Floridi Curves

A methodology for de-risking AI implementation across industries

by **Bruce Benson**, SMD

Foreword **Dr. Luciano Floridi**

John K. Castle Professor in the Practice of Cognitive Science and
the Founding Director of the Digital Ethics Center at Yale University

Artificial intelligence is transforming industries worldwide — from healthcare and finance to automotive and manufacturing. Yet consistent implementation success remains remarkably elusive. Recent comprehensive studies have documented alarming failure rates for AI initiatives, frequently exceeding 80% across various sectors.[1] It appears that these failures rarely stem from algorithmic inadequacies, but rather from fundamental misalignments between AI capabilities and improperly structured supporting environments.

In 2019, I introduced a theoretical framework specifically designed to overcome such implementation challenges (Floridi, 2019).[2] The main suggestion was that successful AI deployment requires methodically converting difficult human-centric tasks into “merely” complex computational problems through strategic environmental restructuring. This transformative process, which I term “enveloping,” involves creating precisely tailored conditions that align with AI’s inherent capabilities, effectively shifting challenges away from human difficulties toward machine-friendly complexity domains.

Bruce Benson’s groundbreaking white paper, which relies on these concepts of complexity and envelopes, represents a critically important contribution to the field by systematically operationalizing these theoretical insights into a comprehensive, pragmatic methodology for practitioners. His work illustrates how organizations across sectors can implement envelope-based solutions through structured processes and leverage what he kindly termed “Floridi Curves” to substantially enhance AI implementation success rates.

This seminal work arrives at a pivotal juncture in AI development, providing organizations worldwide with unprecedented clarity on how rigorous philosophical theory can directly inform and guide practical AI deployment strategies in complex operational environments. It is my distinct privilege to introduce this exceptionally insightful and immediately actionable resource to the professional community.

Introduction

In 2019 Dr. Luciano Floridi published his article entitled "What the Near Future of AI Could Be." He presented a method for turning challenging AI business problems into “merely” complex ones by creating "envelopes" around them to make them feasible for business use. These groundbreaking concepts developed by Luciano are so helpful that FTI Consulting has integrated them into our methodology for designing and implementing AI-powered solutions for our clients.

I was lucky enough to be introduced to Luciano when he and his wife moved from Oxford to New Haven, Connecticut, where I also live, to head Yale’s new Digital Ethics Center. After a succession of coffees and lunches, he agreed to support my work on this paper, which marries his theory with a practical methodology to help businesses solve real-world AI Problems.

The first section of this paper addresses the core concepts of complexity versus difficulty. The second section describes the concept of envelopes and their utility. The final section demonstrates how FTI Consulting’s AI methodology employs these concepts to solve business problems using an example from textbook publishing.

Difficulty vs. Complexity

In day-to-day language, difficulty and complexity are often interchangeable. Luciano, however, uses them to mean distinctly different things. It’s best to begin by quoting from his paper:

> *“AI is best understood as a reservoir of agency that can be used to solve problems. AI achieves its problem-solving goals by detaching the ability to perform a task successfully from any need to be “intelligent” in doing so. The App on my mobile phone does not need to be intelligent to play chess better than I do. Whenever this detachment is feasible, some AI solution becomes possible in principle.*
>
> *…What I would like to suggest is that, for the purpose of understanding AI’s development… it is useful to map problems on the basis of what resources are needed to solve them, and hence how far AI can have such resources. I am referring to computational resources, and hence to degrees of complexity; and to [human] skill-related resources, and hence to degrees of difficulty.”*[3]

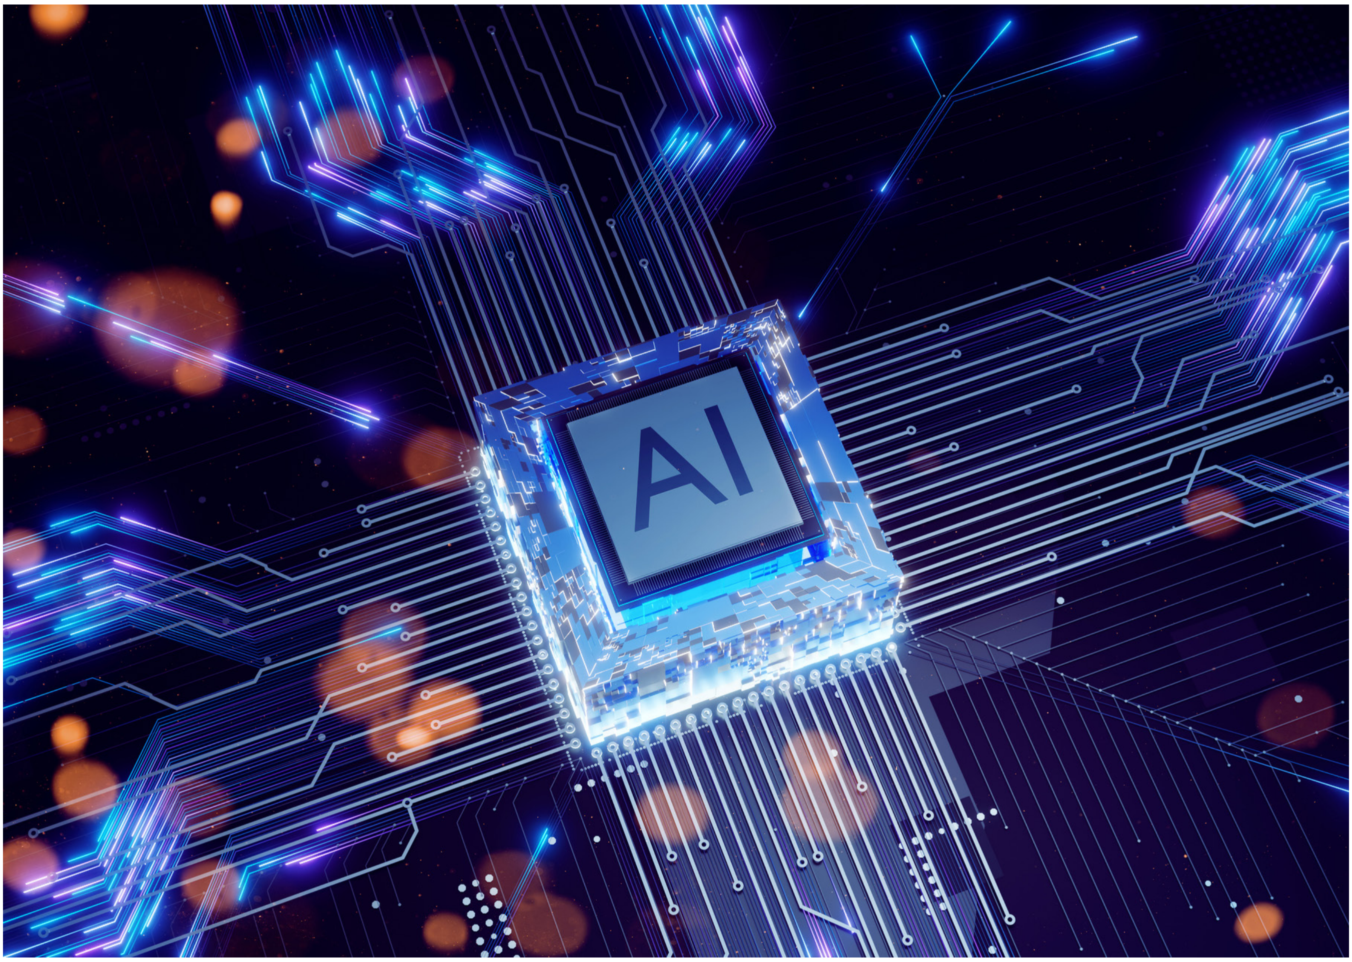

Let's examine Luciano's concepts individually since I have taken this out of context.

Agency and Detachment

Luciano says AI can be best understood as a "reservoir of agency." By this "reservoir," he means the set of available AI "agents" to solve problems. These agents range from simple multivariate regressions to decision trees, techniques like k-nearest neighbors and genetic algorithms, neural networks, generati ve adversarial networks (GANs), large language models and everything in between. At the moment, the development of these agents seems exponential. Also inherent in this reservoir is the use of these agents together. AI agents are inherently compatible since they can feast on related data. For example, one agent can perform natural language processing on a file full of emails and then pass these results into a large language model to expand its vocabulary.

His concept is that "AI achieves its problem-solving goals by detaching the ability to perform a task successfully from any need to be intelligent in doing so." By "intelligent," he means human intelligence or human choice-making. The chess app on his phone is not capable of psychological ploys against his opponent. It is not mulling over its next moves. It simply picks the next move with the highest probability of winning, based on its training.

AI scientists are fully aware of the distinction between intelligence and computational capability, but businesspeople often make the misstep of imbuing hoped-for AI solutions with human abilities.[4] This has important implications when designing AI solutions to business problems. Designers must carefully identify which parts of the overall system require human intelligence and which can be relegated to the reservoir of AI agents. Human intelligence could mean the driver of a car, a data-cleansing organization, an ethics approval committee, and so on.

Difficulty vs. "Mere" Complexity

Now, let's turn to the distinction between difficulty and complexity.

To simplify this, by "difficulty," Luciano means the parts of the problem that require human intelligence, i.e., "skills," and those that can be solved by a combination of (dumb) AI agents with vast computational abilities, which are limited only by available computational resources. Dishwashing requires human skill when we do it, but no skill when we use a dishwasher, which significantly reduces the human skill needed and requires very little AI. Self-driving cars take very little human skills — low difficulty — but enormous amounts of complex AI. Of course, many problems fall somewhere in between.

Figure 1 – Difficulty vs. Complexity

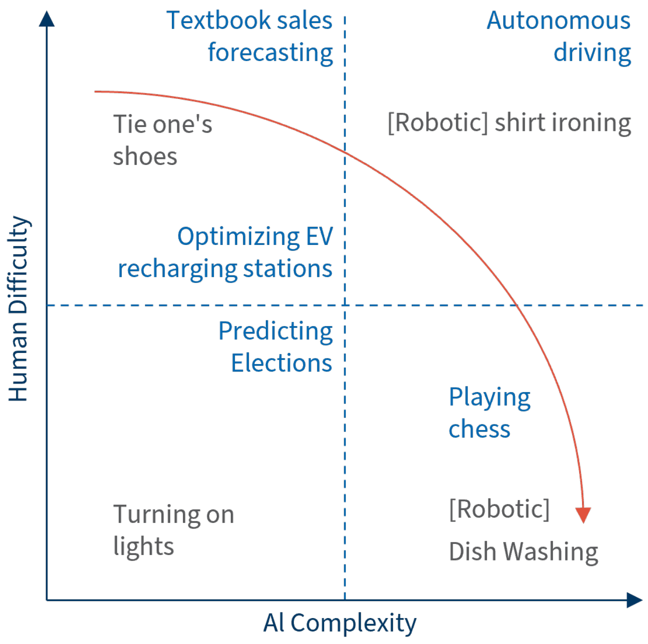

Luciano's paper plots the (human) difficulty of a problem vs. the complexity (amount of AI agency) needed to solve it, as shown in Figure 1.[5]

When analyzing business problems of various kinds, it is helpful to plot the problem (or linked problems) onto this type of matrix. It forces designers to break down a problem into these dimensions and carefully verify where the problem fits.

However, here is Luciano's key point about the graph above: The designer's goal is to shift as much of the AI system they are trying to create down this curve from 'difficult' toward the merely 'complex'. The red Curve indicates this. Why is this critical? The designer aims to understand how much of the problem can be reliably shifted toward AI-powered functionality and how much must remain in human hands. To give credit where credit is due, let's dub the curve above the "Floridi Curve."

An example is the continuous development of the truly self-driving car. Through each model year, these cars rely less and less on the driver's skill while the reservoir of AI embedded in the car's computers, radar, cameras and sensors continues to rise. As designers battle their way down the Floridi Curve with each new model, the car becomes smarter and smarter, requiring more computational power but less human skill.

Our view is that the journey from difficulty toward complexity when developing AI systems should not be haphazard but a rigorous process consciously pursued by the development team. How is the conversion from difficulty toward complexity achieved, and how is rigor introduced into the process? Let's turn to the concept of envelopes.

Envelopes As a Framework for Developing Production AI Systems

Let me quote Luciano again:

> *"How is this translation achieved? By transforming the environment within which AI operates into an AI-friendly environment. Such translation may increase the complexity of what the AI system needs to do enormously but, as long as it decreases the [human] difficulty, it is something that can be progressively achieved more and more successfully. Some examples should suffice to illustrate the point, but first, let me introduce the concept of enveloping."*

> *“In industrial robotics, the three-dimensional space that defines the boundaries within which a robot can work successfully is defined as the robot’s envelop. We do not build droids like Star Wars’ C3PO to wash dishes in the sink exactly in the same way as we would. We envelop environments around simple robots [dishwashers] to fit and exploit their limited capacities and still deliver the desired output…"*

> *“The same applies to Amazon’s robotic shelves, for example. It is the whole warehouse that is designed to be robot-friendly... This is why it is plausible that in an airport, which is a highly controlled and hence more easily ‘envelopable’ environment, a shuttle could be an autonomous vehicle, but not the school bus that serves my village, given that the bus driver needs to be able to operate in extreme and difficult circumstances (countryside, snow, no signals, no satellite coverage, etc.)...”*[6]

The Rigorous Use of Envelopes

We believe that when designing AI-embedded systems, the methodical use of envelopes and their Floridi Curves is an essential new tool to help guide system architecture, prototyping and systems development over time. Here we mean “systems” in all their forms: automated warehouses, cars or AI-enhanced business systems.

How are these new tools used “rigorously”? The key is a systematic progression as designers think through the architecture of the first minimally viable product and what the subsequent iterations might be as they shift a product down its Floridi Curve, requiring fewer human skills but more AI complexity.

We'll first envision how this works conceptually, then, in the last section of this paper, we show an example in practice.

Let’s take the example of a simplistic and fictional semi-autonomous car. We’ll ignore all the many physical problems involved and focus on the software and necessary supporting hardware. We'll use the following as the set of “abilities”[7] for our fictitious car: adaptive (rather than static) cruise control, collision avoidance, the ability to retrieve a car from a parking lot and the ability to parallel park the car without human assistance.

The first step is to develop the current state of our car’s capabilities relative to the desired capabilities above. Figure 2 shows an example. Let’s focus on the driver’s skill difficulties first.

Reading from high to low on the difficulty axis, we have ranked collision avoidance as the most difficult human skill. This is followed by parallel parking, adaptive cruise control and, finally, retrieval of a parked car. (Readers are free to choose their own ranking.)

Figure 2 – Current State (Model Year 0)

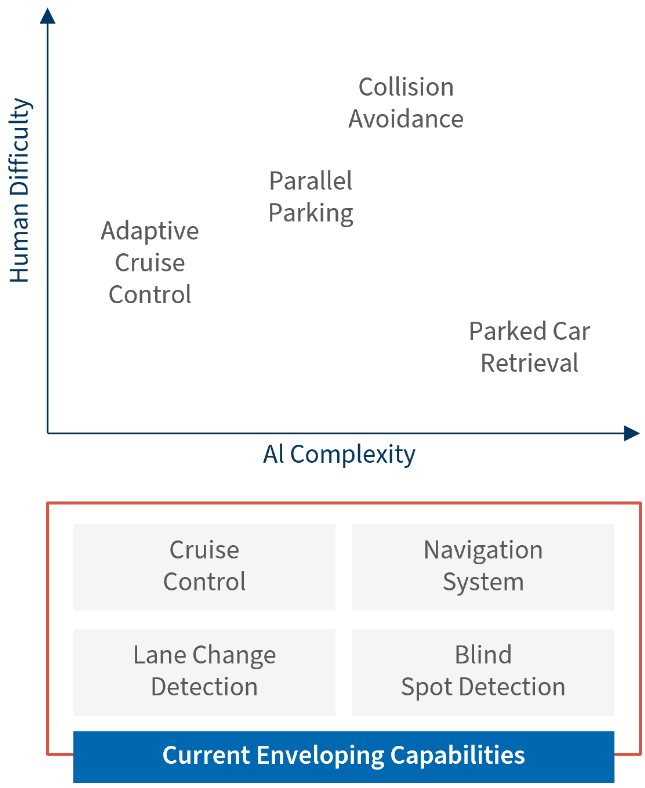

Most drivers find our proposed skills difficult. Take adaptive cruise control. Even with conventional cruise control, the driver is always manually adjusting. She must annoyingly break and reset cruise control whenever there is a situational change in fast-moving traffic. It is hardly worth the effort except on lonely highways. Collision avoidance is manifestly hard, or we wouldn't have collisions (there were over 6 million accidents in the US in 2023).[8] In America, parallel parking is the hardest part of our driving test and feared by all driving students. Parked car retrieval, on the other hand, is easy for us humans but is a notoriously difficult problem for AI.

Rating AI Complexity

Now, where do these abilities stand relative to their AI complexity? This question depends on the capabilities already available to AI developers (what Luciano would call the "reservoir of AI agency"). We've listed these for our fictitious car in the orange box in Figure 2. For example, creating the AI ability to enable adaptive cruise control will be easier to design and build if lane-change sensors and conventional cruise control feedback mechanisms like braking are already available. However, for collision avoidance, none of the existing enveloping capabilities do much to lessen the AI complexity.

In this light, the designer's job is not only to create the AI models — which can be complicated in their own right — but also to create the environmental improvements to the envelope that allow the new models to operate. Hence, the AI models themselves and the enveloping capabilities must be considered simultaneously. An example is the need to install 360-degree cameras in our cars to support AI-based collision avoidance. These cameras are not currently available in our car's capabilities.

In addition, as we saw with adaptive cruise control, our new AI capabilities can stand on the shoulders of prior AI capabilities. So, progressing down the Floridi Curve involves adding successive capabilities to the enveloping environment, thereby increasing AI complexity as we march down the Floridi Curve.

Given this, let's look at the AI Complexity axis in Figure 2. Here, the developer's point of view is to ask the question, "How hard is the AI problem if we can assume various feasible enhancements to the existing envelope?" For example, for adaptive cruise control, the existing mechanical cruise control will give the AI platform access to acceleration and braking.

But that is not enough. Radar must be incorporated into the envelope to make the AI model possible. However, with the assumption of radar, the AI problem is only moderately complex. Hence, we have located it just a third of the way across the axis.

Radar also puts parallel parking and collision avoidance within reach. Parallel parking can be accomplished by teaching our model the same rules learned in driver's ed, along with some collision detection using our radar, and some reinforcement training. Collision avoidance is significantly more challenging but also seems achievable. It would involve radar detection, but AI must also calculate the trajectories and velocities of other cars alongside those of our own vehicle.

Lastly, we face the challenge of retrieving our parked car. As mentioned earlier, this isn't a difficult skill for humans, but it presents a daunting challenge for AI. Of course, radar and the assumption that collision detection capabilities are part of our agency help. However, considering that the car must navigate itself to where we are when we initiate the request, combined with the task of maneuvering through a congested parking garage with imperfect positioning and oncoming traffic, it is quite a challenge. It will require extensive model training to ensure this process is foolproof.

Planning Successive Design Improvements

We now return to the question of using envelopes rigorously. Our design team mapped driver skill difficulty and the current state of AI on the previous graph in Figure 2 (which we'll call Model Year Zero). We'll use Model Year 1 and Model Year 2 to mark the progression of the car's AI capabilities.

In Figure 3, we illustrate how our car's capabilities evolve over the two model years, reflecting designers' ongoing improvements to the envelope and the AI models embedded in the platform. While it may look a little complicated, Figure 3 merely shows the shift down our Floridi Curve through successive planned model years. The distinct colors show in which year the shift occurs.

Figure 3 – Progression from Model Year 0 Through Model Years 1 and 2

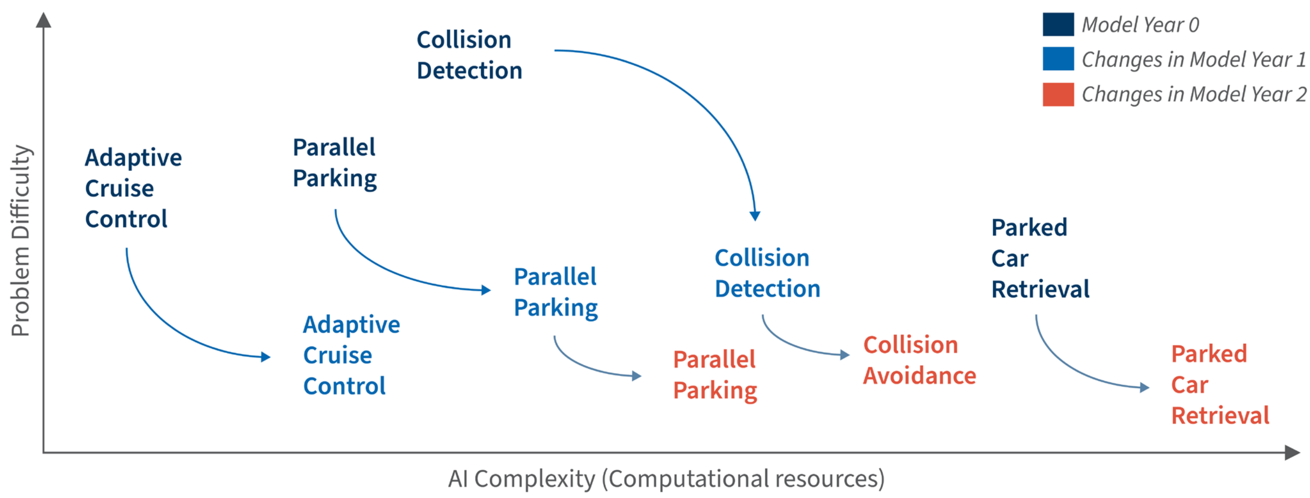

Cruise Control	Navigation System	Lane Change Detection	Blind Spot Detection
Collision Avoidance (AI Skill)	360° Cameras (AI Inputs)	Radar (AI Inputs)	Trajectory Intercept Detection (AI Skill)

Enveloping Capabilities

For example, parallel parking, which had a medium level of human difficulty in the current Model Year Zero (dark blue), has a significantly reduced difficulty after the introduction in Model Year 1 (light blue) of 360 cameras, radar and the parallel parking AI itself. Parallel parking becomes even easier for the driver in Model Year 2 (red) when collision avoidance is added to the envelope. Hence, this ability improves progressively over the course of two model years.

Of note is what happens to our parked car retrieval ability. The graph shows it was not enabled in Model Year 1, because the designers needed time to develop these AI capabilities. Instead, these abilities premier in Model Year 2 when all the enveloping abilities are enabled. The AI resources in the car are significant, but reliance on the driver's skill goes to zero, which is a good thing since there is no driver in the car!

Ability Confirmation

The above procedure has taken us a long way. We now understand the key features of the new envelopes required in each model year, and we know what AI capabilities we need to create. We have also mapped the development to various model years, which gives us the relative timeframes of their development.

However, all of this is necessary but not sufficient. The presumed creation of some of the future abilities of our envelopes in the orange box needs to be verified in order to have confidence in our overall Floridi progression. This does not require absolute certainty, but it does require intellectual confidence that these new capabilities can be achieved. The ability to implement radar and 360° cameras by Model Year 1 should be confirmed with engineers and manufacturers. Equally, the complex AI needed to deliver a parked car to its owner should be evaluated for its feasibility by Model Year 2. Comparing this AI car retrieval problem to other use cases in adjacent disciplines can be affirming. Understanding the approach our AI scientists would take to develop the needed AI can build confidence that these new models can be developed. These are just some of the many ways to achieve this relative certainty.

Let’s now look at a real-world example of how FTI Consulting has built these concepts into our methodology for developing AI systems for our clients.

Putting Theory Into Practice

The following case study from the textbook industry demonstrates how FTI Consulting's methodology for AI implementations incorporates Floridi's concepts.

The Problem of Textbook Demand Forecasting

FTI Consulting’s recent collaboration with a leading educational publisher to improve textbook demand forecasting for the upcoming school year provides a practical illustration of these principles in action. Demand forecasting is a vital issue for the publisher. Textbooks are printed by printing houses, which schedule print runs and assign slots in their schedule to various publishers, and it can take eight to twelve weeks to print a textbook, given the scarcity of these slots. This puts pressure on the publisher to estimate the number of books needed well before the school year begins. Over-forecasting results in wasteful inventory write-offs; under-printing frustrates schools that will require additional books quickly but must wait many weeks due to print cycles.

The company performed well in forecasting aggregate revenue. It would forecast the total number of textbooks it estimated it might sell, multiplied by a set of standard rates, to produce a total revenue forecast for the school year. Back-testing showed that this technique was accurate within 3% of total actual revenues — an extraordinarily good result. However, at the level of the individual textbook, grade levels or states, these forecasts varied wildly. It was only the law of compensating error that made the overall forecast seem accurate.

Textbook forecasting is difficult due to the number of variables involved. First, our client produced hundreds of textbooks in the K-12 market, in almost every subject and every grade level. Consequently, the AI algorithms must consider the demand by title and grade. Then forecasting must be done at the school district level, and there are over 13,300 school districts in the U.S. today.[9] Demand at the district level is also partially determined by their available funding, which can vary dramatically. Recency also plays a part, i.e., how new are the current textbooks the school is using. Schools can often hold onto textbooks for 10 years or more. Another variable is whether a school district is in an “adoption state” or not. Adoption states select textbooks at the state level for buying power and educational consistency. Other exogenous influences such as federal funding and the treatment of creationism also can affect demand.

Humans are very bad at taking all these dimensions into account. Our brains (and our spreadsheets) are just not wired this way. Fortunately, machine learning excels at these types of problems.

Figure 4 shows the key steps in FTI Consulting's methodology for tackling such problems using Floridi Curves and envelopes. Below is a description of each step.

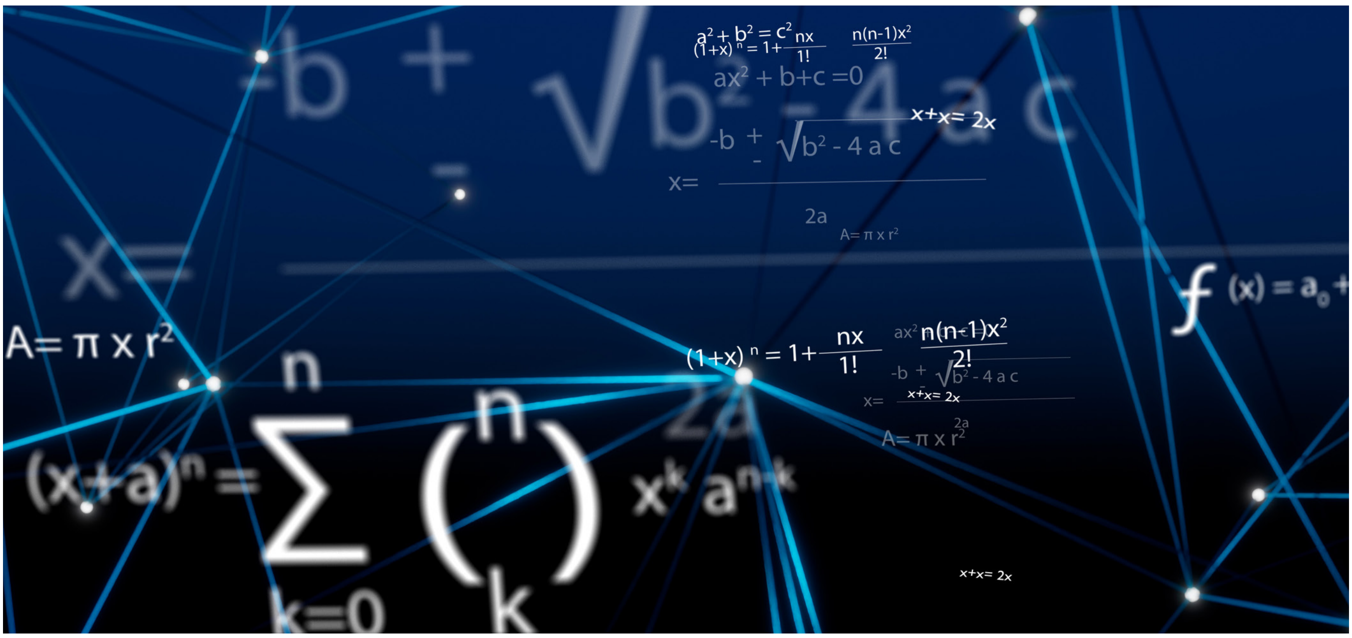

1. **Conceptualizing solutions to the forecasting problem.** Working with our client's team, we first plotted the current state on the X/Y matrix of skills vs. AI complexity. The team was required to analyze everything related to the forecasting process used currently as well as all the data assets available to feed future AI models and any existing enveloping capabilities. The team then theorized the envelopes and AI Models needed to march down the Floridi Curve toward the client's objectives. Once understood, the team developed the X/Y future state map with its needed AI models and envelopes, resulting in a shared understanding between the FTI Consulting team and our client around the approach to achieving the desired future state.
2. **Verifying feasibility of AI models and envelopes.** We developed a proof-of-concept AI model for a fifth-grade math textbook and back-tested its forecast against the company's own forecast to ensure that forecasting improvements were possible. With the client's IT department, we looked at its ability to build the new envelopes (i.e., automated data cleansing routines and pipelines) into a data lake that would feed our AI models. These results created a feedback loop to step 1 where the approach is refined as needed.
3. **Incorporating a checkpoint evaluating the economic feasibility of the approach.** This is the most critical step in the project, defined by three key questions. First, do the model prototypes built in step 2 demonstrate sufficient improvement over current techniques to justify continuing with the envisioned project? Second, does the economic improvement justify the cost of the project? If the textbook modeling shows there would be a 30% reduction in over-printing, saving $100,000 annually, but the project will cost $5 million, then the project has a dubious return on investment, even when amortizing the project cost and savings over five years. Note that the checkpoint may be something other than ROI, depending on the nature of the project. In our car example above, for instance, given that competitors are introducing similar capabilities, staying up with the competition may override any near-term ROI considerations.
4. **Design necessary system components.** For textbooks, this amounts to a routine architecture to design the envelopes and their fit with existing systems. The rollout plan typically consists of project "sprints" common in Agile system development. The AI model building also begins in parallel, not relying on the surrounding systems' designs to be developed.
5. **Implementation.** New system changes, including new enveloping capabilities, are built, tested and then integrated into the company's forecasting workflows. The AI team must also design into the new processes capabilities to evaluate model "drift" over time and the process to correct for it when it occurs.

Figure 4 – FTI Consulting Methodology, A Textbook Forecasting Example

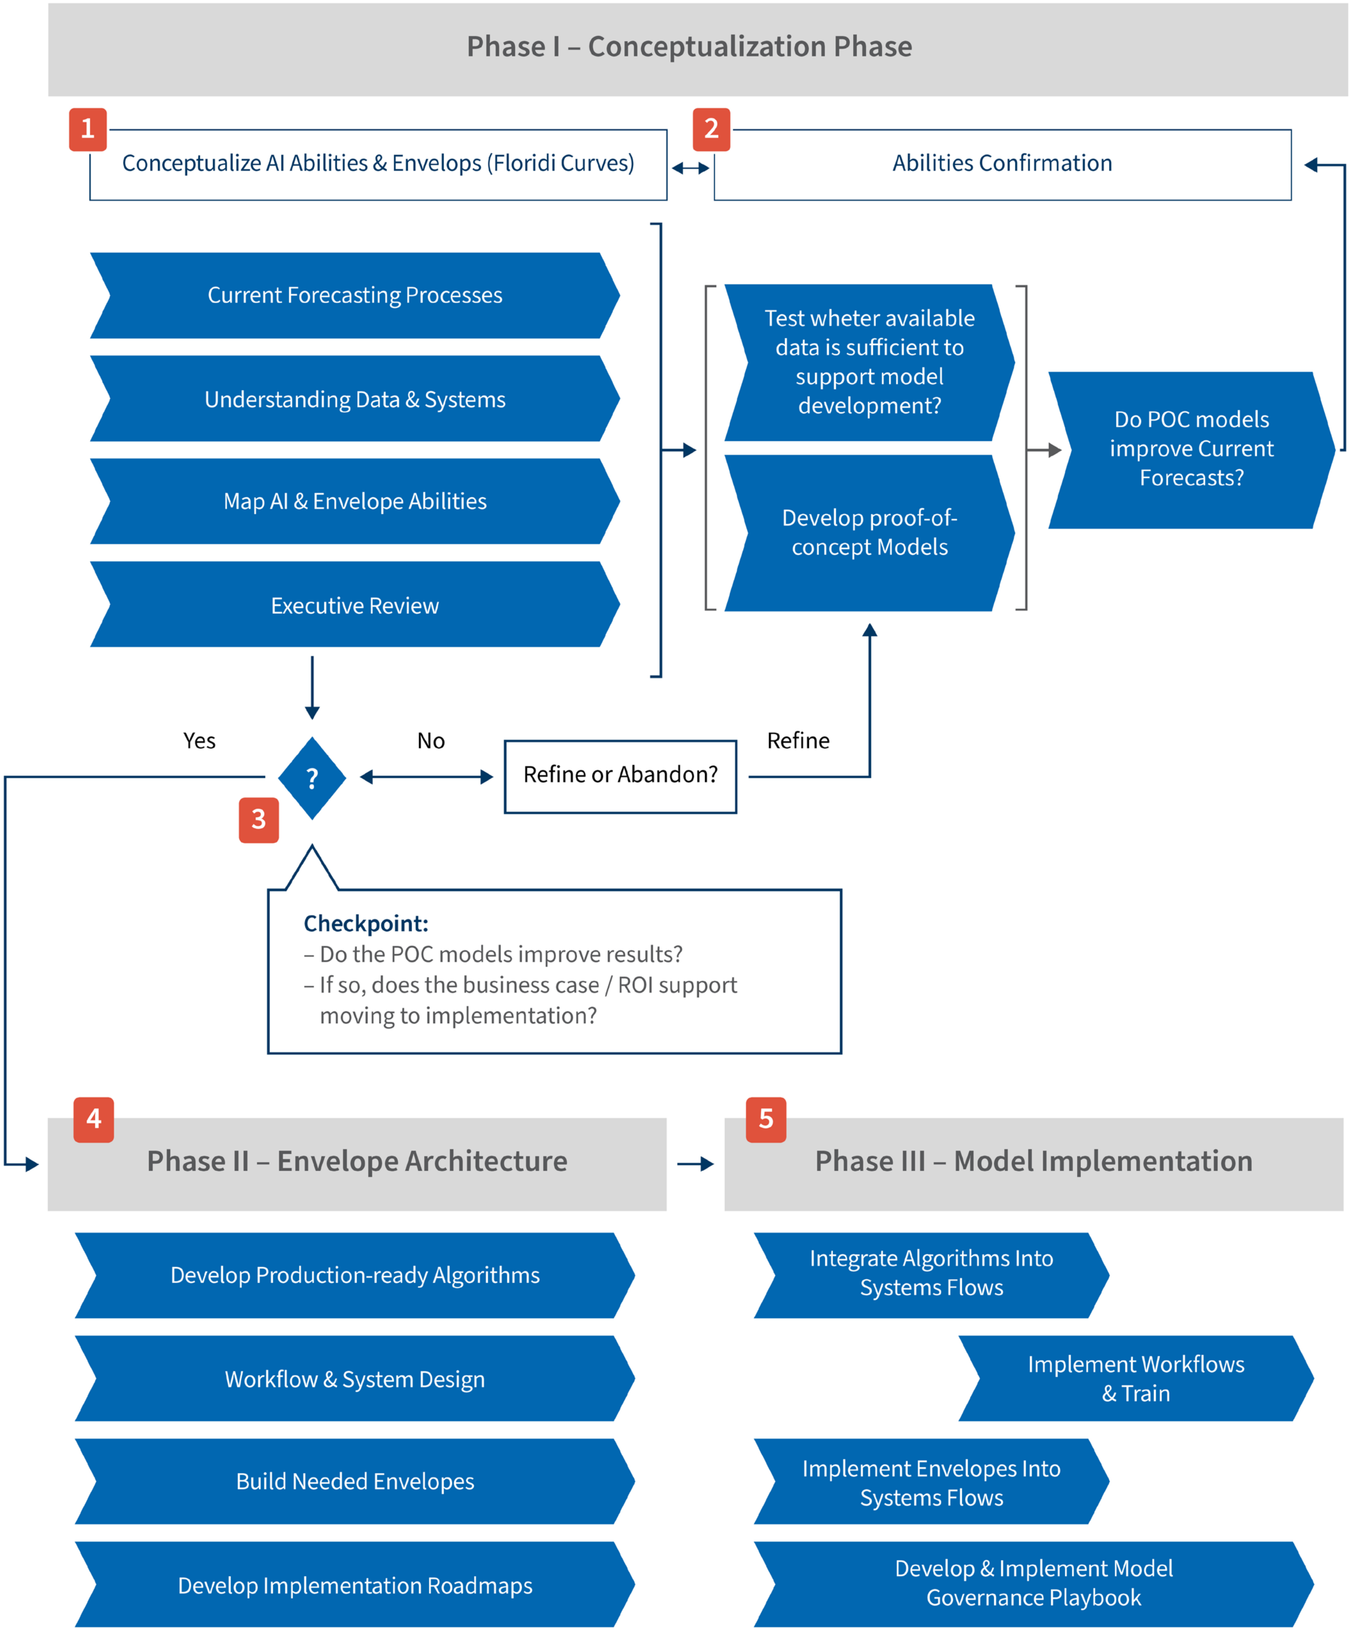

Conclusion

We have covered a lot of new ground in this article. We explored Luciano Floridi's concepts of difficulty versus complexity, Floridi Curves, and envelopes. We then showed how to use these ideas to drive the development of new AI systems to successful completion rigorously.

We believe it is a useful framework for any industry. Of course, each industry will have its own well-honed methods for bringing the actual AI to life. Car companies will collaborate with parts manufacturers, robotic warehouse designers will collaborate with their building architects, and textbook companies will collaborate with their IT departments and data scientists. Yet this framework provides a mutual understanding of the key AI models and envelopes that must be built, and when, in order to achieve the desired outcomes.

Some of the benefits of our framework include:

- Clarity on which AI models and envelopes must be built and when.
- The timeframes for development (model years, etc.) which form the basis of a multiyear rollout plan.
- Strategic budget development, that is, the ability to forecast the conceptual costs in labor, skills, construction, and implementation costs.
- A checkpoint in the methodology for deciding if the gain in AI improvements actually provides the hoped-for return on investment (step 3 in our methodology).
- A framework that can be universally applied across industries.
- A means of achieving a shared vision across disparate groups when developing the Floridi progression (as shown in Figure 3).
- An approach that is easy for any development team to adopt, yet is grounded in the strong theoretical foundations of Floridi's philosophy of AI.
- An easily understood communications framework for boards and CEOs.

In summary, this approach can significantly improve the likelihood of success when implementing AI solutions for business. This is mainly because it sidesteps the biggest risk in AI projects, which is overestimating the return on AI while underestimating the necessary envelopes that must be built to make them successful.

If you have questions about this article, please contact Bruce Benson, whose contact information is below.

About the Collaborators

Bruce Benson, Author

Bruce Benson is a Senior Managing Director and leads FTI's Advanced Technologies and Media Solutions practice. He has over 35 years of professional experience in media, technology, strategy, marketing, advertising, brand strategy, economic analysis, operational effectiveness, complex technology architectures, consumer engagement, intellectual property management and digital transformation.

Mr. Benson is a frequent keynote speaker and the author of many articles on digital strategy and media and has co-led conferences with key Harvard strategists on intellectual property management and copyright law. He was previously appointed as an "industry luminary" to Adobe's product strategy board. As a result of Mr. Benson's work with EW Scripps, FTI Consulting won the Consulting Strategy Firm of the Year award by the Association of Management Consulting Firms. Mr. Benson was the Executive Chairman of Ziff Davis until its sale. Previously, Mr. Benson served as EVP of Corporate Strategy at Young & Rubicam, SVP at Sony Entertainment and Technology Partner at PwC.

Mr. Benson can be reached at bruce.benson@fticonsulting.com

Dr. Luciano Floridi, Advisor

Luciano Floridi is the John K. Castle Professor in the Practice of Cognitive Science and the Founding Director of the Digital Ethics Center at Yale University. Widely recognized as one of the most authoritative voices in contemporary philosophy, he pioneered the philosophy and ethics of information and has become a leading interpreter of the digital revolution. His scholarly output exceeds 400 publications on the philosophy of information, digital ethics, AI ethics, philosophy of technology and history of philosophy, many of which have been translated worldwide. His most recent books are The Ethics of Artificial Intelligence (OUP, 2023) and The Green and The Blue – Naive Ideas to Improve Politics in the Digital Age (Wiley, 2023)". Artificial Agency – Explorations of a Post-AI Culture and Encounters – An Experiment in Distant Writing are forthcoming. Among numerous accolades, in 2022 he was honored as Knight of the Grand Cross of the Order of Merit of the Italian Republic for his foundational contributions to philosophy.

Endnotes

[1] Cooper, R. G., "Why AI projects fail: Lessons from New Product Development," IEEE Engineering Management Review 52 (4): 15–21 (June 26, 2024), https://ieeexplore.ieee.org/document/10572277

[2] Floridi, L., "What the Near Future of Artificial Intelligence Could Be," Philosophy & Technology 32 (1): 1–15 (March 19, 2019), https://link.springer.com/article/10.1007/s13347-019-00345-y.

[3] Ibid. at 9

[4] From an AI scientist's point of view, our reservoir of agency is comprised today solely of "narrow" AI capabilities, which can only perform limited tasks. This is opposed to the ability to perform tasks requiring true artificial general intelligence (human) intelligence, known as AGI, which most scientists agree is at least decades away.

[5] FTI Consulting has enhanced Luciano's diagram.

[6] Floridi, supra n. 2, at 11.

[6] Id.

[7] We use the word "abilities" to distinguish from conventional user requirements, use cases, user stories, and all manner of traditional systems terminology. Abilities are more than these things. They are skills that humans possess, which we seek to replace with AI.

[8] Overview of Motor Vehicle Traffic Crashes In 2023, Traffic Safety Facts, U.S. Department of Transportation, April 2025, https://share.google/R800CZkavdulxYCwg

[9] Claybourn, Clay, Private School vs. Public School, August 19, 2025, U.S. News, https://www.usnews.com/education/k12/articles/private-school-vs-public-school

BRUCE BENSON
Senior Managing Director
+1 203 606 3854
bruce.benson@fticonsulting.com

DR. LUCIANO FLORIDI
John K. Castle Professor in the Practice
of Cognitive Science
Yale University