removes the dependence on random walks and the ellipsoid algorithm; instead, we use samples that are distributed according to the original heavy-tailed distribution but *non-linearly scaled* to lie inside the centroid body. We prove in Lemma 2 that the covariance of this subset of samples is enough to orthogonalize the mixing matrix A . Secondly, we prove that one can, in fact, “forget” that the data is heavy tailed and orthogonalize by using the empirical covariance of the data, even though it diverges, and that this is enough to orthogonalize the mixing matrix A . However, as observed in experimental results, in general this has a downside compared to orthogonalization via centroid body in that it could cause numerical instability during the “second” phase of ICA as the data obtained is less well-conditioned. This is illustrated directly in the table in Figure 3 containing the singular value and condition number of the mixing matrix BA in the approximately orthogonal ICA model.

4.1 Orthogonalization via centroid body scaling

In (Anderson et al. 2015), another orthogonalization procedure, namely *orthogonalization via the uniform distribution*

in the centroid body is theoretically proven to work. Their procedure does not suffer from the numerical instabilities and composes well with the second phase of ICA algorithms. An impractical aspect of that procedure is that it needs samples from the uniform distribution in the centroid body.

We described orthogonalization via centroid body in Section 1, except for the estimation of $p(x)$, the Minkowski functional of the centroid body. The complete procedure is stated in Subroutine 1.

We now explain how to estimate the Minkowski functional. The Minkowski functional was informally described in Section 1. The Minkowski functional of ΓX is formally defined by $p(x) := \inf\{t > 0 : x \in t\Gamma X\}$. Our estimation of $p(x)$ is based on an explicit linear program (LP) (2) that gives the Minkowski functional of the centroid body of a finite sample of X *exactly* and then arguing that a sample estimate is close to the actual value for ΓX . For clarity of exposition, we only analyze formally a special case of LP (2) that decides *membership* in the centroid body of a finite sample of X (LP (1)) and approximate membership in ΓX . This analysis is in Section 5. Accuracy guarantees for the approximation of the Minkowski functional follow from this analysis.

Subroutine 1 Orthogonalization via centroid body scaling

Input: Samples $(X^{(i)})_{i=1}^N$ of ICA model $X = AS$ so each S_i is symmetric with $(1 + \gamma)$ moments.

Output: Matrix B approximate orthogonalizer of A

- 1: **for** $i = 1 : N$ **do**,
 - 2: Let λ^* be the optimal value of (2) with $q = X^{(i)}$. Let $d_i = 1/\lambda^*$. Let $Y^{(i)} = \frac{\tanh d_i}{d_i} X^{(i)}$.
 - 3: **end for**
 - 4: Let $C = \frac{1}{N} \sum_{i=1}^N Y^{(i)} Y^{(i)T}$. Output $B = C^{-1/2}$.
-

Lemma 2. *Let X be a random vector drawn from an ICA model $X = AS$ such that for all i we have $\mathbb{E}|S_i| = 1$ and S_i is symmetrically distributed. Let $Y = \frac{\tanh p(X)}{p(X)} X$ where $p(X)$ is the Minkowski functional of ΓX . Then $\text{Cov}(Y)^{-1/2}$ is an orthogonalizer of X .*

4.2 Orthogonalization via covariance

Here we show the somewhat surprising fact that orthogonalization of heavy-tailed signals is sometimes possible by using the “standard” approach: inverting the empirical covariance matrix. The advantage here, is that it is computationally very simple, specifically that having heavy-tailed data incurs very little computational penalty on the process of orthogonalization alone. It’s standard to use covariance matrix for *whitening* when the second moments of all independent components exist (Hyvarinen, Karhunen, and Oja 2001): Given samples from the ICA model $X = AS$, we compute the empirical covariance matrix $\tilde{\Sigma}$ which tends to the true covariance matrix as we take more samples and set $B = \tilde{\Sigma}^{-1/2}$. Then one can show that BA is a rotation matrix, and thus by pre-multiplying the data by B we obtain an ICA model $Y = BX = (BA)S$, where the mixing matrix BA is a rotation matrix, and this model is then amenable to various algorithms. In the heavy-tailed regime where the second moment does not exist for some of the components, there is no true covariance matrix and the empirical covariance diverges as we take more samples. However, for any fixed number of samples one can still compute the empirical covariance matrix. In previous work (e.g., (Chen and Bickel 2004)), the empirical covariance matrix was used for whitening in the heavy-tailed regime with good empirical performance; (Chen and Bickel 2004) also provided some theoretical analysis to explain this surprising performance. However, their work (both experimental and theoretical) was limited to some very special cases (e.g., only one of the components is heavy-tailed, or there are only two components both with stable distributions without finite second moment).

We will show that the above procedure (namely pre-multiplying the data by $B := \tilde{\Sigma}^{-1/2}$) “works” under considerably more general conditions, namely if $(1 + \gamma)$ -moment exists for $\gamma > 0$ for each independent component S_i . By “works” we mean that instead of whitening the data (that is BA is rotation matrix) it does something slightly weaker but still just as good for the purpose of applying ICA algorithms in the next phase. It *orthogonalizes* the data, that is now BA is close to a matrix whose columns are orthogonal. In other words, $(BA)^T(BA)$ is close to a diagonal matrix (in a sense made precise in Theorem 3).

Theorem 3 (Orthogonalization via covariance matrix). *Let X be given by ICA model $X = AS$. Assume that there exist $t, p, M > 0$ and $\gamma \in (0, 1)$ such that for all i we have*

- (a) $\mathbb{E}(|S_i|^{1+\gamma}) \leq M < \infty$,
- (b) (normalization) $\mathbb{E}|S_i| = 1$, and
- (c) $\Pr(|S_i| \geq t) \geq p$. Let $x^{(1)}, \dots, x^{(N)}$ be i.i.d. samples according to X . Let $\tilde{\Sigma} = (1/N) \sum_{k=1}^N x^{(k)} x^{(k)T}$ and $B = \tilde{\Sigma}^{-1/2}$. Then for any $\epsilon, \delta \in (0, 1)$, $\|(BA)^T BA - D\|_2 \leq \epsilon$ for a diagonal matrix D with diagonal entries d_1, \dots, d_n satisfying $0 < d_i, 1/d_i \leq \max\{2/pt^2, N^4\}$ for all i with probability $1 - \delta$ when $N \geq \text{poly}(n, M, 1/p, 1/t, 1/\epsilon, 1/\delta)$.

In Theorem 3, the diagonal entries are lower bounded, which avoids some degeneracy, but they could still grow quite large because of the heavy tails. This is a real drawback of orthogonalization via covariance. HTICA, using the more

sophisticated *orthogonalization via centroid body scaling* does not have this problem. We can see this in the right table of Figure 3, where the condition number of “centroid” is much smaller than the condition number of “covariance.”

5 Membership oracle for the centroid body

We will now describe and theoretically justify a new and practically efficient ϵ -weak membership oracle for ΓX , which is a black-box that can answer approximate membership queries in ΓX . The formal definition and argument are in the supplementary material. We give an informal description of the algorithm and its correctness.

The algorithm implementing the oracle (Subroutine 2) is the following: Let $q \in \mathbb{R}^n$ be a query point. Let X_1, \dots, X_N be a sample of random vector X . Given the sample, let Y be uniformly distributed in $\{X_1, \dots, X_N\}$. Output YES if $q \in \Gamma Y$, else output NO.

Idea of the correctness of the algorithm: If q is not in $(\Gamma X)_\epsilon$, then there is a hyperplane separating q from $(\Gamma X)_\epsilon$. Let $\{x : a^T x = b\}$ be the hyperplane, satisfying $\|a\| = 1$, $a^T q > b$ and $a^T x \leq b$ for every $x \in (\Gamma X)_\epsilon$. Thus, we have $h_{(\Gamma X)_\epsilon}(a) \leq b$ and $h_{\Gamma X}(a) \leq b - \epsilon$. We have

$$h_{\Gamma Y}(a) = \mathbb{E}(|a^T Y|) = (1/N) \sum_{i=1}^N |a^T X_i|.$$

By (Anderson et al. 2015, Lemma 14), $(1/N) \sum_{i=1}^N |a^T X_i|$ is within ϵ of $\mathbb{E}|a^T X| = h_{\Gamma X}(a) \leq b - \epsilon$ when N is large enough with probability at least $1 - \delta$ over the sample X_1, \dots, X_N . In particular, $h_{\Gamma Y}(a) \leq b$, which implies $q \notin \Gamma Y$ and the algorithm outputs NO, with probability at least $1 - \delta$.

If q is in $(\Gamma X)_{-\epsilon}$, let $y = q + \epsilon \hat{q} \in \Gamma X$.

Claim: For $p \in \Gamma X$, for large enough N and with probability at least $1 - \delta$ there is $z \in \Gamma Y$ so that $\|z - p\| \leq \epsilon/10$.

This claim applied to $p = y$ to get z , convexity of ΓY and the fact that ΓY contains $B \simeq \sigma_{\min}(A) B_n^n$ imply that $q \in \text{conv}(B \cup \{z\}) \subseteq \Gamma Y$ and the algorithm outputs YES.

We conclude with the main formal claims of the argument and a precise description of the oracle below:

Proposition 4 (Dual characterization of centroid body). *Let X be a n -dimensional random vector with finite first moment, that is, for all $u \in \mathbb{R}^n$ we have $\mathbb{E}(|\langle u, X \rangle|) < \infty$. Then $\Gamma X = \{\mathbb{E}(\lambda(X)X) : \lambda : \mathbb{R}^n \rightarrow [-1, 1] \text{ is measurable}\}$.*

Lemma 5 (LP). *Let X be a random vector uniformly distributed in $\{x^{(i)}\}_{i=1}^N \subseteq \mathbb{R}^n$. Let $q \in \mathbb{R}^n$. Then:*

1. $\Gamma X = \frac{1}{N} \sum_{i=1}^N [-x^{(i)}, x^{(i)}]$.
2. Point $q \in \Gamma X$ iff there is a solution $\lambda \in \mathbb{R}^N$ to the following linear feasibility problem:

$$\frac{1}{N} \sum_{i=1}^N \lambda_i x^{(i)} = q, -1 \leq \lambda_i \leq 1 \quad \forall i. \quad (1)$$

3. Let λ^* be the optimal value of (always feasible) linear program

$$\lambda^* = \max \lambda, \text{ s.t. } \frac{1}{N} \sum_{i=1}^N \lambda_i x^{(i)} = \lambda q, \lambda \in [-1, 1]^N \quad (2)$$

with $\lambda^* = \infty$ if the linear program is unbounded. Then the Minkowski functional of ΓX at q is $1/\lambda^*$.

Subroutine 2 Weak Membership Oracle for ΓX

Input: Query point $q \in \mathbb{R}^n$, samples from symmetric ICA model $X = AS$, bounds $s_M \geq \sigma_{\max}(A)$, $s_m \leq \sigma_{\min}(A)$, closeness parameter ϵ , failure probability δ .

Output: (ϵ, δ) -weak membership decision for $q \in \Gamma X$.

- 1: Let $N = \text{poly}(n, M, 1/s_m, s_M, 1/\epsilon, 1/\delta)$.
 - 2: Let $(x^{(i)})_{i=1}^N$ be an i.i.d. sample of X .
 - 3: Check the feasibility of linear program (1). If feasible, output YES, otherwise output NO.
-

Proposition 6 (Correctness of Subroutine 2). *Let $X = AS$ be given by an ICA model such that for all i we have $\mathbb{E}(|S_i|^{1+\gamma}) \leq M < \infty$, S_i is symmetrically distributed and normalized so that $\mathbb{E}|S_i| = 1$. Then, given a query point $q \in \mathbb{R}^n$, $\epsilon > 0$, $\delta > 0$, $s_M \geq \sigma_{\max}(A)$, and $s_m \leq \sigma_{\min}(A)$, Subroutine 2 is an ϵ -weak membership oracle for q and ΓX with probability $1 - \delta$ using time and sample complexity $\text{poly}(n, M, 1/s_m, s_M, 1/\epsilon, 1/\delta)$.*

References

- Amari, S.; Cichocki, A.; Yang, H. H.; et al. 1996. A new learning algorithm for blind signal separation. *Advances in neural information processing systems* 757–763.
- Anderson, J.; Goyal, N.; Nandi, A.; and Rademacher, L. 2015. Heavy-tailed independent component analysis. In *56th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2015)*. arXiv preprint arXiv:1509.00727.
- Bermond, O., and Cardoso, J.-F. 1999. Approximate likelihood for noisy mixtures. In *Proc. ICA*, volume 99, 325–330. Citeseer.
- Bickson, D., and Guestrin, C. 2010. Inference with multivariate heavy-tails in linear models. In *Advances in Neural Information Processing Systems*, 208–216.
- Cardoso, J.-F., and Soudoumiac, A. 1993. Blind beamforming for non-gaussian signals. In *Radar and Signal Processing, IEE Proceedings F*, volume 140, 362–370.
- Cardoso, J. 1989. Source separation using higher order moments. In *International Conference on Acoustics, Speech, and Signal Processing*.
- Chen, A., and Bickel, P. J. 2004. Robustness of prewhitening against heavy-tailed sources. In *Independent Component Analysis and Blind Signal Separation, Fifth International Conference, ICA 2004, Granada, Spain, September 22–24, 2004, Proceedings*, 225–232.
- Chen, A., and Bickel, P. J. 2005. Consistent independent component analysis and prewhitening. *Signal Processing, IEEE Transactions on* 53(10):3625–3632.
- Cichocki, A.; Amari, S.; Siwek, K.; Tanaka, T.; Phan, A. H.; and et al. Icalab toolboxes. <http://www.bsp.brain.riken.jp/ICALAB>.
- Clemencon, S., and Slim, S. 2007. On portfolio selection under extreme risk measure: The heavy-tailed ica model. *International Journal of Theoretical and Applied Finance* 10(03):449–474.
- Comon, P., and Jutten, C., eds. 2010. *Handbook of Blind Source Separation*. Academic Press.
- Donohue, K. D. 2009. <http://www.engr.uky.edu/~donohue/audio/Data/audioexpdata.htm>. Accessed: 2016-05-01.
- Gael, J. V.; Teh, Y. W.; and Ghahramani, Z. 2009. The infinite factorial hidden markov model. In *Advances in Neural Information Processing Systems*, 1697–1704.
- Hyvarinen, A.; Karhunen, J.; and Oja, E. 2001. *Independent Component Analysis*. John Wiley and Sons.
- Hyvarinen, A. 1999. Fast and robust fixed-point algorithms for independent component analysis. *Neural Networks, IEEE Transactions on* 10(3):626–634.
- J. Eriksson, A. K., and Koivunen, V. 2001. Novel characteristic function based criteria for ica. In *Proceedings ICA 2001*.
- Kidmose, P. 2000. Alpha-stable distributions in signal processing of audio signals. In *41st Conference on Simulation and Modelling, Scandinavian Simulation Society*, 87–94.
- Kidmose, P. 2001a. *Blind Separation of Heavy Tail Signals*. Ph.D. Dissertation, Technical University of Denmark.
- Kidmose, P. 2001b. Independent component analysis using the spectral measure for alpha-stable distributions. In *Proceedings of IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing, NSIP 2001*.
- Nolan, J. P. 2015. *Stable Distributions - Models for Heavy Tailed Data*. Boston: Birkhauser.
- Sahmoudi, M.; Abed-Meraim, K.; Lavielle, M.; Kuhn, E.; and Cibat, P. 2005. Blind source separation of noisy mixtures using a semi-parametric approach with application to heavy-tailed signals. In *Proc. of EUSIPCO 2005*.
- Shereshevski, Y.; Yeredor, A.; and Messer, H. 2001. Super-efficiency in blind signal separation of symmetric heavy-tailed sources. In *Statistical Signal Processing, 2001. Proceedings of the 11th IEEE Signal Processing Workshop on*, 78–81. IEEE.
- Wang, B.; Kuruoglu, E. E.; and Zhang, J. 2009. Ica by maximizing non-stability. In *Independent Component Analysis and Signal Separation*, 179–186. Springer.
- Welling, M.; Osindero, S.; and Hinton, G. E. 2002. Learning sparse topographic representations with products of student-t distributions. In *Advances in neural information processing systems*, 1359–1366.