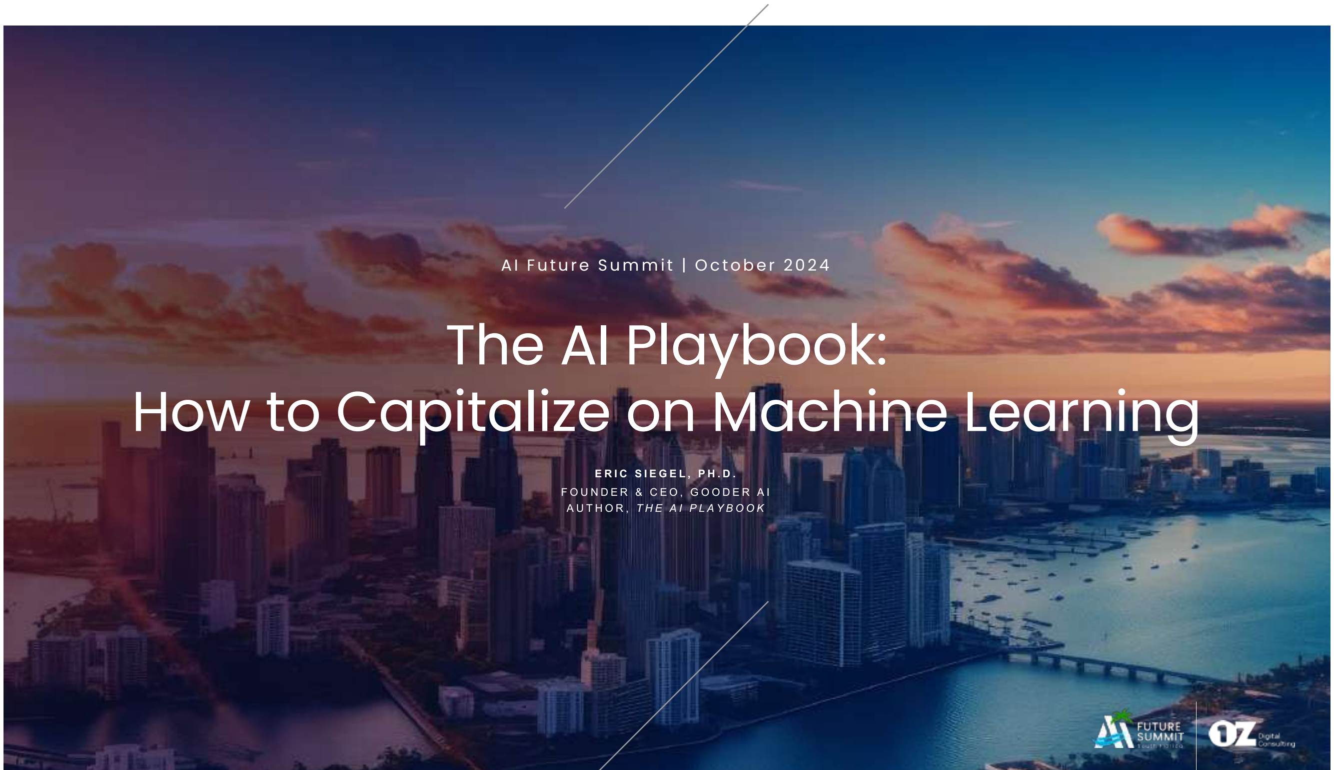

For more information/citations regarding the examples in this presentation, see the Notes, freely-accessible online, for the book "Predictive Analytics" by Eric Siegel (http://www.thepredictionbook.com). Most of the various examples shown are covered in the book (some only briefly, within the book's 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.

To continue your learning:
http://www.MachineLearning.courses

business

data

prediction

machine learning

© 2024 OZ Digital, LLC - All rights reserved.

predictive AI
predictive analytics

© 2024 OZ Digital, LLC - All rights reserved.

boost sales

cut costs

combat risk

prevent fraud

fortify healthcare

streamline logistics

conquer spam

win elections

© 2024 OZ Digital, LLC - All rights reserved.

> Machine learning's practical deployment represents the forefront of human progress: improving operations with science.

Morgan Vawter
Global GP of Data & Analytics, Unilever

© 2024 OZ Digital, LLC - All rights reserved.

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).

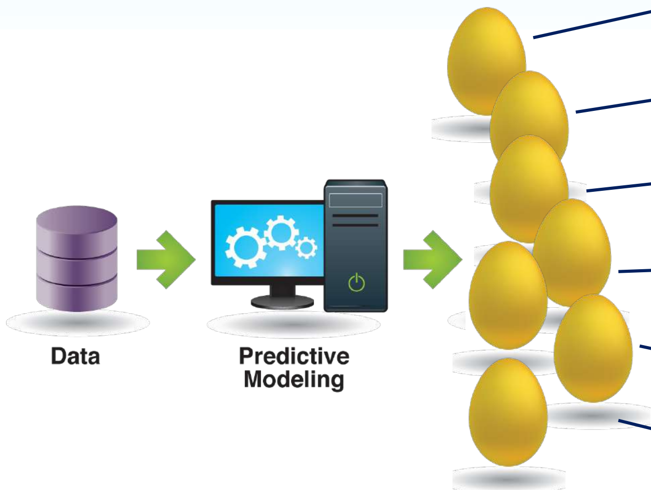

OPERATIONS:

~~Marketing~~

~~Sales~~

~~Fraud detection~~

~~Credit scoring~~

~~Logistics~~

~~Etc.~~

© 2024 OZ Digital, LLC - All rights reserved.

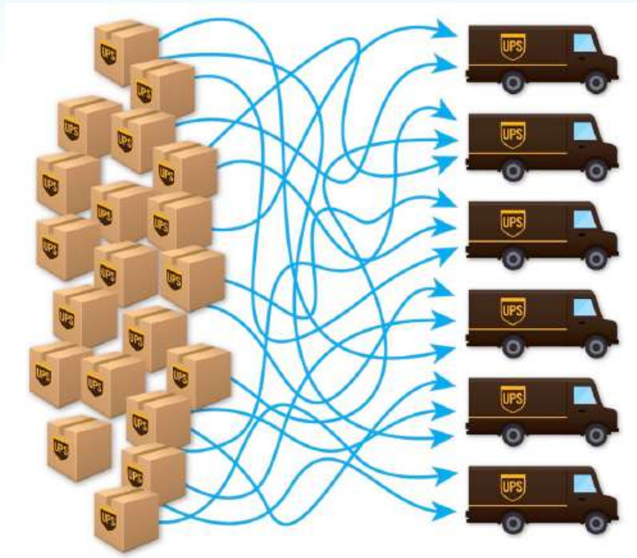

© 2024 OZ Digital, LLC - All rights reserved.

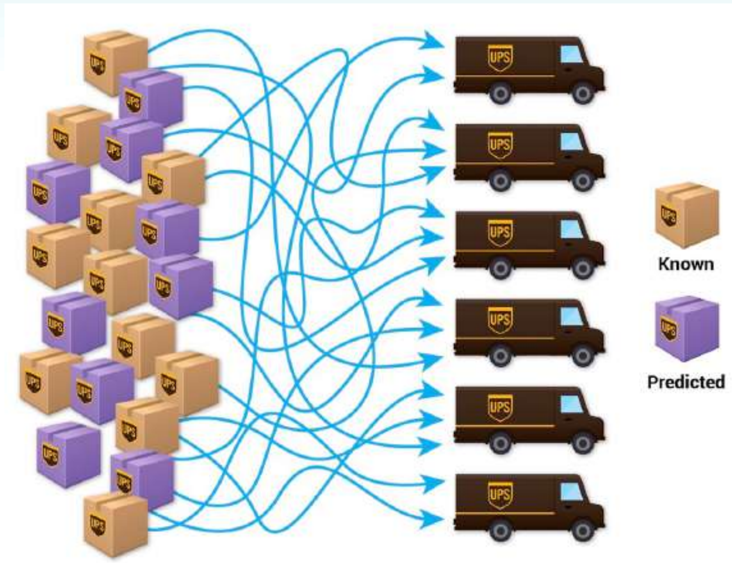

© 2024 OZ Digital, LLC - All rights reserved.

9

Package Flow Technology and ORION

Overall annual savings:

185 million miles

$350+ million

185,000 metric tons of emissions

© 2024 OZ Digital, LLC - All rights reserved.

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

18.5 million miles

$35+ million

GenAI is easier to use.

© 2024 OZ Digital, LLC - All rights reserved.

"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%."

© 2024 OZ Digital, LLC - All rights reserved.

GenAI is easier to use…

…but much harder to use well…

…and oversold.

© 2024 OZ Digital, LLC - All rights reserved.

© 2024 OZ Digital, LLC - All rights reserved.

3 Ways Predictive AI Delivers More Value Than Generative AI
By Eric Siegel, Forbes
https://www.forbes.com/sites/ericsiegel/2024/03/04/3-ways-predictive-ai-delivers-more-value-than-generative-ai/

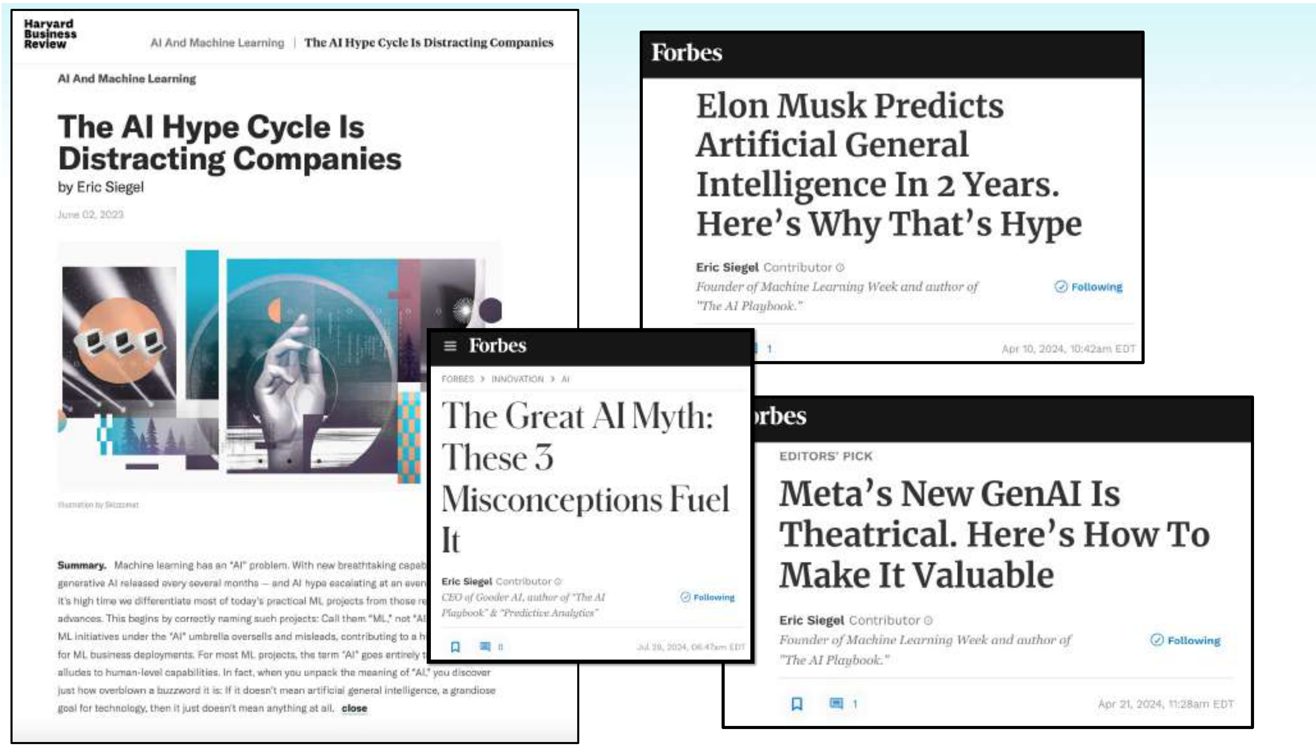

https://hbr.org/2023/06/the-ai-hype-cycle-is-distracting-companies

https://www.forbes.com/sites/ericsiegel/2024/07/29/the-great-ai-myth-these-3-misconceptions-fuel-it/

https://www.forbes.com/sites/ericsiegel/2024/04/10/artificial-general-intelligence-is-pure-hype/

https://www.forbes.com/sites/ericsiegel/2024/04/21/metas-new-genai-is-theatrical-heres-how-to-make-it-valuable/

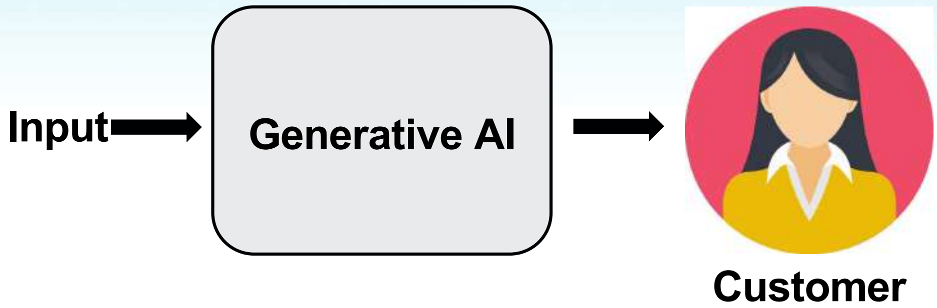

100% autonomous, 95% reliable

© 2024 OZ Digital, LLC - All rights reserved.

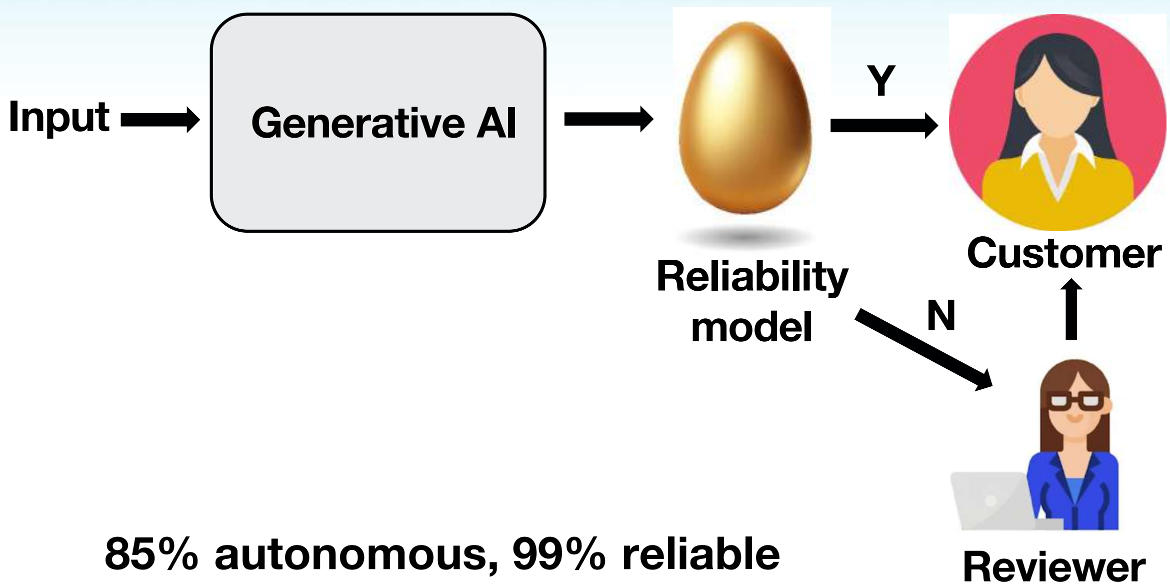

85% autonomous, 99% reliable

© 2024 OZ Digital, LLC - All rights reserved.

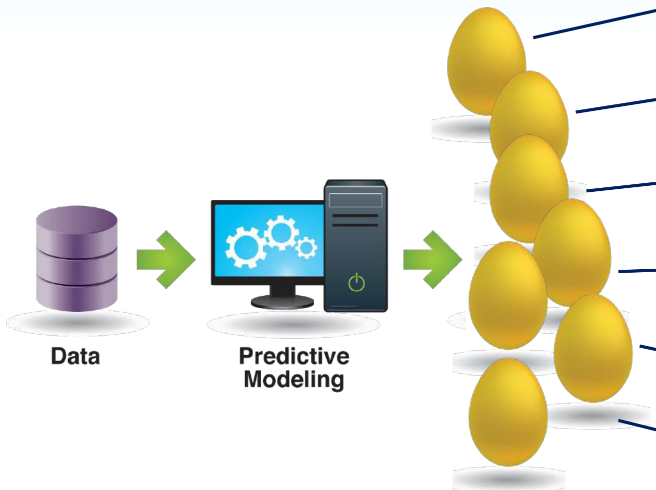

OPERATIONS:

- Marketing
- Sales
- Fraud detection
- Credit scoring
- Logistics
- Etc.

© 2024 OZ Digital, LLC - All rights reserved.

Unmet need: a business playbook for running ML projects

© 2024 OZ Digital, LLC - All rights reserved.

8

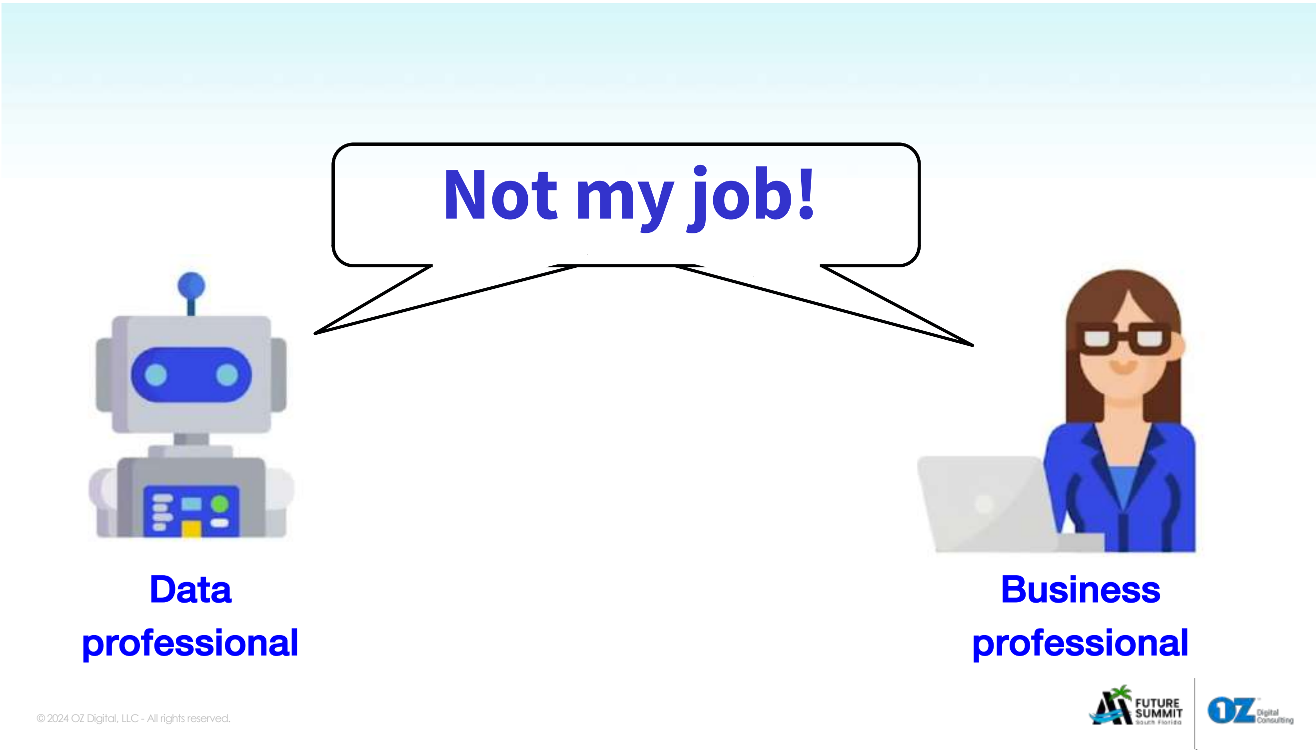

Who's meant to form and participate in the business process? Both sides. And which side takes up the mantle? Often, neither does. As you'll see, both sides have reason to consider it someone else's responsibility.

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

Agenda

- Predictive AI's value proposition
- Why predictive AI projects routinely fail to deploy
- Semi-technical understanding
- The **bizML** playbook

© 2024 OZ Digital, LLC - All rights reserved.

Before:

After:

I am an individual patient, and an individual insurance policyholder.
Risk effects all parties involved.

Knee Walking

© 2024 OZ Digital, LLC - All rights reserved.

ACL replacement surgery choice of graft source influences the risk of long term knee pain when “knee walking”.

© 2024 OZ Digital, LLC - All rights reserved.

© 2024 OZ Digital, LLC - All rights reserved.

Insured “office workers”

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

© 2024 OZ Digital, LLC - All rights reserved.

"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


Risk an individual may:
▪Not respond
▪Defect
▪Default
▪Commit fraud

© 2024 OZ Digital, LLC - All rights reserved.

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

© 2024 OZ Digital, LLC - All rights reserved.

© 2024 OZ Digital, LLC - All rights reserved.

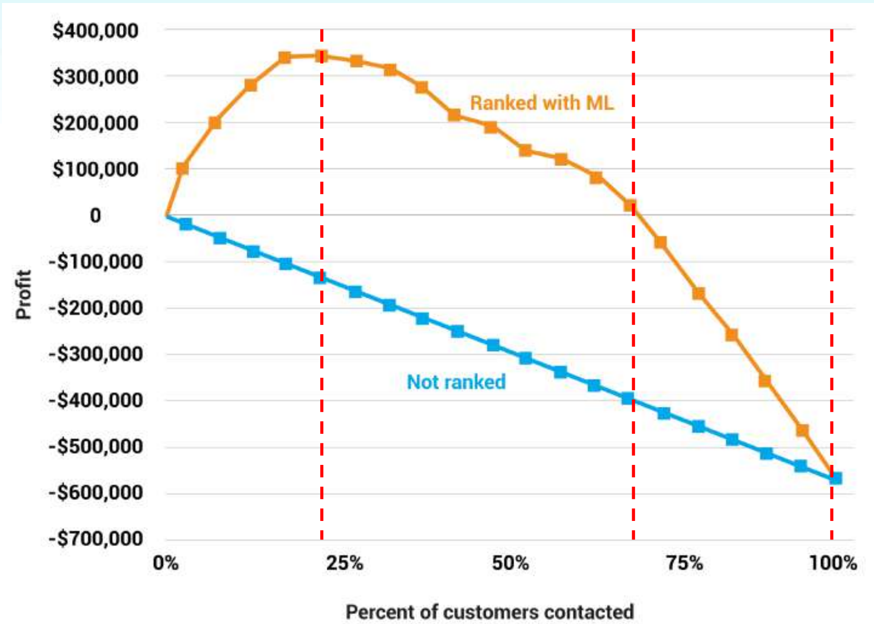

© 2024 OZ Digital, LLC - All rights reserved.

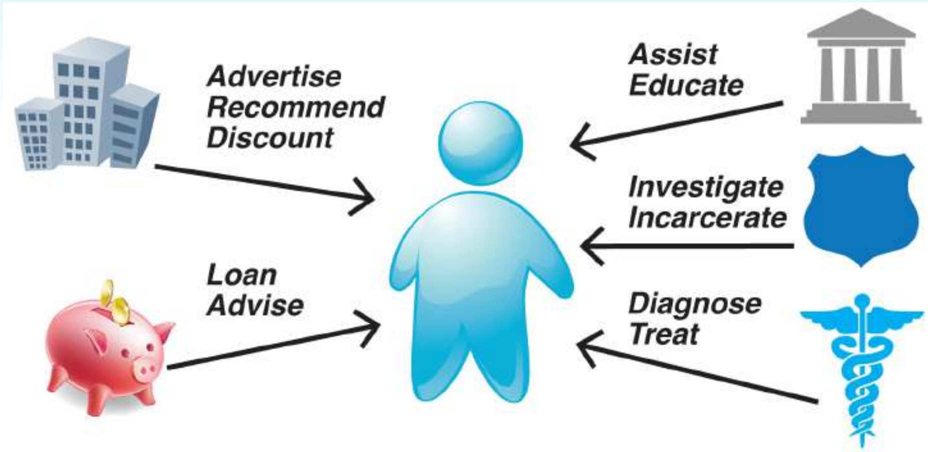

Millions of decisions a day determine whom to *call, mail, approve, test, diagnose, warn, investigate, incarcerate, set up on a date, and medicate.*

© 2024 OZ Digital, LLC - All rights reserved.

© 2024 OZ Digital, LLC - All rights reserved.

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*

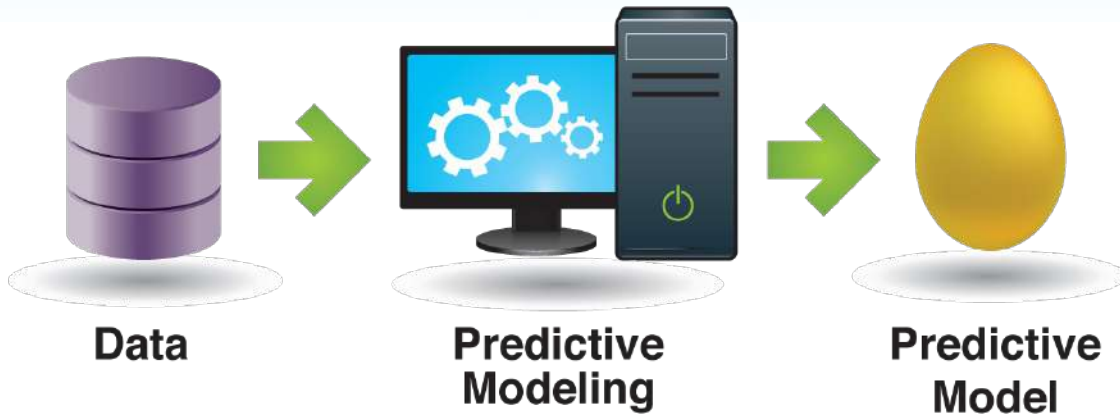

© 2024 OZ Digital, LLC - All rights reserved.

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).

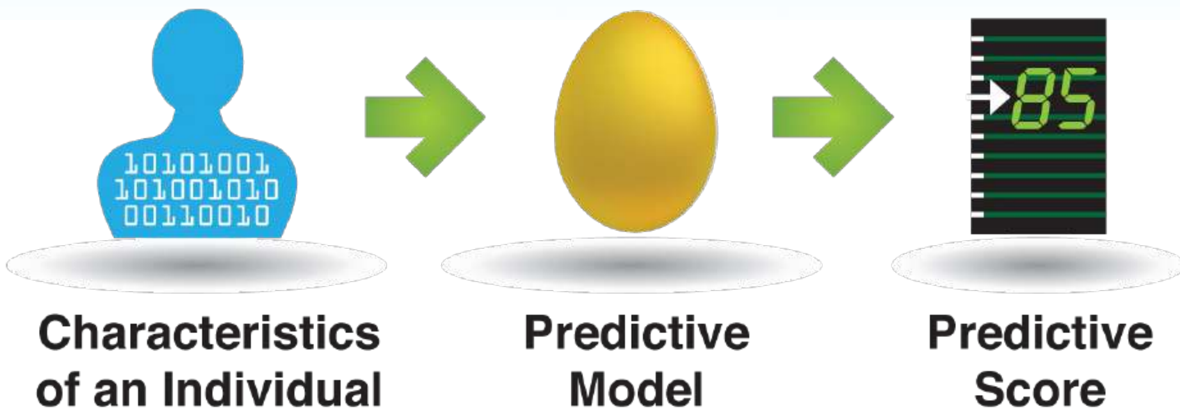

© 2024 OZ Digital, LLC - All rights reserved.

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).

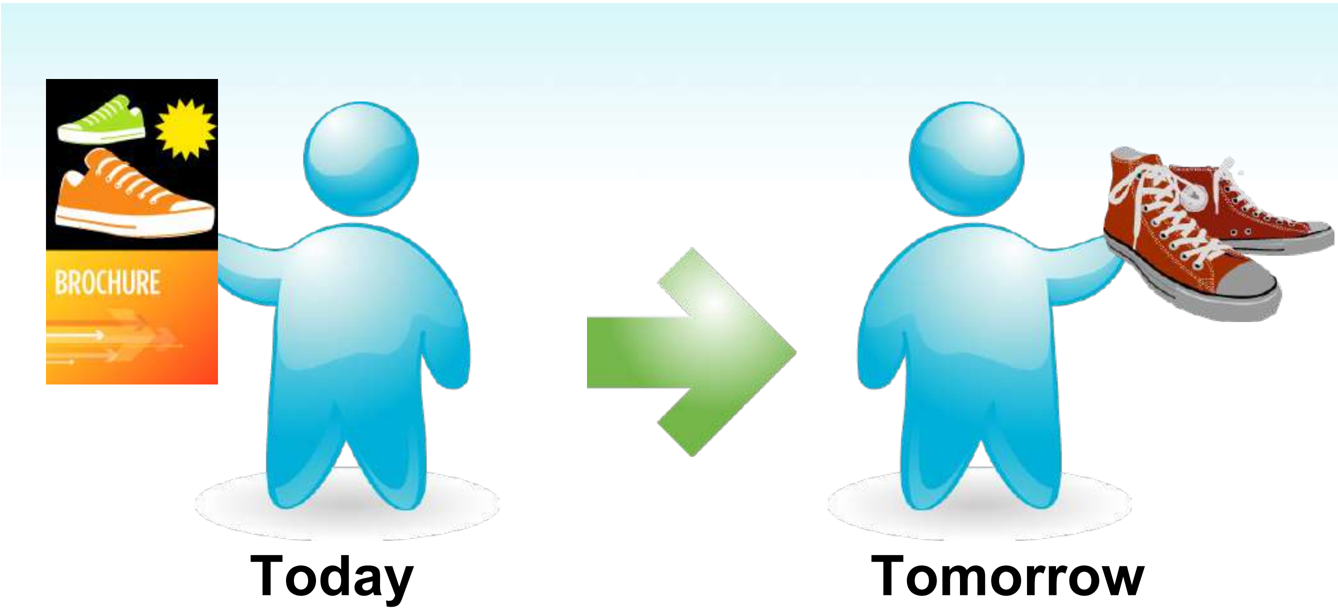

© 2024 OZ Digital, LLC - All rights reserved.

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).

The Challenge of Prediction

Prediction is very difficult, especially if it's about the future.

- Niels Bohr

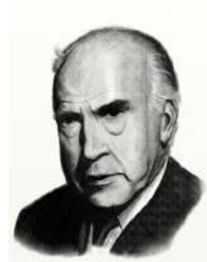

How come you never see a headline like “Psychic Wins Lottery”?

- Jay Leno

© 2024 OZ Digital, LLC - All rights reserved.

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).

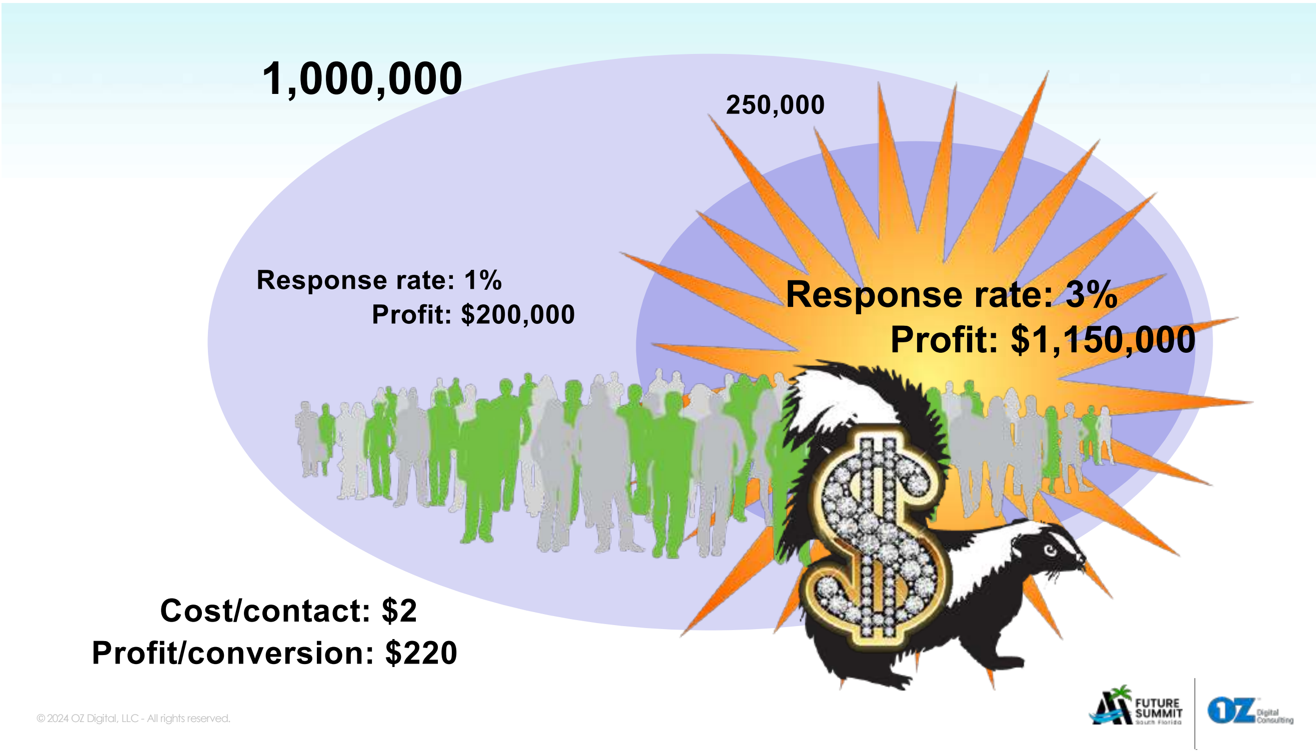

A crummy predictive model delivers big value. It’s like a skunk with bling.

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.

© 2024 OZ Digital, LLC - All rights reserved.

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).

...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/IBMfirstTennBank.pdf

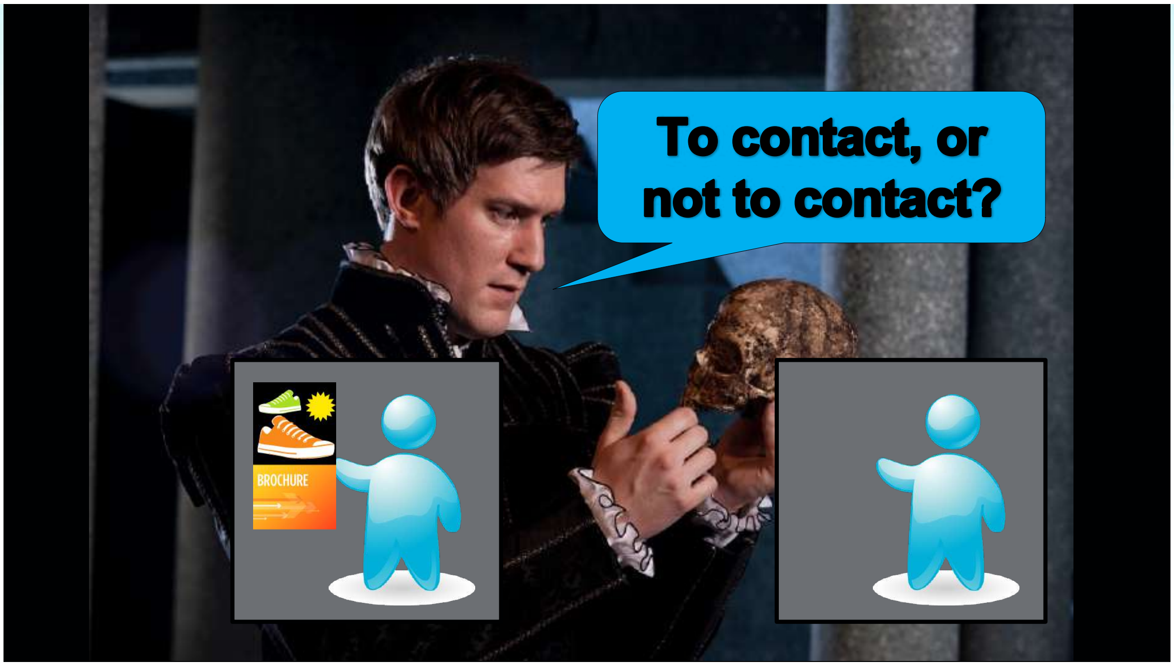

© 2024 OZ Digital, LLC - All rights reserved.

© 2024 OZ Digital, LLC - All rights reserved.

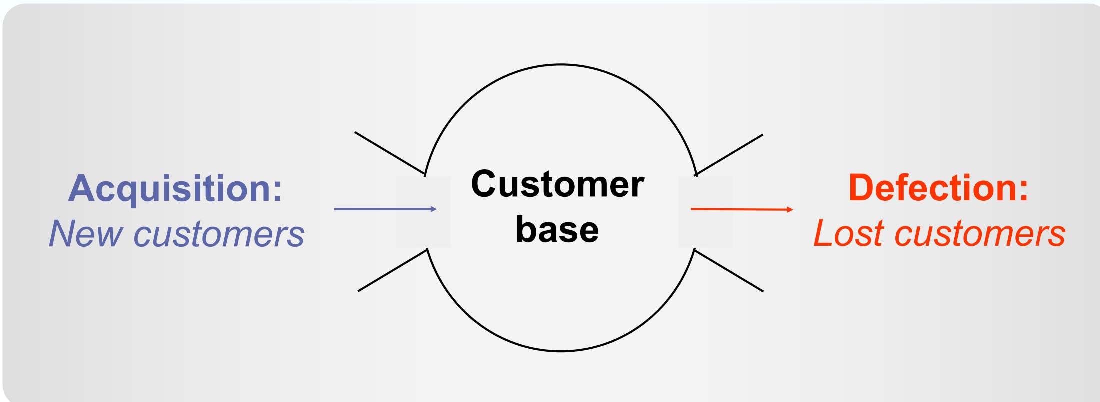

© 2024 OZ Digital, LLC - All rights reserved.

The only way to target a retention campaign precisely where it's needed is with predictive scores that earmark which customers are most likely to leave.

Targeted retention is often the lowest hanging fruit among prospective applications of predictive analytics at an organization.

For more information, see Chapter 7 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

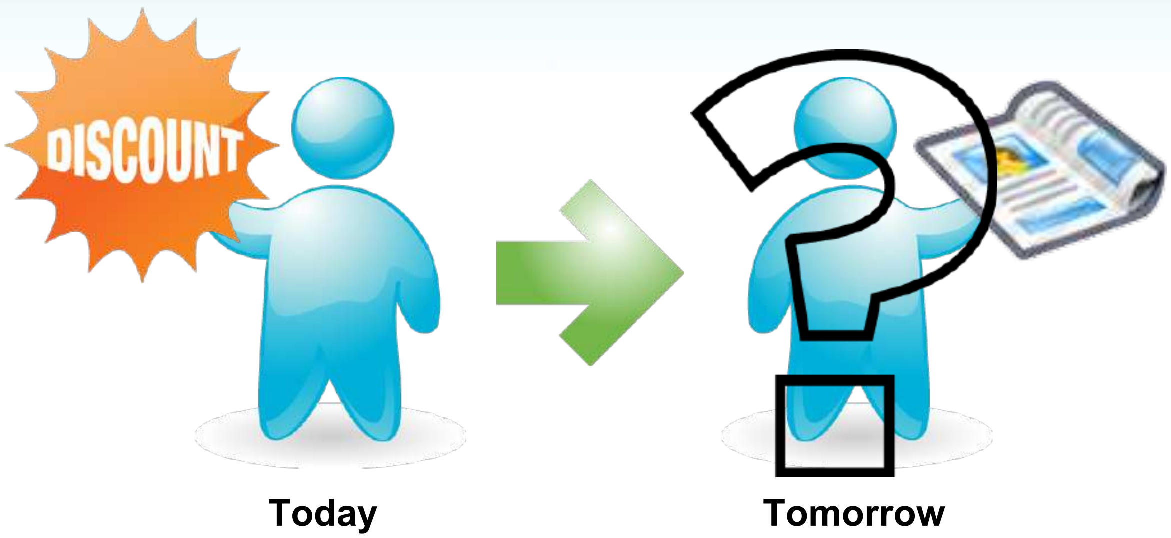

© 2024 OZ Digital, LLC - All rights reserved.

Customer Attrition: Mortgages

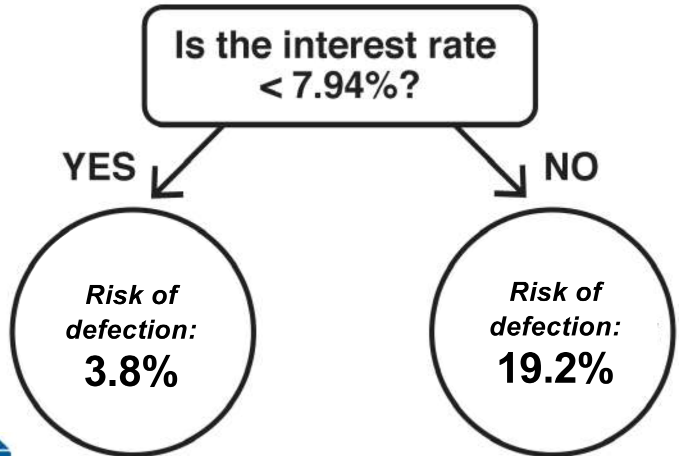

© 2024 OZ Digital, LLC - All rights reserved.

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

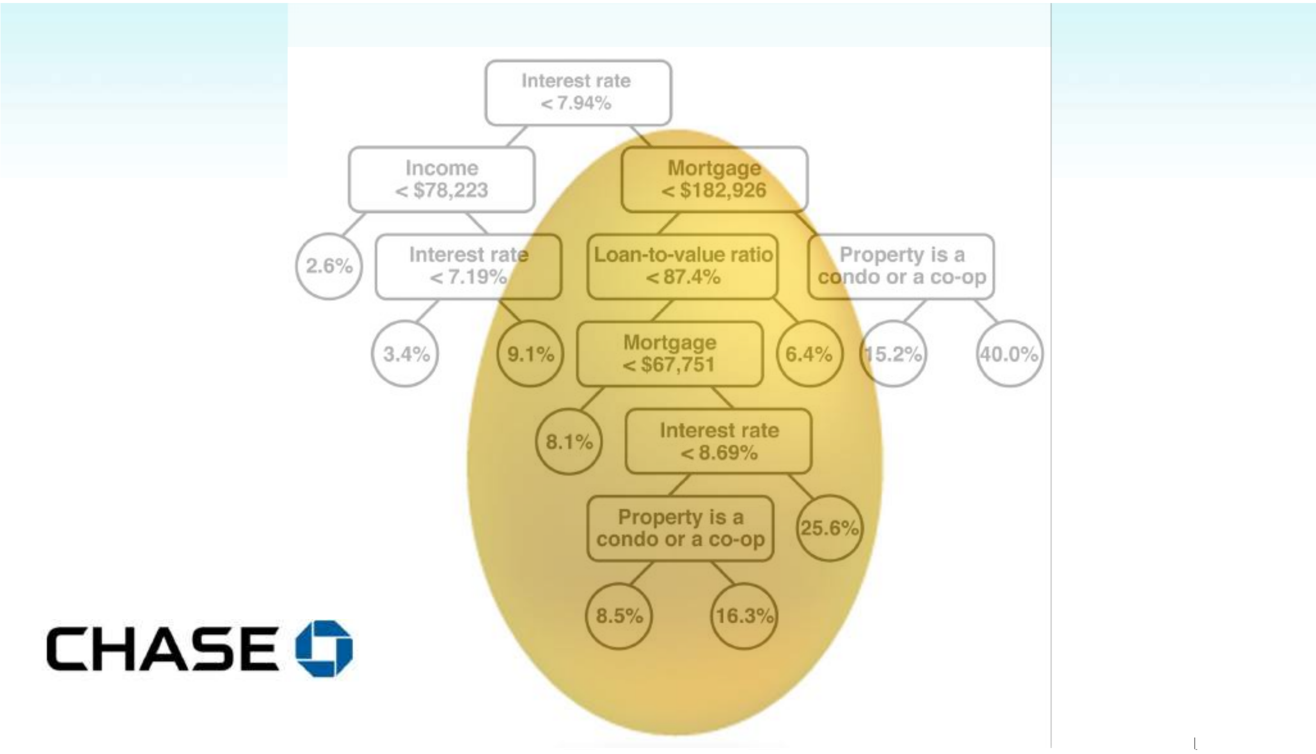

Interest rate
< 7.94%
Income
< $78,223
Mortgage
< $182,926
2.6%
Interest rate
< 7.19%
Loan-to-value ratio
< 87.4%
Property is a
condo or a co-op
3.4%
9.1%
Mortgage
< $67,751
6.4%
15.2%
40.0%
8.1%
Interest rate
< 8.69%
Property is a
condo or a co-op
25.6%
8.5%
16.3%

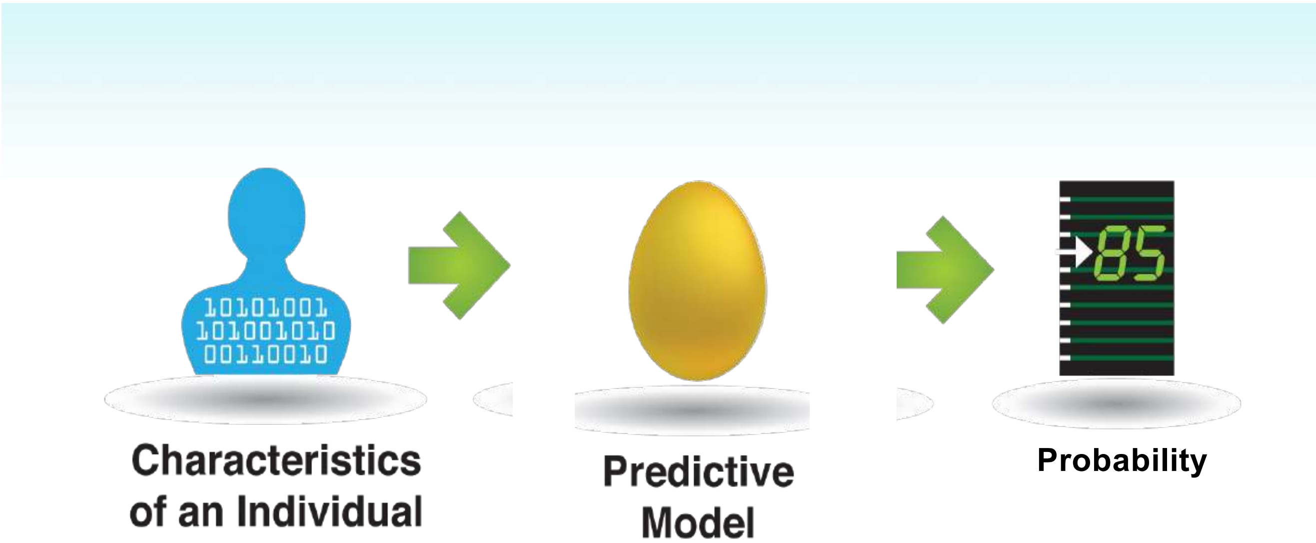
10101001
101001010
00110010
85
Characteristics of an Individual
Predictive Model
Probability

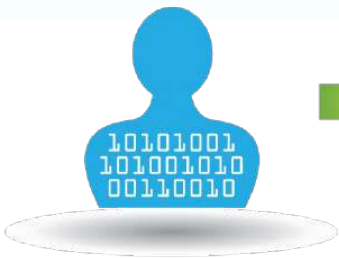

Characteristics of an Individual

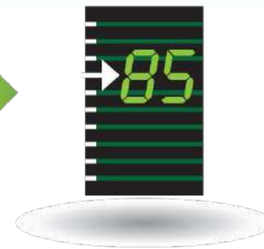

Probability

A probability calculator

© 2024 OZ Digital, LLC - All rights reserved.

8

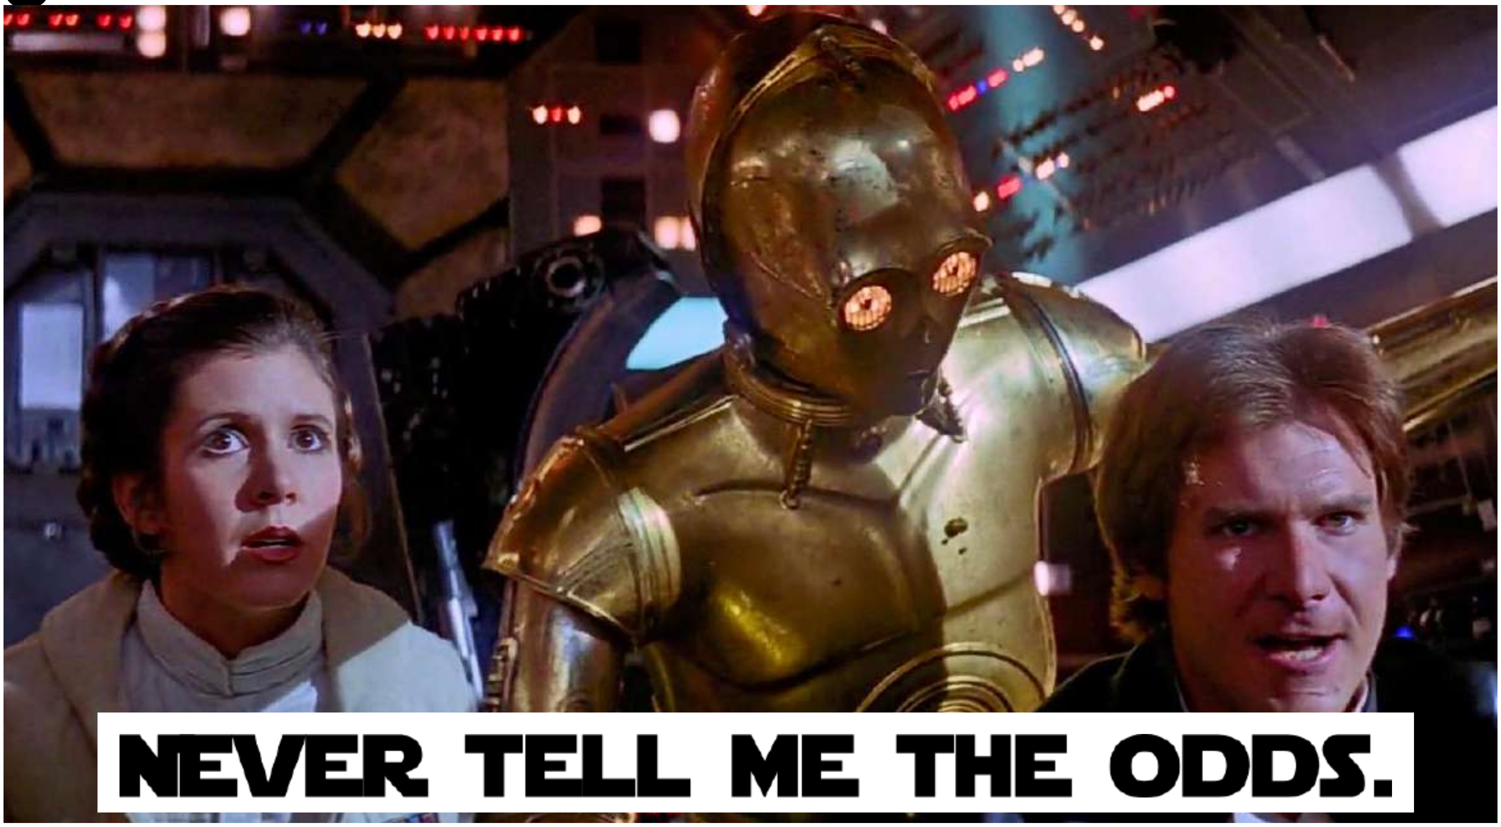

C3P0: The possibility of successfully navigating... 3,720 to 1.

Han Solo: Never tell me the odds.

© 2024 OZ Digital, LLC - All rights reserved.

“Moneyball” celebrates math and yet epitomizes the glossing over of math.

1. **What's predicted**
2. **How well**
3. **What's done about it**

© 2024 OZ Digital, LLC - All rights reserved.

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**

© 2024 OZ Digital, LLC - All rights reserved.

Benchmarks for training and deployment

Technical performance and business performance

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.

© 2024 OZ Digital, LLC - All rights reserved.

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.

© 2024 OZ Digital, LLC - All rights reserved.

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.

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.

© 2024 OZ Digital, LLC - All rights reserved.

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.

 

© 2024 OZ Digital, LLC - All rights reserved.

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**

© 2024 OZ Digital, LLC - All rights reserved.

Benchmarks for training and deployment

Technical performance and business performance

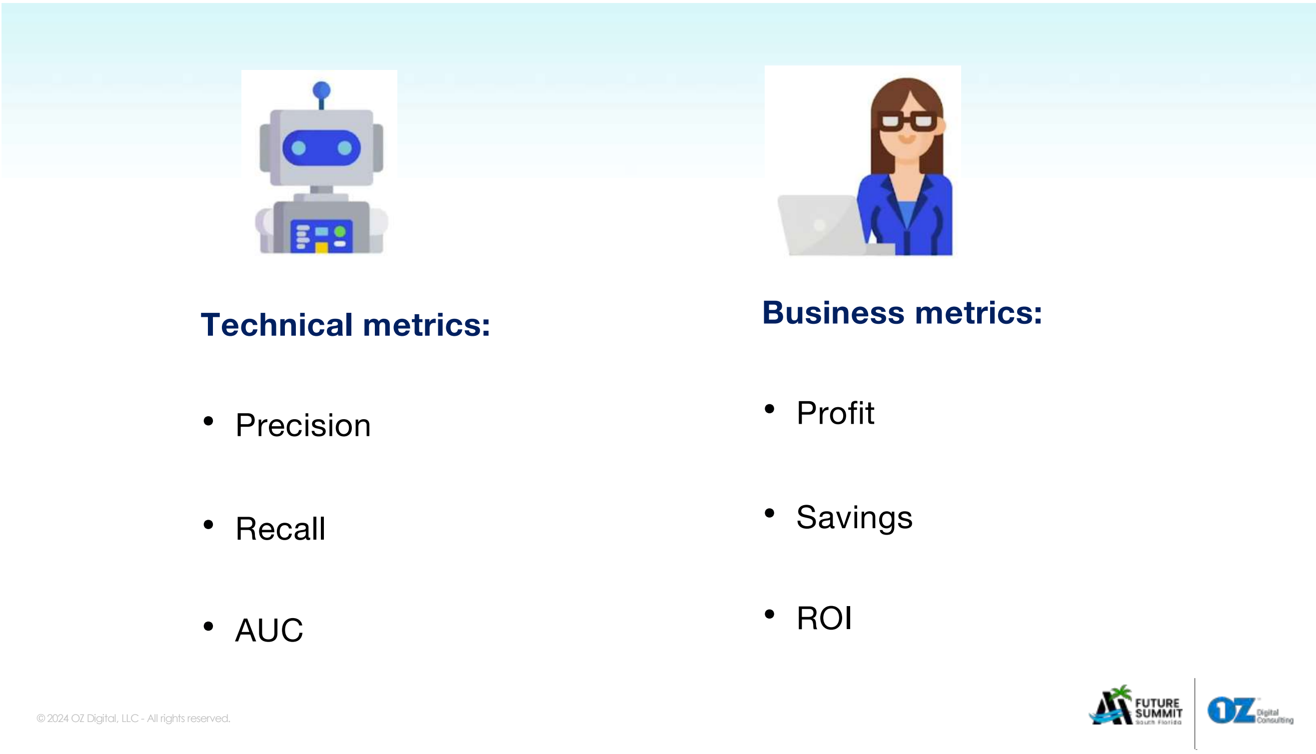

Eric Siegel has co-founded a startup to move from technical ML metrics to business metrics: https://www.Gooder.ai

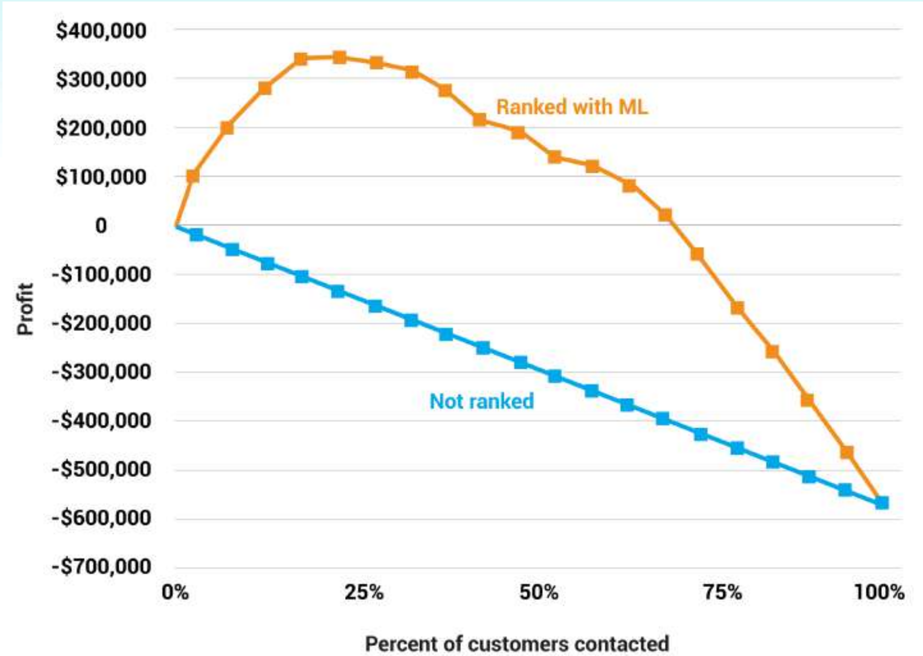

© 2024 OZ Digital, LLC - All rights reserved.

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.

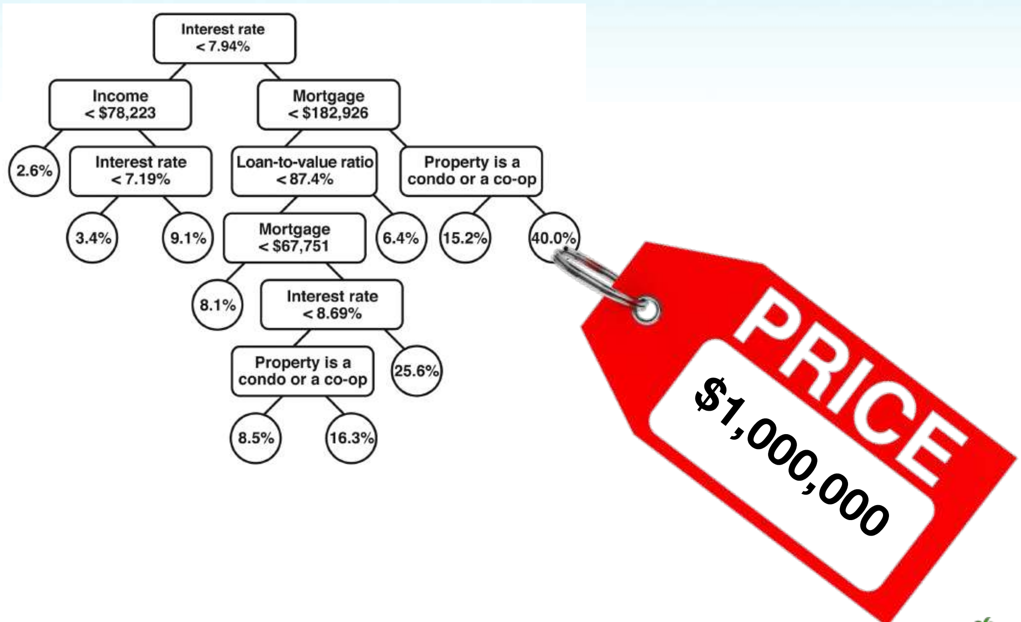

© 2024 OZ Digital, LLC - All rights reserved.

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

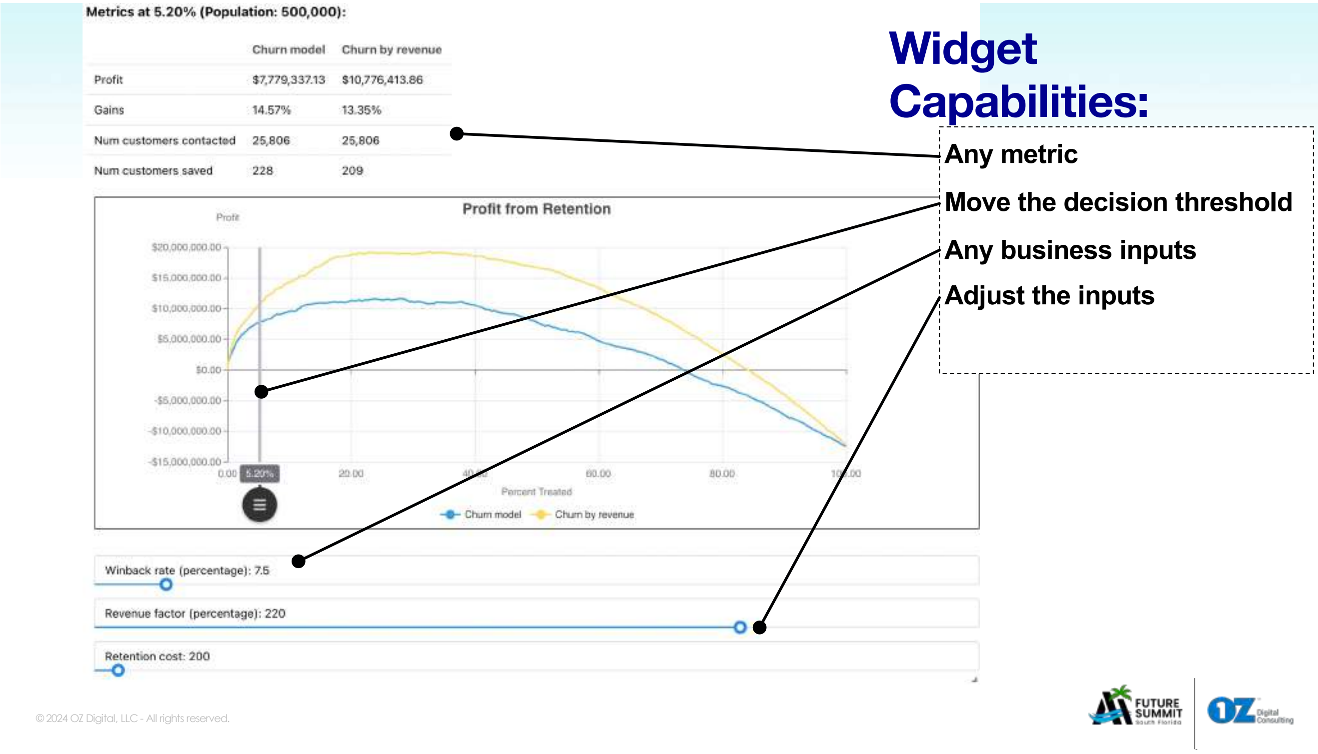

Move the decision threshold:
View the available tradeoffs between competing metrics

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.

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	$10,908,998.54	$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

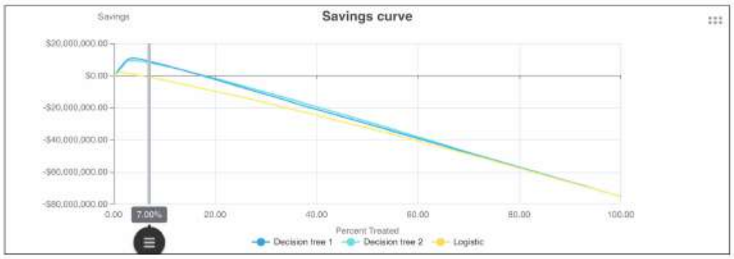

FP-cost: 10

FN-cost: 60

© 2024 OZ Digital, LLC - All rights reserved.

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

© 2024 OZ Digital, LLC - All rights reserved.

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

© 2024 OZ Digital, LLC - All rights reserved.

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

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

© 2024 OZ Digital, LLC - All rights reserved.

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

Data scientist:

Those are management issues.

My model's valuable – of course it'll be deployed.

© 2024 OZ Digital, LLC - All rights reserved.

6

Business professional:

I delegate all that technical stuff to the experts.

© 2024 OZ Digital, LLC - All rights reserved.

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

7

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

8

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.*

Business professional:

I delegate all that technical stuff to the experts.

To drive a car, I don't need to look under the hood.

© 2024 OZ Digital, LLC - All rights reserved.

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

© 2024 OZ Digital, LLC - All rights reserved.

1

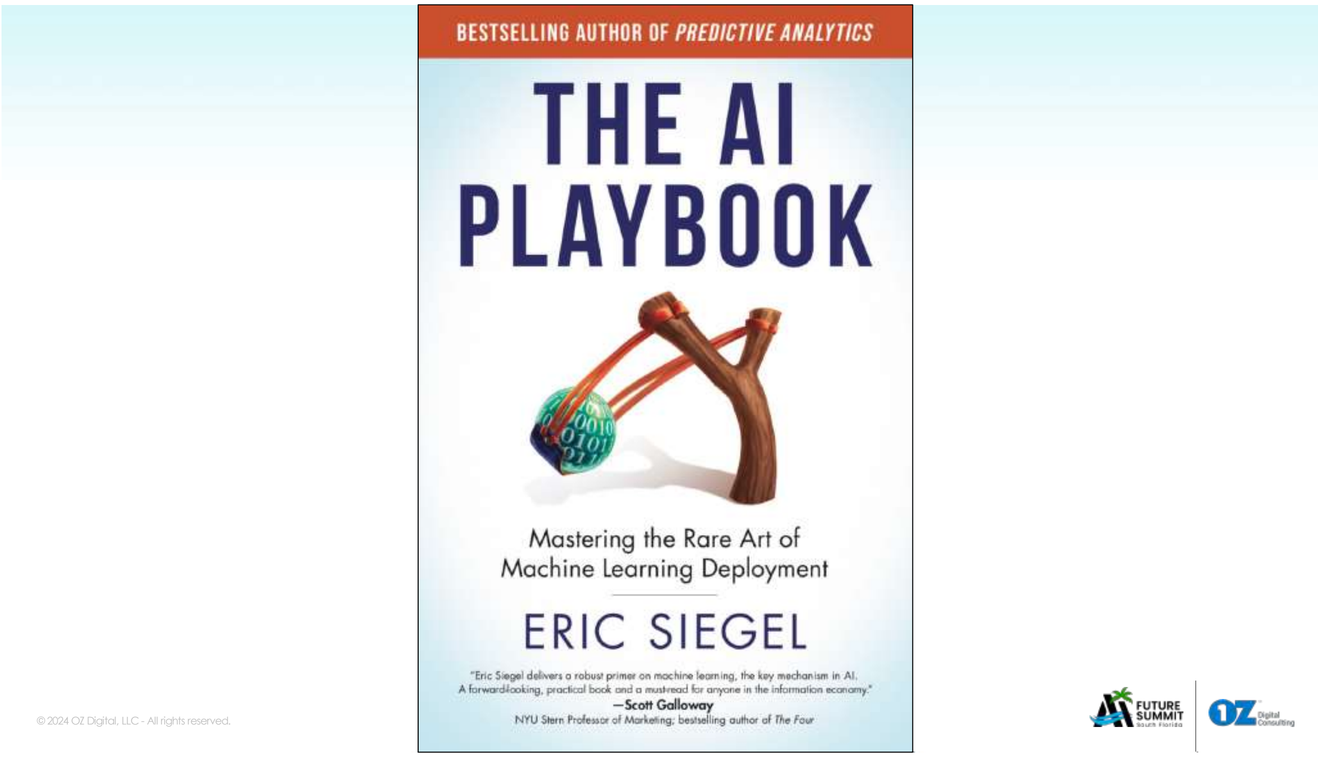

http://www.bizML.com

Conclusions

- Predictive AI is a consulting gig
- Upskill on semi-technical
- A specialized playbook: **bizML**

© 2024 OZ Digital, LLC - All rights reserved.

More content from Eric Siegel

MACHINE LEARNING WEEK

Talks at Google

MACHINE LEARNING Times
The machine learning professionals' premier resource

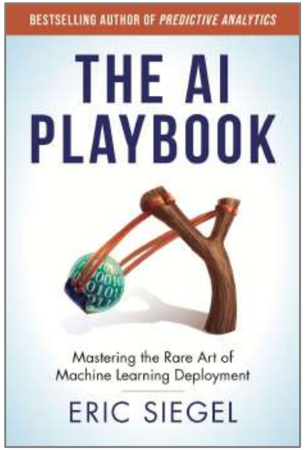

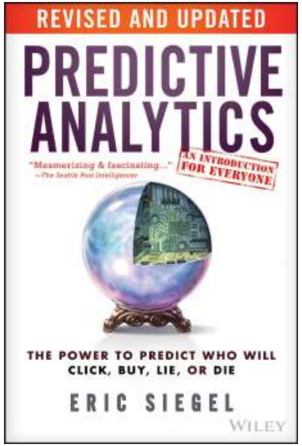

Forbes

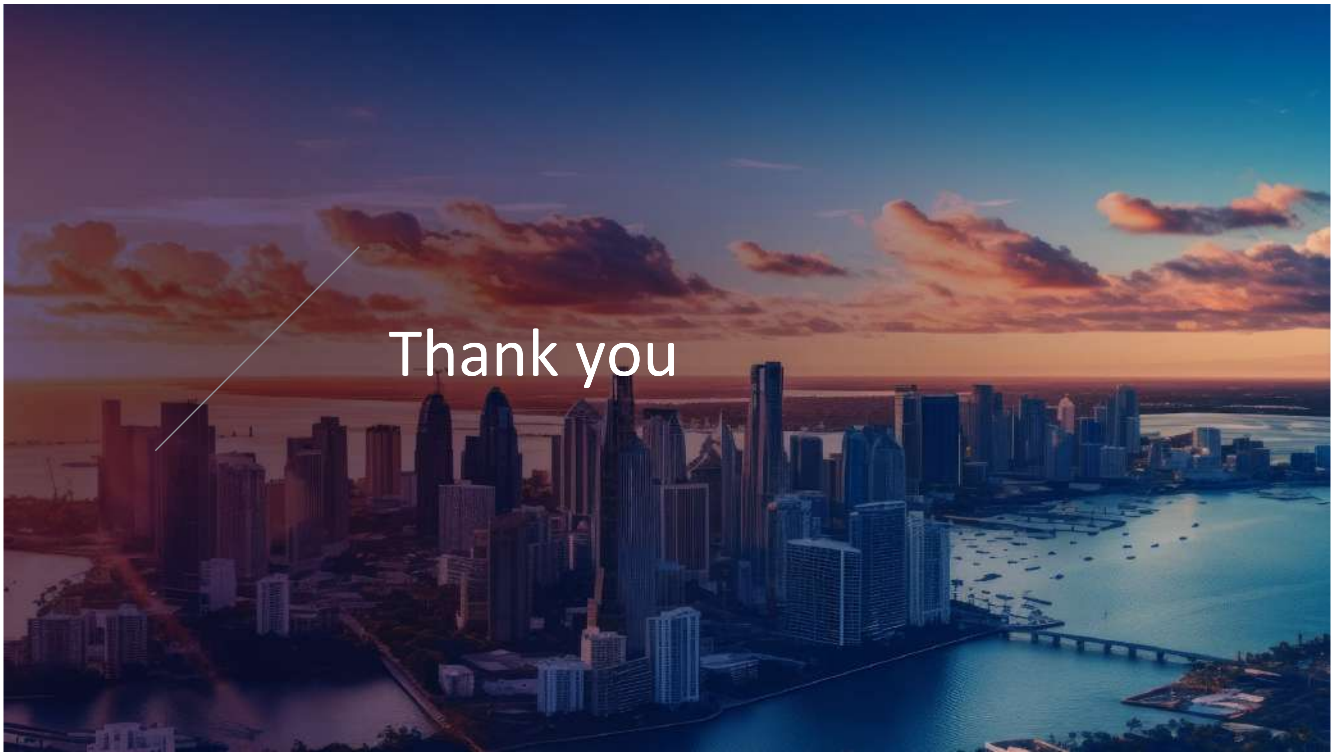
Thank you