

AI for Industrials: Making AI Make Money

A practical guide for industrial companies to use AI to improve core operational processes and increase profits.

CONTENTS



We Are Nearing the End
of the Current AI Hype Cycle



What AI is Good For



Making AI Make Money



Conclusion

ABOUT THE AUTHORS

Jim Manzi

Founder,
Foundry.ai

Jim is a co-founder of Foundry.ai. He was founder, CEO and Chairman of Applied Predictive Technologies, which became the world's largest cloud-based AI software company. Previously, Jim developed pattern recognition software at AT&T Laboratories, and worked as a corporate strategy consultant. He is the author of several software patents, as well as the 2014 Harvard Business Review article "The Discipline of Business Experimentation."

Jim received an SB in mathematics from MIT, and was subsequently awarded a Dean's Fellowship in statistics to the doctoral program at the Wharton School of the University of Pennsylvania.

Scott Setrakian

Vice Chairman,
Foundry.ai

Scott leads Foundry.ai's San Francisco office. Prior to joining Foundry.ai, Scott was co-founder and Managing Director of Applied Predictive Technologies. Previously, Scott sat on the Board of Directors of Mercer Management Consulting and ran the firm's global Oil, Gas, Chemicals, Pharmaceuticals and Process Industries Consulting Group.

Scott received an MBA and an A.B. in Human Biology from Stanford University. He sits on the Board of Directors of the Buena Vista Funds, the William Saroyan Foundation and the San Francisco Zoo.

WE ARE NEARING THE END OF THE CURRENT AI HYPE CYCLE

The term 'AI' (Artificial Intelligence) has gathered momentum over the last few years and is arguably today's most over-hyped business buzz phrase. The vast majority of Global 2000 CEOs are being challenged by their Boards to demonstrate operational and financial benefits from applying AI to their businesses. Almost every large company has established some kind of working team to brainstorm and prioritize AI applications, and many companies have funded some specific pilot projects that have emerged from this process.

In our experience, most of these initiatives will end up disappointing their sponsors. This perspective is confirmed by the results of an extensive 2019 MIT / BCG report¹, in which more than two-thirds of corporate executives at companies with significant AI efforts reported that these efforts have not created material value. The problem is *not* that experienced senior executives are somehow misled by all the buzz surrounding AI. They know a hype cycle when they see one. They also understand the eternal verities of successfully introducing an important new digital technology:



Pilot quickly at low-cost



Demand some measurable short-term impact, and reinforce success



Maintain an unwavering focus on the bottom-line

This has been the playbook for the large companies that have most successfully introduced major new technologies from enterprise data warehouses to CRM systems to the Web, and it is a requirement for successfully introducing AI to create shareholder value.

The problem is that most executives do not have a sufficiently granular understanding of AI to allocate pilot-stage resources well. As with all technological advances, it is unnecessary for CXOs to understand the detailed inner workings of a piece of AI software, just as it is unnecessary for them to understand the detailed engineering of their mobile phones' operating systems. But what they *do* need to know about AI is the answer to one question: *What is it good for?*

Right now, most senior executives don't have a clear answer to this question, and therefore many companies are undertaking pilots addressing the wrong business problems, measuring the wrong short-term metrics, and trying to build platforms and roadmaps which – at this stage of AI's maturity – will often do more harm than good.

WHAT AI IS GOOD FOR

AI on television is robots playing Jeopardy. But in our experience, AI that makes money for a large industrial business almost always follows a specific pattern: it is software that uses data + math to create statistical improvement in a repeating business decision process. Some operational processes, such as predictive maintenance, are already widely-seen as AI targets in industrial companies. But there are numerous others, including:

- Back-office processes such as indirect procurement or financial reporting
- Production processes such as yield improvement or automated quality testing
- Go-to-market processes such as dynamic pricing or deploying industrial sales teams

In fact, once you know what to look for, profitable AI opportunities are hiding in plain sight everywhere inside of a large industrial company. AI can drive outsized profit improvements, but the benefits aren't achieved in a single huge 'transformational' investment. They are realized incrementally, by building a mountain of pebbles. And for large industrial companies, each 'pebble' can be worth millions of dollars of profits.

Isn't this just process automation? Yes, in a sense. But it is automation (and to be more accurate, often semi-automation) of a specific kind of process: cognitive business processes that require decision-making under conditions of uncertainty.

MAKING AI MAKE MONEY

Choosing the right business application areas to address is the starting point for effective executive management of AI. In theory, starting with a list of new technologies and then determining how much money is at stake in applying each of them to various business challenges should get to roughly the same place as starting from business problems and testing the case for applying new technologies against them. In practice, however, we have found companies are much better off starting with the business problems, primarily because evaluating technical feasibility is a far more delegable task than judging where the profit opportunities sit in a business. The MIT / BCG report referenced above also indicates that AI initiatives under the direct purview of the CEO or other line executives are **more than twice as likely** to create value as those treated as a technology project reporting to the CIO.

Specifically, senior executives should identify a short list of core repeated decision processes with high-profit leverage that would be improved with better data utilization. We have rarely found them to be wrong about this. The work of the staff is then to estimate the value-at-stake for each process that is addressable with

the AI technology of today (not the potential technology of five years from now). An AI project should not commence without a clear plan to generate measurable incremental cash flow within 12 months.

Every organization is unique, but there are two core business processes that we believe hold very large AI profit opportunities for most industrial companies: (1) Predictive asset management, and (2) Dynamic price optimization

1 Predictive asset management



Objectives

Predictive maintenance is often cited by industrial companies as the most obvious AI low-hanging fruit, but there is still a very high ratio of talk to success. The objective is to predict mechanical failures before they happen, leading to efficient preemptive maintenance that maximizes output and minimizes maintenance costs and downtime.

We have found it is better to approach this challenge as an asset management one, rather than one focused purely on maintenance. Predicting a future asset's operational performance only matters when there is a recommended intervention with net positive odds-adjusted economics.

Sometimes this will be a preventive maintenance procedure or a component replacement to forestall a much more expensive problem later. Often, however, it will involve changing the machine's operational parameters to sacrifice current production, sometimes in relatively small increments, for future benefits, such as extended unit performance. This means that an economic trade-off engine that calculates the value of each alternative (e.g.,

major maintenance intervention, minor maintenance intervention, altering operating throughputs, doing nothing, etc.) is central to the success of an AI system supporting asset management.



Data and Modeling

In our experience, the significant technical challenge of building the machine learning model that accurately predicts failure modes and probabilities is actually not the hardest part. Typically, the most difficult obstacle is data management.

A core input for any predictive maintenance model is IoT (Internet-of-Things) data from the assets, usually in the form of sensor readings and fault codes. In the real world, this data tends to be massive and messy: petabytes in volume containing tens of thousands of unique fault codes. Very often the relevant data sets are highly error prone, as sensors fail or become mis-calibrated in ways that are hard to measure without extremely expensive sensor duplication. Additionally, the data stream is almost always incomplete, because there are many opportunities to drop data from sensor to SCADA system to data warehouse.

Converting all of this to usable inputs for a prediction model demands the integration of expertise across multiple areas including:

- A modern, usually cloud-based, data pipeline incorporating non-SQL data structures and federated query capabilities;
- Sophisticated AI data engineering for data vectorization, compression and feature extraction; and
- Engineering expertise in the physical asset itself

Integrating data from additional sources can also drive significant improvements in modeling accuracy. Maintenance logs, video (e.g., external inspection video, internal borescope imagery, or even ‘unintentional’ data like plant security video) and audio recordings of the assets, as well as truly external data such as hourly weather, can greatly enhance the usability and fidelity of the IoT data. In each case, AI can be deployed to convert text log notes, imagery or sound files into vectors of numbers that can be used as additional raw materials for machine learning prediction models.

Simply put, chasing the dream of simply pointing a generic AI method at a pile of sensor data to “find patterns without bias” will generally end in disappointment. The substantial benefits available from AI-powered asset management software require careful and expert programming, but are fully achievable.



Implementation

Any new AI capability will inevitably require some change in existing business processes in order to generate added profits, and these changes are usually the greatest barrier to realizing value quickly. We have frequently seen predictive maintenance exercises run aground on the “So what?” problem – essentially, the organization doesn’t respond to model-driven recommendations.

To increase the likelihood that that predictions lead to profitable timely responses, the predictive models should be integrated with sensible response protocols, within the existing business structure, and with the minimum feasible degree of change to existing business processes.

2 Dynamic price optimization

Objectives

Dynamic Pricing is a new phrase describing a long-standing business goal: pricing for profit rather than pricing for volume. It requires real-time awareness of customer demand, competitive pricing, and your own margin economics. The payoff can be significant.

Pricing for profit is hard to argue against in theory. However, it is challenging for the field salesforce to set prices that maximize margin while navigating competitive dynamics, opaque cost economics, negotiating tensions, and revenue incentives.

“Price optimization” software has been on the market for decades to help companies cope with these issues, but despite the economic benefits that these tools promise, an extensive Bain & Co. survey² of more than 1,700 leading B2B companies in 2018 indicated that only about 25% of them are utilizing price optimization software, and that this proportion is about half of what it was 15 years ago³.

In our experience, sales teams often reject the recommendations from these legacy pricing systems as ‘theoretical,’ observing that the models:

- Use limited data, especially external data
- Excessively simplify the estimation of price elasticities
- Present recommendations with limited transparency into their calculations

Rejection of the recommendations nearly always leads to quoting prices that are lower than suggested, to minimize the risk of a lost sale. The cost of lost margin potential is invisible.

Despite these historical challenges, recent developments in AI can address the shortfalls of these legacy systems, and are creating high-profit opportunities for industrial sales forces to realize the benefits of dynamic pricing.



Data and Modeling

Perhaps the most important data required to achieve true dynamic pricing is intelligence into what competitors are offering and at what price, and how this compares to one's own similar-but-not-identical offers. Competitive behavior can provide insight into how much the market will bear, and illuminate opportunities to nudge up price without sacrificing competitiveness.

AI techniques provide the capability to automatically ingest competitive pricing data from a wide variety of sources. Depending on the product market, the sources of this data will often include: (i) competitive product catalogs; (ii) federal, state, local and education department filings, public meeting notes and contract awards notices; (iii) specialized trade journals; and (iv) customer company profile databases.

The basic technique to capture and exploit this data has two steps. The first is to use a combination of scraping, APIs and other search methods to cast a wide net. This is conceptually simple, but it often requires more infrastructure and expertise than companies expect. The second is to find the signal (pricing by product by competitor) within all of the noise present in the data, quickly and reliably. With sufficient investment, this can now be done much more effectively than even a few years ago, using a combination of natural language processing, machine vision and machine learning. (See Insert)

This external data is integrated with internal data — including prior price offers, contract terms, won/lost data, known competitive bids, and cost and margin data — to form the basis for price modeling.

With the integrated internal and external data sets in place, new machine learning models now allow companies to model data in previously impossible ways, using optimized ensembles of deep learning, gradient boosting machines and other algorithms.

These approaches enable modern demand prediction systems to consider the subtle and quickly changing demand patterns that exist in most businesses, and respond in real time. In contrast, legacy models typically use methodologies that require training on multiple years of historical data, and are unable to significantly modify predictions in the face of short-term changes in the market.



Implementation

Implementation of these AI tools should seek to minimize change to existing processes. Essentially, the system needs to provide pricing recommendations that are simultaneously granular, comprehensible, and timely. And most importantly, if accepted they need to increase profit.

The acid test is whether the sales force follows the recommendations, which requires confidence that they are in fact correct. To that end, these AI dynamic pricing tools allow reps to ‘look inside the black box’ to understand what factors are driving price recommendations. Additionally, these systems can incorporate both won/lost results and direct user feedback to self-learn and improve over time.

INSERT

Four Practical AI Building Blocks

There are four key AI capabilities that consistently drive business value. These capabilities are at their most powerful when they are all present in AI software, operating together. Each has a name that is a somewhat intimidating buzz-phrase, but all have practical definitions that executives should understand:

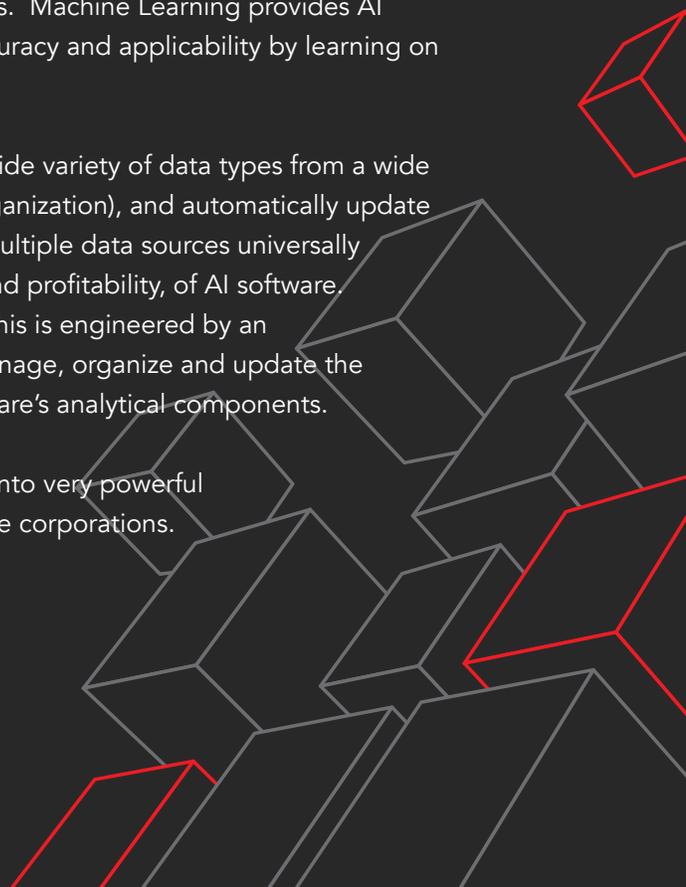
Natural Language Processing is a technology that allows computer software to understand and process the meaning of words and sentences. This supports a wide variety of AI capabilities, from scanning the entire Internet for information on a given topic, to monitoring customer sentiment, to ingesting and summarizing the minutes of a public hearing.

Machine Vision is the technology of automatically converting visual data (such as a picture or a video feed) into numerical inputs that can be incorporated into software models. In essence, Machine Vision enables AI software to integrate and process data from the visual world into decision models, from in-store cameras to equipment monitors to satellite photos to medical x-rays.

Machine Learning is a critical building block because AI modeling often depends on the discovery of patterns between historical data and outcomes. Machine Learning provides AI software with the capability to continuously improve its accuracy and applicability by learning on its own as it receives new data.

Data API Integration describes the ability to integrate a wide variety of data types from a wide variety of data sources (both internal and external to an organization), and automatically update the combined data set on an ongoing basis in real time. Multiple data sources universally improve the predictive accuracy, and therefore the utility and profitability, of AI software. Integrating the data and keeping it updated is essential. This is engineered by an Application Programming Interface ("API") designed to manage, organize and update the multiple data sources, and integrate them into the AI software's analytical components.

Taken together, these four components can be integrated into very powerful software tools to support complex decision making for large corporations.

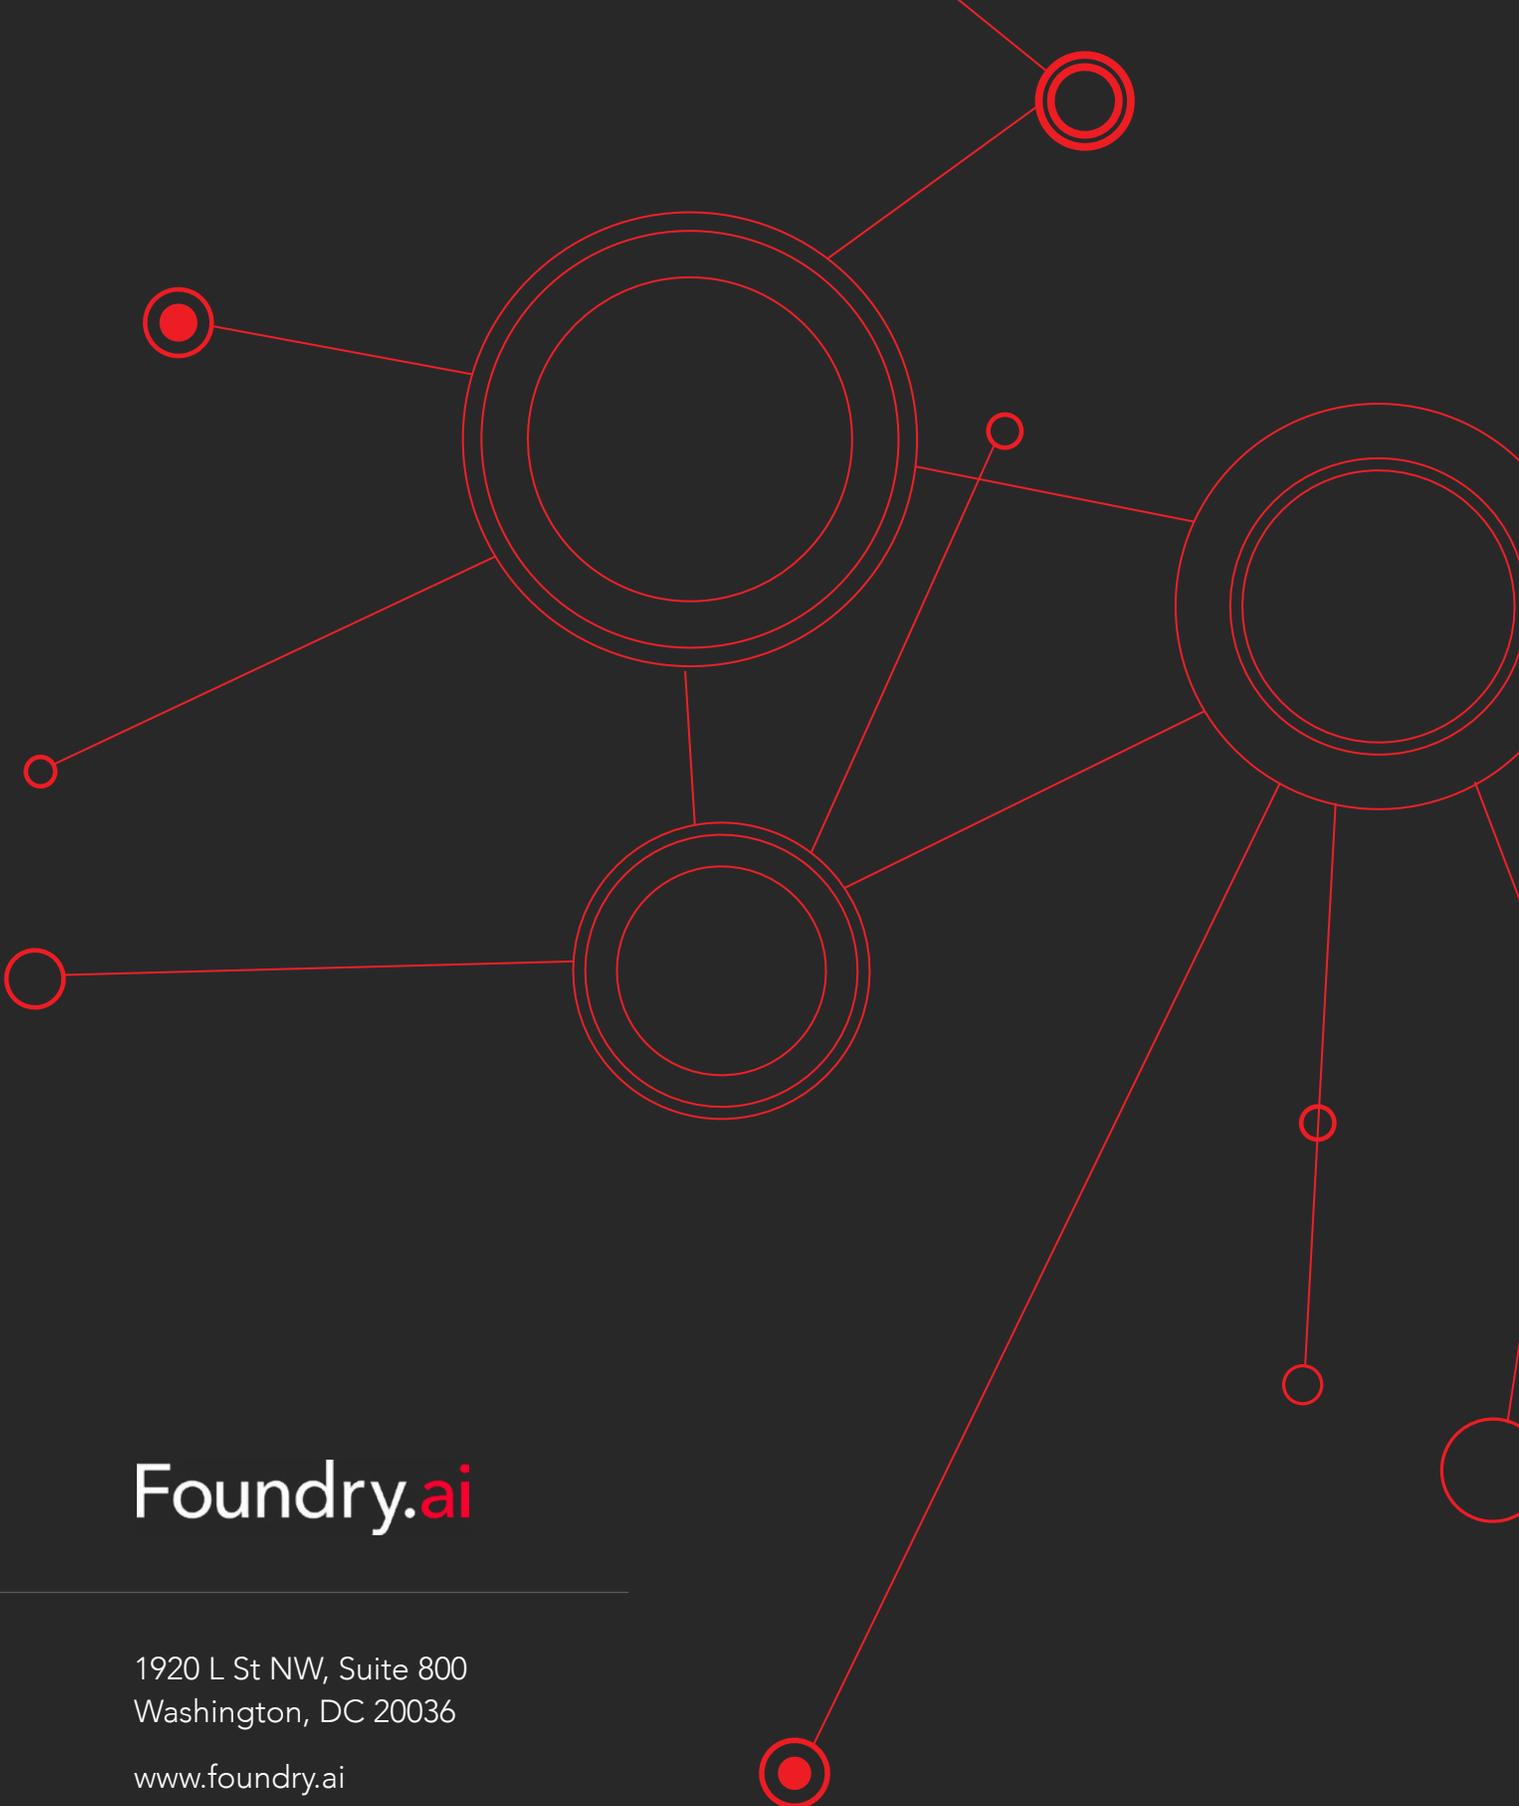


CONCLUSION

While much of today's AI focus centers on flashy use cases and large transformational investments, most of the real successes for industrial companies have come from using AI in highly practical ways to improve core business operational processes. We have seen executives who follow the guidelines in this paper successfully capture this opportunity, creating significant value within surprisingly short timelines. We hope these concepts help you in this important domain and more broadly as you consider how to drive focused, pragmatic, and profitable AI initiatives within your organization.

SOURCES

- 1 S. Ransbotham, S. Khodabandeh, R. Fehling, B. LaFountain, D. Kiron, "Winning With AI," MIT Sloan Management Review and Boston Consulting Group, <https://sloanreview.mit.edu/projects/winning-with-ai>, (October 2019)
- 2 Ron Kermisch and David Burns, "Is Pricing Killing Your Profits? How the best B2B companies set and get the right price." Bain & Company, <https://www.bain.com/insights/is-pricing-killing-your-profits>, (June 13, 2018)
- 3 Darrell K. Rigby, "Price Optimization Models," Bain & Company, <https://www.bain.com/insights/management-tools-price-optimization-models>, (April 2, 2018)



Foundry.ai

1920 L St NW, Suite 800
Washington, DC 20036

www.foundry.ai