

ARTIFICIAL INTELLIGENCE FOR REGULATORY COMPLIANCE: ARE WE THERE YET?

Tom Butler & Leona O'Brien

The world's first digital computer, ENIAC, would have turned 73 this year. Since ENIAC's birth in 1946, we have used computers to create a digital version of our analogue world. Human intelligence (and natural stupidity) evolved for the analogue world; however, human cognitive capabilities are limited when it comes to the complexity of understanding and decision-making in the digital world. This paper explores the capability of Artificial Intelligence (AI) to transform the financial industry. Banks and insurance companies have effectively digitized their businesses, with financial institutions reportedly spending more than any others on data. However, they find themselves caught between the *Scylla* of big regulation and the *Charybdis* of big data, particularly where financial compliance and risk management is concerned. It is no surprise, then, to discover that AI is shaping the FinTech and RegTech landscapes, in addition to related activities in the legal and professional services sectors. In the face of unbridled enthusiasm and unquestioning acceptance of many of the claims made for AI, this paper takes a balanced, critical stance in explaining the *what*, *why*, and *how* of AI in the financial industry with a particular focus on the *art of the possible* in regulatory compliance.



Artificial Intelligence for Regulatory Compliance: Are we there yet?

Tom Butler & Leona O’Brien
University College Cork
tbutler@ucc.ie leona.obrien@ucc.ie

"My colleagues, they study artificial intelligence; me, I study natural stupidity."

– Amos Tversky

Contents

- Abstract3
- Introduction3
- What is Artificial Intelligence?4
 - Understanding the Context for AI.....5
- Understanding AI’s Current Capabilities5
 - Semantics6
 - Perception7
 - Cognition.....7
- Practical AI in FinTech and RegTech: KR, ML, NLP and ANN ..8
- AI for Regulatory Compliance9
- From Single- to Double-Loop Machine Learning..... 11
- A Question Regarding Cognitive Computing in the Financial Industry..... 12
- Conclusions 13
- References 15

Abstract

The world's first digital computer, ENIAC, would have turned 73 this year. Since ENIAC's birth in 1946, we have used computers to create a digital version of our analogue world. Human intelligence (and natural stupidity) evolved for the analogue world; however, human cognitive capabilities are limited when it comes to the complexity of understanding and decision-making in the digital world. This paper explores the capability of Artificial Intelligence (AI) to transform the financial industry. Banks and insurance companies have effectively digitized their businesses, with financial institutions reportedly spending more than any others on data. However, they find themselves caught between the *Scylla* of big regulation and the *Charybdis* of big data, particularly where financial compliance and risk management is concerned. It is no surprise, then, to discover that AI is shaping the FinTech and RegTech landscapes, in addition to related activities in the legal and professional services sectors. In the face of unbridled enthusiasm and unquestioning acceptance of many of the claims made for AI, this paper takes a balanced, critical stance in explaining the *what, why, and how* of AI in the financial industry with a particular focus on the *art of the possible* in regulatory compliance.

Introduction

The digital technologies spawned by ENIAC have, by the end of the second decade of the 21st century transformed work in almost every industry. The potential of *smart machines* to *informate* and *automate* the banking industry was comprehensively explored by Shoshana Zuboff back in the 1980s. Since then, *smart machines* have transformed the financial industry. However, smart machines with artificial intelligence (AI) capabilities offer the potential for greater automation and to enable governments and business organisations to better manage the digital world of big data and

to further transform both the home and work place. In assessing AI's capabilities, Brynjolfsson et al. (2018) argue that AI "is a *general purpose technology*," like the steam engine and electricity, which spawns a plethora of additional innovations and capabilities." Nevertheless, they argue that AI's potential remains unclear, as does the answer to the question 'How best can AI be employed to the benefit of the workforce and economy?'

BigTech companies such as Amazon, Google, Apple, Facebook, Microsoft and IBM are all investing heavily in AI R&D (Mercer & Macaulay, 2018). So too are financial institutions and FinTechs (Noonan, 2018). However, while Google appears to lead the way in some reports with \$3.9 billion invested, with much of that in its 2014 acquisition of Nest Labs (\$3.2 billion) (Krauth, 2018), Kisner et al. (2018) point out that between 2010 and 2015, IBM invested \$15 billion in its version of AI—cognitive computing—and that since 2015, it invested an additional \$5 billion. This is a breath-taking figure, which dwarfs that of the other top 10 tech companies whose collective investment stands at an estimated \$8.5 billion. No surprise then that IBM Watson Financial Services is leading the way in the application of AI in RegTech for regulatory compliance (Groenfeldt, 2018).

All is not what it appears to be, as for the last 62 years, predictions of AI's capabilities and application have, generally, failed to materialise; however, some argue that we may be at a tipping point (Wadhwa, 2016). Nevertheless, others argue that AI will not achieve its potential "without a reengineering of how business organizations operate" (Reeves, 2017; cf. Ransbotham et al., 2017). Organisations that have achieved success in the application of AI, such as Chevron, provide much in the way of lessons for others to learn from, apparently. Buoyed by its success, Chevron's CIO, Bill Braun, informed researchers at MIT Sloan Management Review, "It's springtime for AI, and we're anticipating a long summer"

(Ransbotham et al., 2018). However, keeping with the weather theme, others are speaking about a "third AI Winter": that is, the period of pessimism among researchers, the press, industry and governments that followed the failure of AI to deliver on its promises in the 1970s and late 1980s/early 1990s, which led to a drought in funding on both occasions (see Crevier, 1993).

What is Artificial Intelligence?

The concept of artificial intelligence (AI) was posited by John McCarthy during the Dartmouth Summer Research Project on Artificial Intelligence in 1956. Early computer scientists recognised the potential of computers to perform logical and mathematical computations better, faster and cheaper than humans and at greater scale. Thus began the development of smart machines, capable of what many believed to be AI, with the objective of mirroring and/or outperforming human intelligence and mimicking the behaviour of the human brain. As with all such innovations, early applications were envisaged by both business and the military, principally in the US.

Three types of AI are proposed by researchers.

1. Weak AI or "narrow AI" is non-sentient machine intelligence where R&D "is aimed at creating programs carrying out specific tasks like playing chess, diagnosing diseases, driving cars and so forth (most contemporary AI work falls into this category.)" (Goertzel, 2007, p. vi). This is perceptual not cognitive computing, as explained below, and falls far short of the hype.
2. Strong AI/Artificial General Intelligence (AGI) machines are currently hypothetical, in that they have the ability to apply intelligence to any problem, rather than just one specific domain viz.

"AI systems that possess a reasonable degree of self-understanding and autonomous self-control, and have the ability to solve a variety of complex problems in a variety of contexts, and to learn to solve new problems that they didn't know about at the time of their creation" (Goertzel, 2007, p. vi).

3. Super intelligence is also hypothetical, in that its machines with artificial intelligence surpassing all human semantic, perceptual and cognitive abilities (Legg, 2008). Due to recursive self-improvement, super intelligence is expected to be a rapid outcome of creating AGI.

As Goertzel (2007) pointed out, Weak or Narrow AI is all that is currently possible, at least commercially. There may, however, be a transition point with Explainable AI or XAI, as David Gunning of DARPA argues that *"The current generation of AI systems offer tremendous benefits, but their effectiveness will be limited by the machine's inability to explain its decisions and actions to users...Explainable AI will be essential if users are to understand, appropriately trust, and effectively manage this incoming generation of artificially intelligent partners."* (Gunning, 2017). Regulators in the financial industry will, no doubt, be comforted by this development, as they are concerned about current machines, used as Virtual Assistants or Chatbots by banks, making unexplainable or *black box* decisions in relation to consumers. They are also concerned about firms managing risk and compliance using machine learning and deep learning algorithms. Note Gunning's use of the term *"intelligent partners"*, this gels with the ethical desire to use AI to augment, rather than replace, human intelligence.

Understanding the Context for AI

The digital world created by computer systems attempts to be a mirror image of the analogue reality humans inhabit. Digitising financial data conveys an obvious advantage where sophisticated, routine, or large-scale computation of structured data is concerned. However, in the world of big financial data, where there are truly enormous volumes of heterogeneous structured and unstructured data across siloed data stores, machines need to be really smart to assist human cognition and decision-making. Then there is the challenge of big regulation, where legal obligations have yet to be digitised. Hence, some argue that AI may not be up to the task of dealing with big regulation and big data, as Rene Buest of CIO magazine points out that all it does presently is *digital pattern matching* (Buest, 2017). To understand why, a short non-technical explanation is required.

If you think about it, every data item is simply a digital pattern of 1s and 0s. Remember too that today's computer processors, from those in your smartphone to an IBM supercomputer, just perform mathematical functions (such as add and subtract, multiply and divide) and logical operations (AND, OR, NOT). Numerical data has unique values that make computation or logical operations possible at all scales. However, they must still be labelled as being specific types or categories etc., so too must non-numerical data such as text, audio and images. The digital representation of your bank balance may be identical to your surname, that is they have the same digital pattern of 1s and 0's. However, because they are labelled and structured in context, a machine (computer plus software application) will not attempt to deduct your latest withdrawal from your surname. Likewise, the digital patterns that represent the totality of your facial image may be the same as that of your pet gorilla's, or the instance of a digital smart contract on a

blockchain. Even text, which is unstructured data, where each character or symbol has a pre-defined digital value assigned to it as in UNICODE, requires classification and labelling for the machine. Hence, labels, context and structure must be added in the form of metadata—that is data about data. So in terms of data processing, all AI can presently do is sophisticated digital pattern matching and linking. While Rene Buest argues that this is not enough, it is quite useful when it comes to financial institutions' abilities to address the problems posed by big regulation and big data.

Understanding AI's Current Capabilities

The first thing you need to understand is that AI-based smart machines do not understand. To a machine, one digital pattern is the same as another—they mean nothing to a machine, as machines cannot 'know', in the animal sense. Remember a machine is simply made of a central processing unit (the CPU, that does the maths and logic), a memory unit, (short-term and long-term), paths of communication, and a software algorithm or programme that presents the CPU with instructions and data. There is no ghost in the machine. Neither does it have a 'mind' in the neurological sense. Human created algorithms put the smarts in the machine.

There are four general technological paradigms in AI: knowledge representation (capturing semantics in models such as ontologies), natural language processing (NLP), machine learning (ML) and deep learning (DL) which is a form of ML using artificial neural networks (ANN). Together they offer great promise in making AI a reality, but they are not currently being used productively.

The following is a categorisation of the human attributes and capabilities that AI technologies possess, or will need to possess, if the AI paradigm's promise is

to become reality. This conceptualisation is drawn from the work of Amit Sheth and his colleagues Pramod Anantharam, and Cory Henson (Sheth et al. 2017; Sheth, 2018). Their perspective offers the a truly coherent description of what computer scientists are endeavouring to do with AI.

Semantics

In a smart machine, as with the human brain, semantics captures the meaning attributed to concepts and their relationships—it is the way knowledge and data are represented. Hence, knowledge representation (KR) is all about capturing the semantics of business entities expressed as data. Here knowledge about data is represented in semantic meta-data models expressed as taxonomies/vocabularies of concepts and their relationships. These are captured and stored typically in a knowledge base using an XML formalism called RDF-OWL (Butler, 2017). This network of concepts and relations is used to represent declarative knowledge about the world as graphs. Formal rule languages such as SWRL (Semantic Web Rule Language) have been proposed to augment OWL's declarative axioms and provide sophisticated rule representation. Note that OWL is a machine-readable language which represents declarative knowledge of a domain axiomatically as RDF triples. However, while the RDF/OWL namespace defines predicates with certain semantic rules, as indicated, it, requires extension by formal rule languages, such as SWRL. Our colleague Dr. Elie Abi-Lahoud argues that *"in order to support real-world uses cases, [a regulatory] ontology needs to be augmented, and further supplemented by rules that encode the meaning of regulations and policies"* (Ford et al. 2016). In addition, while W3C standards such as SPARQL, SPIN, and SHACL provide important mechanisms to perform rule checking, more powerful rule

languages than SWRL exist: take for examples, RuleML (Boley et al. 2001), RuleLog (Grosz 2013), and Legal RuleML (Ceci et al. 2016). Two other powerful approaches are Stanford Research Institute's Sunflower, which uses Flora-2 (Ford et al. 2016), or ErgoAI¹. Using such approaches, knowledge may then be utilized for interpreting business data through algorithmic perception and cognition, albeit guided by human subject matter experts (SMEs). While Sheth et al. (2017) refer to human capabilities, the following applies to smart machines: *"Semantic concepts may represent (or unify, or subsume, or map to) various patterns of data, e.g., [a machine] may recognize a person by her face (visual signal) or by her voice (speech signal). Once recognized, however, both the visual and speech signals represent a single semantic concept of a person as recognized by the [machine]. Semantics hide the syntactic and representational differences in data and helps refer to data using a conceptual abstraction. Generally, this involves mapping observations from various physical stimuli, such as visual or speech signals, to concepts and relationships as humans would interpret and communicate them"*.

The EDM Council's Financial Industry Business Ontology is the best known example of an open semantic standard and business conceptual model developed by financial institutions for *semantic computing*. It represents knowledge of how all financial instruments, business entities and processes work in the financial industry (Bennett, 2013; EDM Council, 2018). Semantic meta-data models are relatively mature in several industries and a number of large banks currently employ such models for *semantic computing*.

¹ <http://coherentknowledge.com/comparison-of-ergoai-to-flora-2/>

Perception

Perception involves capturing and interpreting signals from the environment. AI technologies perform excellent pattern recognition and classification, when analogue textual, numerical and a range of other visual and audio signals are digitised. In a human being, sensory inputs generated from physical stimuli influence the development of feelings, beliefs, facts, and ideas. Perception involves sensory recognition, using cognitive filters to capture data of relevance to humans. Today smart machines have a range of data capture mechanisms, from digital video cameras to sophisticated optical sensors, in addition to traditional data input devices. In humans, data is utilized by cognitive facilities to help understand the world. However, human perception involves interpretation based on previous knowledge (e.g. of things and patterns). As Sheth et al. posit "*perception is a cyclical process of interpretation and exploration of data utilizing associated knowledge of the domain. Perception constantly attempts to match incoming sensory inputs (and engage associated cognition) with top-down expectations or predictions.*" AI-based smart machines are based on a family of machine-learning (ML) and natural language processing (NLP) algorithms and artificial neural network (ANN) applications that currently perform these functions. However, all three approaches require significant human intervention. This takes the form of 'training' algorithms to recognise the correct digital patterns by labelling the digital patterns of known digital entities and things or correcting 'errors' in perception. The three categories of training are: Supervised learning, Unsupervised Learning and Reinforced Learning. Natural language processing uses similar approaches and is applied to texts, whether written or audio transcriptions. There are many examples of these approaches in large financial institutions and in new and established FinTech/RegTech vendors in the industry.

However, much more could be achieved if *semantic computing* and *perceptual computing* could be integrated.

Cognition

Cognition builds on perception and semantic links to achieve human understanding. It is especially reliant on previous knowledge or pre-understanding, captured in our mental models, including semantic models, many of which contain natural systematic errors and knowledge gaps, which result in pre-judgments, prejudice or bias towards certain categorizations and interpretations. To make decisions, humans engage in inductive, deductive or abductive reasoning, much of which is proven to be faulty. Thus, there are significant limitations to human cognition, primarily due to informational overload, attentional and cognitive bias, and the limited ability of humans to compute probability and evaluate risk (Levitin, 2012). In smart machines, cognitive computing can enable contextual, machine understanding of data from perception, based on knowledge gained from previous learning and captured in perceptual algorithms and semantic knowledge bases.

IBM Watson Financial Services and its main commercial vehicle, Watson Regulatory Compliance, is a "*cognitive cloud solution ... that streamlines the process of managing regulatory compliance.*" IBM adopted the term *cognitive computing* (Kelly and Hamm, 2013), as AI, as it is currently understood, is an 'urban myth' (Atkinson, 2016). However, it is clear from AI scholars such as former Stanford professor Roger Shank, that all IBM Watson is doing, and indeed all others too, is *perceptual computing*. Thus, success here has been elusive and no cognitive computer system (ML/NLP/ANN) has, to the best of our knowledge, passed the *Turing Test*. This simple test assesses whether a computer's problem solving abilities,

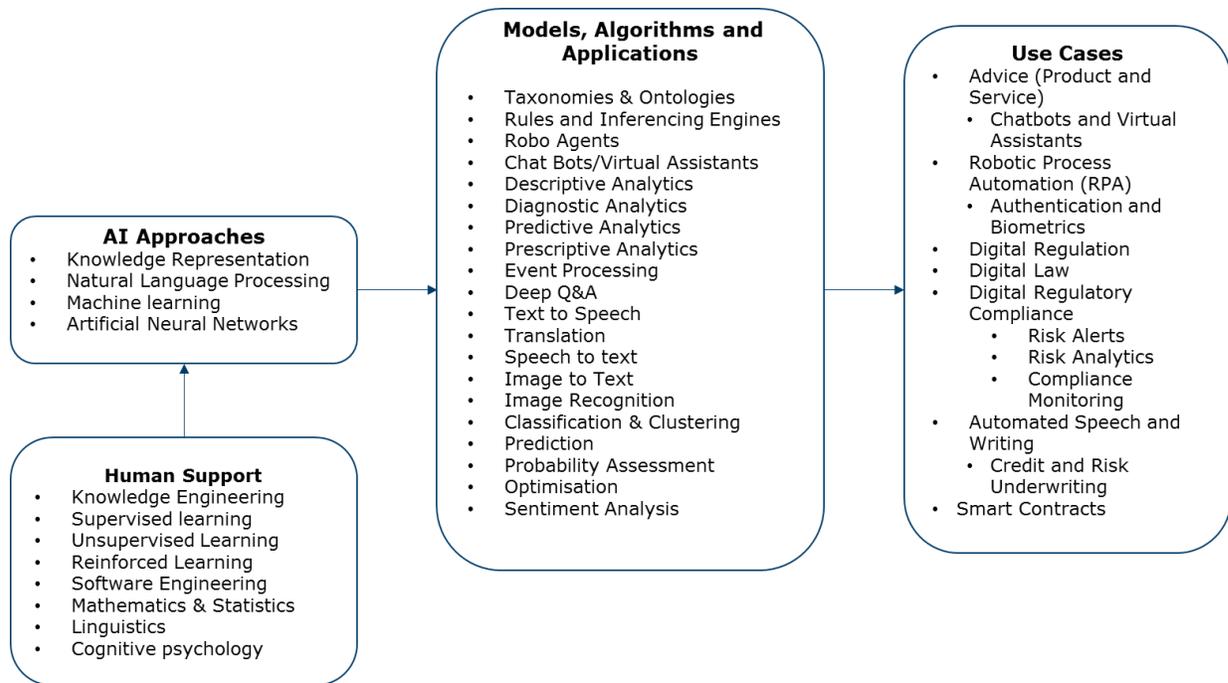


Figure 1 AI Approaches, Models and Use Cases in the Financial Industry

responses (output) and behaviours are viewed as indistinguishable from that of a human. It is clear from this brief analysis that the term *cognitive computing* is a misnomer, what is possible are *semantic and perceptual computing*.

Practical AI in FinTech and RegTech: KR, ML, NLP and ANN

It is clear that the big regulation and big data problem domains are not suited to traditional programming approaches that encode regulatory and business rules and logic into static programmes (cf. Butler, 2017; Butler & O’Brien, 2019). Then there are the questions on the tractability of static software engineering approaches and practices where solving complex knowledge engineering problems is concerned. Current software practices in Silicon Valley and RegTech and FinTech startups has been subject of significant criticism, and the software itself considered high risk, when compared to that being produced by the aviation

industry (see Pein, 2018; Somers, 2017). This may be a serious issue for financial compliance in the industry, given that there is significant evidence that problems with dysfunctional approaches to software development contributed to the financial crisis (Bamberger, 2010).

In order to make AI solutions more tractable, the industry focus is now on machine learning and natural language processing technologies, where developing and improving models using data and supervised (training) and unsupervised learning has proven to be more effective than previous AI approaches. In ML, a machine uses data to make (human) experience-driven predictions and decisions using algorithms and models on data patterns. With deep learning a machine can autonomously *learn* in order to adapt to changing environments. However, while this promises to reduce manual software engineering time, it is still too immature a technology. It is also a Black Box

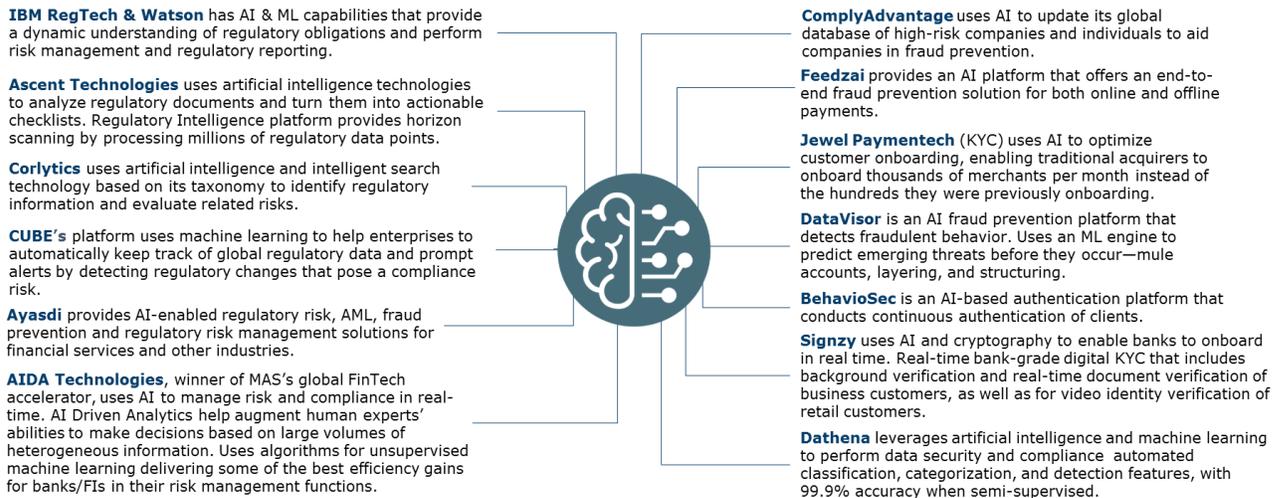


Figure 2 AI Innovators across the Regulatory Compliance Value Chain

approach, and would be currently unacceptable to regulators; however, as indicated, there is a move towards explainable AI. Nevertheless, given the weakness of current software engineering practices, all such models need comprehensive risk assessment.

The following figure captures the current state of practice in AI and indicates their potential as models, algorithms and applications, and more importantly in the context of this paper practical use cases in the financial industry, with an emphasis on financial compliance.

It is important to note that machine learning approaches predominate. These include, Regression (Statistical Models), Support Vector Machine (SVM), Graph Theory, Bayesian (BBN, GNB), Decision Trees, Ensemble, Clustering and Rule System models and algorithms. Using an ML approach reduces development time and costs and is a shift from designing traditional programme-based applications, but is not risk free. Currently, ensemble models and learning algorithms are integrated into applications and presented to users as SaaS applications, for example. It is also clear that much human input is required for the development of all AI approaches,

particularly subject matter expertise. In the domain of financial compliance, this includes legal, business, compliance and risk knowledge.

AI for Regulatory Compliance

Major financial institutions are beginning to scale up investments in the use of AI across the front, middle and back office, to bring efficiencies and cost savings and automate routine processes. In their research article, *Machine Learning and AI for Risk Management*, academic colleagues Aziz and Dowling (2019) present a technical analysis and cite many examples of the application of AI-based machine learning to bring efficiencies and cost effectiveness across business processes in banks and insurance companies. This paper's own research reveals that RegTech innovators have a number of growing but nascent AI-based solutions using KR, ML, and NLP in the regulatory compliance space. These currently focus on the following areas, particularly in retail and wholesale banking and also insurance:

- Horizon Scanning
- Unpacking Regulatory Provisions
- KYC and AML

- Fraud Detection & Prevention
- Consumer Risk Assessment
- Biometrics and Identity
- Cybersecurity
- Compliance Workflows
- Financial Risk Management
- Diligence, Vendor and Third Party Risk
- Surveillance

One recent case study involving Ascent Technologies is instructive of the type of solutions currently available. In early 2018, Commonwealth Bank of Australia (CBA) and ING completed a RegTech pilot with Ascent Technologies and law firm Pinsent Masons. Focusing on Markets in Financial Instruments Directive II (MiFID II), Ascent Technologies, applied NLP and ML to interpret and convert 1.5 million paragraphs of regulation into a series of "bitesize, actionable tasks appropriate for the banks" (Peyton, 2018). The obligations output by Ascent Technologies' legal subject matter experts and algorithms compared favourably to that which was produced by legal teams, but at a significant cost saving of "hundreds of hours" of legal interpretation, disambiguation and consolidation. The scale and scope of such approaches requires further validation. One issue for concern requires comment, and this applies to all such approaches, is that if the output from AI is 'text' to be read, then there is little in the way of cost savings across the regulatory compliance value chain. If, however, the output is expressed in machine-readable as well as a human readable format, then a financial institution can begin the challenge of performing digital regulatory compliance across its internal value chain (Butler and O'Brien, 2019). On a positive note, one of the key findings of the CBA-ING pilot is that collaboration between banks and RegTech innovators, whether start-ups such as Ascent Technologies or established BigTech companies such as IBM Watson Financial Services, is key to reducing the enormous costs of compliance, which Butler and O'Brien

(2019) cite as over \$100 billion per annum.

Figure 2 presents a sample of leading RegTech innovators employing AI technologies across the aforementioned areas. RegTech Associates estimates that there are over 600 vendors. Data on these is available at <http://www.rtdirectory.co/directory/browse/>.

Most of the large financial institutions that have AI R&D programmes in place are either in-house, where adequate IT resources and capabilities exist or can be acquired, or are collaborating with start-up companies through Accelerator Programmes, or pilots such as indicated above. The article by Financial Times investigatory Laura Noonan (2018) *AI in banking: the reality behind the hype* provides balanced, insights into current thinking across the industry. That analysis of 30 of the world's largest banks revealed that robotic algorithms, ML and NLP was being employed for the following:

- Chatbots and virtual personal assistants
- Customer profiling
- Automating and streamlining processes
- Data analytics and spotting patterns in data
- Risk management, chiefly through de-risking first line business processes in the front, middle and back office.

We have seen previously what the BigTech companies are spending. Potential competitors for banks, such as Apple and Amazon are spending \$786 million and \$871 million respectively, with Google at \$3.4 billion (Krauth, 2018). Across the world's top thirty banks the total amount being invested is at least \$90 million and up to \$400 million. So the entire financial industry is, on lower estimates, investing 12.5% of Apple's investment and at best 50% of Amazon's. However, many banks remain tight-lipped about their spending, so the true level of

investment may be much greater than estimated herein. Nevertheless, if one takes IBM's total spend as approx. \$20 billion, we get a sense of how unprepared the banking sector is, and, indeed, vulnerable to BigTech players extending their focus from payments into other banking and insurance services.

We can conclude from the foregoing that there is an imperative in the financial industry, particularly for dominant players, to keep pace with digital innovation. This means applying critical decision-making to avoid the AI hype and vaccinate its C-suite from being enthralled by fantastical fairy tales from the AI community. Enormous strategic, tactical and operational advantages can accrue from the critical application of the technologies mentioned above. The following discussion section will help inform how this may be prudently achieved.

From Single- to Double-Loop Machine Learning

Human actors generally engage in routine decision-making informed by single-loop learning (Argyris 1976)—this applies particularly to organisations. In single loop learning, minor or incremental adjustments are made in response to information feedback. Information which indicates that the fundamental values or assumptions on which a particular decision making strategy is based, are often overlooked and not processed. Static ML solutions are programmed to perform like a human engaged in single-loop learning.

In machine learning algorithms decision-making models and code are designed by data scientists in order to maximise the probability that pattern matching and related decisions-making produces answers with a high probability of accuracy to the limited problems the algorithms are being applied to solve. In such scenarios and use cases, correct patterns, decisions, and/or answers are

already known and the machines trained by humans to encode these patterns. Machine-learning-based AI systems do not know whether their answers are correct, this must be confirmed by a human expert. Thus, machine learning algorithms while being the product of human knowledge from software engineers and data scientists, are also supervised in their learning by human experts. When a human expert is satisfied that the machine algorithms have reached a satisfactory level of confidence in delivering correct decision outcomes, choices, and pattern matching, then training stops. It is clear that such solutions are limited in scope and scalability, and, furthermore, are static and not autonomous.

However, to be intelligent, machine learning algorithms should be capable of double-loop learning (Argyris 1976). Applying Professor Argyris' theory to AI, a ML algorithm should be able to: "(1) *Articulate its goals and objectives*; (2) *advocate for particular information and decisions choices*; (3) *engage in a dialectical approach that incorporates alternative or competing hypotheses and information*; (4) *Be open to change from an established position or decision strategy*; (5) *To make decisions and produce outcomes that are based on the valid information, even where data on certain conditions and factors is unavailable*; (6) *To transcend the constraints of human bounded and social rationality*; and (7) *To enable or aid human decision-making across simple, complicated, complex and chaotic domains*" (Butler et al., 2017). We have noted above that the US Government, in the form of the Defense Advanced Research Projects Agency or DARPA is advocating for XAI or eXplainable AI. As indicated, the ability of AI technologies to meet the above criteria and engage in double-loop learning is some way off. Deep learning algorithms, based on ANN technologies, that are capable of autonomous learning are in development, but for limited applications. Such

algorithms may be capable of unsupervised or adaptive learning, albeit it within very narrow problem spaces. According to Lee Bell of Wired magazine, "Researchers don't know how to do this with machines at the moment, at least not at the level that humans and animals can."

It is clear from reviews of extant research that the performance and characteristics of ML, NLP and ANN models and algorithms are improving. IBM's Watson is the most visible success, but here there are problems, which are worthy of further consideration, if only to understand as to why after \$20 billion of investment, Watson Financial Services are not dominating the industry with its *cognitive computing* platform. Watson's capabilities in *perceptual computing* is, however, proving highly successful in what is essentially digital pattern matching or in automating for labour-intensive, routine low-level legal and compliance data analysis and decision-making.

A Question Regarding Cognitive Computing in the Financial Industry

Hype aside, there is a dearth of evidence, as industry observers point out, for IBM's Watson Health and Watson Financial Services abilities to deliver true *cognitive computing* (Ray, 2017; Freedman, 2017; Shank, 2017). There is nothing wrong with Watson, per se, its strengths are considerable; however, if its critics are to be believed, it is mile-wide and an inch deep, in terms of its cognitive capabilities (Shank, 2017). Its algorithms, like others in the so-called cognitive computing space, require significant training by subject matter experts to produce trustworthy results. The more complex the domain, the greater the level of training and the greater the risk of incorrect decisions. That is why IBM acquired Promontory Financial Group, the risk and regulatory compliance consulting firm, to build Watson Financial Services.

It is significant for the progress of the AI movement in general, that IBM has encountered difficulties with its health offering (Freedman, 2017). The life sciences, and science of medicine, in particular have common, scientifically-based semantics in the form of vocabularies and formal knowledge representations. Given the significant discipline-wide agreement on semantic representations, it would be logical to assume that training machines to perform perceptual and cognitive computing would result in tractable solutions and application. Medicine would, therefore, present a suitable problem domain for AI, due to the level of formal knowledge representation across the field. However, industry commentators cannot reconcile IBM's claims for Watson's success in medicine and health with its achievements (Ray, 2017). Their concerns seemed justified as in May 2018 Watson Health Division shed over 300 workers at three acquired companies Phytel, Explorys, and Truven. Engineers informed industry association the IEEE, that mismanagement by IBM was the root cause (Strickland, 2018). As the IEEE's Eliza Strickland points out "All three companies, acquired in 2015 and 2016, brought with them hefty troves of health care data as well as proprietary analytics systems to mine the data for insights. The companies also brought existing customers: health care providers that used the analytics to improve both their care and their finances." IBM's problems identified in the recent reviews by MIT and Roger Shank should give senior financial executives pause for thought where AI is concerned.

According to Kisner, Wishnow and Ivannikov (2018), "IBM's Watson platform remains one of the most complete cognitive platforms available in the marketplace today. However, many new engagements require significant consulting work to gather and curate data, making some organizations balk at engaging with IBM." Kisner et al. point out that one client, MD Anderson spent over

\$60 million before suspending its Watson project. There was a significant misalignment between Watson and changes to MD Anderson's information systems that may have been outside the control of IBM. However, it is clear that there are many generic problems in enterprise systems integration that may have contributed to the MD Anderson project's failure, which may have had little to do with Watson's capabilities (cf. Gargeya & Brady, 2005, for example). Given, that financial institutions are particularly challenged in this regard in any event, it should give further pause for thought, when introducing any third-party offering, and not just Watson.

AI's lack of success in going beyond 'intelligent' digital pattern matching in record searching, and current limitations of expert system-like functionality in the health sector, raises questions on Watson's application to other domains, such as financial services. Nevertheless, IBM's decision to acquire Promontory Financial Group in order to acquire domain knowledge and expertise to train Watson for regulatory compliance in the financial industry and to address the big regulation problem is significant. It is clear, however, that the industry lacks the much-needed theoretical, conceptual, semantic, disciplinary, and ethical foundations (Wilmott and Orrell, 2017) on which to build the machine learning capabilities for *perceptual*, let alone *cognitive computing*. It also needs to be recognised that the semantic complexities of the regulatory and data domains are considerable, which renders the promise of AI as currently envisaged, intractable and impracticable, unless other fundamental issues are first addressed. Butler (2017) argued that *translation* and *Tower-of-Babel* problems beset in the financial industry and that the absence of a common language impedes industry-wide solutions for regulatory compliance. Such problems can only be solved using semantic standards; hence, we argue semantic meta-data models are a

necessary and sufficient condition for the successful application of AI.

As we have previously argued, "*there are significant limits to AI-based cognitive computing platforms as currently constituted. There is a dearth of evidence that AI is currently capable of interpreting and disambiguating complex legal texts and simplifying them for end-users or clients. Neither can AI explain the meaning of a regulatory rule in terms of compliance requirements and then persist this information in a machine readable format. Furthermore, such technologies cannot identify and map the provisions to impacted activities, without significant subject matter intervention*" (Butler et al. 2017). To be sure, AI algorithms can semantically enrich legal content, if trained to do so, but this comes at a cost and significant human intervention. Beyond that, it is back to the *rough ground* of legal practice and the application of a lawyer's or compliance officer's experiential knowledge and skills. On the data side, it is increasingly being recognised that searching for needles in the big data haystack is unproductive. What is required, after Zuboff (1988), is intellectual skills across the first and second line of defence, to theorize and hypothesize on cause and effect, on correlations, and then to leverage machine learning to capture and analyse data that falsifies or corroborates business theories. Smart machines won't do this, not if they are locked into applying static routines to solve known problems.

Conclusions

In framing this paper's conclusion we draw on Foteini Agrafioti, head of Royal Bank of Canada's AI research arm Borealis, she states in the Financial Times that "*there are too many people making these statements [about big cost and job impacts]...The problems we have solved are very narrow. The misconception is that humans and machines can perform at the same level. There's still a long way*

to go and many challenges we need to solve before a machine can operate [at a level] even near the human mind." She is clearly of the opinion that the benefits of AI as it currently exists are being oversold. If AI is to have the desired impact, financial institutions are going to have to address the *semantic, perceptual* and *cognitive computing* challenges, outlined above, in reaching the goals to automate and informate the enterprise using smart machines.

Forrester Research presents several AI predictions for 2019 (Lo Giudice et al., 2018). These mirror this paper's independent findings and those of our previous research. Putting it all together the following points merit attention.

- The primary challenge for firms adopting AI to perform risk aggregation and compliance reporting is data access and quality, specifically access to and the quality of, heterogeneous data spread across internal and external silos (see Butler, 2017; Butler and O'Brien, 2019). Firms will need to focus on data identification, enrichment, and integration before pattern recognition and intelligent response using AI solutions will be possible. However, semantics and knowledge representation approaches provide an ideal approach to virtualise data (Butler, 2017). Forrester Research argues that "*The tables will turn from AI to IA [Information Architecture] in the majority of firms... irrational exuberance for AI adoption must be equally met with solid efforts on an AI-worthy data environment.*" Knowledge engineering an semantic computing will go a long in achieving this.
- There is a clear shortage of data scientists, particularly where machine learning, NLP, artificial neural networks and knowledge

representation is concerned. This talent shortage will act to constrain many AI initiatives.

- Explainable AI will (XAI) be expected, as Black-Box AI-ML solutions will come under scrutiny from regulators and AI adopters, to say nothing of CROs. The impact of any decisions made by AI-ML-ANN applications that impact clients will need to be transparent as will their use for risk mitigation and risk data aggregation. Research by US DARPA identifies those AI technologies that may provide XAI.
- The need for Humans to be 'in-the-loop' will receive more recognition, given the forgoing, and as the limitations of AI for regulatory compliance become apparent.

Putting all of the above together, the focus will shift back to semantic knowledge representation and machine reasoning, that is *semantic computing*, in addition to machine learning's *perceptual computing*. As Forrester Research's Lo Giudice et al. (2018) conclude: "*In 2019, enterprise AI mavericks will rediscover knowledge engineering and digital decisioning platforms to extract and encode inferencing rules and build knowledge graphs from their expert employees and customers. ML's strength is data. Knowledge engineering's strength is human wisdom. Used together, enterprises can dramatically accelerate the development of AI applications.*" Our previous research has advocated this position for some time. There is no free lunch when it comes to AI. Thus, effective AI will involve hybrid approaches involving *Semantic Knowledge Representation, Machine Reasoning* and *Machine/Deep Learning* technologies that are based on firms applying open standards-based technologies. This offers

the optimal route to sustainably leveraging the many benefits of AI.

Finally, there is one problem. Tool support for semantic knowledge representation by subject matter experts is non-existent, for all practical intents and purposes. Business and legal compliance professionals normally employ the services of a software or knowledge engineer to translate business vocabularies and rules into a knowledge base. Such approaches are considerably error-prone, due to the translation problem (Butler, 2017). Hence, we conducted R&D on the development of a practical application for semantic knowledge representation called SmaRT. SmaRT is an easy-to-use, standards-based,² software application that helps business, legal and compliance officers create business, legal and compliance knowledge bases directly and without formal training in knowledge or ontology engineering or using sophisticated tools such as Stanford University's Protégé. Applications such as SmaRT are necessary if business organisations wish to realize fully the benefits of cognitive computing solutions.

Acknowledgement: Research on this paper and the SmaRT application was supported under the Enterprise Ireland Commercialisation Fund, co-funded by the Irish Government and European Regional Development Fund (ERDF) under Ireland's European Structural and Investment Funds Programme 2014-2020.

References

Argyris, C. (1976). Single-loop and double-loop models in research on decision making. *Administrative Science Quarterly*, Vol. 21, 363-375.

² Semantics Of Business Vocabulary and Rules or SBVR™.

- Atkinson, R. (2016). 'It's Going to Kill Us!' And Other Myths About the Future of Artificial Intelligence. Information Technology and Innovation Foundation. <http://www2.itif.org/2016-myths-machine-learning.pdf>.
- Aziz, S. and Dowling, M. (2019) Machine Learning and AI for Risk Management. In *Disrupting Finance* (pp. 35-50). Palgrave Pivot, Cham.
- Bamberger, K. A. (2010). Technologies of compliance: Risk and regulation in a digital age. *Tex. L. Rev.*, 88, 669-739.
- Bell, I. (2016) Machine learning versus AI: what's the difference? *Wired*, <http://www.wired.co.uk/article/machine-learning-ai-explained>.
- Bennett, M. (2013). The Financial Industry Business Ontology: Best practice for big data. *Journal of Banking Regulation*, 14(3-4), 255-268.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2018). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Boley, H., Tabet, S., & Wagner, G. (2001, July). Design rationale of RuleML: A markup language for semantic web rules. In *Proceedings of the First International Conference on Semantic Web Working* (pp. 381-401). CEUR-WS.org.
- Buest R. (2017) Pattern matching is not enough: Artificial intelligence is much more than just machine learning, CIO Dec 6 2017, <https://www.cio.com/article/3239708/artificial-intelligence/pattern-matching-is-not-enough-artificial-intelligence-is-much-more-than-just-machine-learning.html>

<https://www.omg.org/spec/SBVR/About-SBVR/>

- Butler T., O'Brien, L. and Ceci, M. (2017) Beyond the Hype of AI: A SmARt Approach to Unpacking Regulations Banking & Financial Services Policy Report, 36 (10), 1-12.
- Butler, T. (2017). Towards a standards-based technology architecture for RegTech. *Journal of Financial Transformation*, 45, 49–59.
- Butler, T., & O'Brien, L. (2019). Understanding RegTech for Digital Regulatory Compliance. In *Disrupting Finance* (pp. 85-102). Palgrave Pivot, Cham.
- Crevier, Daniel (1993). *AI: The Tumultuous Search for Artificial Intelligence*, New York, NY: BasicBooks, ISBN 0-465-02997-3
- EDM Council (2018) About FIBO. <https://edmcouncil.org/page/about-fiboreview>.
- Ford, R., Denker, G., Elenius, D., Moore, W., & Abi-Lahoud, E. (2016). Automating financial regulatory compliance using ontology+ rules and Sunflower. In *Proceedings of the 12th International Conference on Semantic Systems* (pp. 113-120). ACM.
- Freedman, D.H. (2017) A Reality Check for IBM's AI Ambitions, MIT Technology Review, June 27, <https://www.technologyreview.com/s/607965/a-reality-check-for-ibms-ai-ambitions>
- Freedman, D.H. (2017) A Reality Check for IBM's AI Ambitions, MIT Technology Review, June 27, <https://www.technologyreview.com/s/607965/a-reality-check-for-ibms-ai-ambitions>
- Gargeya, V. B., & Brady, C. (2005). Success and failure factors of adopting SAP in ERP system implementation. *Business process management journal*, 11(5), 501-516.
- Goertzel, B. (2007). *Artificial general intelligence* (Vol. 2). C. Pennachin (Ed.). New York: Springer.
- Groenfeldt, T. (2018) IBM's Watson Takes On Risk and Regulation In Finance, Forbes, <https://www.forbes.com/sites/tom-groenfeldt/2018/06/21/ibms-watson-takes-on-risk-and-regulation-in-finance/#36bc2c523b3d>.
- Grosz, B. N. (2013, July). Rapid text-based authoring of defeasible higher-order logic formulas, via textual logic and rulelog. In *International Workshop on Rules and Rule Markup Languages for the Semantic Web* (pp. 2-11). Springer, Berlin, Heidelberg.
- Gunning, D. (2017). Explainable artificial intelligence (XAI). Defense Advanced Research Projects Agency (DARPA). <https://www.darpa.mil/attachment/s/XAIProgramUpdate.pdf>.
- Holak, B. (2018) 5 AI predictions for 2019: Pragmatic AI takes hold. TechTarget Search CIO, <https://searchcio.techtarget.com/news/252453560/5-AI-predictions-for-2019-Pragmatic-AI-takes-hold>.
- Kelly III, J. E., & Hamm, S. (2013). *Smart machines: IBM's Watson and the era of cognitive computing*. Columbia University Press.
- Kisner, J., Wishnow, D., and Ivannikov, T. (2018) Creating Shareholder Value with AI? Not so Elementary, My Dear Watson. Jefferies Franchise Note, Jefferies, <https://javatar.bluematrix.com/pdf/f05xcWjc>, Accessed January 2019.
- Krauth, O. (2018) The 10 tech companies that have invested the most money in AI. TechRepublic, January 12, 2018,
- Legg, Shane. 2008. "Machine Super Intelligence." PhD diss., University of Lugano. http://www.vetta.org/documents/Machine_Super_Intelligence.pdf.
- Levitin, D. J. (2014). *The organized mind: Thinking straight in the age of information overload*. Penguin.
- Lo Giudice, D., Goetz, M., Purcell, B., Le Clair, C. and Gualtieri, M. (2018) *Predictions 2019: Artificial Intelligence No Pain, No Gain With Enterprise AI*. Forrester Research.

- Mercer C. & Macaulay, T. (2018) How tech giants are investing in artificial intelligence, <https://www.techworld.com/picture-gallery/data/tech-giants-investing-in-artificial-intelligence-3629737/>. Accessed 23 January 2019.
- Noonan, L. (2018) AI in banking: the reality behind the hype, The industry is taking a cautious approach in spite of excitement about new technology, *Financial Times*, London April, 12, 2018, <https://www.ft.com/content/b497a134-2d21-11e8-a34a-7e7563b0b0f4>.
- Pein, C. (2018). *Live Work Work Work Die: A Journey Into the Savage Heart of Silicon Valley*. Metropolitan Books.
- Peyton, A. (2018) CBA and ING complete RegTech pilot fling, 22nd February 2018. <https://www.bankingtech.com/2018/02/cba-and-ing-complete-regtech-pilot-fling/>
- Ransbotham, S., Gerbert, P., Reeves, M., Kiron, D., & Spira, M. (2018). Artificial Intelligence in Business Gets Real. *MIT Sloan Management Review*, September, 17.
- Ransbotham, S., Kiron, D., Gerbert, P., & Reeves, M. (2017). Reshaping Business With Artificial Intelligence: Closing the Gap Between Ambition and Action. *MIT Sloan Management Review*, 59(1).
- Reeves, M. (2017) Artificial Intelligence: The Gap between Promise and Practice, *Scientific American*, September 21, 2017 <https://blogs.scientificamerican.com/observations/artificial-intelligence-the-gap-between-promise-and-practice/>. Accessed January 24 2019.
- Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited.
- Shank, R. (2017) The fraudulent claims made by IBM about Watson and AI <http://www.rogerschank.com/fraudulent-claims-made-by-IBM-about-Watson-and-AI>.
- Sheth, A. (2016). Internet of Things to smart IOT through semantic, cognitive, and perceptual computing. *IEEE Intelligent Systems*, 31(2), 108-112.
- Sheth, A., Anantharam, P., & Henson, C. (2015). Semantic, cognitive, and perceptual computing: Advances toward computing for human experience. arXiv preprint [arXiv:1510.05963](https://arxiv.org/abs/1510.05963).
- Somers, J. (2017). The coming software apocalypse. *The Atlantic Magazine*. <https://www.theatlantic.com/technology/archive/2017/09/saving-the-world-from-code/540393/>.
- Strickland, E. (2018) Layoffs at Watson Health Reveal IBM's Problem With AI, *IEEE Spectrum*. 25th Jun 2018, <https://spectrum.ieee.org/the-human-os/robotics/artificial-intelligence/layoffs-at-watson-health-reveal-ibms-problem-with-ai>.
- Tiernan R. (2017) IBM's Watson, Despite Hype, Outgunned in A.I., Says Jefferies, *Barron's Tech Trader Daily*, <http://www.barrons.com/articles/ibms-watson-despite-hype-outgunned-in-a-i-says-jefferies-1>
- Wadhwa, V. (2016) The amazing artificial intelligence we were promised is coming, finally. *Washington Post*, June 17, 2016, https://www.washingtonpost.com/news/innovations/wp/2016/06/17/the-amazing-artificial-intelligence-we-were-promised-is-coming-finally/?noredirect=on&utm_term=.fa24a8fa2768
- Wilmott, P. and Orrell, D. (2017) *The Money Formula: Dodgy Finance, Pseudo Science, and How Mathematicians Took Over the Markets*, Wiley, UK.
- Zuboff, S. (1988). *In the age of the smart machine: The future of work and power* (Vol. 186). New York: Basic books.