




EXECUTIVE SUMMARY

Artificial Intelligence: A Framework to Identify Challenges and Guide Successful Outcomes

Lead Analyst: **Cole McCollum**
Analyst

Contributors: **Shriram Ramanathan, Ph.D.**
Director

Kevin See, Ph.D.
VP, Digital Products



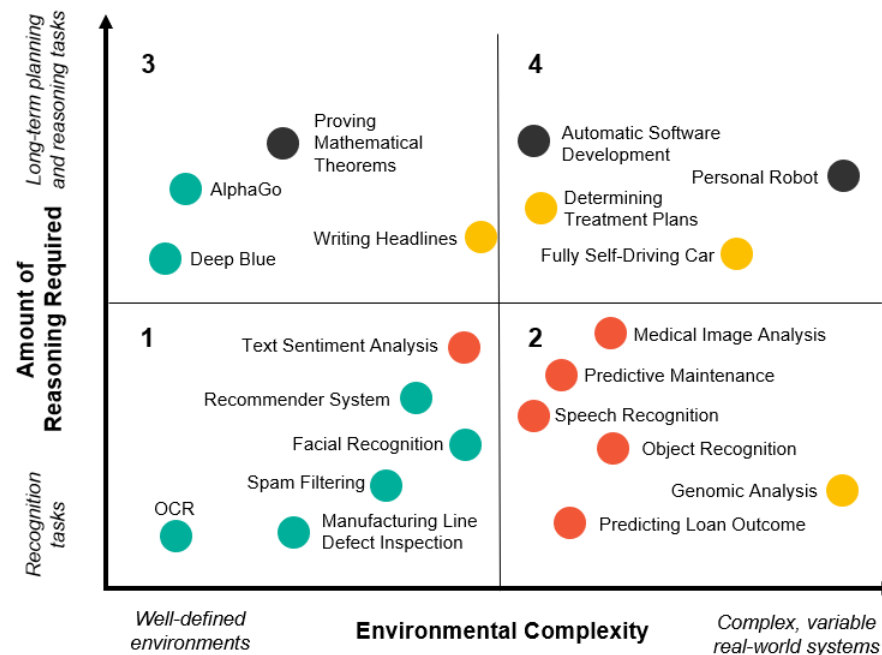
Executive Summary

Given the massive amounts of hype and promise surrounding AI and related technologies like machine learning and deep learning, it's become increasingly difficult to make critical innovation and investment decisions in the space. Taking a technology-first approach (e.g., selecting a vendor simply because they claim to use the latest AI techniques) has led and will lead to many failed companies and projects.

We provide an outcome-focused framework and determined that today's breed of AI techniques is particularly well-suited for pattern recognition tasks. Opportunities exist to leverage this capability, whether for scaling basic human pattern recognition capabilities, emulating expert pattern recognition, or uncovering patterns in data too complex for a human to recognize. As the environment and data get more complex, however, the maturity of today's pattern recognition AI decreases. Tasks that move beyond pattern recognition to tasks that typically require long-term human reasoning to accomplish tend to be more immature or significantly limited in the complexity of environments they can handle.

After an outcome is selected, challenges exist in each phase of AI implementation: data preparation, model selection and training, and deployment. Emerging solutions can mitigate some challenges in implementing AI; however, the impact varies greatly. For instance, tools to handle small datasets like transfer learning and synthetic data are gaining rapid traction, while tools to interpret machine learning algorithms see slower development. Clients should understand the needs of their application and select tools most relevant to the challenge depending on where on the AI spectrum the application falls and the tool's technology readiness level.

Stage of AI Development




AI hype reaches new heights as VC investment pours in globally and tech companies rebrand as AI-first companies

Coming in as the top technology in the [Lux 19 for 2019](#) report, all innovation indicators highlight the massive amount of interest in artificial intelligence (AI) technologies.

In 2019 alone, more than \$30 billion in VC funding went to AI startups, with the bulk of it split between North America and Asia.

Major tech companies like Google have shifted from a mobile-first strategy to an AI-first strategy, prioritizing AI above all other technologies and platforms; likewise, startups with AI as their core value differentiator threaten to displace many incumbents in industries as diverse as automotive and pharmaceuticals.



Andrew Ng: Why AI is the new electricity

Chinese A.I. company SenseTime raises more than \$1 billion in back-to-back funding rounds

Google CEO: A.I. is more important than fire or electricity

'Like A God,' Google A.I. Beats Human Champ Of Notoriously Complex Go Game

Shifting definitions and lack of clarity around what AI really is further add to the hype

Since the field of AI research took off in the summer of 1956, AI has been through several booms and busts, commonly referred to as AI summers and winters, respectively.

- During the last AI winter in the late 1980s and 1990s, researchers used everything besides the term AI to describe what they were working on; **today, everything from a grammar checker to a stock trading application is called AI.** Across industries, everyone from marketers to journalists to CEOs is taking advantage of this confusion by freely using the term AI to describe their research or products.

Unfortunately, AI is a difficult term to classify.

- For starters, the term intelligence doesn't have a universally agreed upon definition (let alone the artificial version).
- Some use AI to describe artificial general intelligence (AGI), an intelligence capable of completing a wide range of tasks like human intelligence, while others use AI to describe narrow AI, an AI that can perform a single task or a few tasks with high competence.
- Some use methods to define AI, while others use outcomes to define AI. For example, after Deep Blue beat the professional chess player Garry Kasparov in 1997, many described the program as lacking "intelligence" because it used brute force search approaches, rather than just observing the outcome.



These tools have been around for decades, but if you don't say AI, customers and investors won't even look at you.

– An oil and gas startup describing its control algorithms



AI and big data – the new oil or the next snake oil?

With all of these advancements in machine learning and data analysis, it easy to understand why AI has become such a hyped technology. However, what is less clear is how these performance gains translate into useful business objectives as well as which applications are achievable with today's AI technology and which are not. While claims from AI companies are nearly limitless and funding seems practically endless, **hints of reality have begun to appear in the AI marketplace:**

- **Marketing meets reality.** Companies that overstated AI's capabilities in complex domains, such as [IBM's Watson for oncology care](#), highlight the somewhat embryonic stage of AI causing expensive and in some cases dangerous failures.
- **Vendors use the AI label simply as a branding strategy.** A [recent study](#) showed that 40% of European "AI" startups actually have do not appear to use any AI. With AI startups attracting 15% to 50% more funding compared to their non-AI peers, many startups are inaccurately rebranding themselves as AI companies in order to attract funding.
- **Early adopters face unintended consequences.** Many early adopters of AI tools have faced major hurdles surrounding issues like bias, the lack of interpretability of AI systems, and cybersecurity threats.



Layoffs at Watson Health Reveal IBM's Problem With AI

AI's failure to live up to the hype is starting to put off investors



Nearly Half Of All 'AI Startups' Are Cashing In On Hype

Amazon scraps secret AI recruiting tool that showed bias against women

To avoid major failures with AI, focus on outcomes, capabilities, implementation, and challenges

As the number of “AI” vendors and projects – many of which claim to use the latest deep learning or machine learning techniques – continues to skyrocket, **clients need an outcome-focused framework in which they can determine which applications to focus on with today’s tools and how to mitigate challenges preventing successful implementations.** At the end of the day, it doesn’t matter what elaborate algorithms a product uses; instead, clients should strive to understand what AI can actually do for them. Clients will be able to use the framework provided in this report to answer key questions like:

With all types of software products being rebranded as AI, how to avoid buying into the latest buzzwords used by a vendor and instead focus on the product’s capabilities?

How to determine an AI application’s maturity level? How to identify opportunities and threats in such a fast-moving field?

What are some of the key challenges to expect while implementing an AI project? How to deal with AI projects that lack large datasets or data science teams? Which challenges will see solutions emerge in the coming years, and which challenges will remain for years to come?

AI is one the most important and impactful sets of technologies available and in development today. In this analysis, we provide an outcome-focused framework that clients can use to maximize gains, minimize threats, and importantly, increase the chances of success for their AI projects.

At a high level, any AI project has four major steps:

PROBLEM SELECTION

Understanding the core use cases and outcomes of AI and mapping those outcomes onto business problems. Included in this step is testing and narrowing down whether an application is possible using today's AI technologies by breaking apart its components and comparing it to other applications on the AI spectrum framework.

DATA PREPARATION

Identifying, cleaning, and wrangling data, whether from sensors, databases, application programming interfaces (APIs), or other sources.

MODEL SELECTION AND TRAINING

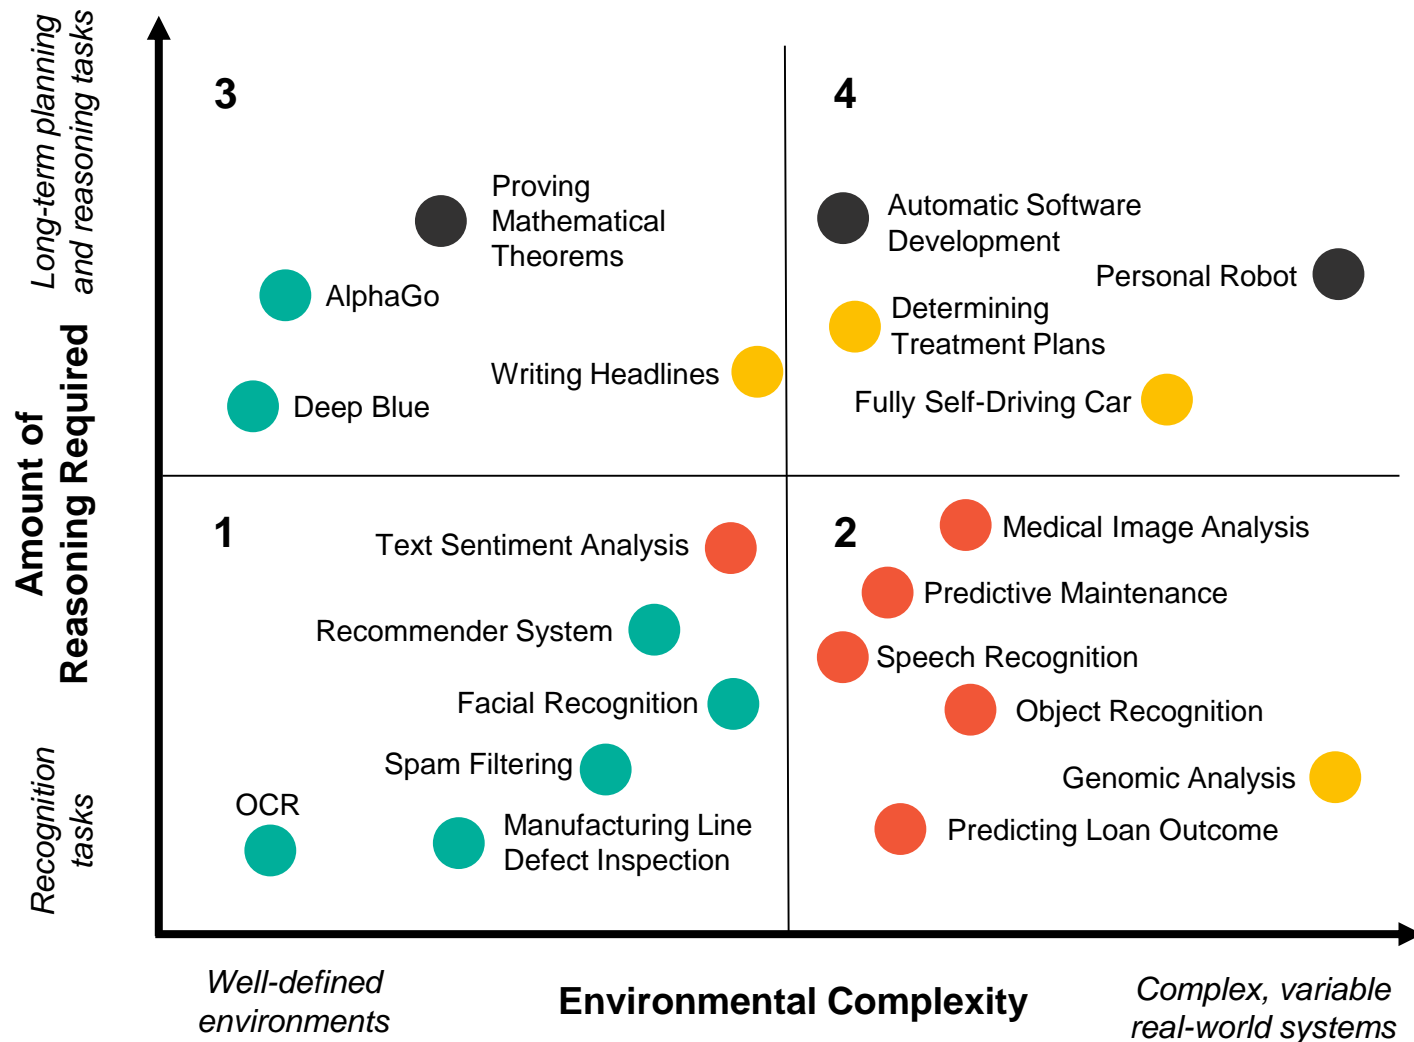
Selecting and testing the right machine learning models for a given problem as well as training the models.

DEPLOYMENT

Utilizing and monitoring machine learning models where they can be used to make predictions with new data.

PROBLEM SELECTION

Depending on the level of AI required, AI applications are all over



Here, we map out examples of AI applications onto the framework and color them by their associated maturity level. In many cases, applications tend to cluster together in different parts of the framework depending on their maturity level.

Stage of AI Development

- Mature
- Late Development
- Early Development
- Research

Machine learning-based data cleaning and wrangling tools can help automate the time-consuming process

WHAT IT IS

Emerging data cleaning and wrangling tools can automatically identify “dirty” data and remove or correct it. Some tools like [Datalogue](#) are designed for the cloud and tabular data like customer data while others like [Foghorn Systems](#) can run at the edge and are designed for time-series sensor data.

EXAMPLE DEVELOPERS



CHALLENGES IT ADDRESSES

Data Cleanliness

STAGE OF DEVELOPMENT



SCALE

LUX TAKE



ENGAGE

Investing in data cleanliness is often an overlooked area when building AI projects; however, it can take up to 80% of a data scientist’s time. Clients should engage with vendors applying automated data cleaning and wrangling approaches, particularly at the edge, where savings in transmission costs and storage can be vast.

Model selection and training are at the heart of any AI project

WHAT IT IS

Model selection involves testing and selecting the best-performing model or group of models. Training involves the process of adjusting parameters to fit a model to a dataset and optimizing its performance.

WHY IT MATTERS

Depending on the type of problem, data (e.g., text, image, numerical, etc.), and other factors, the appropriate model to use varies greatly. Furthermore, training the model properly is key for it to make accurate predictions.

KEY CHALLENGES

Small Datasets

Machine learning approaches, particularly deep learning, can require massive datasets to reach maximum performance. However, many potential applications of AI only have sparse datasets. Furthermore, because most applications of machine learning are supervised, the data must be labeled, which can be highly laborious.

Privacy Concerns

In some applications, such as healthcare and financial services, it can be challenging to access data to train models due to strict privacy regulations. New methods are needed to preserve the privacy of the data while still allowing for valuable insights to be extracted from AI models.

Deployment of AI systems brings a whole new set of challenges

WHAT IT IS

Deployment involves exposing a trained machine learning model to new data for inferencing on an ongoing basis. This process also includes monitoring and in some cases explaining how an AI system arrived at a particular insight.

WHY IT MATTERS

Deployment crosses the line from creating an interesting prediction in a development environment to running models at scale, in real-world environments. While this is, of course, where all the value of machine learning is generated, it also creates new issues, such as the need to explain results and cybersecurity.

KEY CHALLENGES

Interpretability

A fundamental issue surrounding machine learning, particularly deep neural networks, has been that, because of the sheer number of layers, nodes, and connections within a single neural network, it is difficult to understand how these networks arrive at insights. Understanding how these models do so will become increasingly important as AI makes its way into more high-stakes industries like healthcare, defense, and autonomous vehicles.

Edge Requirements

While many current AI deployments operate in the cloud, challenges surrounding latency, bandwidth limitations, and privacy are creating increased demand for edge computing capabilities whereby models can make predictions right where the data is generated.

Focusing on the outcomes and capabilities of today's tools will help minimize the risk of AI failures

Start with the end in mind. Rather than focusing on technologies first (i.e., deploying AI techniques like machine learning for the sake of deploying AI), to reduce the likelihood of an AI failure, start with the end in mind by determining the problem and outcome to solve for. Then, work backwards to determine the level of AI needed to solve that problem and whether it is feasible with today's tools using the AI framework. Following that step, identify the key challenges in development and deployment and determine whether emerging solutions will be capable of mitigating those challenges.

Seek to identify opportunities in pattern recognition. Today's breed of AI techniques – with an emphasis on machine learning – is particularly well-suited for recognizing patterns in data. Many opportunities exist to leverage this capability, whether for scaling basic human pattern recognition capabilities (e.g., speech recognition), emulating expert pattern recognition (e.g., medical image analysis), or uncovering patterns in data too complex for a human to recognize (e.g., gene sequence analysis), and clients should focus significant efforts on these classes of applications. However, clients should be aware that the more complex the environment, the more immature the application will be with today's technologies. Applications that require longer-term planning or reasoning capabilities from a human (e.g., surgery, writing a useful article, or driving a car that can operate in any environment) are likely many years away from full, successful implementations. Clients should be cautious of investing too early in such applications or should consider constraining the desired application to make it a better fit for today's technology.

Outlook

The AI hype will remain, but investors and customers need to increasingly focus on ROI. As AI continues to become a technology that everyone has access to, taking a technology-first approach will lead to many failed companies and efforts. Clients should be cautious of working with companies simply because they claim to use the latest algorithms, and instead focus on the return and value that the solution provides.

Many applications in Quadrant 4 of the AI spectrum are still years away from the full solutions promised. For example, fully autonomous cars capable of driving on any road in any city will take significant further development. Clients should focus on constraining these problems (e.g., self-driving cars on college campuses or retirement communities) to reduce the amount of AI needed in a solution or should plan for long development cycles. Clients seeking nearer-term opportunities with AI should play into the strength of today's machine learning by focusing on simpler pattern recognition applications.

As technology continues to knock down barriers to implementation, cultural elements present major challenges to successful AI deployments. Clients in the manufacturing sector, for example, have indicated that scientists and engineers in their organizations are reluctant to accept insights from black-box AI systems, thereby posing a big cultural barrier to adoption. Clients should seek to get buy-in from the end users of the application as it is being developed to ensure that the end users will both trust and use its predictions. Likewise, because machine learning is an inherently statistical tool, adoption will be more successful in applications that can tolerate some inaccuracy. Clients interested in developing AI solutions for more critical applications should be sure to incorporate fault tolerance into the design (e.g., keeping a human-in-the-loop).



Innovate Smarter & Grow Faster With Lux

Contact us:

www.luxresearchinc.com press@luxresearchinc.com



[@LuxResearch](https://twitter.com/LuxResearch)



[@LuxResearch](https://www.facebook.com/LuxResearch)



[Lux Research, Inc.](https://www.linkedin.com/company/Lux-Research-Inc)

Blog:

[Lux Blog](#)

Free Webinars:

[Lux Webinars](#)

YouTube:

[Lux Research](#)