A Guiding Framework for Applying Machine Learning in Organizations

by

Alan Tham

M.Eng., Electrical and Electronics Engineering
Imperial College, 2003

Submitted to the MIT System Design and Management Program
in partial fulfillment of the requirements for the Degree of

Master of Science in Engineering and Management
at the
Massachusetts Institute of Technology

September 2016

Signature of Author: Signature redacted
Alan Tham
System Design and Management Program
5 August 2016

Certified by: Signature redacted
Patrick Hale
Director, System Design and Management Program
Thesis Advisor

Accepted by: Signature redacted
Stuart Madnick
John Norris Maguire Professor of Information Technology, Sloan School of Management
Professor of Engineering Systems, School of Engineering
Thesis Reader

Accepted by: Signature redacted
Warren Seering
Weber-Shaughness Professor of Mechanical Engineering

A Guiding Framework for Applying Machine Learning in Organizations

by

Alan Tham

Submitted to the MIT System Design and Management Program in partial fulfillment of the requirements for the Degree of Master of Science in Engineering and Management

Abstract

Machine Learning (ML) is an emerging business capability that have transformed many organizations by enabling them to learn from past data and helping them predict or make decisions on unknown future events. While ML is no longer the preserve of large IT companies, there are abundant opportunities for mid-sized organizations who do not have the resources of the larger IT companies to exploit their data through ML so as to gain deeper insights. This thesis outlines these opportunities and provide guidance for the adoption of ML by these organizations.

This thesis examines available literature on current state of adoption of ML by organizations which highlight the gaps that motivate the thesis in providing a guiding framework for applying ML. To achieve this, the thesis provides the practitioner with an overview of ML from both technology and business perspectives that are integrated from multiple sources, categorized for ease of reference and communicated at the decision making level without delving into the mathematics behind ML.

The thesis thereafter proposes the ML Integration framework for the System Architect to review the enterprise model, identify opportunities, evaluate technology adoption and architect the ML System. In this framework, system architecting methodologies as well as Object-Process Diagrams are used to illustrate the concepts and the architecture. The ML Integration framework is subsequently applied in the context of a hypothetical mid-sized hospital to illustrate how an architect would go about utilizing this framework.

Future work is needed to validate the ML Integration framework, as well as improve the overview of ML specific to application domains such as recommender systems and speech/image recognition.

Thesis Advisor
Patrick Hale
Director, System Design and Management Program

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgements

I could not possibly mention all of the many people that have made some contribution to this work, or have helped me to get to and through MIT challenging curriculum. I would not be the person I am today without my wife and my family. I would like to thank my wife, for without her loving support, daily calls and sacrifices, I would not be able to complete this one-year education journey. One year away from her is indeed a long time.

I want to thank my Thesis advisor – Pat Hale. Pat has been instrumental in giving me the liberty to pursue my own professional interest and guiding me in the practical aspects of this work. Pat's support and faith in my work has kept me motivated throughout the Summer of 2016. Thank you Pat and Best Wishes for your health and retirement.

I want to also thank my Thesis reader – Professor Stuart Madnick. Stuart's guidance at the later stages of my thesis was invaluable and I thank him for kindly agreeing to be my reader with such a short notice.

I want to thank all my SDM classmates and friends at MIT for my journey was enriched by the opportunity to experience this MIT fire-hose drinking journey with you. I hope that our paths will cross in the future.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

List of Figures

List of Tables

Chapter 1 Introduction

"In God we trust. Others must bring data"
– Edwards W. Deming, Management Theorist

Over the last decade, the increase in compute power, availability of more data as well as improvements in models and algorithms has led to the expansion of Machine Learning (ML) techniques into a host of domains such as governments, financial institutions and healthcare. These ML tools have enabled these organizations/ businesses to learn from past and present data to help them predict or make decisions on unknown future events. As such, ML, as an emerging business capability, have become increasingly ubiquitous and integral to many business operations. While ML is no longer the preserve of large IT companies, there are abundant opportunities for mid-sized organizations who do not have the resources of the larger IT companies to exploit their data through ML so as to gain deeper insights. This thesis outlines these opportunities and provide guidance for the adoption of ML by these organizations.

This thesis examines available literature on current state of adoption of ML by companies, across geographical regions and highlights the opportunities as well as the gaps that drive the motivation for the thesis: A framework to guide companies on adopting ML for their business, supported by an overview of ML from both technology and business perspectives.

It is noted that ML is a branch of Artificial Intelligence (AI), the latter being the multi-disciplinary field encompassing perception, reasoning, learning and decision-making with action. ML is the "learning" branch of AI and should be viewed as a key stepping stone towards incorporating AI into organizations. That said, this thesis will not focus on AI per se, primarily because the foundation of AI continues to evolve and AI is a concept rather than a technique. Different interpretations abound regarding AI, which makes an analysis of AI unwieldy. This thesis henceforth focusses on ML where the foundations are well understood and where its value have been proven in enhancing businesses.

This first chapter provide the motivations behind the proposed framework for applying ML in organizations and the research investigation underlying this thesis that integrates both engineering and management perspectives. Section 1.1 starts with the the motivations that leads to this thesis. Section 1.2 further details its aim and objectives which underscores the thesis scope. Section 1.3 presents the research objectives and questions that underpin the thesis. At the end of the first chapter, a complete overview of the thesis is presented in Section 1.4.

1.1 Motivations

Machine Learning (ML) is a relatively new field in Data Science that has its roots in the 1950s. It has recently come on its own as a scientific discipline in the late 1990s due to the advances in computing power which allows data scientists and engineers to train computers to learn from the glut of data, without having to extensively feature-engineer the data model. ML has become an integral part to many large IT and data-based companies such as Google, Facebook or Twitter, as well as Banks (companies that are commonly termed as "digital natives") (S. Ransbotham, Kiron, & Prentice, 2015). However, anecdotal evidence as observed in this thesis suggests that the adoption of ML is localized to certain types and size of industries. ML has yet to pick up traction in other industries and smaller companies despite the lauded benefits. The literature review will examine the opportunities offered by ML and challenges to the state of adoption of ML by such companies and industries.

One key observation that is highlighted in the literature review notes that despite the various benefits and opportunities of ML technologies, C-level and mid-management do not completely understand ML such that they could begin exploiting these techniques to enhance their businesses. ML is a recent technology innovation such that many of these decision makers belong to a generation that may not have the academic or professional training to understand ML concepts. As such, they have to depend on in-house experts as well as the surrounding noise and signals around ML to guide their decision making.

Furthermore, ML is a complex, multidisciplinary field that incorporates statistics, probability, calculus and computer science. Its theories, tools and applications are currently fragmented as researchers seek to explore frontiers while developers build newer tools. Given this broad, diverse, complex and dynamic landscape that is identified through the literature review, there is no clear pathway with which organizations can understand and apply ML. Little attention is paid to explain the value proposition, tradeoffs, tools and methods of ML to the management level, while managing and driving various stakeholder needs and requirements in the context of their organizational domains.

The literature review hence outlines the benefits of ML through case studies, which provides the drivers for adoption by these companies. On the flipside, the literature review also highlights the challenges in ML adoption. These observations in the literature review thus underscores the need to provide a broad understanding of ML, pitched at the decision-making level, as well as a framework with which decisions makers can use to launch their organizations into the era of Data and Information Age 2.0. The specific aims and objectives of this thesis is outlined in the next section.

1.2 Aims and Objectives

This thesis aims to provide a framework to guide organizations on integrating ML into their enterprise. Firstly, the thesis outlay an overview of ML, its benefits and challenges, that is communicated at a sufficient level for decision makers to appreciate the benefits of ML to their business as well as understand the subject matter in broader trends. The thesis therefore bridges the gap between science-based ML techniques and management literature. It focuses on ideas and trends, not on the math or models and cover ML from both technology and business perspectives.

Next, the thesis will articulate the ML Integration framework that guides organizations in how they could implement ML effectively. This framework takes a holistic approach to architecting ML into the organization, by considering various technological, environmental and organizational aspects as well as stakeholders needs.

Four different research objectives can be distinguished from the above specified aims:

- **First**: Review the benefits of ML and the challenges of ML adoption by conducting a literature review.
- **Second**: Provide a technology perspective of ML by presenting an overview of the definition and concepts, landscape and techniques.
- **Third**: Provide a business perspective of how industries have been applying ML and in the process highlight the market drivers and barriers and investigate the opportunities, challenges and trends in ML across the industry.
- **Fourth**: Integrate both knowledge in above two research objectives in an interpretative manner to formulate the ML Integration framework that will guide organizations in considering ML for their businesses as well as resources/techniques/methodology to be considered during implementation.

The ML integration framework is thereafter applied to a hypothetical example of a mid-sized hospital, to demonstrate how a System Architect ("architect") would formulate his approach. The next Section 1.5 highlights the specific research objectives and questions that arise from this articulation of the problem statement.

1.3 Research Objectives and Questions

To achieve the above aims and objectives, the key research question was designed using a To-By-Using framework that articulates the System Problem Statement (Crawley, Cameron, & Selva, 2015) as:

System Problem Statement

To: Support organizations in ML adoption and implementation
By: Providing an understanding of ML in the technology and business context, and an ML Integration Framework
Using: Literature Review, Technology and Market Analysis and System Architecting Principles

The following Detailed Research Questions (RQ) are developed to answer the above main Key Research Question. These Detailed RQs are grouped into the four categories shown in Table 1.1, namely – *RQ1: Literature review, RQ2: ML from Technology Perspective, RQ3: ML from Business Perspective* and *RQ4: ML Integration Framework.* Each of these categories represent one major chapter in the thesis and combines to answer the key RQ stated above.

Table 1.1 : Detailed Research Questions (RQ)

Categories	Detailed RQ
RQ1: Literature Review Review the benefits of ML and challenges of ML adoption by conducting a literature review	• What are the Key Benefits of applying ML in various industries? • What are the Key Gaps in understanding ML from the management perspective? • What are the perspectives from Management on ML? • What are the concerns regarding ML? • What are some of the problems facing deployment of ML? • Where are some of the confusion? (Definitions, concepts)
RQ2: ML from Technology Perspective Provide a technology perspective of ML by presenting an overview of the definition and concepts, landscape and techniques	• What is ML? What's the underlying theory? • How can ML be classified in the context of Artificial Intelligence and Big Data • What different techniques/methods exist today and how can we categorize them? • What are the Key Gaps in understanding ML from the management perspective? • What are the challenges of ML techniques and tools? • What are the current and future trends?
RQ3: ML from Business Perspective Provide a business perspective of how industries have been applying ML and in the process highlight the market drivers and barriers	• Industry Trends • What are the Market drivers? • What are the Market barriers? • What strategies do they adopt ML? • Industries that benefits from ML
RQ4: ML Integration Framework Integrate both knowledge in above two research objectives in an interpretative manner and presents the ML Integration framework that will guide organizations in considering ML for their businesses as well as resources/techniques/methodology to be considered during implementation	• Application of ML framework • Elements of the framework ○ Review of Enterprise Model ○ Identification and Assessment of Opportunities ○ Evaluation of Technology Adoption ○ Architecting the ML system

1.4 Outline of Thesis

This thesis consists of seven chapters that describe the background, research methodology, literature review and analysis, current strategies, framework development and conclusions to the above research questions. This Section briefly describes the content of these chapters.

Chapter 2

Reviewing existing literature from journals, magazines and Internet on opportunities of ML and the challenges of ML adoption, which then motivates the subsequent parts of the thesis.

Chapter 3
Outlines research methodology used in this thesis and explains the rationale for the approach.

Chapter 4
Provide an overview of ML from the Technical Perspective, through definitions, concepts and examples. The aim is to present the overview clearly at the management level such that the chapter resolves the confusion and misunderstanding identified in the literature review as well as convey the fundamentals without going into technical details on the theories behind ML.

Chapter 5
Provide an overview of ML from the Business Perspective, through survey of industry that employs ML, identifying market drivers and barriers, their strategies and frameworks for adopting ML. This chapter also highlights potential opportunities in industries that would benefit from ML.

Chapter 6
Proposing a ML Integration framework to support organizations in implementing their ML strategy, as well as the analytical tools required to conduct tradeoff analysis between resources.

Chapter 7
Applying the framework to a hypothetical example of a mid-sized hospital.

Chapter 8
Summary of findings of this research study and recommends area for future work.

Chapter 2 Literature Review

This Chapter highlights the benefits of Machine Learning (ML) as well as the challenges faced by businesses in adopting ML. The observations in this literature review motivates the main thrust of this thesis where a broad understanding of ML, as well as a decision-making and implementation framework is outlined. These are pitched at the management level to foster understanding of this scientific field. This literature review does not delve into details of scientific research; rather it highlights observations based on the review of business literature. As highlighted in the Introduction Chapter, this chapter focusses reviewing literature with ML as its underlying technology theme. It will not review literature that touches on other lateral fields of Artificial Intelligence (AI) such as robotics or reasoning because they do not rely on ML to perform these tasks.

The Research Questions that are addressed in this literature review are:
RQ1: Literature Review

- What are the Key Benefits of applying ML in various industries?
- What are the Key Gaps in understanding ML from the management perspective?
- What are the perspectives from Management on ML?
- What are the concerns regarding ML?
- What are some of the problems facing deployment of ML?
- Where are some of the confusion? (Definitions, concepts)

Section 2.1 discusses briefly the sources of information for the literature review and reasons for the approach. Section 2.2 outlines the benefits provided by ML through common industry case studies. Section 2.3 highlights the observations on the challenges of adoption and proposes recommendations to resolve them. Section 2.3 concludes.

2.1 Sources

There are many types of business-related publications on AI and much less so ML as a specific topic of interest. Due to the dearth of information on ML, as well as the confusion of terms used, the literature review will start by surveying all relevant topics of AI, ranging from big data analytics, predictive analytics, Machine Learning, Deep Learning, Smart Machines, Cognitive Systems and AI in general. The focus of the search should be for information that articulates how these technologies, be it in decision-making, sense-making or automation, enable businesses as well as the challenges faced. This expansionist approach will also highlight the challenges of understanding MI.

An initial analysis of available publications revealed that many of them are conclusions or platitudes without much supporting evidence. In order to form a better data-driven picture of the perception and understanding of ML, information was obtained from credible and publicly available sources. Some of such sources are consultancy reports from large consultancy firms operating in the IT space such as *Accenture* and market research reports from firms such as *Gartner* and *BCC Research*. Some of the information was also obtained from scientific journals and respected magazines but these were infrequent. These reports were often premised on in-house surveys, which are useful proxies to understand the perception and understanding of ML/AI. Even so, many of these reports could not entirely avoid ambiguous statements and conclusions, due to the evolving nature of ML and AI. It is therefore not possible to find literature that solely focusses on Machine Learning.

2.2 Key Benefits of ML

Improved Top Line Profitability and Bottom Line Efficiency

There are many reports that highlight the benefits that ML brings to top line profitability as well as improved business efficiencies. A 2015 Mckinsey report notes that more than a dozen banks in Europe had replaced statistical financial models with ML techniques to accurately forecast the needs of their client base as well as provide targeted retail services through recommendation engines. In some cases, these banks have reported "10 percent increases in sales of new products, 20 percent savings in capital expenditures, 20 percent increases in cash collections, and 20 percent declines in churn". Similarly, General Electric had leveraged on ML techniques to develop insights into their data in order to optimize performance, anticipate breakdowns and streamline maintenance, hence saving operating costs (Pyle & San Jose, 2015). Harvard Business Review (HBR) also notes in a May 2016 report that by studying 168 early adopters, ML has enabled speed improvements of two times or more for most business processes by taking over routine digital tasks involves large amounts of unstructured data (H. J. Wilson, Sachdev, & Alter, 2016)

Value Generation by Making Predictions

The most common application of ML involves making predictions based on prior data, particularly in complex environments where there are multiple variables. There are a few widely publicized examples of such predictive applications, such as self-driving cars ("Google Car"), recommendation engines ("Netflix"), natural language understanding ("Facebook", "Siri") and fraud detection ("Crowdstrike"). These companies (or subsidiaries) have transformed their value propositions by integrating ML as a key technology enabler.

Beyond these applications, there are several other novel applications of ML, depending on the identification of business needs and the executives' understanding of the benefits and limitations of ML. For instance, McKinsey utilized ML to examine scanned resumes and filter potential recruits based on prediction of performance of past recruits. This enabled McKinsey to consider a more diverse range of profiles and counter hidden human bias in their recruitment process (Pyle & San Jose, 2015)

While specific cost/benefits are not articulated in these aggregated reports and the benefits could vary across industries, there are many other application domains where ML could revolutionize hitherto legacy approaches, particularly those where data is readily available for "learning" patterns. The main thrust of this thesis is to identify and classify such application domains, outline the benefits, challenges and trends of ML and subsequently provide a framework for managers to integrate ML into their businesses.

2.3 Challenges of ML Adoption

Confusion over terms

Available surveys suggest that many people in the enterprise think of AI in many different ways. According to Narrative Science where 200 respondents across multiple industries and rank levels were surveyed, 31 percent still think that AI is a "technology that thinks and acts like humans", with varying percentages for different definitions (Narrative Science, 2015).

There is also no definitive way in which these reports define AI, ML or any of related fields such as "cognitive computing", "data analytics", "data science" and "smart machines" (Austin, 2016). Even reports by the same consulting company (Brant & Austin, 2015; Walker, Cearley, & Burke, 2016) refers to ML in different contexts. The 2016 report classifies "Advanced Machine Learning" what the 2015 report defines as "Machine Learning". *Gartner* also classifies technologies like social network analysis and VR – both of which are not viewed as AI technologies – as "Smart Machines".

Given this opacity over the terms used in these commercial publications to differentiate products/applications, technology and the science, it is likely that users (businesses and individuals) will interpret these terms differently, leading to misunderstanding of the value and obstructing the adoption of these new technologies. Hence, there needs to be a set of definition and communication framework such that enterprises can communicate in the same language and avoid misunderstanding within their ranks.

Different attitudes towards Adoption

There is a huge interest by companies willing to invest in ML technologies, particularly by "digital-natives" companies. However, these companies have indicated that the business value from being an early-adopter or early-majority is unclear. In one survey, 76% of 437 companies indicated an interest to invest in such technologies over the next two years (Kart & Heudecker, 2015), while in another survey only 4% of 400 companies said they had the right strategic intent, people, tools and data to draw meaningful insights from that data and to act on them (Sinha & Wegener, 2013). In a recent *MIT Sloan Management Review* (S. Ransbotham et al., 2015) that surveyed 2,719 business executives, managers and analytics professionals, there is a growing gap between the ability of organizations to produce analytical findings and their ability to consume them to create business value. In an *Accenture* survey of 30 pilots in early-adopter companies, while there are reports of improvements of top-line performance by incorporating ML into business processes, it is unclear at which level of employees were these benefits perceived (H. J. Wilson, Alter, & Shukla, 2016). These data suggest that, notwithstanding the interest and the hype, the business value of ML needs to be clearly articulated. This also motivates the need for a framework to drive strategic decisions on investing in a ML capability.

On the other hand, many "analog" companies are likely to maintain a wait-and-see approach towards integrating disruptive innovation to their businesses (Weinelt, Shah, Spelman, Ilangovan, & Gul, 2016). This attitude is likely due to unclear business value from adopting these technologies as well as some of the broader organizational culture challenges that these analog companies face in the face of digital disruptions to their businesses. Such attitudes motivate the need for a framework that provides a roadmap for these analog companies towards adopting ML technologies, that will become useful as these companies seek to transform themselves into digital enterprises.

Differing Trust/Confidence Levels

Over and above unclear business value and challenges faced by companies, there is also mistrust of MI/AI across the ranks of employees. *Gartner* reports that many people believe ML is all about automation and eliminating human input and therefore they mistrust the adoption of these intelligent tools which could make their jobs redundant (Linden et al., 2015). *Accenture* highlighted that while 46% of top-level managers strongly trust the advice of intelligent systems, only 24% of middle-level and 14% of first-line managers had confidence in these systems (Kolbjørnsrud, Amico, & Thomas, 2016). The survey also found out that middle and first-line managers want an understanding of how the system works to generate advice, i.e. the perceived benefits of the system. A 2014 *EIU* survey observed that 30% of the CEO polled continued to rely on their gut instinct to make big decisions, despite the availability of these analytics tools that provide insight into their data. Furthermore 52% of C-suite level mangers and 44% of non C-suite managers stated that they have discounted data that they do not understand (Economist Intelligence Unit, 2014).

The hypothesis is that users' attitudes towards new technology are still centered on being in control and using computers to arrive at the best answers. While they may be comfortable at some levels of automation that makes work and life easier, they are not yet experienced and prepared to work as partners with or even subordinates to intelligent systems. There is therefore a need to bridge the trust gap by considering a technology adoption framework that companies can use to incorporate these new technologies into their business.

Lack of Adoption Frameworks and Practical Guidance for ML

While a large number of generic adoption framework exists for companies to develop technology capabilities, none of them are specifically geared towards developing a capability in ML. The World Economic Forum White paper proposes a digital business transformation approach as well as a priority matrix for companies seeking to transform themselves into digital enterprises (Weinelt et al., 2016). The closest to an available framework is the "Smart Machine Primer" proposed by *Gartner* (Austin, 2016). These generic frameworks focus on companies that have already considered the transformation to digital but none have provided guidance to "analog" companies that have some form of data in their systems but no capability to automatically gain insight from such data through ML.

At the same time as these specific frameworks are useful from the strategic business perspective, a practical guidance for ML from the technical perspective is also important for enterprises to understand the benefits and limitations of ML and the specific domain areas that ML can be applied and where it can fail. This literature review has uncovered several nuggets in the academic community, particularly from Universities, Journals, MOOCs and the online community but there are no specific and practical guidelines available in business and management literature. Businesses have to rely on their in-house talent for advice or failing that, depend on consultancies.

Narrow Implementation Approaches

Even with clear business goals that sets ML as a technology capability to gain competitive advantage, proposed implementation guidelines lacks a deeper analysis of the tradeoffs in resources. *Gartner* have proposed a "Vendor Selection Guidelines" in that summarizes the different options of buying, building or outsourcing these capabilities for text analytics but does not provide a holistic framework for analyzing the different tradeoffs across a wider scope of companies – especially for the "analog" enterprises mentioned earlier (Duncan, Linden, Koehler-Kruener, Zaidi, & Sharma, 2015). This motivates the need for an encompassing implementation framework that companies can consider to develop ML capabilities, in development as well as in a production environment, in tandem with the abovementioned strategic framework for adoption.

Lack of Talent

There is a dearth of talent and expertise across the all levels that challenges the exploitation of ML as a competitive advantage. *Gartner* reported that despite the range of recognized use case, the successful adoption of ML depends on finding talented data scientists that can execute the technology while understanding its pitfalls and limitations (Kart & Heudecker, 2015). *EIU* added that 43% of North American C-suite thinks their senior management colleagues lack sufficient skills or expertise in utilizing big data for decisions making (Economist Intelligence Unit, 2014).

As a result, companies are providing formal or on-the-job training in-house, leveraging on the domain expertise required for understanding how to apply ML. Nevertheless, surveys show that half of respondents cited turning analytical insights into business actions as one of their top analytics challenges, particularly from the management layer (S. Ransbotham et al., 2015). This motivates the need for a comprehensive overview of ML that is pitched at managers so that they can immediately understand the benefits and challenges of integrating ML as a technology capability in their businesses.

2.4 Conclusion

The foremost challenge of the literature review is to separate platitudes and assertions common in these publications from credible data and analysis. Given the hype and interest surrounding AI and ML, the search for literature was particularly focused on respected sources as described in 2.1 Sources. Another challenge is that there is not many scientifically based research that observes and analyzes the trends of ML/AI adoption across industries and geographical regions. While some market research houses and IT consultancy firms do conduct surveys and interviews, there is no also consistent approach to analyze this trend. Compounding this challenge is that different research houses have different interpretations of ML viz AI and Data Science, which are lateral fields but not in the scope of this thesis.

This literature review has highlighted the need for a comprehensive overview of ML from both the technology and business perspectives as well as frameworks of adoption and implementation, specifically catered for "analog" companies as they seek to gain competitive advantage from their data.

Chapter 3 Research Methodology

This chapter describes the research methodology used in this thesis. As this thesis encompasses both socio- and technical domains, it combines both social science as well as scientific research methods to intuit and generalize the trends in ML as well as from the company surveys, in addition to reasoning about the framework for ML applied to businesses. The following Section 3.1 discusses briefly the research approach. Section 3.2 outlines the research design and Methods. Section 3.3 highlights the limitations of this overall research methodology.

3.1 Research Approach

The research takes an inductive method of reasoning as opposed to a deductive approach. Presently, there is no overarching and fundamental scientific or management theory (or set of theories) that underpins these algorithms or trends in ML application in businesses. Rather, initial observations on the literature on ML suggests that research tends to be more applied rather than fundamental theories.

In the absence of a theoretical basis and hypotheses with which to test the observations of ML trends (in literature as well as surveys), a deductive approach is not suitable. Inductive reasoning, on the other hand, allows for the synthesis of the literature review and the conclusions from the company surveys, which forms the likeliest possible explanation for such sets. Inductive reasoning in this instance help to provide a broad based explanation of the current and future trends in ML in both scientific and business domains and address the set of Research Questions RQ1 and RQ2.

In understanding the research findings and observations, the research paradigm chosen is that of an interpretative perspective because the data that is collected is a presentation of other people's construction about what they are up to, influenced by the social interactions between them (Singh & Bajpai, 2008). This approach is useful in the in Chapter 5 where ML is analyzed from the business perspective through survey of market research papers and business reports. Other paradigms such as positivism (where everyone is assumed to experience the same reality) and Critical theory (where other influences such as politics, culture influence the social reality) are not chosen for the reasons stated (Neuman, 2006).

3.2 Research Design and Methods

To answer the Research Questions posed in Section 1. 3, an Exploratory Research is conducted as the thesis does not aim to provide conclusive and final answer to the research questions. This is due in part to the incomplete body of knowledge of all ML techniques and degree of adoption by organizations, in addition to the lack of fundamentals as earlier mentioned. The alternative of a Conclusive Research is not chosen as the overall design because it is difficult to generalize to verifiable insights given the dynamic field of ML.

Secondary Data will be collected from document analysis of available literature from journals, internet, books etc. The procedure for document analysis will be to read these literature and document and describe statistics relevant to the research questions. In contrast, Primary data collection is not feasible in this context, especially in the Chapter 5, because of proprietary information that companies may chose to withhold, which could skew the conclusions.

Qualitative analysis will be conducted on the data collected. The qualitative approach is preferred as opposed to the quantitative approach because there is no academically-agreeable set of heuristics to measure, classify or evaluate the data, particularly in a broad and diverse field of ML (Japkowicz & Shah, 2011). Qualitative analysis is interpretative (parallels the interpretative paradigm) and thus allows for generalized insights to the data in order to address the research questions.

3.3 Limitations

A comprehensive survey of the literature and companies is not likely given that ML is a broad, diverse and complex field that is changing as the research is conducted. A different approach could be to organize future research into the different taxonomies of ML research or conduct the research at a later stage when the evolution of ML techniques has stabilized and companies have fully adopted ML strategies. However, postponing this research will detract from the impetus behind the key research question of this thesis – that is a framework to enable organizations to obtain value by incorporating ML strategies in their businesses.

The choice of the secondary data collection method is at best limited because companies do not articulate their strategy or share their proprietary information online or in other publications. The research can only take an interpretative approach to reason their strategies. Nevertheless, the research can be improved by conducting interviews with selected companies so that better insights into their application of strategy could be obtained. As for the literature survey on ML techniques, many of these information is available in professional journals such as Arxiv (www.arxiv.org), and therefore not a critical concern.

Lastly, the proposed framework for incorporating ML strategies is yet to be tested in a real-world scenario despite two examples provided to exemplify the applicability of the framework. Further research is anticipated for testing the viability of such frameworks in real business cases which will fine tune these guiding principles.

Chapter 4 Machine Learning – Technology Perspective

This chapter provides a broad overview of the multi-disciplinary field of Machine Learning (ML), viewed from a technology perspective that is pitched at the decision-making levels of organizations. Because the field is broad, diverse and complex, this chapter bridges the cognitive gap that many of these decision-makers and managers would face when understanding this technology capability. Many reference textbooks (such as (Bishop, 2006)) or journals (such as IEEE Transactions on Pattern Analysis and Machine Intelligence [PAMI] or Neural Information Processing Systems [NIPS]) have focused on the math and models behind ML but none so far as identified in the literature review has contextualized it succinctly for the manager. This chapter thus highlights the key concepts, taxonomy, driving factors and trends of ML to aid this understanding and facilitate decision making.

The following sections seek to address the following Research Questions that motivates this thesis:
RQ2: ML from Technology Perspective

- What is Machine Learning (ML)? What's the underlying theory?
- How can ML be classified in the context of Artificial Intelligence and Big Data
- What different techniques/methods exist today and how can we categorize them?
- What are the Key Gaps in understanding ML from the management perspective?
- What are the challenges of ML techniques and tools?
- What are the current and future trends?

The first part of this chapter covers the fundamentals of ML. Section 4.1 provides an overview of ML, its core concepts. Section 4.2 expands on the general landscape of Artificial Intelligence (AI) and Data Science and how ML fits into this landscape. Section 4.3 specifies a taxonomy that classifies ML by type, functions and form so that the various techniques can be easily navigated. This taxonomy feeds into Section 4.4 where a model of the training task is presented. Section 4.5 discusses some of the areas where ML will fail.
The second part of this chapter focusses on the applications of ML. Section 4.6 examines the key factors that drive the trends in ML. Section 4.7 surveys the broad application domains in which ML has been deployed while Section 4.8 describes the maturity and progress of ML technology. Section 4.9 concludes the chapter.

4.1 Overview of Machine Learning

This overview of ML and its techniques does not dive into the mathematical concepts and details of these ML algorithms but on the fundamental concepts that can be easily understood by the decision maker. Nevertheless, there exists comprehensive literature available to the practitioner to examine the theory, concepts, application and challenges of these algorithms. Some of these characterize the behavior and suitability of the algorithms across the various domains and datasets. Readers should refer to these literature for mathematical details on ML, for example (Bishop, 2006; Murphy, 2012) for more details.

4.1.1 Definition

In the simplest layman terms, ML is automating automation. It is the training of computer programs by teaching it to learn from examples, so that it can perform cognitive tasks better, such as pattern recognition or prediction. It is about letting the data do the work instead of hardcoding software to process the data. ML Algorithms are also termed as "learners", a term that is used interchangeable for the rest of this thesis. Specifically, in computer science literature, ML is the set of computer algorithms that performs the function of searching through a large space of candidate functions, guided by training data, to find the solution function that optimizes the performance metric (Murphy, 2012). In more formal mathematic terms, given (i) a loss function L and (ii) a sample D from some unknown distribution $\mathcal{D}$, the Machine Learning algorithm computes a solution function f that has low expected error ε over D with respect to L.

There is a fundamental difference between ML and traditional programming in terms of inputs and outputs as illustrated in Figure 4.1. In Traditional programming, the inputs are "Data" and "Programs", which is then processed by the computer to output "Data". In contrast, for ML, the computer takes as input "Data" (represented by training data and test data) and outputs a "Program". The "Program" that is output by the ML can subsequently be used in the traditional computing to process incoming data to predict the output.

As an illustration, consider a recommendation engine for an online store, highlighted in Figure 4.1. During the "Learning" stage, the computer ("Machine Learner") will take as input all of the customer data ("Training Data" as well as held-out "Test Data") in order to formulate the "Solution Function (Program)" that could accurately make recommendations. This solution function is then fed as input during the "Predicting" stage where the computer ("Recommendation Engine") processes "Live Data" and output a "Recommendation" for the online user based on his profile.

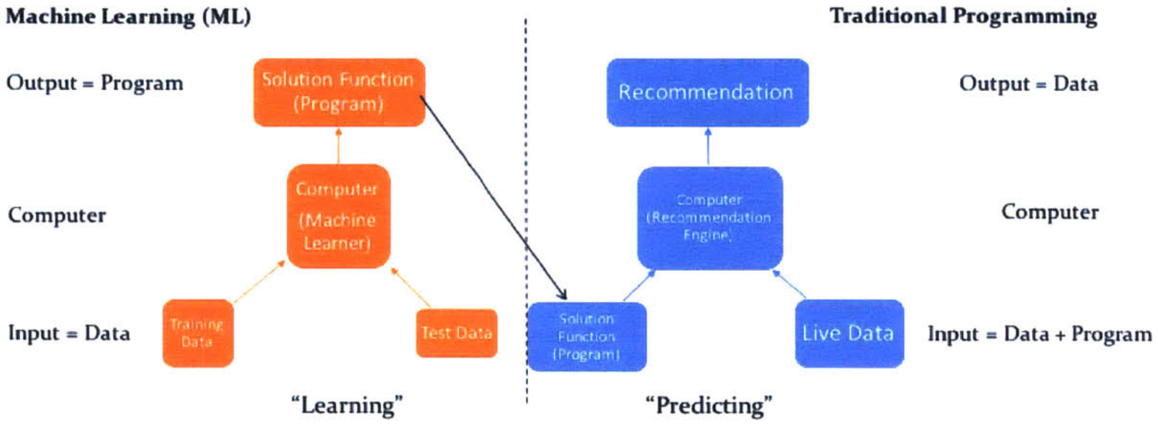

Figure 4.1 : Difference between Traditional Programming and Machine Learning

4.1.2 Concepts and Process

4.1.2.1 Generalization

The machine learning process starts with the first stage of "Learning", where the learner computes the said solution function f, by fine-tuning the parameters of the solution function by learning from examples, otherwise known as "training data". The solution function f will hence adapt and align itself to the training data, and the more training received, the more the solution function is aligned to the training data. In the next stage of the ML process "Predicting", the learner predicts the outcome with unseen data or "test data" as inputs to the trained solution function.

However, if mitigating measures are not taken during the "Learning" stage, the learner will tend to "memorize" the training data whereby the solution function completely aligns to the training data. Hence, when the learner is tested on new and unseen data during the "Predicting" stage, this is no better than random guessing as the learner regurgitates from memory rather than make predictions. This situation is analogous to memorizing exam questions before taking an exam. Accuracy of prediction suffers because the student is unable to answer questions (test data) that are not exact replica of the questions that he had memorized earlier (training data). Some of the questions could be variations of the same topic, but the student is not able to perform well because he will draw a blank or make inaccurate predictions at such unanticipated data that was not previously memorized.

Another illustration to demonstrate this is as follows: A training set of x and y values distributed in an almost 1:1 ratio is shown by the dots in Figure 4.2. The relationship between x and y is actually y = x as indicated by the red line in Figure 4.2. During data collection, random noise was introduced hence the distribution of data as such. If the learner is allowed the memorize the training data due to over-training, the solution function f is shown in blue, where all the training data is memorized. The solution function f (blue) does not accurately reflect the true relationship between x and y (red). Hence during "Predicting", if the test data is provided by x_0, the predicted output y_1 is inaccurate as compared to the correct answer of y_0 that reflect the true relationship of x and y.

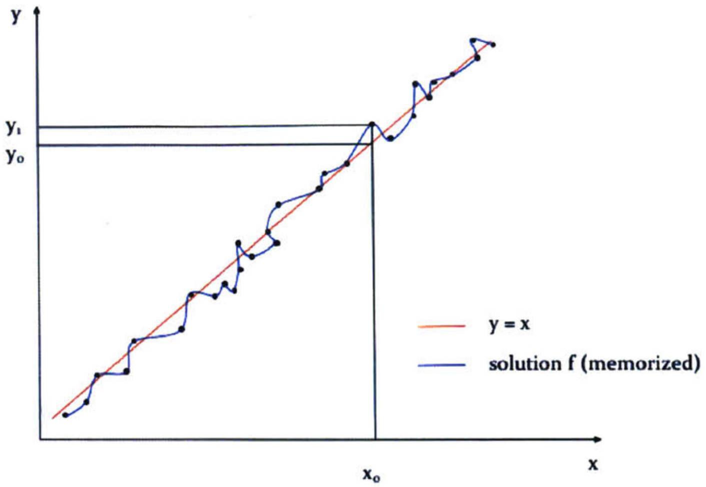

Figure 4.2: Concept of Generalization

The above analogy thus highlights the pitfalls of overtraining on training examples. Thus, the fundamental goal of the learning task is *Generalization* because for every application, it is unlikely that one would obtain exact examples during "Predicting" that was seen during "Learning".

Generalization can be achieved by holding out portions of the training data as "cross-validation" data – which is considered as a representation of test data, only that it is available during the "Learning" stage. Once the learner has learnt from the training data, it is tested on on the cross-validation data to evaluate the learner's prediction accuracy or how well the learner would perform on unseen data during the "Predicting" stage. This is done over several sets of cross-validation data so as to derive the best set of parameters of the solution function, such that it is robust and able to generalize to the test data during the "Predicting" stage. This cross-validation procedure therefore avoids over-training and the problem of memorizing highlighted above as it indicates to the trainer when to stop when the performance of cross-validation deteriorates, the cross-validation being a corollary of test data.

For reasons stated above, cross-validation is considered a compulsory procedure in the ML process. For many applications, up to 20% of the training data is used as cross-validation data so as to avoid the problems of *Generalization*. Typically, if the learner performs well on cross-validation data, it would indicate similar levels of performance on test data (if and only if the test data is in scope – see Section 4.1.2.3 below).

4.1.2.2 Induction

Another core concept is *Induction*. Based on the training data, the learning algorithm induces a function *f* that will map a new example to a corresponding prediction. By training the function, ML invokes the principle of induction that allows for the prediction based on these training data. Induction is a reasoning process that makes generalizations (the solution function *f*) from specific information (training data). Without induction, one would not be able to link the specific examples contained in the training data to the general solution function as there would not be a relationship. Induction provides the logical premise that the training data is a representation of the solution function provided by the learner, and in so could be used to infer or predict the outcome of test data.

A simple example of induction is as follows. Taking the similar example as above, consider a set of training data consisting of a set of *x* and *y* values distributed in an almost 1:1 ratio (indicated by dots in Figure 4.3). From the "Learning" stage, the learner induces a solution function of y = *x* that reflects the training data (indicated by the red line in Figure 4.3). This solution function *y* = *x* can then be used to predict y values where only *x* values are provided, during the "Predicting" stage (indicated by the blue crosses in Figure 4.3). With induction, the y values are interpolated from the y = x axis using x values, such as x_o inducing y_o. However, without induction, one cannot predict y_o values for there is no logical relationship between x_o and y_o, despite the evidence indicating such a pattern shown by the red line. The induction from the red line to the y_o value on y-axis in Figure 4.3 does not exist without the concept of induction, as indicated by the dotted line.

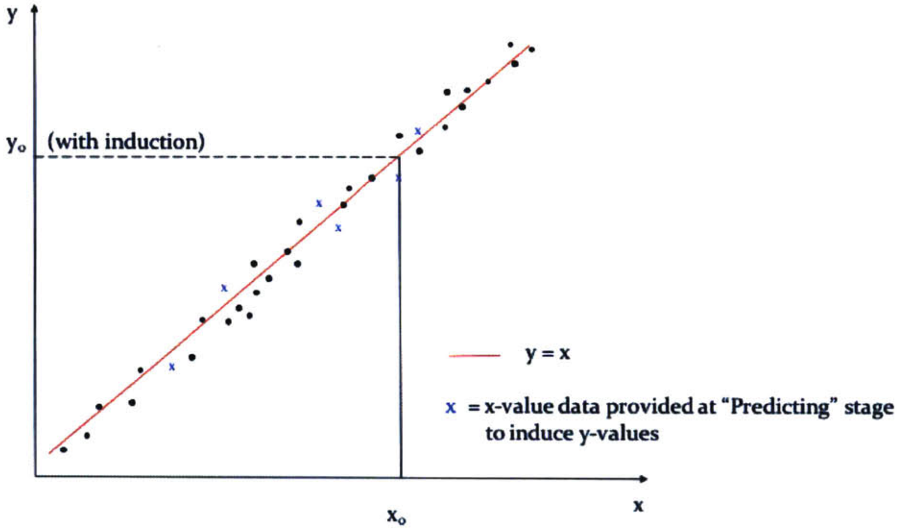

Figure 4.3: Concept of Induction

4.1.2.3 Scope

The third concept is *Scope*. The prediction accuracy of the learner depends on the test data that is used. The learner that is trained on a certain set of examples, for instance images of reptiles would be able to predict accurately only if the test data fall within the scope of reptilian images, i.e. the learner can induce the outcome of the test data based on what it learnt from the training examples. However, ML will fail to induce if images of birds were shown instead to the learner.

To illustrate this concept, the previous example is repeated in Figure 4.4. Note that the training data of *x* and *y* values range from 0 to 1 and the solution function *y* = *x* also ranging in the same limits. The concept of *Scope* indicates that it is not possible to induce accurately if test data is beyond these limits. If the test data is indicated by x_2 (beyond the limits of training data), the predicted outcome y_2 is wrong because the solution function can take on any form beyond the training data – it can be a quadratic function beyond $x = 1$ as indicated by the portion in blue in Figure 4.4, instead of the assumed $y = x$ function indicated by portion in dotted red.

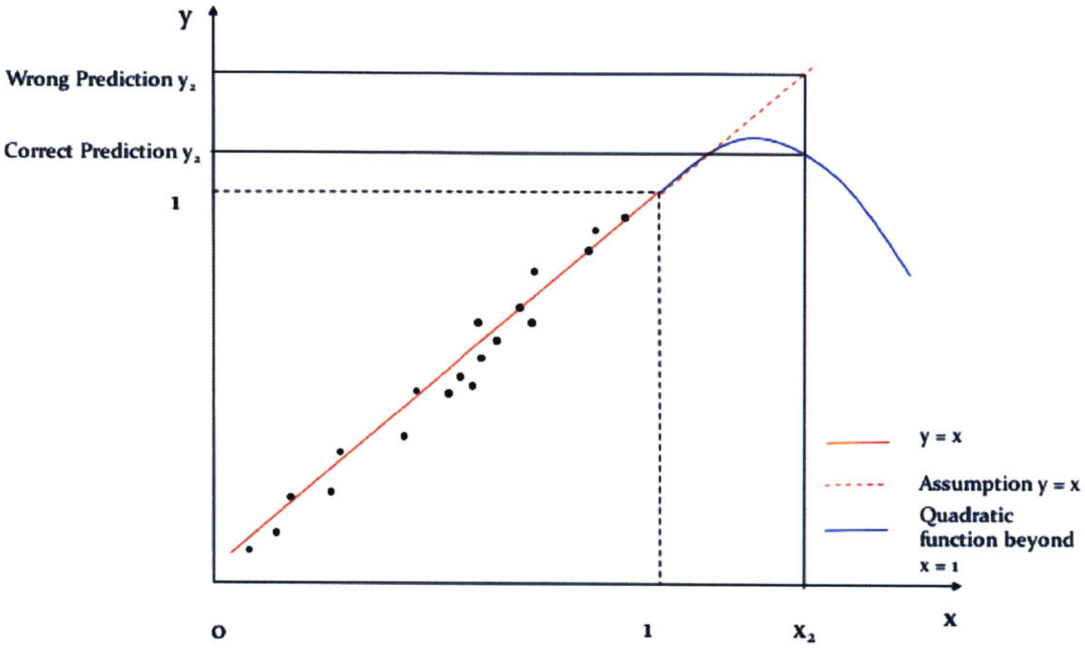

Figure 4.4: Concept of Scope

4.1.2.4 Evaluation

The fourth and final concept is *Evaluation*. While the goal of inductive machine learning is to learn from the training data and induce the function *f*, which will be used to predict test data, the only way to find out if it has indeed learnt is to evaluate its performance of prediction on the test data. The machine learning algorithm is deemed successful if the performance on test data (that is within scope) is high. Evaluation is therefore an ongoing process of adapting the learner to newer and unseen test data.

To illustrate this concept, Figure 4.5a shows a training data of *x* and y values (indicated by dots). The "Learning" process first considers the training and cross-validation data and produces the solution function y = *x* (indicated by the red line). During "Predicting", test data (indicated by circles) is used to ascertain the performance of the learner based on solution function y = *x*. During evaluation, where the prediction is compared against actual y values, the error is small indicating that the solution function is fairly accurate. However, over time, there is drift in New Data as indicated by blue circles in Figure 4.5b. The evaluation error (between these New Data and the predicted values along the red line) has increased. This indicates that the learner is unable to cater to the new variations in test data if it remained with the solution function in red and the learner therefore needs to revise to the new solution function, generating the blue line indicated in Figure 4.5b. Constant evaluation thus allows for the learner to adapt to changes in the test data and revisit the drawing board to develop more robust solution functions.

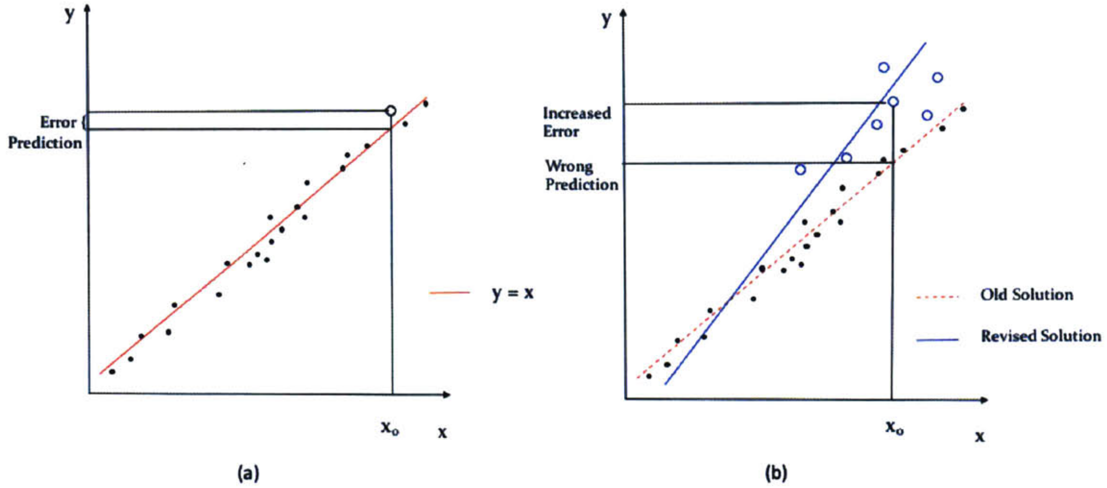

Figure 4.5: Concept of Evaluation

4.2 Landscape

While there is no standard definition of Artificial Intelligence (AI), Russell and Norvig defines AI as the "study of agents that receive percepts from the environment and perform actions", where each of these agents implements a function that maps perceptions or observations to actions. Learning is a part of AI, where it role is to "extend(sic) the reach of the designer into unknown environments."[1] Within this context, Machine Learning (ML) is a sub-discipline of AI that focusses on learning from experience (in this case data) where it seeks to improve different aspects of performance in an automatic fashion. Russell explains further that "an (AI) agent is learning if it improves its performance on future tasks after making observations about the world." (Russell & Norvig, 2010).

At a conceptual level, ML optimizes its performance automatically because computer algorithms ("agents") process data that forms its "experience". A more operational and criterion-based perspective puts forth ML as "(the) computer program is said to learn from experience E with respect to some class of tasks *T* and performance measure *P* if its performance at tasks in *T*, as measured by *P*, improves with experience *E*" (Mitchell, 1997). This operational view distinguishes the ML-computer programs from that of a thinking entity where it is more philosophically challenging to define (How do we define a machine that can think like a human being?). This approach was proposed by Alan Turing in 1950, that serves as an empirical test to reason about the efficacy of these thinking machines (Turing, 1950).

That ML hinges on data to achieve its goals of improving performance is often the reason why ML is often confused with "Data Mining" or "Data Analytics" – sub-fields of Data Science where they similarly depend on the availability of large amounts of data.

[1] There are other sub-fields of AI that studies into machine reasoning and perception, that this study will not delve into.

Dhar defines Data Science as "the study of the generalizable extraction of knowledge from data." and explains that ML is a therefore a method or technique used in the field of Data Science to study and learn patterns in data and make predictions and insights (Dhar, 2012). On the other hand, Big Data Analytics or Data-mining seeks to find relationship in existing data. The key difference between data analytics and ML is that ML learns from previous data and makes predictions from future data whereas Data Analytics uncovers relationships within a static set of data. These techniques of Data Dnalytics are predicated on statistics and take an inward looking approach to analyze data, as opposed to ML which is more optimistic in its goals.

In terms of classification between the two fields of AI and Data Science, ML is situated more in the Data Science field because of two reasons. First, ML is based off data that is both structured and unstructured and which that exists in databases as well as the Internet. Whereas in the context of AI, ML is an agent learns from its experience in the environment which could come from observations via sensors or other forms that do not necessarily need to be tied to databases or networks. ML is thus more closely tied to Data Science based on current trends; yet there could be advances that place ML more closely with AI because of the increasing variety of data that is being collected. Deep Learning, which is a sub-topic in ML, is an example. Second, AI is a more loosely defined field of study which encompasses more areas than just learning. It encapsulates sub-fields like perception, reasoning and understanding and learning is one approach in which machines can "think". Data Science is more well predicated on data and the application of many other disciplines such as Statistics, Probability and Computational Learning to explore and examine data. ML fits better in the field of Data Science than AI.

Figure 4.6 maps the classification between Data Science, ML and AI based on the above discussion. It is observed that Big Data is a description of large volume, velocity and variety of data that has become widely available with advent of internet technologies. It is an enabler and accelerator towards the development of Data Science and ML, but not necessarily AI. Similarly, logic and predicates are the elements of AI but not necessary Data Science (Decision Trees in ML is depends on logic and predicates).

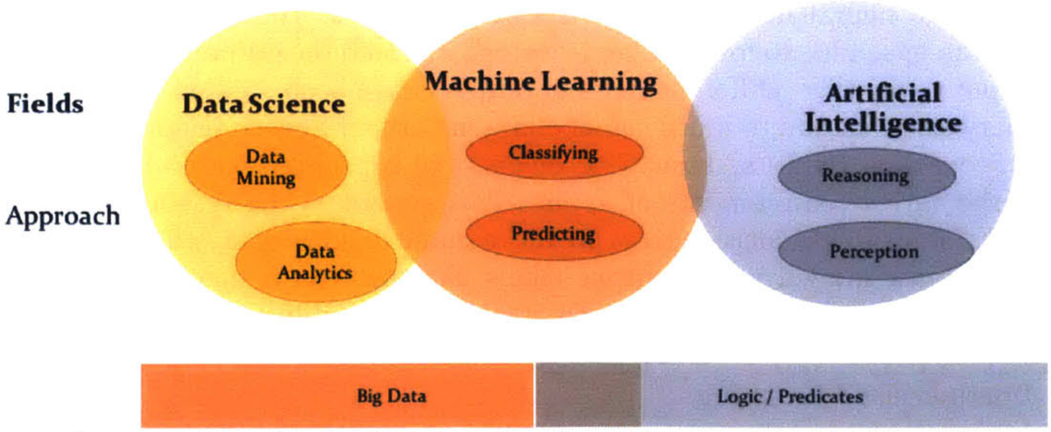

Figure 4.6 : Mapping Machine Learning, Data Science and Artificial Intelligence

4.3 Classification of Machine Learning Systems

Machine Learning systems can be classified by several types: Learning Fields, Learning Model, Output, and Components. By breaking the ML system into these types, one can think about the different tradeoffs that need to be made in order to solve the learning problem on hand, based on the nature of data to be processed or "learnt" from.

4.3.1 Classification: Learning Fields

ML systems can be classified broadly by learning fields into four different categories depending on the nature of the learning: Supervised Learning, Unsupervised learning, Reinforcement Learning (Jordan & Mitchell, 2015) and Deep Learning. The reason why ML is classified by learning fields is because each of these fields require a different type of training data as well as set constraints on the type of prediction that can be obtained.

4.3.1.1 Supervised Learning

Majority of ML systems are based on Supervised Learning. Given an input variable x and output variable y, the goal of the ML systems is to learn the solution function f that maps the input to the output well during the "Learning" stage. At the next stage of "Predicting", the solution function accurately predicts the correct output y1 given a new input x1, so long as generalization and scope concepts are adhered to – see concepts discussed in Section 4.1.2.

This process is called supervised learning because the learning process of obtaining the solution function from training data is akin to that of a teacher supervising the learning process. The correct answers are encapsulated by the training data and the learner iteratively makes predictions on the training data and is corrected by the ML system. When the learner has achieved an acceptable level of performance as evaluated by certain criteria, the "Learning" stage stops.

Many different forms of solution function *f* exist, including decision trees, support vector machines, neural networks and Bayesian classifiers. These functions can be combined in an ensemble to form "super learners", in which the output of one function is fed as input to another. Different types of outputs have been studied, ranging from simple binary classification to multi-set labelling in Parts-of-Speech tagging in Natural Language Processing (NLP) systems. The diversity of supervised learning function is driven by the varied requirements of applications as well as the type of data to be represented. For instance, decision trees are more suited to noisy data and for predicting the output against known set of output labels. Support vector machines are more effective for high dimensional spaces to classify unknown number of output labels.

4.3.1.2 Unsupervised Learning

Unsupervised learning is where there is only input data x with no corresponding output variables. The goal for unsupervised learning is to model the underlying structure or distribution in the data. It involves the analysis of input data under the assumption of structural properties of the data (e.g. the data has an inherent probabilistic density function).

Unsupervised learning is so called because unlike supervised learning, there is no correct answers (no output variable y in the "Learning" stage) and there is no teacher to correct the solution function f. In the absence of a teacher, the objective of unsupervised learning is to optimize a criterion function. The criterion function often takes on the form of statistical methods such as maximum likelihood which approximates the data based on prior assumptions of its distribution. Sampling techniques such as Monte Carlo Markov Chain (MCMC) are then used to optimize the criterion function. In other scenarios, unsupervised ML algorithms such as k-means clustering or Principal Component Analysis (PCA) can organize data by similarity by the principal of dimensionality reduction so that the underlying structure of data can be revealed, without resorting to statistical methods.

Unsupervised learning is less developed than Supervised learning because of the complexities arising from designing algorithms that manipulate data in high dimensional space, but is an active research frontier of ML because the reality is that a large proportion of input data that exists in databases is not labelled and therefore does not have a known result.

4.3.1.3 Reinforcement Learning

Reinforcement learning is the third category of ML. Information available in the training data is between supervised and unsupervised learning. Instead of the training examples indicating an output in supervised learning for every input signal, the output in reinforcement learning is the action taken by the reinforcement learning algorithm across a sequence of inputs. The output action is then evaluated if the action is correct or not. As such, there is no requirement for the algorithm to track which output actions link to which input sequence. The learning tasks of the Reinforcement learning algorithm is to learn a control strategy for an agent acting in an unknown dynamic environment where the algorithm is trained to choose actions for any given input state. Reinforcement learning combines aspects of control theory (policy iteration, variance reductions, Markov decision processes) as well as neuroscience (reward stimulus, feedback mechanisms) to supervised learning.

4.3.1.4 Deep Learning

Deep Learning (DL) is not based on the above approach where the representation of the data at the input or output layer defines its category. DL is a recent ML paradigm that comprise of variations of Neural Networks that have a deep structure, that is using Deep Neural Networks (DNN) that consist of more than one layer. These deeper forms of neural networks are able to represent data in a more abstract form than were possible with other ML algorithms, which had traditionally focused on finding the optimal parameters for the identified function. DNN can be utilized to improve the efficiency of traditional ML models by automatically extracting new features in comparison to ML algorithms where human effort, time and domain expertise are required to feature engineer the parameters. The emergence of the DL paradigm is due to recent progresses (since 2012) particularly in image, speech recognition and natural language processing tasks where the accuracy of such DL algorithms had surpassed that of traditional ML techniques. In this regard, the author of this thesis considers DL to represent an evolutionary shift in ML so much that it it warrants its own categorization apart from the previous three provided in (Jordan & Mitchell, 2015).

4.3.2 Classification: Output

Learners can also be classified by taking a functional perspective, by considering the emergent functions or outputs of a ML system (Russell & Norvig, 2010). The rationale for classifying ML by its output or type of object that it is trying to predict is because of the degree of measuring error of each learner. Since the goal is to build a ML system that can make "good predictions", different approaches will mean different interpretations of "good".

i) **Classification**: Inputs are divided into two or more classes, and the learner must produce a model that assigns unseen inputs from the test dataset to one or more (multi-label classification) of these classes. This can be done in a supervised way using labelled datasets. The output is discrete.

ii) **Regression**: This is similar to classification where outputs are continuous rather than discrete. Regression learners typically are also supervised learners in that they are trained on labelled datasets.

iii) **Clustering**: A set of inputs is divided into groups. Unlike in classification, the groups are not known beforehand, making this typically an unsupervised task.

iv) **Estimation**: Deriving the distribution of inputs in some space, e.g. Gaussian distribution or sparse distribution in order to obtain insights into the data.

v) **Dimensionality reduction**: Mapping inputs into a lower-dimensional space therefore simplifying the representation of these data so that insights can be obtained at these lower-dimensional spaces.

4.3.3 Classification: Components

Learners can also be classified by taking a "Form" perspective, by categorizing them as a combination of the following three components: Representation, Evaluation and Optimization (Domingos, 2012). This is illustrated in Table 4.1 where common examples of these three components are described below. This thesis adapts Domingos's work by incorporating non-classifiers such as in unsupervised learning in Table 4.1.

i) **Representation**: Learners are represented in some formal language or model that the computer can handle and process. For example, k-Nearest Neighbors (k-NN) and k-Means Clustering are Instance-based learners where data can be classified into instances of classes. Hyper-plane learners such as

Logistic Regression or Naïve Bayes optimizes a linear function based on the data. Decision trees iteratively tests the data at each node and makes class predictions at the leaves. Rule-based learners such as Inductive Logic Programming uses logic programming as a uniform representation for input examples, background knowledge, and hypotheses. Neural Networks are suited for modelling complex relationships between the inputs and outputs, finding patterns in data, or deriving the statistics in the input dataset. Graphical models represent a set of random variables and their conditional independencies between one another using a probabilistic graphical approach.

ii) **Evaluation**: An external or internal performance metric to compare the efficacy of candidate algorithms. For instance, the external evaluation metric for machine translation is typically Word Error Rate (WER) or Perplexity. This is different from the internal view of the evaluation/objective function which is used to evaluate the performance of the algorithm, that the algorithm will seek to maximize. Examples of such internal objective functions are maximum likelihood of prediction, mean square error between prediction and labelled outputs.

iii) **Optimization**: Choice of optimization techniques appropriate to the algorithm, which determines the tractability and efficiency of the learner. The choice of optimizers, such as conjugate gradient descent or BatchNorm stochastic gradient descent, will determine the time needed to train the model. More importantly, the proper selection of the optimizer will avoid the problems of local minima which could result in poor prediction and performance.

The method of classification in Table 4.1 allows for the mixing and matching of components that make up these learners and allows the practitioner to sort through the variety of learners that are available. For instance, a ML System can comprise K-nearest neighbors to represent the unsupervised ML system, using likelihood as a criterion function and quadratic programming to optimize the criterion function.

Table 4.1 : Common Examples of Components of ML System. Adapted from (Domingos, 2012)

Representation	Evaluation (Internal and External)	Optimization
Instances • K-nearest Neighbors • Support Vector Machines	• Accuracy / Error Rate • Precision and Recall • Perplexity	**Combinatorial** • Greedy Search • Beam Search • Branch and Bound
Hyper-planes • Naïve Bayes • Logistic Regression	Squared Error	**Continuous (Unconstrained)** • Gradient Descent • Conjugate Gradient • Quasi-Newton methods
Decision Trees	• Likelihood • Posterior Probability	**Continuous (Constrained)** • Linear Programming • Quadratic Programming
Set of Rules • Propositional Rules • Logic Programs	• Information Gain • KL-Divergence	
Neural Networks • Shallow Auto-encoders • Deep Neural Networks	• Cost/Utility • Margin	
Graphical Models • Bayesian Networks • Condition Random Fields • Markov Chains		

4.3.4 Summary

The discussion in Section 4.3 seeks to classify the variety of learning algorithms according to their Learning Fields (Section 4.3.1), Output (Section 4.3.2) and Components (Section 4.3.3). This classification therefore provides a high level intuition which enables the practitioner to consider the appropriate learners based on the application domain as well as the datasets. With this taxonomy of classification, the next section presents a model of the Machine Learning Task which considers the choices presented in this taxonomy.

4.4 A Model of Machine Learning Task

To specify a learning task, one needs a precise *model* that describes what is to be learned and how it is done, and what measures are to be used in analyzing and comparing the performance of different solutions (van Leeuwen, 2004). This model is useful for practitioners in considering how to approach the problem of the ML task and conceptualizing the solution, based on his domain knowledge and experience.

In Table 4.2, the various elements in Leeuwen's model of the ML tasks are highlighted in the left column and demonstrated with an example of a handwriting recognition task that captures handwritten receipts into a bill management system. The Learner specifies the ML system that is doing the learning and where this ML system is applied. The Domain specifies what is being learnt by the ML system, i.e. the solution function as stated in previous Sections. The Goal highlights the overall objective of the learning system. Representation indicates the form that the data would take in the ML system. Information source is another term for the training data. Prior knowledge, as described, refers to some a priori understanding of the information that would aid in the learning process. Algorithmic technology refers to the specific ML algorithm used (see Section 4.3 for classification of these algorithms/learners). Training Scenario describes how the ML system is to be trained with the training data, which is bounded by the Success Criteria/Loss function that determines its performance limits. Lastly, Performance Metrics measures the extent to which the learning is successful.

Table 4.2 : Elements of Machine Learning Task – Handwriting Recognition. Adapted from (van Leeuwen, 2004)

Elements	Example: Handwriting Recognition Task
Learner	Who is doing the learning and where it is applied. In the case of the handwriting task, the software that runs between the image capture and the bill management system, where the software incorporates the ML algorithm as well as its hypothesis learnt during training.
Domain	Specifies what is being learnt, for instance the function that predicts the symbols from handwriting input.
Goal	Specifies the goal of the learning task. For instance, the goal of the handwriting recognition software is to process handwritten receipts into a bill management system.
Representation	How the objects to be learnt is represented. For the handwriting recognition example, input is represented by the images of the handwritten portion of the receipt and the output is represented by the symbol that corresponds to the handwriting.
Information Source	The training data used for training the learner. The information source can take different forms such as a set of (x,y) labelled pairs.
Prior Knowledge	What is known in advance about the domain, that will help to enhance the learning and avoid some of the reasons for failure (see Section 4.5). For the example, the symbols that represent the data are assumed to be numerals.
Algorithmic Technology	The specific ML algorithm used. (see Section 4.3 for classification of these learners). In order cater to the variety of handwriting types, a deep neural network is used in the learning task.
Training Scenario	The description of the learning process. Options include supervised, unsupervised, reinforcement learning or deep learning (see Section 4.3), as well as how the training data is presented during learning, either online (streaming) or offline (batch by batch). The handwriting task suggests a supervised offline learning scenario where labelled training sets are available for training the model.
Success Criteria/Loss Function	The criteria for determining when the learning is completed or has otherwise converged sufficiently. A loss function suitable for the handwriting example would be Mean Squared error between predicted and actual output labels.
Performance Metric	The amount of time, computational power and accuracy required for the learning task. There is often a trade-off between these metrics. For the example using a deep neural network, large training sets are required but the computational time scales exponentially due to the complexity of the network.

4.5 Drivers of Progress in ML

While the previous section examines the technical aspects of Machine Learning, the next sections switches gears to discuss the applications of ML and the recent progress in ML. The phrase "Machine Learning" first appeared in computer science research in the 1950s, pioneered by Arthur Samuel in the first computer learner (Moser, 1990). While initial growth was slow before the 1990s due to insufficient computing power to power the learning process, the usage of ML has been steadily gained traction since the mid-90s when the requisite computational power had become readily available, opening up the way for novel applications such as prediction and fraud detection. Besides advances in computational power, the availability of diverse and large amounts of data that is captured by modern databases, allows for these ML systems to generalize to many different concepts, within feasible amounts of time. The data crunching ability of these ML systems had lead to the exponential growth in ML in the past few years as users began to apply ML to various domains where the ability of ML to "Learn" and "Predict" outshine their human operators. The following discusses the impact of these factors and examines how these factors can drive future growth.

4.5.1 Increase in Computing Power

Data storage prices have decreased and data processing capacities have grown tremendously over the past decades. The smartphone has more storage and compute power than a mainframe in the 1980s. Such advances in computing power allows ML algorithms to learn faster and effectively by processing more data within a realistic timeframe. Data can also be stored cheaply on the cloud or on premise. In tandem, ML algorithms have also become more complex to leverage on the availability of compute power. Today, machine learning models might use millions (or billions for complex neural networks) of parameters to find patterns in data, in comparison to traditional analytical methods that are based on statistics.

The increasing complexity of these algorithms has in turn driven the requirements for more computing power in a positive feedback loop. Several novel approaches have sprung up in recent years. Parallel computing architectures have been modified to meet the processing demands of these learners. Graphical Processing Units (GPUs) that have previously been used for computer visualizations and gaming, have been utilized to train these algorithms because of their efficient parallel architectures. Distributed or Cluster computing, aided in part by cloud-based technologies, have also been innovated to process increasing amounts of data.

That said, there are constraints as to the amount of computing power available to train these learners. There are limits to processor sizes as dictated by Moore's law, thus placing a cap on the computational power that can be obtained from a single processor (and also on smaller companies who could not invest in larger infrastructure). Cluster computing in the Cloud offered by companies such as Amazon Web Services, could overcome this foreseeable future but it too has limits in computing complex neural networks because of the need to parallelize computation at scale. Seeking better Performance-to-watt metrics have also become more important to companies seeking to reduce operating costs while deploying these ML algorithms in production environments. More efficient computation could also enable the personalization of these learners in the hands of consumers through their personal devices, where Learners can learn from the individual's pattern of life, rather than having a company manage clients' private stream of data, as indicated by recent Google patents (Aradhye, Hua, & Lin, 2013).

4.5.2 Wide Availability of Data

The growth of ML is also driven by the wide availability of data through big data technologies as well as Internet of Things (IoT) devices. Learners are able to be more accurate than statistical methods because they are able to learn from more experiences (training data) and therefore able to generalize better to test datasets as the Learner would have "experience" a similar situation before during training. For example, a self-driving car can collect nearly 1GB of data per second and this granularity allows the learner to find patterns in order to make reliable decisions.

A wider spectrum of data from images, text, video or even unstructured data that can be obtained via numerous data collection methods have also driven ML developments in various application fields. ML techniques such as those in deep neural networks have been used in translation of speech that offers competitive performance, while reducing the human effort needed to feature-engineer the model as compared to statistical methods.

Despite the literature stating that these technologies have made large amounts of data available (for example, Facebook processes up to 600TB of data per day (Pamela Vagata & Kevin Wilfong, 2014)), there are several constraints affecting the usefulness of these data in driving better business insights. First, a large proportion of the data could be useful as training data to feed into ML algorithms do not currently exist in fixed structured form that can be processed. Much of these data are unstructured, poorly defined, missing and noisy. Companies have to devote resources to collect, filter and manage the data in a clean format that can be used to train the ML algorithms.

Second, much of the data is unlabeled, i.e. they are observations of environments without indications of which classes they belong to. Much of the progress in ML has been in supervised learning where the algorithm effectiveness benefitted from training on large datasets with millions of labelled training examples. For example, some of the significant progress made by Google in recognizing, classifying and captioning images were enabled by using Amazon Mechanical Turk services to hand label the training dataset (Vinyals, Toshev, Bengio, & Erhan, 2015). However, such an approach does not scale for the larger class of data that is unlabeled. In the quest to design learners that can learn like human beings, where we are able to understand and conclude causes from observations of our environment and adjust our internal model of the world, the next steps in ML is in unsupervised learning. In unsupervised learning, the aim is to design learners that can generate representations of data given observations of the world, without being explicitly trained on labelled output.

4.5.3 Summary

The discussion in Section 4.5 highlights some of the key trends in Machine Learning as well as some of the active research areas. There seems to be no upper bound in the advancement of ML and the speed of experimentation often surpasses the theoretical explanations. As such, many in the industry often regard ML as a combination of "art and science" or "black box", where it becomes difficult to interpret the inner workings of the ML system (Lipton, 2016). The next section describes some of the key application domains of ML, highlighting certain trends that organizations can readily adapt to capture value from their data.

4.6 Application of ML

4.6.1 Appropriate Applications of Machine Learning

ML is suitable for tasks that are too complex to be traditionally programmed. These tasks include those that are performed routinely by humans but unable to be elaborated such that a well defined program can be created. For example, in a machine translation system, while robust statistical methods have been previously developed, translation accuracy frequently suffers because such statistical methods are unable to capture all possible variations in sentence structure as well as colloquial words. ML allows for the system to learn from vast amounts of training data (akin to learning from experience) so that it is able to predict the translation without having to hard-code all possible variations as with the traditional statistical methods. All that is required is a good set of learning model, diverse set of data and sufficient computational power.

Another set of tasks that is too complex are those that involve large amounts of data beyond a human cognitive capability to process, such as weather data, genomics data and medical data from journals. The information contained within these large data sets are often too complex for humans to make sense of, but ML are able to process with ease because of its ability to detect patterns.

ML is also suitable for tasks that require adaptability. Traditional programs are rigid. Once it has been coded and installed, it stays unchanged. However, tasks can change over the course of time and also highly dependent on the environment. ML provides a solution to these requirements because they are able to learn and adapt to their input data. For example, ML is widely used in speech recognition software and its accuracy has improved even after "learning" from noisy input signals.

4.6.2 Application Domains

Although the field of ML is relatively young and developing, ML has found its way into our daily lives through applications such as "Google Now" that scans through your email inbox and other personal data to remind you of upcoming events or occasions, much like a personal assistant. Other examples include "Siri" where it could process natural language and respond in a human-like manner. Broadly, ML can be applied in any areas where patterns in the data can be discerned and used to create value. Patterns in the data can then be used to achieve the goals of the learning task (see Section 4.4 for the definition of "Goal" using Leeuwen's methodology), which are classified below as application domains. For each of these domains, some seminal examples are provided as well as recommendations on how organizations can deploy in these domains. For instance, a relatively quick win for medium size organizations is to consider classification of their data (Section 4.6.2.1) so as to derive deep insights into causality of certain phenomenon like decreasing sales.

4.6.2.1 *Classification*

The patterns in the training data can be used to classify these data into groups, which are subsequently used to derive value. A typical scenario involves a company's customer database which consists of their profile, purchase history and browsing patterns. The ML system processes the customer database and classifies groups of customers which have similar purchase profile, uncovering patterns which are useful for targeted marketing. The ML system is therefore able to predict new products specifically targeted to these groups of customers with a high confidence that they will buy, revealing more cross-sell opportunities.

Classification is a major goal of automatic recommender systems. A widely known example is "Netflix" where its ML systems process the users' browsing history, classify the user into groups by pattern identification and recommend movies based on customer's with similar profiles. This resulted in a high rate of customer satisfaction because the recommended movies did indeed match their interests. In doing so, Netflix is able to create novel value for their customers instead of a static webpage like other competitors (Raphael, 2015).

Similarly, an organization can employ ML to classify its data and uncover patterns hitherto undiscovered, likely using unsupervised learning as a Learning type. Insights gained from such "Learning" process would enable bespoke marketing efforts and improve topline performance.

4.6.2.2 *Prediction/Forecasting*

The input data can be learnt by the ML system to make predictions of future trends or diagnosis, typically using supervised learning as a Learning type. Prediction depends on the diversity of the training data, ability for the ML system to generalize and whether the prediction is in "scope" - concepts that were described in Section 4.1.2. Typical applications of predications are when the ML system identifies patterns during the "Learning" stage and using these attributes to predict/forecast future behavior based on future data where the outcome is unknown.

Such prediction is widely practiced in trading firms where automatic ML systems have taken over the role of traders in its ability to forecast stock prices. Companies such as Two Sigma Investments and Hudson River Trading are consistently successful in applying prediction to their automated trading strategies, generating very high returns for their clients and themselves (Metz, 2016). Another practice area of prediction is in weather service where large amounts of historical weather data is used to predict weather conditions to higher level of accuracy than with traditional forecasting methods (Linn, 2015)

An organization can readily employ ML systems to make predictions based on its data in a variety of ways, such as preventative maintenance of their valuable assets by anticipating likely breakdowns and resolving them before it actually happens. Another area where predictions matter is when trends are spotted by learning from customer data, such as in fashion retail where investments in clothing lines are informed by the prediction.

4.6.2.3 Anomaly Detection

Patterns that are uncovered by ML systems are also used to check against outlier data by testing against the learnt concepts from the "Learning" stage. In many financial institutions, ML systems have been used to ascertain the baseline patterns of their many customers and transactions or what is considered "normal". Any credit card transactions that fall outside of these learned patterns will be immediately flagged for further investigation. A large sampling of training data as well as robust ML system are required to obtain an accurate representation of the baseline pattern, keeping in mind the concepts highlighted in Section 4.1.2. That said, anomaly detection requires more than just a good ML system. Part of discovering what is normal or baseline (which are also subject to changes) depends on human insight. Hence the application of anomaly detection requires operators to be familiar with the ML model and decide what makes sense in the particular situation. (Dunning & Friedman, 2014)

Large financial organizations such as AMEX and Paypal have deployed ML systems to track fraudulent transactions. PayPal utilizes Deep Learning to detect merchant fraud (individuals committing fraud spend money at stores), whereas AMEX uses ML systems trained on a variety of data sources including card membership information, spending details, and merchant information, in order to stop fraudulent transactions before substantial loss is incurred while allowing normal business transactions to proceed in a timely manner. (H20, n.d.) (Friedman, 2015).

While this application domain is somewhat limited, certain types of organizations can benefit from implementing analogy detection using ML, particularly those where the veracity of data is critical to their reputation. Such organizations can extend to medical institutions, government bodies.

4.6.2.4 Understanding

Patterns that are uncovered by ML systems can also be used to discover deeper insights into data hitherto unknown, by revealing the underlying concepts behind different layers of abstraction. While it is not humanly possible to read through copious amounts of newspaper reports on a particular political issue and come up with deep insights, ML systems enable understanding of concepts almost trivially because of their ability to processing large amounts of data quickly and infer patterns computationally. As an application domain, Understanding is seeing widespread deployment in sentiment analysis of product reviews and customer feedback, where companies can take business decisions according to the ML system analysis of users' opinions about their products. Similarly, using sentiment analysis, election candidates can figure out the aggregated people's overall opinion in political debates. Such depends on the ML system conducting Natural Language Processing (NLP) to parse the English text and understand the intent behind the spoken words. (Medhat, Hassan, & Korashy, 2014)

Understanding (sentiment analysis) as an application domain is useful for organizations who want to understand the feedback from their customers in a more insightful manner, without resorting to human operators to parse through copious amounts of information. This has widespread traction in customer-focus organizations where business decisions are predicated on customer's needs and wants.

4.7 Areas where ML will fail

Although ML has been applied in many domains as highlighted in Section 4.6, ML is not a silver bullet that the practitioner can apply to work on his data like a "black box". Besides the problem of unlabeled, unsupervised learning as stated above, below are several common reasons why a ML algorithm might fail on some learning task.

First, the training data could be noisy. There could be noise in the training data both at the feature level (the type of data that is observed) as well as the label level (the classification of the output). For instance, some of the training data used in sentiment analysis of Twitter posts could be noisy because these posts do not really reflect the sentiment of the author, or that the labelling of the posts was wrongly estimated based on narrow set of criteria. Furthermore, the human judgment required to label the output contributes another source of noise at the label level. If there is too much noise in the data, the prediction will be inaccurate and the learning task will fail.

Second, the feature level may be incomplete for the learning task at hand. For example, even with all the medical information that a learner has on hand, such as X-ray, genomic data, personal medical history, the learner may not be able to judge if the patient is likely to have cancer or not. This reason is also known as the problem of causality. The learning task will fail because the prediction is inaccurately premised on the wrong set of features.

Third, the learning task might fail because of the inductive bias of the ML algorithm is too far from the concept that is being learnt (Daume, 2015). This is the problem of "scope" as earlier discussed under the concepts of ML. For example, if the intent of the learning tasks is to classify multiple classes of images, using an ML algorithm such as a Bayesian classifier that does not leverage on the context of data contained within the images will limit the learning performance or even give erroneous predictions. This particular learning algorithm is therefore unable to cope with the data. A prime example of this in 2015 was when Google's photo recognition engine started tagging images of black people as gorillas because the data was beyond the scope of the algorithm (Grush, 2015).

Fourth, and closely related to the problem of inductive bias is the "No Free Lunch Theorem" (Wolpert & Macready, 1997). It states that no learner can succeed on all learnable tasks – every learner has tasks on which it fails while other learners succeed. Therefore, for the learning task to be solve-able, an a priori knowledge of the distribution D of the data is required. For instance, one can make an assumption of a Gaussian distribution over a dataset that is obtained from sampling randomly student grades. The incorporation of prior knowledge is therefore crucial for the success of the ML algorithm, as the stronger the prior knowledge the easier the learner learns from further examples.

Fifth, and one of the biggest challenge in ML is coined "curse of dimensionality" by Bellman (Bellman, 2003). Given a fixed number of training examples, the predictive power of the ML algorithm reduces as the dimensionality or number of features increases. The higher dimension one algorithm goes, the smaller the space the training examples occupy, and this prevents generalization. This implies that ML algorithms will generally break down if one considers a high number of features or dimensions. Fortunately, in many applications, the training examples are not spread out uniformly in the hyperspace as defined by the dimensions. This "blessing of non-uniformity" allows for learners to utilize dimension-reduction methods or operate at a lower dimension in order to generalize correctly.

Lastly, the learning task may fail or perform badly because of underfitting or overfitting. Underfitting is when the ML algorithm fails to exploit the features available in the training set and therefore produces a solution that approximates without providing much value to the prediction. Overfitting is when the ML algorithm pays too much attention to the data and memorizes it, to the extent that generalization cannot occur because it had memorized the data. A model generated by the learner that is expected to do well on future data is one that is neither overfitted nor underfitted.

Of all the six reasons stated, the first two reasons are inherent properties of the data and therefore cannot be resolved. Therefore, it is critical to preprocess the data by selecting the right features as well as cleaning the data prior to the learning task. The latter four reasons for failure can be somewhat resolved by adjusting the learning tasks, but some do remain intractable in certain situations. Inductive bias can be resolved by utilizing a different ML algorithm that exploits the features of the data, while the "no free lunch" theorem can be overcome by understanding the data and its distribution before performing the learning task. The "curse of dimensionality" can be removed by operating at a lower dimensional space, either by removing trivial features or by using dimension-reduction methods. The last reason can be can be solved by introducing several "algorithmic tricks" that has been proven to prevent underfitting or overfitting.

In summary, Machine Learning is not an automatic process. It requires experience and domain knowledge on the part of the practitioner to understand the data, determine the best selection of algorithms and fine-tune its parameters to achieve better prediction. As such ML is considered both an "art" as well as "science".

4.8 Technology Progress of Machine Learning

This section highlights the maturity and progress of ML so that decision makers are more aware of the potential opportunities and risks that may arise and strategize ahead. Several planning tools enable this and these tools will be applied towards the framework in Chapter 6, particularly in conducting effective technology risk assessments.

4.8.1 Technology Readiness Levels

One useful heuristic that enables decision makers to understand the maturity of ML in context of similar technologies is the Technology Readiness Levels (TRL) which was proposed by NASA in 2007 (NASA, 2007). The TRL provides an indication of the risk involved in adopting these technologies as well as the investment costs associated with each level. Referring to Figure 4.7, at the lowest TRL Level 1, the technology is deemed to be proof of concept where principles are demonstrated. Going up the TRL levels, feasible models and prototypes are created and tested, up to TRL Level 9 where actual systems are tested and qualified for deployment in production environment.

The various application types of ML systems (in Section 4.6.2) have been mapped to the TRL levels as shown in the right of Figure 4.7 and described below:

- Classification and Anomaly Detection are at the highest TRL level of 8/9 (Defined as: "Actual system 'flight proven' through successful mission operations") because these ML systems have been used in production and have been proven to be accurate in comparison to human performance.
- Prediction/Forecasting occupies TRL level of 7 (Defined as: "Actual system completed and 'flight qualified'") because these ML systems have been used in production environment. However, there are some kinks to predictive ML systems, as demonstrated by the case of the failed prediction of the Google Image Recognition in tagging black people as gorillas). On balance we are achieving higher levels of accuracy in applications like high frequency trading but we have yet to achieve "successful mission operations"
- Understanding occupies TRL level of 5 (Defined as: "Component validation in relevant environment") because ML systems in Understanding have been limited to the specific use cases only as described in Section 4.6.2.4. It has yet to generalize to other environments where ML systems can be used to understand intent across a wider domain.

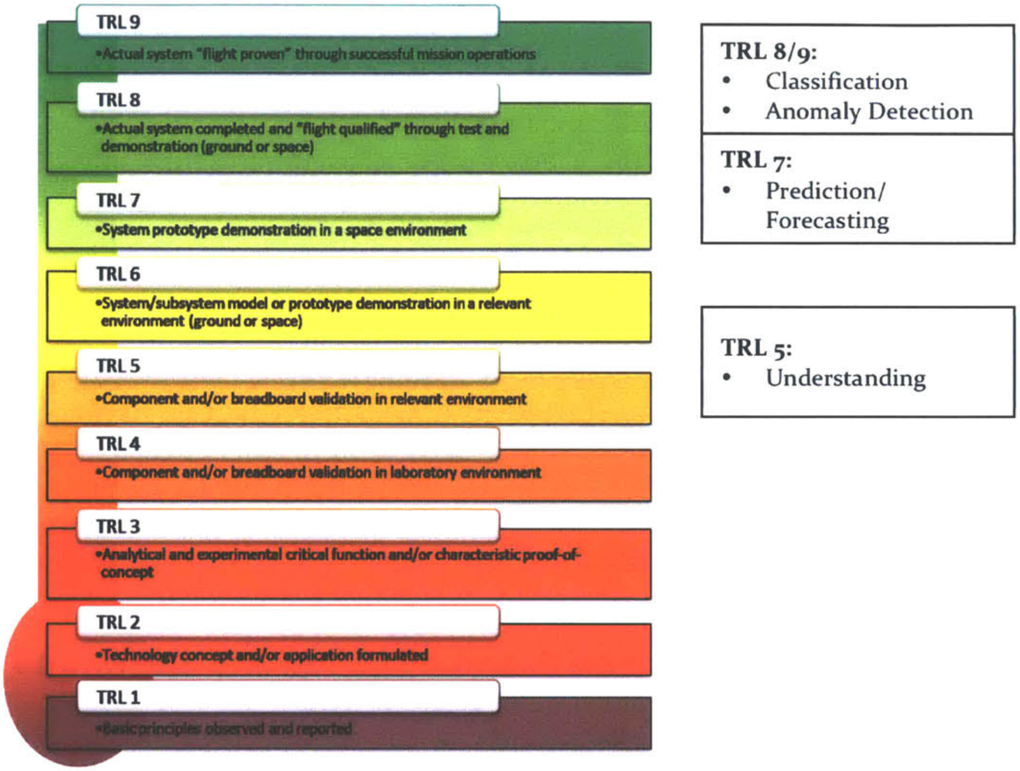
TRL 9
•Actual system "flight proven" through successful mission operations
TRL 8
•Actual system completed and "flight qualified" through test and demonstration (ground or space)
TRL 7
•System prototype demonstration in a space environment
TRL 6
•System/subsystem model or prototype demonstration in a relevant environment (ground or space)
TRL 5
•Component and/or breadboard validation in relevant environment
TRL 4
•Component and/or breadboard validation in laboratory environment
TRL 3
•Analytical and experimental critical function and/or characteristic proof-of-concept
TRL 2
•Technology concept and/or application formulated
TRL 1
•Basic principles observed and reported
TRL 8/9:
• Classification
• Anomaly Detection
TRL 7:
• Prediction/ Forecasting
TRL 5:
• Understanding

Figure 4.7: Technology Readiness Levels (TRL) with ML application domains. Adapted from: (NASA, 2007)

4.8.2 Hype Cycle

Another useful heuristic to understand the maturity of ML in context of similar technologies is the Hype Cycle Methodology proposed by *Gartner*. It represents the maturity and adoption of technologies and applications, giving decision makers a view of how a technology or application will evolve over time and provides insight to manage the technology deployment within the context of specific business goals (Brant & Austin, 2015). Referring to Figure 4.8, while the *Gartner* Hype Cycle is centered on "Smart Machines" that incorporates other technologies beyond ML, relevant ML technologies pertaining to the different application domains specified earlier are boxed in red, indicating their "maturity" and "expectation" levels in comparison with similar technologies.

The Hype cycle indicates expectations from users on how these technologies are likely to evolve and is a subjective assessment rather than a critical evaluation of these technologies. Nonetheless, it allows the decision maker to situate ML technologies with others that could potentially transform the future of the workplace.

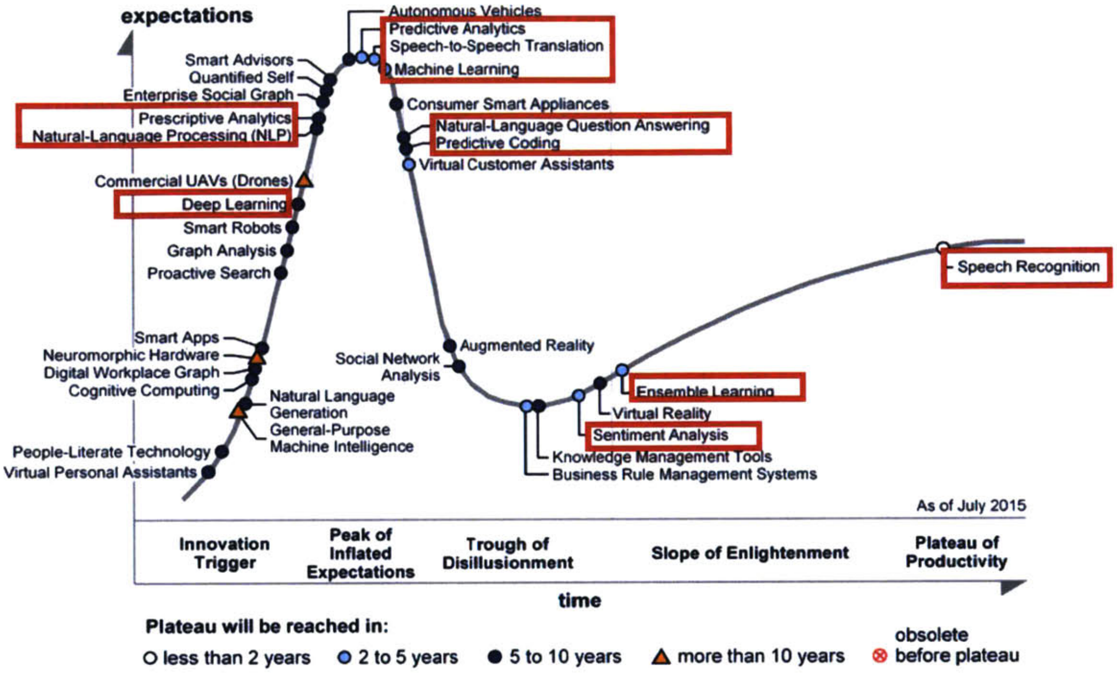

Figure 4.8: Gartner Hype Cycle (Brant & Austin, 2015)

4.9 Conclusion

An overview of ML from the technology perspective has been presented in this chapter, first by defining ML and identifying the key concepts that underpins ML. This is followed by situating ML within the AI landscape so as to clear misconceptions around similar AI technologies, as earlier identified in the literature review. To further aid the decision maker in appreciating the various ML algorithms or learners, these learners are classified according to their Learning Fields, Output or Components. A model of ML is subsequently presented that enables the practitioner to consider the approach to the ML task and conceptualizing the solution. The chapter then switches gears to discuss the applications of ML, first highlighting the drivers for the progress in ML. This is followed by a discussion on the applications domains of ML systems to showcase the diverse fields where ML had made an impact, followed by the section on the areas where ML will fail that serves as a counterpoint that ML cannot be deployed to all fields. The chapter closes with a brief highlight of the progress of ML using heuristic methods such as TRL and *Gartner's* Hype cycle.

Chapter 5 Machine Learning – Business Perspective

This chapter provides an overview of Machine Learning (ML) from a business perspective. Its intent is to enable decision-makers and managers to appreciate and identify opportunities where ML could be applied by observing current trends and the challenges that businesses face in adopting ML into their business strategy. The larger goal is to provide a concise view of ML such that these executives could have the knowledge and confidence to derive competitive advantage using ML as key technology capability.

The following sections seek to address the following Research Questions that motivates this thesis, as follow:
RQ3: ML from Business Perspective

- Industry Trends
- What are the Market drivers?
- What are the Market barriers?
- What strategies do they adopt ML?
- Industries that benefits from ML

Because ML is a set of approach and techniques that have been incorporated into many higher-level systems such as AI, Cognitive Systems and "Smart Machines", the overview and analysis in this chapter will consider any sources as long as ML is used in some form or another. Section 5.1 charts the ML ecosystem as well as how industries have been applying ML. Section 5.2 highlights the key trends from both supply and demand perspectives. Section 5.3 then articulates the challenges facing businesses, which underscores the need to develop the framework in Chapter 6. Section 5.4 concludes.

5.1 Ecosystem

The ML ecosystem consists of both the supply side (Providers of ML technologies) as well as the demand side (Users of ML technologies). (Academicians are not depicted in this ecosystem for brevity and also because they collaborate with both supply and demand side). It is noted that there is no aggregate or comprehensive data available that indicates how users have been partnering with the supply side (vendors that providing ML technologies) or developing an in-house ML capability, likely because these are proprietary information that companies would not share with research houses or at interviews. In order to comprehensive map out the ecosystem, the approach is to separate this industry overview into two perspectives of supply and demand, with the supply perspective highlighting the companies involved in providing ML technologies and demand perspective highlighting the companies that have integrated ML into their processes or products.

5.1.1 Supply Perspective

A common way to chart the supply perspective of ML is to use a hierarchy based on the types of output or service, as shown in Figure 5.1 (Zilis, 2015). Some of these classification types depicted are:

- **Platforms**: Generic platforms that can be used by companies in a wide variety of application domains. It requires close interactions between these vendors and customers in order to deliver a customized solution.
- **Tools**: Technical tools that implements ML components, that can be used by application developers to create platforms or be used as a service by users. Some of these tools are open-sourced.
- **Industry-focused**: Specialized applications that customizes ML to a specific industry and therefore caters to a select group of users, e.g. agriculture or manufacturing facility.
- **Function-focused**: Specialized applications that customize ML to specific functional areas of the business, e.g. Customer Relationship Management (CRM).
- **System-focused**: Specialized form of ML incorporated into a larger system, such as autonomous vehicles and cognitive systems.
- **Agent-focused**: Instead of specifying a platform or service, ML is incorporated into intelligent agents which serve to assist humans in performing tasks. An example is "Siri" and "Amazon Echo".

In a different fashion, Chinese companies are classified by *IResearch* according to hardware/software forms and application/technology/base layers, as shown in Figure 5.2 (IResearch, 2015) . While this classification is by no means similar to Figure 5.1, the intent is to illustrate the crowded landscape of supply-side companies in both Western and Eastern spheres, as well as how segregated these companies are in terms of geographical influence. Figures 5.1 and 5.2 is also the only aggregated illustration of supply-side companies through research conducted by (Zilis, 2015) and (IResearch, 2015), which allows potential competitors to survey the landscape of ML as well as indicate opportunities for Western and Eastern companies to expand their markets into other spheres. For instance, Western companies with hardware technology in ML, such as Nvidia can expand and fill the hardware technology gap in China, as seen by the lack of companies competing in this space in Figure 5.2.

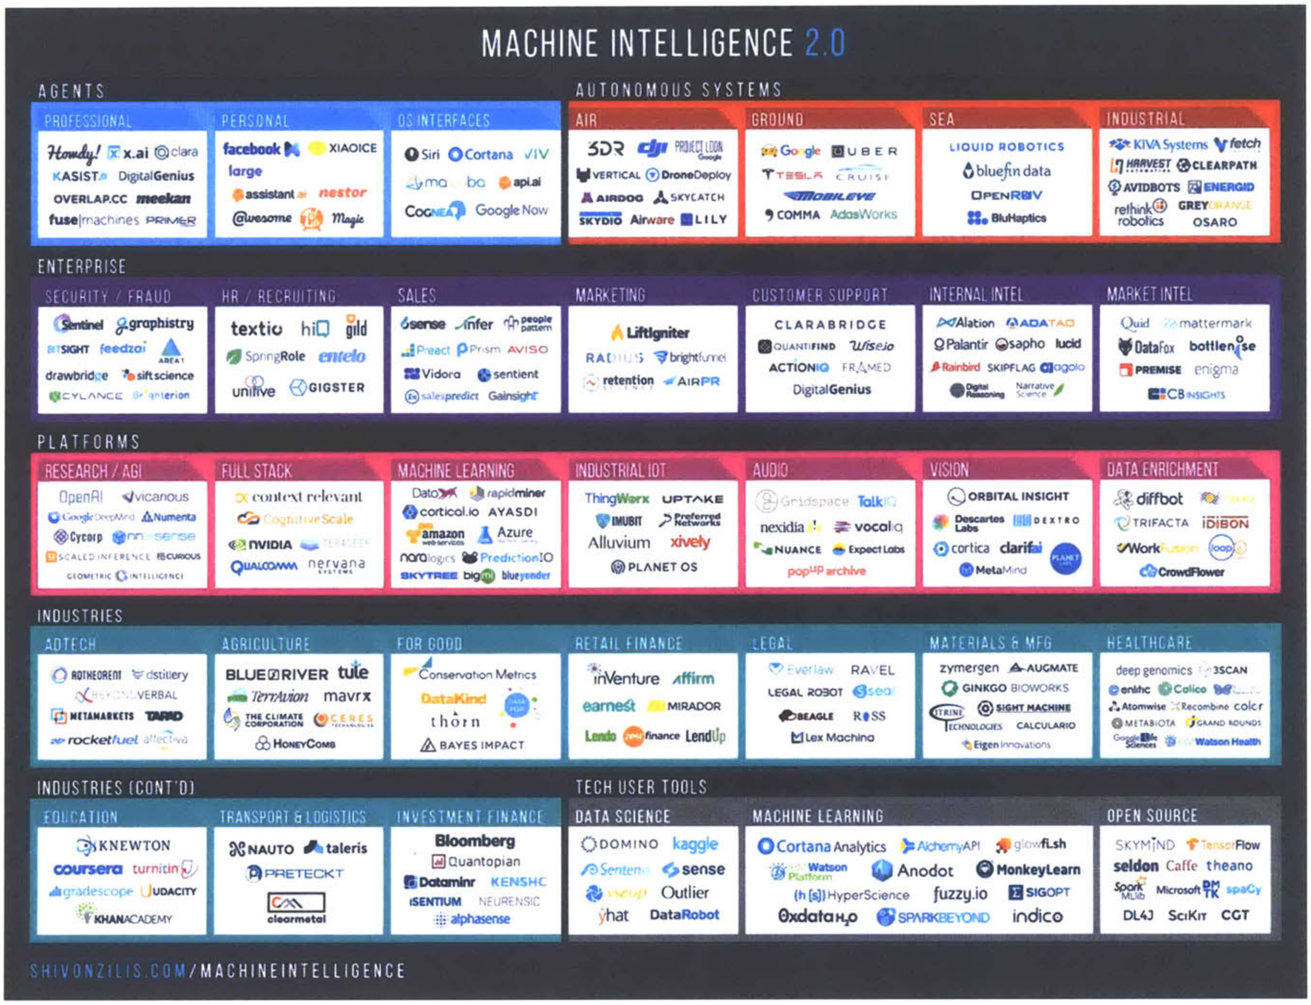
MACHINE INTELLIGENCE 2.0
AGENTS
PROFESSIONAL
Howdy! x.ai clara KASISTO DigitalGenius OVERLAP.CC meekan fuse|machines PRIMER
PERSONAL
facebook M XIAOICE large assistant.ai nestor @wesome Magic
OS INTERFACES
Siri Cortana VIV mabo api.ai Cognea Google Now
AUTONOMOUS SYSTEMS
AIR
3DR dji PROJECT LOON VERTICAL DroneDeploy AIRDOG SKYCATCH SKYDIO Airware LILY
GROUND
Google UBER TESLA CRUISE MOBILEYE COMMA AdasWorks
SEA
LIQUID ROBOTICS bluefin data OPENROV BluHaptics
INDUSTRIAL
KIVA Systems fetch HARVEST CLEARPATH AVIDBOTS ENERGID rethink robotics GREY ORANGE OSARO
ENTERPRISE
SECURITY / FRAUD
Sentinel graphistry BITSIGHT feedzai AREA1 drawbridge sift science CYLANCE brighterion
HR / RECRUITING
textio hiQ gild SpringRole entelo unitive GIGSTER
SALES
6sense infer people pattern Preact Prism AVISO Vidora sentient salespredict Gainsight
MARKETING
LiftIgniter RADIUS brightfunnel retention AIRPR
CUSTOMER SUPPORT
CLARABRIDGE QUANTIFIND Wise.io ACTIONIQ FRAMED DigitalGenius
INTERNAL INTEL
Alation ADATAO Palantir sapho lucid Rainbird SKIPFLAG Gogolo Digital Reasoning Narrative Science
MARKET INTEL
Quid mattermark DataFox bottlenose PREMISE enigma CB INSIGHTS
PLATFORMS
RESEARCH / AGI
OpenAI vicarious Google DeepMind Numenta Cycorp numenta sense SCALED INFERENCE CURIOUS GEOMETRIC INTELLIGENCE
FULL STACK
context relevant CognitiveScale NVIDIA TERADEEP QUALCOMM nervana
MACHINE LEARNING
Dato rapidminer cortical.io AYASDI amazon web services Azure nara logics PredictionIO SKYTREE bigml blueyonder
INDUSTRIAL IOT
ThingWorx UPTAKE IMUBIT Preferred Networks Alluvium xively PLANET OS
AUDIO
Gridspace TalkIQ nexidia vocaliq NUANCE Expect Labs popUp archive
VISION
ORBITAL INSIGHT Descartes Labs DEXTRO cortica clarifai PLANET LABS MetaMind
DATA ENRICHMENT
diffbot TRIFACTA iDIBON WorkFusion loop CrowdFlower
INDUSTRIES
ADTECH
ROTHEORENT dstillery BEYOND VERBAL METAMARKETS TARGD rocketfuel affectiva
AGRICULTURE
BLUE RIVER tule TerrAvion mavrx THE CLIMATE CORPORATION CERES HoneyComb
FOR GOOD
Conservation Metrics DataKind DATAKIND thorn BAYES IMPACT
RETAIL FINANCE
inVenture affirm earnest MIRADOR Lendo finance LendUp
LEGAL
Everlaw RAVEL LEGAL ROBOT seal BEAGLE ROSS Lex Machina
MATERIALS & MFG
zymergen AUGMATE GINKGO BIOWORKS SIGHT MACHINE TRINE TECHNOLOGIES CALCULARIO Eigen Innovations
HEALTHCARE
deep genomics 3SCAN enlitic Calico Atomwise Recombine color METABIOTA GRAND ROUNDS Google Life Sciences Watson Health
INDUSTRIES (CONT'D)
EDUCATION
KNEWTON coursera turnitin gradescope UDACITY KHANACADEMY
TRANSPORT & LOGISTICS
NAUTO taleris PRETECKT clearmetal
INVESTMENT FINANCE
Bloomberg Quantopian Dataminr KENSHO iSENTIUM NEURENSIC alphasense
TECH USER TOOLS
DATA SCIENCE
DOMINO kaggle Sentenai sense yseop Outlier yhat DataRobot
MACHINE LEARNING
Cortana Analytics AlchemyAPI glowfi.sh Watson Platform Anodot MonkeyLearn (h [s]) HyperScience fuzzy.io SIGOPT 0xdata H2O SPARKBEYOND indico
OPEN SOURCE
SKYMIND TensorFlow seldon Caffe theano Spark MLlib Microsoft DMTK spaCy DL4J SciKit CGT
SHIVONZILIS.COM/MACHINEINTELLIGENCE

Figure 5.1 : Machine Intelligence 2.0 (Zilis, 2015)

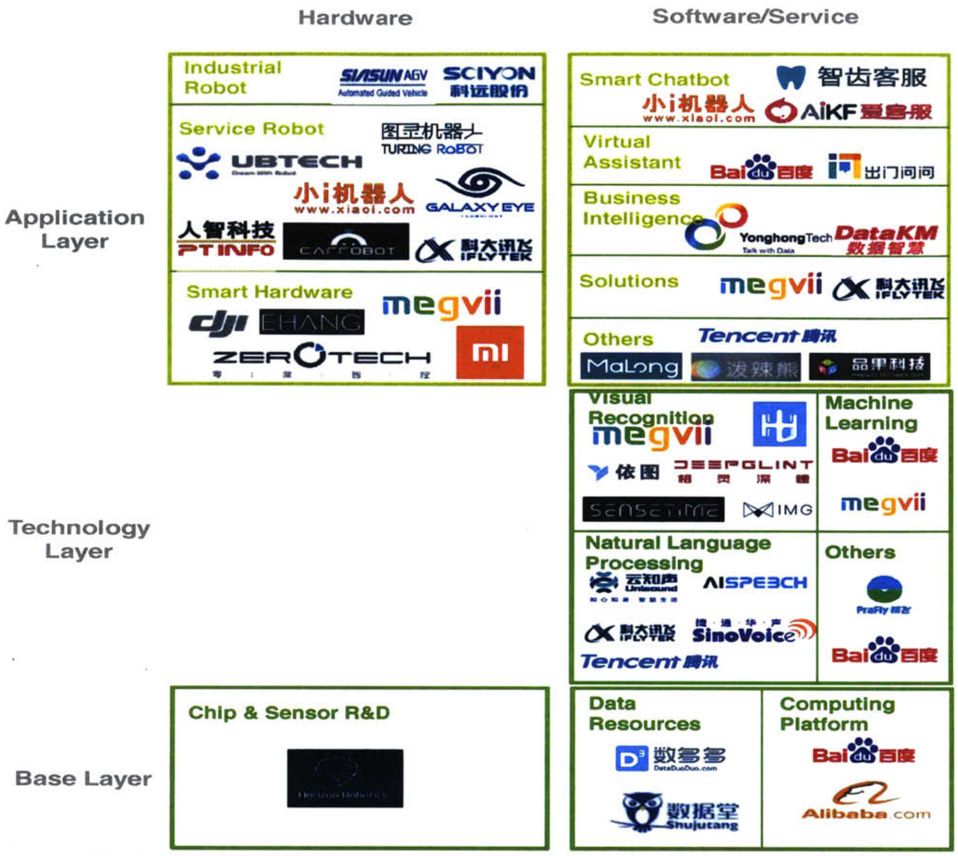

Figure 5.2 : Chinese Machine Learning companies (IResearch, 2015)

5.1.2 Demand (User) Perspective

The demand perspective is even more diverse than the supply perspective, yet it also lends itself to some forms of classification, despite lack of available information on how companies are deploying ML. Users can be grouped by industry areas such as manufacturing and financial services, within which business functions such as demand forecasting or credit worthiness evaluation demonstrate how ML is employed, as illustrated in Figure 5.3. These examples showcase how ML is employed across various industries. Together with the discussion in Section 4.6.2 on the four application domains of ML from the technology perspective, organizations would be more informed when seeking opportunities to integrate ML into their business by referring to these examples.

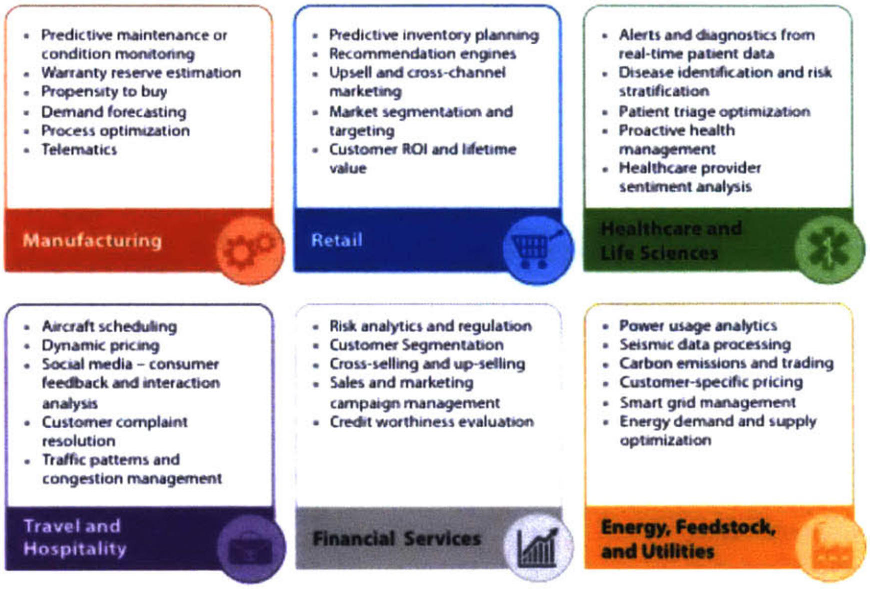
Predictive maintenance or condition monitoring
Warranty reserve estimation
Propensity to buy
Demand forecasting
Process optimization
Telematics
Manufacturing
Predictive inventory planning
Recommendation engines
Upsell and cross-channel marketing
Market segmentation and targeting
Customer ROI and lifetime value
Retail
Alerts and diagnostics from real-time patient data
Disease identification and risk stratification
Patient triage optimization
Proactive health management
Healthcare provider sentiment analysis
Healthcare and Life Sciences
Aircraft scheduling
Dynamic pricing
Social media – consumer feedback and interaction analysis
Customer complaint resolution
Traffic patterns and congestion management
Travel and Hospitality
Risk analytics and regulation
Customer Segmentation
Cross-selling and up-selling
Sales and marketing campaign management
Credit worthiness evaluation
Financial Services
Power usage analytics
Seismic data processing
Carbon emissions and trading
Customer-specific pricing
Smart grid management
Energy demand and supply optimization
Energy, Feedstock, and Utilities

Figure 5.3: Machine Learning Categories (Gadkari & Mohan, 2014)

5.2 Trends

Following the same approach as in Section 5.1, the trends in ML in this section is analyzed from the supply and demand (user) perspectives. This section provides a quick overview of trends and is not intended to replace the content of what many market research reports provide. The intent of this section is to identify key trends that drive the need for the framework that follows in Chapter 6.

5.2.1 Supply Perspective: Market Size

In terms of sizing of the supply-side market, there is no specific data that aggregates the market size of all companies that specifically sells ML as a service or product. Nevertheless, the market size can be generalized to "Smart Machines[2]" or Artificial intelligence where there is overlap with ML systems and applications. *BCC Research* (McWilliams, 2016) estimates the total market for Smart Machines at $6.6 billion in 2015, with projected growth to reach $15 billion in 2021 as indicated in Table 5.1. GII Research (Wood, 2016) estimates that the total market of AI is expected to reach $5.05 billion in 2020. There are other market research companies prognosticating different market values of the AI and Smart Machines. It is noted that analysis of this market is difficult due to buzzwords, vague definitions and difficulty in classifying the applications. For instance, it is difficult to ascribe the fraction of spending on "big data" that is related to ML or other AI application. It is envisioned that with the classification presented in Section 5.1 which could be commonly accepted by market research companies, more

[2] Smart machines as defined by *BCC Research* are hardware or software systems that can accomplish their designated tasks even under conditions of uncertainty and variability. They are also able to operate autonomously, adapt and learn from changing conditions in real time and communicate with other machine (McWilliams, 2016)

fine-grain studies could be conducted that would benefit readers (the decision makers) by clarifying the confusion over these terms.

Table 5.1: Global Market for Smart Machines, Through 2026 ($ Millions) (McWilliams, 2016)

	2015	2016	2021	CAGR% 2016-2021	2026	CAGR% 2021-2026
Total sales of smart machines	6,563.5	7,444.8	14,977.3	15.0	40,277.0	21.9

5.2.2 Supply Perspective: Products/Service Trends

With more computational power, enabled by cloud technologies, faster processors and cheaper data storage, ML algorithms have evolved from plain-vanilla predictive analytics systems towards more complex deep neural networks, as mentioned in Section 4.5. This has allowed for a larger variety of companies that supply ML technologies as well as the exponential increase in startups and related investments in the last five years according to data compiled by researcher CB Insights for Bloomberg News (Clark, 2015) in Figure 5.4.

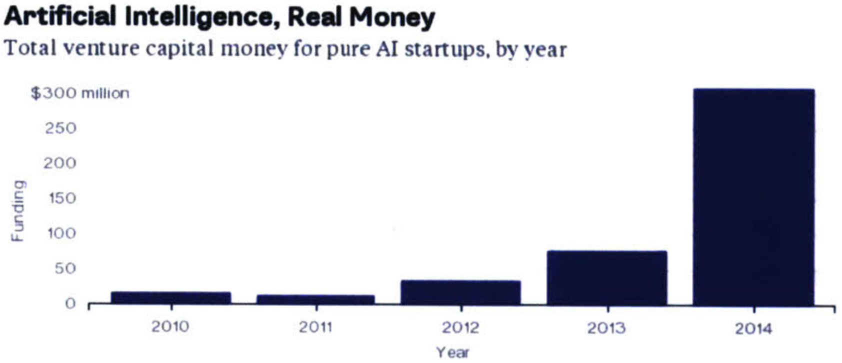

Figure 5.4 : Total Venture Capital Investments in Artificial Intelligence (Clark, 2015)

There are more companies that provide pattern recognition solutions in fields such as text analytics ("Skymind") or image recognition ("Clarifai"), as well as intelligent assistants ("Viv"), where these ML technologies are utilized to perform more abstract learning tasks. The concept of "System of Systems" have also taken root in ML, represented by the rise in Cognitive Systems, where different ML technologies, together with other parts of AI such as perception and reasoning systems, are integrated to form a holistic solution.

There is also a general trend towards accessibility of analytics tools, where companies can deploy these tools easily by procuring a service or deploying them in their production environment, without the need for extensive customization. Companies such as "Tableau" and "Dato" are active in this application area, alleviating the need for companies to rely on data science talent to develop an in-house capability. Furthermore, such Machine-Learning-as-a-Service (MLaaS) concept allows for companies with limited budget to rely on these tools and existing staff to derive insights from available data – that is if these companies already do have mechanisms for collecting data in the first place. Examples of these MLaaS companies include MLaaS by Amazon Web Services, "BigML", "Ersatz" (C. Wilson, 2015).

5.2.3 Supply Perspective: Changing Ecosystem

Given the increasing ubiquity of ML, ML technology is now seen as a complex socio-technical system that affects and is affected by large number of factors. Externally, there is an increasingly vocal concern on data privacy, such as how large companies like "Google" make use of our private data and how these companies accede to these demands (Simonite, 2016). There needs to be a conversation on the tradeoff between aggregated insights and privacy access in a transparent manner. Furthermore, society also needs to decide the scope of decisions that could be predicted/recommended by a machine and to what extent is the machine or its designer culpable for a risky decision that could end in loss of life.

Internally, researchers have pushed further ahead with different concepts of ML, from scaling up ML from just a couple of connected servers to leveraging on cloud platforms to conduct ML at the hyper-scale level. Different paradigms are being explored, such as how to make the ML learn from learning in a continuous manner and how to utilize unique human insights to help make a more accurate and intuitive system (Beyer, 2016). This research thrust is accelerated by large amounts of investments poured into AI (and ML as a related technology) from both the public and private sector. For instance, China has announced in May 2016 that it would speed up the development AI sector and create a market worth more than 100 billion yuan ($15.26 billion) over the next three years ("China rolls out three-year program for AI growth," 2016).

5.2.4 Demand Perspective: Extent of Adoption

From the demand-side, there are no specific reports or surveys indicating the extent of adoption of ML by specific companies. Nevertheless, information can be extrapolated to form a generalized observation of the state of interest or integration of these ML technologies, although there are limitations to such interpretations. According to 451 Research, where a survey of 200 North American IT executives from companies with 500 or more employees across diverse industries was conducted, organizations had consistently showed a great appetite for machine learning, with 41.5% of respondents desiring a program of this nature within 12 months. In addition, 22% of respondents said they already had a machine learning program (Lehmann, Roy, & Winter, 2016). This suggests that a once-TRL 5 technology (using the TRL methodology in Section 4.8) is has now moved toward TRL8/9 and become generally adopted. In particular, predictive analytics and recommendations has emerged as the top-ranked types of machine learning programs as shown in Figure 5.5. Similarly, *Gartner* survey of 437 IT and business leaders indicated that 76% of them planned to invest or had already invested in big data analytics (of which ML is a specialized branch) (Kart & Heudecker, 2015). This survey canvassed respondents from variety of industries globally, with sizes of companies ranging from less then $250 million in revenue to more than $10 billion.

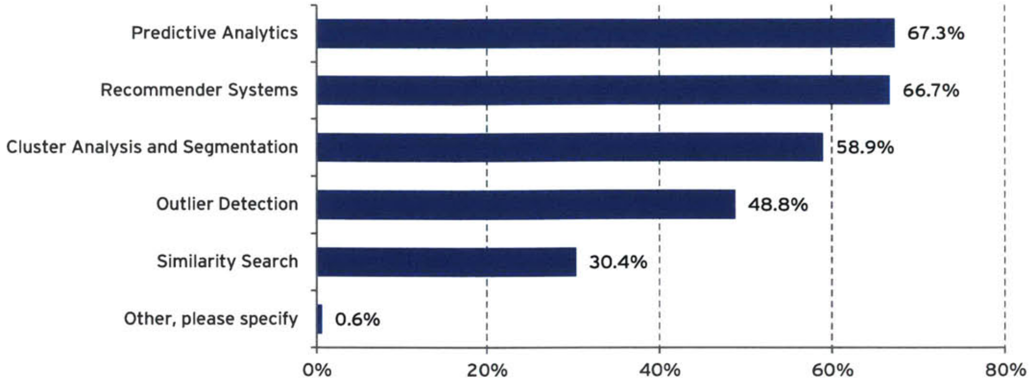

Figure 5.5 : Predictive Analytics top ranked (Lehmann et al., 2016)

However, the above-mentioned market research had only polled larger companies with 500 or more employees. The trend of increase adoption of ML does not seem to be significant in smaller companies. A 2014 *Nielson* study interviewed about 2,000 North American small businesses (no size indicated) and found mixed sentiments or assumptions about predictive analytics. 41% think that conducting market research is too costly and 35% said that they have yet to consider using predictive analytics to gain competitive advantage. These figures indicate a lower level of interest as compared to lager companies, but it remains inconclusive as to whether smaller companies have adopted ML technologies due to the lack of available data on these smaller businesses.

Nevertheless, the same study noted that in the span of 2014 – 2015, there is a significant uptick in data science boutique consultancies and predictive analytics platforms and applications that provides services to these smaller businesses ("How Small Businesses Can Scale the Big Data Barrier," 2014). With the increase in supply-side technology companies, it is likely that the barriers to entry is reduced which would encourage smaller organizations to consider ML as a technology enabler.

5.2.5 Demand Perspective: Organization Trends

The trend of adoption of ML and its effectiveness are examined using the following lenses that comprise an organization: i) Leadership, ii) Decision-making and Structure, iii) Processes, iv) People and culture. Some of these trends were previously identified in Chapter 2: Literature Review – for completeness, these observations are further detailed here to highlight the general trend. As mentioned in Section 2.3, there is confusion of the terminology of "data analytics", "predictive analytics", "Machine learning" which essentially mean the same. These terms have been used interchangeably in the various surveys cited below. For consistency, these terms are taken to mean the same as Machine Learning (ML).

i) Leadership

The previously mentioned 2014 *EIU* survey notes that 64% of 1,135 senior executives had already changed the way they made decisions based on big data, with a majority of these executive in North American organizations. 25% were planning to leverage on data in the next two years (Economist Intelligence Unit, 2014). While the intent to do so is clear, the follow-through seemed to be fraught with challenges. The *Bain's* survey notes that only 23% of companies surveyed have clear strategies for using "predictive analytics" effectively (Sinha & Wegener, 2013). This trend suggests that companies do have some way to go in articulating a vision and strategy that exploits the value that ML or predictive analytics bring to their businesses. *MIT Sloan* research suggests that companies that aligns the strategic plan for predictive analytics with their overall corporate strategy are much more likely to be successful (B. S. Ransbotham, Kiron, & Prentice, 2016).

ii) Decision-making and Structure

Many large corporates have stood up the senior executive post of a Chief Data officer (CDO) or Chief Analytics Officer (CAO) to support the "data analytics" strategy (PwC, 2015). However, there is no consistency in connecting these senior posts to functional lines in the company and fostering collaboration between silos. *Bain* reports that only 36% of companies have a dedicated data-insights team that is able to mix data science, business acumen and technical expertise. Such structural changes must also go hand in hand with changes in the decision making process. Those who frequently make data-driven decisions are majority top performers in *Bain's* survey of 400 companies (Sinha & Wegener, 2013). It is imperative that these companies do so lest they get left behind by forward thinking companies who have outperformed their competitors by leveraging on ML to derive insights into their collected data.

iii) Processes

In a study of more than 30 pilot projects in early-adopter companies of ML, *Accenture* found that companies reported several business processes that have been improved by utilizing ML, as shown in Figures 5.6 (H. J. Wilson, Alter, et al., 2016). For instance, "Predictive analytics" have been found to improve the management of risk compliance, financial resources and ability to market and sell products and services, with the impact depicted by the thickness of the flow lines highlighted in red.

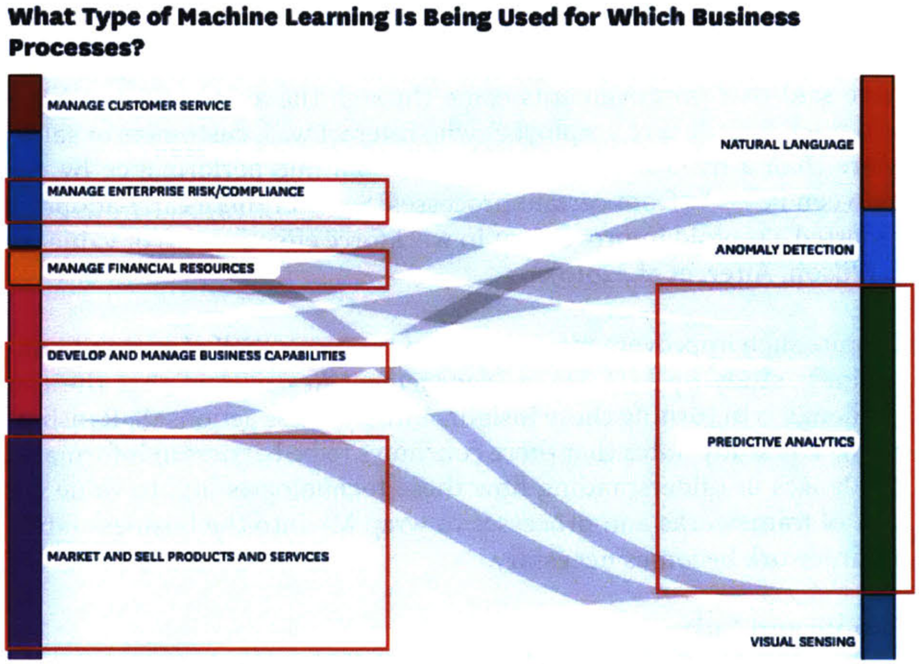

Figure 5.7: ML Benefits to Business Processes (H. J. Wilson, Alter, et al., 2016)

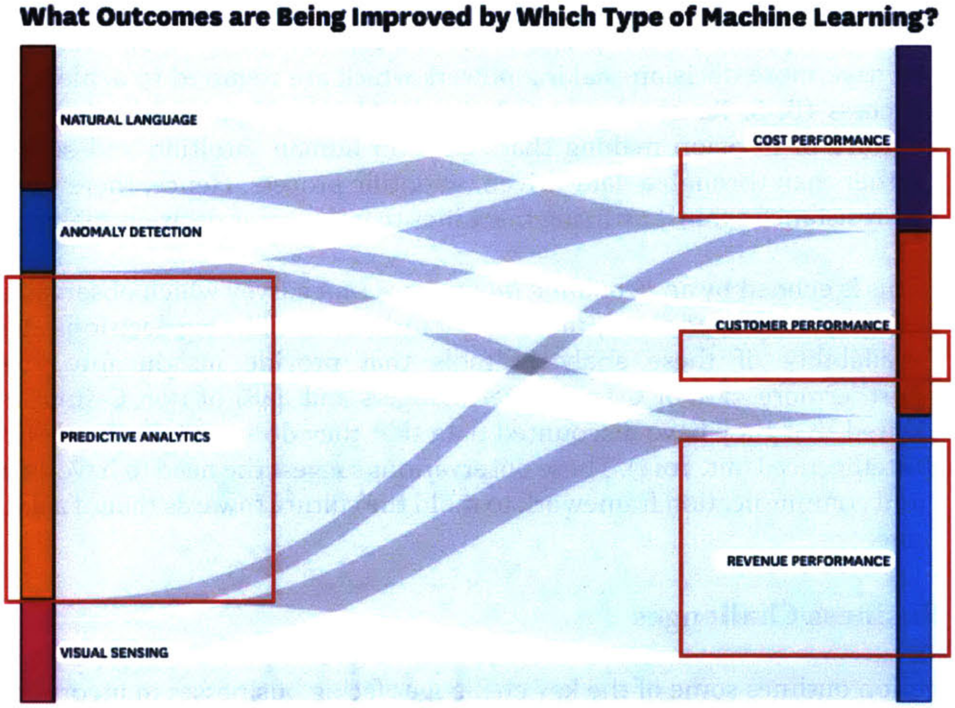

Figure 5.6: ML Benefits to Business Outcomes (H. J. Wilson, Alter, et al., 2016)

Similarly, Figure 5.7 shows the impact on business outcomes from utilizing these ML technologies. Continuing with the previous example of "Predictive Analytics", business outcomes of revenue performance was realized more than customer or cost performance as depicted by the thickness of the flow lines highlighted in red. (H. J. Wilson, Alter, et al., 2016)

In the same *Accenture* report, majority of these companies have reported some improvements to top-line performance. Nearly half of early movers surveyed have said that improvements came through the automatic provision of timely and predictive data to employees who interact with customers or sales prospects. More than a third realized gains in bottom-line performance by cutting costs between 15 – 7-% from certain processes. Some of these early-adopters have also reported a tenfold improvement in workforce effectiveness or value creation (H. J. Wilson, Alter, et al., 2016).

Despite such improvements in business processes, half of respondents of another business study by *MIT Sloan Management Review* said that the top analytics challenge is in turning these insights into business actions (S. Ransbotham et al., 2015). The study notes that there continues to be barriers in information sharing, challenges in understanding how these technologies link to value creation and lack of frameworks and processes to wrap ML into the business fabric. As such, a framework becomes necessary.

iv) People and Culture

Key findings from a recent 2016 survey by *MIT Sloan Management Review* indicated that while the optimism on the potential of analytics remain strong, most companies are not prepared for the cultural change of allowing computers to have more decision-making powers which are required to achieve sustained success (B. S. Ransbotham et al., 2016). Many companies continue to have a culture of decision making that relies on human intuition and sense-making rather than through a data-driven, scientific process. Hence, there continues to be resistance against utilizing machines to guide their decision making.

This is echoed by an *Economic Intelligence Unit* survey which observed that 30% of CEOs continued to rely on their gut instinct to make big decisions, despite the availability of these analytics tools that provide insight into their data. Furthermore 52% of C-suite level mangers and 44% of non C-suite managers stated that they have discounted data that they do not understand (Economist Intelligence Unit, 2014). These observations suggest the need to have an adoption and communication framework to mold the culture towards that of a data-driven one.

5.3 Business Challenges

This section outlines some of the key challenges facing businesses in incorporating ML into their analytic strategy, based on trends identified in the previous sections and the literature review conducted in Chapter 2. These key challenges can be classified by: Organizational, Operational, Infrastructure, Technology and Size.

5.3.1 Organizational

As highlighted in Section 5.2 Demand Perspective: Organization Trends, the key organizational challenge that a business would face is the articulation of a data-driven vision that is in alignment with overall corporate strategy. Developing a culture where data-driven analysis and decision making is the norm, is like any management initiative, where long lasting impact can only be made if it is supported by a fundamental shift in culture through leadership in change management, supported by processes and structure.

The second organizational challenge is the need to adjust business processes and structure to incorporate ML. Many of ML adoption initiatives are likely to fail without methodologies for framing and selecting which of the business processes can and should leverage on ML and frameworks for evaluating these initiatives against business outcomes. A good organizational structure to take ownership of these ML initiatives becomes critical in larger companies where there is a need to collaboration between functional lines such as IT and business strategy.

The third organizational challenge relates to people. The shortage of analytical talent such as data scientists will continue to be a challenge for large corporations. This challenge becomes even pronounced for smaller companies that are not able to afford such expertise and because these analytical professionals tend to be attracted to larger and more reputable companies. Companies therefore need to devise different strategies such as outsourcing or utilizing MLaaS service providers to mitigate these talent shortages.

5.3.2 ML Technology

The complexities in understanding, modelling and fine-tuning ML algorithms have been covered in Section 4.3. There are however further implications on business from the ML technology perspective. Research into new ML algorithms continues, averaging about 120 papers per year in the Journal of Machine Learning Research and 130 papers per year in International Conference on Machine Learning. Organizations therefore must decide if and when to incorporate newer ML algorithms, balancing the tradeoffs between transitioning from existing analytical tools, the benefits of the new model and ultimately the business value created.

The decision is made more difficult because ML systems are challenging to understand. While Chapter 4 had endeavored to articulate a layman perspective of ML, these systems are inherently complex and non-linear. This is a challenge in industries where regulations or even the customers require some sense of how these predictive analytics work, especially in high risk domains such as the military and healthcare. Without a good understanding or intuition of how ML works, chances are that users are less likely to trust the outcome of the system, regardless of how accurate the prediction is. It is only when an organization understands the advantages and constraints of ML and communicates these ideas effectively to both users and employees that ML will eventually be adopted as system of choice for deriving predictive insights from their data.

5.3.3 Operational

As mentioned in Section 4.6, ML depends on data that is of quality, secure and integrated with the rest of the data streams. The quality of the training sets is critical to ensure that the model is accurate. Low quality data could be due to noise, incomplete or missing values and bias inherent in the collection. There is a need to control the data collection process so as to mitigate this. Furthermore, there is a regulatory requirement in many industries to ensure the data is secure and governed with policies in place to manage and control data leaks. For big data, it is also imperative for these governance policies to manage the large volume of data, prevent data inaccuracies and related problems as the data systems scale.

Beyond data collection, there is also a need to integrate data streams from different sources. The challenge is in the synchronization between disparate datasets, transforming them into a form that can be processed by the ML system and to do so effectively without being "cursed by dimensionality". Human intuition is also necessary to link the data exploration to value creation. Operating a ML system should address business needs and yield quantifiable results. There needs to be an iterative process in which data scientists or the analytics team can hypothesize, query, explore and iterate in response to changing business needs. Such data challenges suggest the need for a set of best practices to address the above.

Beyond the data challenge in the operational environment, the organization is likely to face issues in deploying these models in a more complex production environment. Scaling up from development environment is one key challenge as one seeks to operationalize. While the team may develop on open-source tech tools such as "Python" or "R", such tools become unwieldy when dealing with large volumes and velocities of data. Therefore, it is critical to anticipate the end-state requirements for the ML system as well as integration challenges with the existing IT infrastructure and use this basis to create a development roadmap.

Even with a development roadmap that produces valuable models that is able to scale and integrate from development to the production environment, the models will require constant adjustments in response a dynamic production environment. The future could stop behaving like the past. New unexpected changes to the business environment, such as black swan events, will lead to the current model breaking down because the current model had not incorporated these scenarios. This means that the roadmap will need to branch off to different variations of models in response to these environments. A model management system is therefore necessary to organize and manage which models to use in which scenarios in an agile manner.

5.3.4 Infrastructure

Traditional relational databases and legacy software and hardware can fail under the volume of data and processing power required by these ML systems. Big data is increasingly being used to train these learners. For larger organizations, existing data infrastructure such as those based on Hadoop Distributed File System (HDFS) would be sufficient to distribute large amounts of data in clusters. For organizations without the infrastructure, particularly for smaller-sized businesses, they would need to architect an appropriate storage architecture that can scale up easily and incorporate various platforms even before they embark on data collection and adopting ML to process the data.

Adding to this challenge is that ML tasks are computationally intensive. New and more accurate ML algorithms demand an even greater amount of computation, in particular deep learning types. Deep Learning algorithms typically take weeks to months to process data. Building powerful and scalable computer infrastructure is a challenge for companies facing resource constraints.

Fortunately, and as mentioned in Section 4.5, cloud technologies, hardware acceleration and distributed computing are available to many of such companies to scale up and allows flexibility in their adoption approach. While such investments in X-as-a-Services must be appropriately balance against the business outcomes, it also allows companies to experiment with and implement different approaches, even opt out of such services when not required.

5.3.5 Size

In the case of larger companies, size is a trivial factor, because they have the resources to develop an in-house capability or the option of procuring a ML capability, for e.g. IBM acquiring Alchemy API to bolster their product offerings ("IBM Acquires AlchemyAPI," 2015). Small businesses, however, face more acute challenges as they seek to balance limited budget vs the expected value from ML adoption. A large majority of these smaller businesses continue to view market research as too costly or simply lack the time. They still lack expertise and time to leverage on available information to deliver better business outcomes ("How Small Businesses Can Scale the Big Data Barrier," 2014). This trend is likely to be more acute in smaller companies compare to larger ones. However, if these companies do not innovate and create experiences that adjust automatically, respond in real time, predict customer habits and deliver tailored services to their consumers, they will be pushed out by competitors (large and small) who are more willing to take the risk in adopting predictive analytics in ML. Fortunately, there is a trend towards more accessible and holistic tools that enable small businesses to bridge this gap.

5.4 Conclusion

An overview of ML from the business perspective has been presented, by first charting the ecosystem of ML by its constituent "supply" and "demand" perspectives, followed by highlighting some of the key business trends in the same two perspectives. This chapter concludes by observing the challenges faced by large and small businesses from various aspects that ranges from organizational to infrastructure. These challenges highlight the need for a generalized framework for companies in considering, adopting and implementing ML. This framework should cater to the different business context as well as the state of digital innovation of companies.

Chapter 6 Framework

This chapter presents the Machine Learning (ML) integration framework for organizations where it guides users in deciding and architecting ML systems into their businesses in order to build technology capabilities to meet their stakeholder needs. Chapter 5 had previously articulated the business trends regarding ML and highlighted some key challenges in Section 5.3 that this framework will address. That said, the framework makes a distinction between the different profile of companies (small enterprises to large multi-functional industries) because the requirements of stakeholders vary from companies to companies.

This framework is intended to be used by key decision makers and enterprise architects, in conjunction with the technology and business perspectives highlighted above, so that informed decisions could be made. Even then, some assumptions are made prior to the framework. This framework only applies to organizations where there is some representation of data in an easily consumable form, i.e. the companies have a database management system to store the data. This framework does not apply to "analog companies" where there are no infrastructures in place to collect and manage the data. For this, the reader is referred to (Weinelt et al., 2016) for guidance on how to transform their businesses digitally.

The research questions that this chapter seek to address are:
RQ4: ML Integration Framework

- Application of ML framework
- Elements of the framework
    - Review of Enterprise Model
    - Identification and Assessment of Opportunities
    - Evaluation of Technology Adoption
    - Architecting the ML system

6.1 Outline of Framework

The application of the ML Integration framework is best suited after the enterprise decision makers have deliberated and decided on an enterprise strategy for integrating ML. Some enterprises would need to transform itself substantially before even applying ML, while others might find it easier if they already have the right infrastructure, process and culture in-place. Readers are encouraged to refer to frameworks such as ARIES (Architecting Innovation Enterprises) (Nightingale & Rhodes, 2015) for suggestions on how to transform their enterprises in response to changing external and internal ecosystems. Since ML is often considered as a paradigm shift for companies who want to embark on a data-driven approach towards supporting decision making and optimization of scarce resources, the ARIES framework is suitable because it is a systemic and holistic architecting approach to enterprise transformation, where it considers the emergent properties from interrelated enterprise elements in response to dynamic contexts and satisfying diverse sets of stakeholders.

The ML Integration framework subsequently forms the next "stage" of implementation after deciding on an appropriate enterprise strategy. Henceforth, the team/individual responsible for utilizing the framework to transform their enterprise is referred to as the "System Architect". The term "System Architect" (short-form: Architect) is preferred because it suggests that the role consists of conceptualizing a blueprint for operationalizing ML into their enterprise in a holistic and systematic fashion.

The framework is organized into four parts: i) Reviewing the Enterprise Model, ii) Identifying the Opportunities, iii) Evaluating Technology Adoption and iv) Architecting of ML System, as shown in Figure 6.1. The **Review of the Enterprise Model** guides the architect in evaluating the analytical maturity of the company, differentiating them into four levels of analytical organizations – Analog, Challenged, Practitioner, Innovator. With this review of the enterprise as-is, it allows the architect to appreciate the level of opportunities and challenges available. The next part **Identification of Opportunities** uses the Product/Process/Insight framework to explore opportunities in the ML domain and uses the Viable/Valuable/Vital (3V) frameworks to screen these opportunities. Third, in the **Evaluating Technology Adoption**, the enterprise is evaluated using the Technology/Organizational/Environment (TOE) framework to ascertain the level of acceptance of new technology and to identify areas that need to be addressed. Lastly, the **Architecting of the ML System** part guides the architect in determining stakeholder requirements, conceptualizing different architectures, selecting the architecture and designing the ML system using an object-process methodology to frame the architecture to reduce ambiguity and complexity for subsequent implementation efforts.

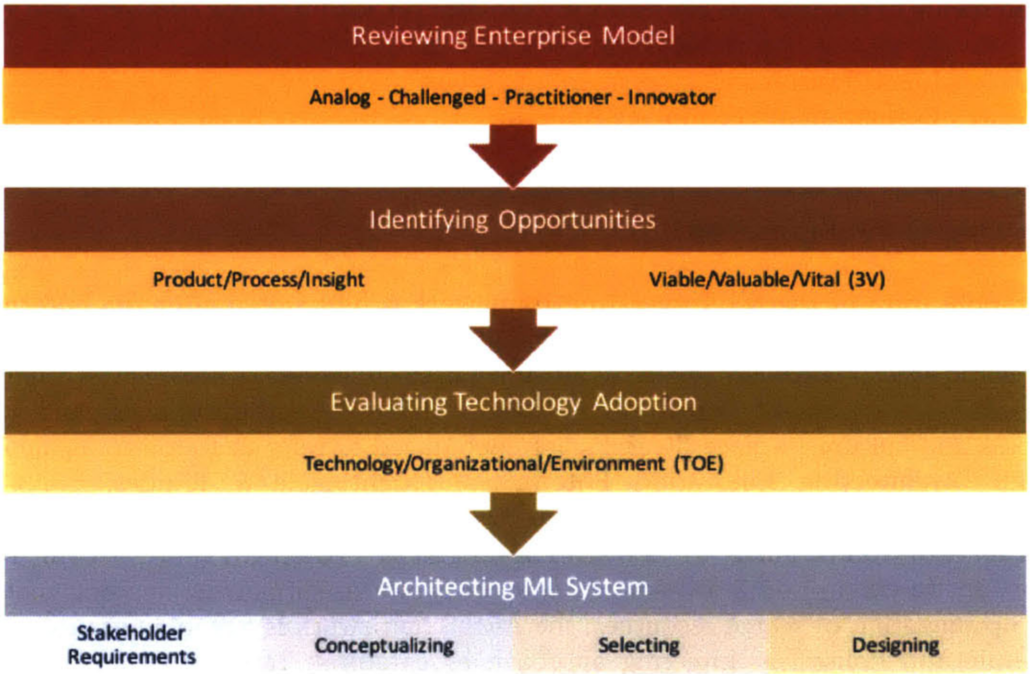
Reviewing Enterprise Model
Analog - Challenged - Practitioner - Innovator
Identifying Opportunities
Product/Process/Insight
Viable/Valuable/Vital (3V)
Evaluating Technology Adoption
Technology/Organizational/Environment (TOE)
Architecting ML System
Stakeholder Requirements
Conceptualizing
Selecting
Designing

Figure 6.1: Machine Learning (ML) Integration Framework

6.2 ML Integration Framework

6.2.1 Reviewing the Enterprise Model

The first part of the ML Integration Framework is the **Reviewing the Enterprise Model** Framework, illustrated in Table 6.1. Four levels of analytical organizations are shown in the leftmost column: Analog, Analytically Challenged, Analytical Practitioner and Analytical Innovators. This enterprise model is adapted from the "Three Levels of Analytical Organization" (Kiron, Prentice, & Ferguson, 2014). Component attributes define which of these four levels that an organization would belong:

- **Decision-making**: How does the organization use data to drive decision making?
- **Strategy**: How is analytics applied – Strategically, Operationally or Tactically?
- **Infrastructure**: How is data collected and processed?
- **People**: What kind of relevant skillsets do they have?

It is up to the architect to determine the "as-is" model of their enterprise, by mapping out their organizations through surveys and interviews. Once the model is determined, the architect needs to consider the implications for subsequent steps of the ML integration framework, as shown in the rightmost column in Table 6.1. If the enterprise is "Analytically Challenged", the architect needs to focus on building an analytics culture in the system design portion by emphasizing on a technology adoption framework of ML. Likewise, if the enterprise is a "Analytically Practitioner", it would be advantageous to focus on optimizing internal business processes using ML and thereafter using this experience to innovate product or service offerings. Lastly, when the enterprise is a "Analytical Innovator", it should focus on creating innovative products, services or make improvements to its processes, through the utilization of ML techniques. As mentioned earlier, the framework will not focus on "Analog companies" where there is no process or structure for analytics. The reader is referred to (Weinelt et al., 2016) and the ARIES framework for advice on how to transform their businesses digitally, in a holistic and systematic fashion.

Table 6.1: Review of Enterprise Model. Adapted from (Kiron et al., 2014)

Enterprise Model	Attributes				Implications when implementing ML
	Decision-Making	Strategy	Infra - structure	People	
No Analytics "Analog Companies"	No data-driven decision making	Analytics not applied in operational or strategic aspects	Lack of data management/ data not collected at all	Absence of analytical talent to manage data	Need to consider holistic transformation of enterprise to become data-driven
Analytically Challenged	Rely more on management experience than data analytics	Focus on cost reduction in use of analytics	Suffer from data quality and access issues	Lack appropriate data management and analytical skills	Need to foster an analytics culture that is open to insights gained from ML

Analytical Practitioners	Working to become more data-driven	Primarily operational in application of analytics	Have "just good enough" data	Have more information to make decisions	Need to optimize internal processes using ML, so as to solidify the culture
Analytical Innovators	Analytics culture driven by senior mandate	More strategic in application of analytics	Place a high value on data	Higher level of data management and analytical skills	Need to create innovative products and services using ML

6.2.2 Identifying the Opportunities

The second part of the ML Integration Framework is **Identifying the Opportunities**. Viewed from the business perspective, this thesis observes that the applications of ML technologies typically fall into three main categories: i) Products, ii) Processes, and iii) Insights, as illustrated in Table 6.2. **Product** applications embeds the ML technology to create products or services to provide benefits to customers, such as virtual assistants. **Process** applications uses ML technologies to improve or automate operations within an organization's workflow. **Insight** applications uses ML to gain insights that can inform strategic and operational decisions across the organization. While these applications were previously covered in Sections 4.6 and 5.1, the questions in Table 6.2 guide the architect in considering the application areas where ML will make an impact. Specifically, these questions require the architect to address the value proposition that ML brings to these applications.

Table 6.2 : Categorizing Opportunities: Products, Processes, Insight

Categories	Questions
Products	• What is the value proposition to the customer by embedding ML technologies into these products/services? • What existing products/services can be made more valuable by making them more effective, convenient, safer, faster or distinctive? • What is the business outcome arising from these ML-enabled products/services?
Process	• What business processes can be automated and what is the expected business outcomes? E.g. increase speed, reduce operational costs. • Of these, which process can be augmented using ML (e.g. clinical decision systems), that allows the work to be done faster, more accurately with more support from ML. Which of these can be replaced entirely using ML (e.g. customer service support) that allows ML to take on all of a worker's responsibilities and by performing the work flawlessly without human error? • Are there any processes/tasks that is at a scale that is impractical with conventional approaches, e.g. processes large amounts of data that a human cannot cognitively process?
Insight	• What conclusions can be drawn from unstructured data and what predictions can be similarly made from operational data? • What are the business outcomes from such prediction, e.g. reduced costs, improved efficiencies, enhanced customer service?

Once the architect has decided on the application(s) where ML technologies will make an impact, he would need to evaluate the business case and assess the opportunities for investing in this technology. A starting framework for assessing the opportunities is the **3V Framework** proposed by Deloitte where it analyses opportunities of cognitive technologies along the three screens of Viable, Valuable and Vital, using i) Indicators and ii) Screening tools of the enterprise (Schatsky, Muraskin, & Gurumurthy, 2015). This thesis adapts the 3V framework to focus on ML technologies, as compared to the generic cognitive technologies that Deloitte referred to in their journal.

i) ML Indicators

The "ML indicators" highlight observations of the enterprise that would make embedding ML technology viable/valuable/vital. As observed in Section 4.6.1 Appropriate Applications of ML, tasks that are too complex to be traditionally programmed or those that require processing large amounts of data beyond a human cognitive capability are "viable" areas to be considered for augmentation or replacement using ML. Similarly, the task that require adaptability over the course of a dynamic environment would derive benefit from using ML techniques, such as pattern recognition in a noisy environment.

However, not all viable opportunities are valuable. The architect has to further consider if the value proposition from integrating ML is valuable to the enterprise, be it reducing high labor costs, utilizing employee for higher cognitive tasks and improving overall performance. For some enterprises, these opportunities might be more than just valuable or viable. It may be vital for a company to integrate ML into its products/service offering in order to compete with other companies already utilizing ML. It might also be vital for companies wishing to scale up to process huge datasets from big data infrastructure.

ii) Screening Tools

Besides observing the "ML Indicators", screening tools are used to uncover opportunities along the business process, staffing models, data assets and markets areas. Process Maps are used to highlight tasks where ML can have viable or valuable applications, particularly tasks that can be automated without human intuitive skills or can be encoded into a set of decision rules. Staffing Models are used to identify roles where cognitive skills are underutilized for the job but where ML can help to provide automation. Data investigation helps to uncover data that may be under-analyzed or un-exploited. Lastly, market analysis helps to map out the competitive landscape and find opportunities to differentiate the company from other competitors. Examples of screening tools for each of the screening levels and ML indicators are shown in Table 6.3.

Table 6.3 : Viable, Valuable, Vital (3V) Framework . Adapted from (Schatsky et al., 2015)

Screen	ML indicators	Example Screening Tools
Viable	Task, job or workflow that requires low or moderate level of human skill plus perception	Process Map
	Human expertise can be expressed as a set of rules (using decision trees	Process Map
	Large datasets available	Data Investigation
Valuable	Employee's cognitive abilities are under-utilized	Staffing Model
	Business process has high labor costs	Staffing Model
	Expertise is scarce; value of improved performance is high	Staffing Model, Market Analysis, Process Map
Vital	Competitive landscape requires use of ML technologies	Market Analysis
	Products/Services cannot scale by relying on human labor exclusively	Staffing Model, Process Map

6.2.3 Evaluating Technology Adoption

The third part of the ML Integration framework is an Evaluation of the current state of Technology Adoption in the enterprise. This evaluation is an adaptation of an existing Technology Adoption Model to ML, where it enhances the architect understanding of the factors that affect the enterprise decision to adopt a give technology/innovation, in a holistic and systematic fashion. Furthermore, the evaluation enables the architect to make a comparison of important attributes against existing technologies as well as identification of areas where improvements are required. These are then the areas that the architect should focus when designing the ML system in the next and final part of the framework.

While many variations of Technology Adoption Models exist, the Technology, Organization, Environment (TOE) Framework (Tornatzky & Fleischer, 1990) is chosen and adapted for ML applications, shown in Figure 6.2. This is because the TOE framework evaluates the enterprise in a system-thinking fashion by considering internal as well as external factors that influence the process by which an enterprise adopts and implements a technological innovation. The three aspects are: i) Technological Context, ii) Organizational Context, and iii) Environmental Context. Technological context describes both the external and internal technologies relevant to the enterprise, including current practices and infrastructure available to the enterprise. Organizational context refers to the attributes relating to the enterprise such as structure, size and processes. Environmental context is the area in which the enterprise conducts its business, where it maps out the industrial trends and competitive landscape.

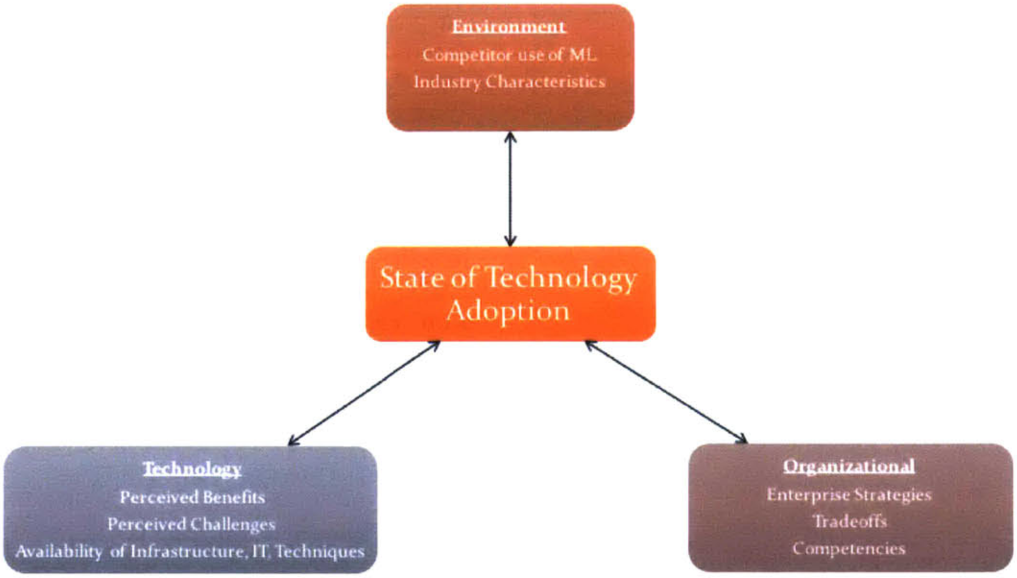

Figure 6.2 : Technology, Organizational, Environment (TOE) Framework. Adapted from (Tornatzky & Fleischer, 1990)

The TOE Framework is thereafter adapted to ML through the following questions in Table 6.4 that guides the architect in evaluating the current state of ML adoption in their enterprise, providing focus on areas when designing the ML system in the next section.

Table 6.4 : Questions to evaluate state of ML adoption using TOE Framework

Aspects	Attributes	Questions
Technology	Perceived Benefits	• How does the enterprise view the benefits of ML in terms of better products, processes, insights, in context of the industry that the enterprise is operating in? • What are their expectations in terms of measurable value?
	Perceived Challenges	• How does the enterprise view the risk and uncertainty arising from using the insights/automation form ML? • How much trust would they place on automated decision making? • In what ways is ML not suited for the industry that the enterprise is in? What are the limitations?
	Availability of Infrastructure, IT, Techniques	• Is the current enterprise data management sufficient to support ML? • Are ML tools sufficient and available for the enterprise to incorporate into existing systems? • What are the perceived integration challenges?
Organizational	Enterprise Strategies	• How does the enterprise regard the value propositions of ML? Is it Viable, Valuable and Vital?

		• Is there a vision of how the enterprise intend to integrate ML? What does it take for the enterprise to execute this vision, e.g. organizational transformation?
	Tradeoffs	• What existing system exist and what is the cost-benefit analysis of replacing/augmenting them with ML • What are the expected benefits expressed in terms of performance metrics, e.g. costs, reliability?
	Competencies	• What level of competencies do employees exhibit with regard to ML techniques as well as understanding them? • What level of competencies do employees demonstrate with regard to implementing ML algorithms, interpreting the results and assessing the quality of these results? • What is the staffing strategy and what are the staffing plans to hire ML experts?
Environment	Competitor use of ML	• How have competitors used and deploy ML to gain a competitive advantage? • What is the performance and limitations from the perspective of these competitors?
	Industry Characteristics	• Is there a trend in the industry to use ML? (e.g. financial fraud detection) • Is it a localized or global trend? (e.g. only limited in North America)

6.2.4 Architecting of ML System

The Architecting of the ML System is the final part of the ML Integration Framework. The previous three parts of the ML Integration Framework: Review of the enterprise model, Identification of the opportunities and Evaluation of the technology adoption formulates the major upstream influences of the ML System architecture, where they support the architect in executing arguably the most important portion of the ML Integration framework. The role of the architect at this point is to reduce ambiguity arising from stakeholder needs, employ creativity to think of concepts of the architecture and manage the complexity of the architecture for subsequent design, implementation and operation. Readers are encouraged to refer (Crawley et al., 2015) for more information on the tools and methodologies used in this part of the framework:

In this part of the framework, the architect conceives and designs the architecture of the ML system according to stakeholder needs. This is done through four activities: i) Understanding Stakeholder Requirements, ii) Conceptualizing the system architecture that integrates ML, iii) Selecting the architecture, and iv) Designing the System from systems, process, organizational and people perspectives. Each of these four activities is detailed below.

The framework concludes at the designing stage and does not proceed with the implementation and operating stage of the entire Product Development Process (PDP)[3] which is left for the architect to deliberated based on specific nuances of the enterprise. This is because the framework is unable to provide a set of definitive guidelines to cover the vast diversity in enterprise hardware and software systems and its processes.

6.2.4.1 *Understanding Stakeholder Requirements*

In this activity, the architect conducts an analysis of the stakeholders needs in a tabulated format as shown in Table 6.5, illustrated with some examples. Stakeholders can be classified as external or internal. External stakeholders are usually customers of the enterprise although regulators, NGOs etc. could be involved in the analysis. The internal stakeholders can comprise investors, management, employees etc. The needs or expectations from stakeholders are typically obtained through interviews as well as from the architect's intuition, although other needs identification methods exist, such as voice of the customer. The architect will thereafter prioritize these needs/expectations according to a set of predetermined criteria (not shown), which will thereafter set the requirements and constraints of the architecture of the ML system.

Table 6.5 : Stakeholder Requirements

Relative to Organization	Stakeholders	Needs/Expectations
Internal	Management	• Meeting business goals • Compliance with policies
	Administrators	• More efficient business processes • Better tracking and monitoring
	Analysts	• Less routine work • Support, teamwork
External	Investors	• Expected ROI • Mitigated risk
	Customers	• Expected product quality • Quality for Cost
	Vendors	• Well defined requirements • Contractual agreements

[3] PDP process consists of CDIO stages: Conceive, Design, Implement, Operate

From Table 6.5, the architect will express a System Problem Statement (SPS) in a "To-By-Using" format which indicates what the overall ML system is intended to do in order to bring value/success to the enterprise. The "Using" portion will only be filled once the System Architecture is selected.

System Problem Statement

To: [Statement of Intent and solution neutral function]

By: [Statement of Function and solution specific]

Using: [Statement of Form, after architecture is selected]

The SPS also indicates the scope and boundary of the problem that must be humanly solvable. The SPS is further constrained by a set of goals each comprising metrics and targets, that the architect plans to be accomplished in order to meet stakeholder needs. It is prioritized by "must", "should" or "shall" model verbs in decreasing priority. The goals must meet the criteria of humanly solvable, complete, consistent, attainable and representative (Crawley et al., 2015). Table 6.6 illustrates some examples of these system goals that are specific to the ML architecture.

Table 6.6 : Example System Goals

Critical	Must cost less than $1m to implement
	Must output accurate prediction to > 80%, 80% of the time
	Must be fully integrated with existing IT systems without requiring changes to existing IT systems
Important	Should learn on-line with new data
	Should take in data sources from silo databases in the enterprise
	Should operate with at most 5 operators/analysts
Desirable	Shall provide predictive information on a common portal
	Shall have regular patches when vulnerabilities are found
	Shall be maintenance free for a period of 5 years

6.2.4.2 Conceptualizing ML System Architecture

Once the stakeholder needs are defined as the SPS and a set of tractable goals, the architect uses his creativity and domain knowledge to conceive the system architecture to "solve" the problem. While doing this task, it is useful to take a solution-neutral approach so that a more diverse set of concepts can be conceived without being fixed at a particular form.

First, the architect starts off by representing the generic ML system as part of the "Whole Product System" using the Object-Process Diagram (OPD) format as shown in Figure 6.3 (Crawley et al., 2015). This would help the architect understand the system boundaries between the ML System and interactions with other related systems that need to be accounted for. In this instance, the ML System is represented as a component system that have interfaces with other accompanying systems such as Data Mining, Big Data Databases and Report Management as shown by the connecting lines in Figure 6.3. These systems form the "Whole Product System" which refers to the Data Analytics System within the general use context of an "Analytical Organization" (refer to Table 6.1 for details on "Analytical Organization").

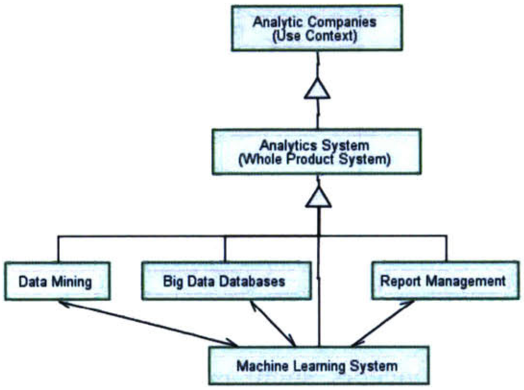
Analytic Companies (Use Context)
Analytics System (Whole Product System)
Data Mining
Big Data Databases
Report Management
Machine Learning System

Figure 6.3 : Whole Product System (Data Analytics System)

Second, after mapping out the "Whole Product System" and understanding the interfaces and use context, the architect will focus on the ML System and map out its functions and form. The externally delivered function that brings value to the ML System is: "Machine-Learning from Data", as represented in OPD format in Figure 6.4. "Machine-Learning from Data" is also the solution-neutral statement as it does not indicate the form needed to execute this function. The externally delivered function consist of the process of "Machine-Learning" (represented by an ellipse and can be further specialized into solution specific processes such as Predicting, Detecting, Classifying, Recognizing) that acts on the value-related operand of "Data" by converting the operand state from "No insight" to "With Insight". The form (represented by a rectangle and can be further specialized into solution specific forms such as algorithms, databases and ML software stacks) enables this externally delivered function of "Machine-Learning from Data" to be carried out.

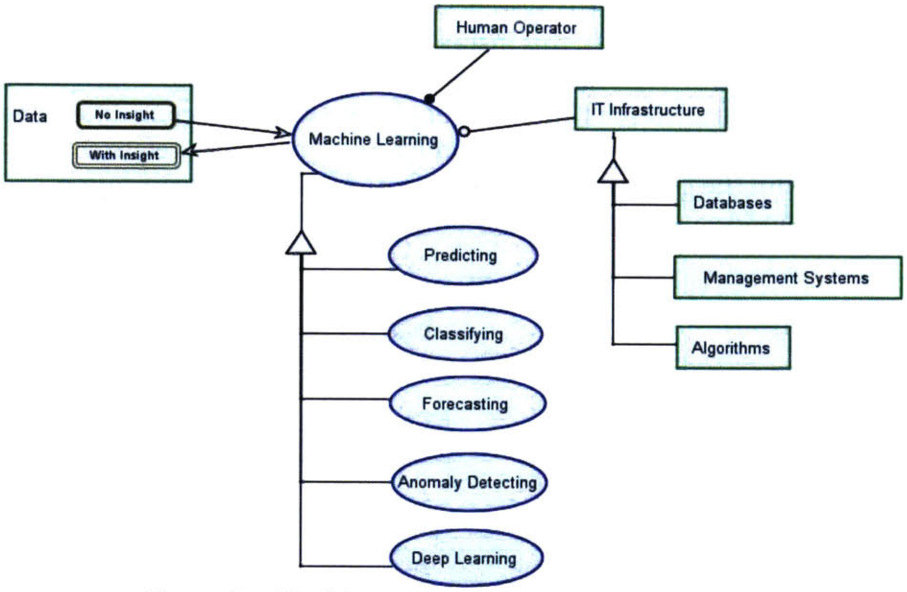

Figure 6.4 : Machine Learning System

Third, after drawing out the concepts such as Predicting Data, Classifying Data (represented by the ellipses and the "Data" rectangle), the architect then focusses on developing the internal value related processes and forms for each of these concepts. These internal value-related process and forms will combine to bring forth the emergent function of that particular concept, such as "Predicting Data".

Figure 6.5 decomposes the concept of "Predicting Data" into constituent internal value-related processes, operands and forms, showing complexities and potential areas of modularity in the formal structure. The process of "Predicting" is decomposed into the constituent processes such as collecting, integrating and preprocessing, as shown in the red box. The operand of "Data" is further specified into its different types as it gets processed, such as collected data and represented data, as shown in the green box. The forms are decomposed into several components, as shown in the boxes to the right of Figure 6.5.

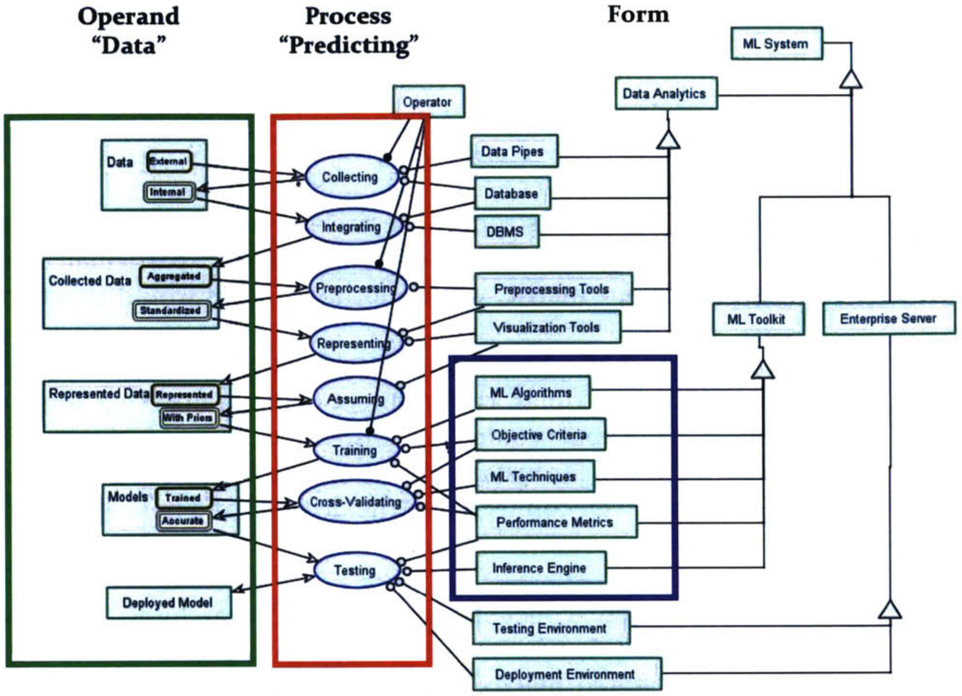

Figure 6.5 : Concept "Predicting Data"

Other concepts of "Classifying Data" or "Forecasting Data" exhibit similar structure as that of Figure 6.5. The key differentiators between these concepts are the type of ML algorithms, objective criteria and performance metrics, highlighted by the blue box. In addition, in the special case of the concept of "Deep Learning", the ML toolkit is more connected with other objects because "Deep Learning" relies formal relationship to big data stores and compute resources, though this is not represented in Figure 6.5.

Nevertheless, Figure 6.5 can be used as a template by the architect to conceive similar concepts to address the SPS and goals identified earlier. It is also at this stage where the architect thinks about the entire product life cycle and also anticipates potential failure modes (e.g. the learner is not performing as expected). These should be reflected as alternatives to the generated concepts.

Finally, for each of the generated concepts, the architect manages the complexity of the processes and forms so as to ensure the goals are delivered and value achieved. One way of managing complexity is by modularizing the various subsystems, as shown in Figure 6.5. Modularization minimizes the interactions between the ML System and other systems such as Data Analytics and Production system, allowing for potential cost reduction in the implementation stage, as well as exploring strategies such as outsourcing instead of in-house designing/building. Another way that the architect can manage complexity is to consider stepwise integration of subcomponents of the ML system, testing at each stage before full integration into the enterprise systems.

6.2.4.3 Selecting Architecture

In the third activity, the architect selects the architecture/ concepts by measuring them against the SPS and goals. The architect evaluates the concepts on how i) Connected and ii) Sensitive they are to the SPS and goals.

i) **Sensitivity** refers to how strongly the architecture decision influence metrics, which are defined by the goals. Most of the time, a particular concept would strongly influence one metric, while another concept would affect another.

ii) **Connectivity** refers to whether substantial rework is required when changing this decision. At other times, the architect might realize that some parts of the whole set of concepts could be abstracted away without affecting the architectural decision, for instance the database portion of a set of ML concepts are similar and not well connected to the decision and hence are not considered in the evaluation.

Therefore, to evaluate these concepts holistically, the architect will then have to conduct a tradespace analysis where these metrics are prioritized and the concepts scored against some weighted matrix of these metrics. The concept is then selected based on the overall score in this analysis. There are other methods for selecting the architecture, such as using a combinatory approach to the metrics or iso-mapping relevant concepts together and indicating the pareto fronts. The architect is encouraged to refer to available literature on this topic.

6.2.4.4 Designing the System (System, Process, Organizational, People Perspectives)

Based on the architecture/concept selected in the previous activity, the architect designs the ML system by looking holistically from the i) System, ii) Process, iii) Organizational and iv) People Perspectives. The guidelines below, while not exhaustive, suggest some ways that the architect can use to execute the selected concept and achieve the goals specified.

i) **System Perspective**

First and foremost, the architect needs to consider the interfaces that the ML system would have with other existing systems. He needs to design the ML system such that these interfaces are minimized to reduce complexity and ensure that the specifications are stated clearly to avoid any ambiguity in the exchange of information between these systems. In the context of a software/hardware solution such as ML systems, the architect would need to pay further attention to the data standards and structures in the information busses between these information systems, lest there be conflict in the systems understanding each other. It is also equally important to investigate the interface between the human operator and the ML system – after all the ML system is a socio-technical system where it requires the domain knowledge and human intuition to fine-tune the parameters of the various sub-systems, particularly those related to the ML toolkit. The architect has to consider user interface designs such that it is intuitive to adjust these parameters, thus abstracting the user away from the complexities underneath these interconnected systems.

Beyond interfaces, the architect also needs to consider the long term value of the ML system beyond what stakeholders might need at the beginning. These are encapsulated as the "Lifecycle Properties" of the system (System Perspective). There are three commonly desired lifecycle properties for ML systems: a) Reliability, b) Scalability and c) Flexibility. For certain applications, d) Safety becomes an important property to be considered.

a) *Reliability*, although traditionally meaning the ability to perform intended function without failure for given period of time, here refers to the reliability of the insight (prediction, detection, forecasting) from the deployed ML system. The lifecycle property of reliability (of insight) is highly dependent on the ability of the model to generalize from training examples, i.e. reliability is not always guaranteed in all production environment. There could be scenarios that the model cannot predict accurately because the trained model was not able to learn to infer from existing training examples. Therefore, it is critical for the architect to design the system such that the model is tested regularly for reliability across dynamic changes in data.

b) *Scalability* refers to the ML System's ability to scale up to ingest larger volume of data and at a faster processing speed. Increasingly, customers desire faster prediction speeds and almost instantaneous operations, such as in the detection of banking fraud. The ML system must be able to scale up to interpret multi-dimensional data from volumes of enterprise information and at near real-time speeds in certain application domains. The faster an insight can be obtained, the faster the enterprise can act on the insight and make a decision, and the more competitive advantage that it can offer to its customers, for example how Netflix operates almost real-time in its movie recommendation system.

c) *Flexibility* refers to the ability of the ML System to adapt to different types of data as well as concepts. The enterprise's requirements will change over time and product offerings evolve. It is critical for the ML System to be able to change its algorithms to provide the insight that is relevant to the desired outcome. For instance, a company may decide to deploy a virtual voice assistant for its website, over and above the predictive analytics that it already performs at the backend. The ML System must be flexible enough to extend its predictive capabilities to feed into the virtual voice assistant without major rework.

d) *Safety* often plays a role in ML systems. In certain applications, such as self-driving cars, ML is a critical technology that enables the application. The architect needs to consider safety analyses or risk assessment methods to ascertain how much to trust the prediction of the ML system versus the safety levels that can be tolerated depending on the identified safety hazards.

ii) **Process Perspective**

The architect needs to design business processes for users of the ML System (users are commonly referred to as Data Analysts or Scientist) that is aligned to business outcomes, so as to derive value from the investment in ML integration. While the process flow in Figure 6.6 articulates the operational process for ML, it still needs to be connected to business outcomes. One useful framework is the Discovery, Insights, Actions, Outcomes (DIAO) Framework defined by *PWC*, where it links the operational processes to business outcomes (Blase & Rao, 2015).

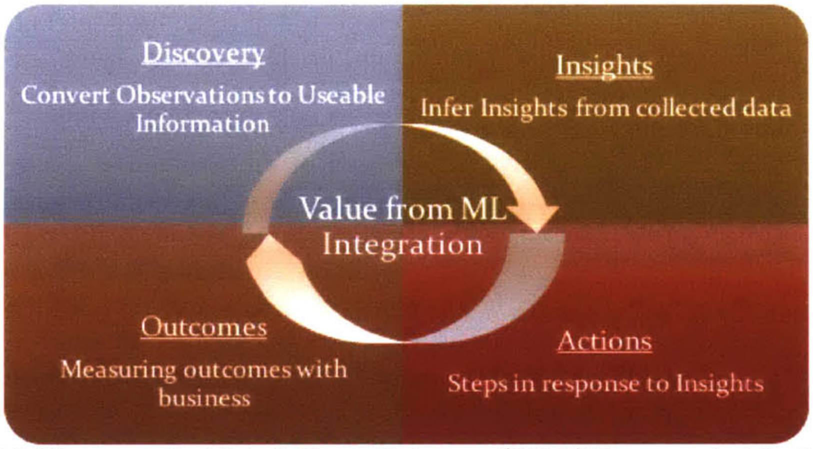
Discovery
Convert Observations to Useable Information
Insights
Infer Insights from collected data
Value from ML Integration
Outcomes
Measuring outcomes with business
Actions
Steps in response to Insights

Figure 6.6: Discovery, Insights, Actions, Outcomes (DIAO) Framework. Adapted from (Blase & Rao, 2015)

a) *Discovery* converts observations into useable information. It requires the establishment of the scope of value that the company wants to inform by exploring new types of data to gain business insights. The scope of value specifies the type and amount of data that should be collected and maintained, while taking into account data usage, privacy and security issues. Discovery is related to the "Data Analytics" components in the Concept "Predicting Data" shown in Figure 6.5.

b) *Insights* is the application of ML techniques to infer insights from the collected data. Insights is based on the stakeholder needs that is expressed as SPS and goal statements. Insights is related to the rest of the "ML Toolkit" components in the Concept "Predicting Data" shown in Figure 6.5.

c) *Actions* are the steps that the organization takes in response to the insights inferred from data. Insights are not useful unless they help the organization make better decisions or take meaningful actions. The enterprise should integrate these insight-driven actions into existing workflow and processes.

d) *Outcomes* measure the impact of ML integration on value metrics, be it financial outcomes or operational efficiencies. This is predetermined by the management level, so that the return on investment of ML integration can be evaluated. Clearly articulated outcomes also helps to inculcate an analytical culture in the enterprise when people can experience the benefits gained by learning from the data.

If the architect chooses to outsource the implementation of the ML system, especially in smaller businesses, the architect would need to consider vendor selection and how these vendors would fit into the overall organization in terms of processes and system interfaces.

iii) Organizational Perspective

The architect needs to design organizational structures to realize the ML integration. Without individuals or teams with the right skillsets to operationalize the ML system, the expected value of ML would not be realized. These teams can be separate analytics arms of the enterprise where they enable the enterprise as a whole or they could be embedded into functional areas where domain knowledge is required. The selection of the organizational model depends on the skillsets of the ML experts and the domain knowledge required in implementing ML. Several organizational models exist, ranging from the centralized to de-centralized forms. This study will not elaborate on these options; the architect is encouraged to refer to management literature on this topic. Even if the enterprise decides to outsource the implementation of the ML system, the architect needs to consider in-house staff who are well versed in ML (or trained in ML) and familiar with their specific domain areas to translate the business needs into system requirements for the vendor.

iv) People Perspective

Even if the ML System has been integrated into the enterprise from system, process and organizational perspectives, the architect needs to consider the perspectives of stakeholders who are using the ML system, to maximize its full potential and derive the business value. It is equally important to obtain organizational buy-in from everyone involved in the ML capability, from data owners to data users. This requires overcoming "trust" issues when it comes to driving business decisions with data, to overcoming the resistance to share information across silos in the organization – all of which were highlighted in Section 5.3 Business Challenges. A simple model for this is to involve the enterprise leader in creating and communicating a vision that encapsulates ML as a business enabler and how it would transform the enterprise. Subsequently, the architect with the management team will encourage stakeholder participation in using the ML capability, embedding change agents at the appropriate levels where necessary to foster collaboration and encourage adoption across silos. Lastly, the architect can design business prototypes that utilize ML to obtain quick buy-in from staff. These are some aspects of technology change management that the architect would need to address from the people perspective. This study will not elaborate on these options; the architect is encouraged to refer to management literature on this topic.

6.3 Conclusion

This chapter presents the framework for any enterprise that has decided to integrate Machine Learning (ML) into their organization, after they have appreciated the opportunities that ML can bring from the technology and business perspectives in Chapters 4 and 5. This framework is not exhaustive – the architect is encouraged to refer to literature indicated in this chapter or relevant texts for in depth advice. Furthermore, the framework concludes at the designing stage and does not provide guidance for the Implement and Operate stages of the CDIO process because the generalized approach to the framework cannot cater to every nuances of individual companies. It is therefore left to the architect to determine how this framework should be implemented and operated according to the concepts and designs discussed earlier. The framework with the component tools and activities is summarized below. The next chapter will apply this framework to a hypothetical mid-sized hospital.

Table 6.7 : Summary of ML Integration Framework

Parts	Activities	Tables/Figures Referenced
Reviewing Enterprise Model	-	Table 6.1
Identifying Opportunities	Categorizing Opportunities (Products, Processes, Insight)	Table 6.2
	Assessing Opportunities (3V Framework)	Table 6.3
Evaluating Technology Adoption	Technology, Organizational, Environment (TOE) Framework	Figure 6.2
	Questions to evaluate state of ML adoption using TOE Framework	Table 6.4
Architecting ML System	Understanding Stakeholder Requirements	Table 6.5 – 6.6
	Conceptualizing the system architecture	Figures 6.3 – 6.5
	Selecting the architecture	
	Designing the System i) System Perspectives ii) Process Perspectives – Discovery, Insights, Actions, Outcomes (DIAO) Framework iii) Organizational Perspectives iv) People Perspectives	Figure 6.6

Chapter 7 Application of ML Integration Framework

This chapter provides a brief example of how the architect might apply the Machine Learning (ML) Integration framework in the context of a hypothetical mid-size hospital. This example scenario is based from the author's own experience with healthcare professionals in the Singapore hospital ecosystem. While some of the representations in this scenario may not be factually correct from a healthcare professional's standpoint and some of the areas seemed incomplete, the aim of this chapter is to present how one would go about applying the ML integration framework from an architect's perspective.

7.1 Background of Hospital

Hospital A is a mid-sized general hospital with 1,500 beds catering to acute-care, with specializations in all aspects of medicine ranging from pediatrics to oncology. It is staffed by 3,000 healthcare professionals, supported by a team of 30 IT professionals to run their database systems. As a recent strategic IT initiative, the hospital has integrated big data systems so that administrators can access detailed operational data to improve healthcare delivery and cut unnecessary expenses. The CEO has decided that much more can be exploited from the big data that consists of Electronics Medical Records (EMR), Operational data and other sources of collection. He thinks that by integrating ML into existing systems, opportunities such as predictive medical diagnosis, recommender systems for medical treatments, and detection of contagious diseases in the community could be made available to the hospital. In particular, the CEO referred to the advances made medical diagnosis in small specialized diagnostic problems (Kononenko, 2001). Hence, the CEO has directed for the architect (a Team Lead in the IT Department) to conceptualize and design the approach for the hospital to integrate ML.

7.2 Application of ML Integration Framework

7.2.1 Reviewing the Enterprise Model

The architect conducts a series of surveys and interviews with both hospital staff across all ranks and with selected patients. He concluded that the hospital is still "Analytically Challenged" (Refer to Table 6.1), despite the recent (2 years ago) IT initiative to move to big data analytics. This is because a large proportion of the hospital administrators still did not use data to drive their decision making and the expected business outcome of big data analytics was unmet due to lack of talent available to interpret the results in a business/operational context. With this review, the architect assess that he should emphasize on a technology adoption framework of ML and the people perspective of his design.

7.2.2 Identifying the Opportunities

Next, the architect identifies the opportunities presented by ML. From his interviews, he observed that doctors continue to be frustrated with their inability to keep up with new and changing treatment options in their field. They lamented the lack of time they have to read through medical journals and wished for an intelligent assistant to help them make better decisions. Quite similarly, the administrators were concerned about minimizing risk inherent in a stressful hospital environment. Their concerns were due to recent incidents where patients were misdiagnosed by their doctors which lead to several costly damages to the hospital.

The architect then brainstorms ideas based on the Product, Processes and Insight categories (Table 6.2), and assesses these opportunities using the 3V framework (Table 6.3). Using some of the screening tools (in particular the process map) in Table 6.3, he determines that it would be valuable for the hospital to concentrate on minimizing the risk of misdiagnosis by supporting the doctors to make better judgments using ML. The architect believes that by using ML to analyze the large amounts of EMR in the databases, the system would learn automatically how patients are treated in real-life scenarios and therefore allow the doctors to make better decisions, thus improving patients' safety and improve healthcare quality. Advances have already been made in cancer prognosis (Kourou, Exarchos, Exarchos, Karamouzis, & Fotiadis, 2015), as well as in solutions provide by ML supply-side companies such as "Enlitic" in diagnostic healthcare (Salim, 2015). However, at this brainstorming stage, the architect is conscious that he should not arrive at a binding specific solution or idea. He would have to conduct a proper stakeholder requirements analysis, derive the problem statement and goals and architect the system so as to have the best possible solution.

7.2.3 Evaluating Technology Adoption

In the next part of the framework, the architect evaluates the hospital's level technology adoption of ML using the TOE (Technology, Organization, Environment) Framework in Table 6.4. He observed that many staff across multiple levels remain skeptical about the benefits that ML can bring, particularly the doctors where a large majority were dismissive that the recommendations from the ML system was going to be more reliable than human judgment. Nevertheless, when the architect highlighted first the value proposition of the ML System to doctors before surveying their perception of such a technology, an increased percentage of them said that they recognized the benefits in terms of reducing misdiagnosis and were willing to trial such a system. The architect also realizes that while the current enterprise system was sufficient for ML to be integrated, the system had not reached full operating capability of the Big Data IT initiative. Therefore, the development of the ML system has to be bootstrapped with the progress of this Big Data initiative such that any interoperability issues between these two enterprise systems can be ironed out during development.

From the organizational lens, the above observation indicated that the architect and the management team would need to pay more attention to the "People Perspective" when designing the ML system such that the technology would be more readily adopted by key stakeholders. The CEO would need to communicate a clear enterprise strategy that uses ML and devise business outcomes to measure the level of integration. At the operational level, the architect would need to re-design business processes of these doctors such that the ML tools is made available to them, which fosters better adoption when these doctors see the ML system in its physical sense and experience direct benefit in terms of more accurate diagnosis and risk reduction. It is also critical for the top management to articulate that the integration of ML into the healthcare industry remains is an ongoing learning process as there remain certain risk should these vaunted promises of ML not met. That said, the architect would also have to pay more attention to hiring new staff who are well-versed in ML and have had experience in the healthcare domain, to support this integration initiative. He would also have to plan training programs to raise the competency of users and improve adoption.

The architect also conducts a scan of the external environment and assesses that competitor hospitals have yet to embark on utilizing ML to make better clinical diagnosis. One reason for this is that it is difficult to obtain an unbiased estimation of the reliability of the diagnosis (Kononenko, 2001) . He thinks there could be further co-opetition with other hospitals so as to make the system more robust (the more training examples, the better the ML algorithm generalizes) and believes that the Singapore population would stand to benefit from ML-driven clinical diagnosis, as long as the challenges in sharing EMR information between hospitals are addressed. This could be addressed at the policy level through the Health Ministry in future stages.

7.2.4 Architecting of ML System

7.2.4.1 Understanding Stakeholder Requirements

The architect conducts a series of interviews as well as observations of the enterprise in order to map out the stakeholder requirements. He then identifies the key stakeholders and beneficiaries and articulates a System Problem Statement (SPS) that asserts what the ML system does to deliver value. Table 7.1 below is a snapshot of some of the stakeholder requirements, with the key needs highlighted in blue.

Table 7.1 : Stakeholder Requirements Analysis

Relative to Organization	Stakeholders	Needs/Expectations
Internal	Hospital Management Team	Mitigate Risk from Misdiagnosis Complete transformation of IT Compliance with policies Exploiting new technologies to make better decisions using data
	Administrators	More efficient hospital administrative processes Faster way of finding data
	Doctors	More accurate and smarter diagnosis from all available information
	Nurses	System alerts before administering wrong drugs dosage/type
External	Healthcare Regulators	Compliance with policies on EMR and Healthcare
	Patients	Most accurate diagnosis Most effective treatment possible with healthcare insurance Good quality of in-patient care in hospital
	IT Vendors	Well defined requirements

The architect then considers all the above needs, prioritizes them and identifies a single System Problem Statement (SPS) that best captures the primary value delivered to the primary beneficiary. In this case the primary beneficiary are the Patients, the doctors being an intermediary. Since the architect knows that the architectural form has to be a Machine Learning System that uses data, he notes this in the "Using" portion but is cognizant that there may be other specific form design choices that he may have to make. The SPS is therefore described as:

System Problem Statement

To: Diagnose medical conditions accurately

By: Predicting the root cause of patient symptoms

Using: Machine Learning Systems and Data (General)

With the stakeholder requirements, the architect also articulates the following system goals, detailing the metrics and targets that must, should or shall be met with the architecture in Table 7.2. These metrics and targets were agreed upon through the various interviews with stakeholders (the metrics are guesstimates for this hypothetical scenario)

Table 7.2 : System Goals (Hospital)

Priority	Goals
Critical	Must output accurate prediction of medical condition > 95%, 95% of the time
	Must cost less than $0.5m to implement
	Must be fully integrated with existing IT systems and ongoing Big Data initiative
	Must be fully secure and private to meet healthcare data policies
Important	Should be able to learn from existing as well as new, online data (< 100TB per day)
	Should be able to scale to incorporate larger sets of data (>100TB per day)
	Should be flexible in terms of the ML algorithm used
	Should have visualization tools with which to measure and evaluate the ML system
	Should operate with at most 5 IT staff with training
	Should explain the rationale for the decision that can be understood by humans
Desirable	Shall have a single common interface/portal for all doctors to use
	Shall be usable for doctors without need to be significantly trained
	Shall be maintenance free for a period of 5 years

7.2.4.2 Conceptualizing the system architecture

With the SPS and set of goals that sets the boundaries of the architecture, the architect uses his domain knowledge and creativity to conceptualize several ideas. Two concepts "Predicting" and "Deep Learning Prediction" are illustrated in Figure 7.1 and 7.2 with form and function decompositions to show the difference in the concepts. The architect notes that there is similarity in the processes but the forms are drastically different between these two concepts.

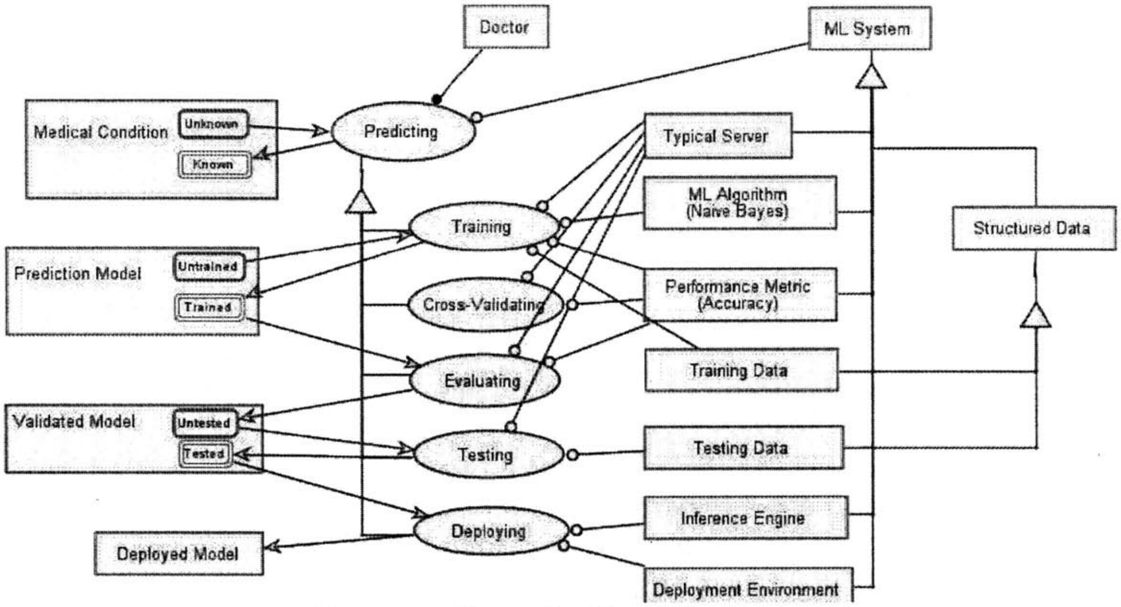

Figure 7.1 : Concept 1 - Predicting

Specifically, he notes that for the Concept of Deep Learning, he would need to pay attention to the instrument objects of GPU processing systems as well as the structured and unstructured data that differentiates the concept from "Predicting". He also notes that not shown in these concepts are the interrelationships between the objects which will need to be factored in during the system design stage. Nonetheless, the externally value related function of "Predicting Medical Condition" remains the same.

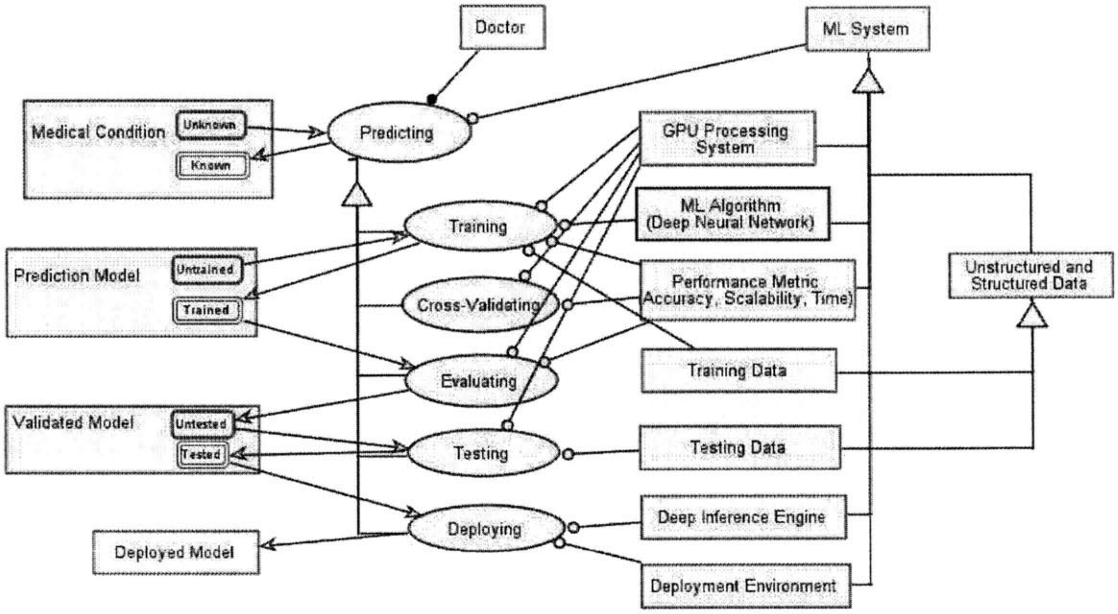

Figure 7.2 : Concept 2 – Deep Learning

7.2.4.3 *Selecting the architecture*

The architect then selects the architecture based on the goals specified earlier. The architect notice that for some architectural decisions such as the ML algorithm, they are strongly connected to and sensitive in response to the goals, in particular the critical ones. For instance, in order to achieve 95% prediction 95% of the time (a critical goal), he would have to select the "Deep Learning Prediction" concept as recent research suggests that it is more accurate than the naïve-Bayes prediction concept for medical diagnosis. At the same time, he would need to balance this concept with the important goal of the system being able to determine the reliability for the prediction made so that doctors can have confidence in the system (one of the key reasons why medical diagnosis using ML has hitherto not taken off). However, for Deep Learning systems, the route to which the prediction is arrived at is not as visible as compared to the naïve-Bayes prediction system. It is difficult for humans to visualize the decision pathways in these Deep Learning algorithm, which could result in the perception that the prediction is not reliable. Thus, the architect would have to communicate the reliability of the Deep Learning system using other metrics for doctors to perceive the prediction as reliable, such as cross-evaluating with other sets of test data and tracking mis-prediction rates.

The architect then uses a scoring system to measure how much influence the concepts have over the metrics, selecting the two or three most important metric. He then selects the concept that scores the best across these metrics and explains the rationale for the selection to his sponsors. In this example, the architect selects the "Deep Learning Prediction" concept for the next stage – Design. This is a descriptive example of the tradespace analysis that the architect would have to make when selecting the right architecture. This example is meant to be illustrative of the creative steps that the architect would take – for in depth examples, the architect is encouraged to refer to literature on this topic of tradespace analysis.

7.2.4.4 *Designing the System*

In this last activity, the architect assesses that it would be useful to get the endorsement from key stakeholders and his sponsors before deep diving into system design. Once the "Deep Learning Prediction" architecture has been endorsed, the architect organizes his design through the following perspectives.

i) **System Perspectives**

First, the architect considers the interconnections that the "Deep Learning Prediction" architecture need to have with the existing IT infrastructure as well as the ongoing big data project. He realizes a need to specify the data format and structures at the interfaces between these systems so that they are interoperable. Furthermore, he endeavors to reduce the complexity of the overall system by deliberating on the number and type of interfaces between these subsystems, in particular how the Deep Learning system interface with the data stores and server racks. He is cognizant of the need to dedicate compute and storage resources for the selected architecture that is based on deep learning. The architect considers the options of building these resources in-house or to utilize cloud computing to process the data. He decides that it would be more secure to implement the system in-house, because of the goal statement on data security and privacy that was specified in the stakeholder requirements. These are critical requirements in the healthcare industry. The architect then finalizes the system infrastructure by charting the "Whole of System" diagram, showing how the ML system interfaces and interacts with the other systems in the hospital. These systems not limited to IT systems but also administration systems and human operators (This is not illustrated).

The architect then considers the "Lifecycle Properties" of the system and decides that safety is of paramount important, closely followed by reliability. The predictions of the "Deep Learning Prediction" architecture must be tested thoroughly before deployment. He conducts a holistic and systemic safety analysis, for instance using System-Theoretic Process Analysis (STPA) (Leveson, 2011) such that hazard areas are identified and mitigating safety measures are taken to prevent occurrence of a wrong prediction. The architect would also map out the risk assessments and mitigating measures of these hazard areas as a form of communication to the stakeholders so that they are aware of the level of risk inherent in these prediction systems for healthcare. The process perspective also need to cater for the potential inaccuracies in the ML system by ensuring that the users (doctors in this example) make the final diagnosis, while supported by the ML system.

Closely related to safety is reliability, where the architect needs to ensure that the system is designed to be reliable across all sorts of data. He does this by cross-validating the trained model with new sets of data and measuring the accuracy in prediction. He also designs for a select group of doctors to validate the reliability of prediction at predefined times of the year. Lastly, on scalability, the architect decides that the EMR is sufficient to train a model that is based on Deep Learning. The data that is available in other formats, such as the doctor notes and doctor-patient communication, is not easily collected and could be staged later for collection when handwriting and speech to text recognition technologies have matured. Nevertheless, he has to contend with the possibility of the ML system to scale up to incorporate such unstructured data in the new future and design the system to address such possibilities.

ii) **Process Perspectives**

The architect then links the operational processes in the ML System to the business outcomes using the DIAO framework, explaining his observations in the Notes column according to Discovery, Insight, Actions and Outcome factors in Table 7.3. He connects both the SPS and system goals to the value-related functions of the selected system architecture in order to gain insights to fill the DIAO framework, which is a useful tool to communicate the value of the ML system chosen. In this example, the architect notes the following:

Table 7.3 : Applying DIAO Framework

DIAO	Notes
Discovery	Predicting medical conditions based on past examples of diagnosis, incorporating EMR and other relevant records into the prediction model, while keep patient confidentiality and data secure. This leverages data-driven medical decision making based on past successful examples of diagnosis and treatment.
Insight	More accurate medical diagnosis which help doctors make informed judgments without human error.
Actions	Hospital can reduce the number of misdiagnosed cases and improve overall healthcare delivery.
Outcome	More accurate diagnosis, improved patient outcomes, better reputation from data-driven decisions.

iii) **Organizational Perspectives**

The architect realizes that the IT team that he has does not have the right competencies to implement the ML system in house. Therefore, he recommends employing vendors to help implement the system in several stages, starting with a trial and eventually transferring the knowledge to the IT department. The architect identifies that the General Medicine department would benefit the most from such a trial and decides to implement the system in the department first to gain familiarity before considering hospital-wide deployment.

He then works with the administrators and doctors in the General Medicine department together with the vendors so that the domain knowledge can be captured into the ML system and work out the procedures that need to be followed when doctors interact with the deployed system. At the same time, he ensures that sufficient knowledge is transferred to the IT department so that they could implement the system in other departments at later stages.

iv) **People Perspectives**

Based on his surveys and interviews, the architect acknowledges that people's perception of the ML system would need to evolved in order to realize the potential of such technology. These was highlighted when the architect was evaluating the state of technology adoption in the hospital. To obtain buy-in from the key stakeholders, he engages the senior management of the hospital to articulate a vision where data-driven decision making is the norm for both medical diagnosis and administration, and the ML system is a key technology enabler. Through a series of internal emails and newsletter, the architect would also highlight successful cases of how the General Medicine department has successfully achieved more accurate diagnoses that has led to the business outcome of more quality healthcare for patients. These communication strategies will help to instill trust from both internal and external stakeholders on the insights gained from such artificial intelligent systems such as ML.

7.3 Conclusion

This brief example shows how an architect would review the state of his hospital in terms of analytical capability, identify and assess opportunities, evaluate the state of technology adoption and conceptualize/design architectures of ML systems that would deliver the expected benefit articulated by stakeholder that is encapsulated in the SPS. This example is not exhaustive as its main intent is to show the steps in applying the framework proposed in Chapter 6.

Chapter 8 Conclusion

Machine Learning (ML) is a diverse and dynamic scientific field that yields much potential for organizations that desire to derive insights into their data and/or yield predictive products and services. Nevertheless, the literature review in Chapter 2 observes that many managers and employees are confused over the benefits and challenges of ML because ML is a complex and broad scientific discipline that does not lend itself to easy understanding, particularly when there is lack of in-house talent to champion and communicate these concepts. There is also a lack of adoption frameworks and practical guidelines for understanding and thereafter integrating ML into organizations.

As such, the thesis has provided an overview of ML from both technical and business perspectives in Chapters 4 and 5 that is intended for the decision maker to appreciate the opportunities and challenges before embarking on such analytical transformation. The thesis then proposes the ML Integration framework in Chapter 6 which guides the System Architect in deciding and architecting ML systems into their businesses in order to build the technology capabilities to meet their stakeholder needs. The ML Integration framework is illustrated with a hypothetical scenario of a mid-size hospital in Chapter 7.

Section 8.1 summarizes the key findings in the Chapters with respect to the four research questions RQ 1 – 4 which had motivated this thesis. Section 8.2 highlights the significance of this thesis and recommends future areas of work.

8.1 Research Questions and Findings

RQ1: Literature Review

Chapter 2 addresses RQ1 by reviewing the state of ML adoption and examining the reasons for such. It is observed that the confusion over ML with regard to its terminology, benefits and challenges from both technology and business perspective stem from lack of a single body of knowledge of ML that is communicated at the decision-making level. Even then, there is also a lack of adoption frameworks and practical guidelines for organizations to integrate ML successfully. These observations subsequently motivate RQ 2 – 4.

RQ2: ML from Technology Perspective

Chapter 4 addresses RQ2 by presenting an overview of the definition and concepts, landscape and techniques of ML that are categorized for ease of reference. In addition, the chapter outlines the application domains and challenges in the application of ML, followed by a brief on the current progress in ML in contrast with similar technologies. This overview therefore gives the reader an easy to understand overview of ML from the technology perspective without diving in mathematical and scientific details and is aimed at addressing the confusion highlighted in the literature review.

RQ3: ML from Business Perspective

Chapter 5 addresses RQ3 by first presenting the ecosystem of ML and then highlighting the key trends and challenges faced by businesses in operationalizing ML, based on surveys that has been conducted by other researchers as well as the author's experience. Theses challenges seemingly stem from multiple aspects of the organization, from leadership to infrastructure. Therefore, there is a need for a holistic approach to integrate ML system into the organization, as detailed in Chapter 6.

RQ4: ML Integration Framework

Chapter 6 addresses RQ4 by detailing the ML Integration framework so that architects can holistically and systematically integrate Machine Learning (ML) into their organization, after they have appreciated the opportunities and challenges of ML from the technology and business perspectives in Chapters 4 and 5. The four parts of the framework are: i) Reviewing the Enterprise Model, ii) Identifying the Opportunities, iii) Evaluating Technology Adoption and iv) Architecting of ML System. The framework starts at the conceptualization stage and ends before the implementation stage (of the CDIO framework), leaving for the architect to consider the options available during implementation and operationalization. Chapter 7 subsequently shows how the ML Integration framework is applied in a hypothetical mid-sized hospital, transforming it into an "analytical innovator" that employs ML to diagnose medical conditions based on data, in combination with doctor's experience and intuition.

8.2 Significance and Future Areas of Work

This thesis is aimed at providing the decision maker, manager or architect with sufficient knowledge to appreciate, consider and implement ML into their organizations. To the extent that it has done so by integrating disparate sources of information into a single body of knowledge so as to clear confusion and the hype surrounding ML, the thesis does not break new ground in terms of ML scientific research because the focus is on the application of ML to organizations.

Similarly, while the thesis combines several frameworks and customizes it specifically in the ML Integration framework, it does not propose any significant advancements in management frameworks. Rather, it has adapted management frameworks in the context of ML, which has hitherto has not been achieved. That said, the interpretative research methodology in this thesis is hampered by the lack of quantitative surveys – interpretations based on third party surveys have had to be made in order to arrive at the findings and observations in Chapters 2, 4 and 5.

That said, there are three areas of future work to advance the body of work in this thesis. In particular, the ML Integration framework needs to be tested in a real world scenario and reviewed for its applicability. Case studies of organizations that have implemented this framework will be useful addenda to the framework by providing domain-specific examples of ML integration. Second, qualitative surveys should be conducted to observe the state of ML adoption in companies globally and examine the functional areas where

ML makes the most impact, as well as account for geographical/economical differences. Lastly, Chapters 4 and 5 can be expanded into a larger and deeper set of guidelines specific to application domains such as recommender systems, predictive analytics or recognition systems. Practitioners can therefore refer to these guidelines for layman view of ML in their specific domain areas without delving into the math and science driving these algorithms and systems.

Bibliography

Aradhye, H., Hua, W., & Lin, R. S. (2013). Native machine learning service for user adaptation on a mobile platform. Google Patents. Retrieved from http://www.google.com/patents/US8429103

Austin, T. (2016). *Smart Machines Primer for 2016.* Retrieved from https://www.gartner.com/doc/3242223/smart-machines-primer-

Bellman, R. E. (2003). *Dynamic Programming.* Dover Publications, Incorporated.

Beyer, D. (2016). *The Future of Machine Intelligence: Perspectives from Leading Practitioners. O'Reilly Media Inc.* Retrieved from http://www.oreilly.com/data/free/future-of-machine-intelligence.csp

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning. Pattern Recognition* (1st ed., Vol. 4). New York: Springer-Verlag. Retrieved from http://www.library.wisc.edu/selectedtocs/bg0137.pdf

Blase, P., & Rao, A. D. (2015). *Data and Analytics: Creating or destroying shareholder value?* Retrieved from https://www.pwc.com/us/en/analytics/publications/assets/pwc-data-analytics-creating-or-destroying-shareholder-value.pdf

Brant, K., & Austin, T. (2015). *Hype Cycle for Smart Machines, 2015. Gartner.* Retrieved from https://www.gartner.com/doc/3099920/hype-cycle-smart-machines-

China rolls out three-year program for AI growth. (2016, May). *Xinhua.* Retrieved from http://english.gov.cn/state_council/ministries/2016/05/23/content_281475355720632.htm

Clark, J. (2015). I'll Be Back: The Return of Artificial Intelligence. Retrieved June 27, 2016, from http://www.bloomberg.com/news/articles/2015-02-03/i-ll-be-back-the-return-of-artificial-intelligence

Crawley, E., Cameron, B., & Selva, D. (2015). *Systems Architecture: Strategy and Product Development for Complex Systems.* Pearson.

Daume, H. (2015). A course in machine learning. Retrieved from http://ciml.info/dl/v0_9/ciml-v0_9-all.pdf

Dhar, V. (2012). Data Science and Prediction. *Communications of the ACM, 56*(12), 64–73. http://doi.org/10.2139/ssrn.2086734

Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM,* 55(10), 78. http://doi.org/10.1145/2347736.2347755

Duncan, A. D., Linden, A., Koehler-Kruener, H., Zaidi, E., & Sharma, S. (2015). *Market Guide for Text Analytics.* Retrieved from http://www.gartner.com/document/3178917

Dunning, T., & Friedman, E. (2014). *Practical Machine Learning: A New Look At Anomaly Detection.* O'Reilly.

Economist Intelligence Unit. (2014). Gut & gigabytes: Capitalising on the art & science in decision making. Economist Intelligence Unit. Retrieved from https://www.eiuperspectives.economist.com/sites/default/files/Gut_&_gigabytes_Capitalising_on_the_art_&_science_in_decision_making.pdf

Friedman, E. (2015). Machine Learning at American Express: Benefits and Requirements. Retrieved July 23, 2016, from https://www.mapr.com/blog/machine-learning-american-express-benefits-and-requirements#.VhXEThNVhBc

Gadkari, S., & Mohan, A. (2014). Using Big Data for Machine Learning Analytics in Manufacturing. *Tata Consulting Services*. Retrieved from http://www.tcs.com/SiteCollectionDocuments/White Papers/Machine-Learning-Analytics-in-Manufacturing-0714-1.pdf

Grush, L. (2015, July). Google engineer apologizes after Photos app tags two black people as gorillas. Retrieved June 27, 2016, from http://www.theverge.com/2015/7/1/8880363/google-apologizes-photos-app-tags-two-black-people-gorillas

H2o. (2016). H2o.ai - Financial Services. Retrieved from http://www.h2o.ai/verticals/financial/

How Small Businesses Can Scale the Big Data Barrier. (2014). *Nielson*. Retrieved from http://www.nielsen.com/us/en/insights/news/2014/how-small-businesses-can-scale-the-big-data-barrier.html

IBM Acquires AlchemyAPI. (2015, March). Retrieved June 27, 2016, from https://www-03.ibm.com/press/us/en/pressrelease/46205.wss

IResearch. (2015). 2015 China' s Artificial Intelligence Report. Retrieved from http://www.iresearchchina.com/content/details8_20728.html

Japkowicz, N., & Shah, M. (2011). *Evaluating Learning Algorithms: A Classification Perspective*. New York, NY, USA: Cambridge University Press.

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning:Trends, perspectives, and prospects. *Science, 349*(6245).

Kart, L., & Heudecker, N. (2015). *Survey Analysis: Practical Challenges Mount as Big Data Moves to Mainstream. Gartner*. Retrieved from http://www.gartner.com/newsroom/id/3130817

Kiron, D., Prentice, P. K., & Ferguson, R. B. (2014). The Analytics Mandate. *MIT Sloan Management Review, 55*(4), 1.

Kolbjørnsrud, V., Amico, R., & Thomas, R. J. (2016). *The Promise of Artificial Intelligence: Redefining management in the workforce of the future. Accenture.* Retrieved from https://www.accenture.com/_acnmedia/PDF-19/AI_in_Management_Report.pdf#zoom=50

Kononenko, I. (2001). Machine learning for medical diagnosis: history, state of the art and perspective. *Artificial Intelligence in Medicine, 23*(1), 89–109. http://doi.org/10.1016/S0933-3657(01)00077-X

Kourou, K., Exarchos, T. P., Exarchos, K. P., Karamouzis, M. V., & Fotiadis, D. I. (2015). Machine learning applications in cancer prognosis and prediction. *Computational and Structural Biotechnology Journal, 13*, 8–17. Retrieved from http://dx.doi.org/10.1016/j.csbj.2014.11.005

Lehmann, C., Roy, K., & Winter, B. (2016). *The State of Enterprise Data Quality: 2016.* Retrieved from http://siliconangle.com/files/2016/01/Blazent_State_of_Data_Quality_Management_2016.pdf

Leveson, N. G. (2011). *Engineering a Safer World: Systems Thinking Applied to Safety. The MIT Press*. The MIT Press.

Linden, A., Tohamy, N., Kart, L., Patrick, C., Jacobson, S. F., & Cappelli, W. (2015). Machine Learning Drives Digital Business, (December). Retrieved from https://www.gartner.com/doc/2820120/machine-learning-drives-digital-business

Linn, A. (2015). How's the weather? Using artificial intelligence for better answers.

Retrieved July 23, 2016, from http://blogs.microsoft.com/next/2015/08/10/hows-the-weather-using-artificial-intelligence-for-better-answers/#sm.000186uhfpo9bebwv1g2rpcfby6w4

Lipton, Z. C. (2016). The Mythos of Model Interpretability. Learning; Artificial Intelligence; Computer Vision and Pattern Recognition; Neural and Evolutionary Computing; Machine Learning. Retrieved from http://arxiv.org/abs/1606.03490

McWilliams, A. (2016). Smart Machines: Technologies and Global Markets. BCC Research. Retrieved from http://www.bccresearch.com/market-research/instrumentation-and-sensors/smart-machines-tech-markets-report-ias094b.html

Medhat, W., Hassan, A., & Korashy, H. (2014). Sentiment analysis algorithms and applications: A survey. *Ain Shams Engineering Journal*, 5(4), 1093–1113. http://doi.org/10.1016/j.asej.2014.04.011

Metz, C. (2016). The Rise of the Artificially Intelligent Hedge Fund. Retrieved July 23, 2016, from http://www.wired.com/2016/01/the-rise-of-the-artificially-intelligent-hedge-fund/

Mitchell, T. M. (1997). *Machine Learning. Machine Learning* (Vol. 1). http://doi.org/10.1007/BF00116892

Moser, G. (1990). In Memoriam, Arthur Samuel: Pioneer in Machine Learning. *AI Magazine (AAAI)*, *11*(3), 10–11.

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. The MIT Press.

Narrative Science. (2015). *State of Artificial Intelligence & Big Data in the Enterprise*. Retrieved from http://resources.narrativescience.com/i/526857-2015-state-of-ai-and-big-data-in-the-enterprise/3

NASA. (2007). NASA Systems Engineering Handbook. *Systems Engineering*, *6105*(June), 360. http://doi.org/10.1016/0016-0032(66)90450-9

Neuman, W. L. (2006). *Social Research Methods: Qualitative and Quantitative Approaches*. Pearson.

Nightingale, D. J., & Rhodes, D. H. (2015). *Architecting the Future Enterprise*. MIT Press. Retrieved from http://www.jstor.org/stable/j.ctt17kk86j

Pamela Vagata, & Kevin Wilfong. (2014). Scaling the Facebook data warehouse to 300 PB. Retrieved June 27, 2016, from https://code.facebook.com/posts/229861827208629/scaling-the-facebook-data-warehouse-to-300-pb/

PwC. (2015). *Great expectations: The evolution of the chief data officer*. Retrieved from https://www.pwc.com/us/en/financial-services/publications/viewpoints/assets/pwc-chief-data-officer-cdo.pdf

Pyle, D., & San Jose, C. (2015). An executive's guide to machine learning. Retrieved July 23, 2016, from http://www.mckinsey.com/industries/high-tech/our-insights/an-executives-guide-to-machine-learning

Ransbotham, B. S., Kiron, D., & Prentice, P. K. (2016). *Beyond the Hype : The Hard Work Behind Analytics Success*. Retrieved from http://sloanreview.mit.edu/projects/the-hard-work-behind-data-analytics-strategy/

Ransbotham, S., Kiron, D., & Prentice, P. K. (2015). *The Talent Dividend. MIT Sloan Management Review* (Vol. 56). Retrieved from http://sloanreview.mit.edu/projects/analytics-talent-dividend/

Raphael, C. (2015). Netflix Recommendations: How Algorithms Keep Customers Watching. Retrieved July 23, 2016, from https://www.rtinsights.com/netflix-recommendations-machine-learning-algorithms/

Russell, S. J., & Norvig, P. (2010). *Artificial Intelligence: A Modern Approach. Artificial Intelligence*. Prentice Hall.

Salim, A. (2015). Deep learning in diagnostic healthcare: The future? Retrieved July 23, 2016, from http://www.idgconnect.com/abstract/9573/deep-learning-diagnostic-healthcare-the-future

Schatsky, D., Muraskin, C., & Gurumurthy, R. (2015). Cognitive technologies: The real opportunities for business. *Deloitte Review*, (16), 114–129. Retrieved from http://www2.deloitte.com/tr/en/pages/technology-media-and-telecommunications/articles/cognitive-technologies.html

Simonite, T. (2016). Microsoft and Google Want to Let Artificial Intelligence Loose on Our Most Private Data. Retrieved June 27, 2016, from https://www.technologyreview.com/s/601294/microsoft-and-google-want-to-let-artificial-intelligence-loose-on-our-most-private-data/

Singh, Y. K., & Bajpai, A. B. (2008). *Research Methodology: Techniques and Trends*. APH Publishing.

Sinha, V., & Wegener, R. (2013). *The Value of Big Data: How Analytics Differentiates Winners. Bain & Company*. Retrieved from http://www.bain.com/publications/articles/the-value-of-big-data.aspx

Tornatzky, L. G., & Fleischer, M. (1990). The processes of technological innovation. *The Journal of Technology Transfer*, *16*(1), 45–46. http://doi.org/10.1007/BF02371446

Turing, A. (1950). Turing. Computing machinery and intelligence. *Mind*, 59(236), 433–460. Retrieved from papers2://publication/uuid/E74CAAC6-F3DD-47E7-AEA6-5FB511730877

van Leeuwen, J. (2004). Approaches in Machine Learning. In W. F. J. Verhaegh, E. Aarts, & J. Korst (Eds.), *Algorithms in Ambient Intelligence* (pp. 151–166). Dordrecht: Springer Netherlands. http://doi.org/10.1007/978-94-017-0703-9_8

Vinyals, O., Toshev, A., Bengio, S., & Erhan, D. (2015). Show and tell: A neural image caption generator. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Vol. 07–12–June, pp. 3156–3164). http://doi.org/10.1109/CVPR.2015.7298935

Walker, M. J., Cearley, D. W., & Burke, B. (2016). *Top 10 Strategic Technology Trends for 2016: Information of Everything. Gartner*. Retrieved from http://www.gartner.com/document/3227533?ref=TypeAheadSearch&qid=9d9c620102a780dd0d7cb3b21b645562

Weinelt, B., Shah, A., Spelman, M., Ilangovan, S., & Gul, M. (2016). World Economic Forum White Paper Digital Transformation of Industries. World Economic Forum. Retrieved from http://reports.weforum.org/digital-transformation-of-industries/wp-content/blogs.dir/94/mp/files/pages/files/digital-enterprise-narrative-final-january-2016.pdf

Wilson, C. (2015). Machine Learning as a Service: Is It Really Here? Retrieved July 23, 2016, from http://blog.syncsort.com/2015/09/big-data/machine-learning-as-a-service-is-it-really-here/

Wilson, H. J., Alter, A., & Shukla, P. (2016). Companies Are Reimagining Business

Processes with Algorithms. Retrieved April 6, 2016, from https://hbr.org/2016/02/companies-are-reimagining-business-processes-with-algorithms

Wilson, H. J., Sachdev, S., & Alter, A. (2016). How Companies Are Using Machine Learning to Get Faster and More Efficient. Retrieved July 23, 2016, from https://hbr.org/2016/05/how-companies-are-using-machine-learning-to-get-faster-and-more-efficient?utm_content=buffer75919&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer

Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, *1*(1), 67–82. http://doi.org/10.1109/4235.585893

Wood, L. (2016). *Artificial Intelligence (AI) Market By Technology (Machine Learning, Natural Language Processing (NLP), Image Processing, And Speech Recognition), Application & Geography - Global Forecast To 2020. GII Research*. Retrieved from http://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-market-74851580.html

Zilis, S. (2015). The current state of machine intelligence 2.0. Retrieved June 27, 2016, from https://www.oreilly.com/ideas/the-current-state-of-machine-intelligence-2-0