

Statistical Inference, Learning and Models in Big Data

Franke, B., Plante, J.-F., Roscher, R., Lee, A., Smyth, C., Hatefi, A., Chen, F., Gil, E., Schwing, A., Selvitella, A., Hoffman, M. M., Grosse, R., Hendricks, D. and Reid, N.

1. Introduction

Big data provides big opportunities for statistical inference, but perhaps even bigger challenges, often related to differences in volume, variety, velocity, and veracity of information when compared to smaller carefully collected datasets. From January to June, 2015, the Canadian Institute of Statistical Sciences organized a thematic program on Statistical Inference, Learning and Models in Big Data. This paper arose from presentations and discussions that took place during the thematic program.

Scientifically, the program emphasized the roles of statistics, computer science, and mathematics, in obtaining scientific insight from “Big Data”. Two complementary strands were introduced: cross-cutting, or foundational, research that underpins analysis of big data, and domain-specific research focused on particular application areas. The former category included machine learning, statistical inference, optimization, network analysis, and visualization. Topic-specific workshops addressed problems in health policy, social policy, environmental science, cyber-security and social networks. These divisions are not rigid, of course, as foundational and application areas live in a synergy where they feed the growth of each other. Several very important application areas where big data is fundamental were not able to be the subject of focussed workshops, but many of these applications did appear in individual presentations. The program started with an opening conference¹ that gave an overview of the topics of the 6-month program². The talks presented at the Fields Institute are available online through FieldsLive³. While there is much important research underway in improving the processing, storing and the rapid accessing of records (Chen, 2014; Schadt, 2010), this program focused on the challenges for inference and analysis.

Two striking features of the presentations during the program were the breadth of topics, and the remarkable commonalities that emerged across this broad range. We attempt here to illustrate many of the common issues and solutions that arose, and provide a picture of the status of Big Data research in the statistical sciences. While the report is largely inspired by

¹ The program for the opening conference is at <http://www.fields.utoronto.ca/programs/scientific/14-15/bigdata/boot/>

² A description of the complete thematic program is at <https://www.fields.utoronto.ca/programs/scientific/14-15/bigdata/>

³ The Fields Institute video archive link is <http://www.fields.utoronto.ca/video-archive>

the series of talks at the opening conference and related references, it also touches on issues that arose throughout the program.

In the next section we identify a number of challenges that are commonly identified in the world of Big Data, and in Section 3 we describe some examples of general inference strategies that are being developed to address these challenges. Section 4 highlights several applications, to make some of these ideas more concrete.

2. “Big Data, it’s not the data”

Many talks and papers on big data begin by trying to define or describe the somewhat over-hyped label “Big Data”. The title of this section is taken from the opening presentation by Robert Bell, who paraphrased Supreme Court Justice Porter Stewart by saying “I shall not today further attempt to define big data, but I know it when I see it”. Bell also pointed out that the standards of big data are evolving, as computational resources increase. However, some common features emerge that distinguish the big data world, and it is popular to describe these briefly by the “Four Vs”: volume, velocity, variety and veracity.

The most obvious challenge of working with big data is the volume: it is, simply put, big. This presents computational challenges, foremost; in managing the data, in finding algorithms that scale well, and in developing specialized computational techniques. The size also presents problems for inference: approaches that work for smaller datasets often simply will not scale, and alternative approaches need to be developed. A multidisciplinary approach is essential, because many different areas of expertise are needed together. As a broad caricature with some basis in reality, statistical scientists are well-trained in inference, and in having a well-developed intuition for bias and variability, and computational scientists are well-trained in fast implementation, in exploration of new ideas, and in managing data. Both sets of skills are essential. It is the case, and was pointed out by several speakers, that the costs of research are in many fields moving from data acquisition to data analysis.

In many areas of application, data is “big” because it is arriving at a high velocity, from continuously operating instrumentation, e.g., when considering autonomous driving scenarios (Geiger et al. 2012), or through an on-line streaming service. Sequential or on-line methods of analysis may be needed, which requires methods of inference that can be computed quickly, that repeatedly adapt to sequentially added data, and that are not misled by the sudden appearance of irrelevant data, as discussed, e.g., by Jain et al. (2006) and Hoens et al. (2012). The areas of applications using sequential learning are diverse. For instance, Li and Fei-Fei (2010) use object recognition techniques in an efficient sequential way to learn object category models from any arbitrarily large number of images from the web. Monteleoni et al. (2011) use sequential learning for climate change modelling to estimate transition

probabilities between different climate models. For some applications, such as high energy physics, bursts of data arrive so rapidly that it is infeasible to even record all of it. In this case a great deal of scientific expertise is needed to ensure that the data retained is relevant to the problem of interest.

In addition to, or perhaps because of, its large volume big data tends to be complex. Or more precisely, the models used to try to make sense of big data tend to be complex. The main manifestation of this complexity is high dimensionality; for a great many types of big data the number of variables p is larger, often much larger, than the number of observations n . This takes us far from the classical regime with a fixed number of variables p and a larger, often much larger, number of observations n . This change in the relationship between n and p can lead to counter-intuitive behaviours, or a complete breakdown, of conventional methods of analysis. Problems in which p scales with n are very common: for example, in models for networks that assign one parameter to each node, there are $p = n$ parameters for a graph with n^2 possible edges (Chung and Lu, 2002).

Bioucas-Dias et al. (2013) discuss hyperspectral remote sensing image data, which may have hundreds or thousands of spectral bands, p , but a very small number of training samples, n . In a recent genomics study reported by the 1000 Genomes Project Consortium (2012), geneticists acquired data on $p = 38$ million single nucleotide polymorphisms in only $n = 1092$ individuals.

The asymptotic properties of classical methods in a regime where p increases with n are quite different than in the classical regime of fixed p . For example, if $X_i, i = 1, \dots, n$ are independent and identically distributed as $N(0, I_p)$, and p/n stays constant as p and n increase, the eigenvalues of the empirical covariance matrix $S = XX'/n$ do not converge to the true value of 1, but rather have a limiting distribution described by the Marcenko-Pastur law. Donoho and Gavish (2014) show how this law is relevant for understanding the asymptotic minimax risk in estimating a low rank matrix from noisy measurements by thresholding the singular values. Donoho et al. (2013) obtain similar results for the problem of matrix recovery.

Big data often entails very complex dependence structures. Temporal and spatio-temporal processes may involve a very large number of parameters, for example. But we also encounter entirely new notions of dependence, such as the relationship between elements modelled as networks. This differs from temporal or spatial dependence in not having a natural geometry with which to measure 'neighbourhoods', and related distances (Kolaczyk, 2009). Diaconis (2007) and Orbanz and Roy (2013) develop an analogy to time-series-type stationarity by considering exchangeable graphs whose edge distribution is invariant to label permutation.

The complexity of many big data problems is closely related to the increasing variety of data types that are available. Much of classical statistical modelling assumes a rectangular structure for the data, with rows indexing observational or experimental units, and columns representing variables. “Short and fat” data arrays have $p \gg n$. Classical arrays are “tall and skinny”, by contrast. Images, both two- and three-dimensional, blog posts, twitter feeds, social networks, audio recordings, sensor networks, are just some examples of new types of data that don’t fit this “rectangular” model. Often rectangular data is created somewhat artificially: for example for two-dimensional images, by ordering the pixels row by row and recording one or more values associated with each pixel.

Another aspect of variety is that big data is often pieced together from several datasets that have not been collected with the current research question in mind: this is sometimes referred to as “found data” or “convenience samples”. This type of data will typically have a heterogeneous structure (Hodas and Lerman 2013), and may well not be representative of the population of interest. Robertson et al. (2012) show how the study of diseases and treatment efficacy can be improved by linking health survey data (Admitted Patient Data Collection) to administrative data (Registry of Births, Deaths and Marriages) - an example showing how separate datasets are being combined to answer a new research question. Linking of databases raises substantial and difficult issues of privacy: in order to link databases accurately, a unique identifier common to the databases is required. If this identifier can be in turn linked to publicly available data, such as media articles, then privacy can be completely compromised. A high-profile example of this was the publication of the health records of the Governor of Massachusetts, which in turn led to strict laws on the collection and release of health data in the United States known as the “HIPAA rules” (Tarran, 2014).

The veracity of big data is a pressing issue which has been widely discussed, and Gaffield (2015) warned that we should be aware of the potential for backlash when, inevitably, one or more high profile analyses of a big database turn out to be incorrect. The vast majority of big data is observational, but more importantly may be collected very haphazardly. Classical ideas of experimental design, survey sampling, and design of observational studies often get lost in the excitement of the availability of large amounts of data. Issues of sampling bias, if ignored, can lead to incorrect inference just as easily with big data as with small. Harford (2014) presents a number of high-profile examples, from the *Literary Digest* Poll of 1936 to Google Flu Trends in 2010. Networks of hard-to-reach respondents are sometimes established using the method of “snowball sampling”, where index subjects are used to recruit additional study subjects from their contacts – this has the potential to lead to a large sampling bias (Kolaczyk, 2009, p. 133).

Even if we could overcome the challenges in extracting information from big data, conveying the extracted information to humans without information overload remains a huge challenge of its own. Visualization methods can provide powerful tools for communication, but it has limited “bandwidth”. One consequence of this is that we often must compromise between the intuitiveness of a visualization and the power of a visualization (Carpendale et al., 2004; Nacenta et al., 2012). Another consequence is that we typically visualize only a small subset of the data, certainly at one time, and this may give a biased view (Kolaczyk, 2009; p. 91).

Big data may not fit into memory or even the available storage on a desktop computer. Tackling big data problems often requires a computer cluster or other distributed compute systems. Two systems widely promoted are Hadoop and MapReduce; these introduce analytic and computational complexity, and aspects of statistical inference in these settings is not very well developed. As the costs of processors, memory and disk space continues to decrease, the limiting piece now is often communication between machines (Scott et al. 2013). Methods to reduce this communication as much as possible during the analysis are being developed; one example is the consensus Monte Carlo algorithm of Scott et al. (2013, 2015).

3. Strategies for big data analysis

In spite of the wide variety of types of big data problems, there is a common set of strategies that seem essential. Aside from pre-processing and multi-disciplinarity, many of these follow the general approach of identifying and exploiting a hidden structure within the data that often is of lower dimension. We describe some of these commonalities in this section.

3.1 Data Wrangling

A very large hurdle in the analysis of data, which is especially difficult for massive data sets, is manipulating the data for use in analysis. This is variously referred to as data wrangling, data carpentry, or data cleaning⁴. Big data is often very “raw”: considerable pre-processing such as extracting, parsing, and storing, is required before it can be considered in an analysis.

Examples of data wrangling include selecting variables or features from image and sound data (e.g. Guyon, 2003), integrating data from heterogeneous sources (Izadi et al., 2013), as well as parsing free-form text for text mining or sentiment analysis (e.g., Gupta, 2009; Pang, 2008). A more advanced approach in text mining uses the semantics of text rather than the simple count of words to describe texts (e.g., Brien et al. 2013).

One example of a data processing problem arises in the merging of several sets of spatio-temporal data, which may have different spatial units – for example one database may be

⁴ The R packages `dplyr` and `tidyr` were expressly created to simplify this step, and are highly recommended. The new “pipe” operator `%>%`, operating similarly to the old Unix pipe command, is especially helpful. (Wickham, 2014).

based on census tracts, another on postal codes, and still another on electoral districts. Brown et al. (2013) have developed an approach to heterogeneous complex data based on a local EM algorithm. This starts by overlaying all partitions that are used for aggregation, and then creating a finer partition such that all elements of that partitions are fully included in the larger ones. The local EM algorithm based on this is parallelizable, and hence provides a workable solution to this complex problem. A related issue arose in the presentation of Gaffield (2015), where historical census data is being used as part of a study of the impact of individuals on the history of Canada. The census tracts have changed over time, and the current database resolves this by mapping backwards from the current census tracts.

3.2 Visualization

After data wrangling, visualization is very often the first formal step of the analysis: the aim is to find efficient graphical representations that summarize the data and emphasise its main characteristics (Tukey, 1977). Visualization can also serve as an inferential tool in different stages of the analysis -- understanding the patterns in the data helps in creating good models (Buja et al, 2009). For example, Albert-Green et. al (2014), Woolford and Braun (2007, 2013) and Woolford et al. (2009) present exploratory data analysis tools for large environmental data sets, which are meant to not only plot the data, but also to look for structures and departures from such structures. The authors focus on convergent data sharpening and the stalagmite plot in the context of forest fire science.

Visualization is “bandwidth-limited”, in the sense that there is not enough screen space to visualize large, or even moderately sized, data. Interactivity provides many ways to help with this; the most familiar is zooming into large visualizations to reveal local structure, much in the manner of Google Maps. New techniques are being developed in the information visualization community: Neumann, Carpendale & Agarawala (2006) described taking cues from nature, and from fractals, in developing *PhylloTrees* for illustrating network graphs at different scales. Carpendale in her bootcamp presentation described several recent innovations in interactive visualization, many of which are illustrated at the website of the Innovis Research Group (Carpendale, 2015).

3.3 Reducing Dimensionality

Direct reduction of dimensionality is a very common technique: in addition to facilitating visualization (Huron et al., 2014; Wickham et al., 2010), it is often needed to address issues of data storage (Witten and Candes, 2013). Hinton and Salakhutdinov (2006) describe the use of direct dimension reduction in classification. Dimensionality reduction also often incorporates an assumption of sparsity, which is addressed in the next subsection.

Principal Component Analysis (PCA) is a classical example of dimensionality reduction, introduced by Pearson (1901). The first two, or three, principal components are often plotted in the hopes of identifying structure in the data. One of the issues with performing PCA on

large data sets is the intractability of its computation, which relies on singular value decomposition (SVD). Witten and Candes (2013), Eldar and Kutyniok (2012), and Halko, (2011), among others, have used random matrix projections. This is a powerful tool behind compressed sensing and a striking example of dimensionality reduction, allowing approximate QR decomposition or singular value decomposition with great performance advantages over conventional methods. Witten and Candes (2013) approximate a very large matrix by the product of two “skinny” matrices that can be processed relatively easily. Using randomized methods one can infer the principal components of the original data from the singular value decomposition of these skinny factor matrices, thus overcoming a computational bottleneck.

Depending on the application, linear dimensionality reduction techniques such as PCA may need to be extended to non-linear ones, e.g., if the structure of the data cannot be captured in a linear subspace (Hinton and Salakhutdinov, 2006). However, they are harder to compute and parameterize.

In network analysis, a novel way to reduce the dimensionality of large and complex datasets is by employing an explanatory tool for data analysis called “network histogram” (Olhede and Wolfe, 2014) illustrated in Figure 1. Sofia Olhede in joint work with Patrick Wolfe suggests a non-parametric approach to estimate consistently the generating model of a network. The idea is to identify groups such that nodes from the same group show a similar connectivity pattern and to approximate the generating model by a piecewise constant function by averaging over groups. The piecewise constant function here is called stochastic block model and it approximates the generating model in the same way as a histogram approximates a pdf or the Riemann sum approximates a continuous function.

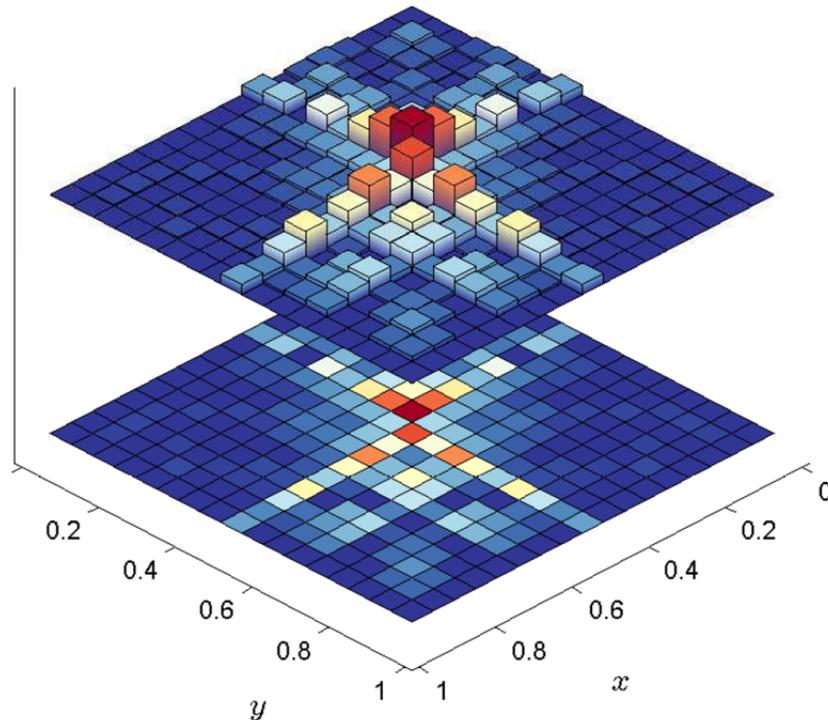


Figure 1. Example of a network histogram from Olhede and Wolfe (2014). A network can be expressed as a matrix where each entry represent the strength of connection between two nodes. Cleverly ordering the nodes uncovers the underlying structure of the network.

3.4 Sparsity and Regularization

Increasing model sparsity enforces a lower-dimensional model structure. In turn, it makes inference more tractable, models easier to interpret and it leads to more robustness against noise (Candes et al. 2008, Candes et al. 2006). Regularization techniques that enforce sparsity have been widely studied in the literature for more than a decade, including the lasso (Tibshirani, 1996), adaptive lasso (Zhou, 2006), the smoothly clipped absolute deviation (SCAD) penalty (Fan and Li, 2001), and modified Cholesky decomposition (Wu and Pourahmadi, 2003). These methods and their more recent extensions are now key tools in dealing with big data.

A strong assumption of sparsity makes it possible to use a higher number of parameters than observations, enabling us to find the optimal representation of the data while keeping the problem solvable. A good example in this context is the winning entry for the Netflix Challenge recommender system, which combined a large number of methods, thereby artificially increasing the number of parameters to improve prediction performance (Bell et al. 2008).

Sparsity-enforcing algorithms combine variable selection and model fitting: these methods typically shrink small estimated coefficients to zero and this typically introduces bias which

must be taken into account in evaluating the statistical significance of the non-zero estimated coefficients. Much of the research in lasso and related methods has focussed on point estimation, but inferential aspects such as hypothesis testing and construction of confidence intervals is an active area of research. For inference about coefficients estimated using the lasso, van de Geer and Bühlmann (2013), and Javanmard and Montanari (2014a, 2014b) introduce an additional step in the estimation procedure for debiasing. Lockhart et al. (2014) in contrast focus on testing whether the lasso model contains all truly active variables. In a survival analysis setting, Fang et al. (2014) test the effect of a low dimensional component in a high dimensional setting by treating all remaining parameters as nuisance parameters. They address the problem of the bias with the aim of improving the performance of the the classical likelihood-based statistics: the Wald statistic, the score statistics and the likelihood ratio statistic, all of these based on the usual partial likelihood for survival data.

To offer a compromise to the common sparsity assumption of zero and large non-zero valued elements, Ahmed (2013) proposes a relaxed assumption of sparsity by adding weak signals back into consideration. His new high-dimensional shrinkage method improves the prediction performance of a selected submodel significantly.

3.4 Optimization

Optimization plays a central role in statistical inference, and regularization methods that combine model fitting with model selection can lead to very difficult optimization problems. Complex models with many hidden layers, such as arise in the construction of neural networks and deep Boltzmann machines, typically lead to non-convex optimization problems.

Sometimes a convex relaxation of an initially non-convex problem is possible: the lasso for example is a relaxation of the problem that penalizes the number of non-zero coefficient estimates, i.e. it replaces an L_0 penalty with an L_1 penalty. While this convex relaxation makes the optimization problem much easier, it induces bias in the estimates of the parameters, as mentioned above.

Non-convex penalties can reduce such bias, but non-convex optimization is usually considered to be intractable. However, statistical problems are often sufficiently well-behaved that non-convex optimization methods become tractable. Moreover, due to the structures of the statistical models, Loh and Wainwright (2014, 2015) suggest that these models do not tend to be adversarial, hence it is not necessary to look at the worst case scenario. Even when non-convex optimization algorithms, such as the Expectation-Maximization algorithm, do not provide the global optimum, the local optimum can be adequate since the distance between the local and global maxima is often smaller than the statistical error (Balakrishnan et al., 2014, Loh and Wainwright, 2015). Wainwright concluded his presentation at the opening conference with a "speculative meta-theorem", that for any problem in which a

convex relaxation performs well, there is a first order simple method on the non-convex formulation which performs as good.

Applications of non-convex techniques include spectral analysis of tensors, analogous to SVD of matrices, which can be employed to aid in latent variable learning (Anandkumar, 2014). Non-convex methods can also be used in dictionary learning as well as in a version of robust PCA, which is less susceptible to sparse corruptions (Netrapalli, 2014). In some cases, convex reformulations have also been developed to find optimal solutions in non-convex scenarios (Guo and Schuurmans, 2007).

Many models used for big data are based on probability distributions, such as Gibbs distributions, for which the probability density function is only known up to its normalizing constant, sometimes called the partition function. Markov chain Monte Carlo methods attempt to overcome this by sampling from an unnormalized probability density function based on a Markov chain that has the target distribution as its equilibrium distribution, but these methods do not necessarily scale well. In some cases, approximate Bayesian computation (ABC) methods provide an alternative approximate, but computable, solution to a massive problem (Blum and Tran 2010).

3.5 Measuring distance

In high-dimensional data analysis, distances between “points”, or observations, are needed for most analytic procedures. For example optimization methods typically move through a possibly complex parameter space, and some notion of step-size for this movement is needed. Simulation methods such as Markov chain Monte Carlo sampling also need to ‘move’ through the parameter space. In both these contexts the use of Fisher information as a metric in the parameter space has emerged as a useful tool – it is the Riemannian metric (Rao, 1945; Amari and Nagaoka, 2000) when regarding the parameter space as a manifold.

Girolami et al. (2011) exploited this in Markov chain Monte Carlo algorithms, with impressive improvements in speed of convergence. Grosse (2015) showed how optimization of multi-layer Boltzmann machines could be improved by using an approximation to the Fisher information, the full Fisher information matrix being too high-dimensional to use routinely. A similar technique emerged in environmental science (Sampson and Guttorp, 1992), where two-dimensional geographical space is mapped to a two-dimensional deformed space in which the spatial structure becomes stationary and isotropic. Schmidt et al. (2011) describe how this approach can be exploited in a Bayesian framework, mapping to a deformed space in higher dimensions, and apply this to analyses of solar radiation and temperature data.

3.6 Representation Learning

Representation learning, also called feature learning or deep learning, aims to uncover hidden and complex structures in data (Bengio et. al, 2013) and frequently uses a layered

architecture to uncover a hierarchical data representation. The model for this layered architecture may have higher dimensionality than the original data. In that case, one may need to combine representation learning with the dimensionality reduction techniques described above in this report.

Large, well-annotated datasets facilitate successful representation learning. But labeling data for supervised learning can consume excessive time and resources, and may even prove impossible. To solve this problem, researchers have developed techniques to uncover data structure with unsupervised techniques (Anandkumar et al. 2013; Salakhutdinov and Hinton, 2012).

3.7 Sequential Learning

A strategy to deal with big data characterized by a high velocity is sequential or incremental learning. This technique treats the data successively, either because the task was designed in this way or the data does not allow any other access, as it is the case for continuous data streams such as long image sequences (Kalal et al. 2012). Sequential learning is also applied if the data is completely available, but has to be processed sequentially. This occurs when the data set is too large to fit into the memory or to be processed in a practically tractable way. For example, Scott et al. (2013) introduce consensus Monte Carlo, which operates by running separate Monte Carlo algorithms on distributed machines, and then averaging individual Monte Carlo draws across machines; the result can be nearly indistinguishable from the draws that would have been obtained by running a single machine algorithm for a very long time.

3.8 Multi-disciplinarity

A common theme in a great many presentations on big data is the necessity of a very high level of multi-disciplinarity. Statistical scientists are well-trained in both planning of studies, and in inference under uncertainty. Computer scientists bring a sophisticated mix of computational strategies and data management expertise. Both these groups of researchers need to work closely with scientists, social scientists, and humanists in the relevant application areas to ensure that the statistical methods and computational algorithms are effective for the scientific problem of interest. For instance, Dean et al. (2012) describe an exciting collaboration in forestry, where they are working on linking statistical analysis, computational infrastructure for new instrumentation, with social agency with the goal of maximizing their impact on the management of forest fires. The synergy space between them is created in an integrative approach involving scientists from all fields. This integration is not an invention of the world of big data by any means -- statistics is an inherently applied field and collaborations have been a major focus of the discipline from its earliest days. What has changed is the speed at which new technology provides new challenges to statistical analysis, and the magnitude of the computational challenges that go along with this.

Much of the big data now available involves information about people -- this is one of the aspects of big data that contributes to its popularity. The social sciences can help to elucidate potential biases in such data, as well as help to focus on the most interesting and relevant research questions. Humanists can provide frameworks for thinking about privacy, about historical data and changes through time, and about ethical issues arising in sharing of data about, for example, individuals' health. Individuals may well be willing to share their data for research purposes, for example, but not for marketing purposes: the humanities offer ways of thinking about these issues.

4. Examples of Applications

The breadth of applications is nearly limitless; we provide here a selection of the application areas discussed during the opening conference, in part to make the ideas outlined in the previous section more concrete.

4.1 Public Health

Diet is an important factor of public health in the study of disabilities (Murray et al., 2013), but assessing people's nutritional behaviour is difficult. Nielsen collects information about all products sold by a large number of groceries and corner stores at the 3-digit postal code level, and purchase card loyalty programs have data on purchases at the household level. Using these sources of information along with a UPC database of nutritional values, Buckeridge et al. (2012) match nutrition variables to existing medical records in the province of Québec, in collaboration with the *Institut de santé publique du Québec*. Since the purchase data is temporal, Buckeridge et al. (2014) measure the decrease in sales of sugary carbonated drinks following a marketing campaign, hence assessing the effectiveness of this public health effort.

Big data is also being used for syndromic surveillance, the monitoring of the syndromes of transmittable diseases at the population level. In the province of Québec, a central system monitors the triage data from emergency rooms including free form text which needs to be treated appropriately. When the automated surveillance system issues an alert, an ad-hoc analysis is performed to decide whether immediate actions are needed. Buckeridge et al. (2006) evaluate the performance of such systems through simulation for different rates of false positive alerts.

Social media produce data whose volume and variety are often challenging. The traditional reporting of flu from physician reports generally introduces a 1-2 week time delay. Google Flu Trends (Ginsberg et. al 2009) speeds up the reporting by tabulating the use of search keywords, such as "runny nose" or "sore throat", in order to estimate the prevalence of

influenza-like illnesses. Predictions can sometimes go wrong, see Butler (2013), but the gain in reaction time is significant. Using the GPS location in the metadata of Twitter posts Ramírez-Ramírez et al. (2013) introduce Simulation of Infectious Diseases (SIMID) that improve geographic and temporal locality beyond Google Flu Trends. They illustrate the use of SIMID to estimate and predict flu spread in the Peel sub-metropolitan region near Toronto.

4.2 Health Policy

Many examples in health policy rely on the linkage of large administrative data sets. In Ontario, the Institute for Clinical Evaluative Sciences (ICES) is a province-wide research network that collects and links Ontario's health-related data at the individual level. Similar databases exist in other provinces. Comparing data on diabetes case ascertainment from Manitoba with those of Newfoundland and Labrador, Lix et al. (2012) evaluate the incompleteness of physician billing claims and estimate to which extent the number of non-fee-for-service acts is underestimated. With the quality of the data in mind, the authors also improve disease ascertainment with model-based predictive algorithms as an alternative to relying on possibly incorrect disease codes.

Health data can be spatio temporal and complications arise quickly when the data are heterogeneous. Postal codes and census tracts, for instance, define areas that do not match, with shapes that vary through time (see e.g. Nguyen et al., 2012). In a study about cancer in Nova Scotia, Lee et al. (2015) have access to location data for the most recent years, but older data only contain the postal codes. To handle the different granularities, they use a hierarchical model where the location of cancer cases are treated as a latent variable and higher levels of the hierarchy are the aggregation units (Li et al., 2012).

4.3 Law and Order

Insurgencies and riots are frequent in South America, with hundreds or thousands of occurrences in each country every year. Korkmaz et al. (2015) have a database of all Twitter messages sent from South America over a period of four years as well as a binary variable named GSR, a gold standard to determine whether an event was an insurgency or not. The occurrence of an insurgency can be predicted by the volume of tweets, the presence of some keywords therein (from a total of 634 who were initially identified as potentially useful), and an increased use of *The Onion Router* (TOR), an online service to anonymize tweets. The massive database containing all the tweets is stored on a Hadoop file system and its treatment is done with Python through mapreduce. Those tools which are not part of a typical statistics cursus are standard in the treatment of massive data sets.

4.4 Environmental Sciences

The world is facing global environmental challenges which are hard to model, given their inherent complexity. To estimate the temperature record since 1000, Barboza et al. (2014) use different proxies including tree rings, pollen, coral, ice cores, hence combining very

heterogeneous data. From earlier work, the authors found that historical records can be much better predicted by incorporating climate forcings in the model, for instance solar irradiance, greenhouse gas concentration and volcanism.

In an example combining health and environment data, Ensor et al. (2013) analyze the location and time of cardiac arrests from 2004 to 2011 in the city of Houston, with respect to air pollution measurements taken by EPA sensors scattered across the city. Instead of considering daily averages of exposition to air pollution, this study measured the exposure of an individual in the three hours preceding a cardiac arrest and find a significant effect on the risk of an out-of-hospital cardiac arrest. In an unpublished follow-up study discussed by Sallie Keller, a synthetic model is used to simulate the activities of 4.4 million people in Houston on a given day. The individual exposure varies greatly, even within a given household, as people's exposure to pollution depends on their movements in the city. The synthetic data helped optimizing CPR trainings geographically in the city.

4.5 Education

In a research project to study the theory of Johnson (2003) about grade inflation, Reese et al. (2014) analyze the history of Brigham Young University student evaluations from 2002 to 2014, matching it with student records. A total of 2.9 million course evaluations from 185,000 student were used. The variables of interest are three questions answered on an ordinal scale and a free text section for which a sentiment analysis determines if the comment is positive or negative. The magnitude of the Bayesian proportional odds model fitted to the ordinal response variables is too big for MCMC methods, but the ABC approach leads to a computationally feasible solution. The data shows empirical evidence of grade inflation.

With an example based on education data, Cook (2014) illustrates how visualization may be used for inference. Based on the 2012 OECD Pisa tests results from 485,900 students in 65 countries, Cook shows a figure presenting the difference in mean of math scores and reading scores between boys and girls by country. While the boys seem to outperform the girls in math, a much greater difference is seen in favour of the girls for the reading scores. To have an idea of the significance of the differences, Cook produces similar plots where the gender labels are randomly assigned to scores within each country. The fact that the true data stands out in comparison to the random assignments indicates that the gap did not occur by chance.

4.6 Mobile Application Security

Conveying information of interest from big data is a challenge which may be addressed by visualization. Hosseinkhani et al. (2014) propose Papilio, a novel graphical representation of a network that displays the different permissions of Android apps. The creation of a tractable representation is made possible by taking advantage of the special structure of the data. The whole representation is still too large for a single screen, so halos (Baudisch, P. and

Rosenholtz, 2003) are used to indicate the position of off screen nodes, as illustrated in Figure 2. The radius of the halos are proportional to the distance to the node.

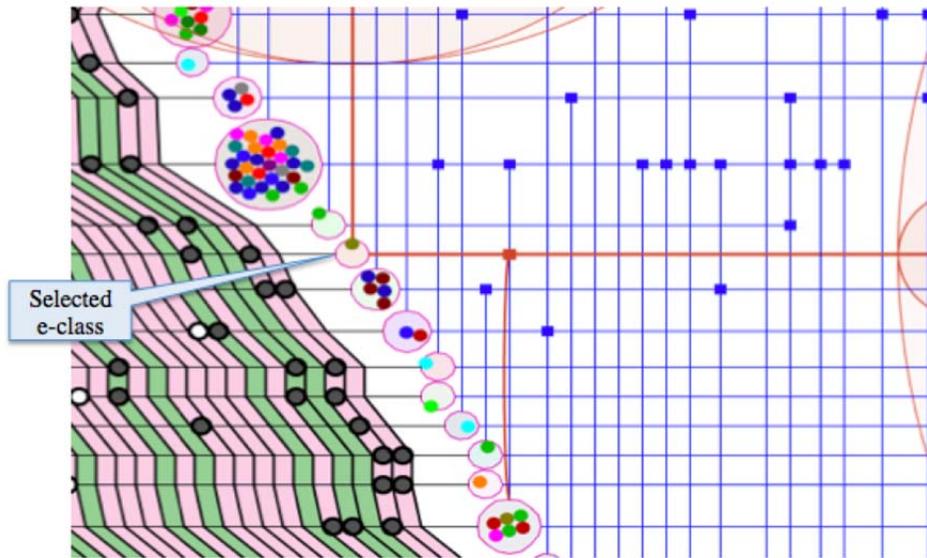


Figure 2: Illustration of the Halos in Papilio, which indicate the position of off screen nodes. This Papilio figure from Hosseinkhani et al. (2014) displays the different permissions of Android apps.

4.7 Image Recognition and Labelling

Images are unstructured data and their analysis progressed significantly with the development of deep learning models such as deep Boltzmann machines that are especially suitable for automatically learning features in the data. The Toronto Deep Learning group offer an online demo⁵ of multi-modal learning based on the work of Srivastava and Salakhutdinov (2012) that can generate a caption to describe any uploaded image. The image search engine of Google also uses deep learning models to find images that are similar to an uploaded picture.

4.8 Digital Humanities

A large body of research in the humanities is built from an in-depth study of carefully selected authors or records. The ability to process large amounts of data is opening new opportunities where social evidence can be gathered at a population level. For instance, Drummond et al. (2006) use the 1901 Canadian census data to revise different theories about bilingualism in Canada and improve their agreement with the data. They also observe that patterns in Québec differ significantly from those in the rest of the country.

⁵ <http://deeplearning.cs.toronto.edu/i2t>

In his book *Capital in the Twenty-First Century*, Piketty (2014) discusses the evolution of references to money in novels as a reflection of inequalities in the society. Underwood et al. (2014) dig further into the question by listing all references to money amounts in English-language novels from 1750 to 1950. They find that the relative frequency of references to money amounts increases through time, but the amount mentioned decreases in constant dollar amount. They do not attempt to draw strong conclusions from this social evidence, mentioning for instance that the changes in the reading audience due to an increased literacy among the working class may also play a role in the evolution of the novels.

On a larger scale, Michel et al. (2011) analyze about 4% of all books ever printed to describe social trends through time as well as changes in linguistics. While patterns in the data match milestones events such as the Great Wars and some global epidemics, the evolution of words such as “feminism” measure important social trends. Local events such as the censorship of different artists for a number of years in different countries are also visible in the data. Those are some examples of ‘culturomics’, the quantitative analysis of digitized text to study cultural trends and human behavior.

4.9 Materials Science

In an application to Materials Science by Nelson et al. (2013), the properties of a binary alloy depend on the relative position of atoms from both materials whose description relies on a very large number of explanatory variables. Surprisingly, as Candès, E. J. et al. (2006) showed, a random generation of candidate materials is more efficient than a carefully designed choice. This is an example where big data shows a counter intuitive behaviour, since one could think that careful design of experiment should be able to beat random selections.

5. Conclusion

While technical achievements make the existence and availability of very large amounts of data possible, the challenges emerging from such data goes beyond processing, storing and accessing records rapidly. As its title indicates, the thematic program on Statistical Inference, Learning and Models for Big Data focussed on the data analysis aspects of big data.

In this report, we identify challenges and solutions that are common across different disciplines working on inference and learning for big data. With new magnitudes of dimensionality, complexity and heterogeneity in the data, researchers are faced with new challenges of scalability, quality and visualization. In addition, big data is most often observational in nature, having been collected for a purpose other than the one intended, and hence making inference even more challenging.

Emerging technological solutions provide new opportunities, but they also define new paradigms for data and models. Making the models bigger with layers of latent variables may provide an insight into the features of the data, while reducing the data to lower dimensional models can be key in getting any calculations through. Lower dimensional structure analysis manifests itself not only as sparsity assumptions and dimensionality reduction, but also in visualization. Other techniques such as non-convex algorithms, sampling strategies and matrix randomization aid in the processing of large data sets. With massive data, settling for an approximate solution seems to be the saving compromise, whether by allowing for more flexibility in an objective function or by replacing a standard algorithm that is infeasible at such a large scale.

Many of the most striking examples of big data analysis presented during this thematic program came from teams of many scientists with complementary backgrounds. Moving forward, research teams need to scale up, inviting talents from different fields to come together and join forces. The term “big data” has gotten a lot of media attention in recent years, creating a hype that will certainly fade. The new problems that emerged, however, are here to stay and the need for innovative solutions will certainly keep our community busy for many years to come.

Acknowledgments

We are grateful for the Fields Institute for Research in the Mathematical Sciences, which provided the major portion of the funding for the thematic program on Statistical Inference, Learning and Models in Big Data. The Fields Institute is funded by the Natural Sciences and Engineering Research Council of Canada and the Ministry of Training, Colleges and Universities in Ontario. Funding was also provided for specific workshops by the Pacific Institute for the Mathematical Sciences, the Centre de Recherche Mathématiques, the Atlantic Association for Advancement in the Mathematical Sciences, the Department of Statistics at Iowa State University, the National Science Foundation, and the Canadian Statistical Sciences Institute.

References

1000 Genomes Project Consortium. An integrated map of genetic variation from 1,092 human genomes. *Nature* **491**(7422) 56-65.

AHMED, S. E. (2014). *Penalty, shrinkage and pretest strategies: Variable selection and estimation*. Springer International Publishing.

- ALBERT-GREEN, A., BRAUN, W. J., MARTELL, D. L. and WOOLFORD, D. G. (2014). Visualization tools for assessing the Markov property: Sojourn times in the Ontario Fire Weather Index. *Environmetrics* **25**(6) 417-430.
- AMARI, SHUN-ICHI and NAGAOKA, I. (2000). *Methods of information geometry*. Translations of Mathematical Monographs, (Vol. 191). American Mathematical Society, Providence, RI; Oxford University Press, Oxford.
- ANANDKUMAR, A., GE, R., HSU, D., KAKADE, S. M. and TELGARSKY, M. (2014). Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research* **15**(1) 2773-2832.
- ANANDKUMAR, A., HSU, D., JANZAMIN, M. and KAKADE, S. M. (2013). When are overcomplete topic models identifiable? Uniqueness of tensor tucker decompositions with structured sparsity. In *Advances in Neural Information Processing Systems* 1986-1994.
- BALAKRISHNAN, S., WAINWRIGHT, M. J. and YU, B. (2014). Statistical guarantees for the EM algorithm: From population to sample-based analysis. *arXiv preprint arXiv:1408.2156*.
- BARBOZA, L., LI, B., TINGLEY, M. P. and VIENS, F. G. (2014). Reconstructing past temperatures from natural proxies and estimated climate forcings using short-and long-memory models. *The Annals of Applied Statistics* **8**(4) 1966-2001.
- BAUDISCH, P. and ROSENHOLTZ, R. (2003, April). Halo: a technique for visualizing off-screen objects. In *Proceedings of the SIGCHI conference on Human factors in computing systems* 481-488. ACM.
- BELL, R. M., KOREN, Y. and VOLINSKY, C. (2008). The bellkor 2008 solution to the netflix prize. *Statistics Research Department at AT&T Research*.
- BENGIO, Y., COURVILLE, A. and VINCENT, P. (2013). Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **35**(8) 1798-1828.
- BIERNACKI, P. and WALDORF, D. (1981). Snowball sampling: Problems and techniques of chain referral sampling. *Sociological methods and research* **10**(2) 141-163.
- BINGHAM, E. and MANNILA, H. (2001, August). Random projection in dimensionality reduction: applications to image and text data. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining* 245-250. ACM.

- BIUCAS-DIAS, J. M., PLAZA, A., CAMPS-VALLS, G., SCHEUNDERS, P., NASRABADI, N. M. and CHANUSSOT, J. (2013). Hyperspectral remote sensing data analysis and future challenges. *Geoscience and Remote Sensing Magazine, IEEE* **1**(2) 6-36.
- BLUM, M. G. and TRAN, V. C. (2010). HIV with contact tracing: a case study in approximate Bayesian computation. *Biostatistics* **11**(4) 644-660.
- BRIEN, S., NADERI, N., SHABAN-NEJAD, A., MONDOR, L., KROEMKER, D. and BUCKERIDGE, D. L. (2013, May). Vaccine attitude surveillance using semantic analysis: constructing a semantically annotated corpus. In *Proceedings of the 22nd international conference on World Wide Web companion* 683-686. International World Wide Web Conferences Steering Committee.
- BROWN, P. E., NGUYEN, P., STAFFORD, J. and ASHTA, S. (2013). Local-EM and mismeasured data. *Statistics and Probability Letters* **83**(1) 135-140.
- BUCKERIDGE, D. L., CHARLAND, K., LABBAN, A. and MA, Y. (2014). A method for neighborhood-level surveillance of food purchasing. *Annals of the New York Academy of Sciences* **1331**(1) 270-277. doi: 10.1111/nyas.12332
- BUCKERIDGE, D. L., IZADI, M., SHABAN-NEJAD, A., MONDOR, L., JAUVIN, C., DUBÉ, L., ... and TAMBLYN, R. (2012). An infrastructure for real-time population health assessment and monitoring. *IBM Journal of Research and Development* **56**(5) 2-1.
- BUCKERIDGE, D. L., OWENS, D. K., SWITZER, P., FRANK, J. and MUSEN, M. A. (2006). Evaluating detection of an inhalational anthrax outbreak. *Emerging Infectious Diseases* **12**(12) 1942-1949.
- BUJA, A., COOK, D., HOFMANN, H., LAWRENCE, M., LEE, E. K., SWAYNE, D. F. and WICKHAM, H. (2009). Statistical inference for exploratory data analysis and model diagnostics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **367**(1906) 4361-4383.
- BUTLER, D. (2013). When Google got flu wrong. *Nature*, **494**(7436), 155.
- CANDÈS, E. J., ROMBERG, J. and TAO, T. (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on* **52** (2) 489-509.

- CANDÈS, E. J., WAKIN, M. B. and BOYD, S. P. (2008). Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier analysis and applications* **14**(5-6) 877-905.
- CARPENDALE, S., LIGH, J. and PATTISON, E. (2004). Achieving higher magnification in context. In *Proceedings of the 17th annual ACM symposium on User interface software and technology* 71-80.
- CARPENDALE, S. (2015). Innovations in visualization.
<http://innovis.cpsc.ucalgary.ca/Research/Research> accessed on June 24, 2015.
- CHEN, M., MAO, S. and LIU, Y. (2014). Big data: A survey. *Mobile Networks and Applications* **19**(2) 171-209.
- COOK, D. (2014). Visiphilia: The Gender Gap in Math is Not Universal. *CHANCE* **27**(4) 48-52.
- CHUNG, F. and LU, L. (2002). The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences* **99**(25) 15879-15882.
- DEAN, C., BRAUN, W. J., MARTELL, D. and WOOLFORD, D. (2012). Interdisciplinary Training in a Collaborative Research Environment. Chrisman and Wachowicz (Eds.), *The added value of scientific networking: Perspectives from the GEOIDE network members 1998-2012*. GEOIDE Network, Quebec, Canada.
- DEAN, J. and GHEMAWAT, S. (2008). MapReduce: simplified data processing on large clusters. *Communications of the ACM* **51**(1) 107-113.
- DIACONIS, P. and JANSON, S. (2007). *Graph limits and exchangeable random graphs* arXiv: 0712.2749. Department of Mathematics, Uppsala University.
- DONOHO, D., GAVISH, M. AND MONTANARI, A. (2013). The phase transition of matrix recovery from Gaussian measurements matches the minimax MSE of matrix denoising. *Proceedings of the National Academy of Sciences* **110**(21) 8405-8410.
- DONOHO, D. and GAVISH, M. (2014). Minimax risk of matrix denoising by singular value thresholding. *The Annals of Statistics* **42**(6) 2413-2440.
- DRUMMOND, C., MATWIN, S., & GAFFIELD, C. (2006). Inferring and revising theories with confidence: analyzing bilingualism in the 1901 Canadian census. *Applied Artificial Intelligence* **20**(1), 1-33.
- ELDAR, Y. C. AND KUTYNIOK, G. (2012). *Compressed sensing: theory and applications*. Cambridge University Press.

- FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association* **96**(456) 1348-1360.
- FANG, E. X., NING, Y. and LIU, H. (2014). Testing and confidence intervals for high dimensional proportional hazards model. *arXiv preprint arXiv:1412.5158*.
- ENSOR, K. B., RAUN, L. H. and PERSSE, D. (2013). A Case-crossover analysis of out-of-hospital cardiac arrest and air pollution. *Circulation* **127** 1192-1199.
- GAFFIELD, C. (2015). Big data vs. human complexity: an early status report on the central question of the 21st century. Presented at the Fields Institute, January 20th 2015. available at <http://www.fields.utoronto.ca/video-archive/2015/01/315-4136>, accessed on June 30, 2015.
- GEIGER, A., LENZ, P. and URTASUN, R. (2012). Are we ready for autonomous driving? The kitti vision benchmark suite. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* 3354-3361.
- GINSBERG, J., MOHEBBI, M. H., PATEL, R. S., BRAMMER, L., SMOLINSKI, M. S. and BRILLIANT, L. (2009). Detecting influenza epidemics using search engine query data. *Nature* **457**(7232), 1012-1014.
- GIROLAMI, M. and CALDERHEAD, B. (2011). Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(2) 123-214.
- GROSSE, R. (2015). Scaling up natural gradient by sparsely factorizing the inverse Fisher information matrix. *International Conference on Machine Learning*, to appear.
- GUO, Y. and SCHUURMANS, D. (2007). Convex relaxations of latent variable training. In *Advances in Neural Information Processing Systems* 601-608.
- GUPTA, V. and LEHAL, G. S. (2009). A survey of text mining techniques and applications. *Journal of emerging technologies in web intelligence* **1**(1) 60-76.
- GUTTORP, P. and SCHMIDT, A. M. (2013). Covariance structure of spatial and spatiotemporal processes. *Wiley Interdisciplinary Reviews: Computational Statistics* **5**(4) 279-287.
- GUYON, I. and ELISSEEFF, A. (2003). An introduction to variable and feature selection. *The Journal of Machine Learning Research* **3** 1157-1182.

- HALKO, N., MARTINSSON, P. G. and TROPP, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review* **53**(2) 217-288.
- HARFORD, T. (2014). Big data: are we making a big mistake? *Financial Times* March 24, 2014.
- HINTON, G. E. and SALAKHUTDINOV, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science* **313**(5786) 504-507.
- HODAS, N. O. and LERMAN, K. (2014). The simple rules of social contagion. *Scientific Reports* **4** 4343. doi:10.1038/srep04343.
- HOENS, T. RYAN, POLIKAR, ROBI, CHAWLA, NITESH V. (2012). Learning from streaming data with concept drift and imbalance: an overview. *Progress in Artificial Intelligence*, 1. Jg., Nr. 1, S. 89-101.
- HOSSEINKHANI, M., FONG, P. W. L. and CARPENDALE, S. (2014). Papilio: Visualizing Android application permissions. In *Proceedings of the EuroVis Conference* **33**(3), 391- 400.
- HSIEH, C. J., SUSTIK, M. A., DHILLON, I. S., RAVIKUMAR, P. K. and POLDRACK, R. (2013). BIG & QUIC: Sparse inverse covariance estimation for a million variables. In *Advances in Neural Information Processing Systems* 3165-3173.
- HURON, S., CARPENDALE, S., THUDT, A., TANG, A. and MAUERER, M. (2014, June). Constructive visualization. In *Proceedings of the 2014 conference on Designing interactive systems* 433-442. ACM.
- IZADI, M., SHABAN-NEJAD, A., OKHMATOVSKAIA, A., MONDOR, L. and BUCKERIDGE, D. L. (2013). Population health record: An informatics infrastructure for management, integration, and analysis of large scale population health data. *Proceeding of AAAI (HIAI 2013)*, Bellevue, Washington, USA July, 14-18.
- JAIN, S., LANGE, S. and ZILLES, S. (2006, January). Towards a better understanding of incremental learning. In *Algorithmic Learning Theory* (169-183). Springer Berlin Heidelberg.
- JAVANMARD, A. and MONTANARI, A. (2014a). Hypothesis testing in high-dimensional regression under the Gaussian random design model: Asymptotic theory, *Information Theory, IEEE Transactions* **60**(10) 6522-6554.
- JAVANMARD, A. and MONTANARI, A. (2014b). Confidence intervals and hypothesis testing for high-dimensional regression. *The Journal of Machine Learning Research* **15**(1) 2869-2909.

- JOHNSON, V. E. (2003). *Grade Inflation: A crisis in college education*, Springer-Verlag, New York.
- KALAL, Z., MIKOLAJCZYK, K. and MATAS, J. (2012). Tracking-learning-detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **34**(7) 1409-1422.
- KASS, R.E. (1989) The geometry of asymptotic inference. *Statistical Science*. **4**(3) 188–234.
- KOLACZYK, E. D. (2009). *Statistical analysis of network data: methods and models*. Springer Science and Business Media.
- KORKMAZ, G., REESE, C.S., HIGDON, D.M., AND KELLER, S.K. (2015). *Insurgency Prediction using Social Media Streams*. In Preparation.
- LEE, J. S., NGUYEN, P., BROWN, P. E., STAFFORD, J. and SAINT-JACQUES, N. (2015). *Local-EM Algorithm for spatio-temporal analysis with application in Southwestern Nova Scotia*, Unpublished Manuscript.
- LI, Y., BROWN, P., GESINK, D. C. and RUE, H. (2012). Log Gaussian Cox processes and spatially aggregated disease incidence data. *Statistical methods in medical research* **21**(5) 479-507.
- LI, L. J. and FEI-FEI, L. (2010). Optimol: automatic online picture collection via incremental model learning. *International journal of computer vision* **88**(2) 147-168.
- LIX, L. M., SMITH, M., AZIMAE, M., DAHL, M., NICOL, P., BURCHILL, C., BURLAND, E., ... and BAILLY A. (2012). *A systematic investigation of Manitoba's provincial laboratory data*. Winnipeg, MB: Manitoba Centre for Health Policy.
- LOCKHART, R., TAYLOR, J., TIBSHIRANI, R. J. and TIBSHIRANI, R. (2014). A significance test for the lasso. *Annals of Statistics* **42**(2) 413-468.
- LOH, P. L. and WAINWRIGHT, M. J. (2014). Support recovery without incoherence: A case for nonconvex regularization. *arXiv preprint arXiv:1412.5632*.
- LOH, P. L. and WAINWRIGHT, M. J. (2015). Regularized M-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *Journal of Machine Learning Research* **16** 559-616.
- MICHEL, J. B., SHEN, Y. K., AIDEN, A. P., VERES, A., GRAY, M. K., PICKETT, J. P., ... & AIDEN, E. L. (2011). Quantitative analysis of culture using millions of digitized books. *Science* **331**(6014), 176-182.
- MONTELEONI, C., SCHMIDT, G. A., SAROHA, S. and ASPLUND, E. (2011). Tracking climate models. *Statistical Analysis and Data Mining: The ASA Data Science Journal* **4**(4) 372-392.

- MURRAY, C. J., VOS, T., LOZANO, R., NAGHAVI, M., FLAXMAN, A. D., MICHAUD, C., ... and BRIDGETT, L. (2013). Disability-adjusted life years (DALYs) for 291 diseases and injuries in 21 regions, 1990–2010: a systematic analysis for the Global Burden of Disease Study 2010. *The Lancet* **380**(9859) 2197-2223.
- NACENTA, M., HINRICHS, U. and CARPENDALE, S. (2012, May). FatFonts: combining the symbolic and visual aspects of numbers. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* 407-414. ACM.
- NELSON, L. J., OZOLIŅŠ, V., REESE, C. S., ZHOU, F. and HART, G. L. (2013). Cluster expansion made easy with Bayesian compressive sensing. *Physical Review B* **88**(15), 155105.
- NETRAPALLI, P., NIRANJAN, U. N., SANGHAVI, S., ANANDKUMAR, A. and JAIN, P. (2014). Non-convex robust pca. In *Advances in Neural Information Processing Systems* 1107-1115.
- NGUYEN, P., BROWN, P. E. and STAFFORD, J. (2012). Mapping cancer risk in southwestern Ontario with changing census boundaries. *Biometrics* **68**(4) 1228-1237.
- OLHEDE, S. C. and WOLFE, P. J. (2014). Network histograms and universality of blockmodel approximation. *Proceedings of the National Academy of Sciences* **111**(41) 14722-14727.
- ORBANZ, P. and ROY, D. M. (2013). Bayesian models of graphs, arrays and other exchangeable random structures. *arXiv preprint arXiv:1312.7857*.
- PANG, B. and LEE, L. (2008). Opinion mining and sentiment analysis. *Foundations and trends in information retrieval* **2**(1-2) 1-135.
- PERSON, K. (1901). On lines and planes of closest fit to systems of points in space. *Philosophical Magazine* **2**(6) 559-572.
- PIKETTY, T. (2014). Capital in the twenty-first century. *Cambridge, MA, London*.
- RAMÍREZ-RAMÍREZ, L. L., GEL, Y. R., THOMPSON, M., DE VILLA, E. and MCPHERSON, M. (2013). A new surveillance and spatio-temporal visualization tool SIMID: SIMulation of Infectious Diseases using random networks and GIS. *Computer Methods and Programs in Biomedicine* **110**(3) 455-470.
- RAO, C. R. (1945). Information and the accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Math. Soc.* **37** 81-89.

- REESE, C.S., TOLLEY, H.D., KEITH, J., and BURLINGAME, G. (2014). *Unraveling the mystery of grade inflation and student ratings: A BYU Case Study*, BYU Department of Statistics Technical Report, TR-14-021.
- ROBERTSON, J., MCELDOFF, P., PEARSON, S. A., HENRY, D. A., INDER, K. J. and ATTIA, J. R. (2012). The health services burden of heart failure: an analysis using linked population health data-sets. *BMC Health Services Research* **12**(1) 103.
- SALAKHUTDINOV, R. and HINTON, G. (2012). An efficient learning procedure for deep Boltzmann machines. *Neural Computation* **24**(8) 1967-2006.
- SCHADT, E. E., LINDERMAN, M. D., SORENSON, J., LEE, L. and NOLAN, G. P. (2010). Computational solutions to large-scale data management and analysis. *Nature Reviews Genetics* **11**(9) 647-657.
- SCHMIDT, A. M., GUTTORP, P. and O'HAGAN, A. (2011). Considering covariates in the covariance structure of spatial processes. *Environmetrics* **22**(4) 487-500.
- SCOTT, S. L., BLOCKER, A. W., BONASSI, F. V., CHIPMAN, H. A., GEORGE, E. I. and MCCULLOCH, R. E. (2013, October). Bayes and big data: The consensus Monte Carlo algorithm. In *EFaBBayes 250 conference* 16.
- SIMPSON, D., ILLIAN, J., LINDGREN, F., SØRBYE, S. and RUE, H. (2011). Going off grid: Computationally efficient inference for log-Gaussian Cox processes. *arXiv preprint arXiv:1111.0641*.
- SRIVASTAVA, N. and SALAKHUTDINOV, R. R. (2012). Multimodal learning with deep boltzmann machines. In *Advances in Neural Information Processing Systems* 2222-2230.
- STUKEL, T. A., GLAZIER, R. H., SCHULTZ, S. E., GUAN, J., ZAGORSKI, B. M., GOZDYRA, P. and HENRY, D. A. (2013). Multispecialty physician networks in Ontario. *Open Medicine* **7**(2) e40.
- TARRAN, B. (2014). Now you see me, now you don't. *Significance* **11**(4) 11-14.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society* **58**(1) 267-288.
- TUKEY, J. W. (1977). *Exploratory data analysis*. Addison-Wesley.
- UNDERWOOD, T., LONG, H. and RICHARD, J.S. (2014). Cents and Sensibility: Trust Thomas Piketty on economic inequality. Ignore what he says about literature. Blog post on Slate.com. Retrieved from

http://www.slate.com/articles/business/moneybox/2014/12/thomas_piketty_on_literature_balzac_austen_fitzgerald_show_arc_of_money.html on June 29th, 2015.

VAN DE GEER, S., BÜHLMANN, P., RITOV, Y. A. and DEZEURE, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics* **42**(3) 1166-1202.

WICKHAM, H., COOK, D., HOFMANN, H. and BUJA, A. (2010). Graphical inference for infovis. *Visualization and Computer Graphics, IEEE Transactions on* **16**(6) 973-979.

WITTEN, R. and CANDÈS, E. (2013). Randomized algorithms for low-rank matrix factorizations: sharp performance bounds. *Algorithmica* 1-18.

WOOLFORD, D. G. and BRAUN, W. J. (2007). Convergent data sharpening for the identification and tracking of spatial temporal centers of lightning activity. *Environmetrics* **18**(5) 461-479.

WOOLFORD, D. G., BRAUN, W. J., DEAN, C. B. and MARTELL, D. L. (2009). Site-specific seasonal baselines for fire risk in Ontario. *Geomatica* **63**(4) 355-363.

WU, W. B. and POURAHMADI, M. (2003). Nonparametric estimation of large covariance matrices of longitudinal data. *Biometrika* **90**(4) 831-844.

ZOU, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association* **101**(476) 1418-1429.