

*Bernoulli* 0(00), 2013, 1–13  
DOI: 10.3150/12-BEJSP17

# On statistics, computation and scalability

MICHAEL I. JORDAN

*Department of Statistics and Department of EECS, University of California, Berkeley, CA, USA.  
E-mail: jordan@stat.berkeley.edu; url: www.cs.berkeley.edu/~jordan*

How should statistical procedures be designed so as to be scalable computationally to the massive datasets that are increasingly the norm? When coupled with the requirement that an answer to an inferential question be delivered within a certain time budget, this question has significant repercussions for the field of statistics. With the goal of identifying “time-data tradeoffs,” we investigate some of the statistical consequences of computational perspectives on scalability, in particular divide-and-conquer methodology and hierarchies of convex relaxations.

The fields of computer science and statistics have undergone mostly separate evolutions during their respective histories. This is changing, due in part to the phenomenon of “Big Data.” Indeed, science and technology are currently generating very large datasets and the gatherers of these data have increasingly ambitious inferential goals, trends which point towards a future in which statistics will be forced to deal with problems of scale in order to remain relevant. Currently the field seems little prepared to meet this challenge. To the key question “Can you guarantee a certain level of inferential accuracy within a certain time budget even as the data grow in size?” the field is generally silent. Many statistical procedures either have unknown runtimes or runtimes that render the procedure unusable on large-scale data. Although the field of sequential analysis provides tools to assess risk after a certain number of data points have arrived, this is different from an algorithmic analysis that predicts a relationship between time and risk. Faced with this situation, gatherers of large-scale data are often forced to turn to ad hoc procedures that perhaps do provide algorithmic guarantees but which may provide no statistical guarantees and which in fact may have poor or even disastrous statistical properties.

On the other hand, the field of computer science is also currently poorly equipped to provide solutions to the inferential problems associated with Big Data. Database researchers rarely view the data in a database as noisy measurements on an underlying population about which inferential statements are desired. Theoretical computer scientists are able to provide analyses of the resource requirements of algorithms (e.g., time and space), and are often able to provide comparative analyses of different algorithms for solving a given problem, but these problems rarely refer to inferential goals. In particular, the notion that it may be possible to save on computation because of the growth of statistical power as problem instances grow in size is not (yet) a common perspective in computer science.

In this paper we discuss some recent research initiatives that aim to draw computer science and statistics closer together, with particular reference to “Big Data” problems. There are two main underlying perspectives driving these initiatives, both of which present interesting conceptual challenges for statistics. The first is that large computational problems are often usefully addressed via some notion of “divide-and-conquer.” That is, the large problem is divided into subproblems that are hopefully simpler than the original problem, these subproblems are solved

(sometimes again with a divide-and-conquer strategy) and the solutions are pieced together to solve the original problem. In the statistical setting, one natural subdivision strategy involves breaking the data into subsets. The estimator of interest is applied to the subsets and the results are combined. The challenge in the statistical setting is that the analysis of subsets of data may present different statistical properties than the overall dataset. For example, confidence intervals based on subsets of data will generally be wider than confidence intervals based on the original data; thus, care must be taken that the overall divide-and-conquer procedure yields a correctly calibrated interval.

The second perspective involves a notion of “algorithmic weakening,” whereby we do not consider a single algorithm for solving an inference problem, but instead consider a hierarchy of algorithms that are ordered by computational complexity. As data accrue, we want to back off to cheaper algorithms that run more quickly and deliver a result that would be viewed as being of poorer quality from a classical algorithmic point of view. We hope to do this in a way such that the increasing statistical strength of the data compensate for the poor algorithmic quality, so that in fact the overall quality of inference increases as data accrue, even if we impose a computational budget. The challenge is to do this in a theoretically sound way.

The remainder of the paper is organized into three subsections, the first two concerned with divide-and-conquer algorithms, and the third concerned with algorithmic weakening.

## 1. Bag of little bootstraps

In this section we consider the core inferential problem of evaluating the quality of point estimators, a problem that is addressed by the bootstrap [Efron (1979)] and related resampling-based methods. The material in this section summarizes research described in Kleiner *et al.* (2012).

The usual implementation of the bootstrap involves the “computationally-intensive” procedure of resampling the original data with replacement, applying the estimator to each such bootstrap resample, and using the resulting distribution as an approximation to the true sampling distribution and thereby computing (say) a confidence interval. A notable virtue of this approach in the setting of modern distributed computing platforms is that it readily parallelizes – each bootstrap resample can be processed independently by the processors of a “cloud computer.” Thus in principle it should be possible to compute bootstrap confidence intervals in essentially the same runtime as is required to compute the point estimate. In the massive data setting, however, there is a serious problem: each bootstrap resample is itself massive (roughly 0.632 times the original dataset size). Processing such resampled datasets can overwhelm available computational resources; for example, with a terabyte of data it may not be straightforward to send a few hundred resampled datasets, each of size 632 gigabytes, to a set of distributed processors on a network.

An appealing alternative is to work with subsets of data, an instance of the divide-and-conquer paradigm. Existing examples of this alternative approach include subsampling [Politis, Romano and Wolf (1999)] and the  $m$ -out-of- $n$  bootstrap [Bickel, Götze and van Zwet (1997)]. In both cases, the idea is that a dataset of size  $n$  can be processed into multiple sets of size  $m$  (there are  $\binom{n}{m}$  such subsets in the case of sampling without replacement), and the estimator computed on each set. This yields fluctuations in the values of the point estimate. The challenge is the one referred to earlier – these fluctuations are on the wrong scale, being based on datasets that are

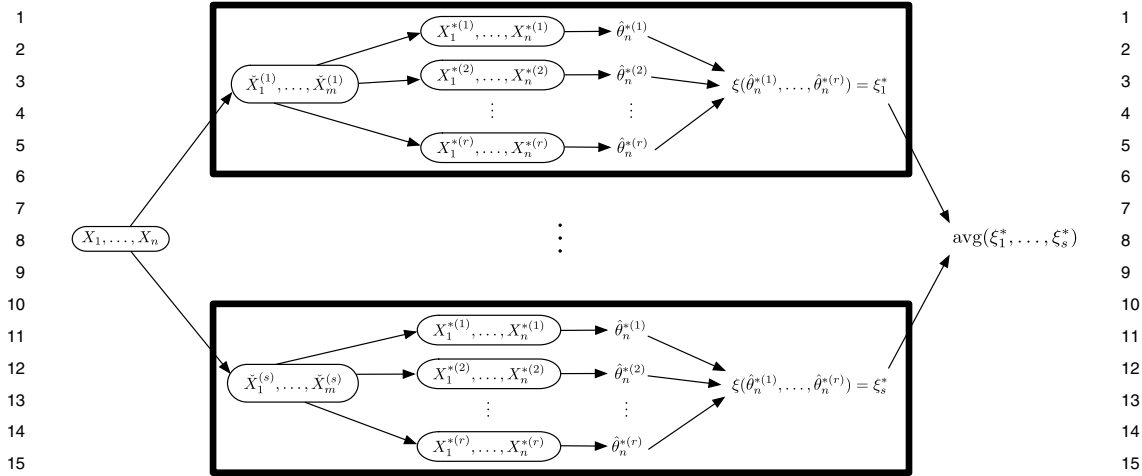
1 smaller than the original dataset. Both subsampling and the  $m$ -out-of- $n$  bootstrap assume that an  
2 analytical correction factor is available (e.g.,  $\sqrt{m/n}$ ) to rescale the confidence intervals obtained  
3 from the sampled subsets. This renders these procedures somewhat less “user-friendly” than the  
4 bootstrap, which requires no such correction factor.

5 There are situations in which the bootstrap is known to be inconsistent, and where subsam-  
6 pling and the  $m$ -out-of- $n$  bootstrap are consistent; indeed, the search for a broader range of  
7 consistency results was the original motivation for exploring these methods. On the other hand,  
8 finite sample results do not necessarily favor the consistent procedures over the bootstrap [see,  
9 e.g., Samworth (2003)]. The intuition is as follows. For small values of  $m$ , the procedure per-  
10 forms poorly, because each estimate is highly noisy. As  $m$  increases, the noise decreases and  
11 performance improves. For large values of  $m$ , however, there are too few subsamples, and per-  
12 formance again declines. In general it is difficult to find the appropriate value of  $m$  for a given  
13 problem.

14 In recent work, Kleiner *et al.* (2012) have explored a new procedure, the “Bag of Little Boot-  
15 straps” (BLB), which targets computational efficiency, but which also alleviates some of the  
16 difficulties of subsampling, the  $m$ -out-of- $n$  bootstrap and the bootstrap, essentially by combining  
17 aspects of these procedures. The basic idea of BLB is as follows. Consider a subsample of size  $m$   
18 (taken either with replacement or without replacement). Note that this subsample is itself a ran-  
19 dom sample from the population, and thus the empirical distribution formed from this subsample  
20 is an approximation to the population distribution. It is thus reasonable to sample from this em-  
21 pirical distribution as a plug-in proxy for the population. In particular, there is nothing preventing  
22 us from sampling  $n$  times from this empirical distribution (rather than  $m$  times). That is, we can  
23 implement the bootstrap on the correct scale using this subsample, simply by using it to generate  
24 multiple bootstrap samples of size  $n$ . Now, the resulting confidence interval will be a bona fide  
25 bootstrap confidence interval, but it will be noisy, because it is based on a (small) subsample.  
26 But we can proceed as in subsampling, repeating the procedure multiple times with randomly  
27 chosen subsamples. We obtain a set of bootstrap confidence intervals, which we combine (e.g.,  
28 by averaging) to yield the overall bootstrap confidence interval.

29 The procedure is summarized in Figure 1. We see that BLB is composed of two nested pro-  
30 cedures, with the inner procedure being the bootstrap applied to a subsample, and the outer  
31 procedure the being the combining of these multiple bootstrap estimates. From a computational  
32 point of view, the BLB procedure can be mapped onto a distributed computing architecture by  
33 letting each subsample be processed by a separate processor. Note that BLB has the virtue that  
34 the subsamples sent to each processors are small (of size  $m$ ). Moreover, although the inner loop  
35 of bootstrapping conceptually creates multiple resampled datasets of size  $n$ , it is not generally  
36 necessary to create actual datasets of size  $n$ ; instead we form weighted datasets of size  $m$ . (Also,  
37 the weights can be obtained as draws from a Poisson distribution rather than via explicit multi-  
38 nomial sampling.) Such is the case, for example, for estimators that are plug-in functionals of the  
39 empirical distribution.

40 An example taken from Kleiner *et al.* (2012) serves to illustrate the very substantial compu-  
41 tational gains that can be reaped from this approach. Consider computing bootstrap confidence  
42 intervals for the estimates of the individual components of the parameter vector in logistic re-  
43 gression, where the covariate vector has dimension 3000 and there are 6 000 000 data points, and  
44 where a distributed computing platform involving 80 cores (8 cores on each of ten processors)

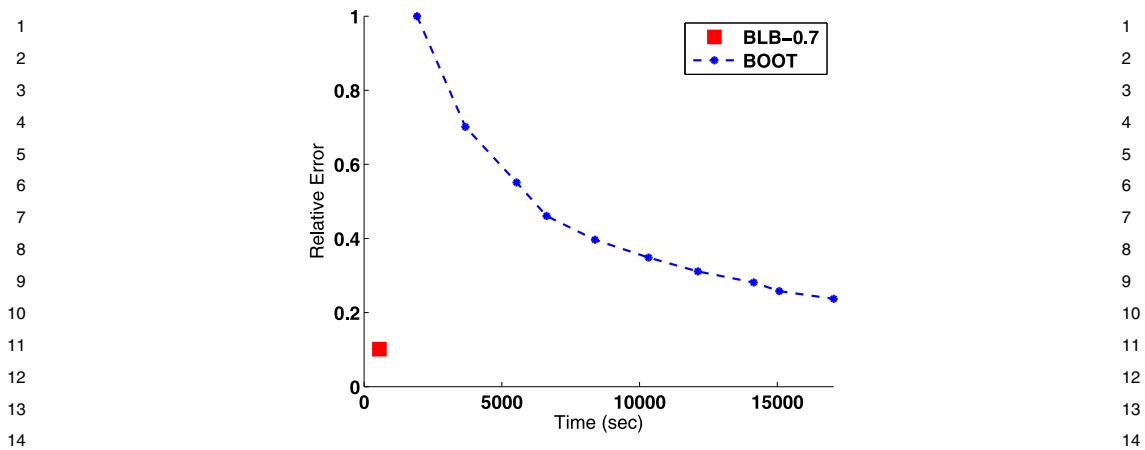


**Figure 1.** The BLB procedure. From the original dataset,  $\{X_1, \dots, X_n\}$ ,  $s$  subsamples of size  $m$  are formed. From each of these subsamples,  $r$  bootstrap resamples are formed, each of which are conceptually of size  $n$  (but would generally be stored as weighted samples of size  $m$ ). The resulting bootstrap estimates of risk are averaged. In a parallel implementation of BLB, the boxes in the diagram would correspond to separate processors; moreover, the bootstrap resampling within a box could also be parallelized.

is available. To implement the bootstrap at this scale, we can parallelize the logistic regression, and sequentially process the bootstrap resamples. Results from carrying out such a procedure are shown as the dashed curve in Figure 2, where we see that the processing of each bootstrap resample requires approximately 2000 seconds of processing time. The other natural approach is to implement a parallel version of BLB as we have discussed, where each processor executes the bootstrap on  $m(n) = n^\gamma$  points via weighted logistic regressions. The results are also shown in Figure 2, as a single dot in the lower-left corner of the figure. Here  $\gamma$  is equal to 0.7. We see that BLB has finished in less time than is required for a single iteration of the bootstrap on the full dataset, and indeed in less than 750 seconds has delivered an accuracy that is significantly better than that obtained by the bootstrap after 15 000 seconds.

Thus we see that there is a very strong synergy between a particular way to organize bootstrap-style computation and the capabilities of modern distributed computing platforms. Moreover, although the development of BLB was motivated by the computational imperative, it can be viewed as a novel statistical procedure to be compared to the bootstrap and subsampling according to more classical criteria; indeed, Kleiner *et al.* (2012) present experiments that show that even on a single processor BLB converges faster than the bootstrap and it is less sensitive to the choice of  $m$  than subsampling and the  $m$ -out-of- $n$  bootstrap.

There is much more to be done along these lines. For example, stratified sampling and other sophisticated sampling schemes can likely be mapped in useful ways to distributed platforms. For dependent data, one wants to resample in ways that respect the dependence, and this presumably favors certain kinds of data layout and algorithms over others. In general, for statistical inference



**Figure 2.** Univariate confidence intervals for a logistic regression on 6 000 000 data points using a 80-core distributed computing platform. The data were generated synthetically and thus the ground-truth sampling distribution was available via Monte Carlo; the y-axis is a measure of relative error with respect to this ground truth. See Kleiner *et al.* (2012) for further details.


to not run aground on massive datasets, we need for statistical thinking to embrace computational thinking.

## 2. Divide-and-conquer matrix factorization

Statistics has long exploited matrix analysis as a core computational tool, with linear models, contingency tables and multivariate analysis providing well-known examples. Matrices continue to be the focus of much recent computationally-focused research, notably as representations of graphs and networks and in collaborative filtering applications. At the basis of a significant number of analysis procedures is the notion of *matrix factorization*, with the singular value decomposition (SVD) providing a canonical example. The SVD in particular yields low-rank representations of matrices, which often maps directly to modeling assumptions.

Many matrix factorization procedures, including the SVD, have cubic algorithmic complexity in the row or column dimension of the matrix. This is overly burdensome in many applications, in statistics and beyond, and there is a strong motivation for applied linear algebra researchers to devise efficient ways to exploit parallel and distributed hardware in the computation of the SVD and other factorizations. One could take the point of view that this line of research is outside of the purview of statistics; that statisticians should simply keep an eye on developments. But this neglects the fact that as problems grow in size, it is the particular set of modeling assumptions at play that determine whether efficient algorithms are available, and in particular whether computationally-efficient approximations can be obtained that are statistically meaningful.

As a particularly salient example, in many statistical applications involving large-scale matrices it is often the case that many entries of a matrix are missing. Indeed, the quadratic growth in



**Figure 3.** The DFC pipeline. The matrix  $M$  is partitioned according to its columns and the resulting submatrices  $\{C_i\}$  are factored in parallel. The factored forms,  $\{\hat{C}_i\}$ , are then transmitted to a central location where they are combined into an overall factorization  $\hat{L}^{\text{proj}}$ .

the number of entries of a matrix is often accompanied by a linear growth in the rate of observation of matrix entries, such that at large scale the vast majority of matrix entries are missing. This is true in many large-scale collaborative filtering applications, where, for example, most individuals will have rated a small fraction of the overall set of items (e.g., books or movies). It is also true in many graph or network analysis problems, where each node is connected to a vanishingly small fraction of the other nodes.

In recent work, Mackey, Talwalkar and Jordan (2012) have studied a divide-and-conquer methodology for matrix factorization that aims to exploit parallel hardware platforms. Their framework, referred to as *Divide-Factor-Combine* (DFC), is rather simple from an algorithmic point of view – a matrix is partitioned according to its columns or rows, matrix factorizations are obtained (in parallel) for each of the submatrices using a “base algorithm” that is one of the standard matrix factorization methods, and the factorizations are combined to obtain an overall factorization (see Figure 3). The question is how to design such a pipeline so as to retain the statistical guarantees of the base algorithm, while providing computational speed-ups.

Let us take the example of *noisy matrix completion* [see, e.g., Candès and Plan (2010)], where we model a matrix  $M \in \mathbb{R}^{m \times n}$  as the sum of a low-rank matrix  $L_0$  (with rank  $r \ll \min(m, n)$ ) and a noise matrix  $Z$ :

$$M = L_0 + Z,$$

and where only a small subset of the entries of  $M$  are observed. Letting  $\Omega$  denote the set of indices of the observed entries, the goal is to estimate  $L_0$  given  $\{M_{ij} : (i, j) \in \Omega\}$ . This goal can be formulated in terms of an optimization problem:

$$\begin{aligned} \min_L \quad & \text{rank}(L) \\ \text{subject to} \quad & \sum_{(i,j) \in \Omega} (L_{ij} - M_{ij})^2 \leq \Delta^2, \end{aligned} \tag{2.1}$$

for a specified value  $\Delta$ . This problem is computationally intractable, so it is natural to consider replacing the rank function with its tightest convex relaxation, the nuclear norm, yielding the following convex optimization problem [Candès and Plan (2010)]:

$$\begin{aligned} \min_L \quad & \|L\|_* \\ \text{subject to} \quad & \sum_{(i,j) \in \Omega} (L_{ij} - M_{ij})^2 \leq \Delta^2, \end{aligned} \tag{2.2}$$

1 where  $\|L\|_*$  denotes the nuclear norm of  $L$  (the sum of the singular values of  $L$ ). 1

2 Candès and Plan (2010) have provided conditions under which the solution to equation (2.2) 2  
3 recovers the matrix  $L_0$  despite the potentially large number of unobserved entries of  $M$ . These 3  
4 conditions involve a structural assumption on the matrix  $L_0$  (that its singular vectors should not 4  
5 be too sparse or too correlated) and a sampling assumption for  $\Omega$  (that the entries of the matrix 5  
6 are sampled uniformly at random). We will refer to the former assumption as “ $(\mu, r)$ -coherence,” 6  
7 referring to Candès and Plan (2010) for technical details. Building on work by Recht (2011), 7  
8 Mackey, Talwalkar and Jordan (2012) have proved the following theorem, in the spirit of Candès 8  
9 and Plan (2010) but with weaker conditions: 9

10  
11 **Theorem 2.1.** Suppose that  $L_0$  is  $(\mu, r)$ -coherent and  $s$  entries of  $M$  are observed at locations 11  
12  $\Omega$  sampled uniformly without replacement, where 12

$$13 \quad s \geq 32\mu r(m+n)\log^2(m+n). \quad 14$$

15  
16 Then, if  $\sum_{(i,j) \in \Omega} (M_{ij} - L_{0,ij})^2 \leq \Delta^2$  a.s., the minimizer  $\hat{L}$  of equation (2.2) satisfies 16

$$17 \quad \|L_0 - \hat{L}\|_F \leq c_e \sqrt{mn} \Delta, \quad 18$$

19  
20 with high probability, where  $c_e$  is a universal constant. 20

21  
22 Note in particular that the required sampling rate  $s$  is a vanishing fraction of the total number 22  
23 of entries of  $M$ . 23

24 Theorem 2.1 exemplifies the kind of theoretical guarantee that one would like to retain under 24  
25 the DFC framework. Let us therefore consider a particular example of the DFC framework, 25  
26 referred to as “DFC-PROJ” by Mackey, Talwalkar and Jordan (2012), in which the “divide” 26  
27 step consists in the partitioning of the columns of  $M$  into  $t$  submatrices each having  $l$  columns 27  
28 (assuming for simplicity that  $l$  divides  $n$ ), the “factor” step involves solving the nuclear norm 28  
29 minimization problem in equation (2.2) for each submatrix (in parallel), and the “combine” step 29  
30 consists of a projection step in which the  $t$  low-rank approximations are projected onto a com- 30  
31 mon subspace.<sup>1</sup> Retaining a theoretical guarantee for DFC-PROJ from that of its base algorithm 31  
32 essentially involves ensuring that the  $(\mu, r)$ -coherence of the overall matrix is not increased very 32  
33 much in the random selection of submatrices in the “divide” step, and that the low-rank approxi- 33  
34 mations obtained in the “factor” step are not far from the low-rank approximation that would be 34  
35 obtained from the overall matrix. 35

36  
37 In particular, Mackey, Talwalkar and Jordan (2012) establish the following theorem:<sup>2</sup> 37

38  
39  
40 <sup>1</sup>In particular, letting  $\{\hat{C}_1, \hat{C}_2, \dots, \hat{C}_t\}$  denote the  $t$  low-rank approximations, we can project these submatrices onto the 40  
41 column space of any one of these submatrices, for example,  $\hat{C}_1$ . Mackey, Talwalkar and Jordan (2012) also propose an 41  
42 “ensemble” version of this procedure in which the low-rank submatrices are projected onto each other (i.e., onto  $\hat{C}_k$ , for 42  
43  $k = 1, \dots, t$ ) and the resulting projections are averaged. 42

43  
44 <sup>2</sup>We have simplified the statement of the theorem to streamline the presentation; see Mackey, Talwalkar and Jordan 43  
(2012) for the full result. 44



**Theorem 2.2.** Suppose that  $L_0$  is  $(\mu, r)$ -coherent and that  $s$  entries of  $M$  are observed at locations  $\Omega$  sampled uniformly without replacement. Then, if  $\sum_{(i,j) \in \Omega} (M_{ij} - L_{0,ij})^2 \leq \Delta^2$  a.s., and the base algorithm in the “factor” step of DFC-PROJ involves solving the optimization problem of equation (2.2), it suffices to choose

$$l \geq c\mu^2 r^2 (m+n)n \log^2(m+n) / (s\varepsilon^2)$$


columns in the “divide” step to achieve

$$\|L_0 - \hat{L}\|_F \leq (2 + \varepsilon)c_e \sqrt{mn} \Delta,$$

with high probability.

Thus, the DFC-PROJ algorithm achieves essentially the rate established in Theorem 2.1 for the nuclear norm minimization algorithm. Moreover, if we set  $s = \omega((m+n) \log^2(m+n))$ , which is slightly faster than the lower bound in Theorem 2.1, then we see that  $l/n \rightarrow 0$ . That is, DFC-PROJ succeeds even if only a vanishingly small fraction of the columns are sampled to form the submatrices in the “divide” step.

Figure 4 shows representative numerical results in an experiment on matrix completion reported by Mackey, Talwalkar and Jordan (2012). The leftmost figure depicts the accuracy (measured as root mean square error) achieved by the ensemble version of DFC-PROJ is nearly the same as that of the base algorithm. The rightmost figure shows that this accuracy is obtained at a fraction of the computational cost required by the baseline algorithm.



**Figure 4.** Numerical experiments on matrix completion with DFC. (a) Accuracy for a state-of-the-art baseline algorithm (“Base-MC”) and DFC-PROJ (see Mackey, Talwalkar and Jordan (2012) for details), as a function of the percentage of revealed entries. (b) Runtime as a function of matrix dimension (these are square matrices, so  $m = n$ ).



### 3. Convex relaxations

The methods that we have discussed thus far provide a certain degree of flexibility in the way inferential computations are mapped onto a computing infrastructure, and this flexibility implicitly defines a tradeoff between speed and accuracy. In the case of BLB the flexibility inheres in the choice of  $m$  (the subsample size) and in the case of DFC it is the choice of  $l$  (the submatrix dimension). In the work that we discuss in this section the goal is to treat such tradeoffs explicitly. To achieve this, Chandrasekaran and Jordan (2013) define a notion of “algorithmic weakening,” in which a hierarchy of algorithms is ordered by both computational efficiency and statistical efficiency. The problem that they address is to develop a quantitative relationship among three quantities: the number of data points, the runtime and the statistical risk.

Chandrasekaran and Jordan (2013) focus on the denoising problem, an important theoretical testbed in the study of high-dimensional inference [cf. Donoho and Johnstone (1998)]. The model is the following:

$$\mathbf{y} = \mathbf{x}^* + \sigma \mathbf{z}, \quad (3.1)$$

where  $\sigma > 0$ , the noise vector  $\mathbf{z} \in \mathbb{R}^p$  is standard normal, and the unknown parameter  $\mathbf{x}^*$  belongs to a known subset  $\mathcal{S} \subset \mathbb{R}^p$ . The problem is to estimate  $\mathbf{x}^*$  based on  $n$  independent observations  $\{\mathbf{y}_i\}_{i=1}^n$  of  $\mathbf{y}$ .

Consider a shrinkage estimator given by a projection of the sufficient statistic  $\bar{\mathbf{y}} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i$  onto a convex set  $\mathcal{C}$  that is an outer approximation to  $\mathcal{S}$ , that is,  $\mathcal{S} \subset \mathcal{C}$ :

$$\hat{\mathbf{x}}_n(\mathcal{C}) = \arg \min_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{2} \|\bar{\mathbf{y}} - \mathbf{x}\|_{\ell_2}^2 \quad \text{s.t.} \quad \mathbf{x} \in \mathcal{C}. \quad (3.2)$$


The procedure studied by Chandrasekaran and Jordan (2013) consists of a *set* of such projections,  $\{\hat{\mathbf{x}}_n(\mathcal{C}_i)\}$ , obtained from a hierarchy of convex outer approximations,

$$\mathcal{S} \subseteq \dots \subseteq \mathcal{C}_3 \subseteq \mathcal{C}_2 \subseteq \mathcal{C}_1.$$

The intuition is that for  $i < j$ , the estimator  $\{\hat{\mathbf{x}}_n(\mathcal{C}_i)\}$  will exhibit poorer statistical performance than  $\{\hat{\mathbf{x}}_n(\mathcal{C}_j)\}$ , given that  $\mathcal{C}_i$  is a looser approximation to  $\mathcal{S}$  than  $\mathcal{C}_j$ , but that  $\mathcal{C}_i$  can be chosen to be a simpler geometrical object than  $\mathcal{C}_j$ , such that it is computationally cheaper to optimize over  $\mathcal{C}_i$ , and thus more samples can be processed by the estimator  $\{\hat{\mathbf{x}}_n(\mathcal{C}_i)\}$  in a given time frame, offsetting the increase in statistical risk. Indeed, such *convex relaxations* have been widely used to give efficient approximation algorithms for intractable problems in computer science [Vazirani (2004)], and much is known about the decrease in runtime as one moves along the hierarchy of relaxations. To develop a time/data tradeoff, what is needed is a connection to statistical risk as one moves along the hierarchy.

Chandrasekaran and Jordan (2013) show that convex geometry provides such a connection. Define the *Gaussian squared-complexity* of a set  $\mathcal{D} \in \mathbb{R}^p$  as follows:

$$g(\mathcal{D}) = \mathbb{E} \left[ \sup_{\delta \in \mathcal{D}} \langle \delta, \mathbf{z} \rangle^2 \right],$$



**Figure 5.** (left) A signal set  $\mathcal{S}$  containing the true signal  $\mathbf{x}^*$ ; (middle) Two convex constraint sets  $\mathcal{C}$  and  $\mathcal{C}'$ , where  $\mathcal{C}$  is the convex hull of  $\mathcal{S}$  and  $\mathcal{C}'$  is a relaxation that is more efficiently computable than  $\mathcal{C}$ ; (right) The tangent cone  $T_{\mathcal{C}}(\mathbf{x}^*)$  is contained inside the tangent cone  $T_{\mathcal{C}'}(\mathbf{x}^*)$ . Consequently, the Gaussian squared-complexity  $g(T_{\mathcal{C}}(\mathbf{x}^*) \cap B_{\ell_2}^p)$  is smaller than the complexity  $g(T_{\mathcal{C}'}(\mathbf{x}^*) \cap B_{\ell_2}^p)$ , so that the estimator  $\hat{\mathbf{x}}_n(\mathcal{C})$  requires fewer samples than the estimator  $\hat{\mathbf{x}}_n(\mathcal{C}')$  for a risk of at most 1.

where the expectation is with respect to  $\mathbf{z} \sim \mathcal{N}(0, I_{p \times p})$ . Given a closed convex set  $\mathcal{C} \in \mathbb{R}^p$  and a point  $\mathbf{a} \in \mathcal{C}$ , define the *tangent cone* at  $\mathbf{a}$  with respect to  $\mathcal{C}$  as

$$T_{\mathcal{C}}(\mathbf{a}) = \text{cone}\{\mathbf{b} - \mathbf{a} \mid \mathbf{b} \in \mathcal{C}\}, \quad (3.3)$$

where  $\text{cone}(\cdot)$  refers to the conic hull of a set obtained by taking nonnegative linear combinations of elements of the set. (See Figure 5 for a depiction of the geometry.) Let  $B_{\ell_2}^p$  denote the  $\ell_2$  ball in  $\mathbb{R}^p$ . Chandrasekaran and Jordan (2013) establish the following theorem linking the Gaussian squared-complexity of tangent cones and the statistical risk:

**Theorem 3.1.** For  $\mathbf{x}^* \in \mathcal{S} \subset \mathbb{R}^p$  and with  $\mathcal{C} \subseteq \mathbb{R}^p$  convex such that  $\mathcal{S} \subseteq \mathcal{C}$ , we have the error bound

$$\mathbb{E}[\|\mathbf{x}^* - \hat{\mathbf{x}}_n(\mathcal{C})\|_{\ell_2}^2] \leq \frac{\sigma^2}{n} g(T_{\mathcal{C}}(\mathbf{x}^*) \cap B_{\ell_2}^p).$$

This risk bound can be rearranged to yield a way to estimate the number of data points needed to achieve a given level of risk. In particular, the theorem implies that if

$$n \geq \sigma^2 g(T_{\mathcal{C}}(\mathbf{x}^*) \cap B_{\ell_2}^p), \quad (3.4)$$

then  $\mathbb{E}[\|\mathbf{x}^* - \hat{\mathbf{x}}_n(\mathcal{C})\|_{\ell_2}^2] \leq 1$ . The overall implication is that as the number of data points  $n$  grows, we can back off to computationally cheaper estimators and still control the statistical risk, simply by choosing the largest  $\mathcal{C}_i$  such that the right-hand side of equation (3.4) is less than  $n$ . This yields a time/data tradeoff.

To exemplify the kinds of concrete tradeoffs that can be obtained via this formalism, Chandrasekaran and Jordan (2013) consider a stylized sparse principal component analysis problem, modeled using the following signal set:

$$\mathcal{S} = \{\Pi M \Pi' \mid \Pi \text{ is a } \sqrt{p} \times \sqrt{p} \text{ permutation matrix}\},$$

**Table 1.** Time-data tradeoffs for the sparse PCA problem, expressed as a function of the matrix dimension  $p$ . See Chandrasekaran and Jordan (2013) for details

$\mathcal{C}$	Runtime	$n$
conv( $\mathcal{S}$ )	super-poly( $p$ )	$\sim p^{1/4} \log(p)$
Nuclear norm ball	$p^{3/2}$	$\sim p^{1/2}$

where the top-left  $k \times k$  block of  $M \in \mathbb{R}^{\sqrt{p} \times \sqrt{p}}$  has entries equal to  $\sqrt{p}/k$  and all other entries are zero. In Table 1 we show the runtimes and sample sizes associated with two different convex relaxations of  $\mathcal{S}$ : the convex hull of  $\mathcal{S}$  and the nuclear norm ball. The table reveals a time-data tradeoff – to achieve constant risk we can either use a more expensive procedure that requires few data points or a cheaper procedure that requires more data points. As a second example, consider the cut-matrix denoising problem, where the signal set is as follows:

$$\mathcal{S} = \{\mathbf{a}\mathbf{a}' \mid \mathbf{a} \in \{-1, +1\}^{\sqrt{p}}\}.$$

Table 2 displays the runtimes and sample sizes associated with three different convex relaxations of this signal set. Here the tradeoff is in the constants associated with the sample size, favoring the cheaper methods. Chandrasekaran and Jordan (2013) also consider other examples, involving variable ordering and banded covariance matrices. In all of these examples, it seems to be the case that the cheaper methods achieve the same risk as more expensive methods with not very many additional data points.

## 4. Discussion

We have reviewed several lines of research that aim to bring computational considerations into contact with statistical considerations, with a particular focus on the matrix-oriented estimation problems that arise frequently in the setting of massive data. Let us also mention several other recent theoretical contributions to the statistics/computation interface. Divide-and-conquer

**Table 2.** Time-data tradeoffs for the cut-matrix denoising problem, expressed as a function of the matrix dimension  $p$ , where  $c_1 < c_2 < c_3$ . See Chandrasekaran and Jordan (2013) for details

$\mathcal{C}$	Runtime	$n$
Cut polytope	super-poly( $p$ )	$c_1 p^{1/2}$
Elliptope	$p^{7/4}$	$c_2 p^{1/2}$
Nuclear norm ball	$p^{3/2}$	$c_3 p^{1/2}$

1 methodology has been explored by Chen and Xie (2012) in the setting of regression and clas- 1  
2 sification. Their methods involve estimating parameters on subsets of data in parallel and using 2  
3 weighted combination rules to merge these estimates into an overall estimates. They are able to 3  
4 show asymptotic equivalence to an estimator based on all of the data and also show (empirically) 4  
5 a significant speed-up via the divide-and-conquer method. The general idea of algorithmic weak- 5  
6 ening via hierarchies of model families has been explored by several authors; see, for example, 6  
7 Agarwal *et al.* (2011) and Shalev-Shwartz, Shamir and Tromer (2012), where the focus is model 7  
8 selection and classification, and Amini and Wainwright (2009), where the focus is sparse covari- 8  
9 ance matrix estimation. In all of these lines of work the goal is to develop theoretical tools that 9  
10 explicitly reveal tradeoffs relating risk, data and time. 10

11 It is important to acknowledge the practical reality that massive datasets are often complex, 11  
12 heterogeneous and noisy, and the goal of research on scalability is not that of developing a single 12  
13 methodology that applies to the analysis of such datasets. Indeed, massive datasets will require 13  
14 the full range of statistical methodology to be brought to bear in order for assertions of knowledge 14  
15 on the basis of massive data analysis to be believable. The problem is that of recognizing that 15  
16 the analysts of massive data will often be interested not only in statistical risk, but in risk/time 16  
17 tradeoffs, and that the discovery and management of such tradeoffs can only be achieved if the 17  
18 algorithmic character of statistical methodology is fully acknowledged at the level of the founda- 18  
19 tional principles of the field. 19  
20

## 21 Acknowledgements 22

23 I wish to acknowledge numerous colleagues who have helped to shape the perspective pre- 24  
25 sented here, in particular Venkat Chandrasekaran, Ariel Kleiner, Lester Mackey, Purna Sarkar 25  
26 and Ameet Talwalkar. 26  
27

## 28 References 29

- 30  
31 Agarwal, A., Duchi, J., Bartlett, P. and Levrard, C. (2011). Oracle inequalities for computationally budgeted 31  
32 model selection. In *24th Annual Conference on Learning Theory, Budapest, Hungary.* 32 <unstr>  
33 Amini, A.A. and Wainwright, M.J. (2009). High-dimensional analysis of semidefinite relaxations for sparse 33  
34 principal components. *Ann. Statist.* **37** 2877–2921. [MR2541450](#) 34 <mr>  
35 Bickel, P.J., Götze, F. and van Zwet, W.R. (1997). Resampling fewer than  $n$  observations: Gains, losses, 35  
36 and remedies for losses. *Statist. Sinica* **7** 1–31. Empirical Bayes, sequential analysis and related topics 36  
37 in statistics and probability (New Brunswick, NJ, 1995). [MR1441142](#) 37 <mr>  
38 Candès, E.J. and Plan, Y. (2010). Matrix completion with noise. *Proceedings of the IEEE* **98** 25–936. 38 <auto>  
39 Chandrasekaran, V. and Jordan, M.I. (2013). Computational and statistical tradeoffs via convex relaxation. 39  
40 *Proc. Natl. Acad. Sci. USA* **13** E1181–E1190. 40 <auto>  
41 Chen, X. and Xie, M. (2012). A split-and-conquer approach for analysis of extraordinarily large data. Tech- 41  
42 nical Report 2012-01, Dept. Statistics, Rutgers Univ. 42 <unstr>  
43 Donoho, D.L. and Johnstone, I.M. (1998). Minimax estimation via wavelet shrinkage. *Ann. Statist.* **26** 879– 43  
44 921. [MR1635414](#) 44 <mr>  
45 Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *Ann. Statist.* **7** 1–26. [MR0515681](#)

1	Kleiner, A., Talwalkar, A., Sarkar, P. and Jordan, M.I. (2012). A scalable bootstrap for massive data. Available at <a href="https://arxiv.org/abs/1112.5016">arXiv:1112.5016</a> .	1
2		2 <unstr>
3	Mackey, L., Talwalkar, A. and Jordan, M.I. (2012). Divide-and-conquer matrix factorization. Available at <a href="https://arxiv.org/abs/1107.0789">arXiv:1107.0789</a> .	3
4		4 <unstr>
5	Politis, D.N., Romano, J.P. and Wolf, M. (1999). <i>Subsampling</i> . Springer Series in Statistics. New York: Springer. <a href="https://doi.org/10.1007/978-1-4939-9730-1">MR1707286</a>	5
6		6 <mr>
7	Recht, B. (2011). A simpler approach to matrix completion. <i>J. Mach. Learn. Res.</i> <b>12</b> 3413–3430. <a href="https://doi.org/10.1214/11-AOS1173">MR2877360</a>	7
8		8 <mr>
9	Samworth, R. (2003). A note on methods of restoring consistency to the bootstrap. <i>Biometrika</i> <b>90</b> 985–990. <a href="https://doi.org/10.1093/biomet/90.4.985">MR2024773</a>	9
10		10 <mr>
11	Shalev-Shwartz, S., Shamir, O. and Tromer, E. (2012). Using more data to speed up training time. In <i>Conference on Artificial Intelligence and Statistics, La Palma, Canary Islands</i> .	11 <auto>
12	Vazirani, V. (2004). <i>Approximation Algorithms</i> . New York: Springer.	12 <mr>
13		13
14		14
15		15
16		16
17		17
18		18
19		19
20		20
21		21
22		22
23		23
24		24
25		25
26		26
27		27
28		28
29		29
30		30
31		31
32		32
33		33
34		34
35		35
36		36
37		37
38		38
39		39
40		40
41		41
42		42
43		43
44		44

## The original reference list

The list of entries below corresponds to the original Reference section of your article. The bibliography section on previous page was retrieved from MathSciNet applying an automated procedure. Please check both lists and indicate those entries which lead to mistaken sources in automatically generated Reference list.

- Agarwal, A., Duchi, J., Bartlett, P., and Levrard, C. (2011) Oracle inequalities for computationally budgeted model selection. *24th Annual Conference on Learning Theory*, Budapest, Hungary.
- Amini, A. and Wainwright, M. (2009) High-dimensional analysis of semidefinite programming relaxations for sparse principal component analysis. *Annals of Statistics*, 37, 2877–2921.
- Bickel, P. J. Gotze, F., and Van Zwet, W. (1997) Resampling fewer than  $n$  observations: Gains, losses, and remedies for losses. *Statistica Sinica*, 7, 1–31, 1997.
- Candès, E. J. and Plan, Y. (2010) Matrix completion with noise. *Proceedings of the IEEE*, 98, 25–936.
- Chandrasekaran, V. and Jordan, M. I. (2013) Computational and statistical tradeoffs via convex relaxation. *Proceedings of the National Academy of Sciences*, 13, E1181–E1190.
- Chen, X. and Xie, M. (2012) A split-and-conquer approach for analysis of extraordinarily large data. Technical Report 2012-01, Department of Statistics, Rutgers University.
- Donoho, D. L. and Johnstone, I. M. (1998) Minimax estimation via wavelet shrinkage. *Annals of Statistics*, 26, 879–921.
- Efron, B. (1979) Bootstrap methods: Another look at the jackknife. *Annals of Statistics*, 7, 1–26.
- Kleiner, A., Talwalkar, A., Sarkar, P. and Jordan, M. I. (2012) A scalable bootstrap for massive data. *arXiv:1112.5016*.
- Mackey, L., Talwalkar, A. and Jordan, M. I. (2012) Divide-and-conquer matrix factorization. *arXiv:1107.0789*.
- Politis, D., Romano, J., and Wolf, M. (1999) *Subsampling*. New York: Springer.
- Recht, B. (2011). A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12, 3413–3430.
- Samworth, R. (2003) A note on methods of restoring consistency to the bootstrap, *Biometrika*, 90, 985–990.
- Shalev-Shwartz, S., Shamir, O., and Tromer, E. (2012) Using more data to speed up training time. In *Conference on Artificial Intelligence and Statistics*, La Palma, Canary Islands.
- Vazirani, V. (2004) *Approximation Algorithms*. New York: Springer.

## 1 Meta data in the pdf file

### 3 Following information will be included as pdf file Document Properties:

4 **Title** : On statistics, computation and scalability

5 **Author** : Michael I. Jordan

6 **Subject** : Bernoulli, Vol. 0, No. 00, 2013, 1-13

7 **Keywords**:

## 9 The list of URI addresses

12 Listed below are all uri addresses found in your paper. The non-active uri addresses, if any, are indicated as  
13 ERROR. Please check and update the list where necessary. The e-mail addresses are not checked – they are listed  
14 just for your information. More information can be found in the support page:

15 <http://www.e-publications.org/ims/support/urihelp.html>.

16 200 [http://www.bernoulli-society.org/index.php/publications/  
17 bernoulli-journal/bernoulli-journal](http://www.bernoulli-society.org/index.php/publications/bernoulli-journal/bernoulli-journal) [2:pp.1,1] OK  
18 404 <http://dx.doi.org/10.3150/12-BEJSP17> [2:pp.1,1] Not Found  
19 --- <mailto:jordan@stat.berkeley.edu> [2:pp.1,1] Check skip  
20 301 <http://www.cs.berkeley.edu/~jordan> [2:pp.1,1] Moved Permanently  
21 200 <http://arxiv.org/abs/arXiv:1112.5016> [2:pp.13,13] OK  
22 200 <http://arxiv.org/abs/arXiv:1107.0789> [2:pp.13,13] OK