



“On Granular High-Resolution Analytics in Big Data Applications”

Irada Ben-Gal

Paris, April 9, 2014

<http://www.eng.tau.ac.il/~bengal/>

<http://www.linkedin.com/pub/irad-ben-gal/4/582/4b4>

bengal@tau.ac.il

Agenda

- Granular Big Data Analytics Opportunity
- Pattern-Based Graphical Modeling
- Academic Example & Real-world Applications
 - Anomaly Detection in Cellular Nets
 - Learning Mobility Patterns
 - Predictive maintenance (vehicle battery)
 - Suspect Tracking
 - User-Oriented Cyber Security
 - Demand Sensing & Inventory Control
 - e-Commerce analysis
- Q&A

1.8
zettabytes
in 2011



In 2011, humans created 1.8 zettabytes of data, the equivalent of an HD movie 43 million years long [2]



x 10,000

The volume of data created by U.S. companies alone each year is enough to fill ten thousand Libraries of Congress [3]

How can businesses use that data?

- Transform decision-making
- Discover new insights
- Optimize commerce
- Innovate their industries
- Target marketing efforts [4]

Still, managing all that information can be problematic...



68%

By 2015, 68% of the world's data will be created by consumers

48 hours

of video are uploaded to YouTube every minute [5]

54k

searches are queried every second in Google [7]

1 million

customer transactions are uploaded per hour into Walmart Databases [8]

4.4 million

By 2015, 4.4 million IT jobs

+60%

Effective data use can boost retail profit margins by more than 60% [10]



Consumer data has a big impact on retailers' bottom lines

Poor data quality costs U.S. businesses \$600 billion annually [11]

-\$600,000,000,000



But who actually owns all that personal information?

"Data brokers have developed hidden dossiers on almost every U.S. consumer...[which] raises a number of serious privacy concerns." - U.S. Congress [12]



Over time, companies can create an intimate portrait of a person's familial and professional associations, political religious beliefs, and health status [13]

A number of high-profile companies have faced legal action over this issue

\$22.5 million

In 2012, Google paid a record \$22.5 million fine to settle allegations that it broke a privacy promise by secretly tracking millions of Web surfers who use Apple's Safari browser [14]

Google Inc. isn't admitting any wrongdoing in the settlement [14]

\$20 million

In the same year, a lawsuit alleged that Facebook users' pictures were used to endorse products they had "liked" without their permission. Facebook settled the lawsuit for slightly more than \$20 million [15]



PROS

Big data helps companies turn information into revenue [16]

Those earnings will accelerate growth in the global economy and create jobs [16]

CONS

This multi-billion dollar industry is largely unregulated [12]






After consumers give their information away, they have no idea what companies do with it [12]

Granular “High Resolution” Analytics



Granularity vs. Segmentation

~~Granularity~~



GAP=
Opportunity!

Hidden Patterns Example: Virtual & Smart TV

Products (8)



(e)Locations (6)



Week day (7)



Weather (5)



Customer types (6)



In this example:
4,294,967,296
potential patterns!!!

Sample pattern:
“On rainy Tuesdays
customer A look for
service B via
smartphone if he is
in downtown”

C-B4 identifies
the significant
patterns and maps
them into business
opportunities.


Granular High-Resolution Analytics

- Volume
- Velocity
- Variability
- on Information Quality (Kenett and Shmueli 2013)

Apply to the principles to BD modeling!

Data Compression Graphical models

Instated of library of statistical models



Data



h	l	l	m
l	h	m	h
m	m	h	m
h	l	m	l
m	h	l	h
l	m	h	m

Pattern-Based Monitoring & Predictions per entity

Cellular {ID}+Location {A}+Time slot {Noon} Service request {A} → Anomaly {True} Priority {High}

IP {#136.55.65.02}+Time {Sunday noon}+View {Page A} → Alert {True}, Probability{>0.8}

SP user {Type A} + TV Program B {Low rate} + Promotion {True} + Weather {Rain} → Program C {High} > 0.8


BN Generalization [Ben-Gal I., et al, ~ 5 publications & patents]

$P(X_1 X_2 X_3 X_4 X_5) = ?$ Joint distribution of a sequence of RV




Independent: marginal distributions

$$P(X_1 \cdots X_5) = P(X_1)P(X_2)P(X_3)P(X_4)P(X_5)$$



Markov Model: fixed-order dependency

$$P(X_1 \cdots X_5) = P(X_1)P(X_2 | X_1)P(X_3 | X_2)P(X_4 | X_3)P(X_5 | X_4)$$







BN Tree: Direct dependencies

$$P(X_1 \cdots X_5) = P(X_1)P(X_2 | X_3)P(X_3 | X_1)P(X_4)P(X_5 | X_3)$$



The model can be further generalized to VOBN

VOBN Model (inhomogeneous & state-dependent net of variable-order trees)

Input data: 5 Products, 3 levels of sales {Low, Medium, High}



Zoom in




Note: extension of BN & DTs


Agenda

- Big Data Analytics Opportunity
 - Pattern-Based Graphical modeling
 - **Academic Example**
 - Real-world applications
 - Anomaly Detection in Cellular Nets
 - Learning Mobility Patterns
 - Predictive maintenance (vehicle battery)
 - Suspect Tracking
 - User-Oriented Cyber Security
 - Demand Sensing & Inventory Control
 - e-Commerce analysis
 - Q&A
-


TFBS Identification (Ben-Gal et al., Bioinformatics 2005)




TFBS Identification



TFBS Identification (Ben-Gal et al., Bioinformatics 2005)



Academic Example (Ben-Gal et al., 2005)



1

A	C	Q	T
0.10	0.10	0.12	0.68

2

A	C	Q	T
0.05	0.08	0.09	0.78
1			
A Q			
A	C	Q	T
0.20	0.08	0.20	0.52
A	C	Q	T
0.13	0.17	0.03	0.67

3

A	C	Q	T
0.12	0.14	0.51	0.24
5			
A Q			
A	C	Q	T
0.07	0.07	0.76	0.10
A	C	Q	T
0.03	0.17	0.70	0.10

4

A	C	Q	T
0.51	0.23	0.12	0.13
3			
A T			
A	C	Q	T
0.72	0.08	0.08	0.12
A	C	Q	T
0.58	0.06	0.12	0.24

5


A	C	Q	T
0.18	0.56	0.09	0.17
Q			
C			
A	C	Q	T
0.06	0.39	0.28	0.28

6

A	C	Q	T
0.46	0.13	0.17	0.24
2			
A C Q			
A	C	Q	T
0.68	0.11	0.16	0.05
A	C	Q	T
0.41	0.36	0.18	0.05
A	C	Q	T
0.67	0.04	0.17	0.13

Academic Example 2 (Ben-Gal et al., Bioinformatics 2005)

Sensitivity: 47.8%



Fixed specificity (TN rate) = 99.9%, Std. error = 0.01%

Academic Example 2 (Ben-Gal et al., Bioinformatics 2005)

	TP (for TN=99.9%)
PWM/Markov(2)	43.5%
VOBN models	47.8%
Overall Improvement	+4.3%

E. Coli Genome (~ 4000 genes) $\Rightarrow$ + 172 potential TFBS


Human Genome (~25,000 genes) $\Rightarrow$ + 1075 suspected TFBS

Agenda

- Big Data Analytics Opportunity
- Granular Pattern-Based Graphical modeling
- Real-world Applications
 - Anomaly Detection in Cellular Nets
 - Learning Mobility Patterns
 - Predictive maintenance (vehicle battery)
 - Suspect Tracking
 - User-Oriented Cyber Security
 - Demand Sensing & Inventory Control
 - e-Commerce analysis
- Q&A



Learning Mobility Patterns (LMP)


- Analyzing Mobility Patterns in city areas
 - Using massive data sets
 - Privacy-preserving methods
- Funded by the Ministry of Science Infrastructure grant.
- Join work with Dr. Eran Toch (TAU) Prof. Boaz Lerner (BGU).




LMP: Modeling People & Places

Recognizing meaningful places, analyzing their use, and using probabilistic models to predict people behavior.






Friday, 22:00




Saturday, 03:00




Saturday, 16:00




Sunday, 09:00



Sunday, 19:00



Sunday, 22:00





B


- User 3470

	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2																		
	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##	##																			
0		A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
1		A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
2		A	A	A	A	A			A	A	A	A	A	A	A	A	A	A	A	A	C	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
3		A	A	A	A	A			A	A	A	A	A	A	A	A	A	A	A	C	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
4		A	A	A	A				A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
5		A	A	A	A				A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
6		A	A	A	A				A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
7	C	C	A	A	A	C			A	A	A	A	A	A	A	A	A	A	A	A	C	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
8	C	A	A	C	B	B	C		A	A	C	F	A	C	C	A	A	B	B	B	C			A	C	B	A	B	A	A	C	B	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																	
9	C	A	A	B	B	B	B		A	A	B	B	A	B	C	A	A	B	B	B	B			A	C	B	A	B	A	B	B	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																	
10	B	C	A	B	B	B	B		A	A	B	B	A	B	C	A	A	B	B	B	B			A	A	B	B	A	B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																	
11	B	B	C	A	B	B	B		A	A	B	B	F	B	F		A	B	B	B	B			A	A	B	B		B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																	
12	B	B	C	A	B	B	B		F		A	B	B	B	B		A	F	B	B	B	B			F	A	B			A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																	
13	B	B	F	A	B	B	B		A	A	B	B	B	B		A	D	B	B	B	B			A	A	B	B		B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																
14	B	B	F	A	B	B	B		A	A	B	B	B	B		F	A	D	B	B	B	B			A	F	B	B		B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																	
15	B	B	A	B	B	B	B		A	A	B	B	B	B	B	F	A	A	B	B	B	B	B			A	D	B	B		B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																
16	B	B	A	A	B	B	B		E		A	B	B	B	B	F	C	A	B	B	B	B	B			A	C	B	B		B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																
17	B	B	A	A	B	B	B		E		A	B	B	B	B	A	A	A	B	B	B	B	B			A	B	B		B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																
18	B	B	A	A	B	B	B		C		A	B	B	B	B		A		A	B	B	B	B			A	B	B		B	A	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B	B																
19	C	A	A	A	B	G	B		A	A	B	B	G	C	C	A	A	A	B	B	C	F	A	F		C	B		A	A		A	A	H	B	B	B	B	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
20	E	A	E	A	C		C		A	A	B	A	C	A	C	A	A		F	A		F	A	F		C	F		A	A	C		A	A	H		B	B	C	C	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
21	C	A		A	A		A		A	A	B	A	G	A	A	A	A	A	F		A		F	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A	A																
22		A		A	A	C	A		A	A	C	F	A	G	A	A	A	A	A	E	A			G	A	A	A																																														
23																																																																									

- User 3465

[illegible]

Customers Churn (benchmarked vs. SPSS/R)



The screenshot shows the CLASS software interface with a table of model performance metrics. A red oval highlights the 'Overall Accuracy' column. The table includes columns for 'Use?', 'Graph', 'Model', 'Build Time (mins)', 'Max Profit', 'Max Profit Occurs in', 'Lift{Top ...', 'Overall Accuracy', 'No. Fields Used', and 'Area Under'.

Use?	Graph	Model	Build Time (mins)	Max Profit	Max Profit Occurs in	Lift{Top ...	Overall Accuracy	No. Fields Used	Area Under
☒		C&R Tree 1	< 1	1,268.788	42	1.501	68.789	8	0.707
☒		Quest 1	< 1	1,265	39	1.533	68.72	9	0.716
☒		Discriminant 1	< 1	1,245	47	1.562	68.443	9	0.726
☒		C5 1	< 1	1,245.909	53	1.5	68.235	9	0.734
☒		Logistic regression 1	< 1	1,225	47	1.55	67.82	9	0.726
☒		Neural Net 1	< 1	1,240	45	1.548	67.82	9	0.724
☒		Bayesian Network 1	< 1	1,180	42	1.523	66.782	9	0.713
☒		Decision List 1	< 1	1,070.01	31	1.528	66.09	5	0.667
☒		CHAID 1	< 1	1,123.793	57	1.526	65.952	8	0.727

	R / SPSS	C-B4
Total Accuracy:	69.0%	77.8%
False Alarm:	30.3%	20.2%
False Positive	28.8%	18.5%
Churn Prediction	66.9%	74.1%
Engaged Prediction:	71.2%	81.5%

Gaining accuracy by using personalized granular modeling (per each user!)

LMP Algorithmic Infrastructure

1. Predicting personal and general mobility.



2. Automatically recognizing places, their function and activities.



Understanding urban land-use, eg creating context-aware mobile apps



3. Modeling large-scale social patterns through their mobility footprint.




Mobility for disease spread

Mobility patterns in wars and crisis




Research: Suspects Tracking (leading Intl' Security Company)

Requirement >50% Recall with 1% False Alarm at most




TBNL: Bayesian Network




Network Data
(No text /content Mining)

Research: Suspects Tracking (leading Intl' Security Company)



Pattern Example:



Call/SMS ratio during the day
Suspects vs. **Non-suspects**

Research: Suspects Activities (using C-B4 SW)

CBC performance: >90% Recall with 1% False Alarm at most
(significantly better than all other ML models 60%)



Automated profiling & classification (in natural language)– e.g.:

- “When In-coming call duration is low (0-29sec) and sms to call ratio is high (>23) and fixed, mail activity hour (23:00-02:30); then likelihood of being suspect is high (>87%)”



Cellular Net Usage Monitoring

 **Cellcom** Largest cellular company in Israel

Improved service Levels by 15% and saving millions of dollars

Network data


User Data

Other Data:

- Weather
- Dates
- Promotions
- Location




Output



- Early alerts
- Anomaly detection
- New demand/usage patterns
- Root cause analysis

Predictive Maintenance and Usage Monitoring


 Largest Cellular company in Israel



Predictive Maintenance and Usage Monitoring



Largest Cellular company in Israel



Thank You