# Big Data is Low Rank 

Madeleine Udell<br>Operations Research and Information Engineering Cornell University

UC Davis, 11/5/2018

## Data table

| age | gender | state | income | education | $\cdots$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 29 | F | CT | $\$ 53,000$ | college | $\cdots$ |
| 57 | $?$ | NY | $\$ 19,000$ | high school | $\cdots$ |
| $?$ | M | CA | $\$ 102,000$ | masters | $\cdots$ |
| 41 | F | NV | $\$ 23,000$ | $?$ | $\cdots$ |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ |  |  |

- detect demographic groups?
- find typical responses?
- identify related features?
- impute missing entries?


## Data table

$m$ examples (patients, respondents, households, assets)
$n$ features (tests, questions, sensors, times)

$$
\left[\begin{array}{l} 
\\
A
\end{array}\right]=\left[\begin{array}{ccc}
A_{11} & \cdots & A_{1 n} \\
\vdots & \ddots & \vdots \\
A_{m 1} & \cdots & A_{m n}
\end{array}\right]
$$

- ith row of $A$ is feature vector for ith example
- jth column of $A$ gives values for $j$ th feature across all examples


## Low rank model

given: $m \times n$ data table $A, k \ll m, n$ find: $X \in \mathbf{R}^{m \times k}, Y \in \mathbf{R}^{k \times n}$ for which

$$
[X]\left[\begin{array}{lll}
{[ } & Y & ] \approx\left[\begin{array}{ll} 
& A
\end{array}\right]
\end{array}\right.
$$

i.e., $x_{i} y_{j} \approx A_{i j}$, where

$$
[x]=\left[\begin{array}{c}
-x_{1}- \\
\vdots \\
-x_{m}-
\end{array}\right] \quad[\quad Y \quad]=\left[\begin{array}{ccc}
\mid & & \mid \\
y_{1} & \cdots & y_{n} \\
\mid & & \mid
\end{array}\right]
$$

## interpretation:

- $X$ and $Y$ are (compressed) representation of $A$
- $x_{i}^{T} \in \mathbf{R}^{k}$ is a point associated with example $i$
- $y_{j} \in \mathbf{R}^{k}$ is a point associated with feature $j$
- inner product $x_{i} y_{j}$ approximates $A_{i j}$


## Why?

- reduce storage; speed transmission
- understand (visualize, cluster)
- remove noise
- infer missing data
- simplify data processing


## Outline

## PCA

## Generalized low rank models

Applications
Impute missing data
Dimensionality reduction
Causal inference
Automatic machine learning

Why low rank?

## Principal components analysis

PCA: for $A \in \mathbf{R}^{m \times n}$,

$$
\operatorname{minimize}\|A-X Y\|_{F}^{2}=\sum_{i=1}^{m} \sum_{j=1}^{n}\left(A_{i j}-x_{i} y_{j}\right)^{2}
$$

with variables $X \in \mathbf{R}^{m \times k}, Y \in \mathbf{R}^{k \times n}$

- old roots [Pearson 1901, Hotelling 1933]
- least squares low rank fitting
- (analytical) solution via SVD of $A=U \Sigma V^{T}$
- (numerical) solution via alternating minimization


## Outline

PCA

Generalized low rank models

## Applications

Impute missing data
Dimensionality reduction
Causal inference
Automatic machine learning

Why low rank?

## Generalized low rank model

$$
\operatorname{minimize} \sum_{(i, j) \in \Omega} L_{j}\left(x_{i} y_{j}, A_{i j}\right)+\sum_{i=1}^{m} r_{i}\left(x_{i}\right)+\sum_{j=1}^{n} \tilde{r}_{j}\left(y_{j}\right)
$$

- loss functions $L_{j}$ for each column
- e.g., different losses for reals, booleans, categoricals, ordinals, ...
- regularizers $r: \mathbf{R}^{1 \times k} \rightarrow \mathbf{R}, \tilde{r}: \mathbf{R}^{k} \rightarrow \mathbf{R}$
- observe only $(i, j) \in \Omega$ (other entries are missing)

Note: can be NP-hard to optimize exactly...

## Matrix completion

observe $A_{i j}$ only for $(i, j) \in \Omega \subset\{1, \ldots, m\} \times\{1, \ldots, n\}$ $\operatorname{minimize} \quad \sum_{(i, j) \in \Omega}\left(A_{i j}-x_{i} y_{j}\right)^{2}+\lambda \sum_{i=1}^{m}\left\|x_{i}\right\|^{2}+\lambda \sum_{j=1}^{n}\left\|y_{j}\right\|^{2}$
two regimes:

- some entries missing: don't waste data; "borrow strength" from entries that are not missing
- most entries missing: matrix completion still works!

Theorem ([Keshavan Montanari 2010])
If $A$ has rank $k^{\prime} \leq k$ and $|\Omega|=O\left(n k^{\prime} \log n\right)$ (and $A$ is incoherent and $\Omega$ is chosen UAR), then matrix completion exactly recovers the matrix $A$ with high probability.

## Maximum likelihood low rank estimation

Choose loss function to maximize (log) likelihood of observations:

- gaussian noise: $L(u, a)=(u-a)^{2}$
- laplacian (heavy-tailed) noise: $L(u, a)=|u-a|$
- gaussian + laplacian noise: $L(u, a)=\operatorname{huber}(u-a)$
- poisson (count) noise: $L(u, a)=\exp (u)-a u+a \log a-a$
- bernoulli (coin toss) noise: $L(u, a)=\log (1+\exp (-a u))$


## Maximum likelihood low rank estimation works

## Theorem (Template)

If a number of samples $|\Omega|=O(n \log (n))$ drawn UAR from matrix entries is observed according to a probabilistic model with parameter $Z$, the solution to (appropriately) regularized maximum likelihood estimation is close to the true $Z$ with high probability.
examples (not exhaustive!):

- additive gaussian noise [Candes Plan 2009]
- additive subgaussian noise [Keshavan Montanari Oh 2009]
- gaussian + laplacian noise [Xu Caramanis Sanghavi 2012]
- 0-1 (Bernoulli) observations [Davenport et al. 2012]
- entrywise exponential family distribution [Gunasekar Ravikumar Ghosh 2014]
- multinomial logit [Kallus Udell 2016]


## Losses

$$
\operatorname{minimize} \quad \sum_{(i, j) \in \Omega} L_{j}\left(x_{i} y_{j}, A_{i j}\right)+\sum_{i=1}^{m} r_{i}\left(x_{i}\right)+\sum_{j=1}^{n} \tilde{r}_{j}\left(y_{j}\right)
$$ choose loss $L: \mathbf{R} \times \mathcal{F} \rightarrow \mathbf{R}$ adapted to data type $\mathcal{F}$ :

| data type | loss | $L(u, a)$ |
| :---: | :---: | :---: |
| real | quadratic | $(u-a)^{2}$ |
| real | absolute value | $\|u-a\|$ |
| real | huber | huber ( $u-a)$ |
| boolean | hinge | $(1-u a)_{+}$ |
| boolean | logistic | $\log (1+\exp (-a u))$ |
| integer | poisson | $\exp (u)-a u+a \log a-a$ |
| ordinal | ordinal hinge | $\begin{aligned} & \sum_{a^{\prime}=1}^{a-1}\left(1-u+a^{\prime}\right)_{+}+ \\ & \sum_{a^{\prime}=a+1}^{d}\left(1+u-a^{\prime}\right)_{+} \end{aligned}$ |
| categorical categorical | one-vs-all multinomial logit | $\begin{gathered} \left(1-u_{a}\right)_{+}+\sum_{a^{\prime} \neq a}\left(1+u_{a^{\prime}}\right)_{+} \\ \frac{\exp \left(u_{a}\right)}{\sum_{a^{\prime}=1}^{d} \exp \left(u_{a^{\prime}}\right)} \end{gathered}$ |

## Regularizers

$$
\operatorname{minimize} \sum_{(i, j) \in \Omega} L_{j}\left(x_{i} y_{j}, A_{i j}\right)+\sum_{i=1}^{m} r_{i}\left(x_{i}\right)+\sum_{j=1}^{n} \tilde{r}_{j}\left(y_{j}\right)
$$

choose regularizers $r, \tilde{r}$ to impose structure:

| structure | $r(x)$ | $\tilde{r}(y)$ |
| :--- | :---: | :---: |
| small | $\\|x\\|_{2}^{2}$ | $\\|y\\|_{2}^{2}$ |
| sparse | $\\|x\\|_{1}$ | $\\|y\\|_{1}$ |
| nonnegative | $\mathbf{1}(x \geq 0)$ | $\mathbf{1}(y \geq 0)$ |
| clustered | $\mathbf{1}(\operatorname{card}(x)=1)$ | 0 |

## Outline

PCA
Generalized low rank models
Applications
Impute missing data
Dimensionality reduction
Causal inference
Automatic machine learning

Why low rank?

## Impute missing data

impute most likely true data $\hat{A}_{i j}$

$$
\hat{A}_{i j}=\underset{a}{\operatorname{argmin}} L_{j}\left(x_{i} y_{j}, a\right)
$$

- implicit constraint: $\hat{A}_{i j} \in \mathcal{F}_{j}$
- when $L_{j}$ is quadratic, $\ell_{1}$, or Huber loss, then $\hat{A}_{i j}=x_{i} y_{j}$
- if $\mathcal{F} \neq \mathbf{R}, \operatorname{argmin}_{a} L_{j}\left(x_{i} y_{j}, a\right) \neq x_{i} y_{j}$
- e.g., for hinge loss $L(u, a)=(1-u a)_{+}, \hat{A}_{i j}=\boldsymbol{\operatorname { s i g n }}\left(x_{i} y_{j}\right)$


## Impute heterogeneous data



## Impute heterogeneous data


qpca rank 10 recovery



## Impute heterogeneous data


qpca rank 10 recovery



## Hospitalizations are low rank

## GLRM outperforms PCA

hospitalization data set

|  | single hospitialization |  |  |
| :---: | :---: | :---: | :---: |
|  | $\begin{aligned} & \check{u} \\ & \stackrel{0}{0} \\ & \text { ơ } \end{aligned}$ | $\begin{aligned} & \text { u} \\ & \stackrel{y}{0} \\ & \stackrel{0}{U} \\ & 0 \\ & 0 \end{aligned}$ | $n$ 0 0 0 0 0 0 0 0 |

(a) Tailored models outperform PCA
on imputation in the hospitalization dataset

[Schuler Udell et al., 2016]

## American community survey

2013 ACS:

- 3M respondents, 87 economic/demographic survey questions
- income
- cost of utilities (water, gas, electric)
- weeks worked per year
- hours worked per week
- home ownership
- looking for work
- use foodstamps
- education level
- state of residence
- ...
- $1 / 3$ of responses missing


## Fitting a GLRM to the ACS

- construct a rank 10 GLRM with loss functions respecting data types
- huber for real values
- hinge loss for booleans
- ordinal hinge loss for ordinals
- one-vs-all hinge loss for categoricals
- scale losses and regularizers by $1 / \sigma_{j}^{2}$
- fit the GLRM
in 2 lines of code:
glrm, labels = GLRM(A, 10, scale = true)
X,Y = fit! (glrm)


## American community survey

most similar features (in demography space):

- Alaska: Montana, North Dakota
- California: Illinois, cost of water
- Colorado: Oregon, Idaho
- Ohio: Indiana, Michigan
- Pennsylvania: Massachusetts, New Jersey
- Virginia: Maryland, Connecticut
- Hours worked: weeks worked, education


## Low rank models for dimensionality reduction ${ }^{1}$

U.S. Wage \& Hour Division (WHD) compliance actions:

| company | zip | violations | $\cdots$ |
| :---: | :---: | :---: | :---: |
| Holiday Inn | 14850 | 109 | $\cdots$ |
| Moosewood Restaurant | 14850 | 0 | $\cdots$ |
| Cornell Orchards | 14850 | 0 | $\cdots$ |
| Lakeside Nursing Home | 14850 | 53 | $\cdots$ |
| $\vdots$ | $\vdots$ | $\vdots$ |  |

- 208,806 rows (cases) $\times 252$ columns (violation info)
- 32,989 zip codes...

[^0]
## Low rank models for dimensionality reduction

ACS demographic data:

| zip | unemployment | mean income | $\cdots$ |
| :---: | :---: | :---: | :---: |
| 94305 | $12 \%$ | $\$ 47,000$ | $\cdots$ |
| 06511 | $19 \%$ | $\$ 32,000$ | $\cdots$ |
| 60647 | $23 \%$ | $\$ 23,000$ | $\cdots$ |
| 94121 | $4 \%$ | $\$ 178,000$ | $\cdots$ |
| $\vdots$ | $\vdots$ | $\vdots$ |  |

- 32,989 rows (zip codes) $\times 150$ columns (demographic info)
- GLRM embeds zip codes into (low dimensional) demography space


## Low rank models for dimensionality reduction

Zip code features:

Archetype Representation of Zip Code Tabulation Areas


## Low rank models for dimensionality reduction

build 3 sets of features to predict violations:

- categorical: expand zip code to categorical variable
- concatenate: join tables on zip
- GLRM: replace zip code by low dimensional zip code features
fit a supervised (deep learning) model:

| method | train error | test error | runtime |
| :---: | :---: | :---: | :---: |
| categorical | 0.2091690 | 0.2173612 | 23.7600000 |
| concatenate | 0.2258872 | 0.2515906 | 4.4700000 |
| GLRM | 0.1790884 | 0.1933637 | 4.3600000 |

## Causal inference with messy data


[Kallus Mao Udell, 2018]

## Matrix completion reduces error



Method - Lasso $\rightleftharpoons$ Ridge - OLS + MF (Proposal) - Oracle

[Kallus Mao Udell, 2018]

## Low rank method for automatic machine learning



Low-Dimensional Features


- yellow blocks are fully observed: results of each algorithm on each training dataset
- last row is a new dataset
- blue blocks are observed results of algorithms on new dataset
- white blocks are unknown entries
- fit a low rank model to impute white blocks
[Yang Akimoto Udell, 2018]


## Low rank fit correctly identifies best algorithm type

Correct algorithm type

[Yang Akimoto Udell, 2018]

## Experiment design for timely model selection

Which algorithms to use to predict performance?

$$
\begin{array}{ll}
\underset{v_{j}}{\operatorname{maximize}} & \log \operatorname{det}\left(\sum_{j=1}^{n} v_{j} y_{j} y_{j}^{T}\right) \\
\text { subject to } & \sum_{j=1}^{n} v_{j} \hat{t}_{j} \leq \tau \\
& v_{j} \in[0,1] \quad \forall j \in[n] .
\end{array}
$$

- $\hat{t}_{j}$ : estimated runtime of each machine learning model
- $\tau$ : runtime budget
[Yang Akimoto Udell, 2018]


## Oboe: Time-constrained AutoML


(a) OpenML

(b) UCl

- Oboe
-_ auto-sklearn
- random

(c) OpenML

(d) UCl

Figure: In 1a and 1b, shaded area $=75$ th -25 th percentile. $\ln 1 \mathrm{c}$ and 1 d , rank 1 is best and 3 is worst.

## Outline

PCA
Generalized low rank models
Applications
Impute missing data
Dimensionality reduction
Causal inference
Automatic machine learning

Why low rank?

## Latent variable models

Suppose $A \in \mathbf{R}^{m \times n}$ generated by a latent variable model (LVM):

- $\alpha_{i} \sim \mathcal{A}$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \mathcal{B}$ iid, $j=1, \ldots, n$
- $A_{i j}=g\left(\alpha_{i}, \beta_{j}\right)$


## Latent variable models: examples

inner product:

- $\alpha_{i} \sim \mathcal{A} \subseteq \mathbf{R}^{k}$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \mathcal{B} \subseteq \mathbf{R}^{k}$ iid, $j=1, \ldots, n$
- $A_{i j}=\alpha_{i}^{T} \beta_{j}$


## Latent variable models: examples

inner product:

- $\alpha_{i} \sim \mathcal{A} \subseteq \mathbf{R}^{k}$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \mathcal{B} \subseteq \mathbf{R}^{k}$ iid, $j=1, \ldots, n$
- $A_{i j}=\alpha_{i}^{T} \beta_{j}$
- rank of $A$ ?


## Latent variable models: examples

inner product:

- $\alpha_{i} \sim \mathcal{A} \subseteq \mathbf{R}^{k}$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \mathcal{B} \subseteq \mathbf{R}^{k}$ iid, $j=1, \ldots, n$
- $A_{i j}=\alpha_{i}^{T} \beta_{j}$
- rank of $A$ ? $k$


## Latent variable models: examples

inner product:

- $\alpha_{i} \sim \mathcal{A} \subseteq \mathbf{R}^{k}$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \mathcal{B} \subseteq \mathbf{R}^{k}$ iid, $j=1, \ldots, n$
- $A_{i j}=\alpha_{i}^{T} \beta_{j}$
- rank of $A$ ? $k$
univariate:
- $\alpha_{i} \sim \operatorname{Unif}(0,1)$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \operatorname{Unif}(0,1)$ iid, $j=1, \ldots, n$
- $A_{i j}=g\left(\alpha_{i}, \beta_{j}\right)$


## Latent variable models: examples

inner product:

- $\alpha_{i} \sim \mathcal{A} \subseteq \mathbf{R}^{k}$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \mathcal{B} \subseteq \mathbf{R}^{k}$ iid, $j=1, \ldots, n$
- $A_{i j}=\alpha_{i}^{T} \beta_{j}$
- rank of $A$ ? $k$
univariate:
- $\alpha_{i} \sim \operatorname{Unif}(0,1)$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \operatorname{Unif}(0,1)$ iid, $j=1, \ldots, n$
- $A_{i j}=g\left(\alpha_{i}, \beta_{j}\right)$
- rank of $A$ ?


## Latent variable models: examples

inner product:

- $\alpha_{i} \sim \mathcal{A} \subseteq \mathbf{R}^{k}$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \mathcal{B} \subseteq \mathbf{R}^{k}$ iid, $j=1, \ldots, n$
- $A_{i j}=\alpha_{i}^{T} \beta_{j}$
- rank of $A$ ? $k$
univariate:
- $\alpha_{i} \sim \operatorname{Unif}(0,1)$ iid, $i=1, \ldots, m$
- $\beta_{j} \sim \operatorname{Unif}(0,1)$ iid, $j=1, \ldots, n$
- $A_{i j}=g\left(\alpha_{i}, \beta_{j}\right)$
- rank of $A$ ?
- can be large!
- if $g$ is analytic with $\sup _{x \in \mathbf{R}}\left|g^{(d)}(x)\right| \leq M$, entrywise $\epsilon$-approximation to $A$ has rank $\mathcal{O}(\log (1 / \epsilon))$


## Nice latent variable models

We say LVM is nice if

- distributions $\mathcal{A}$ and $\mathcal{B}$ have bounded support
- $g$ is piecewise analytic and on each piece: for some $M \in \mathbf{R}$,

$$
\left\|D^{\mu} g(\alpha, \beta)\right\| \leq C M^{|\mu|}\|g\|
$$

$\left(\|g\|=\sup _{x \in \operatorname{dom} g} g(x)\right.$ is sup norm. $)$
Examples: $g(\alpha, \beta)=\operatorname{poly}(\alpha, \beta)$ or $g(\alpha, \beta)=\exp (\operatorname{poly}(\alpha, \beta))$

## Rank of nice latent variable models?

Question: Suppose $A \in \mathbf{R}^{m \times n}$ is drawn from a nice LVM. How does rank of $\epsilon$-approximation to $A$ change with $m$ and $n$ ?

$$
\begin{array}{ll}
\operatorname{minimize} & \operatorname{Rank}(X) \\
\text { subject to } & \|X-A\|_{\infty} \leq \epsilon
\end{array}
$$

## Rank of nice latent variable models?

Question: Suppose $A \in \mathbf{R}^{m \times n}$ is drawn from a nice LVM. How does rank of $\epsilon$-approximation to $A$ change with $m$ and $n$ ?

$$
\begin{array}{ll}
\operatorname{minimize} & \operatorname{Rank}(X) \\
\text { subject to } & \|X-A\|_{\infty} \leq \epsilon
\end{array}
$$

Answer: rank grows as $\mathcal{O}\left(\log (m+n) / \epsilon^{2}\right)$

## Rank of nice latent variable models?

Question: Suppose $A \in \mathbf{R}^{m \times n}$ is drawn from a nice LVM. How does rank of $\epsilon$-approximation to $A$ change with $m$ and $n$ ?

$$
\begin{array}{ll}
\operatorname{minimize} & \operatorname{Rank}(X) \\
\text { subject to } & \|X-A\|_{\infty} \leq \epsilon
\end{array}
$$

Answer: rank grows as $\mathcal{O}\left(\log (m+n) / \epsilon^{2}\right)$

Theorem (Udell and Townsend, 2017)
Nice latent variable models are of log rank.

## Proof sketch

- For each $\alpha$, expand $g$ around $\beta=0$ by its Taylor series

$$
\begin{aligned}
g(\alpha, \beta)-g(\alpha, 0) & =\langle\nabla g(\alpha, 0), \beta\rangle+\left\langle\nabla^{2} g(\alpha, 0), \beta \beta^{\top}\right\rangle+\ldots \\
& =\left[\begin{array}{c}
\nabla g(\alpha, 0) \\
\operatorname{vec}\left(\nabla^{2} g(\alpha, 0)\right) \\
\vdots
\end{array}\right]^{\top}\left[\begin{array}{c}
\beta \\
\operatorname{vec}\left(\beta \beta^{\top}\right) \\
\vdots
\end{array}\right]
\end{aligned}
$$

collect terms depending on $\alpha$ and on $\beta$

- Notice: very high dimensional inner product!
- apply Johnson Lindenstrauss lemma to reduce dimension


## Johnson Lindenstrauss Lemma

Lemma (The Johnson-Lindenstrauss Lemma)
Consider $x_{1}, \ldots, x_{n} \in \mathbb{R}^{N}$. Pick $0<\epsilon<1$ and set $r=\left\lceil 8(\log n) / \epsilon^{2}\right\rceil$.
There is a linear $\operatorname{map} Q: \mathbb{R}^{N} \rightarrow \mathbb{R}^{r}$ such that for all $1 \leq i, j \leq n$,

$$
(1-\epsilon)\left\|x_{i}-x_{j}\right\|^{2} \leq\left\|Q\left(x_{i}-x_{j}\right)\right\|^{2} \leq(1+\epsilon)\left\|x_{i}-x_{j}\right\|^{2} .
$$

(Here, $\lceil a\rceil$ is the smallest integer larger than a.)

## Johnson Lindenstrauss Lemma

Lemma (The Johnson-Lindenstrauss Lemma)
Consider $x_{1}, \ldots, x_{n} \in \mathbb{R}^{N}$. Pick $0<\epsilon<1$ and set $r=\left\lceil 8(\log n) / \epsilon^{2}\right\rceil$.
There is a linear $\operatorname{map} Q: \mathbb{R}^{N} \rightarrow \mathbb{R}^{r}$ such that for all $1 \leq i, j \leq n$,

$$
(1-\epsilon)\left\|x_{i}-x_{j}\right\|^{2} \leq\left\|Q\left(x_{i}-x_{j}\right)\right\|^{2} \leq(1+\epsilon)\left\|x_{i}-x_{j}\right\|^{2} .
$$

(Here, $\lceil a\rceil$ is the smallest integer larger than a.)
Lemma (Variant of the Johnson-Lindenstrauss Lemma)
Under the same assumptions, there is a linear map $Q: \mathbb{R}^{N} \rightarrow \mathbb{R}^{r}$ such that for all $1 \leq i, j \leq n$,

$$
\left|x_{i}^{T} x_{j}-x_{i}^{T} Q^{T} Q x_{j}\right| \leq \epsilon\left(\left\|x_{i}\right\|^{2}+\left\|x_{j}\right\|^{2}-x_{i}^{T} x_{j}\right)
$$

## Summary

big data is low rank

- in social science
- in medicine
- in machine learning
we can exploit low rank to
- fill in missing data
- embed data in vector space
- infer causality
- choose machine learning models


## References

- Generalized Low Rank Models. M. Udell, C. Horn, R. Zadeh, and S. Boyd. Foundations and Trends in Machine Learning, 2016.
- Revealed Preference at Scale: Learning Personalized Preferences from Assortment Choices. N. Kallus and M. Udell. EC 2016.
- Discovering Patient Phenotypes Using Generalized Low Rank Models. A. Schuler, V. Liu, J. Wan, A. Callahan, M. Udell, D. Stark, and N. Shah. Pacific Symposium on Biocomputing (PSB), 2016.
- Low rank models for high dimensional categorical variables: labor law demo. A. Fu and M. Udell. https://github.com/h2oai/h2o-3/blob/master/h2o-r/demos/ rdemo.census.labor.violations.large.R
- Causal Inference with Noisy and Missing Covariates via Matrix Factorization. N. Kallus, X. Mao, and M. Udell. NIPS 2018.
- OBOE: Collaborative Filtering for AutoML Initialization. C. Yang, Y. Akimoto, D. Kim, and M. Udell.
- Why are Big Data Matrices Approximately Low Rank? M. Udell and A. Townsend. SIMODS, to appear.


[^0]:    ${ }^{1}$ labor law violation demo: https://github.com/h2oai/h2o-3/blob/ master/h2o-r/demos/rdemo.census.labor.violations.large.R

