Preemptive Shipping

How Gilt Predicts Which Customers Will Buy Products It Has Never Sold Before

Igor Elbert, Principal Data Scientist, Gilt.com

GILT

Before Gilt – sample sales

Gilt pioneered online "flash sales" in US

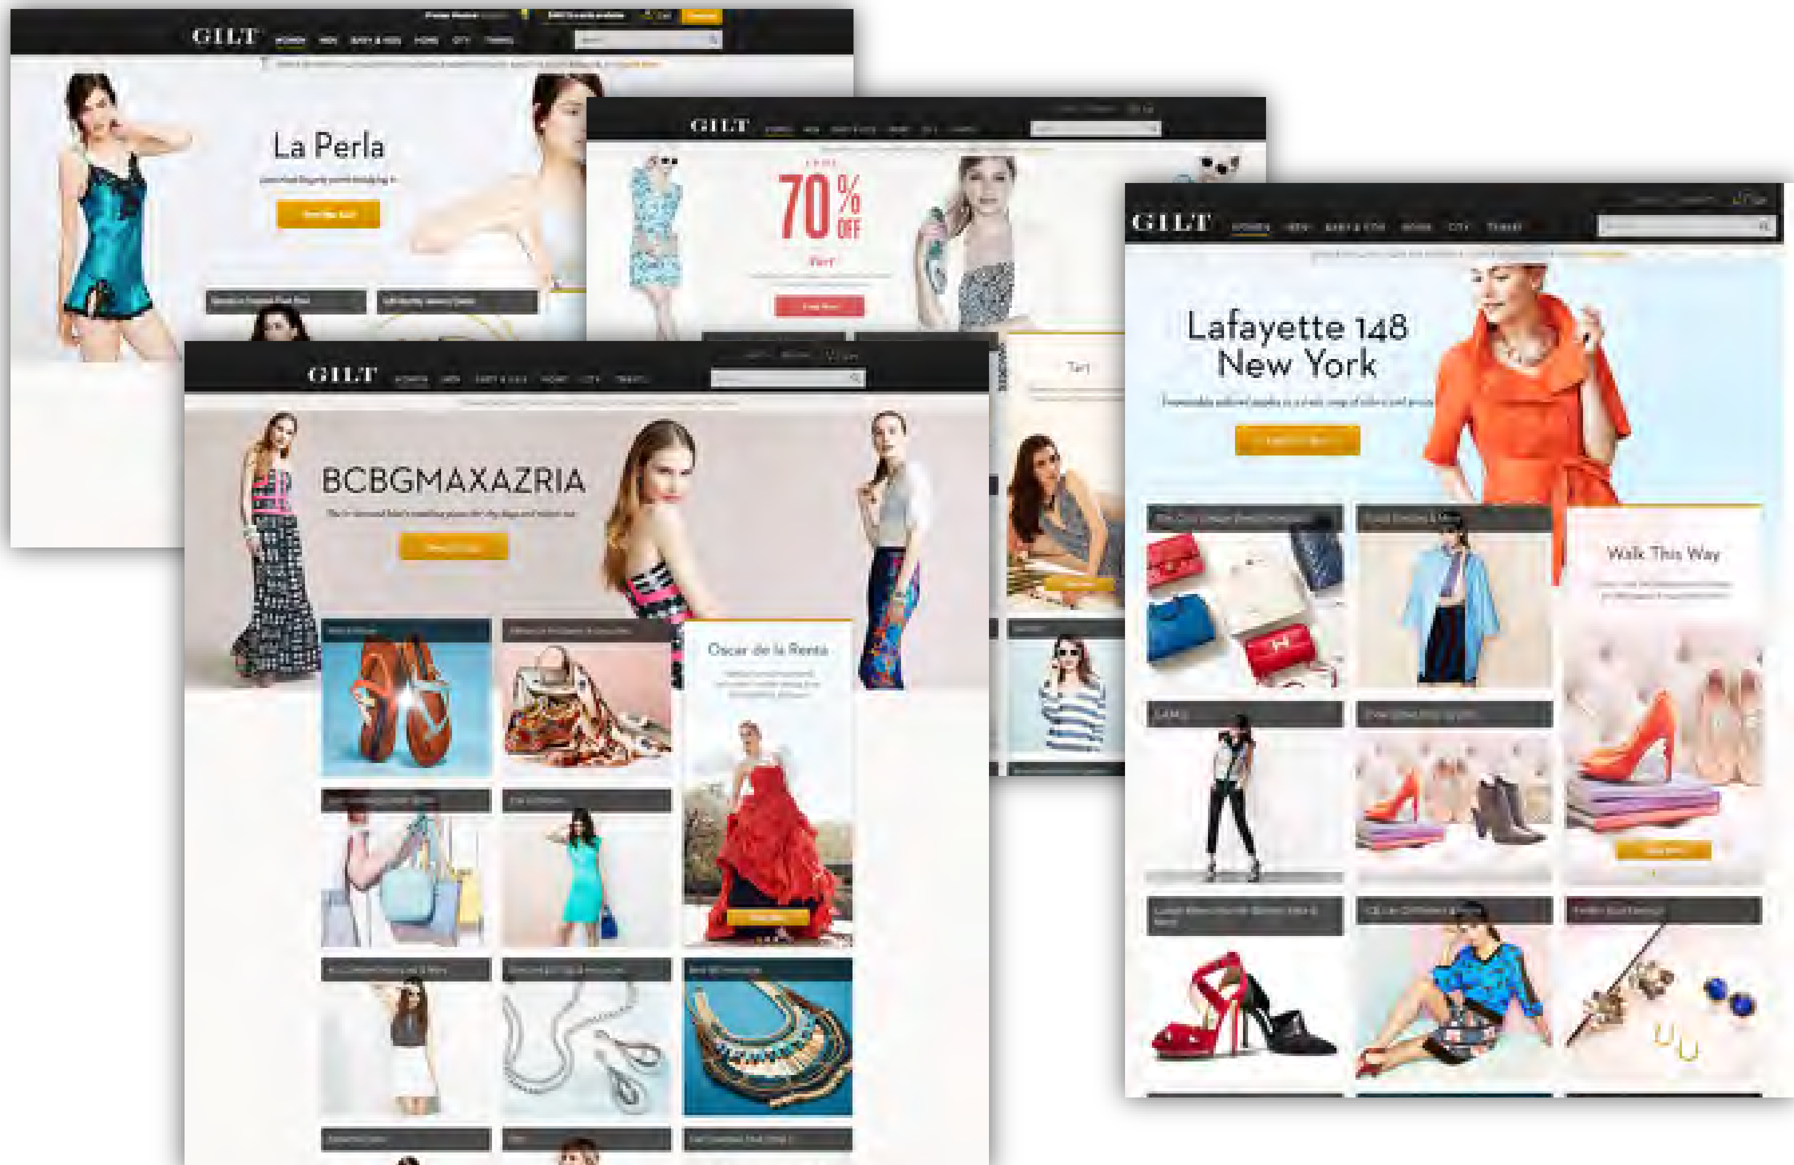

GILT

Can we shorten shipping time?

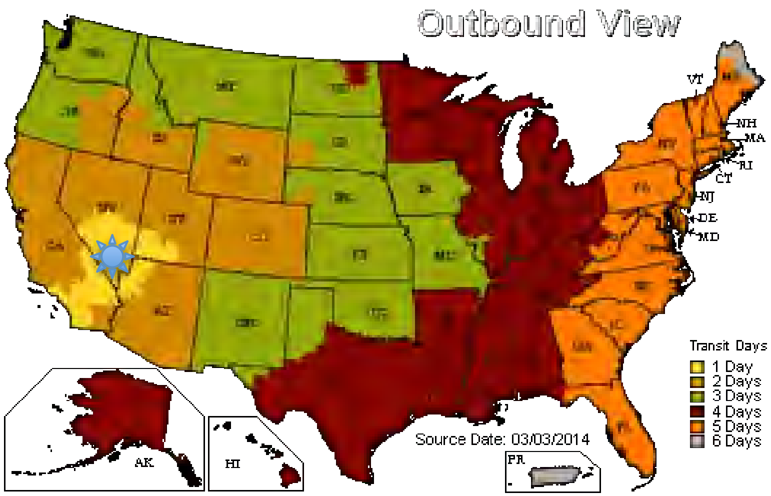

Shipping times from Las Vegas

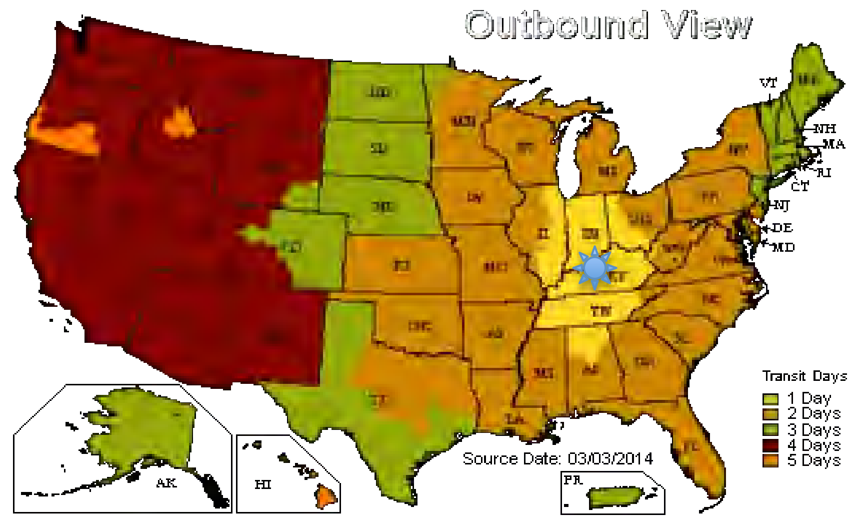

Shipping times from Kentucky

What if we move items closer to potential buyers?

GILT

Goal: Predict which items will be sold 'West'

How is it different from what retailers have always been doing?

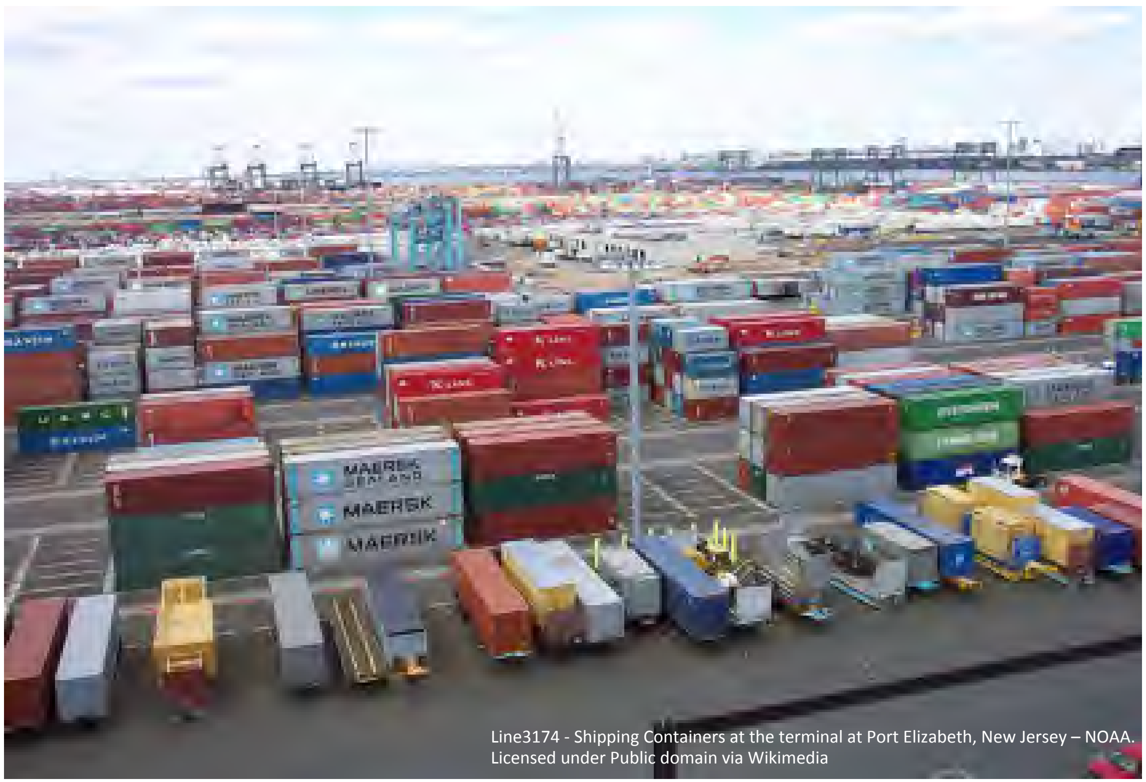

Line3174 - Shipping Containers at the terminal at Port Elizabeth, New Jersey – NOAA. Licensed under Public domain via Wikimedia

GILT

Challenges:

- **Volatile preferences, context**
- **Regional differences**

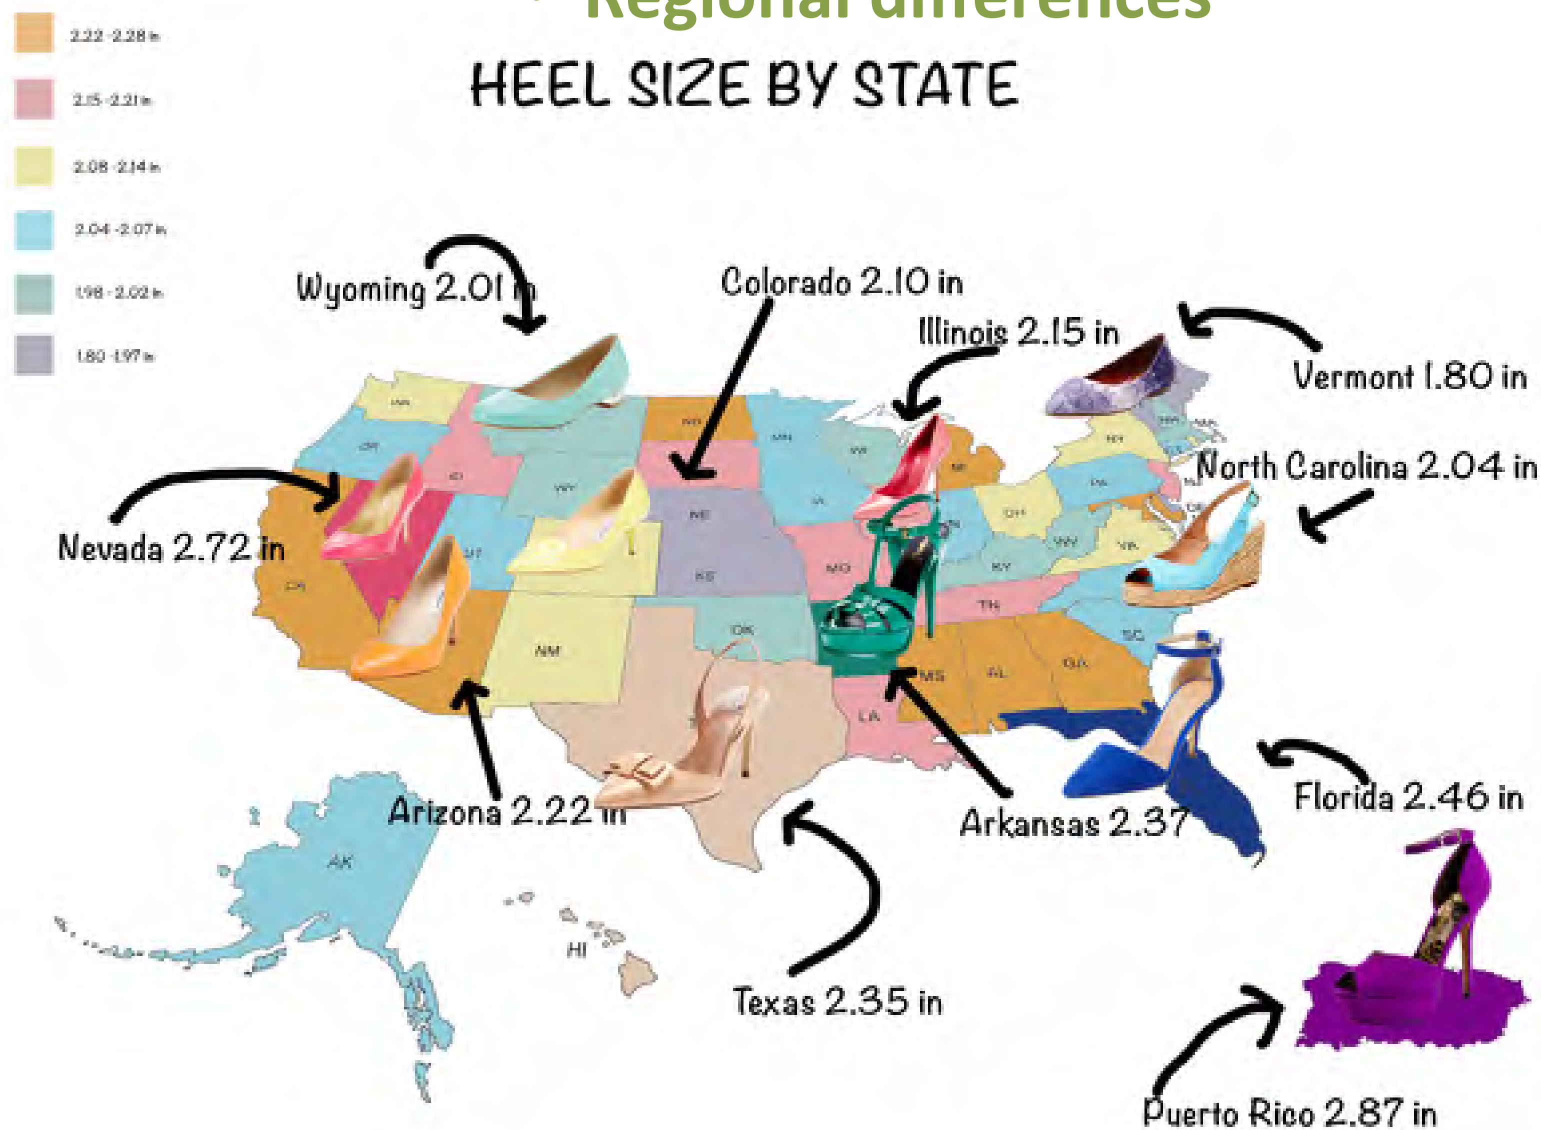

GILT

Challenges:

- Narrow decision window
- No sales history for new products

GILT

But we have a lot of data…

- Orders
- Product Information
- Clickstream

and tools to handle it:

TERADATA ASTER DISCOVERY PLATFORM

What would be a good model?

		Predicted	
		‘West’	‘non-West’
Actual	**‘West’**	50 ☺	15 😐
	‘Non-West’	15 $ ☹	20

For this use-case Precision is more important than Recall

What would be a good model?

		Predicted		
		‘West’	‘East’	No-sale
	‘West’	50 ☺	6 😐	9 😐
Actual	**‘East’**	3 $ ☹	8	6
	No-sale	12 $$	2	4

For this use-case Precision is more important than Recall

Maximize ratio of True-Positives to False-Positives-No-Sale

GILT

Propensity thresholds

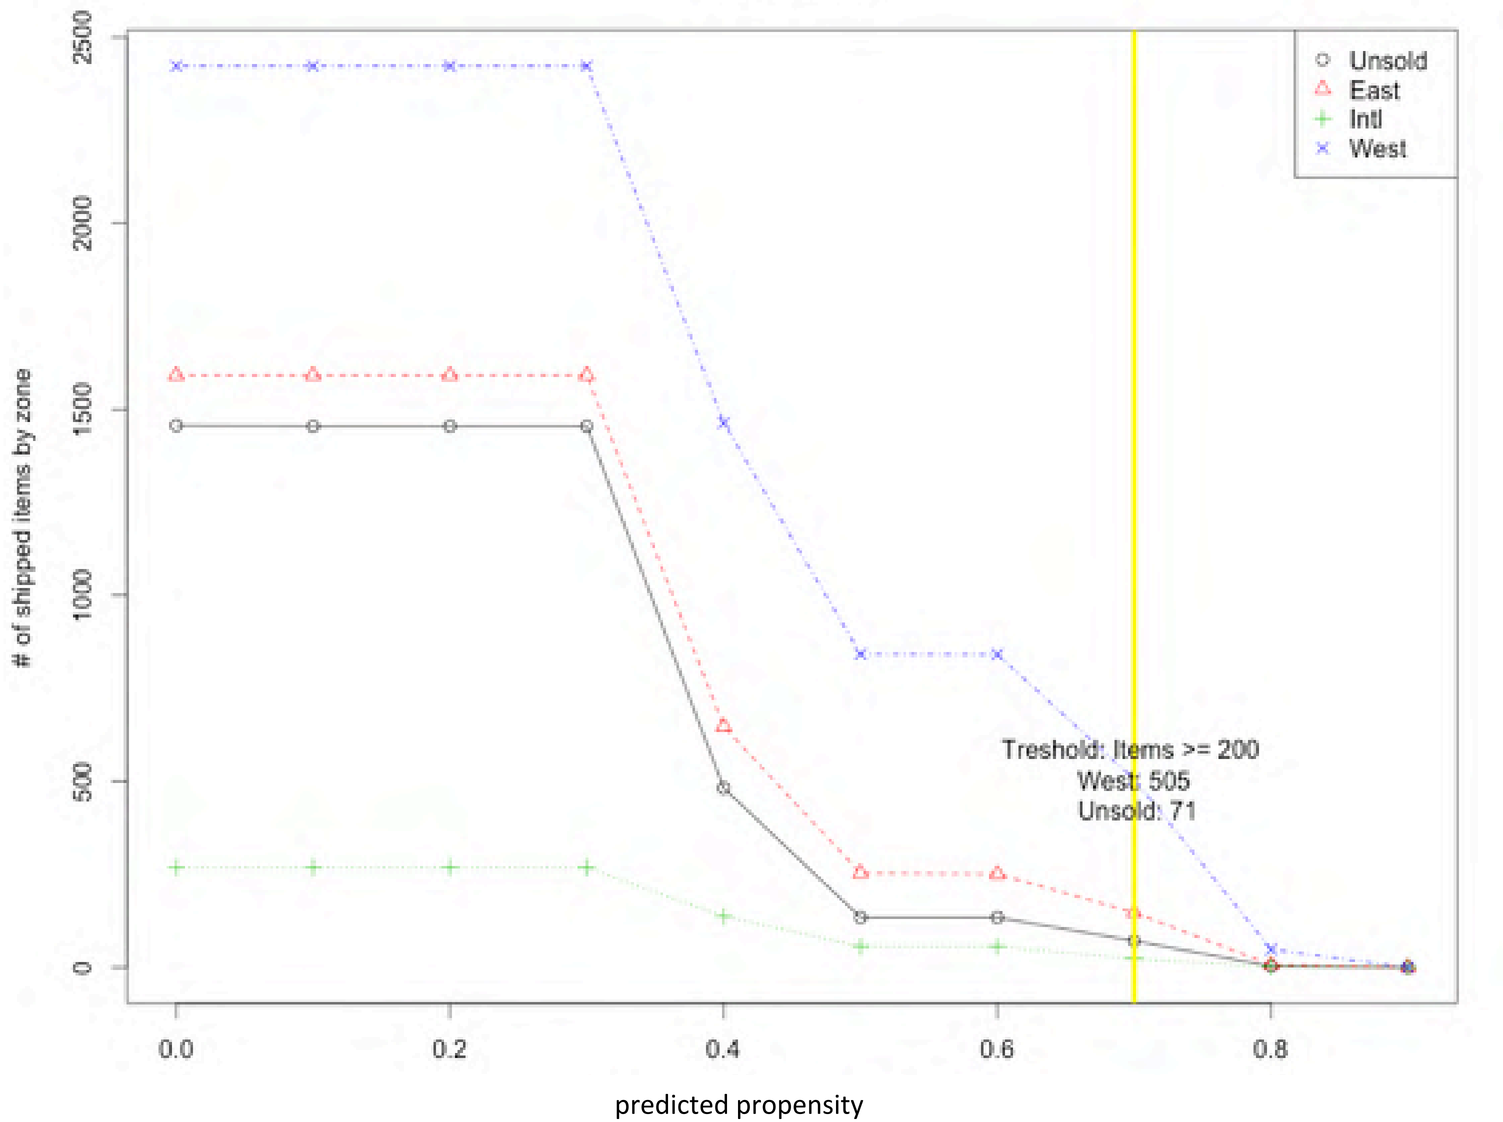

GILT

Initial approach:

Color

Size

Material

Brand

Category

Price

Margin

Discount

Other product metadata

GILT

Data Mining and the Big Picture

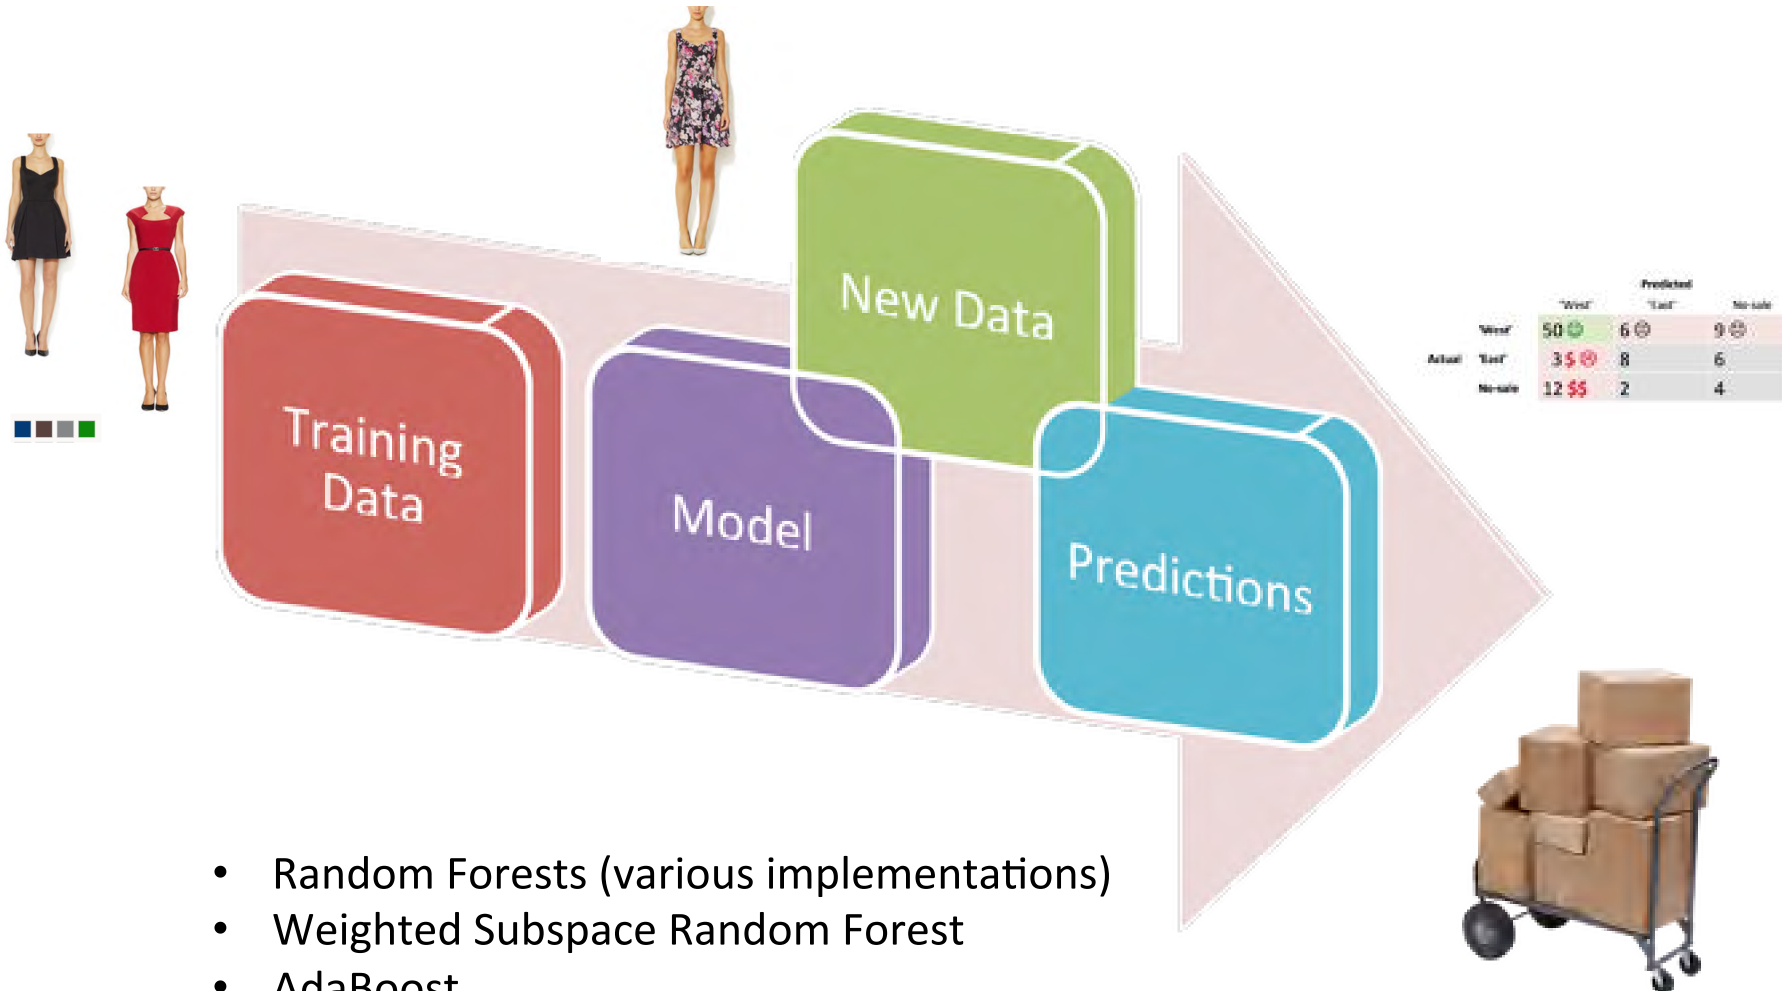

- Random Forests (various implementations)
- Weighted Subspace Random Forest
- AdaBoost
- NNet
- Others

GILT

Bigger picture – Data is the King

a. Dealing with high cardinality
b. When throwing away data is a good thing
c. Wonders of cheating
d. Let's be subjective here

GILT

High cardinality attributes

Error in randomForest.default(m, y, ...) :
Can not handle categorical predictors with more than **53** categories.

Reducing cardinality:

0. **Clean thy data**:

```
REGEXP_REPLACE(TRIM(COALESCE(LOWER(material), '')),
          '(elastace|elatsane|elastine?|elastan[^e]|elastan$)',
          'elastane', 'g')
```

1. **Clean it again**: Levenshtein Distiances:

target_cnt	source_cnt	target	source	distance
40	1	100% nylon	1005 nylon	1
1368	3	100% cotton	100 % cotton	1
60	4	95% cotton 5% elastane	95 cotton 5% elastane	1
1368	4	100% cotton	1005 cotton	1
….	….	….	….	….
21	11	leather suede	leather or suede	3
80	15	suede leather	suede or leather	3

GILT

High cardinality attributes (cont.)

2. **Cut the long tail**. Algorithm specific.
If algorithm does not complain about hard cardinality it does not mean it benefits from it. Use either top N attribute values by number of cases or as many values as needed to cover X% of cases.

3. **Cheat** – picking into test data is now allowed.
Before: Model training took a long time. So model was built to score many data sets.
Now: Model training is quick. It can be adopted to the testing set. Knowing which values are in the test set we can build better model.

4. **Slice the data set** – each subset will have its own set of attribute values

5. **Rough it up**: 'royal blue' is 'blue'
Clustering helps

GILT

Clustering brands

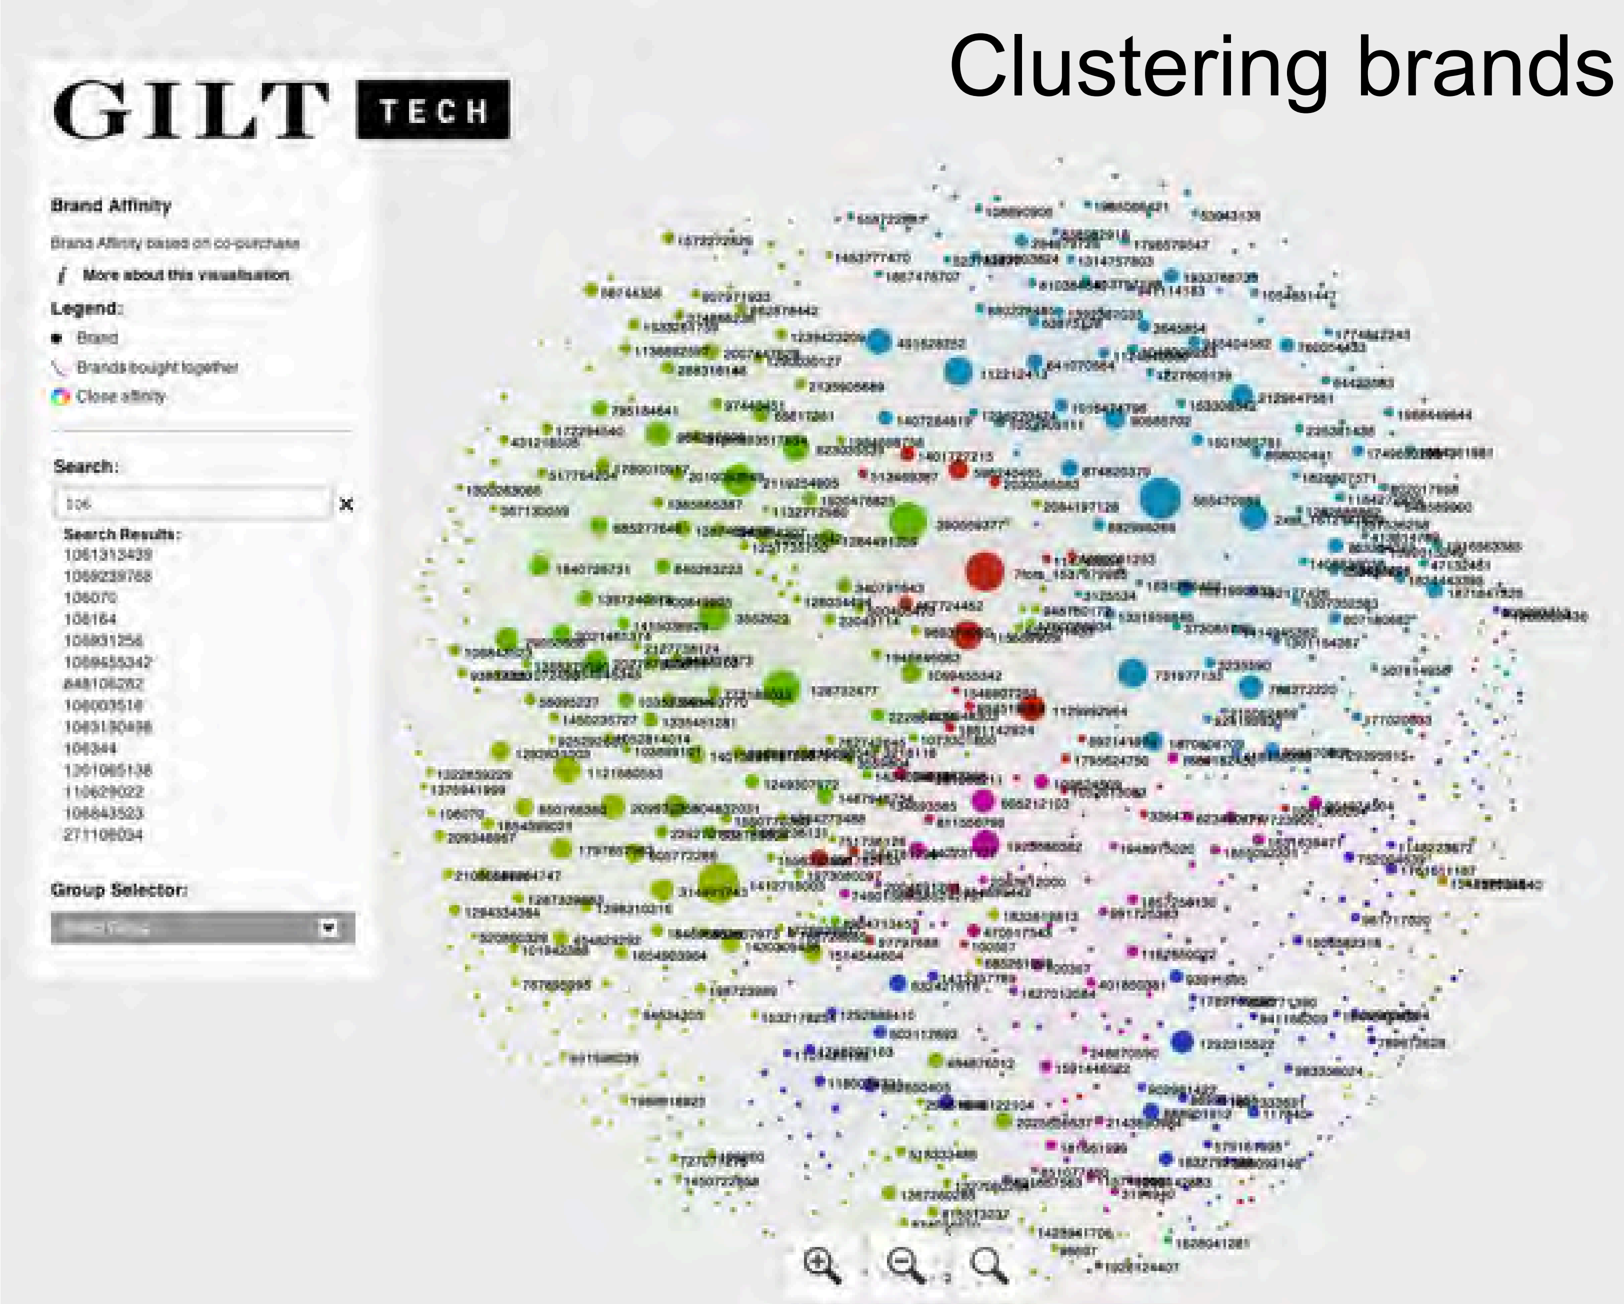

…helps with new brands too

‘Subjective’ attributes:

- Who is Gilt’s buyer for the item?
- Who curated the sale?
- Who set the price?
- Who is the photographer?
- …

Side effects: business insights, best practices

(use simple methods do describe and explain)

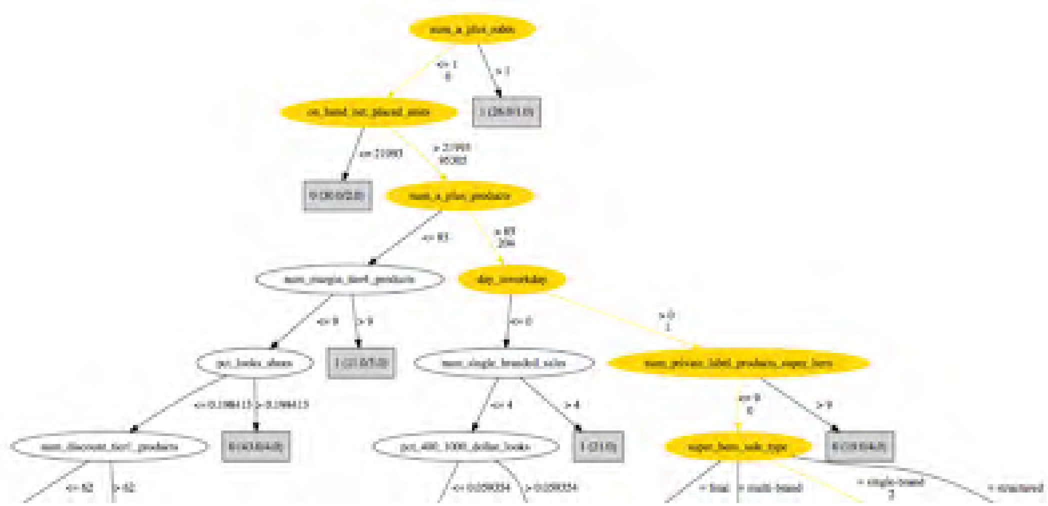

More subjectivity: asking strangers for data

What do you think about this dress?

Material: 100% silk

1. This dress is appropriate for... (select ONE OR MORE)

- Prom
- Graduation
- Birthday Party
- This is not a dress
- Elegant Dinner Party
- Cocktails
- College
- Date
- Girl's Night
- Saturday Night Style
- Work
- Wedding Guest
- Daytime
- Destination
- Bride Dress
- Mother of the Bride
- Bridesmaid
- Bridal Shower
- Engagement Party
- Rehearsal Dinner
- Wedding Brunch
- Other

2. I think this dress is... (select ALL THAT APPLY)

- Pretty
- Fancy
- Unusual
- Beautiful
- Sexy
- Casual
- Weird
- Frivolous
- Formal
- Other

3. I think it's:

- Extremely Expensive
- Expensive
- Moderately priced
- Inexpensive

4. I guess this dress is sold for:

from ____ to ____

(put numbers in US dollars)

5. I think this dress will:

- Will be in strong demand
- Will be in moderate demand
- Only few people will buy it

Template for Amazon's Mechanical Turk

GILT

Using Mechanical Turk:

- Good quality responses
- Quick turn-around
- Produced several good predictors
  - expected and unexpected
- User-generated content (waiting or analysis):

> "This dress looks like it would work well mainly for hourglass figures. The banding looks very atypical, so it would probably get the dress more attention than without the banding."

- Easy to automate (MTurkR R package)

This dress will sell
…. on a good day

- Day of the week
- Day of the year
  - previous, next holiday
- What else is on sale – possible ‘halo effect’
- Total number of items on sale
- Price line-up
- Proxies of traffic – more visitors, higher chances to sell
- Time of sale – competing with other shoppers

Rewards:

- Dry runs look promising – we expect significant reduction in days-in-transit
- Useful side effects:
  - Propensity to sell – pick-and-pack optimization
  - Propensity to not sell – insights on pricing, merchandizing, inventory
- 'Stress-test' for company's logistics – order routing, sale creation, shipping

Conclusions:

- Shop at Gilt - we need data
- See big picture
- Clean and trim the data
- Use subjective data

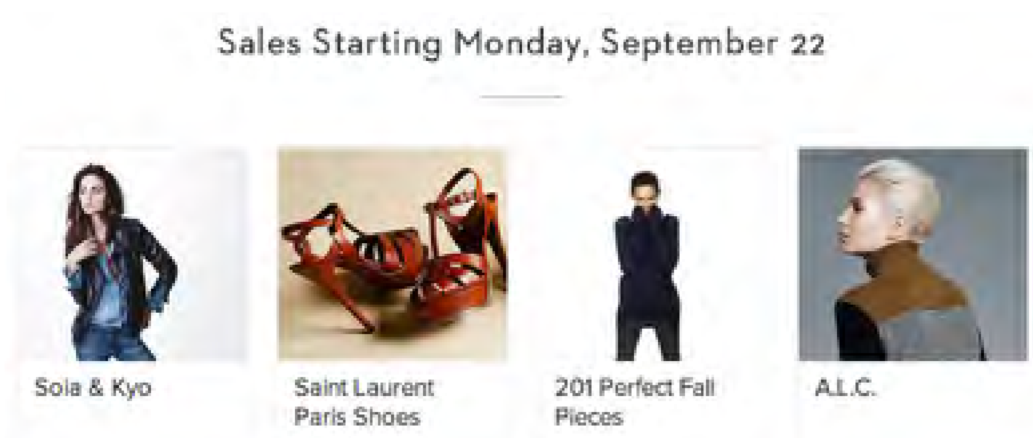

GILT

Questions?

Igor Elbert
ielbert@gilt.com

GILT