COGNITIVE FINANCE

THIS ISSUE

NEXT ISSUE

What is A.I.?

A.I. and Board of Directors: a strategic alliance

A.I. ethical considerations specific to financial services

Investments in A.I. companies

Companies we follow: YNOME and others

Welcome to the Cognitive Finance Paper

At Cognitive Finance Group, our objective is to encourage **Intelligent Business Growth** in financial services using artificial intelligence solutions.

We also aim to achieve success with our growing international community and talented team. We want to solve complex problems, implement solutions and help drive business growth for the clients we advise and the start-ups we support.

This paper aims to foster inclusive, influential yet accessible conversations about artificial intelligence in financial services. In this issue, we are privileged to have excellent contributions from **Greg B. Davies** and **Michael E. Stewart** in the Thought Leadership section, in addition to Clara's. In the Companies' section you'll meet **Imandra** by Aesthetic Integration (formal verification for financial algorithms). *"Dark Pools"* by Scott Patterson is a factual account as to why Imandra is a groundbreaking and much needed technology for fair and safe financial markets.

On the final page we outline **Cognitive Finance Group**. We are a company with a clear vision. We draw inspiration and momentum from everyone we work with: our business associates and the companies we help grow. Our journey is inspired daily by YOU. We value you, your opinions and your ideas.

In the spirit of collaboration, we welcome constructive comments, ideas and topics you might want us to cover in more detail. We look forward to hearing from you. Do email us at *alex@cognitivefinance.ai*.

We hope you find The Cognitive Finance Paper of interest.

Yours truly,

Cognitive Finance Group

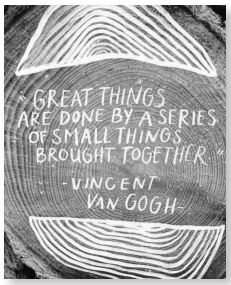

Great things are done by a series of small things brought together

Thought Leadership

The modern alchemy for business growth

Clara Durodié

Founding Partner, Cognitive Finance Group

clara@cognitivefinance.ai

I always have a good disposition and the unshakable belief that no matter how difficult a situation might be, there is always a good outcome to it. So, I'll start by saying that the wealth & asset management sector is profitable. Furthermore, we know that good margins always attract competition. According to Boston Consulting Group, in 2014 the annual global industry profits were $102 billion, generated from operating margins of around 39%. Financial technology venture capital is exploding: CB Insights reports that $19.1b was invested in FinTech companies in 2015, vs. $3.9b in 2013.

This sector shows the traditional signposts of an industry ripe for disruption: unhappy customers and very profitable incumbents.

In my opinion, there are **7 main forces** grouped as economic, socio-political and regulatory trends which are driving a subtle, incremental yet substantial power shift away from incumbents. These forces are redefining the business model in this sector and ultimately will determine who stays and who leaves.

1. Rapid development of Artificial Intelligence solutions. Like any fundamental shifts in established ecosystems, it will provide an opportunity to differentiate for intermediaries and asset managers. Artificial Intelligence will drive greater transparency, ease of use, and efficiency. It will also enable asset managers to find new tools to generate alpha. Generally, it is expected that technology will both raise funds under management, reduce investor's fees, minimise business operating costs and increase client reporting transparency.

2. Slow economic growth has a negative impact on most of the asset management industry and is forcing accountability on asset managers. One of the biggest implications for asset management is that underfunded pension funds will have a difficult time meeting their return expectations.

3. Women and Millennials are a new group of under-served customers with new values and expectations. Women's $14 trillion in assets in 2015 is projected to reach $22 trillion by 2020, according to a white paper by Family Wealth Advisors Council. Over $30 trillion is expected to be passed on from 'baby boomers' to 'millennials'. So, the question isn't if you speak *Millennial*. The real question is how fluent is your business in Millennial?

Having got used to ever evolving analytics from Facebook and Snapchat, millennials will expect nothing less technologically performant from their wealth manager. Women and millennials tend to invest differently. [1]According to the Spectrum Group, millennials are both more risk averse and more socially conscious than past generations when selecting investments.

[1] Who Will Disrupt Asset Management, and How by Katina Stefanova, David Teten, and Brent Beardsley, 2016

4. Geopolitical risk[2] around the world leads to capital flight to safe havens. Regional political instability exists in China, Russia, the Middle East, and South America. These countries now have millions of well educated, technologically savvy and newly wealthy citizens that look to protect themselves and their savings. The IMF recently reported that we are seeing the first net private capital outflows in emerging markets since 1984. For example, according to Bloomberg, money is leaving China at the fastest pace in a decade: an estimated $300 billion in financial outflows in the six months through March 2015. The wealth of the newly prosperous emerging markets will seek refuge in the United States, the UK as well as other perceived safe havens such as Singapore and Switzerland.

5. Transparency. There is a recent regulatory, political and media focus on charges. This has brought scrutiny to if and how this sector is fiduciary-compliant. Unprecedentedly, the corporate governance surrounding this topic has permeated each layer in organisations from their Boards down. The need for more transparency has increased particularly since 2008, as regulators and incumbents also realised the necessity of managing the "black box" of systemic risk.

6. Compression of margins in traditional wealth & asset management. While this sector remains profitable, margins have been under constant pressure. Wealth managers, in particular, have had a large share of regulatory updates to digest and implement. Their compliance departments have grown in size faster than their client-facing teams. Generally, compliance has become the second largest cost centre. The investment teams spend approximately 70% of their office time dealing with regulatory reports, updates and filling forms instead of managing investments, looking after current clients and developing new business. These circumstances have compressed profit margins and forced consolidation in this sector. They have also forced businesses to reconsider cost management frameworks, their operating and revenue model. Enter *digital in wealth management* otherwise referred to as *'robo-advisers'.*

About 18 months ago, after I completed scoping stage 1 of my PhD research, having more time on my hands, I began my presence on LinkedIn with a post claiming that the future of wealth management is digital. Unsurprisingly, some of my traditional wealth management contacts dismissed my post with Olympian conviction, stating that clients *"will always crave human interaction. Technology has no place in wealth management"*. Also, unsurprisingly I replied that "*If you do nothing to improve your digital delivery, you increase the chances of undermining your business*." It didn't take long to prove my point, sadly in 2016 they had to sell their business. Ignorance can be expensive.

Also in 2016, a notable digital wealth manager was born: NetWealth. It offers a mix of stellar teams from Goldman Sachs, UBS and J.P. Morgan private banks. Automation has enabled NetWealth to charge highly competitive fees starting at 65bps for minimum £50,000 which drop to only 35bps for portfolios over £500,000, an indicator that this service is also appealing to high net worth investors.

"Technology is changing the way financial services are delivered"

UBS Wealth, COO, Dirk Klee (2016)

No traditional wealth manager can compete with those fees! NetWealth raised £6.5m for their launch and have the backing of City of London veterans. Therefore, this marks the formal acceptance of digital wealth management in the

[2] Who Will Disrupt Asset Management, and How by Katina Stefanova, David Teten, and Brent Beardsley, 2016

City of London. Moreover, in October 2016, UBS announces the launch of UBS SmartWealth.

Wealth Chief Operating Officer Dirk Klee said it is a "strategically important move for UBS [which] enables us to bring our advice and expertise to a much wider audience. Technology is changing the way financial services are delivered. UBS SmartWealth shows that digital innovation is not the sole reserve of start-ups".

"We must become fintech companies"
Martin Gilbert, CEO Aberdeen Asset Management (2016)

Equally notable, another City veteran and CEO of Aberdeen Asset Management, said that "we must become fintech companies". Aberdeen is a listed company plagued by massive outflows in 2015 which pushed it out of the FTSE 100, a bruising wake-up call.

7. The Board of Directors and their understanding of technology and particularly of how artificial intelligence enables businesses to redesign the revenue generating model in this sector. As I have been building my Non-Executive Director portfolio, I researched in detail the UK Boards. I noticed that most Boards don't have specialists in operations **and** technology or indeed Non-Executive Directors who are knowledgeable to instigate the conversation on advancements in artificial intelligence and machine learning.

Why? Because there are few business technologists with asset management experience, who are able to solve strategic and operational problems in an integrated way, across multiple technology domains. This must be done while maintaining focus on governance, leadership, processes, people and client engagement. Moreover, even when visionary Boards might consider investing in artificial intelligence systems, there is little understanding of what the options are, such as what the artificial intelligence systems are most suited to solving their problems.

Why is it important? Most Boards tend to brush over the artificial intelligence impact and the value to the leadership derived from processing business data. Often, this is due to a lack of knowledge of what artificial intelligence is. This creates a domino effect: without the Board's support to explore the strategic impact of artificial intelligence applications, no headway is made. If Boards don't sign off budgets, usually nothing gets done.

An independent survey conducted by The Economist Intelligence Unit (EIU) in 2016 reveals that a flexible data strategy governed by strong leadership, facilitates business decisions and helps asset managers capture new business from their data. The survey was conducted across the United States and Europe amongst more than 200 executives in asset management and insurance. It highlighted that many businesses struggle to make sense of the deluge of digital data they receive on a daily basis.

These 7 forces have formed a unique environment which exposes structural problems. They also highlight unique opportunities which can be captured with the help of artificial intelligence systems for enterprise growth and profitability. In the coming issues of this Paper, I'll aim to discuss each of them in more detail. In addition, it might be worth mentioning that the next 5 years are likely to see artificial intelligence and Blockchain change the wealth and asset management business model. Artificial intelligence can transform the Front and Middle office, Blockchain the Back office. There's a lot of upheaval ahead and "if you are not a disruptor, then you are, eventually, going to be disrupted."[3]

[3] The Harvey Nash/ Alumni Board report 2016/2017 " Beyond governance - How Boards are changing in a diverse, digital world" in association with London Business School's Leadership Institute

Learning from other sectors

Artificial intelligence is the next big thing after the mobile and social changes which have happened over the past 5 years. Companies like SalesForce rely on artificial intelligence to keep growing past $10bn in sales and $8bn in revenues. Uber has shown that algorithms can successfully disrupt an entire industry.

Thanks to substantial investments in artificial intelligence companies, we are set to see more performant artificial intelligence systems at work. In 2016, Google's algorithm AlphaGo marked a significant turning point when it won against the South Korean Go Grandmaster. It represented a new form of artificial intelligence which is remarkably more advanced.

This enhanced use of artificial intelligence and machine learning is happening at a staggering pace. These are just three examples:

1- organisations like Associated Press use algorithms to write financial reports at a rate of 2,000 stories per minute.

2 - in 2014, the Hong Kong VC firm *Deep Knowledge Ventures* appointed Vital, an algorithm, to its Board of Directors. Vital analyses data from multiple sources and gets to vote on investments. This is a giant leap forward in how venture capital firms approach investments.

3 - a significant algorithm comes from IBM. ROSS the lawyer, is an advanced research resource for the 50 lawyers in the bankruptcy division of a leading US law firm. ROSS reads and draws inferences from existing laws to answer questions about specific cases or guidelines that users can ask using natural language. Its machine learning capabilities allow it to continue improving its answers.

Artificial Intelligence is a strategic capability

In order to stay ahead of their competition, wealth & asset management firms must implement artificial intelligence solutions to solve some of their current problems instead of a mere IT upgrades. Artificial intelligence is a strategic capability, it goes to the core of problems and redefines their solutions. The move towards intelligent growth, powered by artificial intelligence, is happening fast in our sector. It is much faster than we expect or would like to accept.

The question remains *'How can you apply artificial intelligence to gain competitive advantage, secure business growth and ensure your business doesn't become redundant?'.* This is the question our Cognitive Finance advisory team strives to answer for our clients. We help our clients make informed decisions on which artificial intelligence systems to acquire and then help with the implementation.

We also help Boards and senior management design a sound business strategy to incorporate artificial intelligence. We solve real problems.

About the author: *Clara Durodié is a business technologist and Board member with leadership roles in asset & wealth management in UK, Switzerland and Luxembourg. Clara is the founder of Cognitive Finance Group, an advisory and investment company specialised in A.I. solutions for financial services. Clara works with Boards, senior management, academia and A.I. start-ups. Her focus is risk management and business intelligence for profit and growth. Clara serves as mentor at leading Fintech accelerators in London, where she focuses on A.I. powered WealthTech.*

Clara is a member of the Chartered Institute for Securities and Investment (CISI) London UK, has a Certificate in Investment Management and holds an MBA from University of Oxford. When not collecting fine art, Clara is pursuing research for a part-time PhD at the intersection of neuroscience and machine learning interrogating how episodic memory informs savings decisions.

Thought Leadership

Shiny new toys. Does Wealth Management Need A.I.?

Greg B. Davies PhD

Founder, Centapse

GregBDavies@centapse.com

How to get ahead with A.I.

There are few things more enticing to an ambitious banking executive than a new technology that is poorly understood by the industry they're in. It offers the opportunity to appear innovative, cutting-edge, and entrepreneurial through nothing more than frequent jargon dropping. This is particularly true when this new technology is surrounded by popular media exposure and closely associated with the buzzy tech startup world. Current exemplars are *blockchain*, *Big Data*, *agile*, and, of course, *A.I.* and *machine learning*.

Not that A.I. is exactly new – Alan Turing proposed his eponymous test for A.I. in 1950, and it was nearly 20 years ago that IBM's Deep Blue finally beat chess grandmaster Gary Kasparov.

But advances in machine learning, and the rise of the data and tech unicorns has raised the tantalising promise of harnessing it commercially; something the burgeoning fintech bubble is increasingly forcing financial institutions to pay attention to… and that ambitious executives are keen to be associated with.

The problem with this enthusiasm is that it leads to *Shiny New Toy* syndrome. The recipe for career progression: find a hot new technology from outside your firm; persuade budget holders that this is the future and you need to invest in a pilot project or your firm will be left behind; spend innovation budget bringing outside technology in to solve, as yet unspecified, problem; look for problem to solve; ignore existing internal solutions to this problem (they don't have your name attached to them); build expensive pilot solution to said non-problem; await promotion; rinse and repeat.

Picking the right problems

I exaggerate somewhat, but it seems almost self-evident that a sensible approach to innovation is to carefully define the problem that you need to solve and then pick the technology best suited to solve it. This may in some cases be A.I., but often the best (and usually more cost-effective) solution, is something more mundane, quite possibly some capability or design you already have in-house, lying half-developed for lack of senior support and, hence, budget. In wealth management there do seem to be a lot of shiny new A.I. hammers looking for nails.

Accepting that A.I. does offer novel solutions that traditional approaches won't be able to provide, where should you seek to apply it? Which opportunities have the potential to sufficiently enhance your proposition to warrant pouring your organisational resources into the A.I. drain? The important distinction is determining the *nature of the problem.*

A first observation is that many problems in wealth management are simply not complex enough, or do not require enough data processing to warrant the sledgehammer sophistication and expense of A.I. solutions.

Unfortunately, the fact that these problems are easier means that they are often the first choice for proof-of-concept problems to allow the firm to test out the shiny new toy they have just committed to.

Some of these are complicated: they have many moving parts which are difficult for humans to mentally process. This makes them seem ripe for A.I. tools. However, complicated is not the same as complex. If these moving parts can all be effectively handled in a big well-designed decision tree to arrive at the right answer, then there is no need for A.I. Just build the decision tree, which can be entirely deterministic. Indeed, A.I. might not only be unnecessary, but may actually deliver an inferior solution.

Suitability **is the wrong problem**

The paradigm example here is *suitability* – determining the best investment solution for each client based on his or her individual circumstances.

There are many details of each client that affect the solution, and this is an area where humans struggle to comprehend how these details interact, let alone integrate them all flawlessly: getting it wrong can have enormous regulatory and reputational costs.

But it is not a problem that needs A.I. Suitability merely requires mapping a (relatively small) set of information defining each client's circumstances and preferences, to a (relatively small) set of different investment solutions.

The suitability decision tree that connects the two, if well designed, is certainly complicated, but nonetheless just a mapping algorithm that can be easily built, run, and monitored. For each unique set of client characteristics there *should* be a clear, auditable, deterministic process that points to the space of suitable solutions, and ideally, picks one. There is no need to over-engineer a solution to this problem, and indeed A.I. systems that pick solutions probabilistically, and which potentially evolve over time to provide unpredictably different solutions for the same client characteristics, would be a compliance nightmare, and potentially dangerous for the client.

So is Risk Profiling

Another example where there is frequently a misplaced desire to deploy A.I. is in risk profiling. This is over-engineering a (dangerous) solution to a problem that really doesn't exist. Risk Tolerance, correctly understood and measured, is a simple, stable psychometric trait[4], and best assessed with simple psychometric questions. These could be presented in slick, tech-enabled formats – but this is UX, not A.I.

Attempts to assess Risk Tolerance with real time processing of investors' behaviour, social media footprint, or video monitoring of their facial expressions are misguided. Such 'revealed preferences' are extremely unstable and reflect all sorts of short-term behavioural biases and influences of context that should not be used as a foundation for the long-term Risk Profile of an investor's portfolio.

Opportunities for A.I.

Behavioural decision support

However, both these examples bring us to areas where A.I. *could* be useful. For example, the same revealed 'preferences' that are often wrongly confused with Risk Tolerance can be very useful in other ways: in particular they give us important signals of when individual clients might be *anxious* along the investment journey, or where they are in danger of making decisions that deviate from good practice.

[4] Davies & Brooks (2014). Risk Tolerance: essential, behavioural, and misunderstood. Journal of Risk Management in Financial Institutions.

This shouldn't be used to change the Risk Profile, but *should* be used to help guide the client to better decisions in line with their stable Risk Profile. This is definitely a problem that could benefit from extensive data processing.

A.I. approaches could be deployed to look for the costly behavioural patterns in each individual investor, and connect with behaviourally designed systems of targeted communication, news feeds, and nudges to help them control these detrimental impulses; or integrate with CRM systems, sending real time signals to advisors to help them navigate exactly what to say, to whom, and when.

"Decision prosthetics can help improve human decision-making"

A.I. systems thus should not be used to determine suitability, but could be used to dynamically guide advisor and client behaviour *towards* the suitable solution. This problem, unlike risk profiling, has characteristics that A.I. is suited to: rapid processing of large quantities of disparate data; complex pattern recognition; and probabilistic responses, offering guidance where there is not necessarily a right answer, but a range of more or less likely responses (much as IBM's Watson approached the questions in *Jeopardy*).

Dynamic suitability

A further opportunity is to use A.I. to move towards *dynamic suitability* – that is, constantly updating suitability in response to changing client circumstances and preferences. This is not about changing the suitability framework used to determine the right answer for each client, but instead using A.I. to constantly update the data used as *inputs* to the suitability framework. If the client's balance sheet, circumstances, and goals and preferences are continually in flux, then so should be the suitable solution.

Applied well, such technology could increasingly blur the distinction between a) engagement, b) profiling, and c) suitability. Processing real-time data will allow A.I. systems to continually refine and update the client's behavioural profile, use this to prompt appropriate engagement and interaction, which in turn provides profiling information for the suitability inputs, and permits constantly updated solutions that serve the client's constantly changing needs.

There are, of course, many related use-cases that could draw on the same systems and technology: *just-in-time education* providing targeted information and education nuggets to clients when they are most receptive, enhancing understanding; client prospecting using data processes to identify life-events; finding appropriate cross-sell opportunities; emotional recognition using biometric data; and finding patterns of bank-client interaction that the organisation cannot itself discern because of the omnipresent barrier of organisational siloes.

However, it is very likely that these solutions would need be developed within an organisation, building the intelligence to suit the specifics of the proposition, client base, sales processes, and culture. This may require external technology, but a simple 'lift and drop' of technology from external vendors is unlikely to be successful.

Centaur design

Note that in all of these examples machines are not being used to make the decision, but as *decision prosthetics* to help improve human decision-making in pursuit of suitable solutions that are driven by much more traditional, stable models. This is another general feature of problems in wealth management that are most suited to A.I. solutions: they are more about decision-support than automation.

After Gary Kasparov was defeated by Deep Blue in 1997, he did something very novel – he started competing *with* machines, rather than against them, creating 'centaurs' of man and machine. And, despite the dramatic advances in A.I. in the last 20 years, the best chess player in the world today is not a machine, it is a man-machine centaur.

This is because humans and machines think very differently: there are things that machines are good at (data processing; pattern recognition; consistency; low error rates), and there are things humans are good at (empathy; coping with unstructured problems; creativity; generating insights from association across completely different problems; coping with dynamic environments and multiple objectives).

Because these skill sets are different there is value to combining the two – our ideal should be *centaur design*: applying A.I. to the right parts of the right problems, using them as decision prosthetics to help with what humans are bad at, but recognising that ultimately humans are required to make the decisions and navigate the deep uncertainty of a changing environment, and their own unstable preferences. Only when systems are designed to enable both parts of the centaur to work together will A.I. change financial decision making for the better.

There are parts of these systems, like suitability frameworks, where A.I. is not the right tool to provide the right answer and we need something more intentionally designed, engineered, and stable. However, A.I. might help to guide behaviour towards the suitable answer, or ensure the inputs to the stable system are much more current and error free, or even serve as a natural language (and inexpensive) assistant to help human users navigate the process and user interface.

But fundamental to using A.I. well is to first define the problem, and then look for an appropriate technology to solve it, rather than buying expensive shiny new toys and throwing them hopefully at every problem.

And those ambitious executives looking to associate themselves with a rising technology? They would do well to consider this serenity prayer of A.I.:

Grant me:

The serenity to accept what computers do better than people,

The courage to let people do what they do better than computers, and…

The wisdom to know the difference.

From Andreas Weigend

About the author: *Greg B. Davies is a specialist in applied decision science and behavioural finance, turning academic insight into practical applications. In April 2016 he founded Centapse, a firm dedicated to applying sophisticated behavioural insight to design, develop, and deploy solutions across industry to help people (and organisations) make better decisions.*

Greg started, and for a decade built and led, the banking world's first behavioural finance team as Head of Behavioural-Quant Finance at Barclays. He was the architect of Barclays' behavioural profiling tools and holistic Wealth Philosophy, delivering solutions tailored to both financial and emotional investment needs; and he designed Barclays' innovative behavioural approach to impact investing and philanthropy.

He holds a PhD in Behavioural Decision Theory from University of Cambridge; is an Associate Fellow at Oxford's Saïd Business School; a lecturer at Imperial College London; and author of 'Behavioral Investment Management'. Greg is also the creator of Open Outcry, a 'reality opera' premiered in London in 2012, creating live performance from a functioning trading floor.

Thought Leadership

Differences in modern Artificial Intelligence systems

Michael E. Stewart

Chairman and CEO of Lucid

michael.stewart@lucid.ai

All well-known A.I. systems today (IBM Watson, Google Deep Mind, Microsoft Cognitive Services/C, Cycorp's Cyc, et al) are actually only "partial" A.I.-capable systems. None of these systems possesses all of the various characteristics of a true, human-like intelligent brain. Thus, none are (as of yet) true "Artificially Intelligence" if that term is meant to describe a computing system capable of mimicking in a continuous behavior all of the cognitive attributes going on minute-by-minute, day-by-day in an actual human brain

Instead, each of these *partial* A.I. systems is architecturally different from the others, wherein each has been designed and built to focus on solving various different types of problems and/or capabilities which are part of human cognition. These "cognitive-like" problems and capabilities include such attributes as: autonomous discovery (self-learning), automated cognitive focus (continual guided awareness), deep statistical analysis (mathematical/algorithmic discovery), remembrance/knowing (stored reusable "Knowledge"), and other similarly important characteristics of human intelligence.

Although none of today's popular AIs do ALL othese things, the fact that each can perform impressive amounts of SOME of these capabilities allows them to be considered "AI". Most of these AIs leverage one or more types of "discovery" or "learning" capabilities called "machine learning", "neural net", or "autonomous learning". These capabilities are effected by mathematical algorithms autonomously applied to large datasets to *discover* interesting patterns in the data, such as valuable types of trends, inflections, correlations, and other statistically-derivable understandings that could not have been readily discoverable without the aid of those autonomous systems.

These types of AI (IBM's Watson, Deep Mind, Cortana, et al) are loosely referred to as being "right- brain" pattern-istic replicas of humans' (and many other species') abilities to *perceive* and *delineate* patterns in large data sets which we need in order to better understand our realities. Because these AI systems mimic how human minds *perceive* such patterns or trends, but do so autonomously against far larger datasets than humans can cognize, they are said to be artificially-like human intelligence... thus...AI.

But...and this is a very important BUT...these types of so-called "right-brained systems" have several key drawbacks. They create very valuable findings by overcoming human cognitive limitations; discovering valuable patterns in data in autonomous ways, but they lack many other human - intelligence capabilities. They are only partially AI.

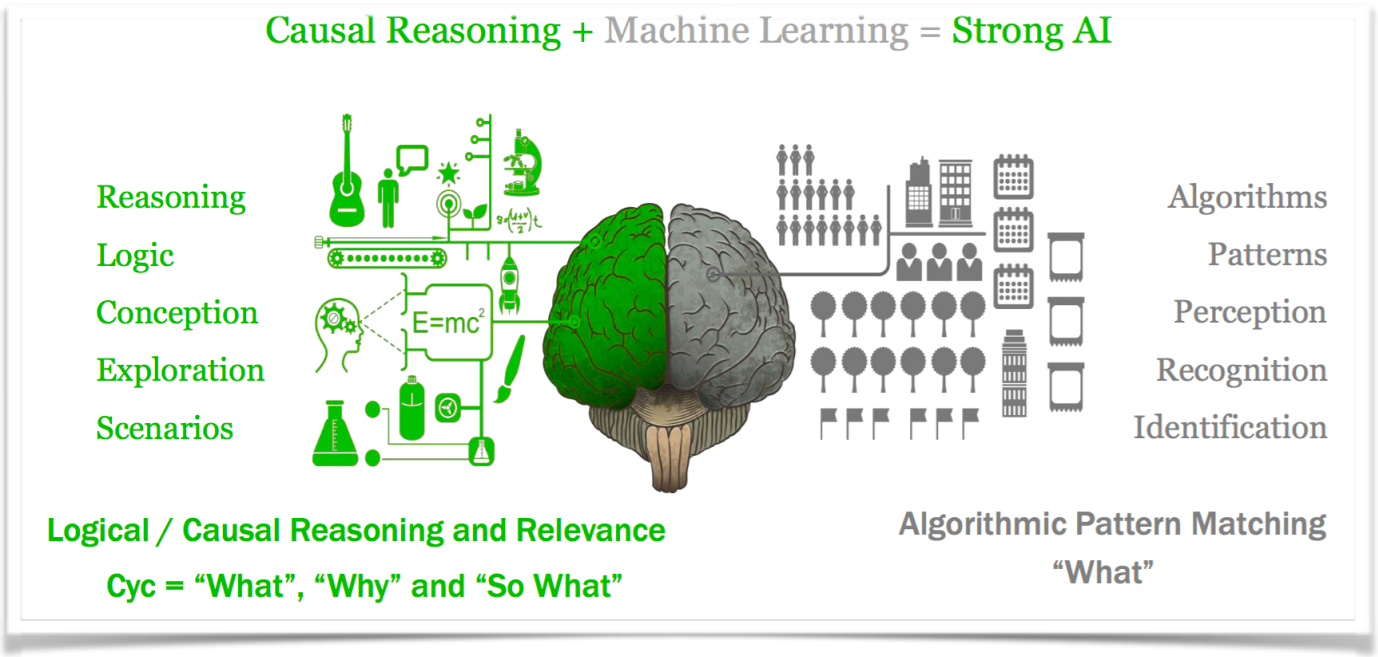

These Machine Learning/Neural-type of AI systems:

- do not actually KNOW what they are studying in any classical sense...meaning they aren't actually leveraging any sort of *conception* of pre-assembled underlying knowledge and/or logic structures to guide their statistical interrogations (humans instead have to use their conceptual reasoning powers to "train" these statistical approaches up-front);
- nor are they storing any of their self-discovered new-knowledge *concepts* into any sort of combined Knowledge-structure or "Knowing repository" (building a 'brain');
- nor are they able to "infer" with inferencing-powers (reasoning) lots of other Knowledge that might not have been in the data at all but could have been "non-statistically derivable" (obtained by using reasoning) from the data via the powers of inference/ reasoning,
- nor can they "explain" their findings in any logics-based format. Instead, they are relying only on the statistical interrogation of the data to create the "veneer" of Intelligence (think of Siri on a phone as *seeming* to be "smart" while not actually being Intelligent at all). Essentially, these systems fake being smart while not knowing anything.

These right-brain types of AI systems like Watson and Deep Mind are looking to answer the question of "WHAT is interesting" to find hidden within the data (correlations, trends, inflections, patterns etc.), but they are not built to answer "WHY" such findings might or might not occur (Causation of the phenomenon).

Also, they cannot extrapolate the "So-What" of their findings, knowing why those findings might be important or not to the Users' needs (Relevance). Causation and Relevance are essential bedrocks, not only of human-like Intelligence, but also of ensuring that AI is at its most useful to users.

Essentially, they are required to take information about a User's data to the next level, and can tell the User about that information in Natural Language and how the information was derived. This capability would offer the *transparency* of an AI which is called for in, for example, the **European "General Data Protection Regulation" 2016.**

Such Causation and Relevance capabilities in AI systems require a completely different type of system-architecture than using algorithms to discover correlations.

The ability of an AI system to be able to derive accurate understandings of Causation and Relevance about its findings requires a huge amount of “Prior Knowledge” and also “Subject-specific Knowledge”, leveraged by various forms of “Reasoning and Inferencing” capabilities about that Knowledge in order for the system to “cognize” those causal/relevant understandings. Causation and Relevance, and the ability for the system to be able to explain what it knows to the User, are powerful, knowledge-based/ reasoning- based capabilities which those right-brained forms of AI are not yet architected to try to solve.

A simple example helps explain these concepts, so consider two short fact-sentences of data:

1. Michael stole Tom’s car.
2. After being paroled, Michael went on to have a good life.

Without any intervening additional data or facts, existing right-brained statistical AIs like Watson or Deep Mind could trivially correlate to high accuracy that Michael stole the car, that he earned parole, and that he went on to have a good life (given that these were definitively stated facts). Even if hidden within huge piles of data, these systems could find and report these facts with high accuracy.

BUT, most all the world's existing AIs (apart from the Cyc Technologies) could not further INFER that Michael not only got caught, but that he also had a trial or he pled guilty in lieu of one, that there was a judge and likely a jury at that trial, that the judge likely had a black robe and kept order with a gavel, that the trial happened in a courtroom possibly with a bailiff, that Michael actually lost the trial and was found guilty, that he went to jail and served more than a few days on his sentence of jail-time, that he likely ate terrible food and had cold and dangerous showers while there in jail, and much more.

Humans can infer hundreds of thousands or maybe even millions of inference-able facts or truths about those two sentences which were NOT written at all anywhere in the data, by reasoning about prior knowledge that is somehow applicable to these facts.

The power of human-like inferencing to know facts or truths which are NOT included in the data is enormously valuable, creating the most leverage possible from users of AIs.

Cyc is the only existing deeply-"Causal type" of Knowledge and Inference-Reasoning-based, *commonsense* A.I. which can answer the original statistical correlations correctly from those two sentences just like other AIs, but which can also INFER all of those additional facts or truths from those sentences.

- Moreover, Cyc would explain WHY it knows all of those other facts or truths to be true, and in which perspectives or circumstances they might be true or not true (thinking about temporal, geographic, circumstantial, and/or hundreds of other constraints/perspectives).

- Furthermore, Cyc could likely infer thousands or tens of thousands of other facts which humans might have forgotten or not thought through completely. And, Cyc can offer detailed knowledge and reasoning to support each one of its findings, explaining also the Relevance of the findings, and doing so in plain Natural Language English to the User.

Like other AI systems, Cyc also has to be “taught” new Knowledge about any new processes it’s being focused upon. However, Cyc has already been taught in its 35+ years a

huge amount of core underlying "commonsense" (literally the massive amounts of "sense" i.e. knowledge that is "common" among all humans to "make sense of" the world).

"Cyc is the only commonsense A.I."

Thus, much like a generally smart human, Cyc can be rapidly taught new Knowledge about new things in short timeframes to create new Applications. This App-construction process now typically takes less than 3/4 months to reach a successful POC, and less than 9/12 months to reach V1.0...with ongoing learning by the App and *only marginal additional investments needing to be made by the User to own a hyper-intelligent process App which transcends human capabilities and continuously improves.*

In fact, in more than 100+ documented applications of Cyc, all have overcome that initial cost/benefit justification hurdle needed to justify any such capital investment.

In each of those applications, Users were typically adding much less than 2% of new Knowledge into Cyc to build each new App because of how much Cyc knew about the world in general and about their industry/processes prior to engaging in the App building process.

Currently, work is underway to make Cyc autonomous; this capability will enable Cyc to intelligently and autonomously "logically unify" new knowledge and "dynamically-orchestrate" disparate, heterogeneous computing systems for global corporations as if they were one unified, common system. We know HOW to do so, we specialize in doing large "pieces" of such engagements already via Cyc's Semantic Knowledge Source Integration (SKSI) capabilities, and these progresses are expected to emerge in the coming year or two at most.

While all other forms of AI on Earth have only required thousands of hours to research, build, and deploy, The Cyc Technology is now more than 35 years old, encompassing far more than 2 million hours of ceration and ontological work by employees with PhD level qualifications to enable its singularly unique commonsensical, causal reasoning capabilities.

Note: *We use the simplistic 'left/right' brain analogy to help readers understand AI while recognizing that modern cognitive scientists see this explanation of the human brain as a fallacious characterization of how our brains actually work.*

About the author: *Michael E. Stewart is the Chairman, CEO and a Founder of Lucid. He has a long career as an entrepreneur and innovator. As a partner in the Torii Group (2005 to 2009), Mr. Stewart performed consulting services in business change management, enterprise risk management and IT systems for the commercial real estate and insurance industries.*

As Chairman and CEO of Express Star Systems (1990 to 1998), Mr. Stewart designed software to pioneer the first integrations between CAD systems, relational databases, and business analytic algorithms. These software solutions optimized the operations and asset management needs of large commercial real estate businesses and other alternative asset classes. In the late '80s and early '90s and between 2001-2005 Mr. Stewart served as a principal systems design specialist with the U.S. government, pioneering next-generation Knowledge Computing and Intelligent Software Agent systems.

Companies we follow

Aesthetic Integration

IMANDRA : Formal Verification for Financial Algorithms

Denis Ignatovich

Co-Founder and Co-CEO, Aesthetic Integration

denis@aestheticintegration.com

Grant Passmore

Co-Founder and Co-CEO, Aesthetic Integration

grant@aestheticintegration.com

Many deep issues plaguing today's financial markets are symptoms of a fundamental problem: The complexity of algorithms underlying modern finance has significantly outpaced the power of traditional tools used to design and regulate them.

When it comes to exhaustively reasoning about the behaviour of complex algorithms, the only viable solution is formal verification, the use of deep advances in mathematical logic and artificial intelligence to automatically reason about algorithms and prove properties of programs.

Related techniques are already *relied* upon by other safety-critical industries such as avionics and microprocessor manufacturing. Aesthetic Integration is bringing formal verification to financial algorithms for the first time.

Imandra by Aesthetic Integration is groundbreaking formal verification technology that helps ensure financial algorithms are designed and implemented safely and fairly. From dark pools and trading strategies to blockchain smart contracts, Imandra analyses algorithms to help you understand what they do, why they do it, and what can possibly go wrong.

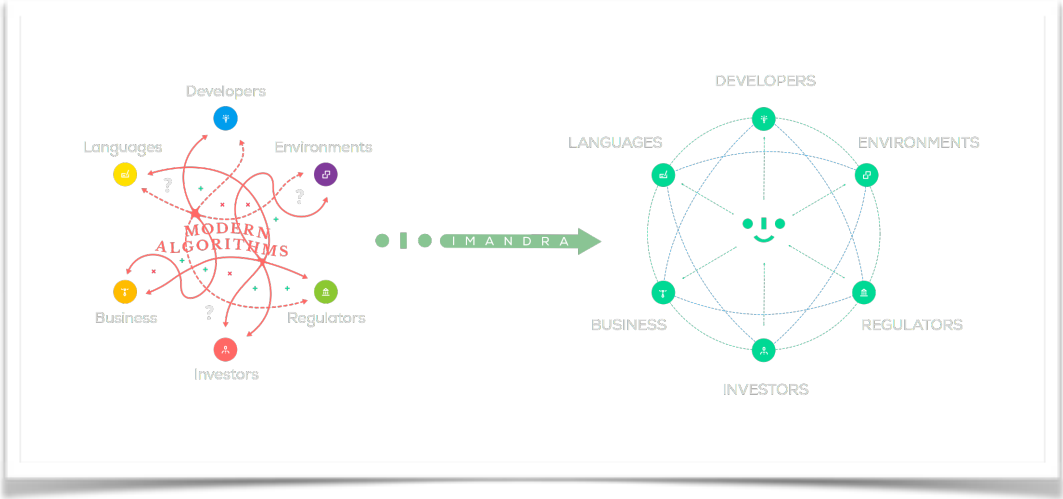

Introduction

Modern financial markets are built on a staggeringly complex tangle of algorithms. Competitive pressures and economic recession (e.g., decreasing margins and shrinking commission pools) have led to increasingly opaque and unstable markets. The effects of glitches and unfair advantages can be devastating, cratering the confidence of investors and hurting the general public.

In recent years, regulators and the industry have made tremendous progress in defining what safe and fair markets are. What's been missing is a way to analyse and regulate the complex algorithms underlying them.

Flash crashes, questions of fairness and a lack of transparent trading logic within dark pools are all symptoms of a fundamental problem: When it comes to designing and regulating electronic trading systems, financial firms and regulators have not had the right tools for the job.

The solution is *formal verification*, deep advances at the intersection of mathematical logic and artificial intelligence that allow us to automatically reason about algorithms and prove properties of programs. Powered by recent breakthroughs, we can at last scale formal verification to the complex software systems used in financial markets.

Aesthetic Integration's Imandra product is software that brings cutting edge formal verification to the design and regulation of complex financial algorithms. Imandra empowers a broad range of stakeholders — from traders, engineers and compliance officers inside financial firms to economists and enforcement teams inside regulatory agencies — with the proper tools to automatically analyse deep properties of safety, fairness and transparency of critical financial algorithms.

Managing the Infinite

Real-world financial algorithms are unfathomably complex. A typical trading system may, at any given time, accept hundreds of inputs and compute hundreds of outputs. The set of its possible configurations — its *state space* — is enormous. Faced with such a set of possible scenarios, how can we even begin to grasp whether a trading system's logic is robust enough to protect itself from making bad decisions? We must find a way to consider *all* possible behaviours of the algorithm to determine *what can possibly go wrong*, and to fix breaches of safety and fairness before they affect markets.

The unprecedented power of formal verification stems from its ability to automatically reason about such enormous state spaces, even *infinitely* large ones. It is quite remarkable, but mathematicians have been reasoning about the behaviour of algorithms over in nite state spaces

for a very long time. To gain some intuition, consider a sorting algorithm

$$F : \text{list } \mathbb{Z} \rightarrow \text{list } \mathbb{Z}$$

that accepts a list of integer values as input and returns as output the input list with its elements sorted in ascending order. How can we *prove* this algorithm will work correctly for all possible inputs? Certainly, we can test *F* on finitely many cases. But there are an infinite number of possible integer lists. Thus with testing, there will always be some cases (in fact, infinitely many cases) that we miss. Testing gives us no guarantee that bugs do not exist; they may be hidden in difficult to find corner cases not considered by our tests.

With formal verification, we can do (infinitely) better: We can use the proof method of *structural induction* to reason about *F* over the entire in nite state space induced by the datatypes involved in its execution.

To prove *F* is correct for all possible inputs, it suffices to prove two properties:

- P_1 : The output of *F* is always sorted.
- P_2 : The output of *F* is always a permutation of its input.

To prove both properties P_1 and P_2, we can use a particular structural induction principle, *list induction*, arguing as follows:

- Base case: P_1 holds of the simplest list.
- Induction step: If P_1 holds for an arbitrary list *X*, then P_1 will also hold for a new list (*n* :: *X*) obtained by prepending an arbitrary integer *n* to X. Here, both n and X are symbolic constants.

If we mathematically prove these two statements, then we have established that the sorting function will work for all possible inputs. With suitable automated theorem proving techniques, the construction of such proofs can often be completely automated. Moreover, if *F* is buggy (and thus no proof of correctness exists), we can instead automatically derive *counterexamples*, i.e., concrete input values that cause *F* to fail to meet its specification. **Please see our white-paper *Creating Safe and Fair Markets*** for a detailed discussion. Now contrast this type of rigorous mathematical reasoning with that of presenting several concrete “test cases” for which the function *F* works and then claiming that, since it works for those few, it should work for all the other infinitely many cases. Such an argument is clearly fallacious. Nevertheless, such “testing” is currently common practice in finance. Its obvious lack of scientific rigour is precisely why systems break down.

The bottom line: To analyse safety and fairness properties of complicated algorithms, we need powerful tools that perform complex mathematical reasoning to prove properties of computer programs automatically. That is, we need the latest advances in formal verification.

Let us examine formal verification’s use in other safety-critical industries. Then we shall discuss how related techniques can empower designers and regulators with the proper tools for ensuring the safety and fairness of algorithms underlying modern electronic financial markets.

How Other Industries Deal with Complex Algorithms

From the safety of autopilot systems navigating commercial jets and self-driving cars to the correctness of microchips in mobile phones, companies and governments worldwide rely on

formal verification to design and regulate safety-critical hardware and software.

Historically, formal verification has been used most in hardware (e.g., microprocessor) design and aerospace (e.g., autopilot) software safety. With recent advances in automated reasoning, it's become possible to scale automatic formal verification to reason about large- scale software systems. For example, Microsoft now requires device driver code for a piece of hardware to pass Microsoft's formal verification toolchain (the Static Driver Verifier) before the hardware can be "Windows Certified."

Companies like Intel, AMD and Centaur use formal verification in nearly every step of their design process. Much early momentum stems from a major debacle in 1994 when Intel released their Pentium® microprocessor with a bug in its floating point division (FDIV) instruction. A massive recall and subsequent re-fabrication cost Intel nearly $500,000,000. With the stakes so high, Intel competitor AMD took the pioneering step of engaging formal verification practitioners to verify the correctness of their new K5® processor FDIV design before fabrication, to great success. Today, major hardware companies have large in-house formal verification teams and the technology is integral to their design and development cycles.

In aerospace, formal verification is typically used to verify the safety of complex software systems underlying Air Traffic Management and on-board Collision Avoidance for autonomous aircrafts and autopilots. The NASA/NIA formal methods program is one of the leading forces. The aerospace regulatory bodies (FAA in the USA, EASA in Europe) specify use of FV-based ('formal' and 'semi-formal') methods via the DO-178C and Common Criteria software certification levels for safety-critical systems. The US Department of Transportation has recently commissioned related work for autonomous robots and self-driving cars.

Introducing Imandra

We designed Imandra from the ground-up specifically for financial algorithms, building upon decades of formal verification research and designing many new proprietary, patent-pending techniques for automated reasoning about financial algorithms.

Imandra is now in use at top global investment banks, empowering heavily regulated business units with unprecedented powers of governance, transparency and compliance for their critical algorithms. We're proud of Imandra's recent achievements. To list a few:

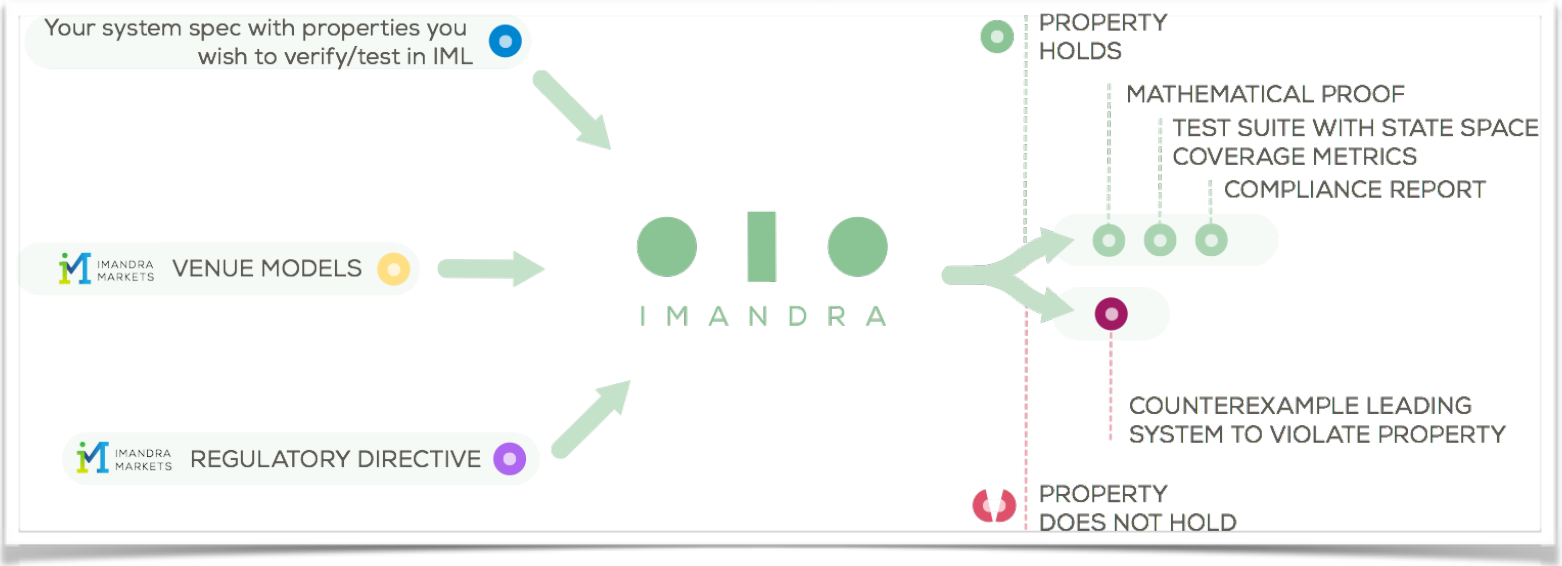

We won first prize in the 2015 UBS Future of Finance Challenge, a global ~5 month competition between 620 companies from 52 countries.

Our entry in the competition used Imandra to analyse the design of UBS's US dark pool with respect to a multitude of regulatory properties.

For key fairness and transparency regulations, we showed how Imandra could be used to automate analysis of the dark pool design and implementation for regulatory compliance. The previous approach to regulatory analysis (resulting in a $14M settlement with the SEC) had taken upwards of four years. See our white-paper Case Study: *2015 SEC Fine Against UBS ATS* for the full details.

With another top global investment bank, we have recently completed the world's first formally verified dark pool design. That is, their matching engine has been formally verified to satisfy over 100 key safety and fairness properties derived from US and EU regulations. Now, whenever they wish to add a new order type, they can use Imandra to specify the new feature and have the compliance of their new design be analysed automatically.

We've recently launched two cloud-based products: Imandra Markets & Imandra Contracts.

Imandra Markets is a cloud-based ecosystem for sharing client-facing algorithm specifications, radically simplifying the process of design, testing and compliance for your trading systems. Trading system connectivity is expensive.

Your team must decipher long texts and attempt to test systems without full understanding of how the venues operate.

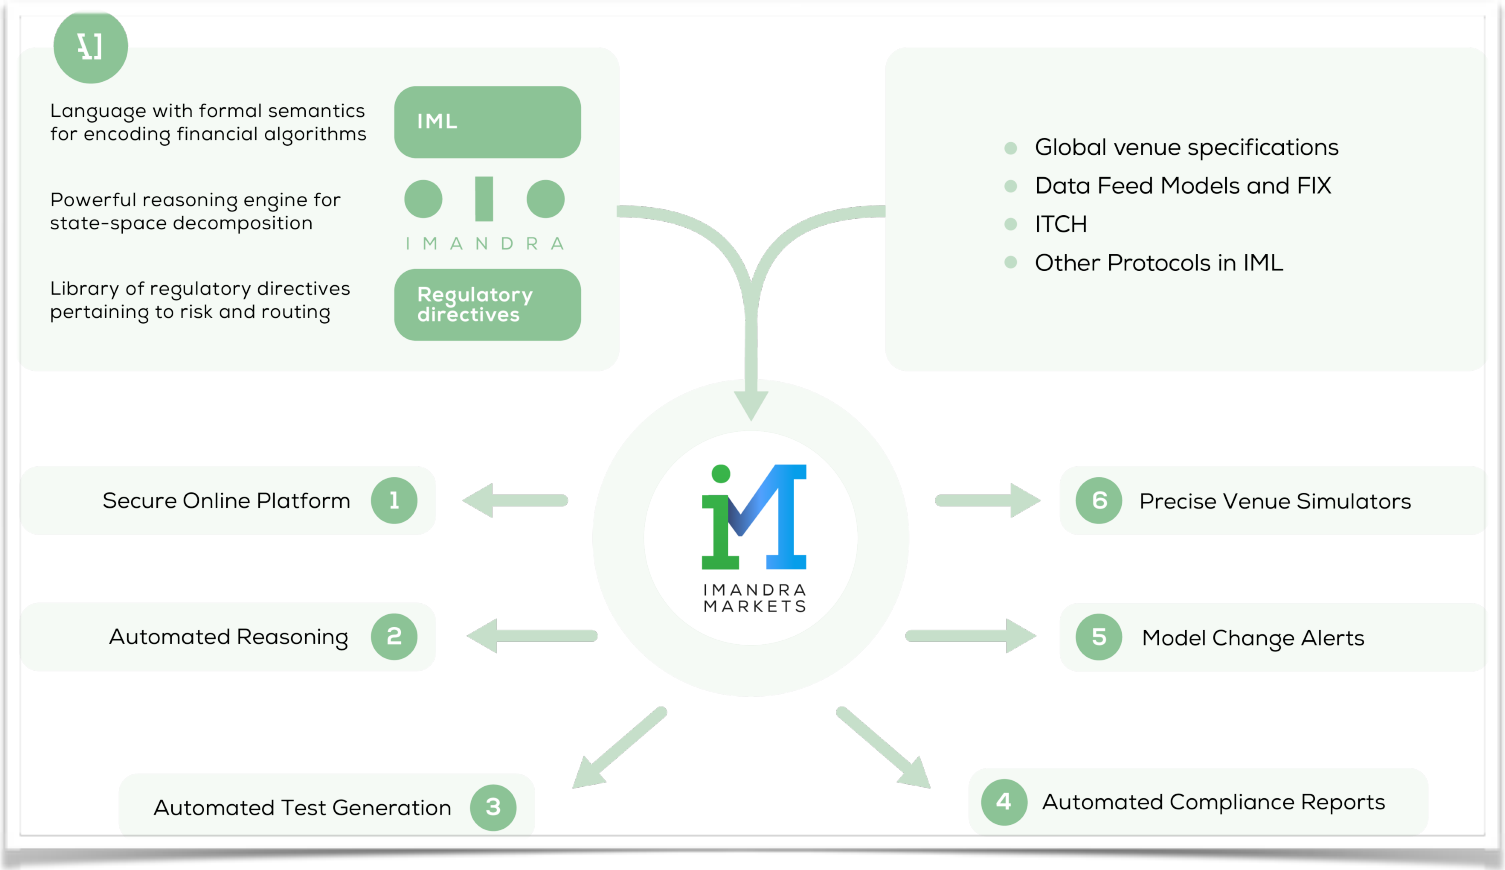

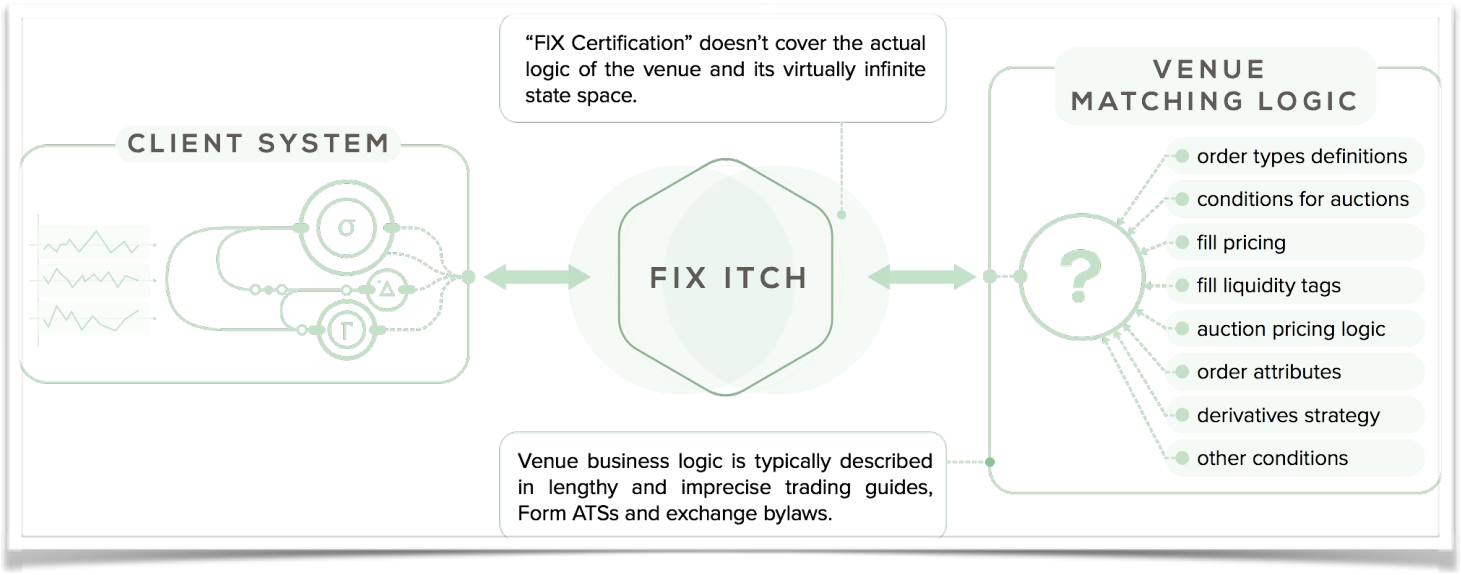

Why is this the case? Most trading systems connect to many different venues.

Modern exchanges and dark pools support many different order types and implement different messaging protocols. As a result, market participants end up wasting time on ***no-value-added*** tasks:

- Deciphering non-trivial matching logic designs from reading ambiguous guides (e.g. BATS's 200+ page description).
- Attempting to reason about corner-cases of routing decisions.
- Demonstrating system compliance to regulators.

Powered by a cloud-based ecosystem of precise venue and trading system specifications, Imandra Markets analyses connectivity between trading systems, provides up-to-date venue simulators, generates test suites and compliance reports, and more.

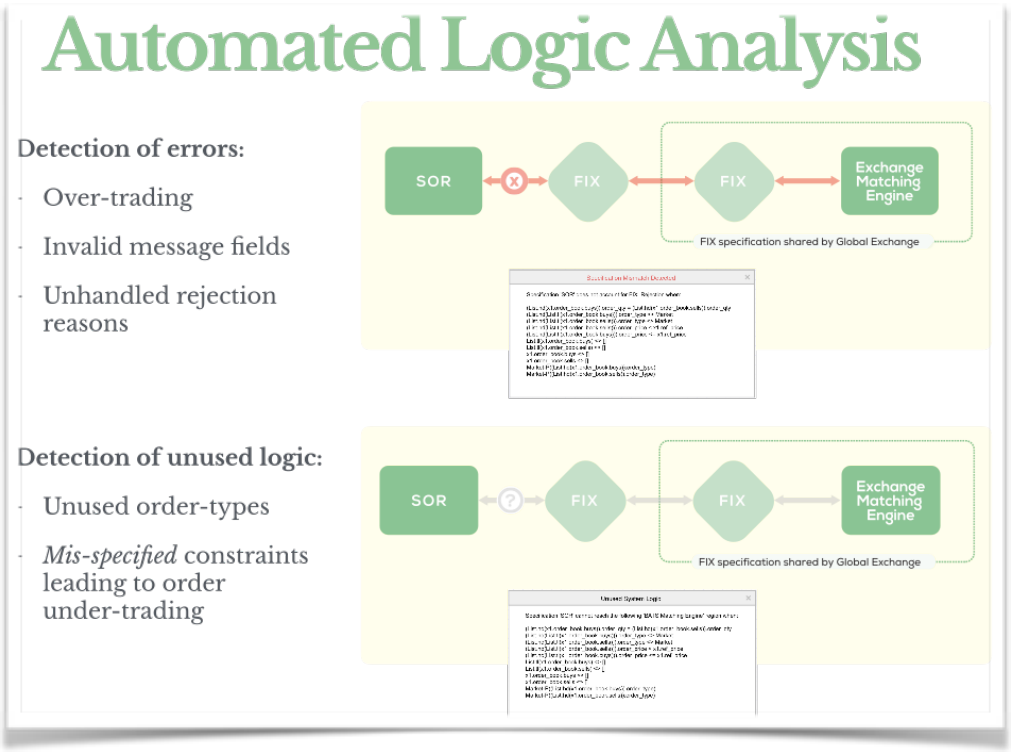

Imandra Contracts analyses blockchain smart contracts, helping you ensure they perform as advertised, and to understand and hedge risks associated with your smart contract positions.

Closing Remarks

Our mission is to provide financial markets and regulators with powerful tools for managing the complex algorithms underlying modern trading systems and venues.

Imandra by Aesthetic Integration brings revolutionary advances in formal verification to bear on financial algorithms, at last allowing us to scale robust engineering methods used in other safety-critical industries to finance.

We are driven by the fundamental improvements these latest advances are bringing to global financial markets.

Formal verification will eliminate significant portions of the costs and resources required to operate and regulate trading businesses. Precision and systematic rigour will replace ambiguous and ad hoc approaches to managing complicated trading systems.

Imandra will help you build safer, more stable and compliant businesses. Together let's make financial markets safe and fair.

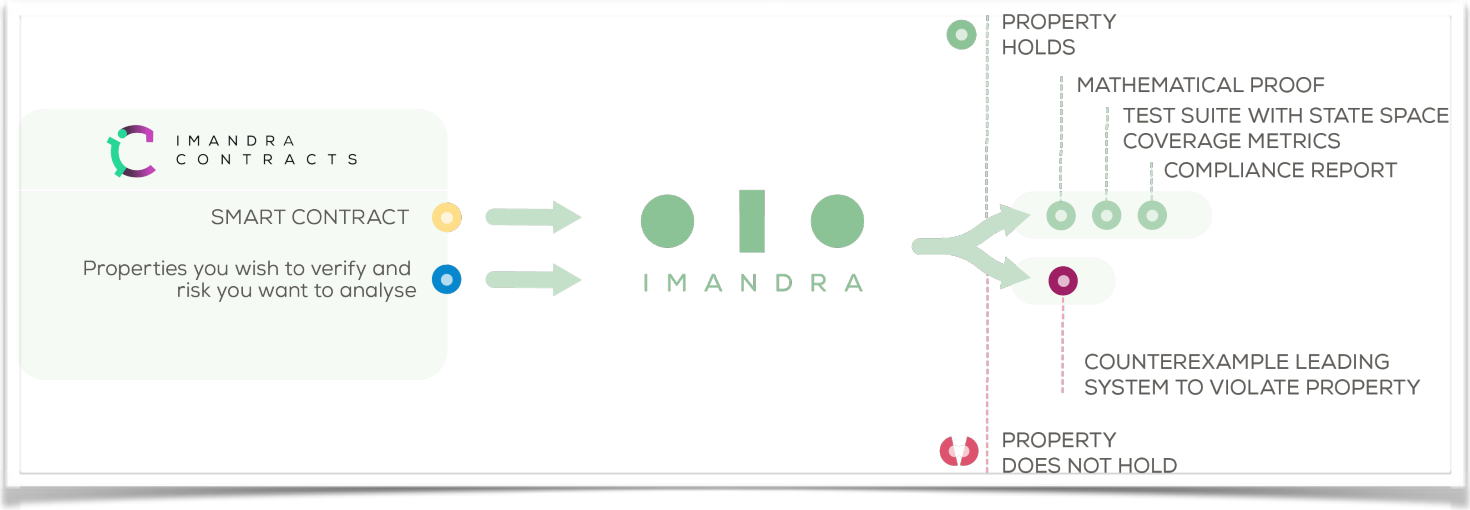

About the authors:

Denis Ignatovich has over a decade of experience in trading, risk management, quantitative modelling and complex trading system design. Prior to joining AI, he was head of the central risk trading desk at Deutsche Bank London. He holds an MSc in Finance from the London School of Economics and degrees in CS and Finance from UT Austin.

Grant Passmore has more than ten years' industrial formal verification experience. He has been a key contributor to safety verification of algorithms at Cambridge, Carnegie Mellon, Edinburgh, Microsoft Research and SRI. He earned his PhD from the University of Edinburgh and is a Life Member of Clare Hall, University of Cambridge.

COGNITIVE FINANCE GROUP

is an advisory & investment company specialised in artificial intelligence (A.I.)

Through our advisory work and public speaking engagements, we are committed to educating and informing financial services sector to the benefits of using A.I. in **building the next level of competitive advantage.**

ADVISORY

FOR financial services firms

Cognitive Finance Advisory partners to deliver strategic change. We strive to be the trusted adviser to financial services' Boards and senior management for selecting and implementing A.I. systems in their businesses.

Our team draws on years of experience driving technology change in blue-chip financial institutions. We apply this depth of experience in delivering strategic programmes for the adoption of A.I.

We have deep knowledge of the A.I. systems vendors. *We remain impartial and will always recommend the most suitable A.I. system.*

In addition to the A.I. industry and regulatory knowledge, we leverage a network of smart and experienced consultants to help prepare for and implement A.I. in your organisation.

We deliver

1. short & focused strategic analysis of business problems, examining how A.I. can be used to address concrete issues faced by your organisation
2. structured programmes or work with your teams to design change programmes for the adoption of A.I.
3. implementation of A.I. focused programmes and projects

contact:

michael@cognitivefinance.ai

INVESTMENTS

FOR professional investors

Cognitive Finance Investments is an institutional grade investment platform investing in A.I. companies.

In order to access the next generation of business growth, we are currently developing a unique investment proposition to serve institutional investors

- pension funds
- university endowment funds
- charitable foundations
- family offices
- investment funds and trusts

contact:

clara@cognitivefinance.ai

COGNITIVE FINANCE - SPEAKING ENGAGEMENTS 2016

WEALTH MANAGEMENT

31 Oct	Barclays annual digital conference, London
9 Nov	A.I. in Finance Masterclass (an ECN event), London
17 Nov	Jersey, AI in wealth management
24 Non	Board governance and A.I. @ Boards Intelligence, London
1 Dec	Parkwell Wealth Management conference, London

REGULATORY, PUBLIC POLICY and GOVERNANCE

24 Oct	Intel, IOT Solutions Conference, Barcelona
25 Oct	The Turing Institute, London
15 Nov	Speech, IEEE, Brussels
16 Nov	Masterclass for European MPs on A.I., Brussels
17 Nov	The Alan Turing Institute, London
18 Nov	Royal Society of Engineers, London
6 Dec	Machine Learning and the Law, NIPS, Barcelona