Compressive Big Data Analytics

Ivo D. Dinov

Statistics Online Computational Resource
Health Behavior & Biological Sciences
Computational Medicine & Bioinformatics
Michigan Institute for Data Science

University of Michigan

www.SOCR.umich.edu

SCHOOL OF NURSING
STATISTICS ONLINE COMPUTATIONAL RESOURCE (SOCR)
UNIVERSITY OF MICHIGAN

Outline

- Driving biomedical & health challenges
- Common characteristics of Big Biomedical Data
- Data science & predictive analytics
- Compressive Big Data Analytics (CBDA)
- Case-studies
  - Applications to Neurodegenerative Disease
  - Data Dashboarding

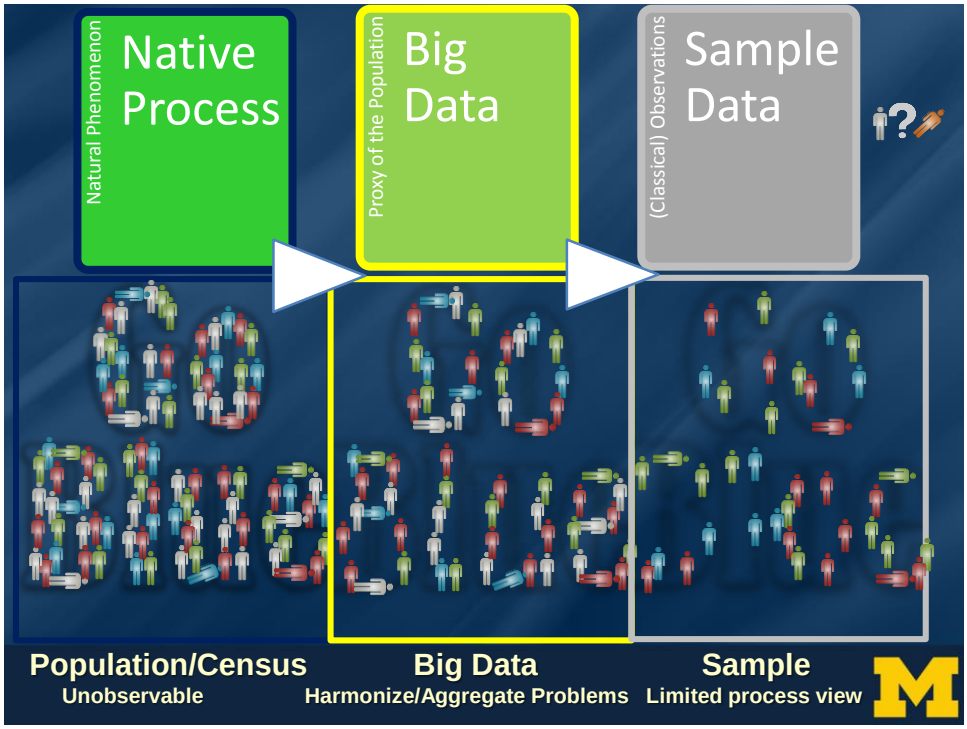

Driving Biomedical/Health Challenges

- **Neurodegeneration**: Structural Neuroimaging in Alzheimer's Disease illustrates the Big Data challenges in modeling complex neuroscientific data. 808 ADNI subjects, 3 groups: 200 subjects with Alzheimer's disease (AD), 383 subjects with mild cognitive impairment (MCI), and 225 asymptomatic normal controls (NC). The 80 neuroimaging biomarkers and 80 highly-associated SNPs.

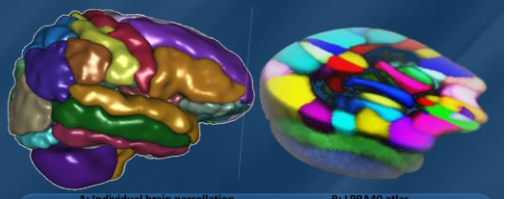

A: Individual brain parcellation B: LPBA40 atlas

Index	Volume Intensity	ROI Name	Index	Volume Intensity	ROI Name
1	21	L superior frontal gyrus	29	82	L inferior occipital gyrus
2	74	R middle frontal gyrus	30	164	R putamen
3	50	R precuneus	31	65	L superior occipital gyrus
4	101	cerebellum	32	30	R middle orbitofrontal gyrus
5	47	L angular gyrus	33	42	R postcentral gyrus
6	122	R cingulate gyrus	34	27	L precentral gyrus
7	83	L middle temporal gyrus	35	32	R lateral orbitofrontal gyrus
8	90	R lingual gyrus	36	121	L cingulate gyrus
9	81	L superior temporal gyrus	37	31	L lateral orbitofrontal gyrus
10	91	L fusiform gyrus	38	92	R fusiform gyrus
11	44	R superior parietal gyrus	39	45	L supramarginal gyrus
12	66	R inferior occipital gyrus	40	88	R parahippocampal gyrus
13	87	L parahippocampal gyrus	41	22	R superior frontal gyrus
14	162	R caudate	42	29	L middle orbitofrontal gyrus
15	85	L inferior temporal gyrus	43	68	R cuneus
16	182	brainstem	44	62	R superior occipital gyrus
17	43	L superior parietal gyrus	45	33	L gyrus rectus
18	28	R precentral gyrus	46	48	R angular gyrus
19	23	L middle frontal gyrus	47	64	R middle occipital gyrus
20	89	L lingual gyrus	48	84	R middle temporal gyrus
21	41	L postcentral gyrus	49	49	L precuneus

http://DSPA.predictive.space

Moon, Dinov, *et al.* (2015)

Characteristics of Big Biomed Data

IBM Big Data 4V's: Volume, Variety, Velocity & Veracity

Big Bio Data Dimensions	Tools
Size	Harvesting and management of vast amounts of data
Complexity	Wranglers for dealing with heterogeneous data
Incongruency	Tools for data harmonization and aggregation
Multi-source	Transfer and joint modeling of disparate elements
Multi-scale	Macro to meso to micro scale observations
Incomplete	Reliable management of missing data

Example: analyzing observational data of 1,000's Parkinson's disease patients based on 10,000's signature biomarkers derived from multi-source imaging, genetics, clinical, physiologic, phenomics and demographic data elements

Software developments, student training, service platforms and methodological advances associated with the Big Data Discovery Science all present existing opportunities for learners, educators, researchers, practitioners and policy makers

Dinov, *et al.* (2016) PMID:26918190

Data Science & Predictive Analytics

- **Data Science**: an emerging extremely transdisciplinary field - bridging between the theoretical, computational, experimental, and applied areas. Deals with enormous amounts of complex, incongruent and dynamic data from multiple sources. Aims to develop algorithms, methods, tools, and services capable of ingesting such datasets and supplying semi-automated decision support systems
- **Predictive Analytics**: process utilizing advanced mathematical formulations, powerful statistical computing algorithms, efficient software tools, and distributed web-services to represent, interrogate, and interpret complex data. Aims to forecast trends, cluster patterns in the data, or prognosticate the process behavior either within the range or outside the range of the observed data (e.g., in the future, or at locations where data may not be available)

http://DSPA.predictive.space

Dinov, Springer (2018)

Big Data	Information	Knowledge	Action
Raw Observations	Processed Data	Maps, Models	Actionable Decisions
Data Aggregation	Data Fusion	Causal Inference	Treatment Regimens
Data Scrubbing	Summary Stats	Networks, Analytics	Forecasts, Predictions
Semantic-Mapping	Derived Biomarkers	Linkages, Associations	Healthcare Outcomes

Dinov, et al. (2016) PMID:26918190

Case-Studies – ALS

- Identify predictive classifiers to detect, track and prognosticate the progression of ALS (in terms of clinical outcomes like ALSFRS and muscle function)
- Provide a decision tree prediction of adverse events based on subject phenotype and 0-3 month clinical assessment changes

Data Source	Sample Size/Data Type	Summary
ProAct Archive	Over 100 variables are recorded for all subjects including: Demographics: age, race, medical history, sex; Clinical data: **Amyotrophic Lateral Sclerosis** Functional Rating Scale (ALSFRS), adverse events, onset_delta, onset_site, drugs use (riluzole) The PRO-ACT training dataset contains clinical and lab test information of 8,635 patients. Information of 2,424 study subjects with valid gold standard ALSFRS slopes used for processing, modeling and analysis	The time points for all longitudinally varying data elements are aggregated into signature vectors. This facilitates the modeling and prediction of ALSFRS slope changes over the first three months (baseline to month 3)

Tang, et al. (2018), in review

Case-Studies – ALS

- Detect, track, and prognosticate the progression of ALS
- Predict adverse events based on subject phenotype and 0-3 month clinical assessment changes

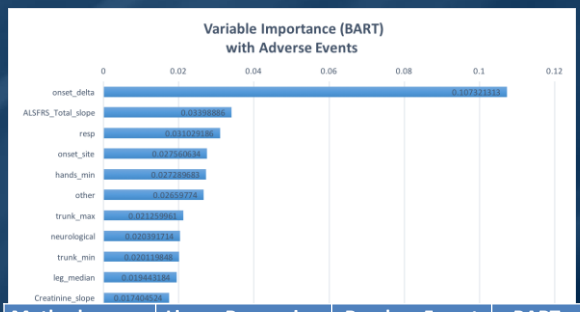

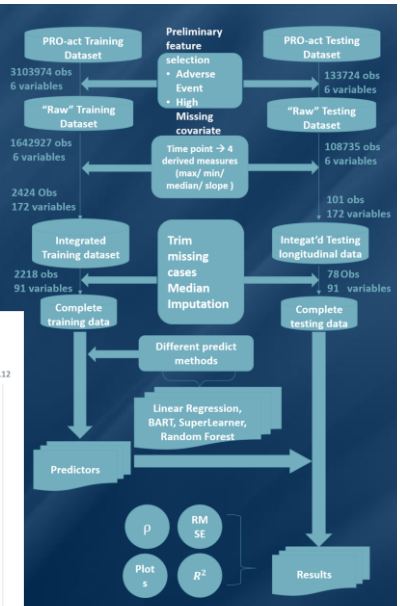

Methods	Linear Regression	Random Forest	BART	SuperLearner
R-squared	0.081	0.174	**0.225**	0.178
RMSE	0.619	0.587	**0.568**	0.585
Correlation	0.298	0.434	**0.485**	0.447

Case-Studies – ALS

- **Main Finding**: predicting univariate clinical outcomes may be challenging, the (information energy) signal is very weak. We can cluster ALS patients and generate evidence-based ALS hypotheses about the complex interactions of *multivariate factors*
- **Classification vs. Clustering**:
  - Classifying univariate clinical outcomes using the PRO-ACT data yields only marginal accuracy (about 70%).
  - Unsupervised clustering into sub-groups generates stable, reliable and consistent computable phenotypes whose explication requires interpretation of multivariate sets of features

Data Representation Fusion Harmonization Aggregation → Cleaning Imputation Wrangling Synthesis → Model-based, Model-free, Classification, Clustering, Inference

Cluster	Consistency	Variance	Cluster-Size	Silhouette
1	1	0	565	0.58
2	0.986	0.018	427	0.63
3	0.956	0.053	699	0.5
4	0.985	0.018	733	0.5

Tang, et al. (2018), in review

Case-Studies – ALS – Explicating Clustering

Feature Name	Between Cluster Significant Differences					
	1-2	1-3	1-4	2-3	2-4	3-4
...			...			
onset_delta.x	1	1	1	1	1	1
...			...			
Q9_Climbing_Stairs_slope	1			1		
...			...			
leg_max		1	1	1	1	
...			...			

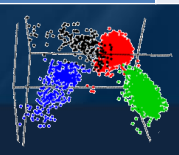

Tang, et al. (2018), in review

Case-Studies – Parkinson's Disease (1)

- **Investigate falls in PD patients** using clinical, demographic and neuroimaging data from two independent initiatives (UMich & Tel Aviv U)
- Applied **controlled feature selection** to identify the most salient predictors of patient falls (gait speed, Hoehn and Yahr stage, postural instability and gait difficulty-related measurements)
- **Model-based** (e.g., GLM) and **model-free** (RF, SVM, Xgboost) analytical methods used to forecasts clinical outcomes (e.g., falls)
- Internal statistical cross **validation** + external out-of-bag validation
- Four specific **challenges**
  - Challenge 1, harmonize & aggregate complex, multisource, multisite PD data
  - Challenge 2, identify salient predictive features associated with specific clinical traits, e.g., patient falls
  - Challenge 3, forecast patient falls and evaluate the classification performance
  - Challenge 4, predict tremor dominance (TD) vs. posture instability and gait difficulty (PIGD).
- **Results**: model-free machine learning based techniques provide a more reliable clinical outcome forecasting, e.g., falls in Parkinson's patients, with classification accuracy of about 70-80%.

Gao, et al. SREP (2018), *in press*

Case-Studies – Parkinson's Disease (1)

Falls in PD are extremely difficult to predict …

Case-Studies – Parkinson's Disease (1)

Method	acc	sens	spec	ppv	npv	lor	auc
Logistic Regression	0.728	0.537	0.855	0.710	0.736	1.920	0.774
Random Forests	**0.796**	**0.683**	**0.871**	**0.778**	**0.806**	**2.677**	**0.821**
AdaBoost	0.689	0.610	0.742	0.610	0.742	1.502	0.793
XGBoost	0.699	0.707	0.694	0.604	0.782	1.699	0.787
SVM	0.709	0.561	0.806	0.657	0.735	1.672	0.822
Neural Network	0.699	0.610	0.758	0.625	0.746	1.588	
Super Learner	0.738	0.683	0.774	0.667	0.787	1.999	

Results of binary fall/no-fall classification (5-fold CV) using top 10 selected features (gaitSpeed_Off, ABC, BMI, PIGD_score, X2.11, partII_sum, Attention, DGI, FOG_Q, H_and_Y_OFF)

Gao, et al. SREP (2018), *in press*

Open-Science & Collaborative Validation

End-to-end Big Data analytic protocol jointly processing complex imaging, genetics, clinical, demo data for assessing PD risk

- Methods for rebalancing of imbalanced cohorts
- ML classification methods generating consistent and powerful phenotypic predictions
- Reproducible protocols for extraction of derived neuroimaging and genomics biomarkers for diagnostic forecasting

https://github.com/SOCR/PBDA

Case-Studies – General Populations

2	20005	Ongoing characteristics	Email access
2	110007	Ongoing characteristics	Newsletter communications, date sent
100	25780	Brain MRI	Acquisition protocol phase.
100	12139	Brain MRI	Believed safe to perform brain MRI scan
100	12188	Brain MRI	Brain MRI measurement completed
100	12187	Brain MRI	Brain MRI measuring method
100	12663	Brain MRI	Reason believed unsafe to perform brain MRI
100	12704	Brain MRI	Reason brain MRI not completed
100	12652	Brain MRI	Reason brain MRI not performed
101	12292	Carotid ultrasound	Carotid ultrasound measurement completed
101	12291	Carotid ultrasound	Carotid ultrasound measuring method
101	20235	Carotid ultrasound	Carotid ultrasound results package
101	22672	Carotid ultrasound	Maximum carotid IMT (intima-medial thickness) at 120 degrees
101	22675	Carotid ultrasound	Maximum carotid IMT (intima-medial thickness) at 150 degrees
101	22678	Carotid ultrasound	Maximum carotid IMT (intima- degrees
101	22681	Carotid ultrasound	Maximum carotid IMT (intima- degrees
101	22671	Carotid ultrasound	Mean carotid IMT (intima-mec
101	22674	Carotid ultrasound	Mean carotid IMT (intima-mec
101	22677	Carotid ultrasound	Mean carotid IMT (intima-mec
101	22680	Carotid ultrasound	Mean carotid IMT (intima-mec
101	22670	Carotid ultrasound	Minimum carotid IMT (intima- degrees
101	22673	Carotid ultrasound	Minimum carotid IMT (intima- degrees
101	22676	Carotid ultrasound	Minimum carotid IMT (intima-medial thickness) at 210 degrees
101	22679	Carotid ultrasound	Minimum carotid IMT (intima-medial thickness) at 240 degrees
101	22682	Carotid ultrasound	Quality control indicator for IMT at 120 degrees
101	22683	Carotid ultrasound	Quality control indicator for IMT at 150 degrees
101	22684	Carotid ultrasound	Quality control indicator for IMT at 210 degrees

- UK Biobank – discriminate between HC, single and multiple comorbid conditions
- Predict likelihoods of various developmental or aging disorders
- Forecast cancer

Data Source	Sample Size/Data Type	Summary
UK Biobank	**Demographics**: > 500K cases **Clinical data:** > 4K features **Imaging data**: T1, resting-state fMRI, task fMRI, T2_FLAIR, dMRI, SWI **Genetics data**	The longitudinal archive of the UK population (NHS)

http://www.ukbiobank.ac.uk
http://bd2k.org

Case-Studies – UK Biobank (Complexities)

Missing Clinical & Phenotypic data for 10K subjects with sMRI, for which we computed 1,500 derived neuroimaging biomarkers.

Including only features observed >30% ($9{,}914 \times 1{,}475$)

Zhou, et al. (2018), pending

Case-Studies – UK Biobank – NI Biomarkers

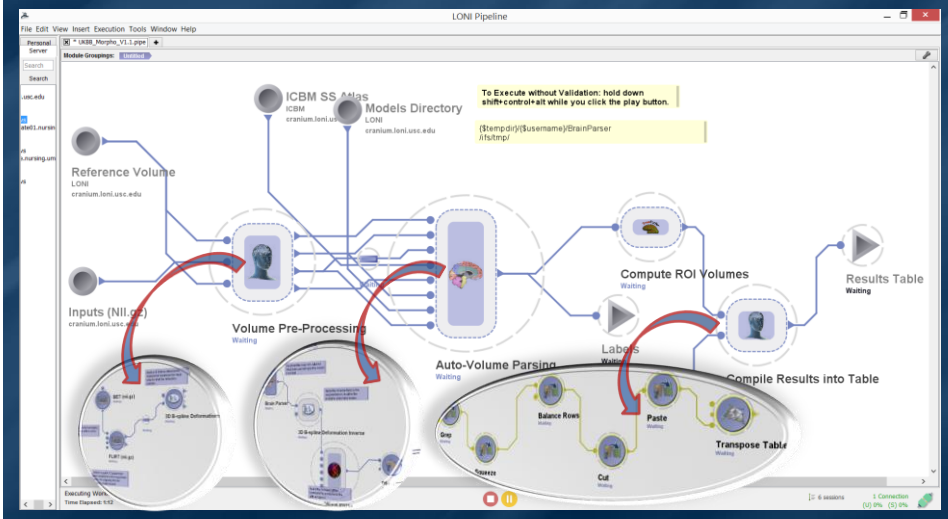

Case-Studies – UK Biobank – Successes/Failures

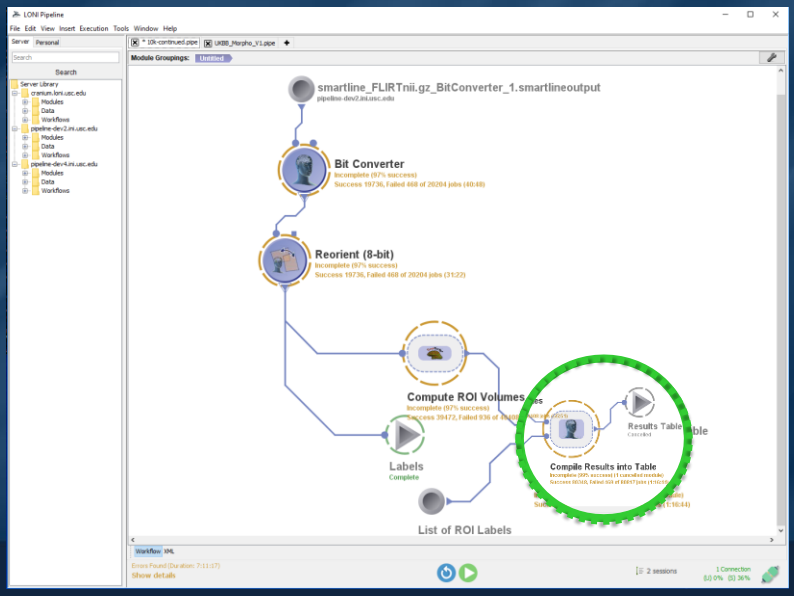

Case-Studies – UK Biobank – Results

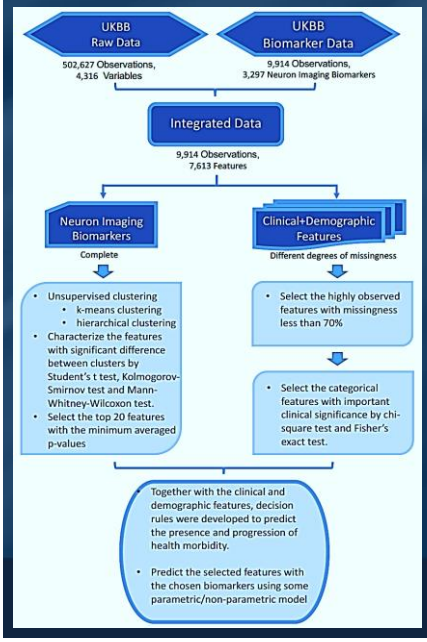

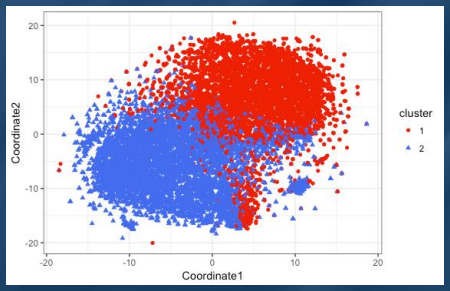

t-SNE plot of the brain neuroimaging biomarkers

k-means clustering			
		Cluster 1	Cluster 2
Hierarchical clustering	Cluster 1	3768 (38.0%)	528 (5.3%)
	Cluster 2	827 (8.3%)	4791 (48.3%)

Cluster	Consistency	Variance	Cluster-size	Silhouette
1	0.997	0.001	5344	0.09
2	0.934	0.001	4570	0.05

Case-Studies – UK Biobank – Results

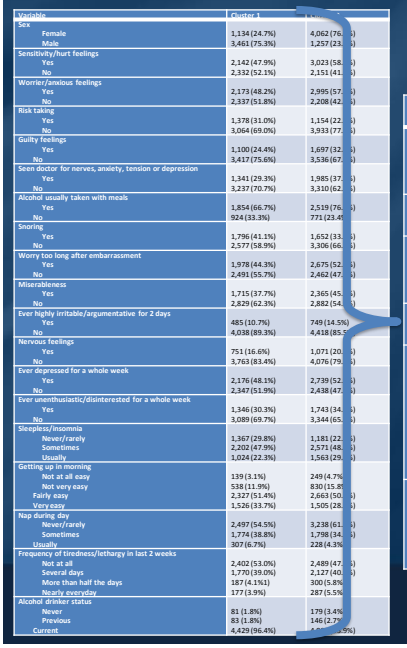

Variable	Cluster 1	Cluster 2
Sex		
Female	1,134 (24.7%)	4,062 (76.4%)
Male	3,461 (75.3%)	1,257 (23.6%)
. . .	. . .	
Nervous feelings		
Yes	751 (16.6%)	1,071 (20.8%)
No	3,763 (83.4%)	4,076 (79.2%)
. . .	. . .	
Frequency of tiredness/lethargy in last 2 weeks		
Not at all	2,402 (53.0%)	2,489 (47.8%)
Several days	1,770 (39.0%)	2,127 (40.9%)
More than half the days	187 (4.1%1)	300 (5.8%)
Nearly everyday	177 (3.9%)	287 (5.5%)
Alcohol drinker status		
Never	81 (1.8%)	179 (3.4%)
Previous	83 (1.8%)	146 (2.7%)
Current	4,429 (96.4%)	4,992 (93.9%)

Case-Studies – UK Biobank – Results

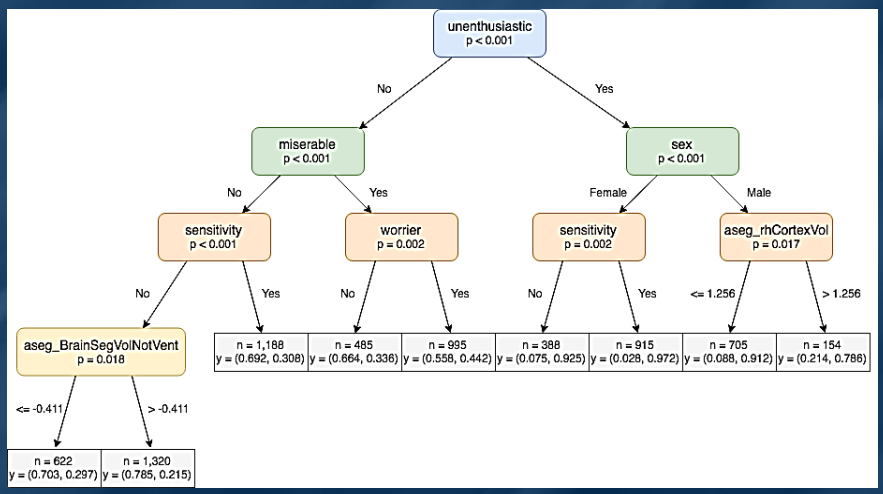

Decision tree illustrating a simple clinical decision support system providing machine guidance for identifying **depression feelings** based on categorical variables and neuroimaging biomarkers. In each terminal node, the y vector includes the percentage of subjects being labeled as "no" and "yes", in this case, answering the question "Ever depressed for a whole week." The p-values listed at branching nodes indicate the significance of the corresponding splitting criterion.

Case-Studies – UK Biobank – Results

	Accuracy	95% CI (Accuracy)	Sensitivity	Specificity
Sensitivity/hurt feelings	0.700	(0.676, 0.724)	0.657	0.740
Ever depressed for a whole week	0.782	(0.760, 0.803)	0.938	0.618
Worrier/anxious feelings	0.730	(0.706, 0.753)	0.721	0.739
Miserableness	0.739	(0.715, 0.762)	0.863	0.548

Cross-validated (random forest) prediction results for four types of mental disorders

End-to-end Pipeline Workflow Solutions

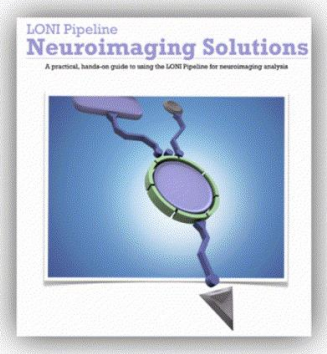

Dinov, *et al*., 2014, Front. Neuroinform.;

Dinov, *et al*., 2013, Brain Imaging & Behavior

Compressive Big Data Analytics (CBDA)

- Foundation for Compressive Big Data Analytics (CBDA)
  - Iteratively generate random (sub)samples from the Big Data collection
  - Then, using classical techniques to obtain model-based or non-parametric inference based on the sample
  - Next, compute likelihood estimates (e.g., probability values quantifying effects, relations, sizes)
  - Repeat – the process continues iteratively until a criterion is met – the (re)sampling and inference steps many times (with or without using the results of previous iterations as priors for subsequent steps)

Dinov, *2016, PMID: 26998309;* *Marino, et al., (pending)*

Synergies with Compressive Sensing

- Define the nested sets

$$S_k = \{x: \|x\|_o \stackrel{\text{def}}{=} |supp(x)| \leq k\},$$

where the data x, as a vector or tensor, has at most k non-trivial elements. Note that if $x, z \in S_k$, then $x + z \in S_{2k} \supseteq S_k$

- If $\Phi_{n\times n} = (\varphi_1, \varphi_2, \varphi_3, \ldots, \varphi_n)$ represents an orthonormal basis, the data may be expressed as $x = \Phi c$, where $c_i = \langle x, \varphi_i \rangle$, i.e., $c = \Phi^T x$, and $\|c\|_o \leq k$. Even if x is not strictly sparse, its representation c may be sparse. For each dataset, we can assess and quantify the error of approximating x by an optimal estimate $\hat{x} \in S_k$ by computing

$$\sigma_k(x)_p = \min_{\hat{x}\in S_k} \|x - \hat{x}\|_p$$

Synergies with Compressive Sensing

- In compressive sensing, if $x \in R^n$, and we have a data stream generating m linear measurements, we can represent $y = Ax$, where $A_{m \times n}$ is a dimensionality reducing matrix ($m \ll n$), i.e.,

$$A_{m \times n}: R^n \rightarrow R^m$$

- The null space of A,

  $N(A) = \{z \in R^n : Az = 0 \in R^m\}$.

  A uniquely represents all $x \in S_k \Leftrightarrow N(A)$ contains no vectors in S_{2k}.

- The spark of a matrix A represents the smallest number of columns of A that are linearly dependent. If $A_{m \times n}$ is a random matrix whose entries are independent and identically distributed, then $spark(A) = m + 1$, with probability 1.

Synergies with Compressive Sensing

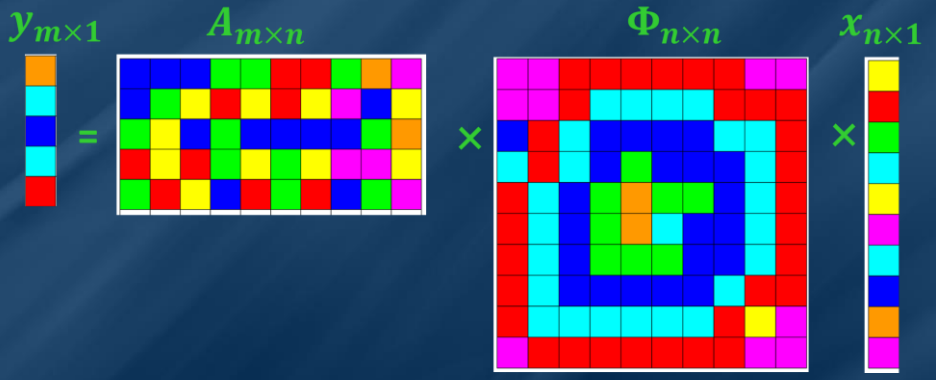

Synergies with Compressive Sensing

- If the entries of A are chosen according to a sub-Gaussian distribution, then with high probability, for each k, there exists $\delta_{2k} \in (0,1)$ such that

$$(1-\delta_{2k})\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1+\delta_{2k})\|x\|_2^2 \quad (1)$$

for all $x \in S_{2k}$ (RIP=Restricted isometry property)

- When we know that the original signal is sparse, to reconstruct x given the observed measurements y, we can solve the optimization problem:

$$\hat{x} = \arg\min_{z:Az=y} \|z\|_0$$

Synergies with Compressive Sensing

- Linear programming may be used to solve the optimization problem if we replace the zero-norm by its more tractable convex approximation, the l_1-norm, $\hat{x} = \arg\min_{z:Az=y} \|z\|_1$
- Given that $A_{m\times n}$ has the above property and $\delta_{2k} < \sqrt{2}-1$, if we observe $y = Ax$, then the solution $\hat{x}$ satisfies

$$\|\hat{x}-x\|_2 \leq C_0 \frac{\sigma_k(x)_1}{\sqrt{k}}$$

- Thus, in compressive sensing applications, if $x \in S_k$ and A satisfies the RIP, condition (1), we can recover any k-sparse signal x exactly (as $\sigma_k(x)_1 = 0$) using only $O(k\log(n/k))$ observations, since $m = O\left(\frac{k\log(n/k)}{\delta_{2k}^2}\right)$
- Finally, if $A_{m\times n}$ is random (e.g., chosen according to a Gaussian distribution) and $\Phi_{n\times n}$ is an orthonormal basis, then $A_{m\times n} \times \Phi_{n\times n}$ will also have a Gaussian distribution, and if m is large, $A' = A \times \Phi$ will also satisfy the RIP condition (1) with high probability.

Compressive Big Data Analytics (CBDA)

- Start with a generic **dataset**: $[X, Y, X_{val}, Y_{val}]$, let $C^j = [X^j, Y^j, X^j_{val}], j = 1, \ldots, M$.
- Define machine learning **predictor** as ML: $ML(C^j): R^{n_j \times k_j} \times R^{n_j} \times R^{m \times k_j} \to R^m$, $ML([X^j, Y^j, X^j_{val}]) = \hat{Y^j}$.
- Define **performance metric** as: $\tau(ML(C^j), Y_{val}): R^m \times R^m \to R$, $\tau(\hat{Y^j}, Y_{val}) = c_j, j = 1,2, \ldots, M$.
- **Rank the samples**: $\{c_{(j)}\} = O^q(\{c_j\}), j = 1,2, \ldots, q$.
- Set the **feature values**: $s_j^{(i)} = Dirichlet_j(f_i) = \begin{cases} 0, & f_i \notin sample\ j \\ 1, & f_i \in sample\ j \end{cases}$.
- Set a **metric** as: $F \in R^{q \times K}, \forall b_{ji} \in F, b_{ji} = s_j^{(i)}$.
- Count the **occurrence**: $S_i = \sum_j^q b_{ji}, i = 1,2, \ldots, K$.
- **Feature mining**: $S_{(i)} = O^K(S_i), i = 1,2, \ldots, K, \Omega^* = \{f_1, \ldots, f_i, \ldots, f_K\}$.
- **Inference**: let $C^p = [\phi_p X, \phi_p Y, \phi_p X_{val}]$ and $[ML_p(C^p): R^{n \times k_{p^*}} \times R^{n \times 1} \times R^{m \times k_{p^*}} \to R^{m \times 1}$, $ML([\Phi_p X, \Phi_p Y, \Phi_p X_{val}]) = Y^p_{val}, p = 5,10, \ldots, \hat{p}, \tau(Y^p_{val}, Y_{val}): R^m \times R^m \to R$, $\tau(ML([\Phi_{p^*} X, \Phi_{p^*} Y, \Phi_{p^*} X_{val}]), Y_{val}) = \underset{p}{best}\, \tau(ML([\Phi_p X, \Phi_p Y, \Phi_p X_{val}]), Y_{val})$.
- Then Φ_{p^*} is the final **dictionary** we need.
- **Performance assessment**: Once we obtain ϕ_{p^*}, and have X_{exp} which is under prediction. Then perform ϕ_{p^*} and SuperLearner algorithm: $Y_{exp} = ML([\Phi_{p^*} X, \Phi_{p^*} Y, \Phi_{p^*} X_{exp}])$, where Y_{exp} represents the expected results.

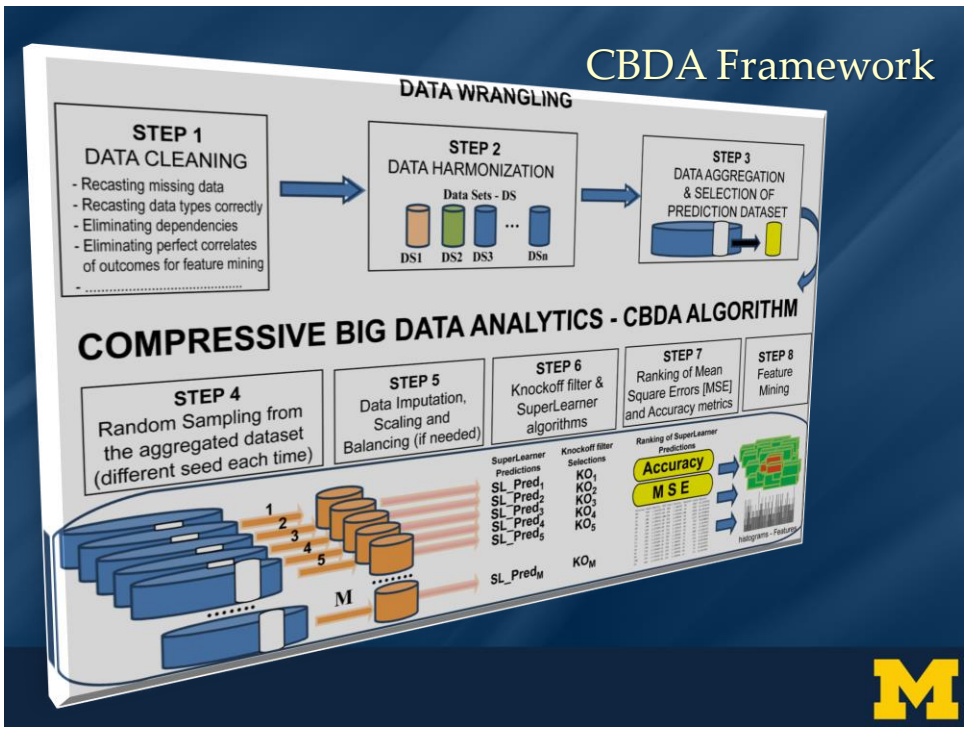

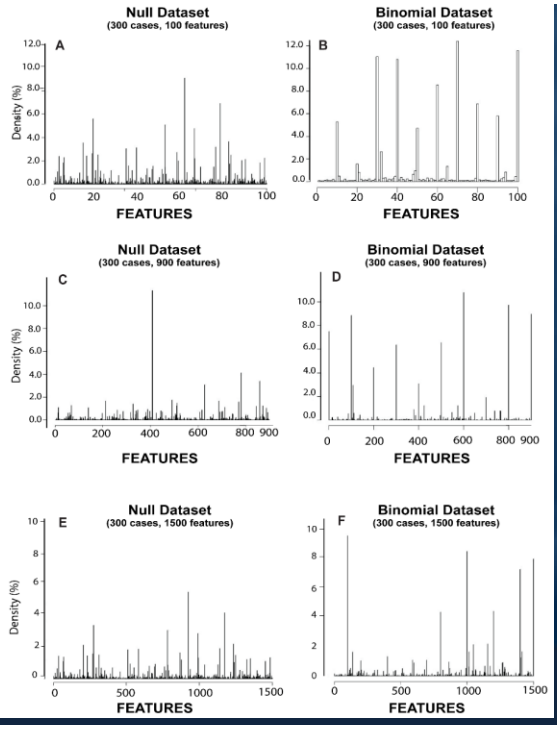

CBDA Results: Simulated Data

Knockoff of Null (left) vs Binomial (right) Data

Panels A, C and E show the correspondent histograms generated from the Knockoff Filter algorithm on the three Null datasets.

Panels B, D and F show the correspondent histograms generated from the Knockoff Filter algorithm on the three Binomial datasets.

Performance metric: MSE

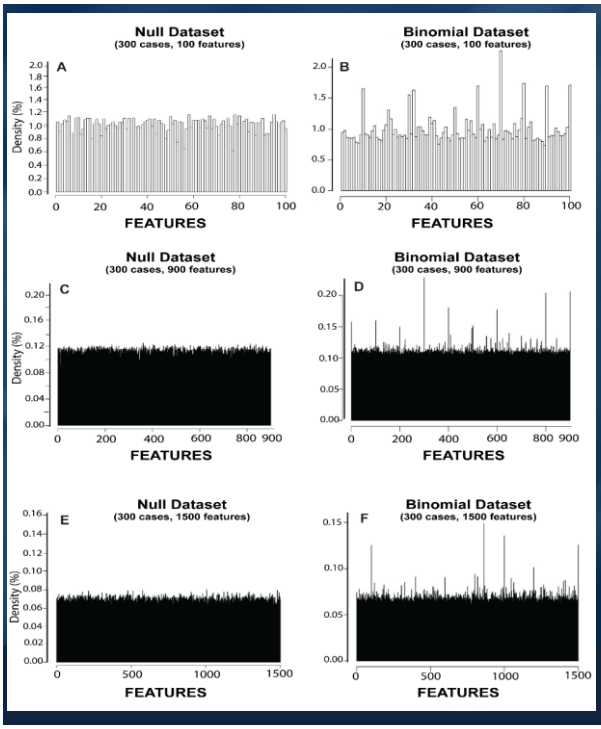

CBDA Results: Simulated Data

CBDA Results: Null (left) vs Binomial (right) Data

Panels A, C and E show the correspondent histograms of the CBDA Results on the three Null datasets.

Panels B, D and F show the correspondent histograms of the CBDA results on the three Binomial datasets.

Performance metric: MSE

CBDA Results: Biomed Data (ADNI)

	Reference		
Prediction	AD	MCI	Normal
AD	69	17	1
MCI	12	243	8
Normal	0	9	140
Overall Statistics			
Accuracy	0.9058 [95% CI = (0.8767, 0.93)]		
No Information Rate	0.5391		
P-Value [Acc > NIR]	<2e-16		
Kappa	0.8426		
McNemar's Test P-Value	0.589		
Statistics by Diagnostic Class			
	AD	MCI	Normal
Sensitivity	0.8519	0.9033	0.9396
Specificity	0.9569	0.9130	0.9743
Positive Pred Value	0.7931	0.9240	0.9396
Negative Pred Value	0.9709	0.8898	0.9743
Prevalence	0.1623	0.5391	0.2986
Detection Rate	0.1383	0.4870	0.2806
Detection Prevalence	0.1743	0.5271	0.2986
Balanced Accuracy	0.9044	0.9082	0.9569

CBDA multinomial classification results (ADNI)

Acknowledgments

Funding

NIH: P20 NR015331, U54 EB020406, P50 NS091856, P30 DK089503, P30AG053760, UL1TR002240

NSF: 1734853, 1636840, 1416953, 0716055, 1023115

The Elsie Andresen Fiske Research Fund

http://SOCR.umich.edu

Collaborators

- **SOCR**: Alexandr Kalinin, Selvam Palanimalai, Syed Husain, Matt Leventhal, Ashwini Khare, Rami Elkest, Abhishek Chowdhury, Patrick Tan, Gary Chan, Andy Foglia, Pratyush Pati, Brian Zhang, Juana Sanchez, Dennis Pearl, Kyle Siegrist, Rob Gould, Jingshu Xu, Nellie Ponarul, Ming Tang, Asiyah Lin, Nicolas Christou, Hanbo Sun, Tuo Wang
- **LONI/INI**: Arthur Toga, Roger Woods, Jack Van Horn, Zhuowen Tu, Yonggang Shi, David Shattuck, Elizabeth Sowell, Katherine Narr, Anand Joshi, Shantanu Joshi, Paul Thompson, Luminita Vese, Stan Osher, Stefano Soatto, Seok Moon, Junning Li, Young Sung, Carl Kesselman, Fabio Macciardi, Federica Torri
- **UMich MIDAS/MNORC/AD/PD Centers**: Cathie Spino, Chuck Burant, Ben Hampstead, Stephen Goutman, Stephen Strobbe, Hiroko Dodge, Hank Paulson, Bill Dauer, Brian Athey

SOCR STATISTICS