NEXUS

LEADERSHIP+

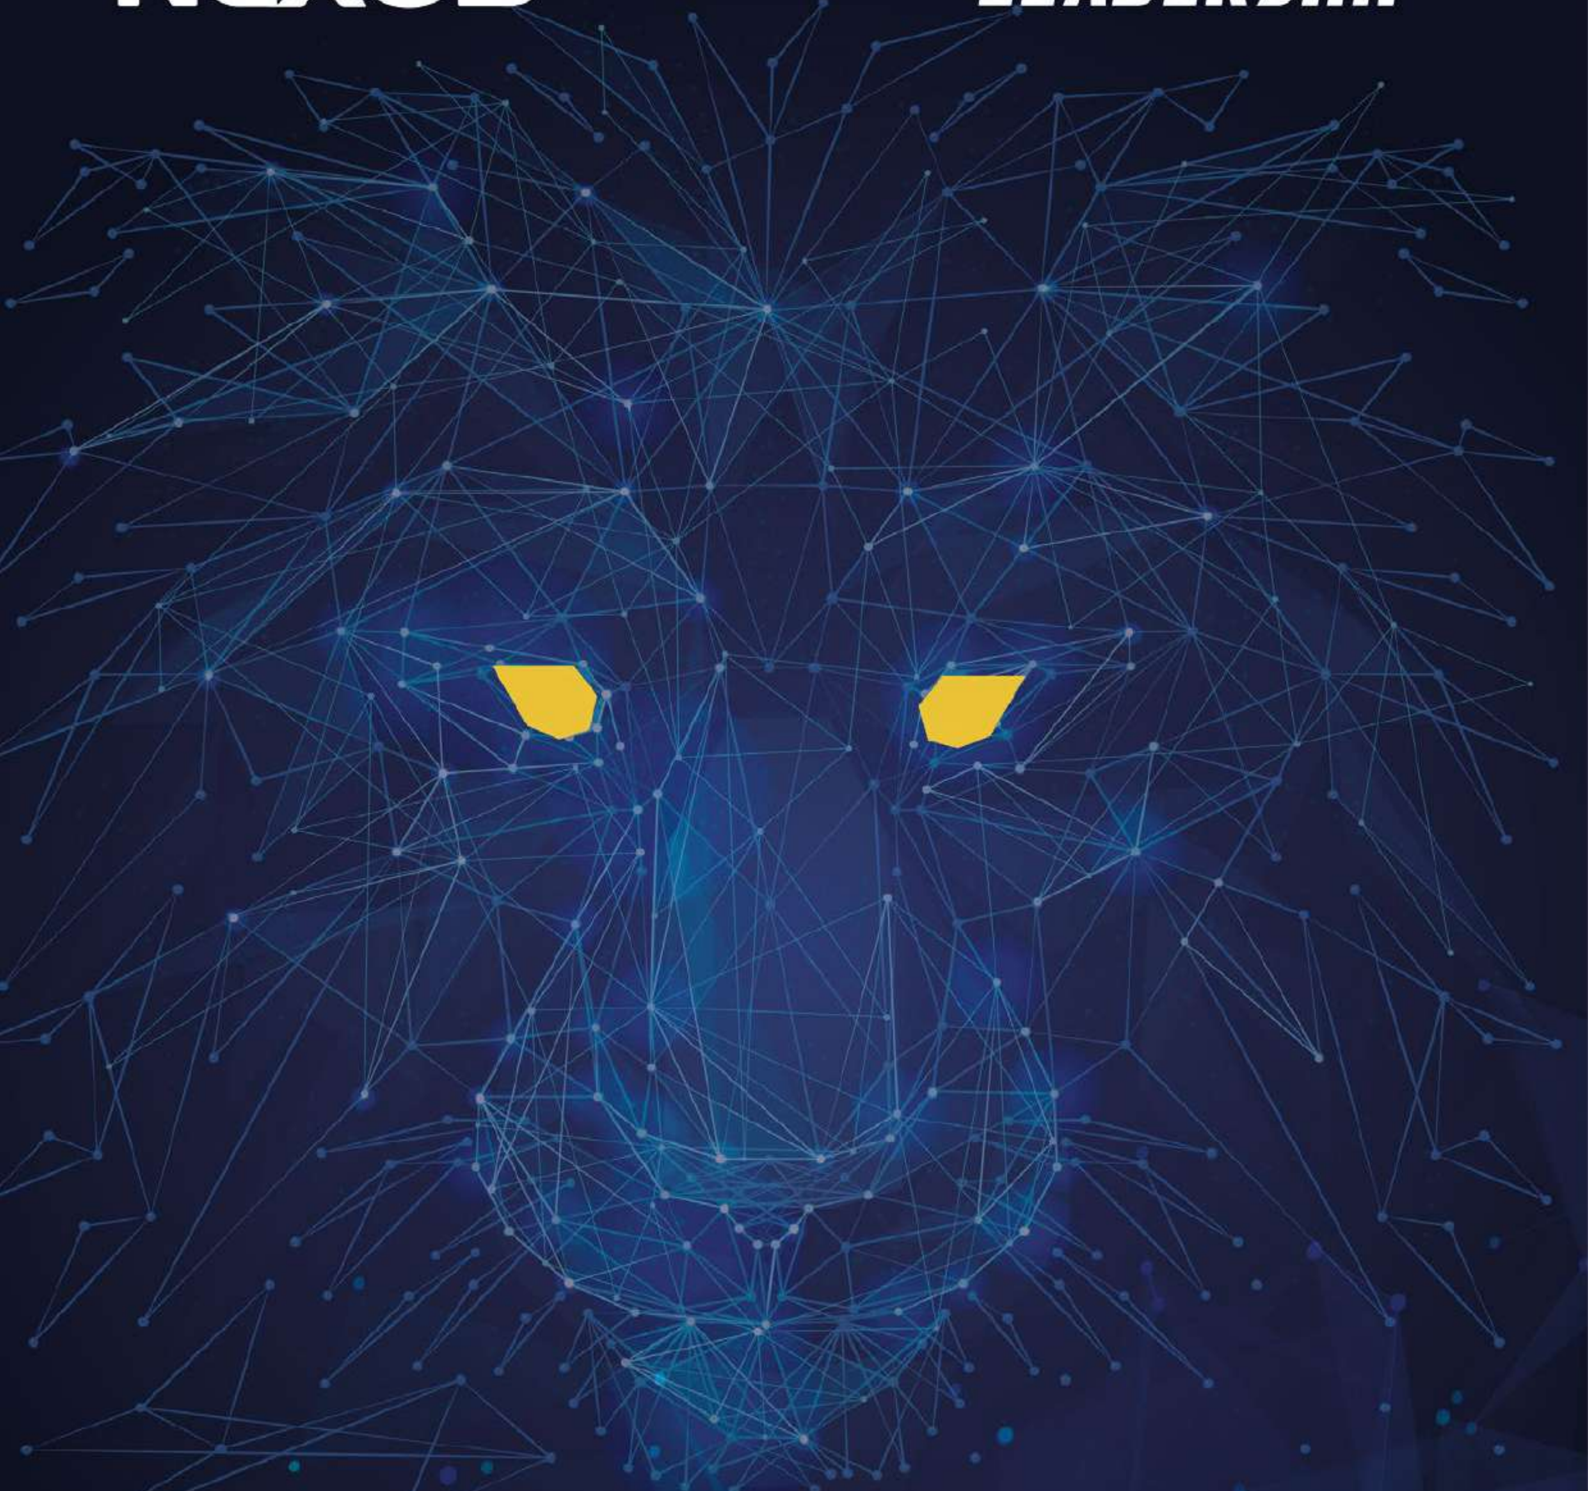

THE GAME CHANGERS

The AI Playbook

Global Intelligence Association
May 2024

Eric Siegel, Ph.D.
Founder, Machine Learning Week
Founder & CEO, Gooder AI
Author, ***The AI Playbook***

To continue your learning beyond this keynote:

Machine Learning Leadership and Practice – End-to-End Mastery

This end-to-end, three-course series will empower you to launch machine learning. Accessible to business-level learners and yet vital to techies as well, it covers both the state-of-the-art techniques and the business-side best practices.

http://www.MachineLearning.courses

Also, for more information/citations regarding the examples in this presentation, see the Notes, freely-accessible online, for the book "Predictive Analytics" by Eric Siegel (http://www.thepredictionbook.com). Most of the various examples shown are covered in the book (some only briefly, within the book's Central Tables of 182 mini-case studies, so not necessarily with more detail there than in this presentation). So, for greater detail about each case study named, see its reference/citation - search by organization name within the book's Notes PDF, available online at http://www.PredictiveNotes.com.

business

data

prediction

machine learning

predictive AI

predictive analytics

boost sales

cut costs

combat risk

prevent fraud

fortify healthcare

streamline logistics

conquer spam

win elections

Prediction as a capability -- calculating probabilities -- is the Holy Grail for improving large-scale operations.

Quote from the foreword of *The AI Playbook*, by Eric Siegel (MIT Press, 2024).

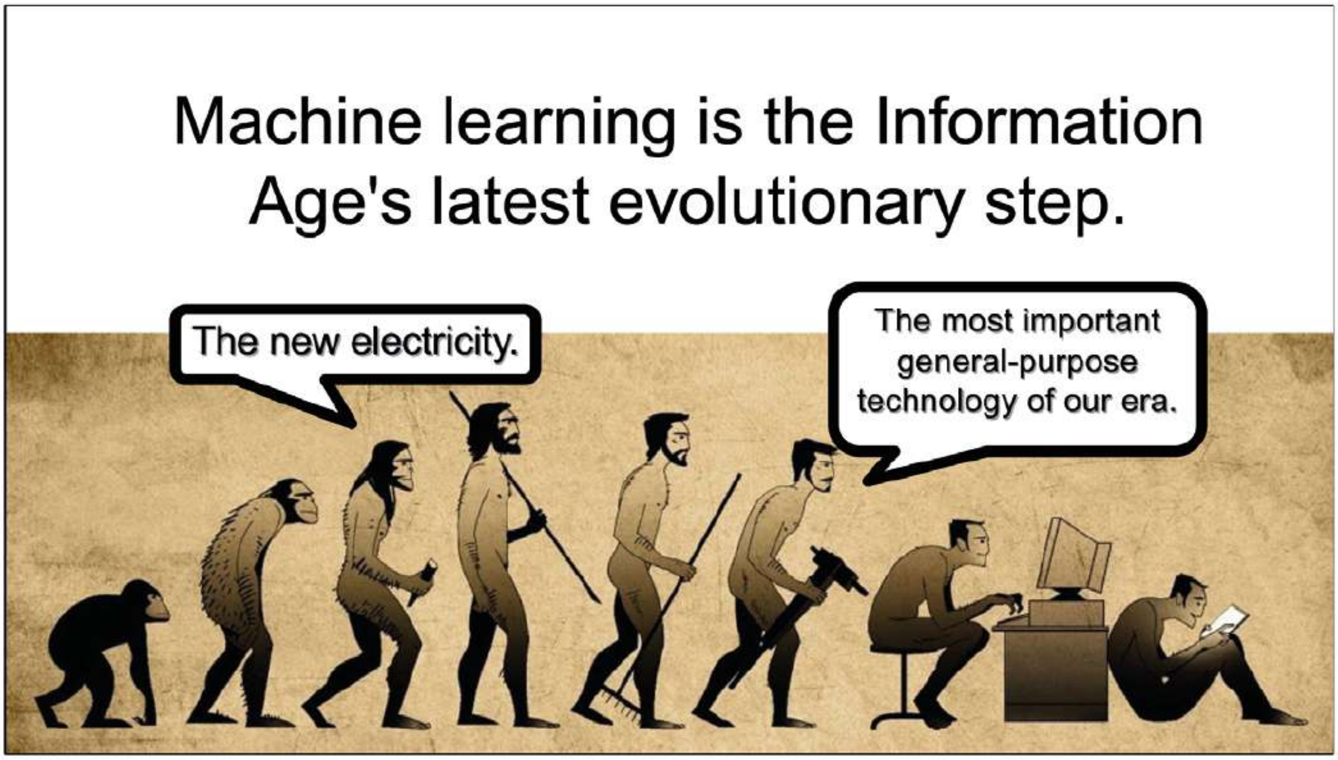

Machine learning is:
"The most important general-purpose technology of our era." –Harvard Business Review (https://hbr.org/2017/07/deep-learnings-next-frontier)
"The new electricity." –Andrew Ng, Stanford University professor, co-founder and chairman of Coursera
"The Information Age's latest evolutionary step" – Eric Siegel

This singular, universally applicable force improves every large-scale thing we do—how we build things, sell things, and prevent bad things from happening—because every function benefits from prediction. As its deployment takes hold across industries, we have moved beyond engineering the infrastructure that manages big data to implementing the science that taps its contents for more intelligent

Has emerged as a commonplace business practice necessary to sustain competitive advantage.

$6.5 billion within a couple years
84% of marketing leaders intend to increase the role of predictive analytics in marketing over the next 12 months (Forbes)
U.S. shortage of analytics experts at 140,000 in the near-term
LinkedIn's first/second "Hottest Skills That Got People Hired" is "statistical analysis and data mining"

Isn't prediction impossible?

predictive AI

OPERATIONS:

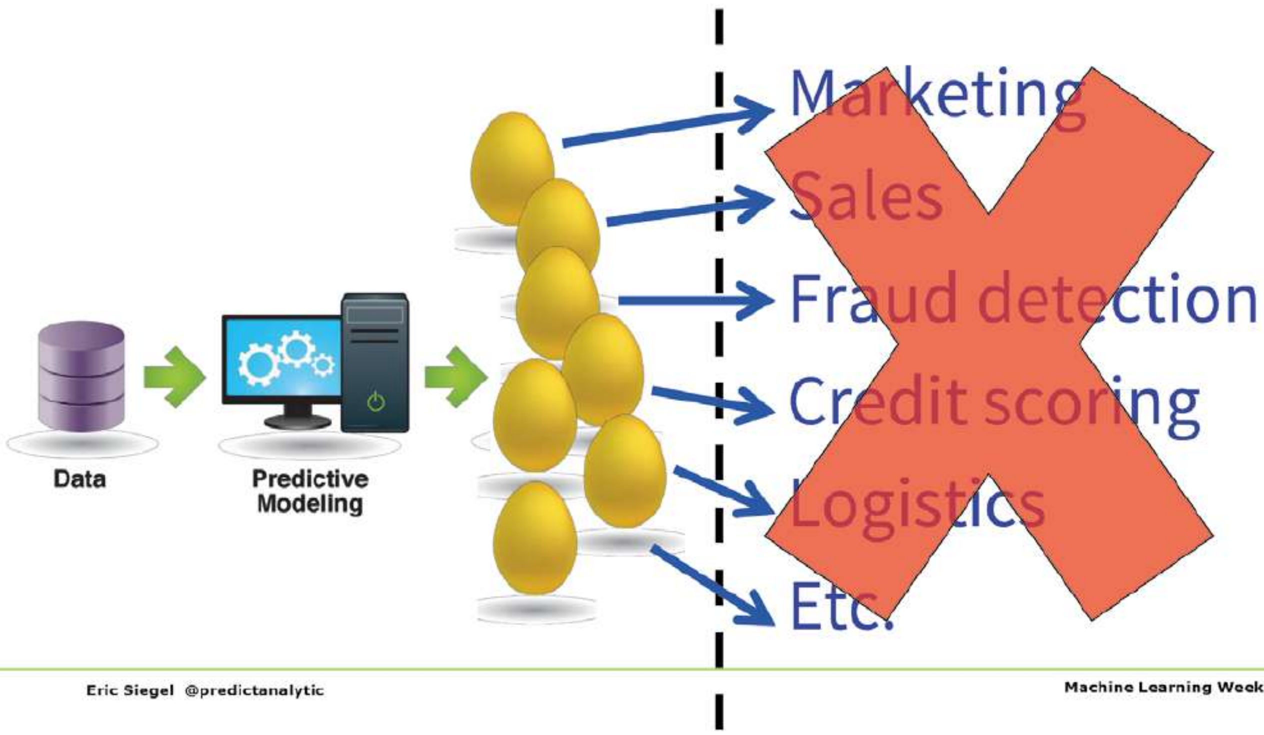

There is no established standardized business practice/paradigm/playbook for running machine learning projects that's well known to leaders, managers, and other business professionals.

However, there is an informally-established paradigm that's widely understood among senior, experienced data professionals. So, first and foremost, the greatest outstanding need is: **Semi-technical understanding by business stakeholders**.

We must greatly alleviate the *knowledge asymmetry*.

Who's meant to form and participate in the business process? Both sides. And which side takes up the mantle? Often, neither does. As you'll see, both sides have reason to consider it someone else's responsibility.

The strategic playbook for running ML projects often defaults to a neglected "no man's land."

Agenda

- Define machine learning / predictive AI
- How it delivers value
- Why ML projects routinely fail to deploy
- Semi-technical understanding
- The bizML paradigm
- Generative AI

Eric Siegel @predictanalytic | Machine Learning Week

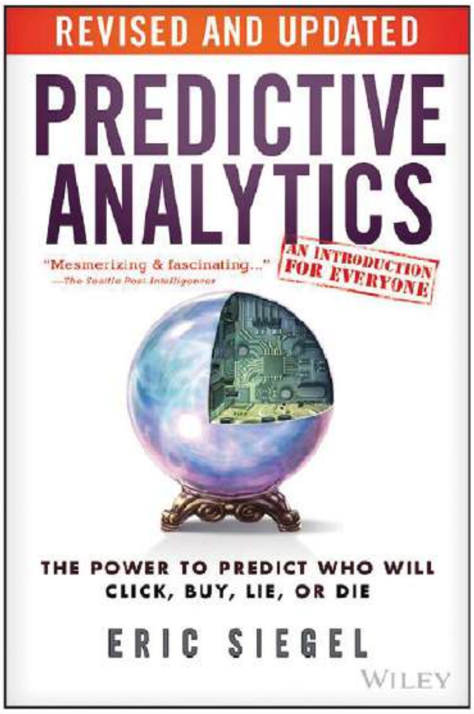

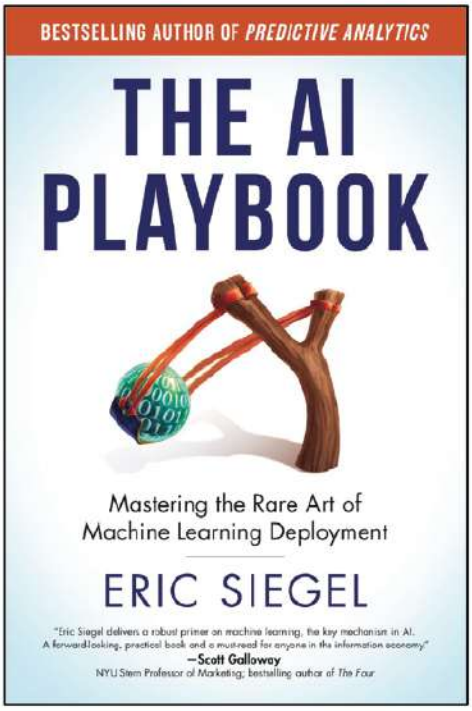


the PREMIER MACHINE LEARNING CONFERENCE
MACHINE LEARNING WEEK
JUNE 4-7, 2024
EXPO: JUNE 6-7
SHERATON PHOENIX DOWNTOWN
PHOENIX, AZ
Generative AI
APPLICATIONS SUMMIT
www.MLweek.com


GOODER AI
Maximize ML's value
by testing and visualizing its business performance.

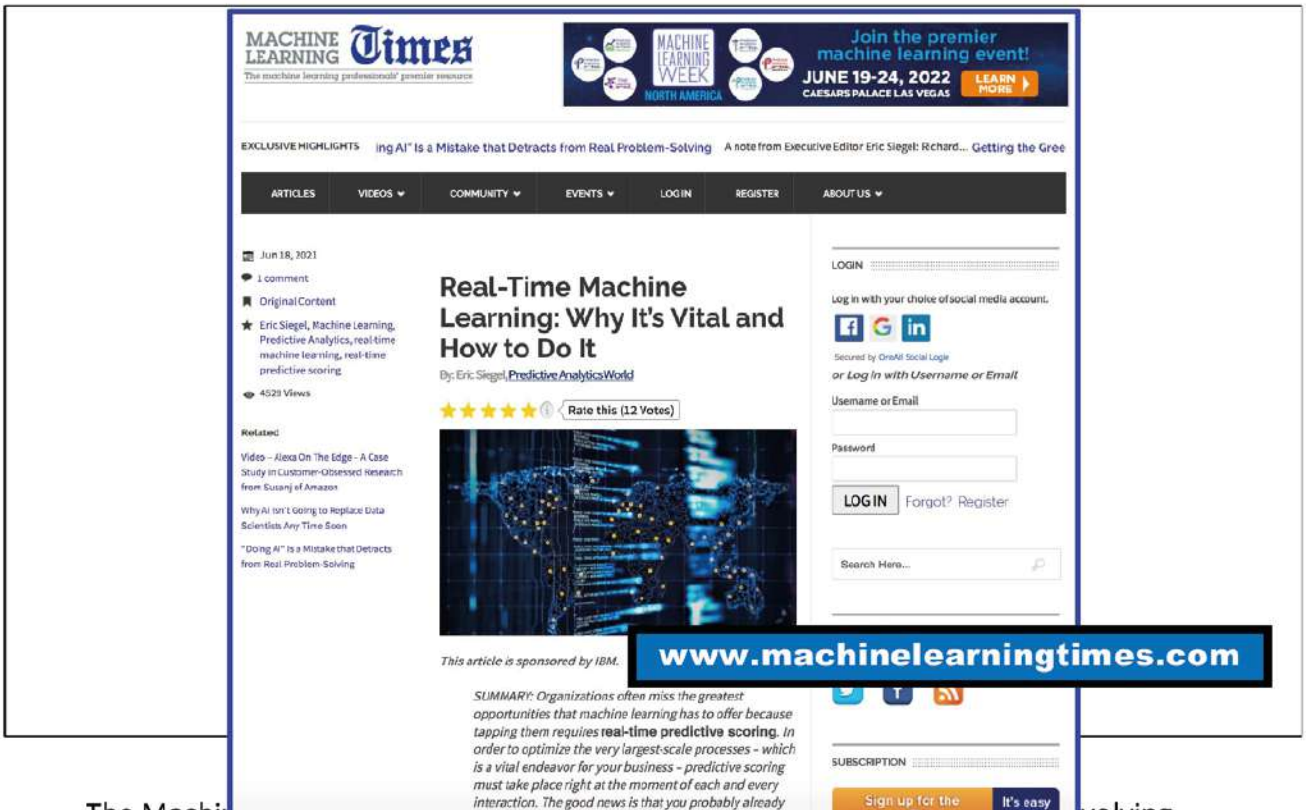

The Machine Learning Times has become the destination for the constantly evolving machine learning community.

http://www.machinelearningtimes.com

This site is the machine learning professionals' premier resource, delivering timely, relevant industry-leading content: articles, videos, events, white papers, and community. The only full-scale content portal devoted exclusively to predictive analytics and its commercial deployment, the Predictive Analytics Times has become a standard must-read.

For the article shown above, see:

https://www.predictiveanalyticsworld.com/machinelearningtimes/real-time-machine-learning-why-its-vital-and-how-to-do-it/12166/

Machine Learning Leadership and Practice – End-to-End Mastery

This end-to-end, three-course series will empower you to launch machine learning. Accessible to business-level learners and yet vital to techies as well, it covers both the state-of-the-art techniques and the business-side best practices.

http://www.MachineLearning.courses

@predictanalytic

Knee Walking

> *Insurance is nothing but management of information. It is pooling of risk, and whoever can manipulate information the best has a significant competitive advantage.*

Eric Webster
VP Marketing, State Farm

"Discussion with State Farm's Eric Webster: Insurance and Data Mining," Gregory Piatetsky, Ph.D., KDNuggets. http://www.kdnuggets.com/news/2009/n08/3i.html

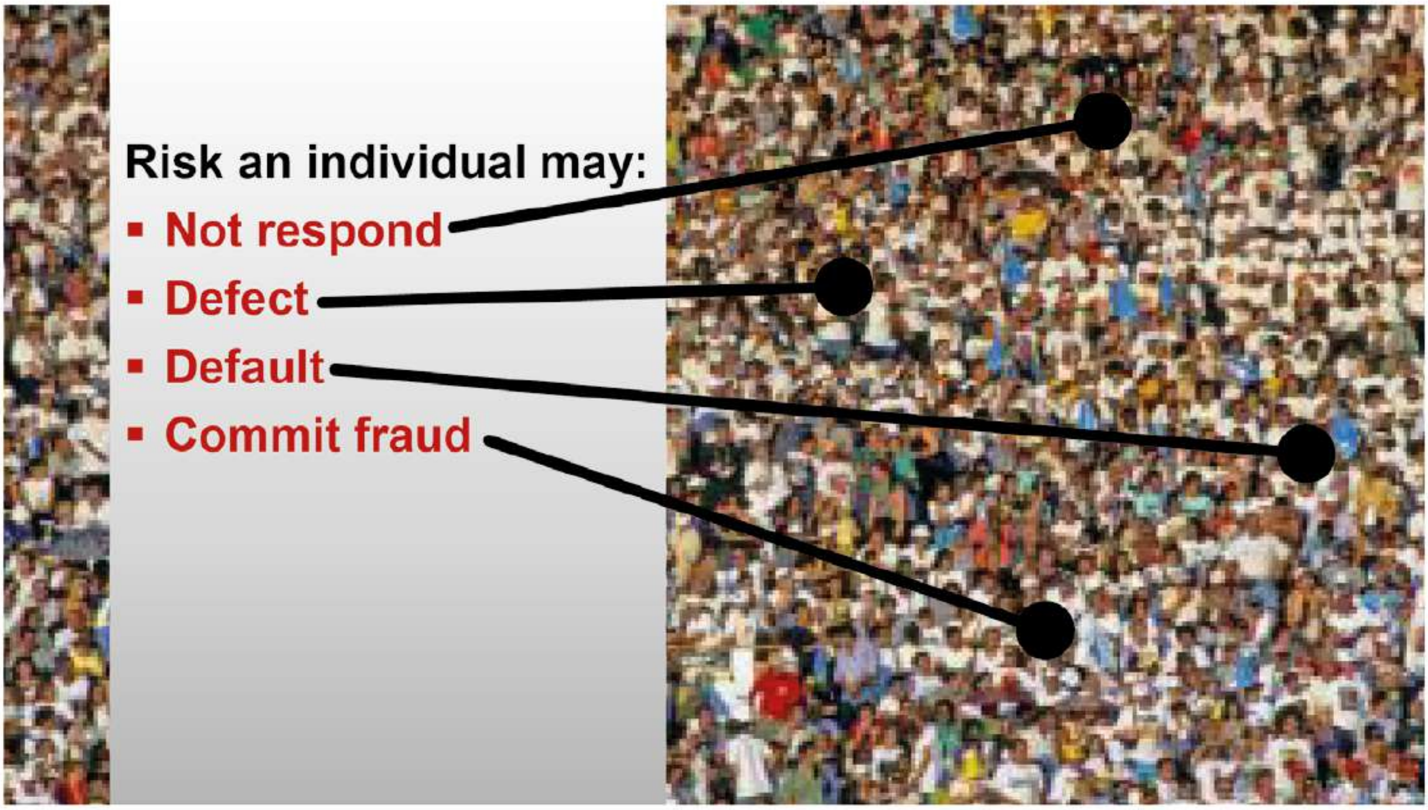

> *No certified, regulated profession like the actuarial practice exists outside of what is strictly considered insurance.*

Douglas Hubbard
The Failure of Risk Management

Eric Siegel @predictanalytic

Machine Learning Week

Also: http://www.prnewswire.com/news-releases/deloitte-poll-37-percent-of-companies-report-having-no-consistent-risk-data-mining-reporting-and-analysis-116644679.html

Triage

Across many use cases, the predictive model proactively targets according to risk or opportunity. It earmarks the individuals with the highest risk or potential gain -- those worthy of investing limited time and resources.

- Infrastructure risk
- Fault detection
- Sales leads
- Fraud detection
- Law enforcement
- Search results

The Antidote to Information Overload

- Google search results
- facebook orders news feed
- Email filters spam

NETFLIX

Spotify

amazon

match

More examples:
- person to tag in a photo (Facebook)
- ad that interests you (also Facebook)
- people to "friend" (Facebook and LinkedIn)
- least-congested driving route (Waze and IBM)
- optimal shipping routes (for UPS drivers).


85 85 84 84 84 75 70 70 70 65 65 65 65 65 64 60 60 60 60
55 50 50 50 50 43 32 29 29 29 24 22 19 19 0 0 0 0 0


85 85 84 84 84 75 70 70 70 65 65 65 65 65 64 60 60 60 60
55 50 50 50 50 43 32 29 29 29 24 22 19 19 0 0 0 0 0

First 100 cases:

10111111110111101111001111110111111011111111111110111110011111101
11111011110111110011101111110110111

Middle 100 cases:

1011010001011010111000100010111101110000011010111010011011010
10001011101101000101101110101101110000

Last 100 cases:

11000000001010010000100000001000000000100001000000100000011000
0000100100001000011000010000100000100100

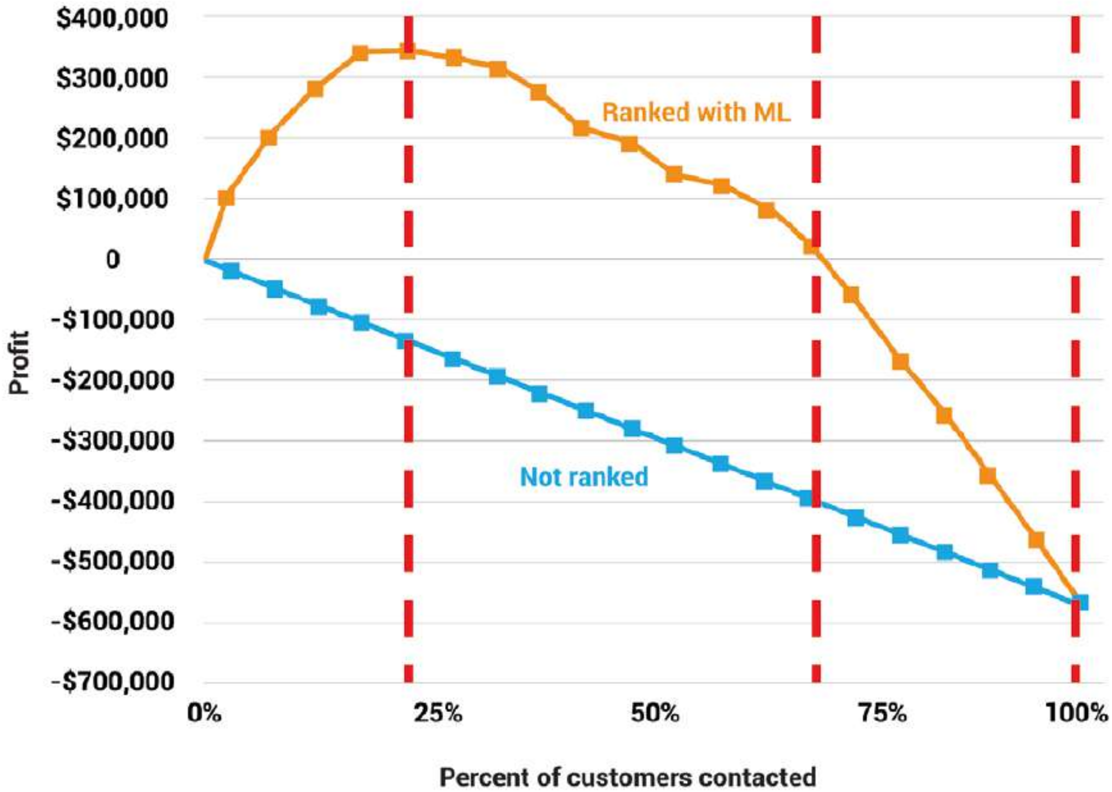

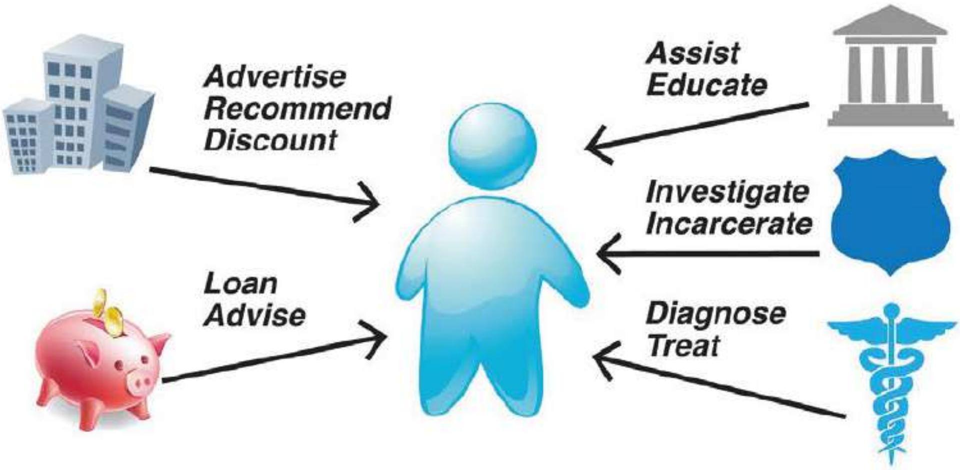

Millions of decisions a day determine whom to ***call, mail, approve, test, diagnose, warn, investigate, incarcerate, set up on a date, and medicate.***

Eric Siegel @predictanalytic

Machine Learning Week

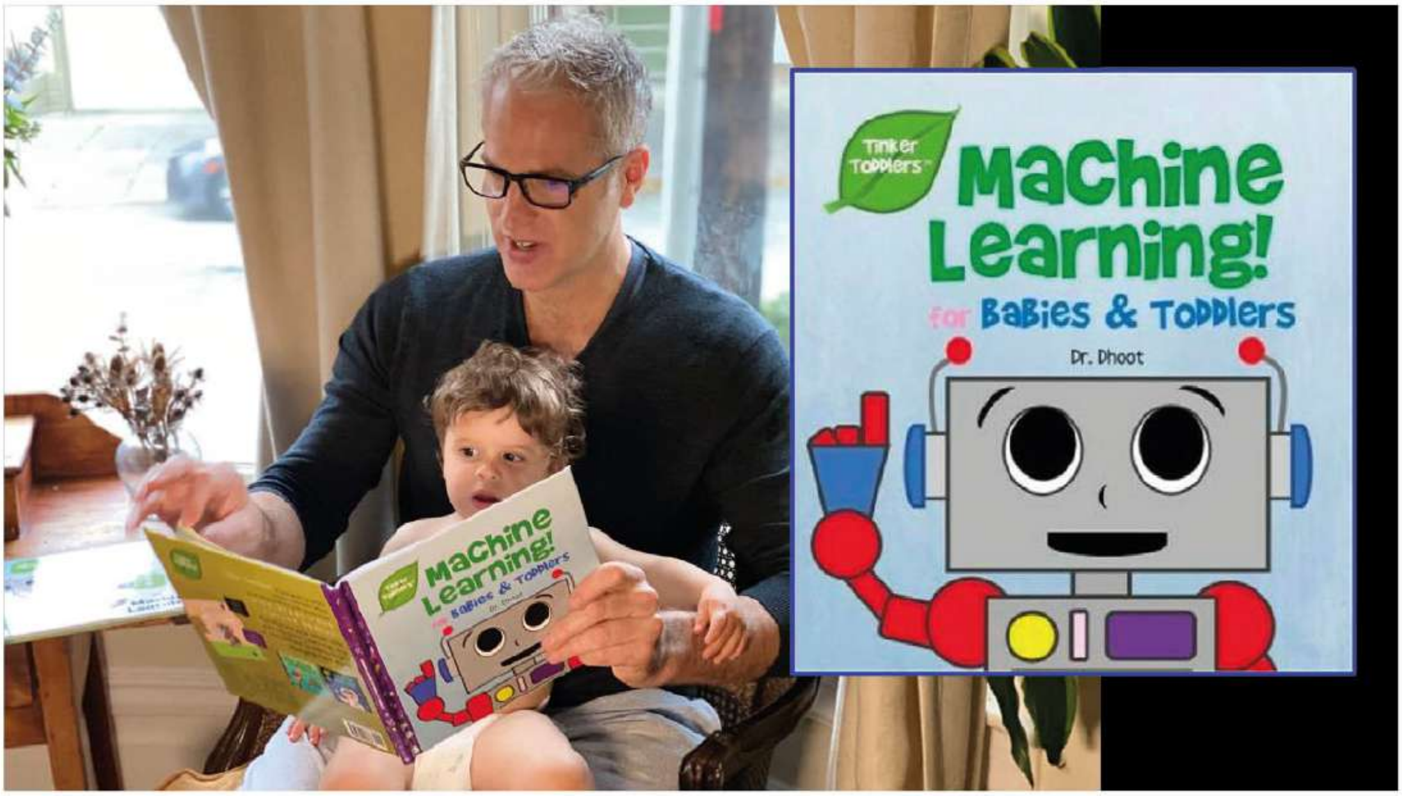

Machine learning:

(data)

Technology that learns from experience ^ to predict the outcome or behavior of each customer, patient, business, vehicle, image, piece of equipment, or other individual unit

… in order to drive better decisions.

This talk is about machine learning in the above practical, applied sense.

A.k.a. *predictive analytics, predictive AI*

Predictive modeling learns from data in order to generate a predictive model. For details on how this works, see Chapter 4 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

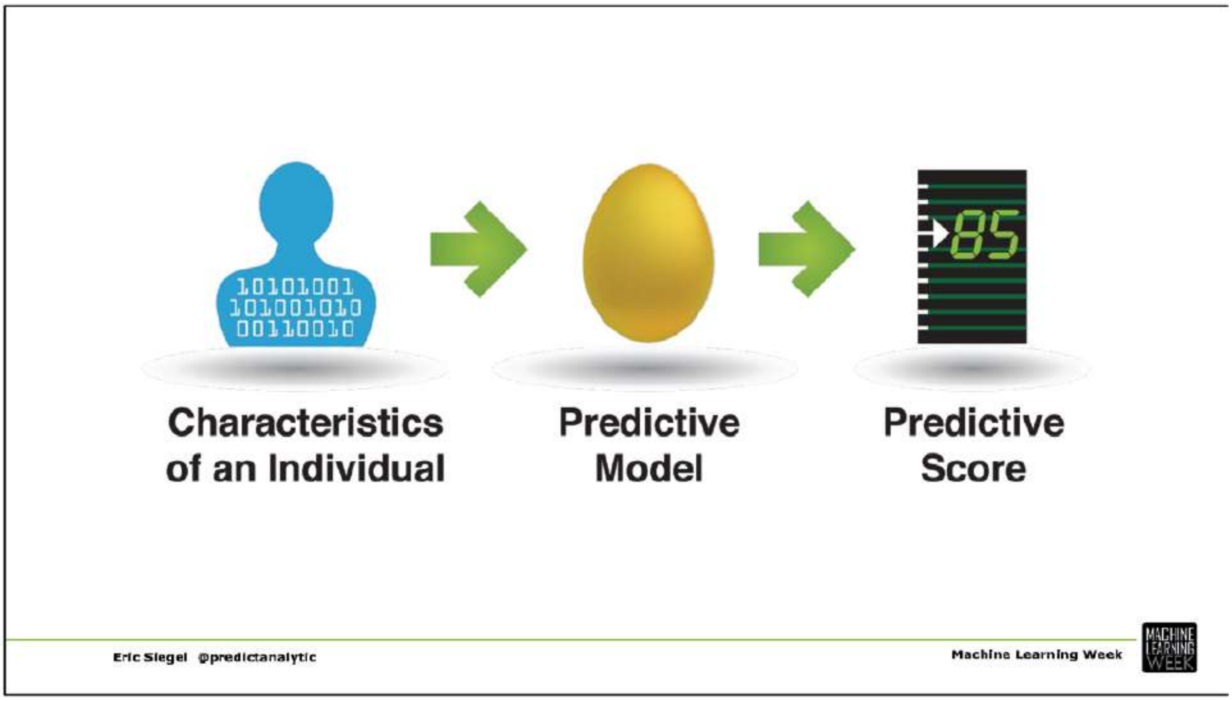

A predictive model generates a predictive score for an individual. For details on how this works, see Chapters 1 and 4 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

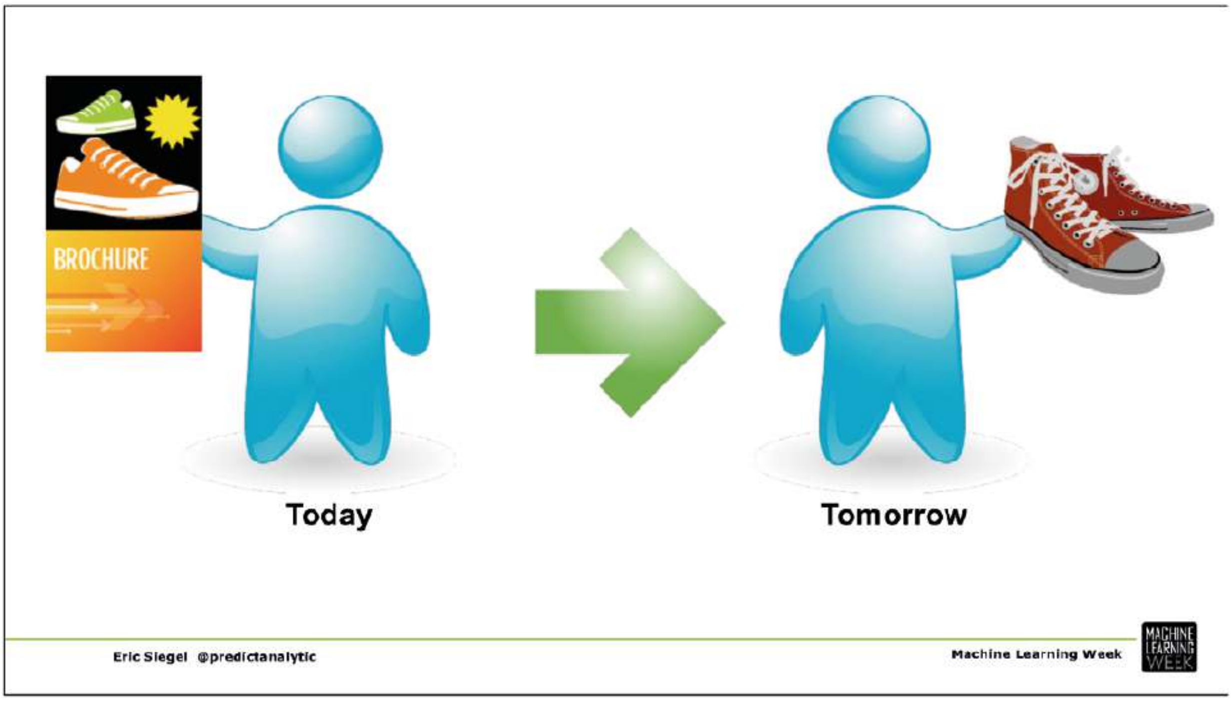

Marketing targets an individual predicted as likely to buy. For details on how this works see the Introduction and Chapter 1 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

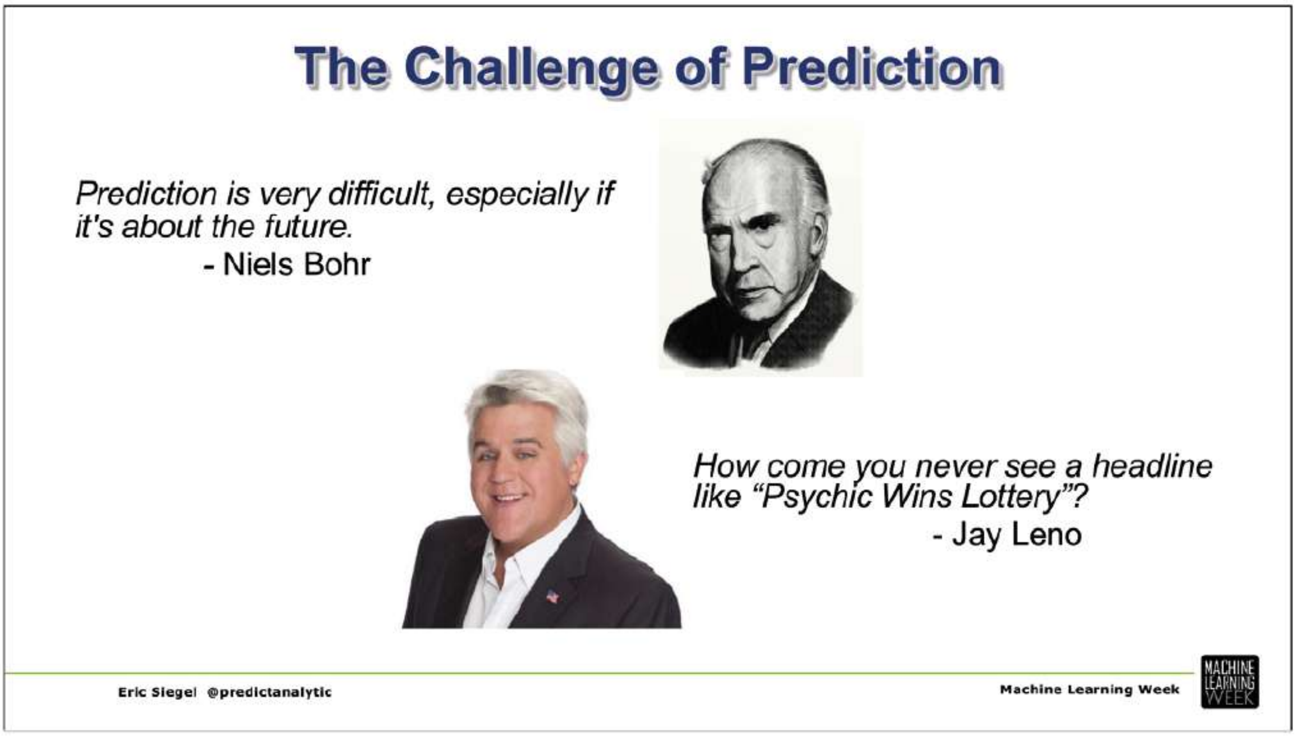

Is prediction an audacious goal? Isn't prediction impossible? For details on how why predictive analytics predicts well enough, see the Introduction and Chapter 1 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (free to access as a PDF on the "Excerpts" page of http://www.thepredictionbook.com).

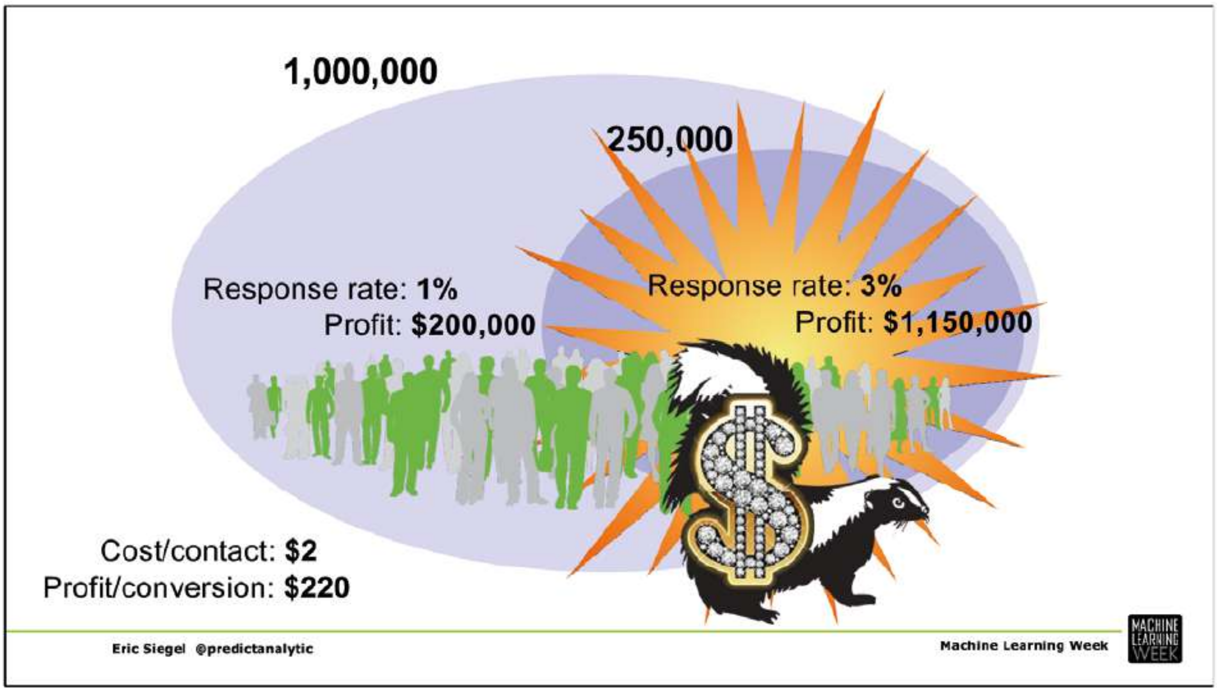

A crummy predictive model delivers big value. It's like a skunk with bling.

here is a spreadsheet detailing the calculations, which you may copy and toy with at will.
https://docs.google.com/spreadsheets/d/1sLp2sGxTZKH0FW4x-ViukfZ-RsCYDrD9B-ReuMZus8M/edit?usp=sharing

Simple arithmetic shows the bottom line profit of direct mail, both in general and then improved by predictively targeting (and only contacting 25% of the list). The less simple part is how the predictive scores are generated for each individual in order to determine exactly who belongs in that 25%. For details on how this works, see Chapter 1 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

The Prediction Effect:

A little prediction goes a long way.

Put another way, predicting better than guessing is often sufficient to generate great value by rendering operations more efficient and effective. For details on how this works, see the Introduction and Chapter 1 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (available for free on the Excerpts page of http://www.thepredictionbook.com).

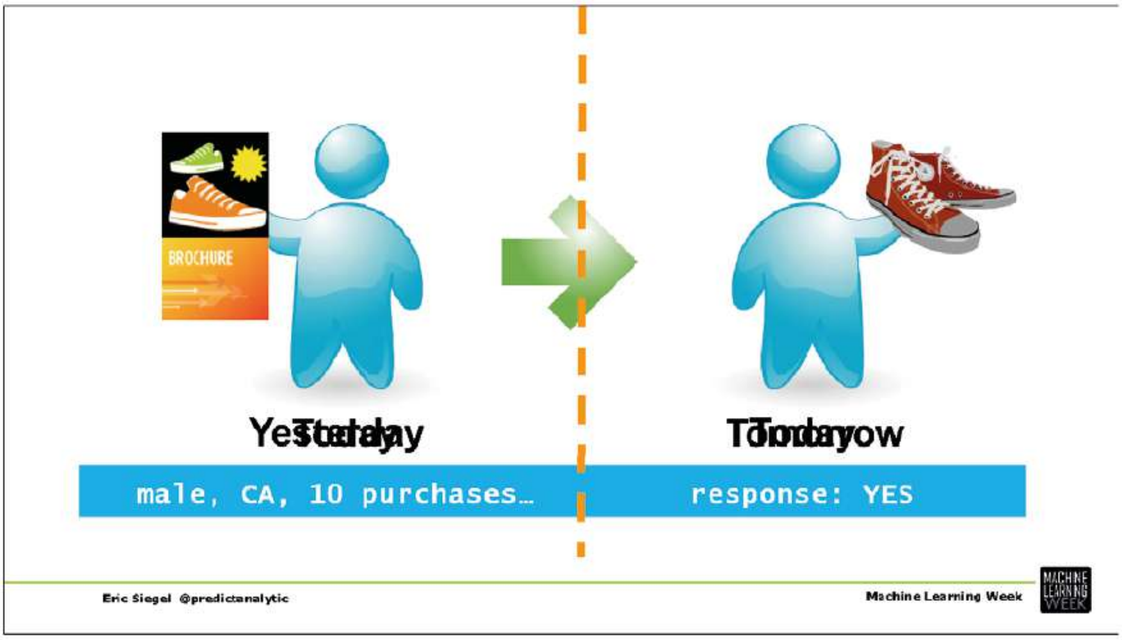

Each row of *training data* corresponds to one individual – first the individual's facts and figures are listed (predictor variables, aka independent variables), and then the target variable (aka dependent variable) – ie, the thing you're trying to predict – is listed.

A table of such rows composes the training data, on which predictive modeling operates.

Wide data has more information about each case

E-commerce	$125	Not-present	$250/day	... ...	Yes
Grocery	$17	Chip	$700/day	... ...	No
Clothing	$275	Swipe	$25/day	... ...	No
Pharmacy	$27	Tap	$150/day	... ...	Yes
Utility	$59	Not-present	$75/day	... ...	No
Airline	$782	Not-present	$35/day	... ...	Yes
Hotel	$1,221	Chip	$100/day	... ...	No
Restaurant	$76	Tap	$40/day	... ...	No
Pharmacy	$32	Swipe	$275/day	... ...	No
Grocery	$112	Tap	$400/day	... ...	No
E-commerce	$43	Not-present	$80/day	... ...	No
Restaurant	$82	Chip	$30/day	... ...	No
Utility	$26	Not-present	$100/day	... ...	No

Long data has many cases

Eric Siegel @p

Learning Week MACHINE LEARNING WEEK

Canadian Tire examples, from "What Does Your Credit-Card Company Know About You?" New York Times, May 12, 2009.
http://www.nytimes.com/2009/05/17/magazine/17credit-t.html

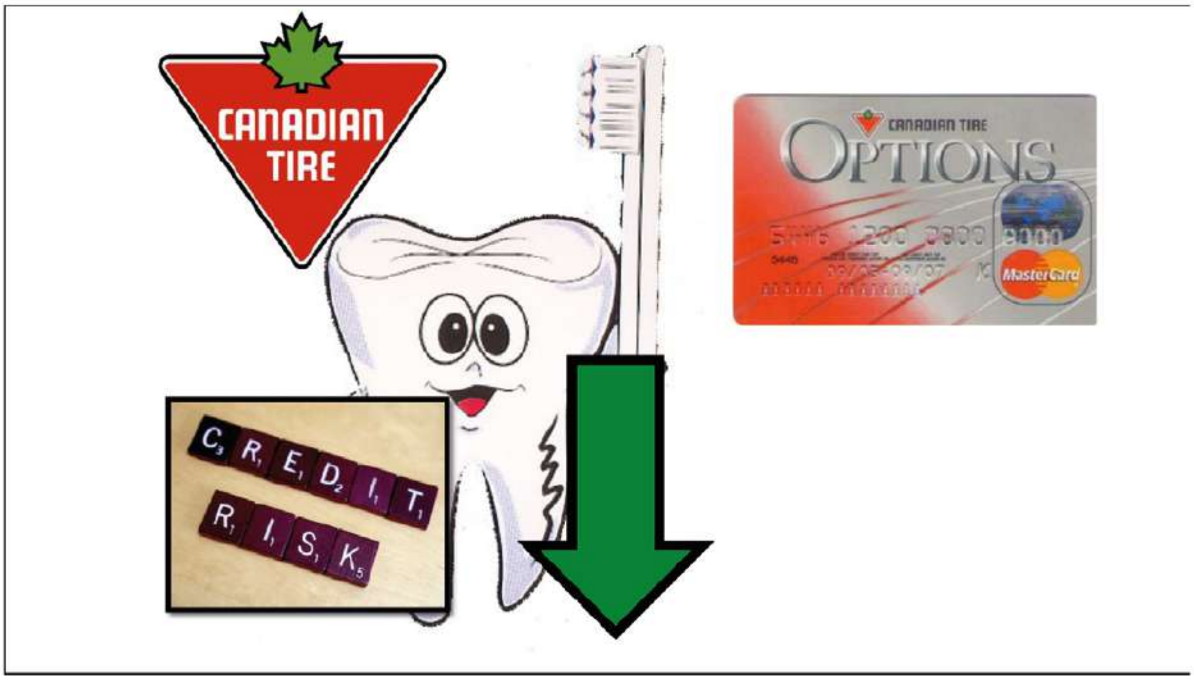

Canadian Tire examples, from "What Does Your Credit-Card Company Know About You?" New York Times, May 12, 2009.
http://www.nytimes.com/2009/05/17/magazine/17credit-t.html

Canadian Tire examples, from "What Does Your Credit-Card Company Know About You?" New York Times, May 12, 2009.
http://www.nytimes.com/2009/05/17/magazine/17credit-t.html

People who “like” curly fries on Facebook are “more intelligent.”

Reference for most examples in this presentation are in the Notes PDF for Eric Siegel's book, "Predictive Analytics." For each example's reference/citation, search by organization name within the book's Notes PDF, available at www.PredictiveNotes.com

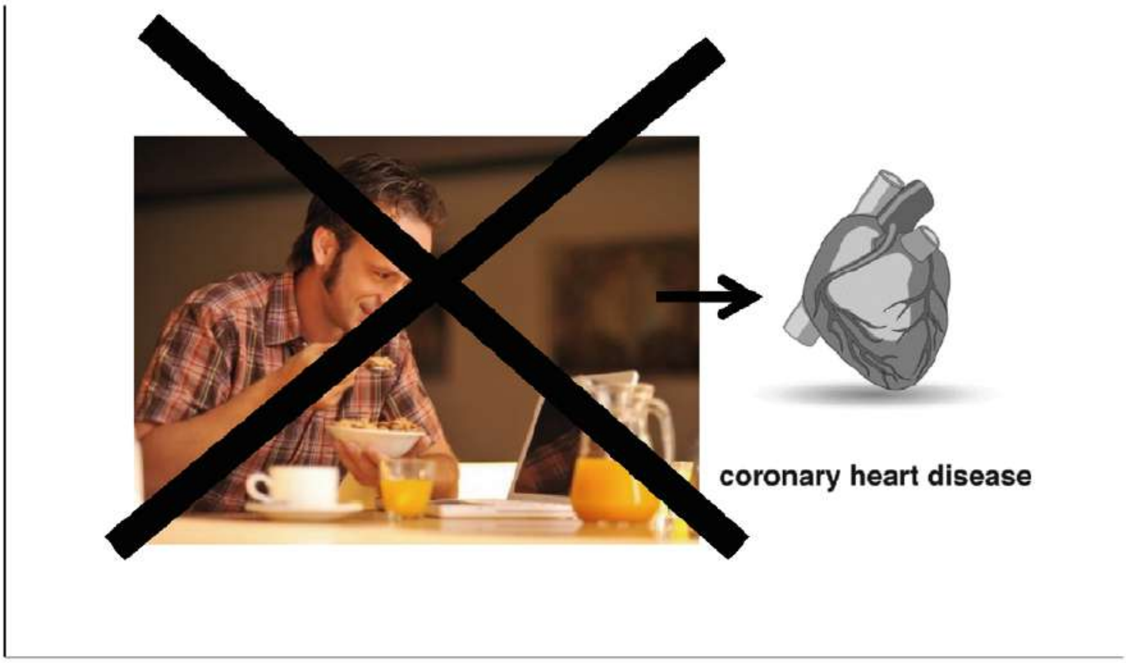

Men who skipped breakfast showed a 27% higher risk of coronary heart disease.

"Prospective Study of Breakfast Eating and Incident Coronary Heart Disease in a Cohort of Male US Health Professionals," by Cahill et al.
http://circ.ahajournals.org/content/128/4/337.full.pdf

The Data Effect:

Data is always predictive.

Eric Siegel @predictanalytic
Machine Learning Week

Correlation does not entail causation. For more information, see Chapter 3 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

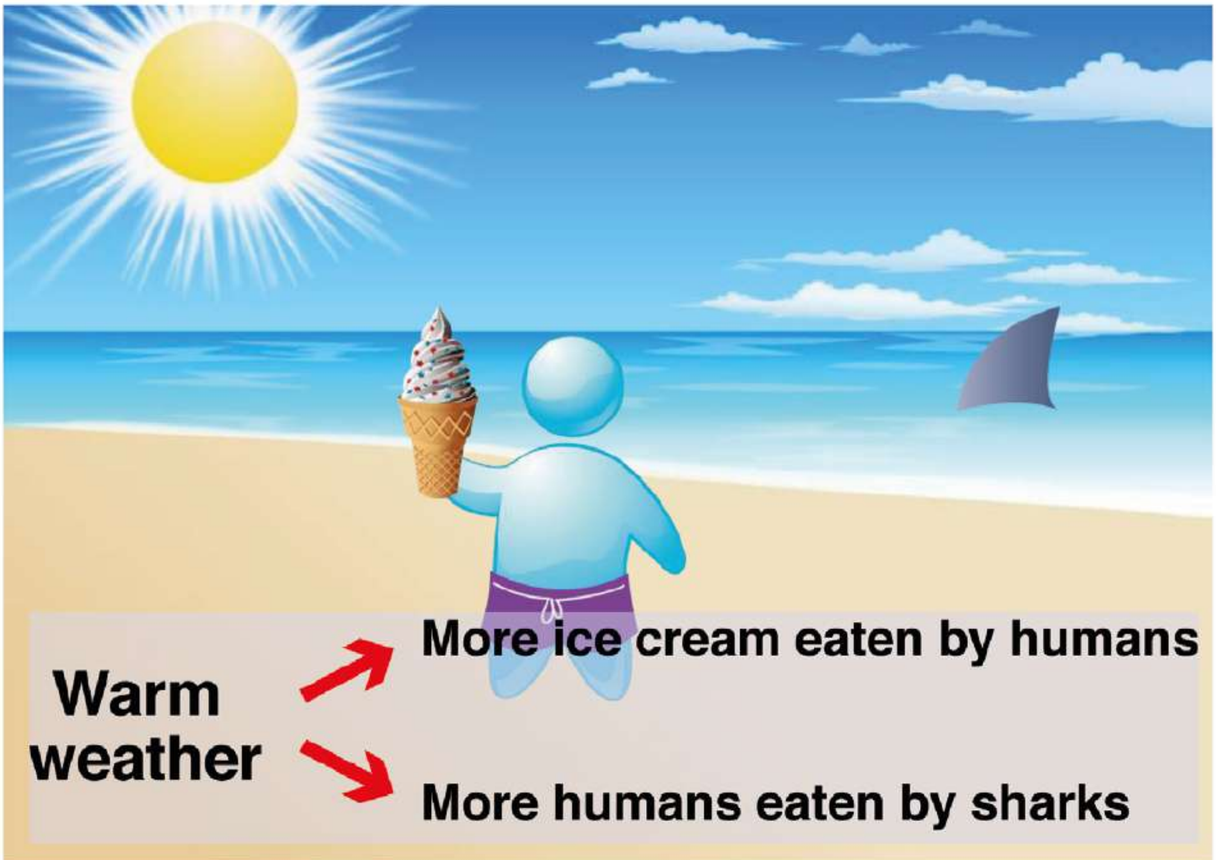

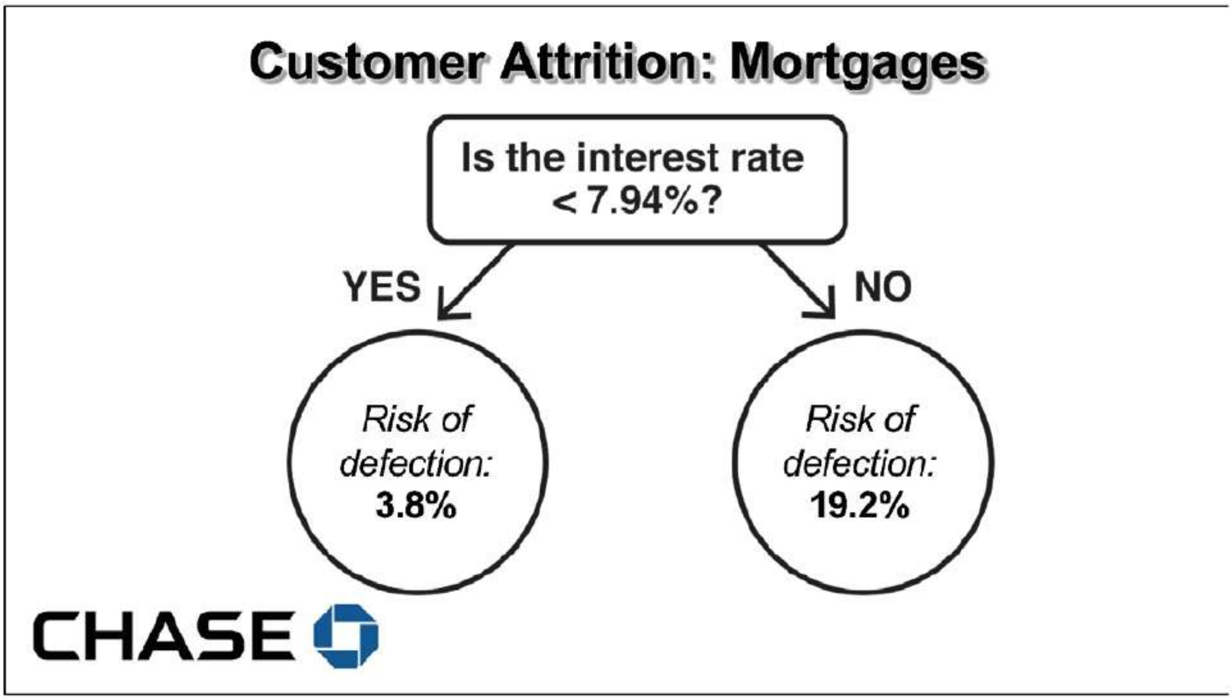

Reference for most examples/case studies in this presentation are in the Notes PDF for Eric Siegel's book, "Predictive Analytics." For each example's reference/citation, search by organization name within the book's Notes PDF, available at www.PredictiveNotes.com

IF:

the mortgage is greater than or equal to $67,751 and less than $182,926

AND:

the interest rate is greater than or equal to 8.69 percent

AND:

the loan-to-value ratio is less than 87.4 percent

THEN:

the probability of prepayment is 25.6 percent.

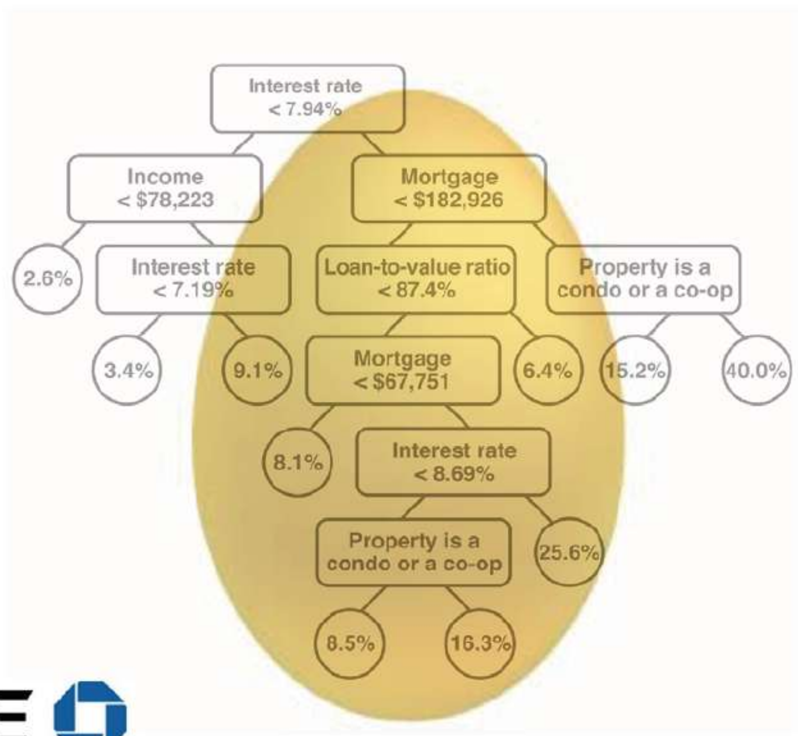

CHASE

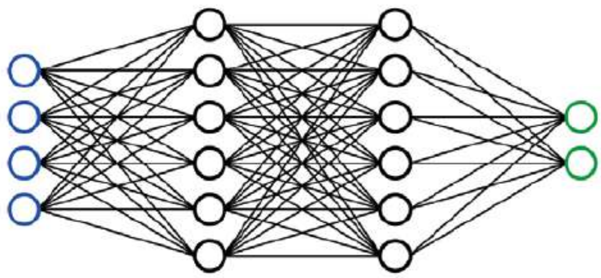

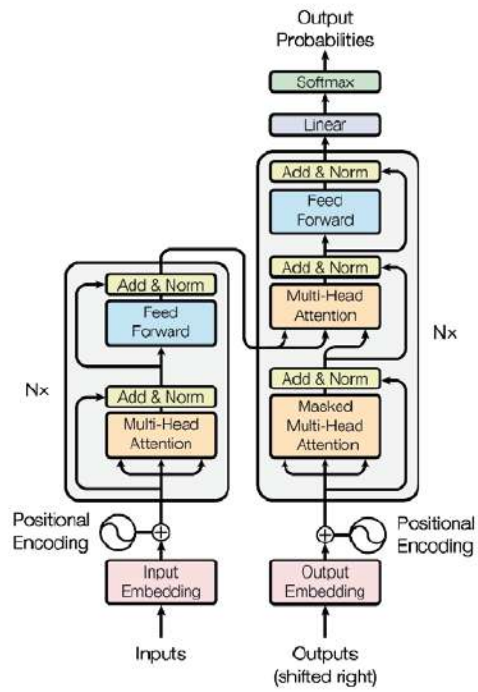
Output Probabilities
Softmax
Linear
Add & Norm
Feed Forward
Add & Norm
Multi-Head Attention
Add & Norm
Masked Multi-Head Attention
Nx
Add & Norm
Feed Forward
Add & Norm
Multi-Head Attention
Nx
Positional Encoding
Input Embedding
Inputs
Positional Encoding
Output Embedding
Outputs (shifted right)

...and many more, such as Cox Communications, FedEx, Sprint, etc. - see the book "Predictive Analytics" (www.thepredictionbook) for many case studies, including a central compendium of 147 mini-case studies, of which 37 are examples in marketing applications of predictive analytics.

Reference for most examples/case studies in this presentation are in the Notes PDF for Eric Siegel's book, "Predictive Analytics." For each example's reference/citation, search by organization name within the book's Notes PDF, available at www.PredictiveNotes.com

PREMIER Bankcard also lowered delinquency to increase net by over $10 million

More information about First Tennessee Bank and other case studies are available at http://tinyurl.com/PAExamples

Dan Marks, First Tennessee Bank, "First Tennessee Bank: Analytics Drives Higher ROI from Marketing Programs," IBM.com, March 9, 2011. www.ibm.com/smarterplanet/us/en/leadership/firsttenbank/assets/pdf/IBM-firstTennBank.pdf

HEALTHCARE

Refined their outreach to increase response by 38%

SMALL BUSINESS

Earned 40% response rate amongst lapsed buyers

NON-PROFIT

Discovered how to earn $650,000 in donations

...and many more, such as Cox Communications, FedEx, Sprint, etc. - see the book "Predictive Analytics" (www.thepredictionbook) for many case studies, including a central compendium of 147 mini-case studies, of which 37 are examples in marketing applications of predictive analytics.

PAYROLL SERVICES

Decreased by 40% the number of phone calls needed in order to book each sales meeting

UNIVERSITY

Scores prospective students

HIGH TECH

Doubled number of leads per phone call

Reference for most examples/case studies in this presentation are in the Notes PDF for Eric Siegel's book, "Predictive Analytics." For each example's reference/citation, search by organization name within the book's Notes PDF, available at www.PredictiveNotes.com

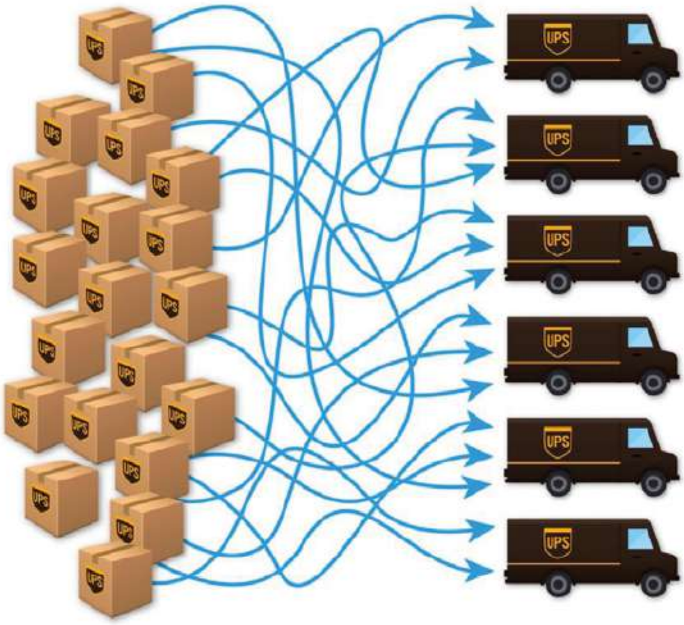

16 million packages/day
1000 shipping centers
55 trucks
300 packages

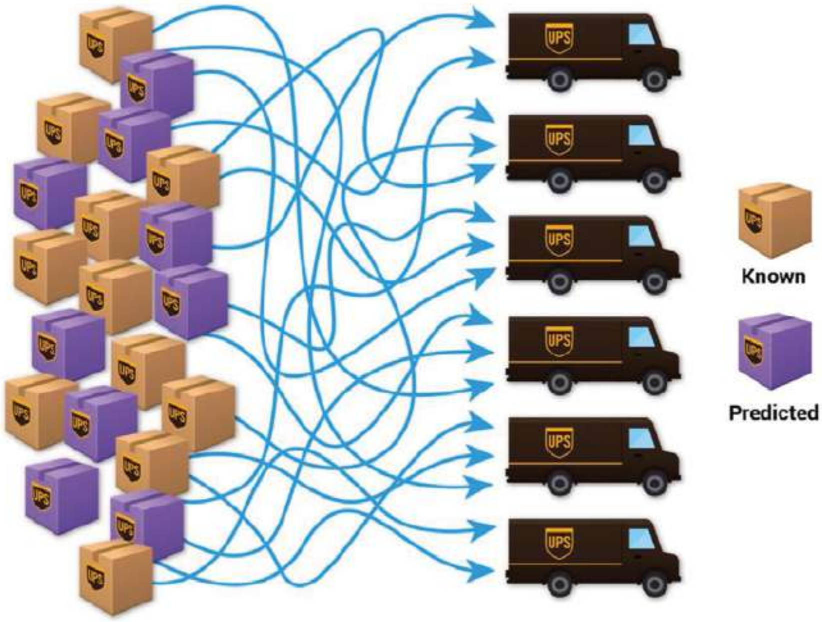

Application	What's predicted	What's done about it
Delivery prediction *to plan for more efficient delivery*	*Will the address receive a package delivery?*	Plan the delivery truck assignments of predicted packages alongside known ones.

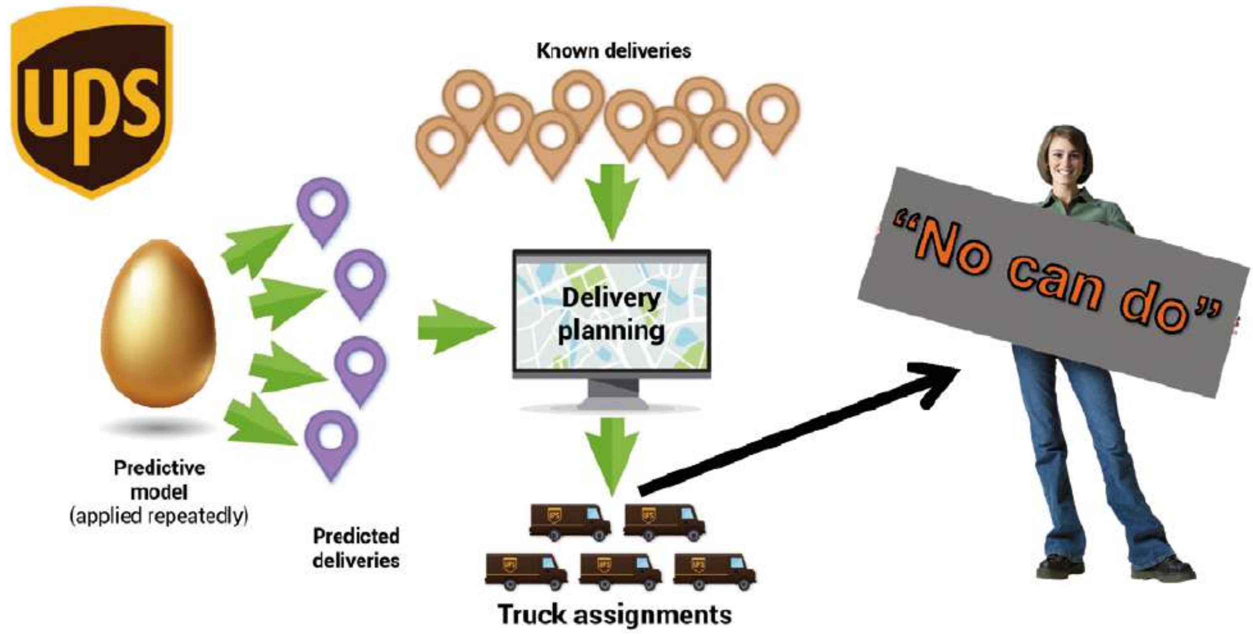

Package Flow Technology and ORION

Overall annual savings:

- 185 million miles
- $350+ million
- 8 million gallons of fuel
- 185,000 metric tons of emissions

Annual savings due specifically to delivery-prediction (estimated):

- 18.5 million miles
- $35+ million

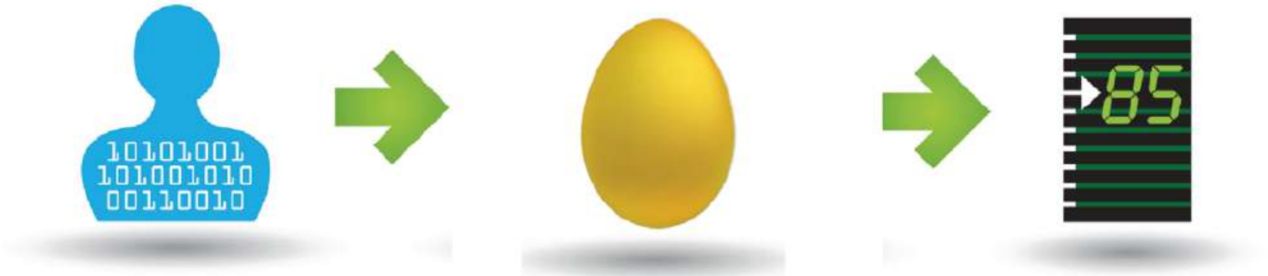

Characteristics of an Individual **Predictive Model** **Probability**

A probability calculator

Predict:

To put a probability on an unknown.

- *The future*
- *Detecting / diagnosing*

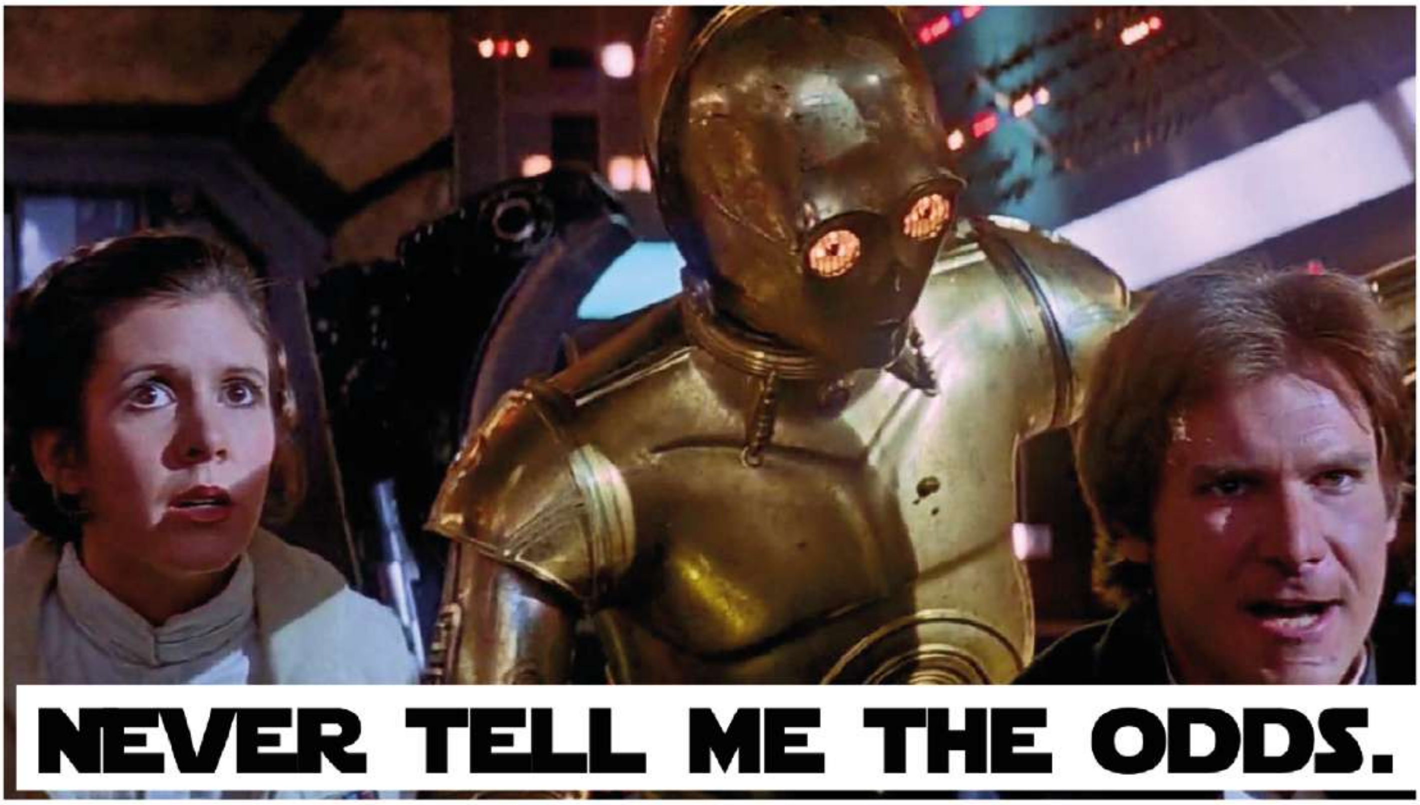

C3P0: The possibility of successfully navigating... 3,720 to 1.

Han Solo: Never tell me the odds.

Eric Siegel @predictanalytic

Machine Learning Week

MACHINE LEARNING WEEK

66

"Moneyball" celebrates math and yet epitomizes the glossing over of math.

1) What's predicted

2) How well

3) What's done about it

Eric Siegel @predictanalytic Machine Learning Week

To transfer business expertise into to technical requirements is to get business professionals ramped up on:

- What's predicted, How well, What do about it

- What the model does, with what success, how use it

- Dependent variable, metrics, deployment: how predictive probabilities actively change business operations in order to improve them. Stakeholders must understand change in order to manage it.

Dependent variable: 1) What's predicted

Metrics: 2) How well

Deployment: 3) What's done about it

Eric Siegel @predictanalytic Machine Learning Week

Application	What's predicted	What's done about it
Response modeling *to increase the marketing response rate*	*Will the customer buy if contacted?*	Mail a brochure to those likely to buy.

Application	What's predicted	What's done about it
Response modeling *to increase the marketing response rate*	***Will the customer buy if contacted?***	Mail a brochure to those likely to buy.

Application	What's predicted	What's done about it
Response modeling *to increase the marketing response rate*	*If sent a brochure, will the customer buy within 13 business days with a purchase value of at least $125 after shipping and not return the product for a refund within 45 days?*	Mail a brochure to those likely to buy.

Dependent variable: 1) What's predicted

Metrics: 2) How well

Deployment: 3) What's done about it

Benchmarks for training and deployment

Technical performance and business performance

The Accuracy Fallacy

THE SPECTATOR: *Linguistic Analysis Can Accurately Predict Psychosis*

THE DAILY MAIL: *AI-Powered Scans Can Identify People at Risk of a Fatal Heart Attack Almost a Decade in Advance*

THE NEXT WEB: *This Scary AI Has Learned How to Pick Out Criminals by Their Faces*

Eric Siegel @predictanalytic Machine Learning Week

For more on the accuracy fallacy, see: "The Media's Coverage of AI is Bogus", by Eric Siegel, Scientific American, November 2019.
https://blogs.scientificamerican.com/observations/the-medias-coverage-of-ai-is-bogus/

Accuracy:

The proportion of cases a predictive model predicts correctly.

Eric Siegel @predictanalytic Machine Learning Week

Accuracy: The proportion of cases a predictive model predicts correctly, that is, how often the model is correct. Accuracy does not differentiate between how often the model is correct for positive and negative examples. This means that, for example, a model with high accuracy could in fact get none of the positive cases correct, if positive examples are relatively rare.

- Predicting bestselling books
- Spotting legal issues in non-disclosure agreements
- IBM's claim that they can predict which employees will quit with 95% accuracy
- Classifying which news headlines are "clickbait"
- Detecting fraudulent dating profile scams
- Spotting cyberbullies
- Predicting earthquake aftershocks
- Predicting the need for first responders after an earthquake
- Detecting diseases in banana crops
- Distinguishing high and low-quality embryos for in vitro fertilization
- Predicting heart attacks
- Predicting heart issues by eye scan
- Detecting anxiety and depression in children
- Diagnosing brain tumors from medical images
- Detecting brain tumors with a new blood test
- Predicting the development of Alzheimer's

Eric Siegel @predictanalytic | Machine Learning Week

For our latest, updated list of all such accuracy fallacy examples, click through from this article to the blog article it links:

https://blogs.scientificamerican.com/observations/the-medias-coverage-of-ai-is-bogus/

... to be specific, this one:

https://www.predictiveanalyticsworld.com/blog/more-examples-of-the-medias-bogus-claims-of-high-accuracy-ai/

https://www.predictiveanalyticsworld.com/machinelearningtimes/accuracy-fallacy-the-medias-coverage-of-ai-is-bogus/10652/

Original (shorter version and behind paywall):
https://blogs.scientificamerican.com/observations/the-medias-coverage-of-ai-is-bogus/

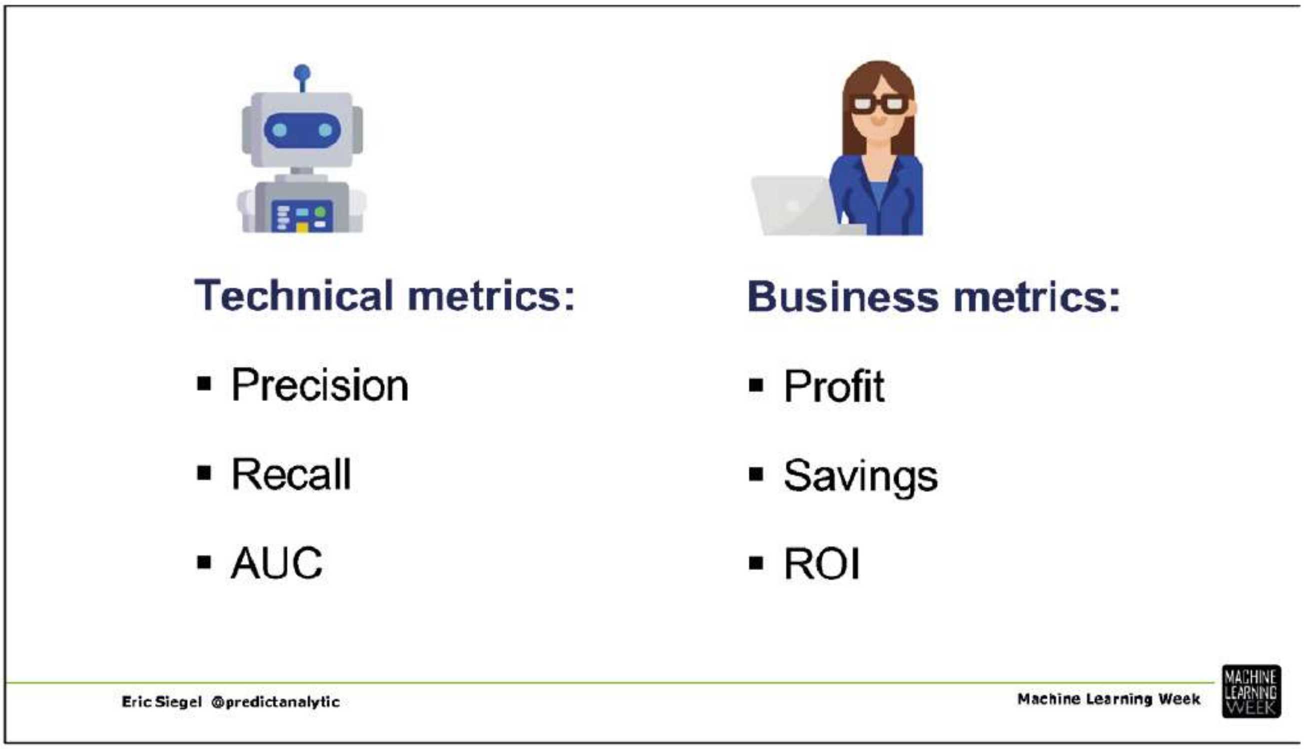

Eric Siegel has co-founded a startup to move from technical ML metrics to business metrics: https://www.Gooder.ai

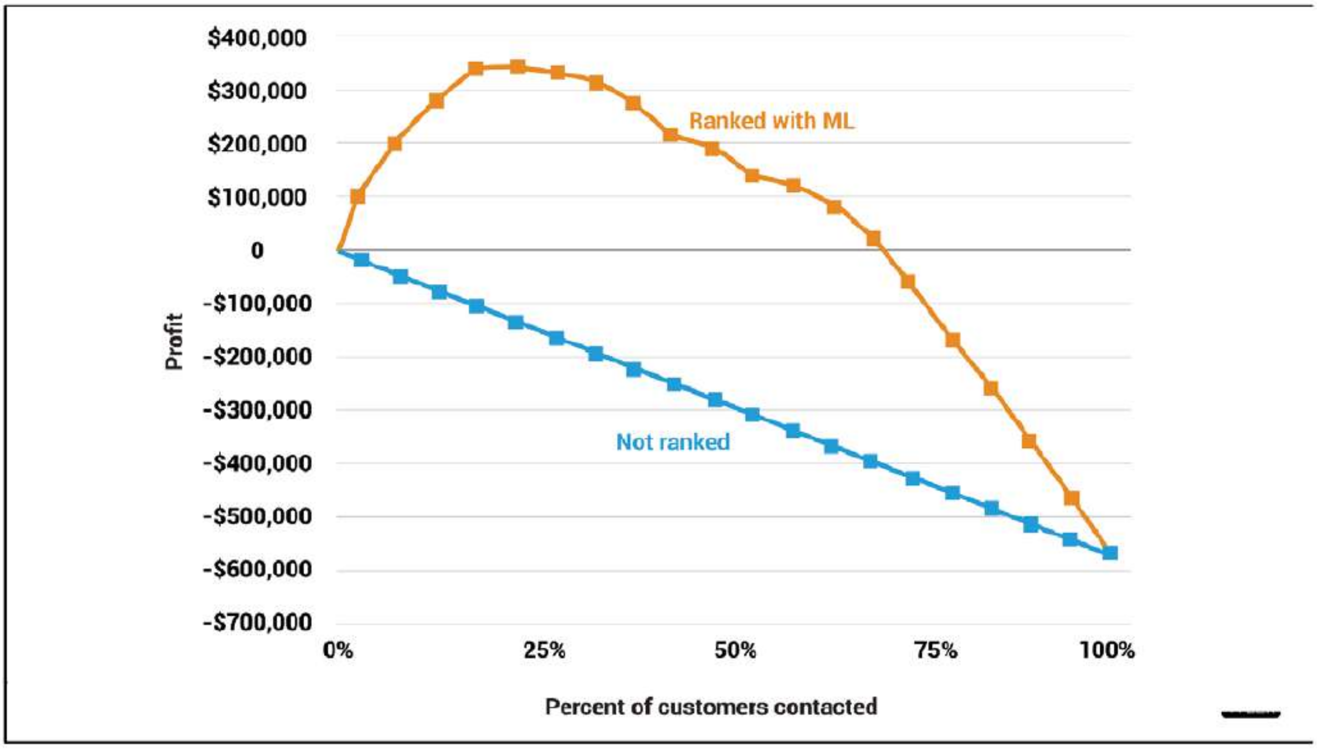

The Cost of Fraud

100,000 cards
1,000 transactions each
100 million transactions per year

Fraud: 0.1%
100,000 fraudulent transactions
Cost per fraudulent transaction: $500

Annual loss:
100,000 * $500
$50 million

Savings with Fraud Detection

Blocked: 2 per 1,000
Blocked: 200,000 per year
Percent fraud: 30%
Blocked fraud: 60,000

False positives: 140,000
False negatives: 40,000

Cost of FP: $100
Cost of FN: $500

Total cost: **$34 million (versus $50 million)**

Cost of a FP	$100	*Average cost of each false positive - inconvenience to cardholder*					
Cost of a FN	$500	*Average cost of each false negative - fraud that goes undetected*					
Fraud rate	0.1%						
Num transactions	100,000,000	of which	100,000	are fraudulent	which cost	$50,000,000	if none are detected
Transactions held	0.2%	as potentially fraudulent, which is		200,000	transactions		
Lift	300	... so the held transactions are		30.00%	positive, that is	60,000	detected cases of fraud
	FP count	FP cost	FN count	FN cost	Total cost		
Cost with model	140,000	$14,000,000	40,000	$20,000,000	$34,000,000		
					$16,000,000	**savings compared to no fraud detection**	

bit.ly/3LrsIHT

To try this spreadsheet yourself, go to:

https://docs.google.com/spreadsheets/d/12dKRw4TYmqwbAkp3VRTwt5gUe1rnLweg0EHfRBQHcIc/

Instructions: First, make a copy of this spreadsheet (do not request edit access). Then change the light yellow cells as you wish. The blue cells are the bottom line.

Consider a bank that has issued 100,000 credit cards and each sees an average of 1,000 transactions per year, with one in a thousand being fraudulent.

I consider FICO Falcon to be one of the world's most successful and widely impactful commercial deployment of ML. It screens all the transactions for 2.6 billion payment cards worldwide. That's two thirds of the world's cards, including about 90% of those in the U.S. and the U.K. 17 of the top 20 international credit card issuers, all of the U.S.'s 100 largest credit-card issuers, and 95 of the U.S.'s top 100 financial institutions use Falcon.

Since its introduction, Falcon has reduced card fraud losses by more than 70% in the U.S. With the U.S. currently suffering around $10 billion in annual fraud losses, that reduction is saving us more than $20 billion per year.

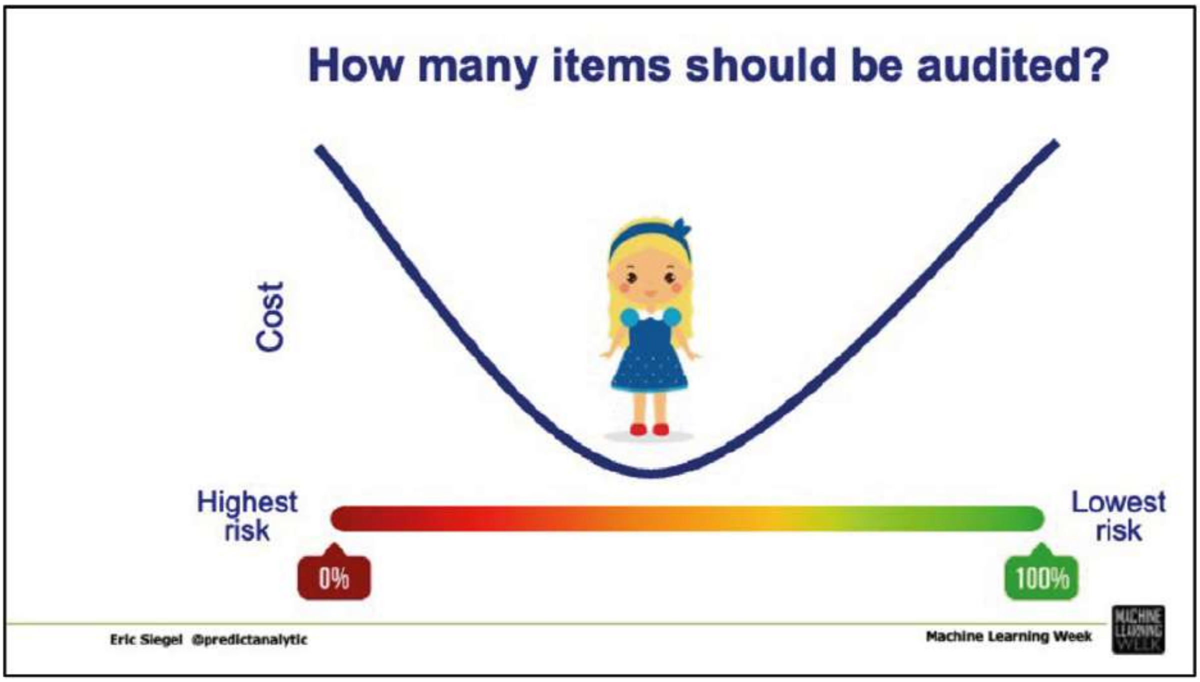

You can't manually audit every single post/transaction -- that's too expensive.
Where do you draw the line? How do you strike that balance?
There's no existing solution for navigating that decision.

Will show a $400k weekly cost savings in the demo – using GPT-3.
And a $1.7M profit improvement for marketing.

The same concept applies to decide how many to:
- Audit (fraud and misinformation)
- Contact (marketing)
- Approve (loan applications)
- Treat or test (healthcare)

The model doesn't autonomously or unilaterally tell you what to do at each case. It depends on business context, strategic decisions, and how good the model is. This is where we navigate whether and *how* to use a model. This visual represents the options and helps us navigate the deployment decisions. Yet tools almost never show this view.

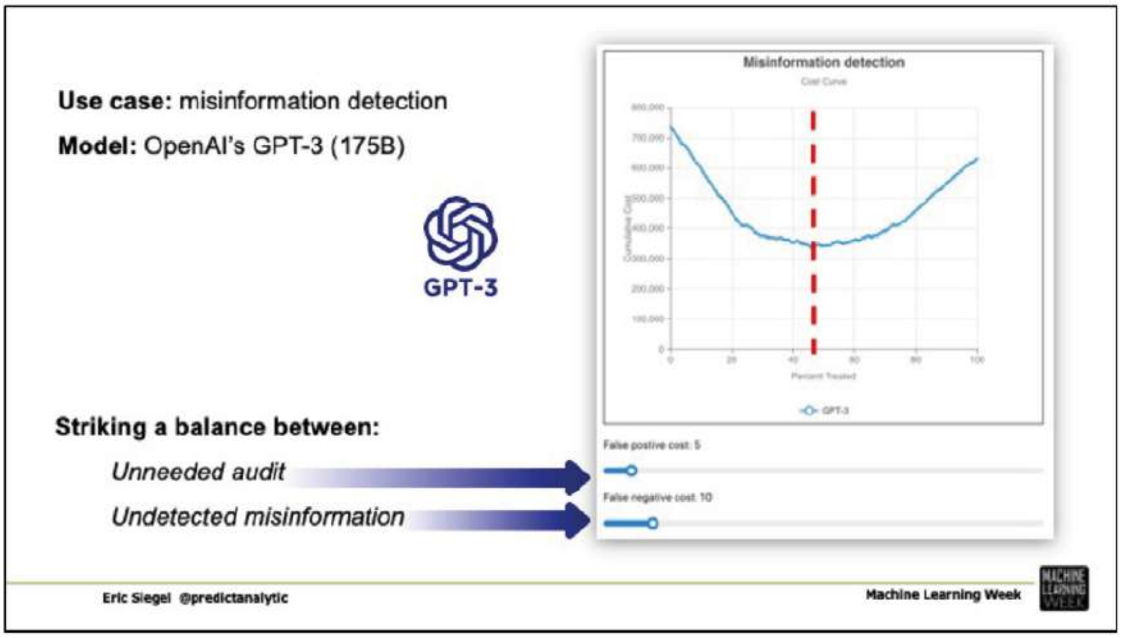

Gooder AI goes beyond technical metrics like AUC to focus on business metrics like profit and costs.

Dataset: boolQ 1000 (5000)
3155 yes
1845 no ("misinformation")

Misinformation detection
OpenAI / davinci (175B) == GPT-3
https://crfm.stanford.edu/helm/latest/?runs=1&runSpec=boolq%3Amodel%3Dopenai_davinci%2Cdata_augmentation%3Dcanonical

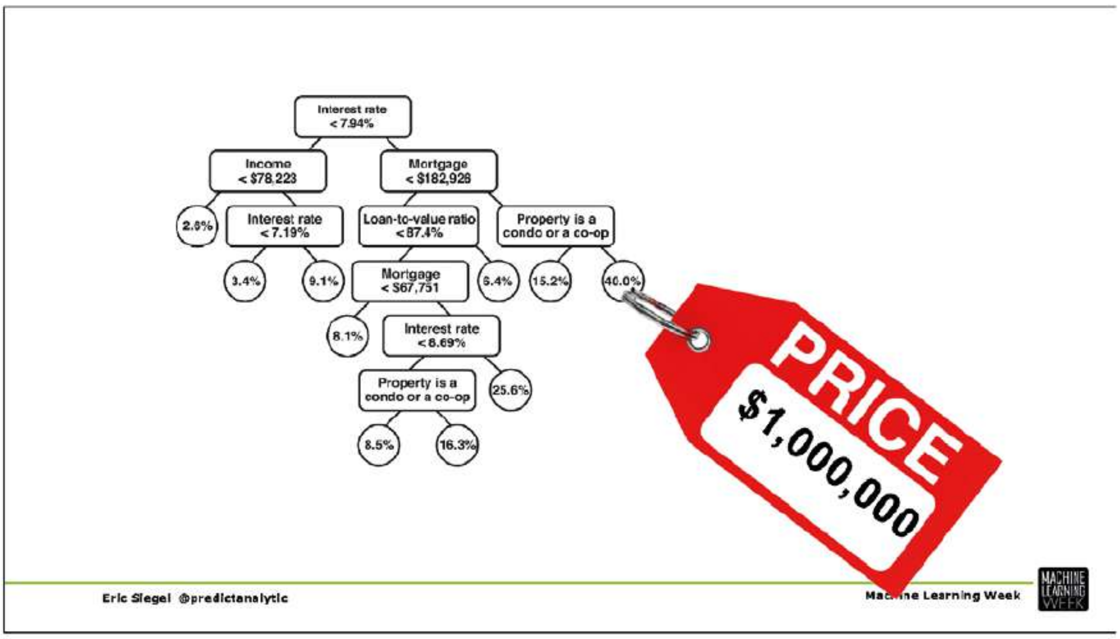

You can't just say, "This model is worth a million bucks" -- its value depends on how you use.

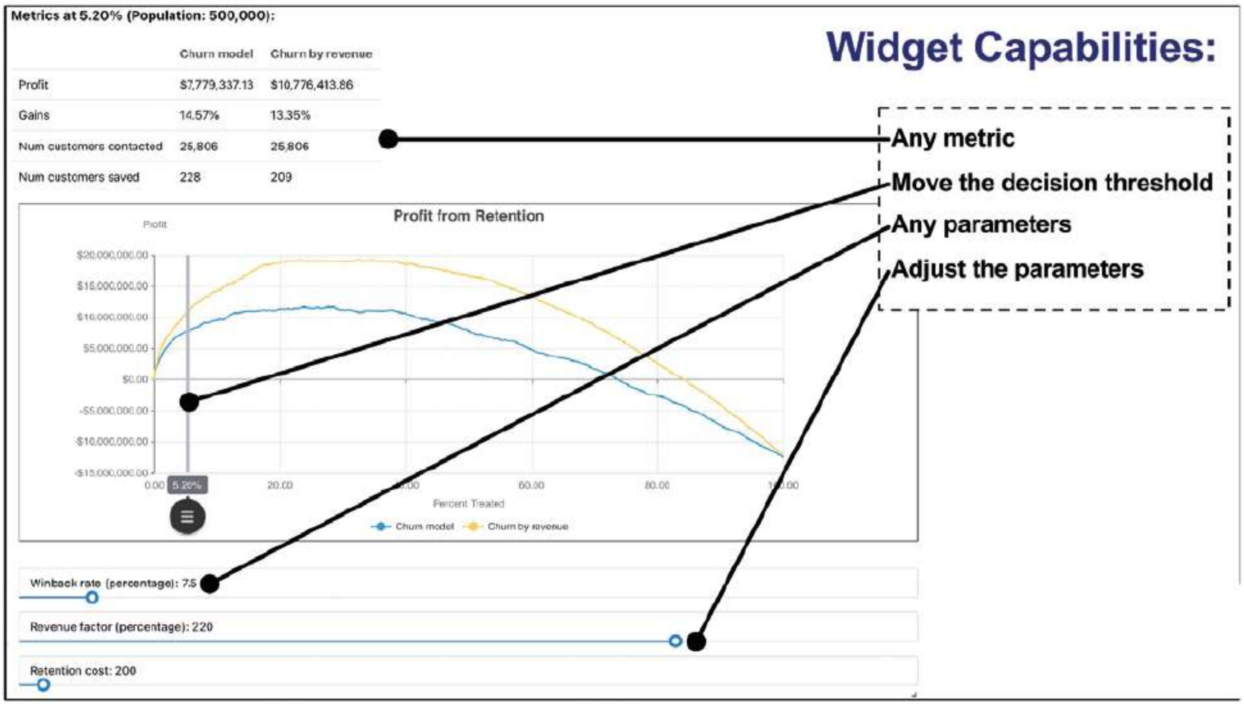

Move the decision threshold:
View the available tradeoffs between competing metrics

Adjust the parameters:
Visualize the effects of changing parameters

A key capability offered by Gooder AI is to contextualize – and visualize the effect of – setting the confidence threshold (aka decision threshold). Doing so isn't the "rocket science" part of an ML project – it's the part that aims the rocket. In making this choice, every predictive AI project must wrestle with striking practical tradeoffs between competing objectives.

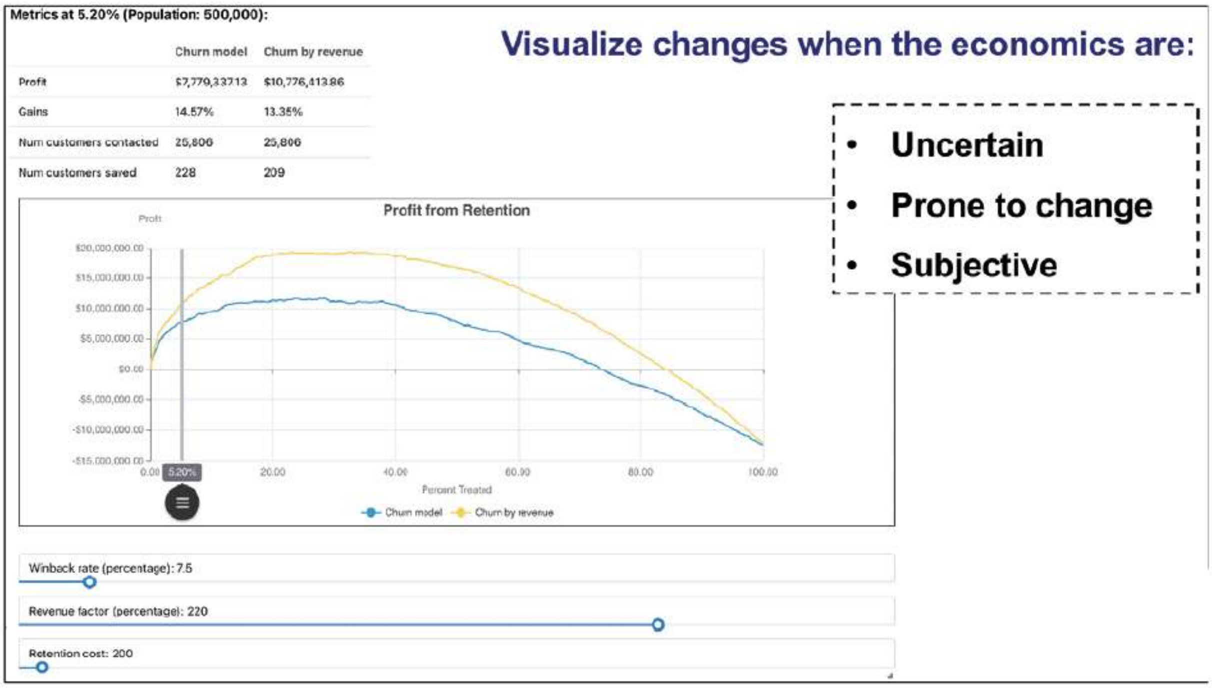

You must visualize the effects of parameter changes when the economics are:

Not yet known
Estimated
Prone to change
- Change cost of contact (black and white)
- Change profit margin (new vendor)
- Mitigate fraud FP cost
- Tactical changes, e.g., switching vendors

Subjective
- Medical diagnosis
- Misinformation detection

GOODER AI

Payment Card Fraud Detection

Metrics at 7.00% (Population: 10,000,000):

	Decision tree 1	Decision tree 2	Logistic
Gains	66.33%	61.28%	24.53%
Loss	$10,908,996.54	$12,153,622.11	$20,714,600.20
Savings	$8,880,346.80	$7,906,661.70	-$977,749.18
False positives	469,367	483,320	610,187
Blocked fraud	224,845	209,300	82,501

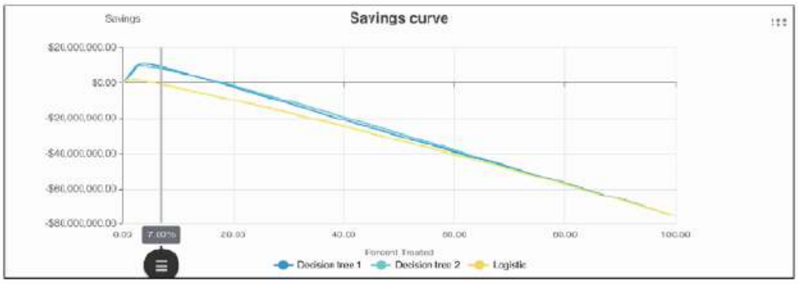

FP-cost: 10

FN-cost: 60

MITSloan Management Review MENU SEARCH SUBSCRIBE

What Leaders Should Know About Measuring AI Project Value

Most AI/machine learning projects report only on technical metrics that don't tell leaders how much business value could be delivered. To prevent project failures, press for business metrics instead.

Eric Siegel • February 07, 2024 Reading Time: 11 min

SUBSCRIBE PERMISSIONS AND PDF SHARE

https://sloanreview.mit.edu/article/what-leaders-should-know-about-measuring-ai-project-value/

ML Lifecycle

Data → Train → **Evaluate** → Deploy → Monitor

GOODER AI

New category: *ML value-capture*

Gooder AI maximizes the value of machine learning by testing and visualizing its business performance.

bizML

1) Establish the **deployment goal**
2) Establish the **prediction goal**
3) Establish the **metrics**
4) Prepare the data
5) Train the model
6) Deploy the model

Eric Siegel @predictanalytic Machine Learning Week

BizML: The strategic playbook for machine learning deployment

1) Establish the deployment goal (value)

Define the business value proposition: how ML will affect operations in order to improve them.

2) Establish the prediction goal (target)

Define what the ML model will predict for each individual case.

3) Establish the evaluation metrics (performance)

Determine the salient benchmarks to track during both model training and model deployment and determine what performance level must be achieved for the project to be considered a success.

4) Prepare the data (fuel)

Define what the training data must look like and get it into that form.

5) Train the model (algorithm)

Generate a predictive model from the data.

6) Deploy the model (launch)

Use the model to render predictive scores and then use those scores to improve business operations.

After step 6: Maintain the model (upkeep)

Monitor and periodically refresh the model as an ongoing process.

Key execution strategy:

All steps require deep collaboration with business stakeholders.
Business stakeholders must hold a semi-technical understanding of ML.
The steps are not executed linearly – backtracking prevails.

From *The AI Playbook* by Eric Siegel

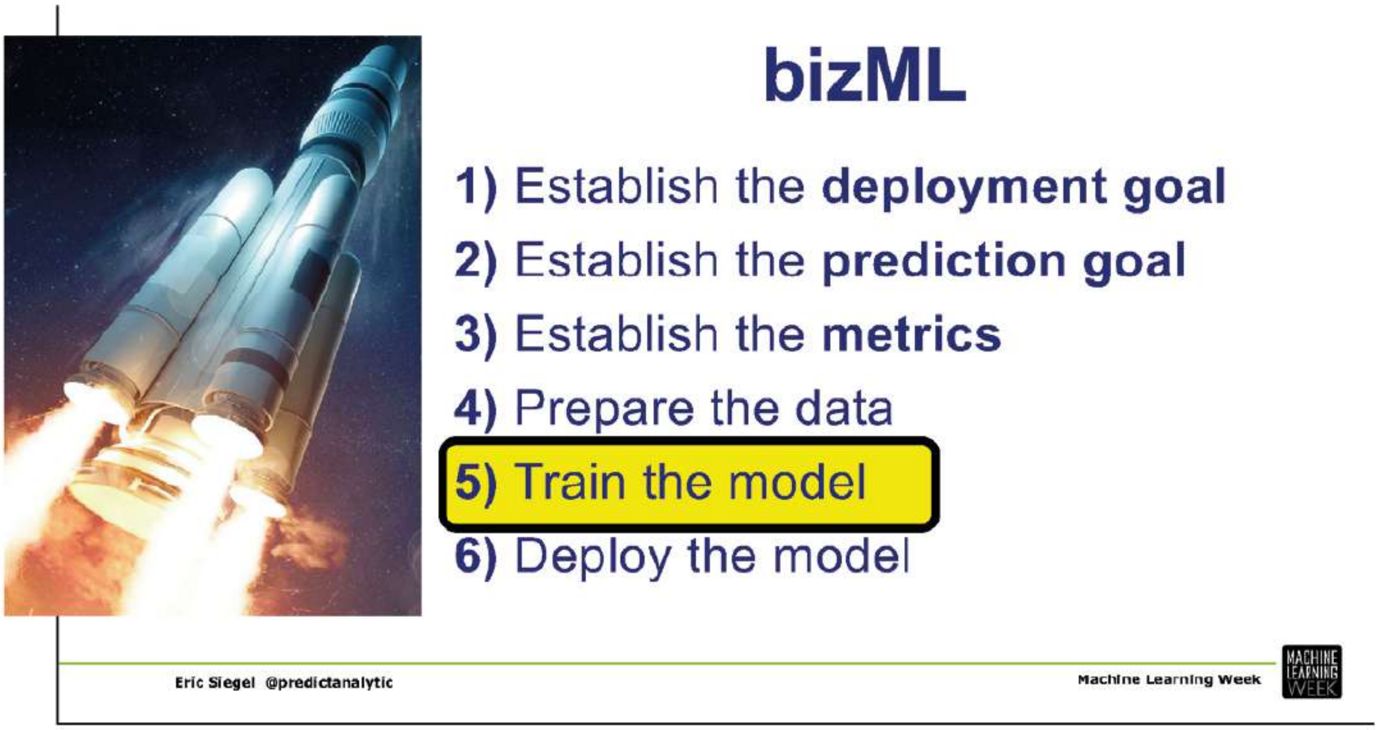

BizML: The strategic playbook for machine learning deployment

1) Establish the deployment goal (value)

Define the business value proposition: how ML will affect operations in order to improve them.

2) Establish the prediction goal (target)

Define what the ML model will predict for each individual case.

3) Establish the evaluation metrics (performance)

Determine the salient benchmarks to track during both model training and model deployment and determine what performance level must be achieved for the project to be considered a success.

4) Prepare the data (fuel)

Define what the training data must look like and get it into that form.

5) Train the model (algorithm)

Generate a predictive model from the data.

6) Deploy the model (launch)

Use the model to render predictive scores and then use those scores to improve business operations.

After step 6: Maintain the model (upkeep)

Monitor and periodically refresh the model as an ongoing process.

Key execution strategy:

All steps require deep collaboration with business stakeholders.
Business stakeholders must hold a semi-technical understanding of ML.
The steps are not executed linearly – backtracking prevails.

From *The AI Playbook* by Eric Siegel

As a result, nobody is connected the hose to the faucet.

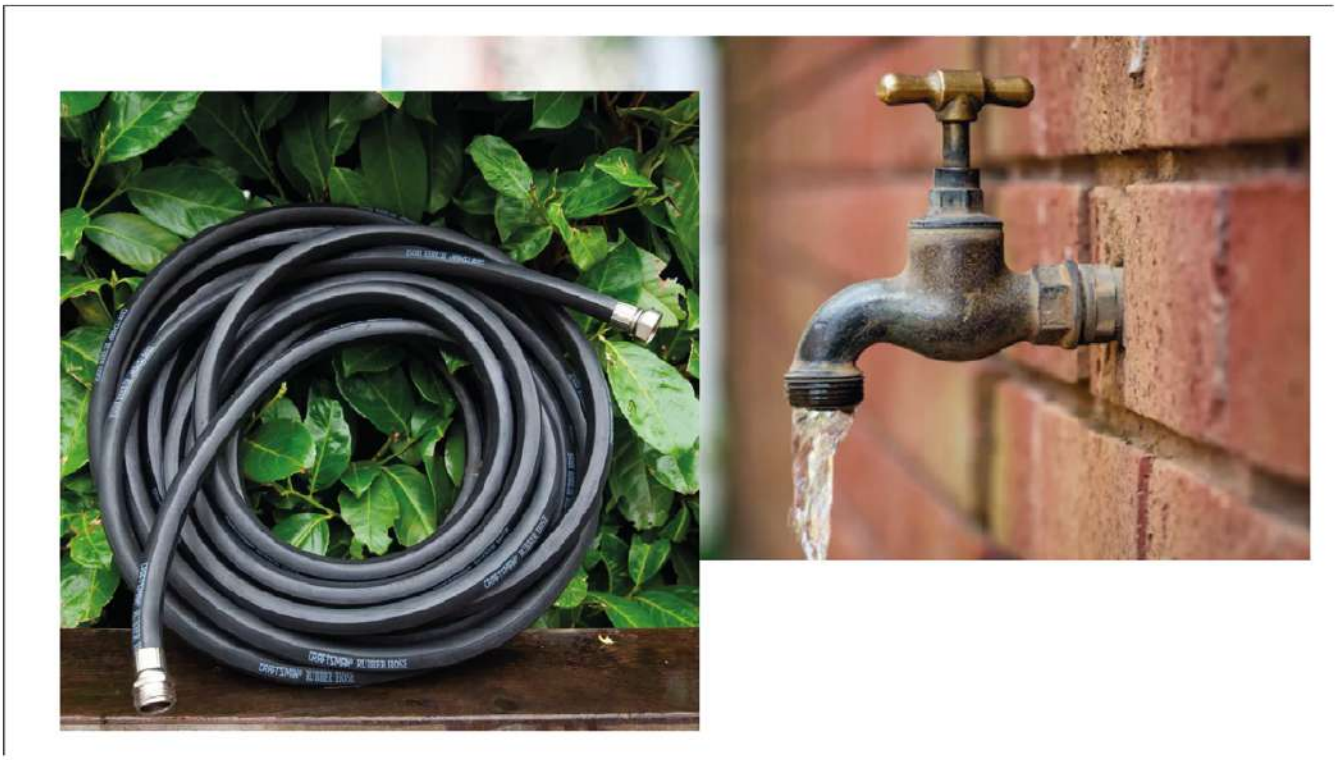

With no one taking proactive ownership, the hose and the faucet fail to connect.

The irony is undeniable: All parties tend to focus more on the technology itself than how it should deploy. *This is like being more excited about the development of a rocket than its launch.*

As a result, nobody is connected the hose to the faucet.

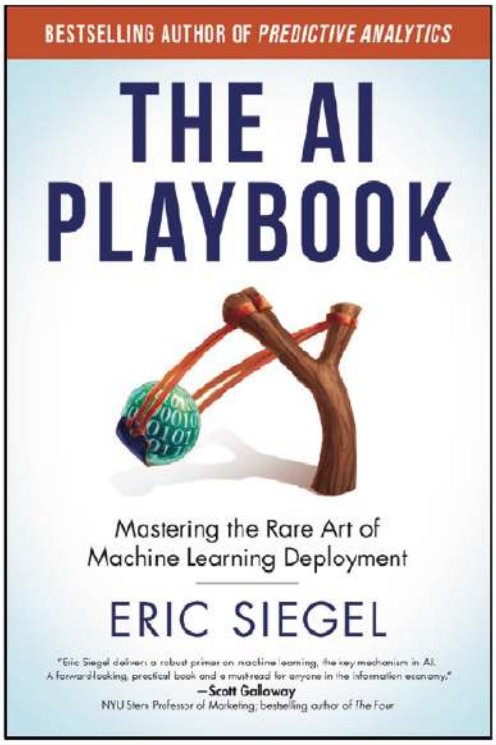

http://www.bizML.com

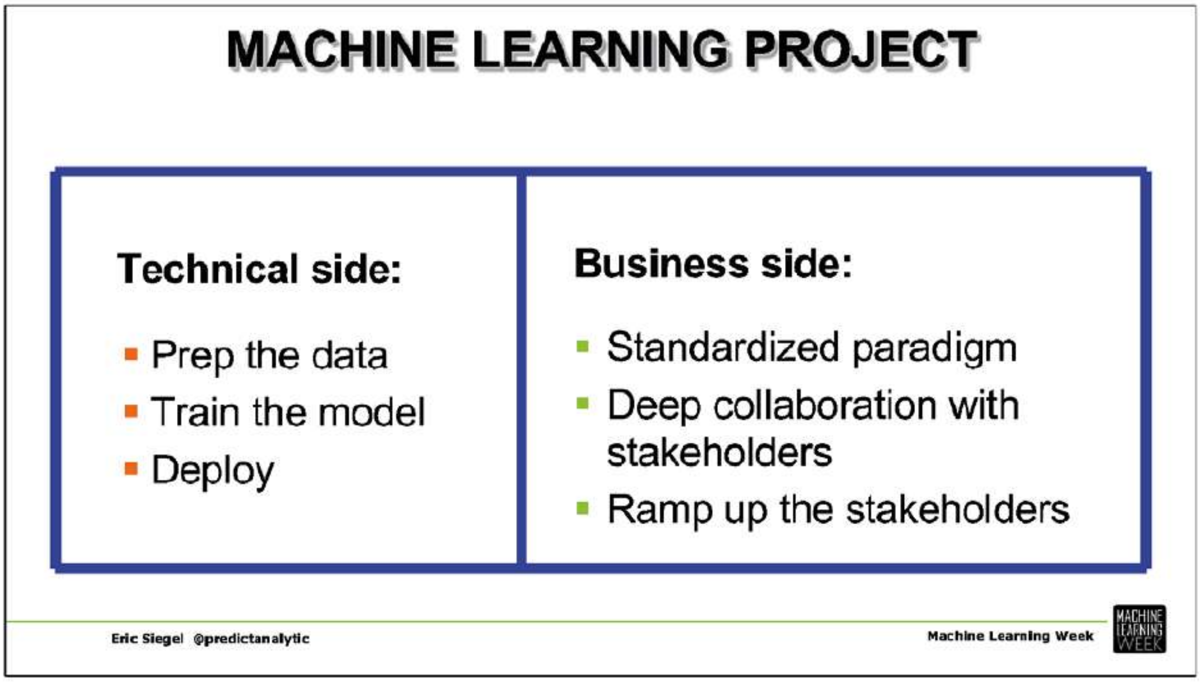

This is the much-neglected "other half" of what it takes to run an ML project.

BizML takes on a very particular change-management process.

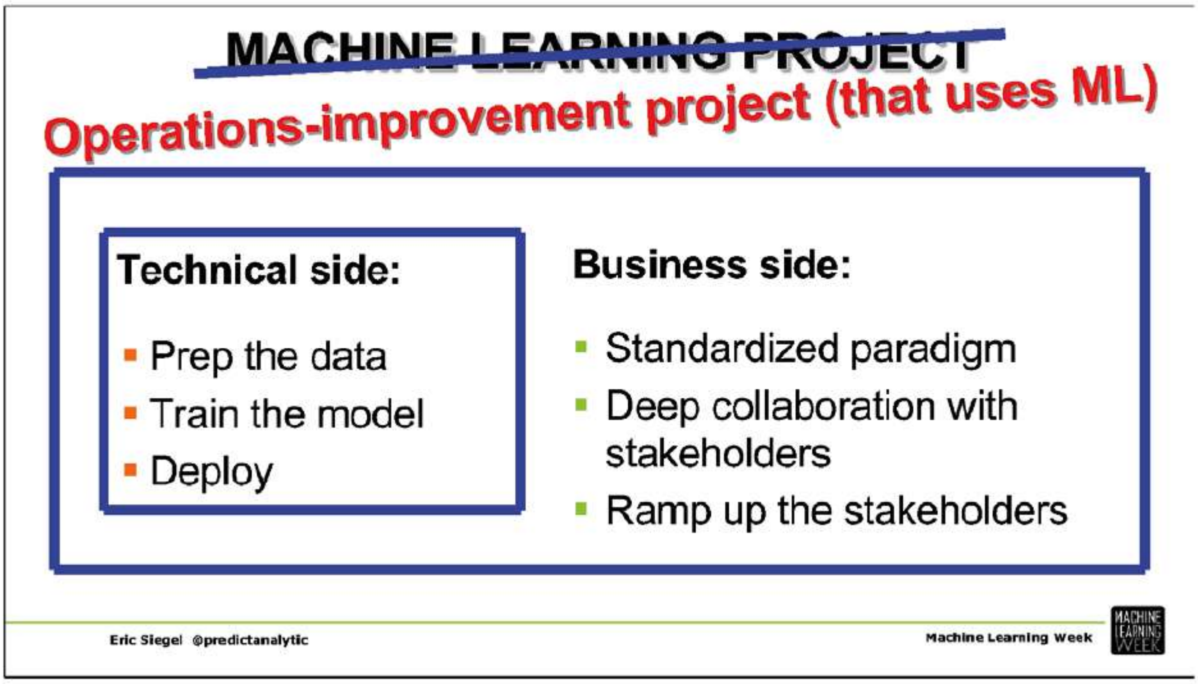

This is the much-neglected "other half" of what it takes to run an ML project

BizML takes on a very particular change-management process.

Generative AI:

The use of ML to generate new content items, such as text, images, video, sound, or code.

Use cases:

- First drafts
- Brainstorming / planning / advice
- Summarization
- Review / proofread
- Chatbots
- Search
- Entertainment (games, creative writing)

Use cases that may be more questionable:
- Tutor
- Therapist
- Employee performance reviews

https://old.reddit.com/r/ArtificialInteligence/comments/1ceaftk/whats_the_most_practical_thing_you_have_done_with/

Application	What's predicted	What's done about it
Generative AI *to generate natural language*	*What is the next word (simplified gist)?*	Write one word at a time by predicting each one (simplified gist).

CHATGPT
OpenAI

Eric Siegel @predictanalytic

Machine Learning Week

Large language models are optimized per-word (simplified gist) -- not for a higher-order objective.

Conclusions

1) Focus on value-driven, concrete use cases

2) Upskill on semi-technical know-how to leverage predictive AI

3) BizML: a specialized playbook to deploy ML

Generative AI:

The use of ML to generate new content items, such as text, images, video, sound, or code.

1) Concrete use case / measure value
2) GenAI easier to use; harder to use WELL
3) Predictive AI is often more valuable
4) #noAGI

Predictive and generative:
- Both use ML (for different purposes)
- Both are branded as "AI"
- Both require a value-focused practice

Overlap: LLMs can also be used as predictive models.

Application	What's predicted	What's done about it
Generative AI *to generate natural language*	*What is the next word (simplified gist)?*	Write one word at a time by predicting each one (simplified gist).

CHATGPT OpenAI

Not designed to be correct.

Large language models are optimized per-word (simplified gist) -- not for a higher-order objective.

"By using Generative AI, marketers were able to reduce the time needed to produce creative campaigns and content by up to 2-3 weeks, resulting in an average time savings of 34%."

... which could equate to an annual savings of almost 3,000 hours across the participant team.

Ally Financial: a top 25 U.S. financial holding company managing $196B in assets and serving over 11 million customers.

One of the nation's leading automotive lenders, as well as providing deposit, lending, mortgage, and credit card services as the nation's largest all-digital bank.

https://www.forbes.com/sites/randybean/2023/11/27/how-ally-financial-is-delivering-business-value-from-generative-ai/?sh=166e03999a86

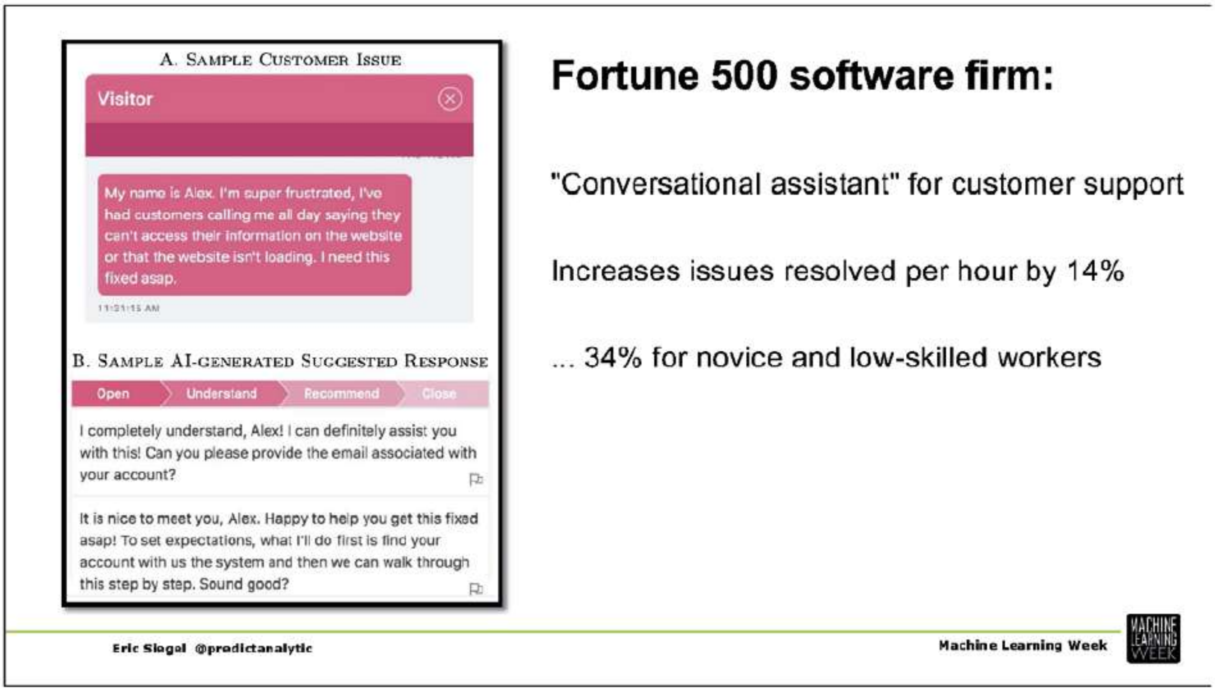

Two main types of output: 1) real-time suggestions for how agents should respond to customers and 2) links to the data firm's internal documentation for relevant technical issues.

GENERATIVE AI AT WORK
by Erik Brynjolfsson et al
NATIONAL BUREAU OF ECONOMIC RESEARCH
https://www.nber.org/system/files/working_papers/w31161/w31161.pdf

https://mitsloan.mit.edu/ideas-made-to-matter/workers-less-experience-gain-most-generative-ai

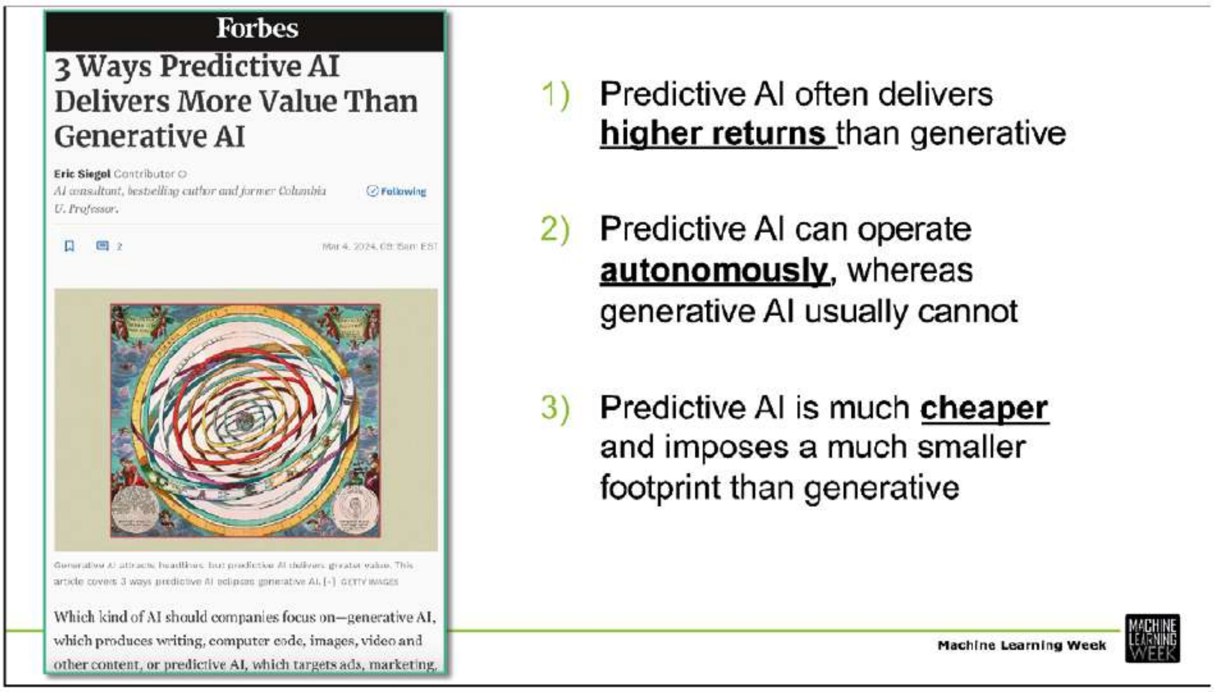

https://www.forbes.com/sites/ericsiegel/2024/03/04/3-ways-predictive-ai-delivers-more-value-than-generative-ai/?sh=7c028f584e84

> *AI strategies fail because AI is a means, not an end. 'Do you have an AI strategy?' makes as much sense as asking, 'Do we have an Excel strategy?'.*

Mihnea Moldoveanu
University of Toronto management professor

Eric Siegel @predictanalytic — Machine Learning Week

https://hbr.org/2019/03/why-ai-underperforms-and-what-companies-can-do-about-it

OTHER ARTICLES TEMPERING AI HYPE:

https://hbr.org/2023/06/the-ai-hype-cycle-is-distracting-companies

https://www.forbes.com/sites/ericsiegel/2024/04/10/artificial-general-intelligence-is-pure-hype/?sh=40e2a52a73c5

https://www.wsj.com/articles/cfos-tackle-thorny-calculus-on-gen-ai-whats-the-return-on-investment-24ebf435

https://twitter.com/binarybits/status/1778432250436046961

https://sloanreview.mit.edu/article/dont-get-distracted-by-the-hype-around-generative-ai/

"Will A.I. Ever Live Up to Its Hype?"
https://www.nytimes.com/2024/05/15/opinion/artificial-intelligence-ai-openai-chatgpt-overrated-hype.html

"Companies have been timid to adopt generative AI technologies on a wide scale."
https://outthinker.us15.list-manage.com/track/click?u=0c3f7e67b4ffa9d91ff7d63f9&id=fa11a89152&e=a5e154a5bc

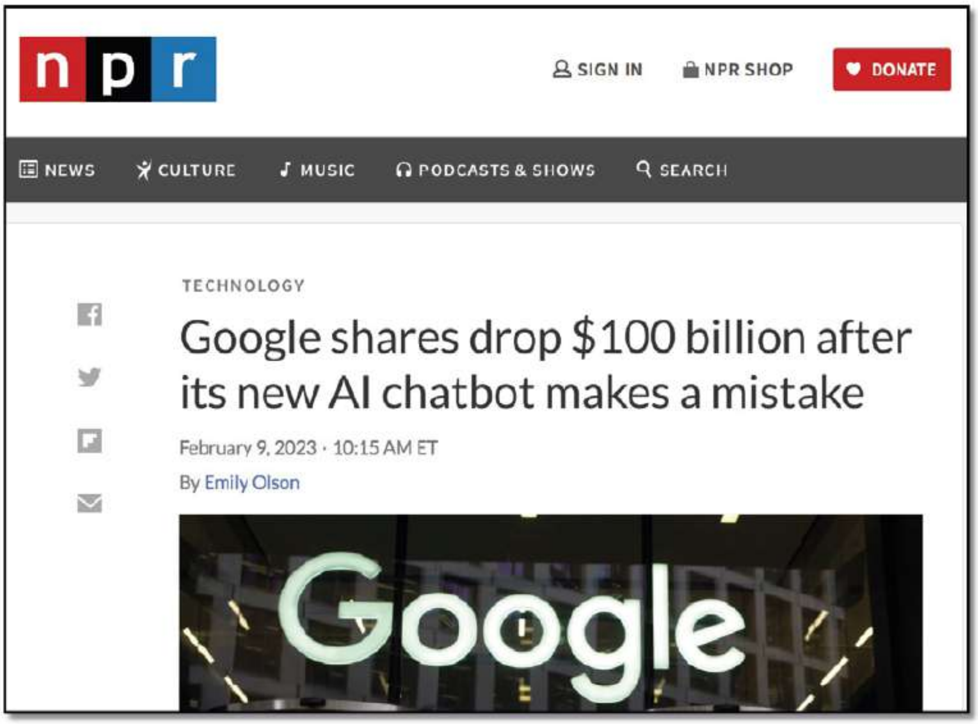

npr

SIGN IN NPR SHOP DONATE

NEWS CULTURE MUSIC PODCASTS & SHOWS SEARCH

TECHNOLOGY

Google shares drop $100 billion after its new AI chatbot makes a mistake

February 9, 2023 · 10:15 AM ET

By Emily Olson

https://www.npr.org/2023/02/09/1155650909/google-chatbot--error-bard-shares

Harvard Business Review

Don't Let Gen AI Limit Your Team's Creativity

From the Magazine (March–April 2024)

Stanford researchers:

Comparing teams with/without ChatGPT

Business problems:

- Develop internal training resources
- Scale up B2B sales

Hypothesis: Teams with ChatGPT would generate vastly more and better ideas

Results: GenAI led to more average ideas

Eric Siegel @predictanalytic — Machine Learning Week

GenAI by definition generates "average" inside the box responses.
GenAI can be overly trusted - more than human.
Slightly fewer D's, mores B's, fewer A's.
The researchers weren't expecting this, and are still bullish, suggesting guidance for more well homes prompt engineering.

https://hbr.org/2024/03/dont-let-gen-ai-limit-your-teams-creativity

https://docsend.com/view/7xyy84yi9kyn8ngr

https://howtofixit.ai/

Final Conclusions

1) **GenAI is easier to use, but harder to use *well***
2) **Focus on value-driven, concrete use cases**
3) **Upskill on semi-technical know-how to leverage predictive AI**
4) **BizML: a specialized playbook to deploy ML**

Eric Siegel @predictanalytic — Machine Learning Week

https://machinelearningweek.com/

https://generativeaiapplicationssummit.com

@predictanalytic

Eric Siegel @predictanalytic

Machine Learning Week

NEXUS

LEADERSHIP+

Strategic Partner

BUSINESS بالعربي VLOGS

CanCham Canadian Chamber of Commerce MIDDLE EAST

webkeyz

ON AIR GROUP

Geminimedia

ADVERTISING MARKETING

WADI TOURS
الوادي للسياحة
Your Real Partner

THE GAME CHANGERS