



Making AI Make Money

A Practical Guide for CEOs to Cut Through the
AI Hype Cycle and Drive Profitable Initiatives

By Jim Manzi and Ned Brody



CONTENTS

- Artificial Intelligence:
Background for Senior Executives
- What Is AI and Why Does It Matter Now?
- Artificial Intelligence:
An Action Plan for Senior Executives
- Conclusion
- Four Practical AI Building Blocks

ABOUT THE AUTHORS

Jim Manzi


Partner, Foundry.ai

Jim was founder, CEO, and Chairman of Applied Predictive Technologies (APT), which became the world's largest purely cloud-based predictive analytics software company. Prior to founding APT, Jim worked at AT&T Laboratories and was a partner at Mercer Management Consulting. Jim received a B.S. in mathematics from the Massachusetts Institute of Technology. He was subsequently awarded a Dean's Fellowship in statistics to the doctoral program at the Wharton School of the University of Pennsylvania.

Ned Brody

Partner, Foundry.ai

Ned has led a variety of technology businesses over the last 20 years, having been Yahoo's Head of the Americas, CEO of AOL Networks, Chief Revenue Officer of AOL, and CFO of LookSmart. Previously, he started and led Mercer Management Consulting's Internet Practice and San Francisco Office. Ned holds a MBA and a B.S. from the Wharton School of the University of Pennsylvania with concentrations in decision theory and finance.



ARTIFICIAL INTELLIGENCE: BACKGROUND FOR SENIOR EXECUTIVES

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 is because they tend to veer in one of two directions: either they are so grandiose that they fail under their own weight, sometimes at great expense (example: create an integrated data architecture that breaks down silos and build a suite of AI analytical services for the collective enterprise), or they pursue interesting and useful technologies that in the end produce too little incremental profit to justify all the cost and management attention they require.

This should sound familiar to any experienced executive. While some companies succeeded, many others wasted enormous amounts of time and money trying to exploit over-hyped emerging technologies over the last few decades, ranging from the World Wide Web to CRM to the Enterprise Data Warehouse to Big Data.

But that doesn't mean these innovations were just meaningless fads. With each, as with AI, there was real substance at the root of the hype. Senior executives who followed a consistent set of practices outlined below were able to successfully leverage these digital technologies when they were at the same early stage that AI occupies now:



Insist on understanding the technology at a practical level, and refuse to be snowed by technical jargon or conventional wisdom.



Act quickly and at low cost. Learn what works from trial-and-error experience and then reinforce successes.



Maintain an unwavering focus on profits as the success criterion for any proposed investment of time and money.

The executives who adopted these practices were able to realize value from these innovations early and – fast forwarding to today – each one of these technologies plays a critical role in the operations of almost every Global 2000 enterprise. Similarly, AI will be a big part of running almost any large corporation in the future. The question is: how can senior executives apply these same lessons to realize the benefits of AI today?

WHAT IS AI AND WHY DOES IT MATTER NOW?

Applying these lessons to AI should begin with defining what we mean by the term.

At its most expansive, AI is the capability of a machine to imitate intelligent human behavior. But what we think of as ‘artificial intelligence’ tends to change over time – once the capability becomes familiar and reliable, we consider it just plain old software. In 1958, the idea of computers playing chess was AI; in 1998, recognizing a face in a photo was AI; today, they are standard software elements.

From a practical business perspective, the most valuable ‘intelligent human behavior’ that AI can replicate is the capacity to recognize patterns and create knowledge from data and then use this understanding to make profitable decisions.

Importantly, the digital capabilities required for machines to make such decisions are growing at breathtaking speed. Everybody knows Moore’s Law – crudely stated, that information processing productivity doubles every 18 months. The less well-known Kryder’s Law and Nielsen’s Law observe that digital data storage and transmission productivity also double every 12 – 24 months. These extraordinary growth rates have been compounding for decades, driving innovations from the origins of the Internet to the explosion in Big Data processing. As a result, each year the feasibility frontier for digital applications advances and software engineers race to build tools that can exploit this burgeoning capacity.

Beyond just these technological advances, the nature and potential sources of data are also changing rapidly. Companies have always relied on data to make decisions, but the advent of systemic data availability, especially external sources of data not typically stored ‘in-house’ at a company, are starting to significantly enhance decision making potential. Weather forecasts, competitive marketing calendars,


consumer preferences and local events are all examples of available external data sources that can be brought to bear on decisions by using AI tools to process them within the required time frames.

These collective advances have now reached a point where technology development is crossing the threshold of human cognition in specific areas. In plain English, many useful business decisions can now be made by software – better, faster and more efficiently than by us humans. At the edge of the frontier, AI is creating immense business value by automating and optimizing business decisions that cannot be made reliably and cost-efficiently by people.

Typically, these decisions have 3 characteristics:



They require complex analysis of significant amounts of data, often drawn from multiple sources.



The decision type occurs in high volume – i.e., there are many individual decisions of the type that occur over time.



There is a tangible financial impact from making each one of these individual decisions better, but adding enough human-based analysis to do so consistently would cost more than the resulting value increase.

With the advancement of the feasibility frontier, examples of AI that can support these types of decisions are suddenly everywhere:

TRADITIONALLY

NOW

Corporate procurement departments manage multi-million dollar purchases through complex sourcing events. These processes are expensive, but economically justifiable due to the large scale of the planned spend.



AI software can be used to identify additional relevant bidders for \$10K purchases in the same way that humans do for \$10MM sourcing events, unlocking huge value when deployed across thousands of such purchases.

Online retailers employ human photo editors to carefully select images for their most important products in an effort to drive up sales volumes. However, this same image selection process is not economically justifiable for lower-sales products.



AI software can efficiently select the highest impact images for e-commerce product displays for the thousands of products that comprise the huge majority of a company's catalog, driving up sales across the entire line.

Health insurance companies employ expert personnel who review large claims to ensure they are properly coded so that significant dollar amounts are paid out correctly. However, the vast majority of claims are low cost, and painstaking human review is not economically viable for these smaller submissions.



AI software can automate the process of detecting improperly coded medical claims for lower-dollar claims in the same way reviewers do today for large claims.

In each case, the pattern is the same. A company uses costly personnel for high-dollar decisions in a business process (e.g., major procurement events, ‘hero’ product images, major medical claims) but cannot cost-effectively deploy people of this quality for the much longer list of smaller decisions of the same type. Often the collective profit impact of these smaller decisions can be enormous – in some cases bigger in aggregate than the total of the large decisions.

By deploying AI systems, the corporation gets several benefits:

- Adding the incremental profit stream of human-like decision making for this larger set of decisions.
- Freeing up the expensive people working on the ‘edge cases’ to focus on the smaller number of highest-impact and most complex decisions.
- Providing additional low-cost support to the people making the highest impact decisions – i.e., ‘helping the B players perform more like A players’.

We have consistently observed that AI initiatives that make money for large corporations are much more prosaic than robots playing Jeopardy. They represent the application of data plus math to create tangible and measurable improvement in a repeated business decision process. And importantly, successful AI programs tend not to create these benefits through a single integrated transformation project, but rather through discrete initiatives focused on a series of individual processes. In essence, you build a ‘mountain of pebbles,’ and in a Global 2000 organization, each ‘pebble’ can be worth millions of dollars per year.

ARTIFICIAL INTELLIGENCE: AN ACTION PLAN FOR SENIOR EXECUTIVES

Above we observed the consistent set of practices followed by the senior executives that most successfully exploited digital technologies such as the World Wide Web and CRM when they were in the early 'hype' stage. To recap: obtain a practical understanding of the technology; act quickly at low cost and learn from experience; and maintain a relentless focus on profits.

Our experience is that, within those guidelines, the following more specific practices for executive management of AI initiatives will most reliably lead to success:

1 Work Backward From Business Problems, Not Forward From New Technologies

In theory, starting with a list of new technologies and then rigorously 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 rigorously testing the business case for applying new technologies against them. But in practice 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.

Specifically, senior executives should identify a short list of core 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).

The senior executives should then select 1 - 3 of these as pilot initiatives and directly sponsor the pilot projects.

2 Focus Relentlessly on Financial Value

Every one of these AI pilot projects should create incremental free cash flow within 12 months, or else be terminated. Corporations pursue strategies that build value over time, and an overall AI strategy should do the same. But in this case, the “Trust me boss, we need to spend money for the next 4 years, but we will have a huge payout starting in year 5” approach almost always ends in disappointment.

Each pilot should have a practical theory of the case, including:

- Clarity on the stream of decisions to be changed.
- Identification of the analytical methods and at least some of the datasets that will be used to improve these decisions.
- A simple analysis demonstrating that a sensible degree of decision improvement will create at least several million dollars per year of pre-tax operating profit gain.
- A reliable method to measure the actual dollar value of business improvement created at the conclusion of the pilot. The measurement of value creation should be as rigorous as is consistent with the business process, and should ideally be an A/B test or other controlled experiment.

3 Minimize Change to Existing Business Processes

Any AI pilot will inevitably require some change in business processes in order to generate more profits. These changes are usually the greatest actual barrier to realizing value quickly. Process changes should be consciously kept to the minimum degree consistent with achieving a significant fraction of the available business benefit. This applies both to the operational units that are executing the target decision process and to the IT teams that support it.

The IT to support these AI applications should be lightweight, cloud-based and interact with existing corporate data stores via simple API. The AI system should also interact with preexisting operational software via API using established interfaces. These interactions should focus on directly driving decisions, rather than just providing yet more information that a human needs to analyze and process – nobody needs another dashboard, however well-intended.

With respect to data, the AI system should impose zero formatting demands on internal data systems and take data exactly as-is. It should also automatically extract and integrate external data sources, as these other data classes tend to be major sources of prediction advantage. No pilot should attempt to build or be part of a detailed IT roadmap, but if successful, it can help to influence the future roadmap.

CONCLUSION

While the AI opportunity is enormous for almost any large business, addressing it as a senior executive can seem both daunting and exasperating. You listen to a lot of big talk, but getting answers to straightforward questions such as, “So, exactly what is this going to change, and how does that make me more money?” is surprisingly difficult. You don’t want to over-constrain a process that might be extremely valuable, but you also don’t want to give AI a pass from the need to generate clear benefits in excess of costs.


And the race to capture real value from AI is officially on. A recent Deloitte study revealed that while 25% of companies have implemented AI and machine learning technologies over the past 24 months, 63% expect to do so within the next 24 months. It is likely that many of these organizations will be disappointed by their initial efforts, but thinking about the groundbreaking technologies that came before, there are countless examples of companies who successfully adopted these innovations early and created a competitive advantage that drove material shareholder value for years. If past is prologue, the difference between the winners and the losers will be how they approach the AI opportunity.

PAST 24 MONTHS



Implemented AI
& Machine Learning

NEXT 24 MONTHS




Planning to Implement AI
& Machine Learning

Source: Deloitte Consulting LLP, *Save-to-Transform as a Catalyst for Embracing Digital Disruption*, Deloitte's Second Biennial Global Cost Survey, April 2019.

In our experience, real-world AI applications tend to consistently reuse software components in four key areas – machine learning, machine vision, natural language processing, and data integration (see insert). By working with experienced AI engineering and data science teams that have access to these core software components and an orientation to creating profits, a pilot AI software tool designed to address a specific decision challenge can be created, implemented, and have its value case proven through controlled experiments within 6 - 9 months.

Senior executives who have followed the guidelines in this paper have consistently succeeded in navigating the challenges of introducing AI into their large corporations. We hope they assist you in driving rapid and significant value for your organization as you consider the next step in your AI journey.



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.



Foundry.ai

1920 L St NW, Suite 800
Washington, DC 20036

<https://grok.ai>

