HARVARD | BUSINESS | SCHOOL

Two AI Cases.
Amgen and DeepMap

Shane Greenstein

HARVARD | BUSINESS | SCHOOL

Thanks for your attention

I joined Harvard Business School six years ago. Part of the unit called Technology Operations and Management.

Today I teach a course called “Digital Innovation and Transformation.” Marco Iansiti, Karim Lakhani, Feng Zhu have also taught it.

It is a case-based course. 20 cases. 4 blog posts. 1 simulation.

It contains four modules:

- Internet basic
- Platforms
- Managing big data
- Commercializing AI

Today: two cases from the last module. I wrote both.

Amgen

Digital Manufacturing at Amgen, HBS case N9-621-008. February 1, 2021.

The setting: biotech production and batch processing

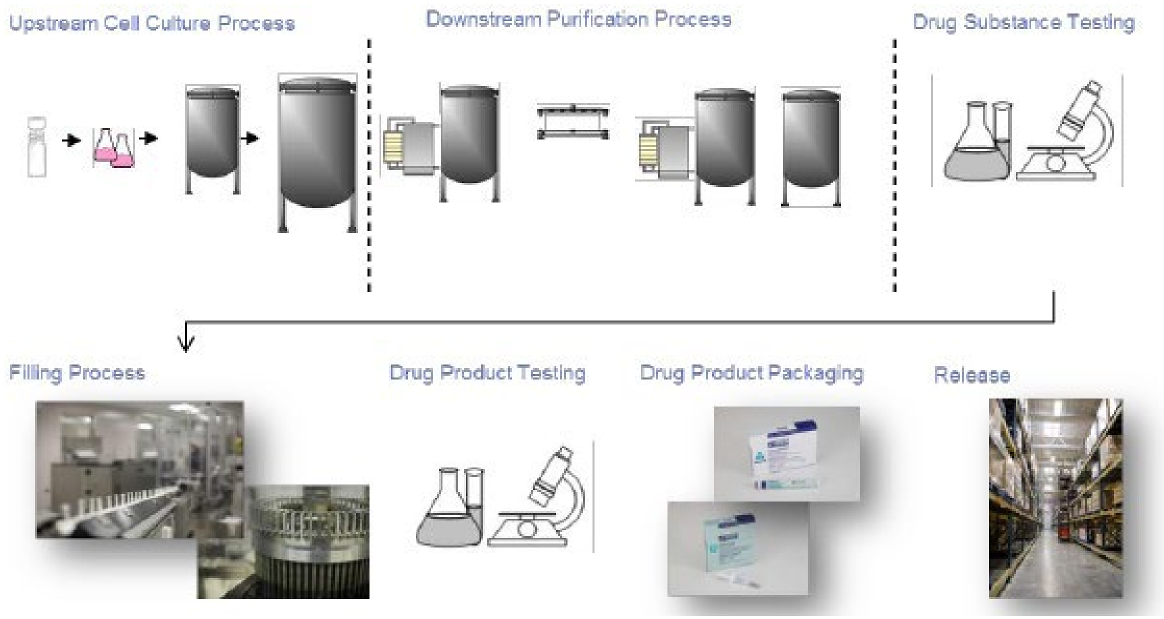

Students view videos ahead of class. Amgen makes many available.
Biotech is just like brewing. (Students like this joke.)
No mistakes are allowed. FDA regulation is extremely tight.

Let's walk through a case discussion..

The case dilemma

Dilemma: Set in fall 2020. Amgen has a long standing program to install sensors with statistical models. With the cooperation of production floor the ROI is positive and high. Still plenty to do throughout Amgen's production facilities across the globe.

- Have begun to experiment ML models.
- The ROI on ML sensors is not high in short run. Perhaps even negative.

There are only so many personnel to go around. Case exams the perspective of statistical team who supports production across the globe. What are the priorities? (a) More sensors in a site that has none or (b) experiments with ML in a site that is willing to experiment?

Begin with a homework that simulates the gains/costs to installing a sensor.

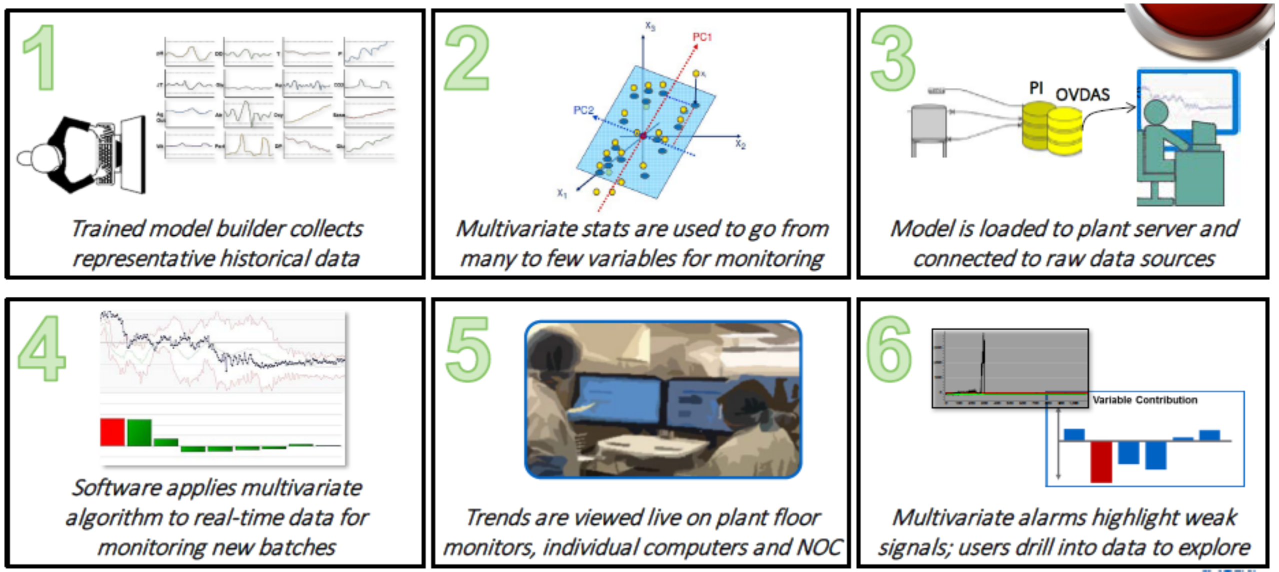

A statistical sensor trigger alarms/lights, which alert staff to check production. Model development requires back and forth between statistical team and floor managers and employees.

A sensor with machine learning resembles a statistical model with one key difference. The algorithm can initiate action, such as open or close a valve, or change temperature settings.

How does a sensor lead to lower costs & higher revenue? Saves time? Reduces failure?

The first program – with only statistical models – installs a statistical model that lowers the failure rate. ROI is positive.

The second program -- with an ML algorithm – saves time but increases the failure rate to avoid false positives. ROI is borderline.

Sensor with lower failure rate	Sensor that saves time with higher failure rate
More capacity from lower failure rate	More cautious. No false positives!
More output	Savings on waste
Lower waste	Should we optimistic or pessimistic
Overwhelming	multiyear
Quantitative experienced	New skills
	Cultural change
	Learn from data

What factors could help a program for installing statistical sensors at a greenfield facility?

AMGEN

Development of models	Initial installation of sensors	Extending use

What factors could help a program for installing statistical sensors at a greenfield facility?

AMGEN

Development of models	Initial installation of sensors	Extending use
Improve upon existing productivity	Repeated and reliable	Value to process
Good relations with floor employees	Within tolerances that are acceptable	Economic justification
Lots of data from production	Alarm fatigue reduced….	Tell them things they did not know
Experiment and tweaking	Clear process for intervention	Training.
	Execution challenges are addressed	Not flavor of the month
	Getting buy-in again	
	Leadership support	
	Value-added is more obvious	
	Training.	
	Complements existing work flow	

What ***additional*** factors could help an effort to install machine learning (ML)?

Development of models	Initial installation of sensors	Extending use of sensors

What ***additional*** factors could help an effort to install machine learning (ML)?

AMGEN

Development of models	Initial installation of sensors	Extending use of sensors
Start simple. Hard enough.	Save time and apply elsewhere	Over time as accumulates
Pilot, recognition from floor	Find right opportunities, low downside risk	Crucial for model to work and learn
Team needs to be on board with experiments	Credible of sensor	Incentives of managers at plants – want to prevent failure, not try something new
	Need regulatory buy-in	
	Need to explain the black box	
	Need on intellectual property	
	Get over proven and trust	
	Enhances work. Perception.	
	Nobody will lose their job tomorrow. Perception.	

H

What challenges will shape the success of an operational strategy for implementing digital technologies across all of Amgen? Priority?

AMGEN

Options	Reasons for making it a priority	Challenges & approach to them
P.R.		
Greenfield		
RT-MSPM		
R.I.		
Experienced		
ML		

What challenges will shape the success of an operational strategy for implementing digital technologies across all of Amgen? Priority?

Options	Reasons for making it a priority	Challenges & approach to them
P.R.	Can also dip your toes in	Unseen anomalies
Greenfield	Hub for more testing more alarm	Low cost of failure
RT-MSPM	NPV positive	Selection of members
	Enthusiasm?	
R.I.	Preserve culture	Worry about labor force, train, manage
Experienced	Financial risks	Top product, relationship, observe
ML	Innovation going forward, upside	Learn from failure, false positives concern
	More data, leverage	Commitment? Selection of team members
	Enthusiasm?	

Takeaways

AMGEN

Biotech production contains a unique cost structure due to its rigid production requirements. “Its just like brewing, but with no errors.”

Calculations of ROI seem straightforward, but are not everything. The value of ML involves more nuance. Additional strategic priorities and adoption challenges must play a role in decision making.

Digitization of internal processes cannot solely be top-down. Adoptions experiences go far better with cooperation from the floor.

Automation takes place amidst a workforce of difference skillsets. In the short run ML enhances existing processes, but –irrespective of the facts – every worker believes it will become a substitute at some point.

The situation requires restructuring of development efforts. Should Amgen create new assignments and jobs devoted to facilitating adoption?

DeepMap

DeepMap: Charting the Road Ahead for Autonomous Vehicles.
HBS Case 9-620-047, November 1, 2019.

DeepMap

Founded in 2016 by veterans of mapping efforts at other firms.

- Anticipated safety will require high resolution maps for L4 & L5.
- New thinking (2021) suggests L2 and L3 may need high resolution.

Business approach to gain OEM trust.

- Customer owns the data/map.
- Long run SaS model of $ per mile. “An SAP/Oracle for OEMs.”

Doing well. 190 employees in spring 2019. 240 today (at fifth bday).

Students to watch a video ahead of time about AVs.

.

The case dilemma

Dilemma: Most potential partners are initiating prototyping efforts., but despite great funding, DeepMap is a young & resource-constrained firm and cannot take all potential partners.

- Why? Immature technology requires months of customization.
- In practice, several engineers dedicated to each customer until the customization is done. May require more later.
- Only so many engineers to go around. Time is precious.
- DeepMap aspires to be important player in future supply chains, and desires to take actions today to raise that probability.

Which partners? Why? What questions should management ask when having conversations with potential partners?

Let's walk through a case discussion.

What are the consensus forecasts for prototypes for L4 AVs?

	Robotaxis	Last Mile Delivery	Long Haul Trucking	Mining, Agriculture
Time to first major deployment?				
How does the service work?				
Value proposition?				
Primary challenge?				

	Robotaxis	Last Mile Delivery	Long Haul Trucking	Mining, Agriculture
Time to first major deployment?	Five years or sooner?	Five years	Two year	A couple years, some already now
How does the service work?	Uber & so on.	Pizza! Groceries.	Straight highways	Operate tractors
	Many firms.	Little robot dudes?	Easier technical	Operate trucks
		Not cars?		Specializing.
				Underground
Value proposition?	Heavy traffic	Safety.	Lots of data. Repetition.	Labor intensive. Efficiency
	More data	Large volumes	Caravan.	More accurate
	Controlled environment	Save money?	Reduce drive errors	Less complexity.
				Map contributes right away
Primary challenge?	People safety & experience.	Competition.	Technical tough.	Fewer people, less risk, injury
	Onboard & off	Crowded space	Heavy, less flexible	Truck size
		populated area		Risk/capital/$$$$

What are the distinctive features of DeepMap's service?

Distinctive DeepMap feature	Appeal of feature	How DeepMap operations supports

What are the distinctive features of DeepMap's service?

Distinctive DeepMap feature	Appeal of feature	How DeepMap operations supports
Update frequently	Correct quickly	
Human checkers in the operations	ML/AI not safe yet, and no prospects of it w/o 3D maps	Staffing. Staffing. For how long?
Software as service. API at boundary	Adjust to use cases. Adjust to hardware.	
Some customization for client.	At this stage. Make it work. Proof of concept.	Many engineers. Expertise.
Communications platform, cloud based	Share data with ops, vehicles, and others. Learn updates.	
Variable updates for the applications. Different pacing.	Customizable to aps.	
Customers own the data.	Value chain can chain. The customer wants the data. It might simplify regulatory issues.	Not a technical decision. A bus decision.

What is a realistic best case and worst case scenario in the next year with a new customer?

Best case	Worst case

What is a realistic best case & worst case scenario in the next year with a new customer?

Best case	Worst case
Long term commitment. Serious.	Very short term focused.
Principles.	Hardware will change a lot. Do not want to be stuck changing a lot.
Not too much customization to win their business.	Stretch to 10 to 15 years. Takes too much funding.
Need this technology for the long term.	Very risk averse investors.
Customers with application to deploy sooner.	Limited infrastructure, depending on the apps.
Limited co-innovation risks. What are the priorities? Will they be risk averse? Want them charge charge ahead.	Not on the verge financial risk.
Propitious location. China? Supportive infrastructure.	Poor infrastructure
Limited area for applications. Fenced areas.	
Big company. Big R&D budgets. Can survive issues.	Small R&D
Density of setting, where data adds value.	
Big fleet.	Small fleet.

What to say to a prospective partner? What should DeepMap think about when takings prospective partners in China and Europe?

What makes a partner attractive at early stage?	What question should DeepMap ask of a partner?

What to say to a prospective partner? What should DeepMap think about when takings prospective partners in China and Europe?

What makes a partner attractive at early stage?	What question should DeepMap ask of a partner?
Integrating into project for core project.	Get a sense of commitment: size of budget, number of people, strategy vision.
Lots of data. Plays to their strengths.	Who else are they talking to? How transparent?
	Are they aspiring to bring it in-house? Short term.
	Gauge customer needs, and the customization requirement.
LET'S ASK OUR GUEST	LET'S ASK OUR GUEST

Wrap for the DeepMap case.

It is fun and interesting to learn about AVs, but for those in the business it is essential to lose the starry-eyed views. Commercialization strategy requires same type of sober thinking seen in other nascent markets.

AVs will be a big user of ML someday, but right now prototyping is still the bulk of activity. Like any nascent market, there are a variety of forecasts for the likely paths that product markets can take.

Firms aspiring to occupy key positions in future supply chains must maintain a portfolio of options. Like any portfolio, make a deliberate choice over a balance of risk and return among partners.

Expect learning and adjustment as ML and AVs evolve. The portfolio is not static, and must be consciously updated.

Thanks for your attention

Happy to send teaching materials to anybody who asks for it.

Happy to address questions offline.

Any questions for now?

Extra slides

Amgen videos

Take a walk through the biologics manufacturing unit operations

- Clean Environment (Links to an external site.)
- Cell Culture (Links to an external site.)
- Purification (Links to an external site.)
- Testing (Links to an external site.)
- Fill, Finish and Packaging (Links to an external site.)
- Biotechnology by Amgen (Links to an external site.) [detailed virtual tour]
- The Challenges in Manufacturing Biologics (Links to an external site.)
- Quality Processes (Links to an external site.)
- Innovation (Links to an external site.)

Amgen Homework

Option 1: Building an RT-MSPM Model for mAbX

Goal: Reduce the batch failure rate. (Note: Reducing the batch failure rate would be only one possible outcome; the model would also likely detect additional anomalies throughout the manufacturing process.)

Setting: You must decide whether to build a statistical model to reduce the failure rate of a batch of mAbX. Each batch requires 10 processes and in total takes three days.

A plant runs at full capacity if it operates 365.25 days a year (the extra .25 comes from leap years). A plant can sell anything it produces. Each batch yields products worth $3.25 million in sales.

Amgen Homework

Production for each batch consists of 10 steps. The last five steps are delicate, and each has a 1% failure rate – i.e., 1 in 100 batches fail at one of these points in the process. Failures in one step are statistically independent of every other step, so the failure rate for a batch is 5%. It is not possible to tell whether the entire batch is lost until the end. A quality assurance test tells the plant manager whether the batch is good or bad.

Each of the 10 steps takes seven hours. The test must be performed in a lab and takes one hour. After a successful test, all equipment is cleaned, which takes one hour. An entire successful cycle takes 72 hours, after which the cycle can be started again. After a failed test, all equipment is cleaned for 5 hours. An entire unsuccessful cycle takes 76 hours. After cleaning, the cycle is started again.

Amgen Homework

Some cost is incurred no matter what the plant produces (i.e., plant labor, electricity, maintenance, plant security, debt for the structure, etc.), which we call "overhead costs." The overhead costs are $120 million per year. Other costs depend on production. Each step incurs $200,000 of material costs, so a batch costs $2 million.

Amgen Homework

Question 1a: How much will the plant be willing to spend on R&D to introduce a statistical process that improves the failure rate of each of the last five steps from 1% to 0.8%? Break it down to several distinct questions. How much does capacity utilization change? How much cost would be saved on wasted material? How much additional profit would the plant generate?

Question 1b: It takes one person-year, or 2,000 hours, to develop and perfect the model that reduces the failure rate. If PhD statisticians require a minimum of low-to-middle six-figure salaries, is it ever worthwhile to hire a full-time employee? When is it not? What if it took 4,000 hours instead, or two full-time employees? Is it still worth hiring both employees?

Amgen Homework

Option 2: Automating One Step in the Manufacturing Process for a Top-Selling Product

Goal: Automate a task.

Setting: Assume that the manufacturing process for the top-selling drug in question is similar to the process outlined in **Option 1** above. Once again, start from a setting with 10 processes, where the last five each display a failure rate of 1%. Once again, the entire success cycle takes 70 hours, plus one hour of testing for quality and one hour of cleaning, and the unsuccessful cycle takes 70 hours, plus one hour of testing and five hours of cleaning.

At present, it is somebody's job to do these tests for quality. The statisticians propose installing an additional sensor on its already extensive array of sensors. The sensor automates the decision around quality. It costs $500,000 to purchase and install.

Amgen Homework

Adding this new sensor would eliminate the one-hour wait for the lab results. It would reduce the time for each successful batch from 72 hours to 71 hours. The time for an unsuccessful batch also declines from 76 to 75 hours.

The sensor needs a model to perform the equivalent of a test for quality. That model uses machine learning. It provides a decision—green/red for success/fail. The model attempts to replicate what the person would have decided after viewing the lab results. Like all machine learning, however, these automated tests face challenges precisely replicating the judgment of humans. They produce a tiny but non-trivial rate of false positives and a tiny but non-trivial rate of false negatives. False positives give a green light even though the batch spoiled, and false negatives give a red light even though it truthfully would have been okay.

Amgen Homework

Statisticians can set parameters in the decision rules that determine the green/red to increase or decrease the probability of false-negatives and false-positives. Because the underlying biochemistry is so complex, there is a trade-off between false-negatives and false-positives, whereby the parameters that lower the probability of one increase the probability of the other, and vice versa.

The firm's senior management believe their reputation with their customers crucially depends on never delivering ineffective output. That belief aligns with government regulations in many countries, which heavily penalize the firm for shipping sub-par products. So the plant's managers instruct the statisticians to reduce the probability of false positives to its lowest possible point. Due to this change, the statisticians estimate that installing the automated sensor will increase the failure rate to "more or less 5.5%." They cannot be sure until they operate for a while.

Amgen Homework

Question 2a: What are the quantifiable gains from installing this sensor? Based on the numbers, is it worth the expense?

Question 2b: Improvements in cycle time generate a range of gains to the firm that cannot be quantified. List gains from automation for the operating strategy of the firm, which you cannot quantify, and list as many as possible.

Video

https://www.youtube.com/watch?v=LiF5RLql6PY

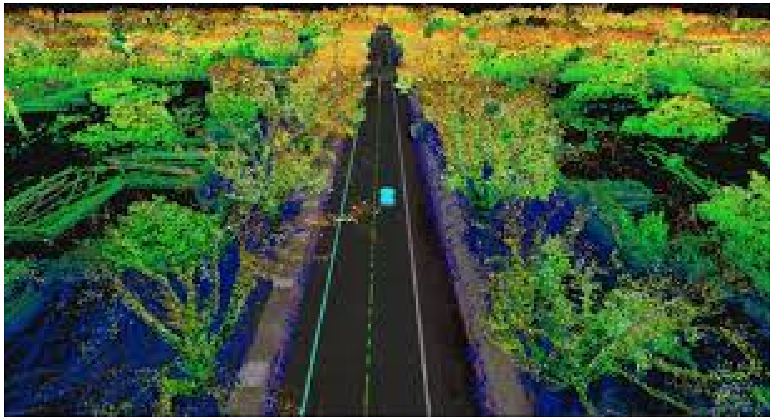

That is a simulation of "what the car sees." This helps humans diagnose whether the system correctly perceives the environment.

Setting: AVs in prototype

Set in spring of 2019. Easy to bring to present.

Two different arrangements of supply chains

- "Full stack DIY." One firm makes software, makes vehicle & operates for own use – e.g., Waymo.
- "OEM-oriented." One firm designs the final product, and others operate it, and others supply components.

Deep disagreements over technical progress & scaling.

- Some anticipate learning from experience. Progress gradually…
- Some anticipate technical discontinuities. Leap frog to frontier.

DeepMap supplied one component, and initially bet on leapfrogging (and later hedged that bet).

Setting: Levels of autonomous vehicles

Level 0. No automation.

Level 1. Drive assistance: At least one system, such as cruise control.

Level 2. Partial automation: two or more systems for acceleration, steering, and braking. Driver must monitor at all times.

Level 3. Conditional automation: Self driving under good weather and/or geo-fencing. Driver prepared to take over on short notice.

Level 4. High automation: Vehicle self-drives in majority of circumstances, and driver is passenger. Driver takes over in exceptional circumstances, such as poor weather or new environment.

Level 5. Full automation: Self-driving in all conditions and environments. Design presumes no human – e.g., no steering wheel.

List of course cases

Digital Building Blocks	
Tech with a Side of Pizza: How Dominos Rose to the Top	421035-PDF-ENG
Korea Telecom: Building a GiGaTopia	617014-PDF-ENG
Streaming over Broadband: Why Doesn't my Netflix Work?	616007-PDF-ENG
Lakes Banking Group: Data Management	618021-PDF-ENG
Houston, We Have a Problem: NASA and the Open Innovation (A)	414044-PDF-ENG

Platforms	
Pinduoduo	620040-PDF-ENG
Lexoo: Building a Long-Lasting Platform	619019-PDF-ENG (99A)
ZBJ: Building a Global Outsourcing Platform for Knowledge Workers (A)	618044-PDF-ENG
Apple Pay and Mobile Payments in Australia (A)	619010-PDF-ENG
Threadless: the Renewal of an Online Community	621056-PDF-ENG
Platform Strategy Simulation Exercise Instructions	620077-PDF-ENG

List of Course Cases

Organizing Digital	
Viacom: Democratization of Data Science	618016-PDF-ENG
TSG Hoffenheim: Football in the Age of Analytics	616010-PDF-ENG
Booking.com	619015-PDF-ENG
Moderna (A)	621032-PDF-ENG
Commcercializing AI	
Digital Manufacturing at Amgen	621008-PDF-ENG
Lemonade: Disrupting Insurance with Instant Everything, Killer Prices, and a Big Heart	519078-PDF-ENG
IBM Watson MD Anderson Cancer Center	621022-PDF-ENG
Twiggle: E-commerce with a semantic search	620025-PDF-ENG
Feeling Machines: Emotion AI at Affectiva	620058-PDF-ENG
DeepMap: Charting the Road Ahead for Autonomous Vehicles	620047-PDF-ENG