How Machine Learning Delivers on the Promise of AI

Edge to Cloud AI
May 2025

Eric Siegel, Ph.D.
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).

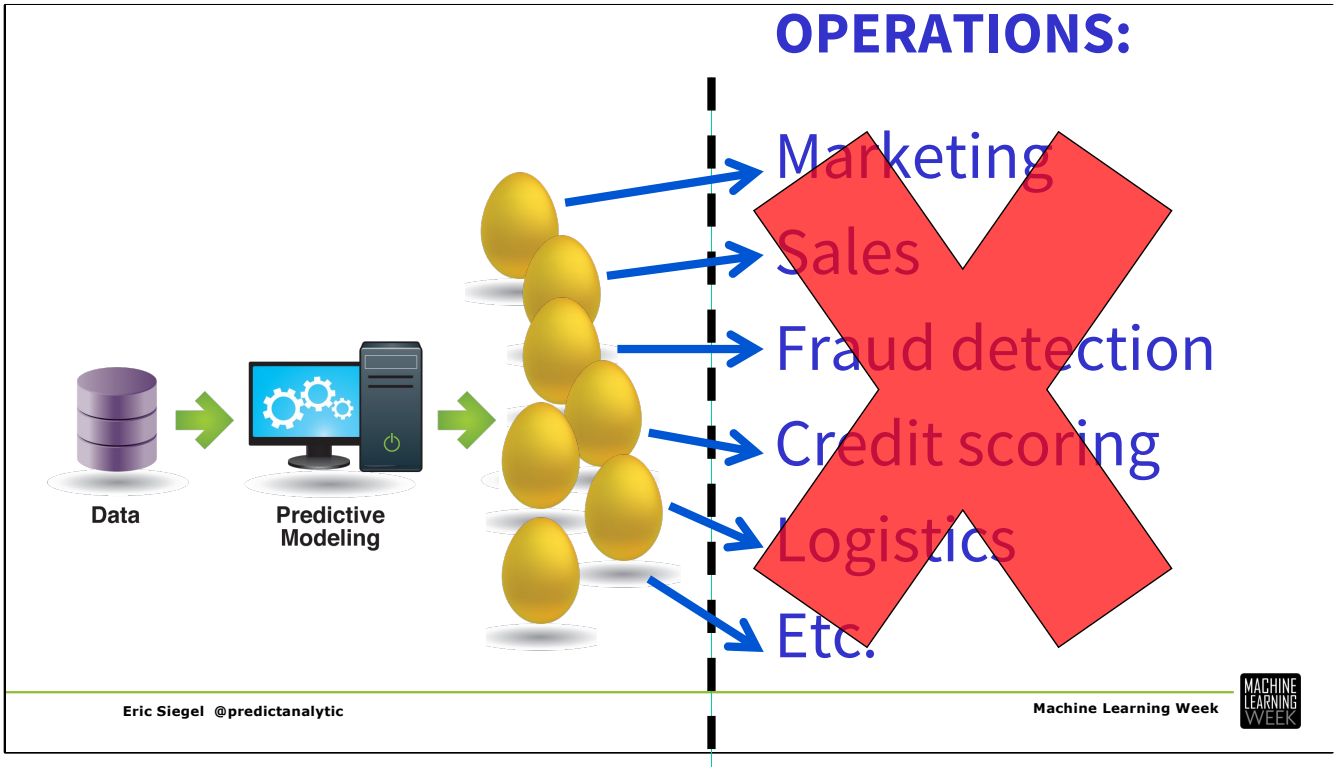
OPERATIONS:
Marketing
Sales
Fraud detection
Credit scoring
Logistics
Etc.
Data
Predictive Modeling
Eric Siegel @predictanalytic
Machine Learning Week
MACHINE LEARNING WEEK

Unmet need: a business playbook for running predictive AI projects

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*.

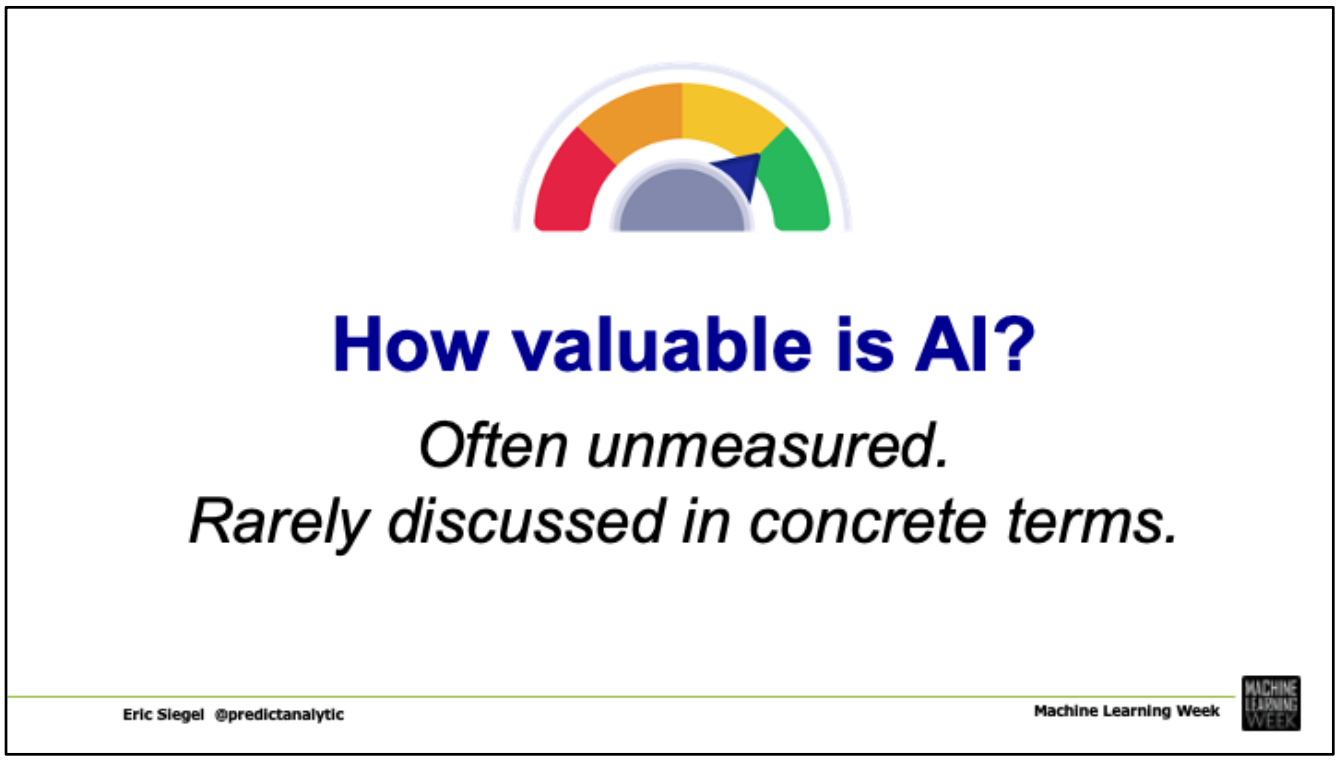

Despite the hoopla, AI's enterprise value goes unmeasured and unmentioned. It's the elephant in the room.

How well does it work? **A model's validation is as important as its training.**

After you've built something, you must assess how good it is – its potential value – before you use it.

Model performance is the whole point:

how well will it solve the business problem?
how will it pan out in deployment?
how much value it will generate?
how do you optimize its deployment?

The evaluation gap is an unfilled industry hole. It's as if "all" subject matter experts (data scientists) who potentially could recognize it and might know how to fill it are the same folks so accustomed to the status quo that they're blind to the hole.

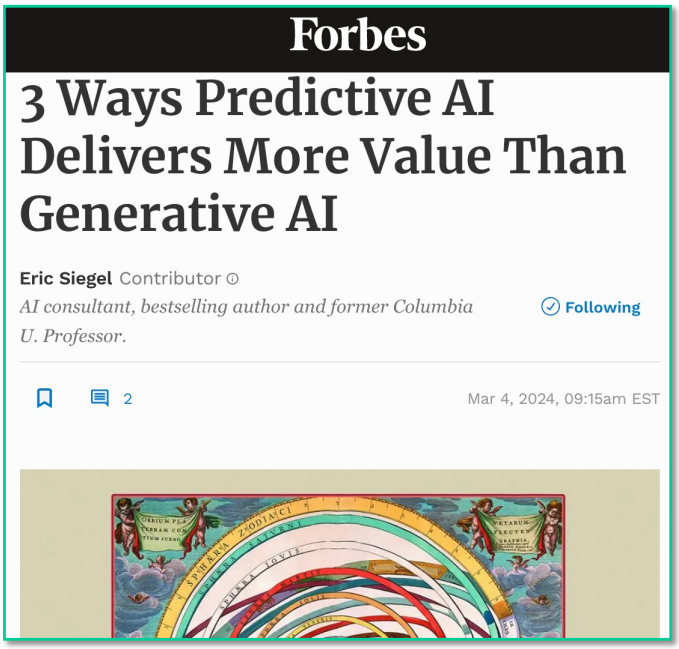

Forbes

3 Ways Predictive AI Delivers More Value Than Generative AI

Eric Siegel Contributor

AI consultant, bestselling author and former Columbia U. Professor.

Following

2

Mar 4, 2024, 09:15am EST

Eric Siegel @predictanalytic

Machine Learning Week

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

1) Predictive AI often delivers **higher returns** than generative

2) Predictive AI can operate **autonomously**, whereas generative AI usually cannot

3) Predictive AI is much **cheaper** and imposes a much smaller footprint than generative

Agenda

- Predictive AI's value proposition
- **PROBLEM:** The field is stuck on technical metrics
- **SOLUTION: ML valuation**

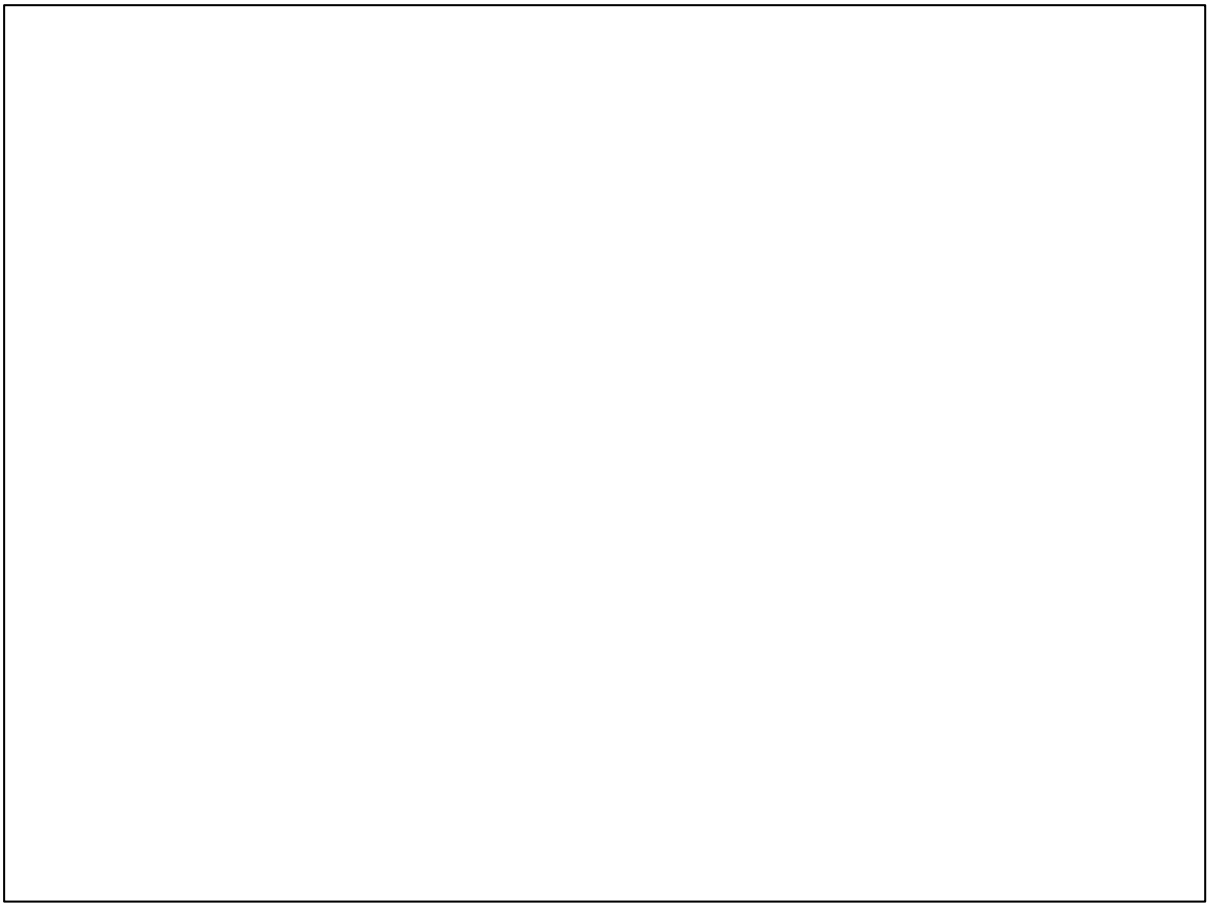

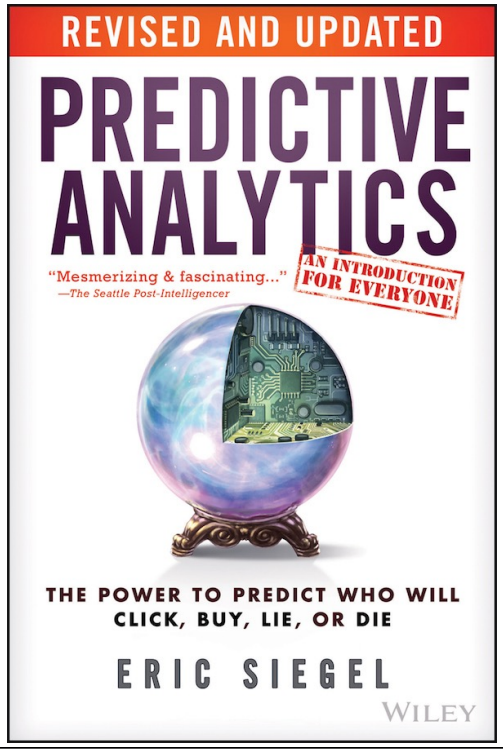

http://www.ThePredictionBook.com

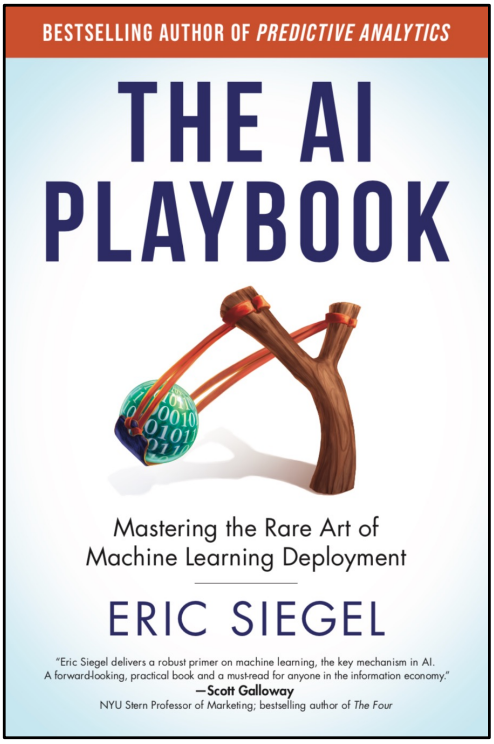

#bizML

http://www.bizML.com

https://machinelearningweek.com

https://machinelearningweek.eu

https://www.gooder.ai/

Today's model development is missing the ability to test and visualize the business performance – so it doesn't maximize for business objectives.

Gooder AI addresses a fundamental flaw with the typical model development process: It evaluates models in terms of business metrics like profit/ROI/savings, rather than only technical metrics like AUC/precision/recall.

Product overview document:
https://docs.google.com/document/d/1AVzEH9vo1pMExZ7V1Uwyf-pnnvX6kga9vEzhuwCQzkg/edit?usp=sharing

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

Validate any ML model for max value.

Universal ML validation to maximize value capture.

Empowers data scientists and orgs to optimize for the business objective.

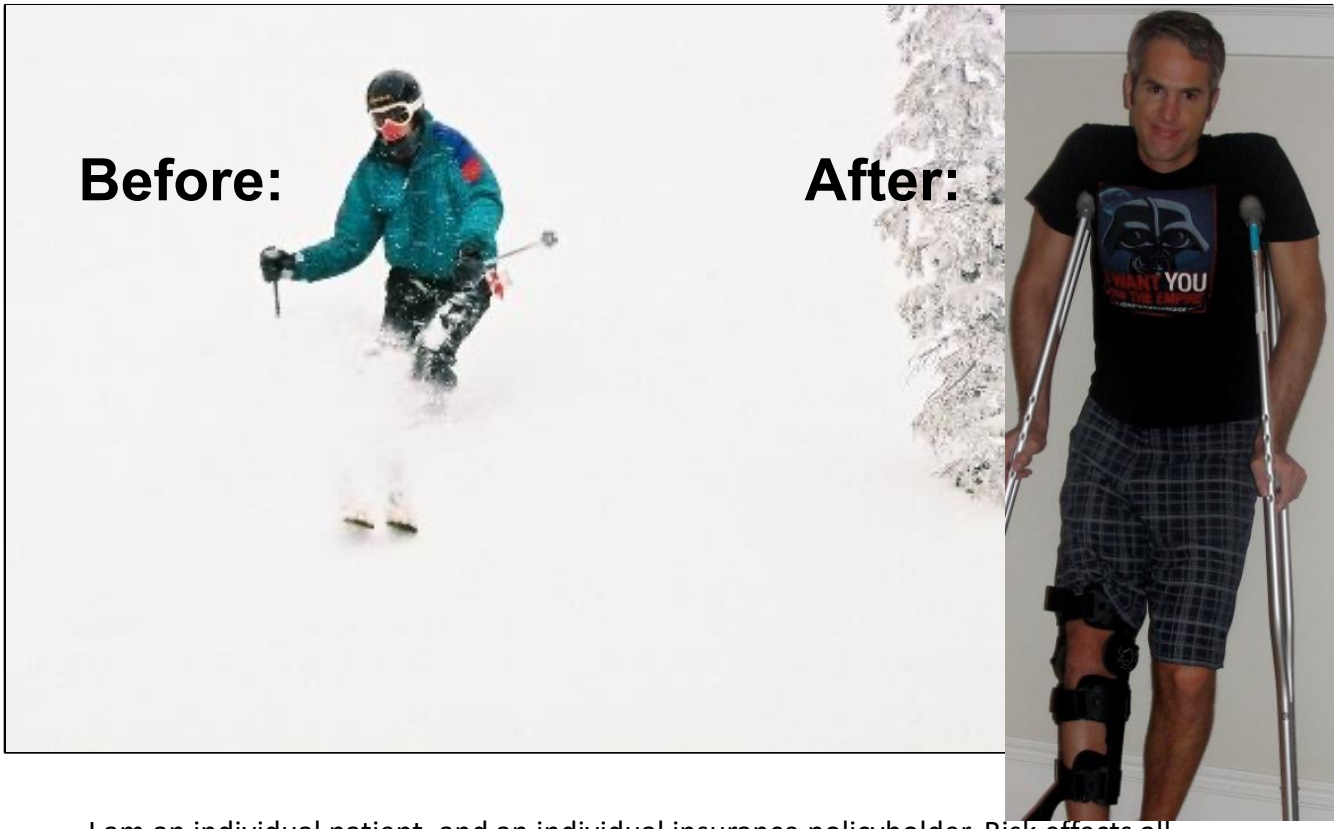

I am an individual patient, and an individual insurance policyholder. Risk effects all parties involved.

Knee Walking

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

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

"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

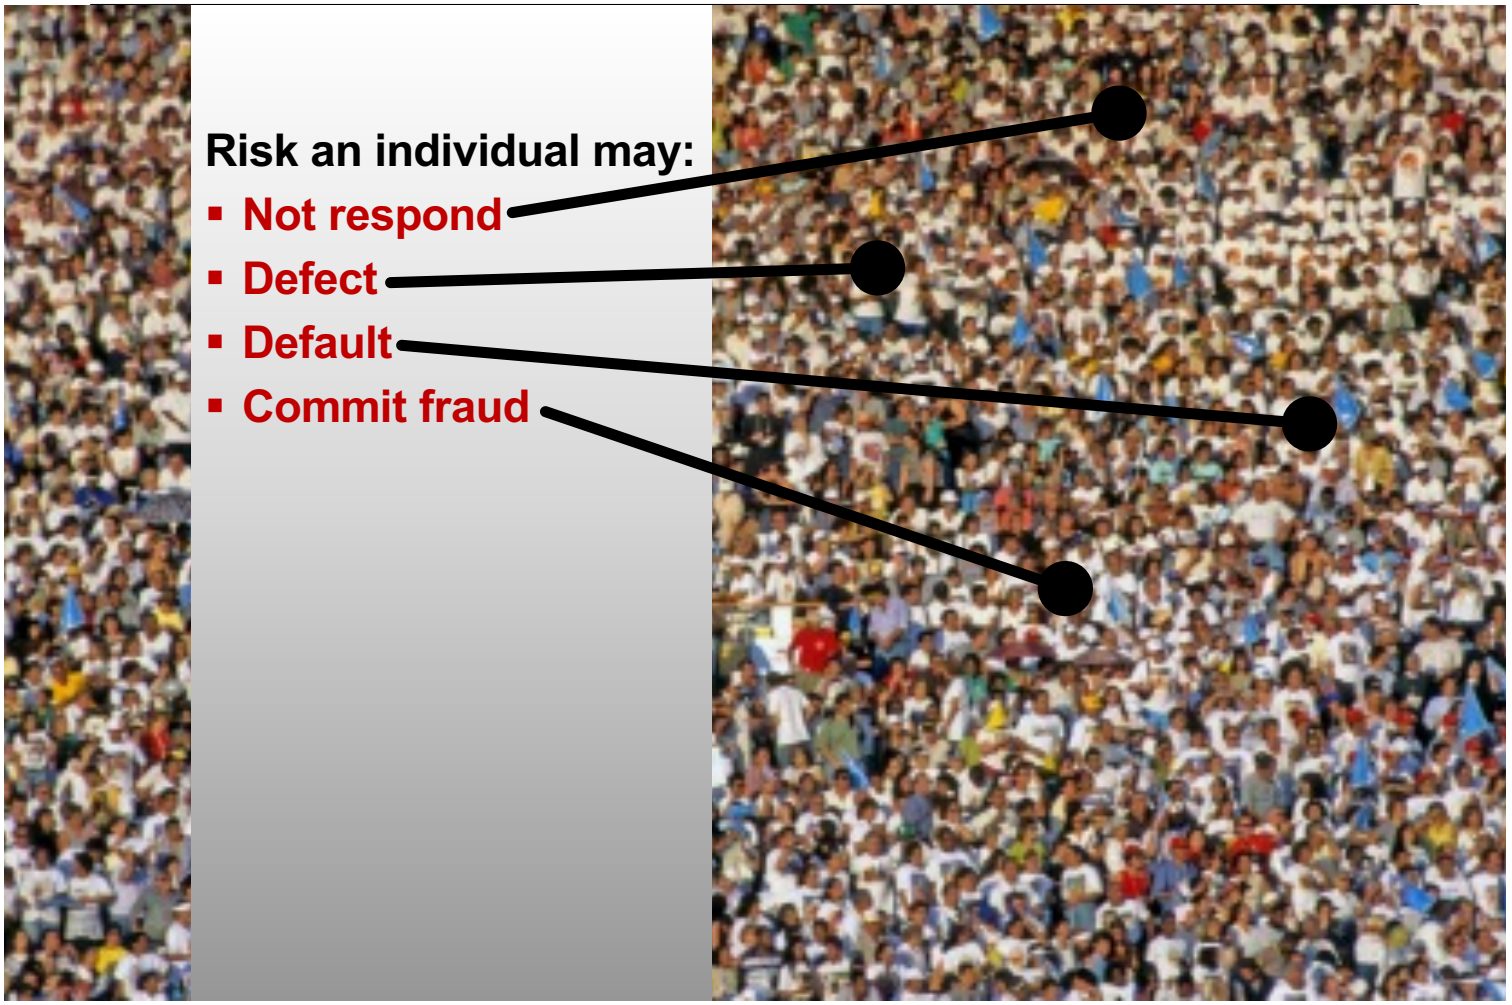
Risk an individual may:
Not respond
Defect
Default
Commit fraud


Tinker Toddlers
Machine Learning!
for Babies & Toddlers
Dr. Dhoot

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*

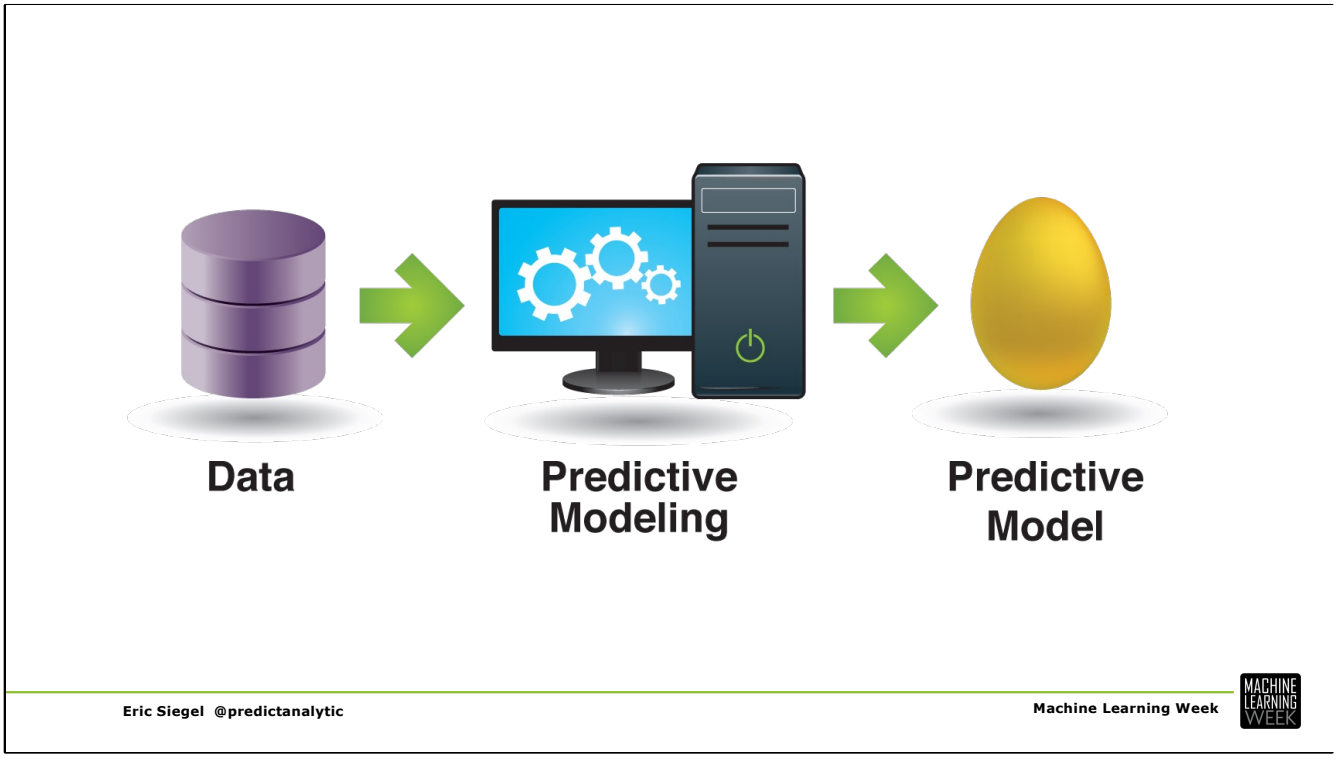

Predictive modeling learns from data in order to generate a predictive model. For details on how this works, see Chapter 4 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

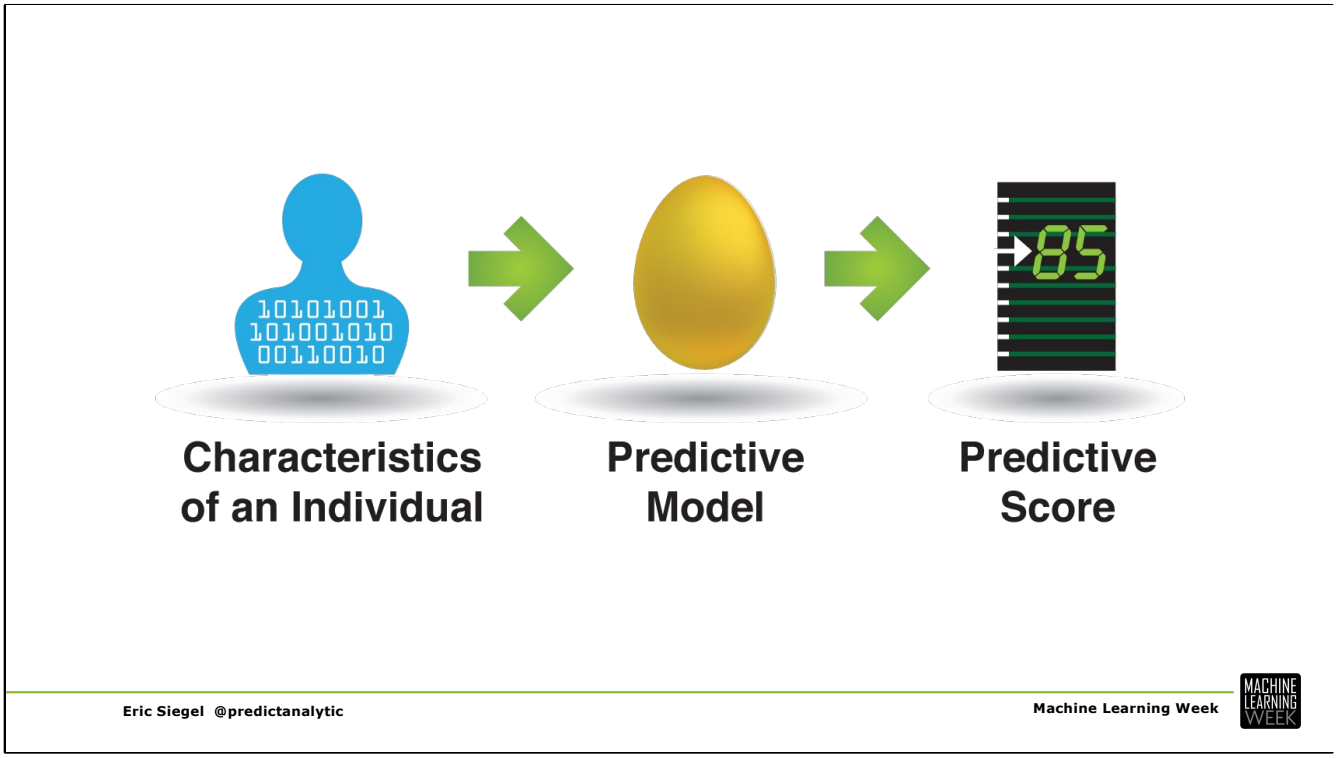

A predictive model generates a predictive score for an individual. For details on how this works, see Chapters 1 and 4 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

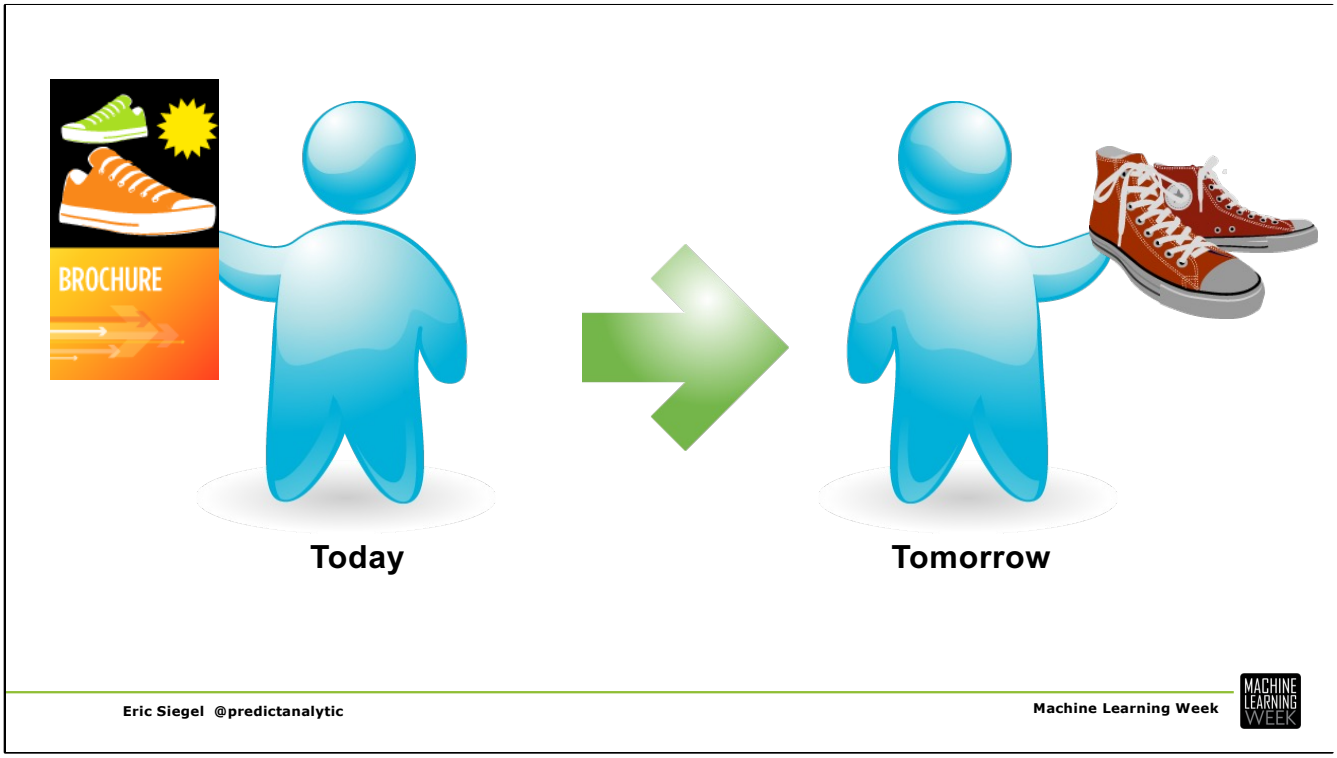

Marketing targets an individual predicted as likely to buy. For details on how this works see the Introduction and Chapter 1 of the book "Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die" (http://www.thepredictionbook.com).

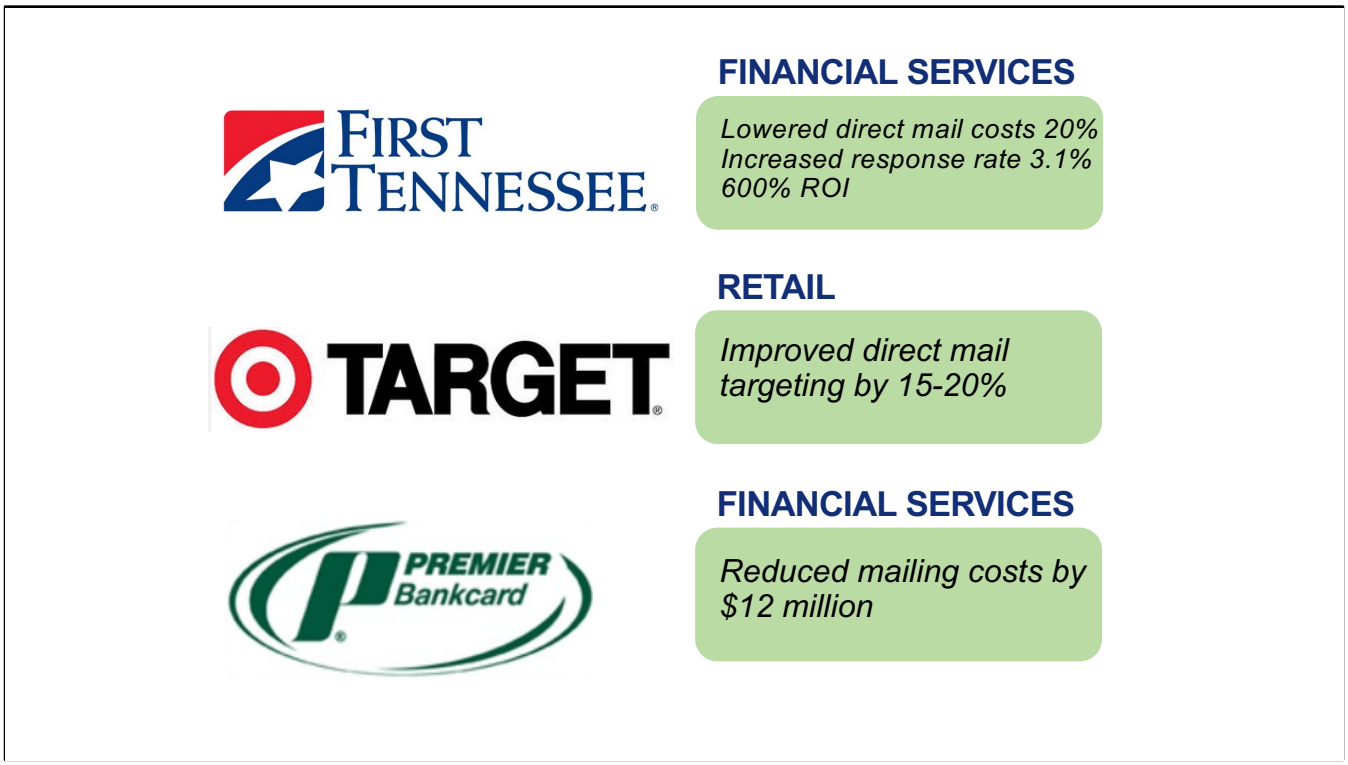

...and many more, such as Cox Communications, FedEx, Sprint, etc. - see the book "Predictive Analytics" (www.thepredictionbook) for many case studies, including a central compendium of 147 mini-case studies, of which 37 are examples in marketing applications of predictive analytics.

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

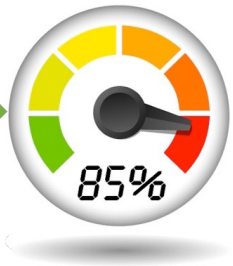

1) What's predicted

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

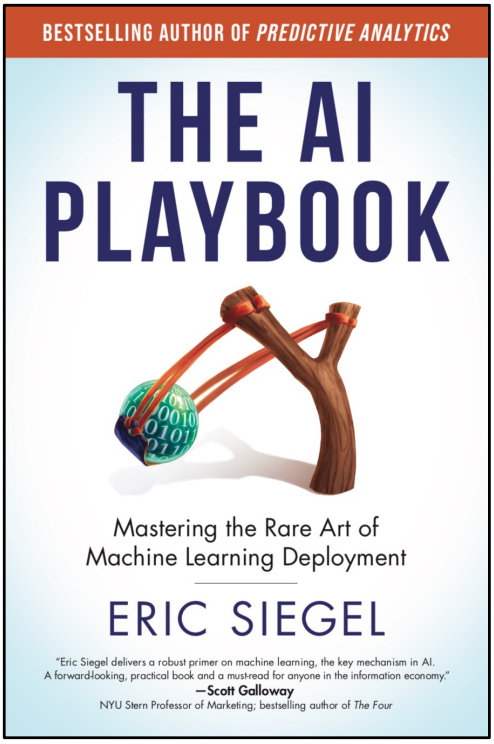

#bizML

http://www.bizML.com

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.

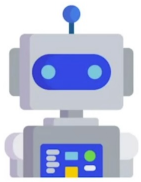

Technical metrics:

- Precision
- Recall
- AUC

Business metrics:

- Profit
- Savings

ML evaluation → ML valuation

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


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

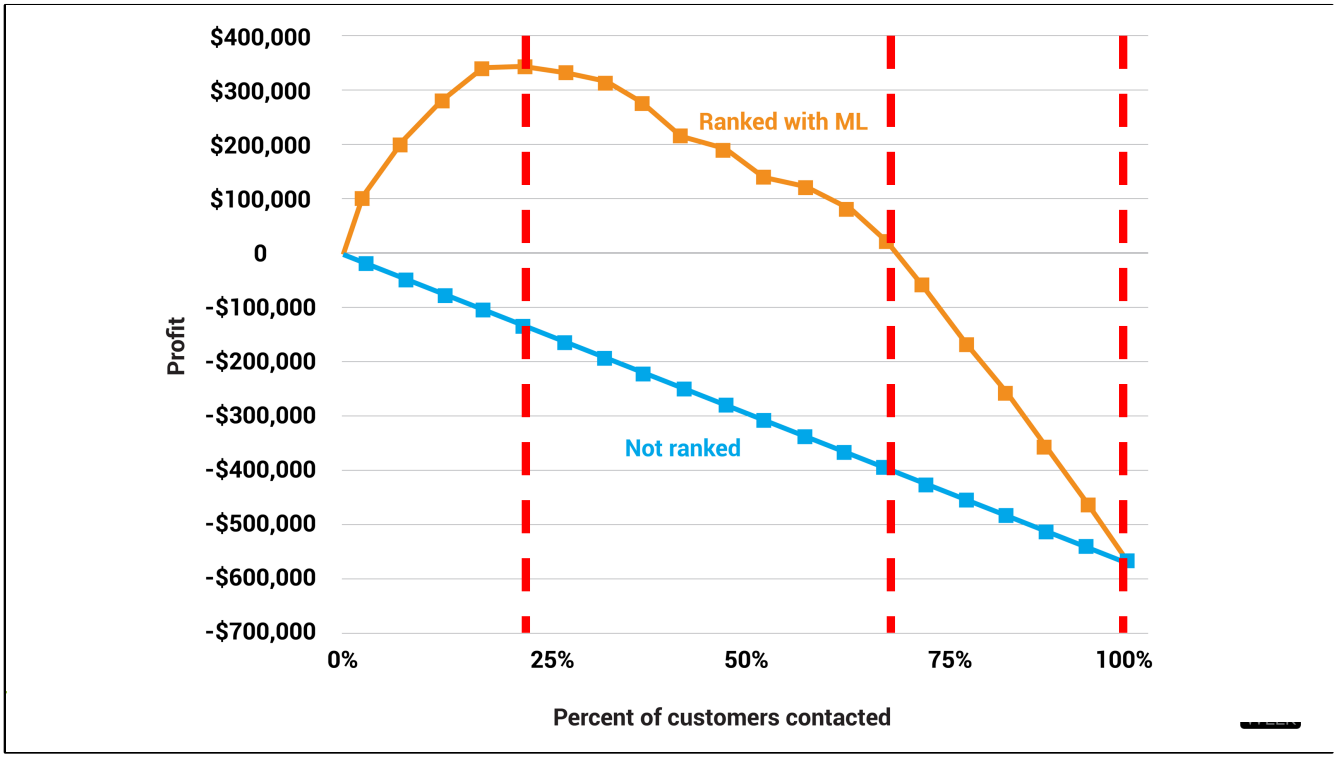
$400,000
$300,000
$200,000
$100,000
0
-$100,000
-$200,000
-$300,000
-$400,000
-$500,000
-$600,000
-$700,000
Profit
Ranked with ML
Not ranked
0%
25%
50%
75%
100%
Percent of customers contacted

GOODER AI™

Targeting marketing

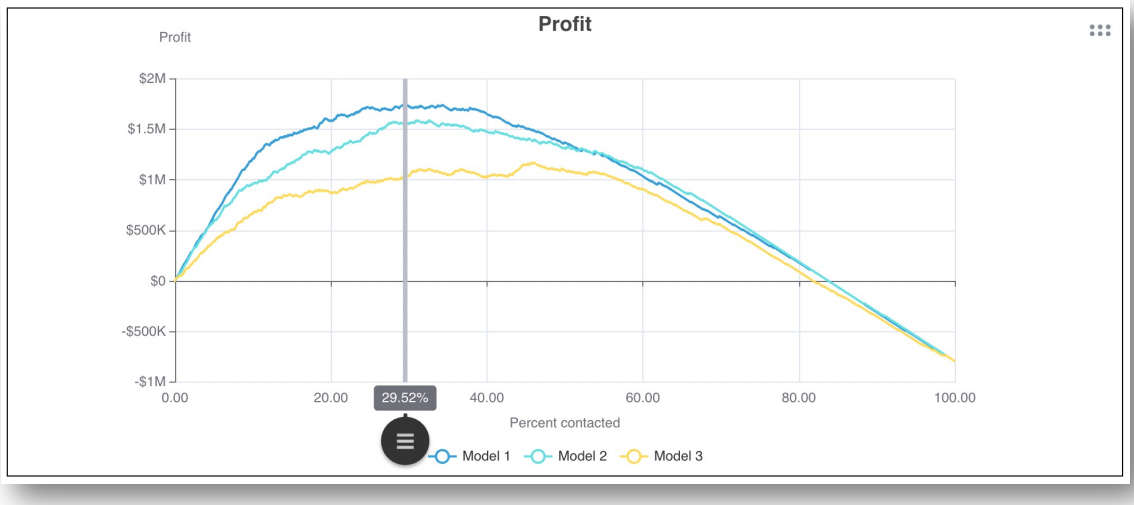

GOODER AI™

Nonprofit fundraising

A large nonprofit hospital network


3%
25%
50%
75%
100%

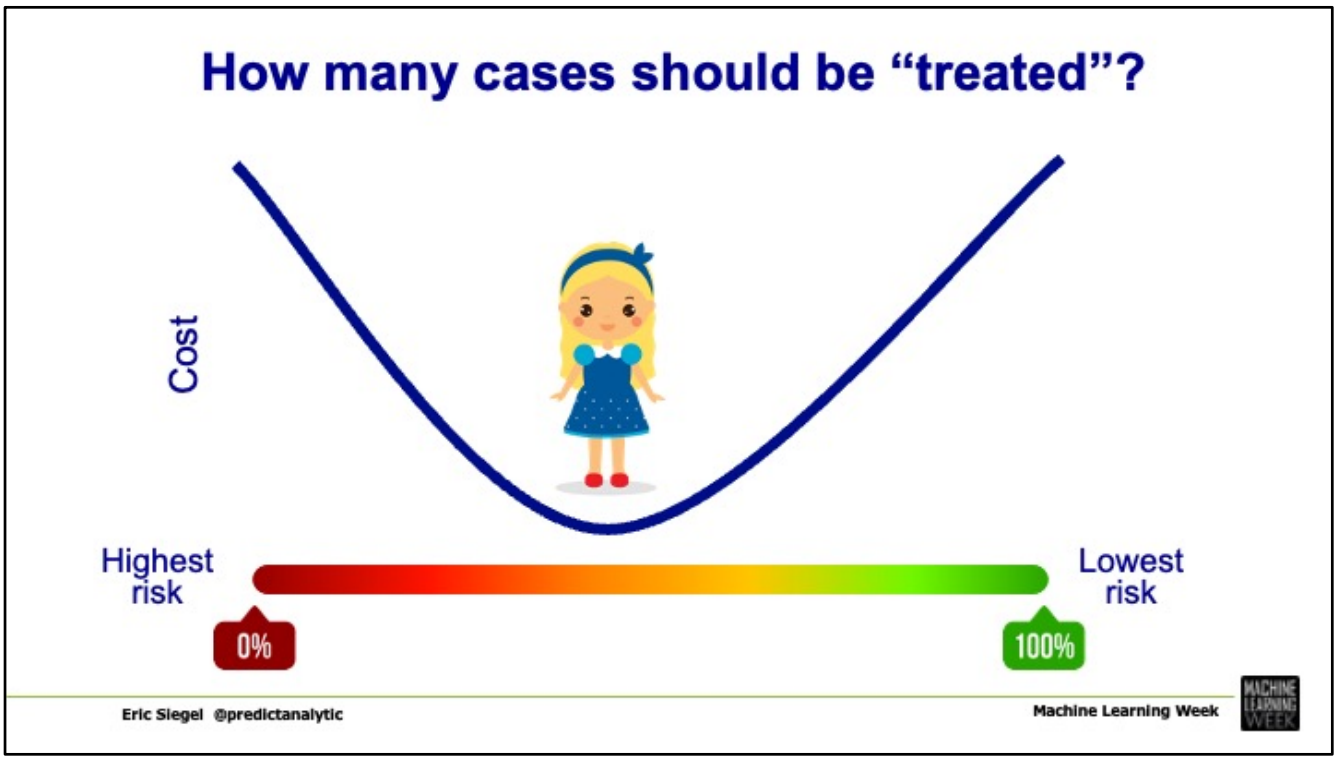

You can't manually audit every single post -- that's too expensive.
Where do you draw the line? How do you strike that balance?
There's no existing solution for navigating that decision.

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.

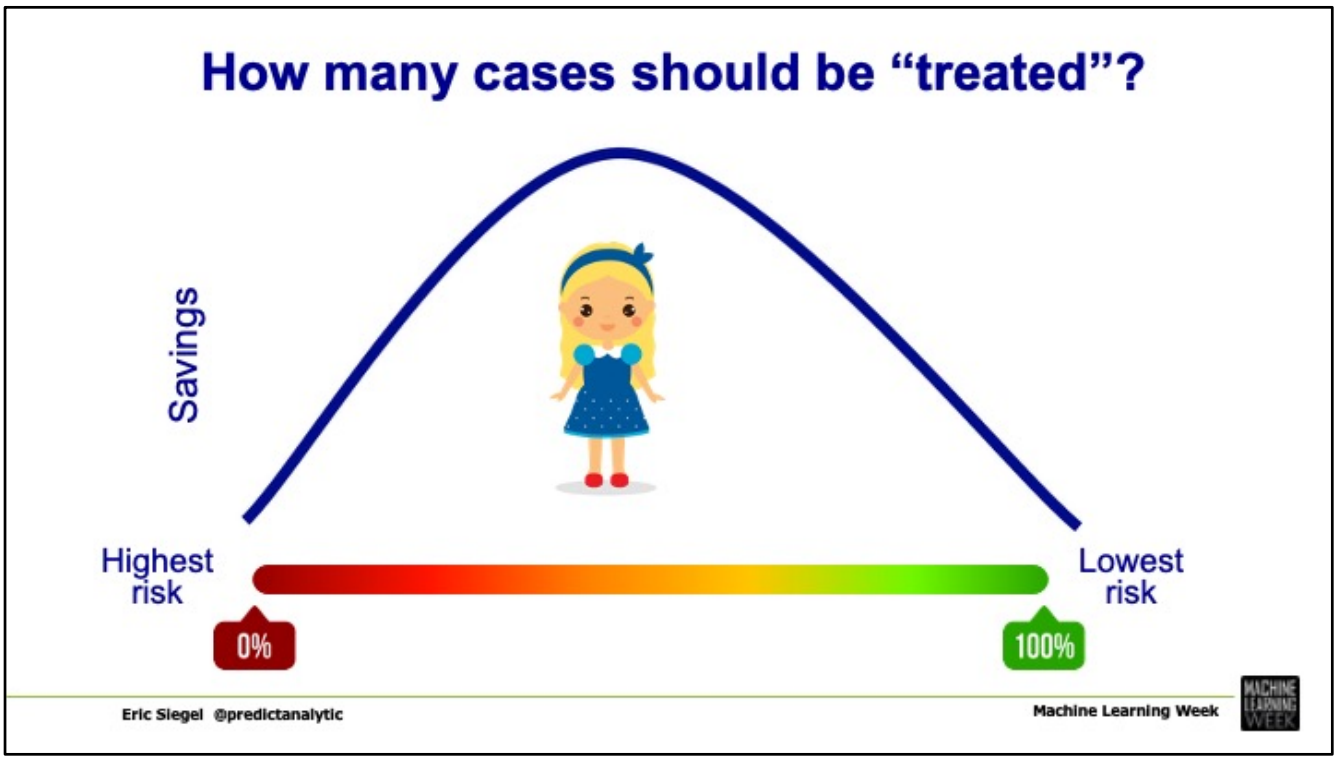

You can't manually audit every single post -- that's too expensive.
Where do you draw the line? How do you strike that balance?
There's no existing solution for navigating that decision.

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™

Fraud detection

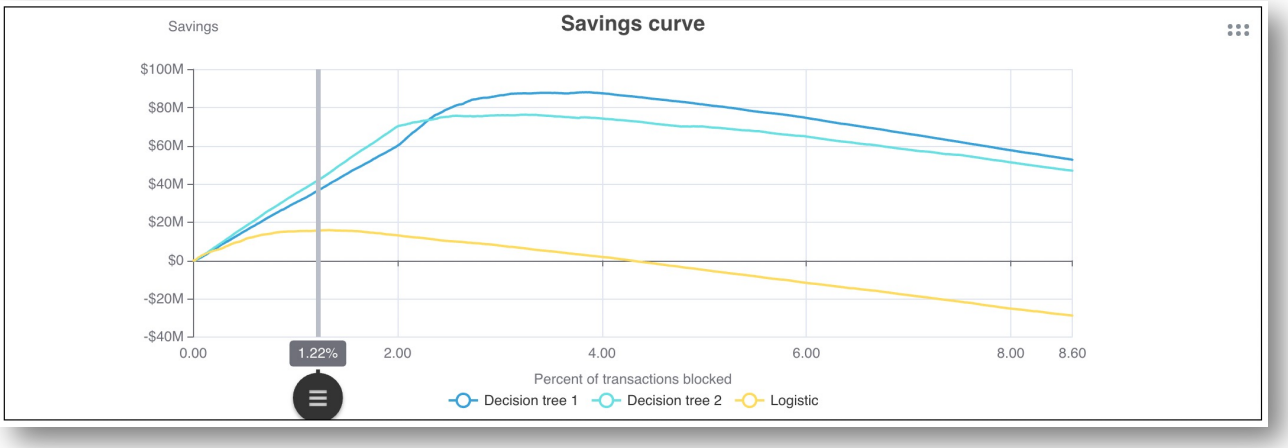

GOODER AI™

Credit scoring

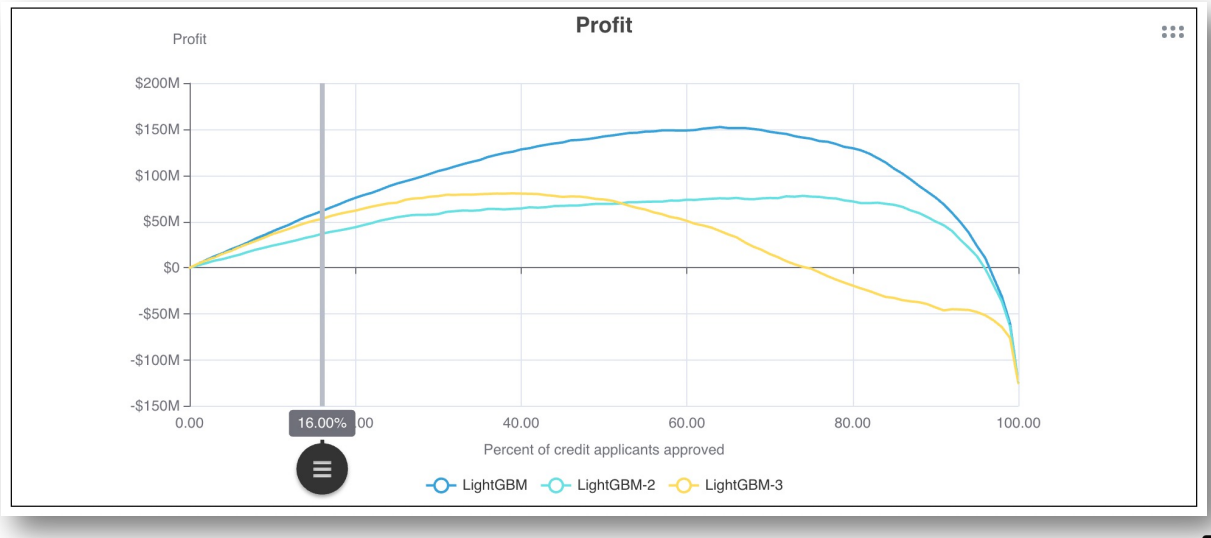

GOODER AI™

Shopping cart abandonment

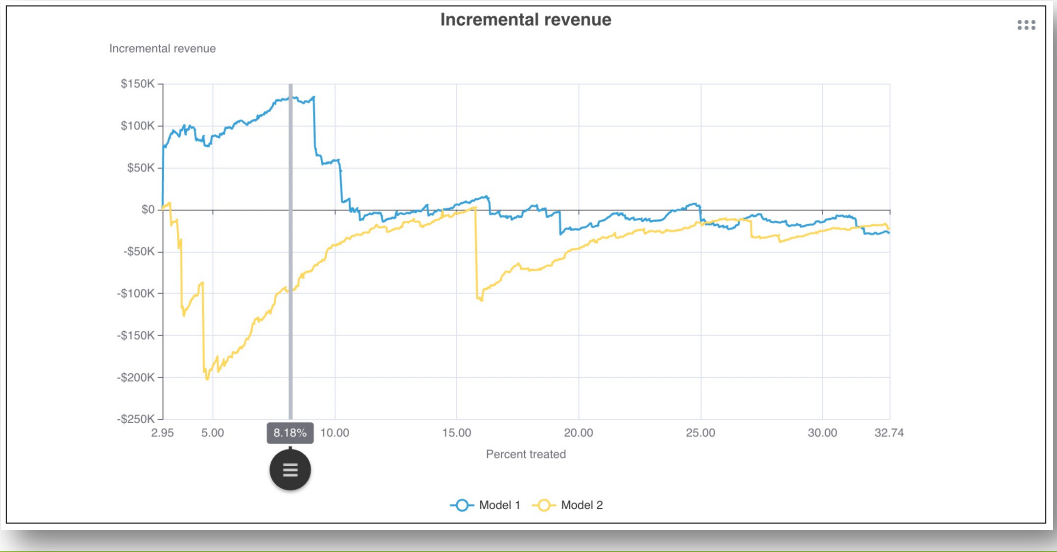

GOODER AI™

Misinformation detection

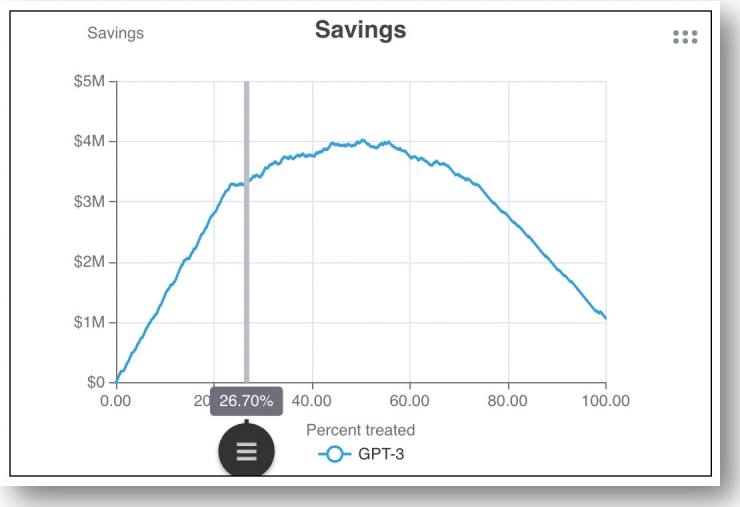

Checking fraud

Email marketing response

E-commerce discount targeting

Microloan risk

Collections pay-backs

Financial donations

Blood donations

Healthcare patient no-shows

Customer attrition

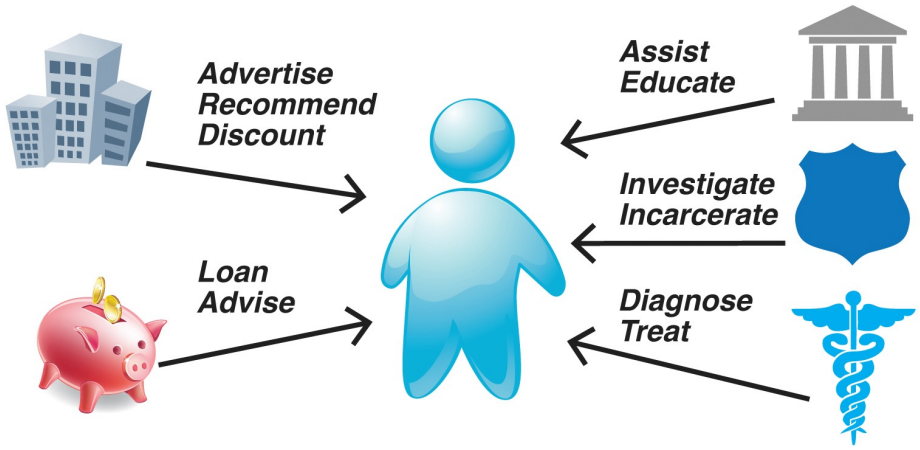

Millions of decisions a day determine whom to ***call, mail, approve, test, diagnose, warn, investigate, incarcerate, set up on a date, and medicate.***

Conclusions

- If you're not measuring value, you're not pursuing value
- ML evaluation → ***ML valuation***
- Gooder AI is the first OTS solution

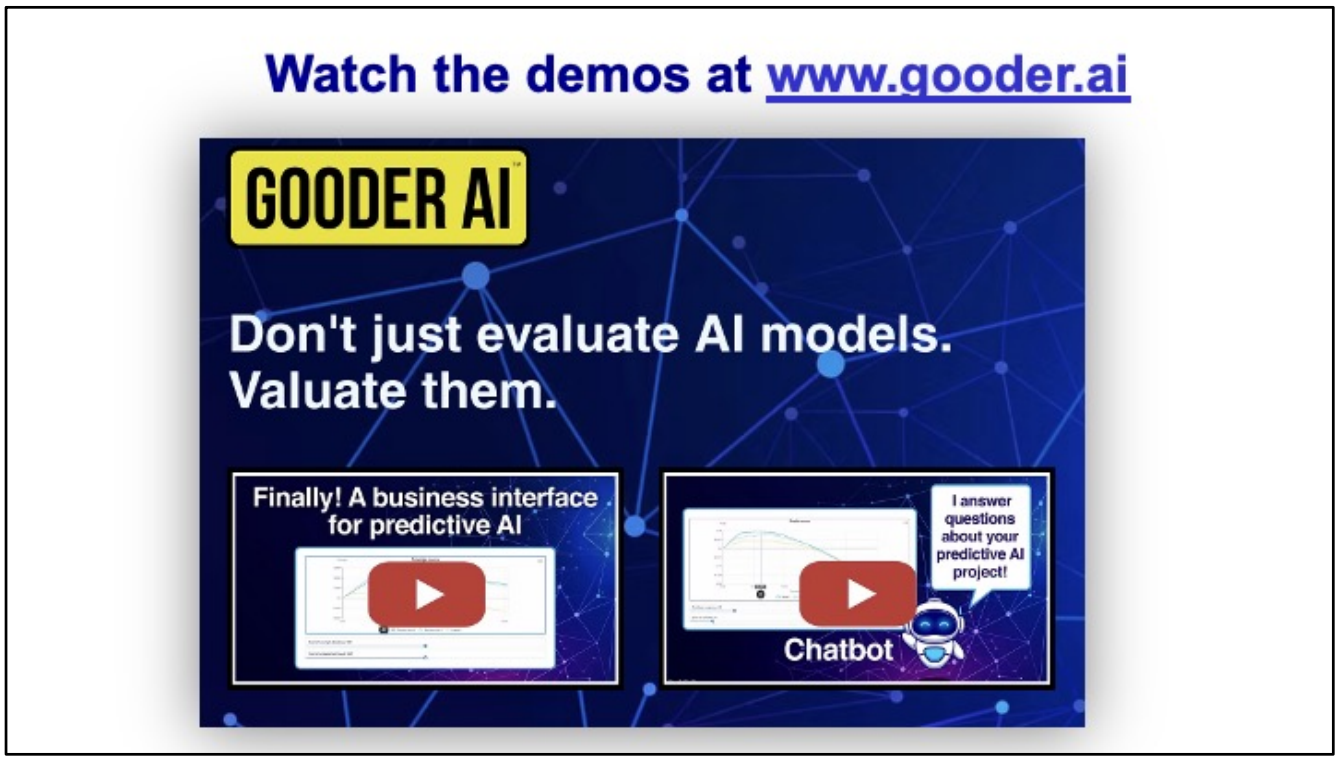

https://www.gooder.ai/

More content from Eric Siegel

MACHINE LEARNING WEEK

Talks at Google

MACHINE LEARNING Times
The machine learning professionals' premier resource

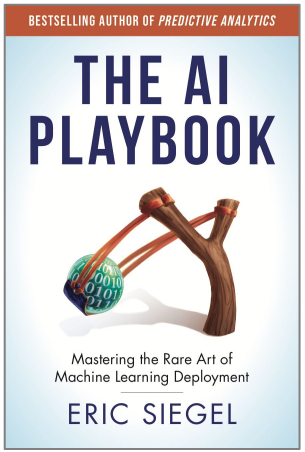

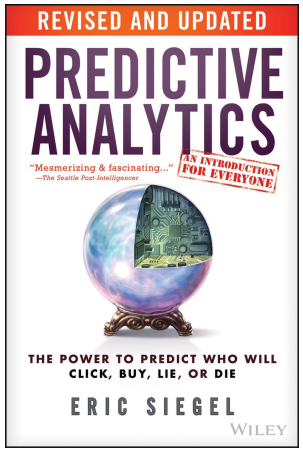

Forbes