Assurance in the age of AI

The impact of emerging technology on assurance approaches and implications for assurance leaders

EY
Building a better working world

Contents

Executive summary

Emerging technologies have quickly created a situation where traditional approaches to assurance are increasingly inadequate to address the new risks these technologies create.

Emerging technologies such as Blockchain, Artificial Intelligence (AI), Internet of Things (IoT) and Robotic Process Automation (RPA) present significant opportunities for both improving our world and creating competitive advantage but they all bring with them new risks that need to be understood, managed and assured.

The speed, ubiquity, complexity and invisibility of technological change has driven holes through and paths around our traditional three lines of defence. Without new approaches to assurance there is the danger of a breakdown in the willingness of people to engage with technology and to share data — an erosion of the 'digital trust' which is increasingly important to the success of our organisations, economies and societies.

This paper uses the example of machine learning (an area of AI) to illustrate examples of some of the new risks that come with emerging technologies. We identify two areas where assurance approaches need to change:

- Firstly we outline necessary changes to existing assurance approaches to make them more timely, relevant and capable of addressing the risks emerging technology creates, and
- Secondly we outline an approach to ethical assurance which is an area we believe will be increasingly important if assurance is to remain relevant to organisations, investors and society more broadly.

To implement these changes we outline a number of practical steps assurance leaders can take to work out where to focus, to upskill their team, and to continue to appropriately adapt their functions in the future.

We conclude that if assurance is to continue to contribute to building a better working world for our companies, our stakeholders and ultimately the next generation then action needs to be taken now by assurance leaders to engage with their stakeholders in this area.

We hope that the suggestions in this paper provide a useful starting point for this conversation and we would be delighted to discuss any of the topics covered.

What assurance challenges does emerging technology create?

Much has been written (and many beautiful PowerPoint slides created) about emerging technologies such as Blockchain, RPA and AI. However, approaches to assuring these have often been slow to emerge and where they have these are mostly technology specific.

There are four common characteristics of emerging technology that have made designing appropriate assurance techniques increasingly challenging:

1.

Speed

The pace at which new technologies such as Blockchain and AI are evolving drives three main challenges:

- 'Pilots', 'proof of concepts', 'agile' and other quick ways of implementing emerging technology means that it has often landed and is in use inside an organisation before the assurance implications have been considered.
- By the time technical assurance training has been developed and rolled out (with equally beautiful PowerPoint slides), the technology has often moved on. Traditional methods for developing and delivering training haven't kept pace with the rate at which technology is evolving.
- Regulators and professional bodies have yet to develop frameworks and approaches for guiding how these should be considered, implemented and assured.

2.

Ubiquity

The extent of the potential, and in some cases actual, adoption of these technologies creates a further challenge. Simply put both the likelihood and impact of emerging technology risks are increasing:

- The likelihood increases as the breadth of adoption increases. For example Gartner predicts that AI will be in almost every new software product by 2020[1].
- The impact increases as the depth of adoption increases. For example, IoT technologies are increasingly used to control and protect national infrastructure and AI is being used in healthcare both for diagnosis and recommendation of treatment.

3.

Complexity

Emerging technologies aren't impacting organisations in nice bite-sized chunks:

- Convergence means these technologies interact (for example, there is no reason you can't use AI to process Blockchain transactions on IoT). The ever increasing interactions between autonomous computer systems may lead to unpredictable and potentially untraceable outcomes and as such technology specific assurance approaches are of limited value.
- Extended enterprises mean that these technologies are not controlled exclusively by the organisation and are often adopted through the use of third party services or dictated by the supply chain. Increasingly, the data that is used by emerging technologies is shared between organisations.

4.

Invisibility

There is a danger that risks and therefore the need for assurance goes unnoticed:

- The very existence of the emerging technology components may be unclear when it is embedded into things we use. Software may include things such as machine learning and a service may be delivered using automation e.g. chat bots. Even where this use is clear, there is often no transparency around the level of assurance that has been already been performed over it.
- The need for assurance may be less visible to teams where the risks created by emerging technology initially impact stakeholders outside of the organisation. For example profiling based on observed data (collected through online activity or cctv), derived or inferred data could cause significant unwarranted reputational damage for an individual.

[1] https://www.gartner.com/newsroom/id/3763265

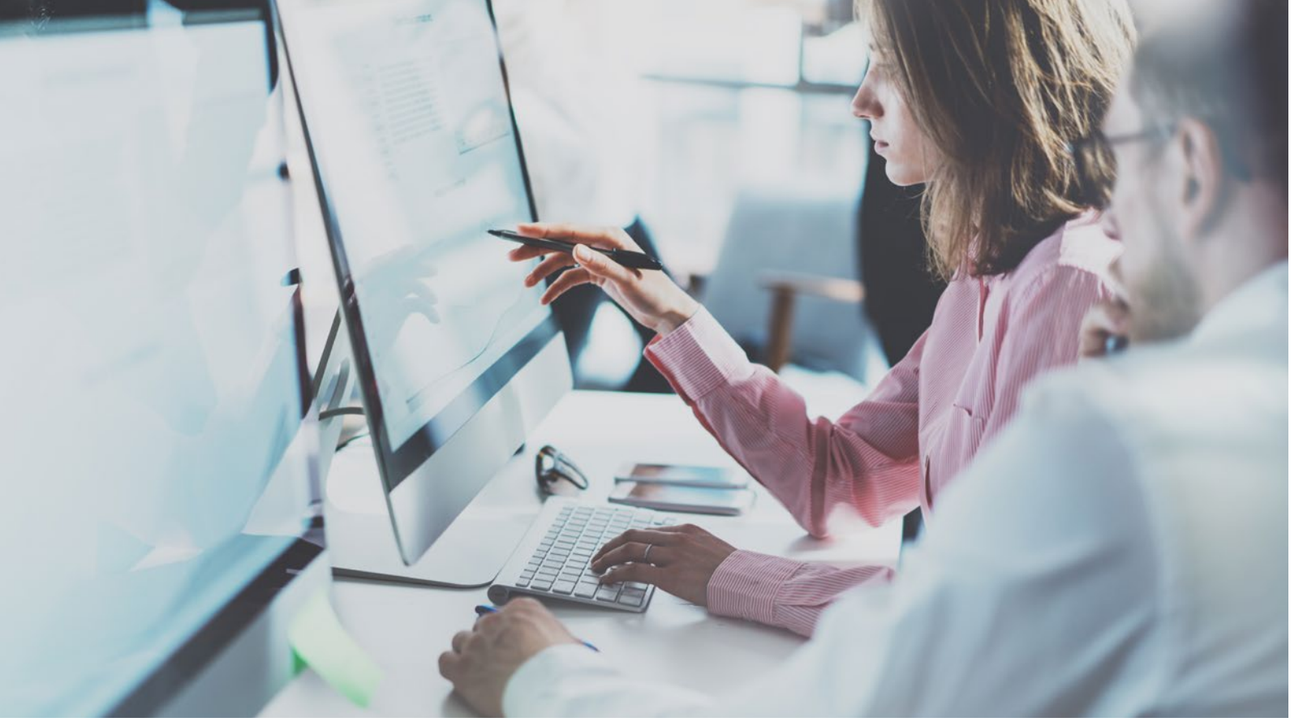

So what happens if we can't find ways of addressing these new assurance challenges? Without sufficient scrutiny, it is not beyond the realms of possibility that companies may unwittingly make decisions or undertake actions that are at best sub-optimal and at worst a violation of law. For example a lack of transparency and rigour over how an organisation obtains, uses and gains comfort over data used by AI to categorise, rank and sort individuals will increasingly be untenable.

Rapidly there will be an erosion in 'digital trust' and given that trust is the cornerstone of the digital economy getting this wrong could result in a loss of customers, market share and brand value.

Conversely those that get it right will be able to differentiate themselves from their competitors in the digital economy as they look to disrupt themselves and enter new markets. Digital trust will support the adoption of new products and services by customers and afford organisations the ability to reach more people and iterate more quickly. As these emerging technologies move us to an increasingly digital economy, the role of trust becomes more, not less, relevant to the success of companies, economies and societies.

Example: Machine Learning – the need for algorithmic assurance

To bring this to life let's have a look at the real life example of machine learning. Using algorithms that iteratively learn from data, machine learning allows computers to find hidden insights or optimise data sets without explicitly being programmed where to look or how to do this.

Below is the simplest representation of how machine learning works:

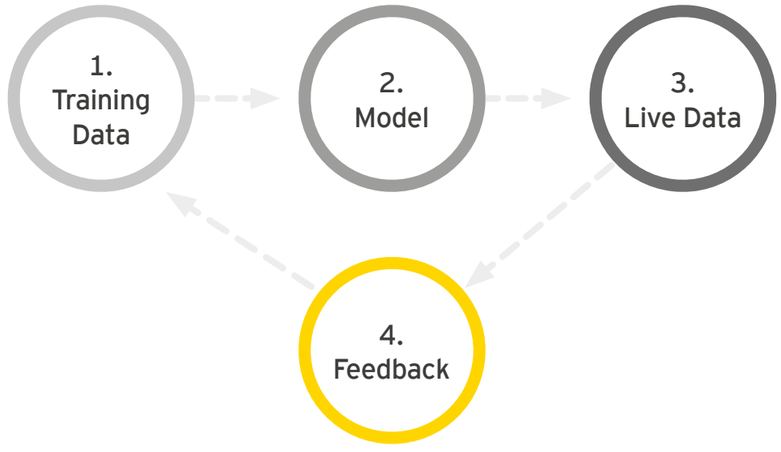

Over and above traditional assurance concerns such as confidentiality (who can access each stage of the above), integrity (how do I know data flows between each stage as intended) and availability (how do I know this is resilient and restorable) we have provided some of the potential risks at each stage of the machine learning process in a table over the next two pages.

Whilst the table summarises example risks for machine learning, there are likewise new risks for other emerging technologies (e.g. RPA, Blockchain, IoT) and the increasingly agile methods by which they are deployed (e.g. cloud, devsecops) that traditional assurance approaches do not address.

Example risks in machine learning

Stage	Example risks
1: Training data Data is provided to train the machine to teach it: ▸ what are the inputs (also known as features) that should be considered and ▸ what are the associated outputs for those inputs.	▸ **Data gaps.** There is a risk that historic data which is used for training the model is incomplete and does not correlate to the environment that the model will operate in (for example if new customer types are encountered). It may also be that what has been a good predictor of outputs in the past (e.g. gender for candidates for promotion) may not be a desired input to assist in generating outputs in the future. ▸ **Data bias.** If the training data contains biases, these can be inherited and repeated by the machine. Additionally, human labelling of training data can also introduce opportunities for bias to creep in. For example an entry system classifying anyone with the title of Dr as male was reported at a gym in the UK and even crowd-sourcing attempts to collect training data have been found to contain significant bias. ▸ **Data collection.** Proliferation of the IoT has meant that sensors are making their way into more and more of the everyday things we use producing vast amounts of data that can be used to understand more about us. Data collated in this way and scraped from online activity can be sold and used to profile an individual without them being aware that their data is being collected and used. Lack of transparency and rigour over the method by which data is collected or where it is being purchased from can present reputational and regulatory risk. ▸ **Feature selection risks.** The feature selection process presents a number of risks. There is a risk of selecting too many (which may provide a false sense of accuracy), too few (meaning the model can't deal with nuances or optimises on one feature to the extreme) or the wrong ones (even if they can provide a good explanation of historic patterns, they may not be a good predictor of future outcomes). There is also a risk of unintended cross-feature correlation to introduce unintended features (e.g. education history and family income can be a proxy for ethnicity in some areas).
2: Model Algorithms are used to interpret the training data and create a model which will predict an output when given an input.	▸ **Algorithmic explainability.** Even where a model offers a high level of accuracy the way in which it achieves this may not be easily explained given the complex and evolutionary nature of the algorithms used. Where models are used for things such as credit scoring there is a risk that unless appropriate ways of generating details of how it reaches decisions are included, it may be difficult to justify decisions or demonstrate their correctness if challenged. Guidelines in the new General Data Protection Regulation (GDPR) coming into force in May 2018 cite that, individuals should have the right to obtain an explanation of how a decision based on automated processing was reached[2]. ▸ **Algorithmic auditability.** As live models constantly evolve, learning from new data, if auditability is not designed in from the start it will be almost impossible to assure these models and the decisions they make after the event. There are currently no consistent guidelines in place for what should be done to preserve an audit trail (e.g. storing interim versions and details of random seeds used in training the model). ▸ **Algorithmic consensus.** For many applications it may be difficult to reach a consistent (and appropriately justifiable to all stakeholders) consensus as to what features are relevant (e.g. what features make a good candidate for promotion or when ranking universities, what makes a good university?). ▸ **Algorithmic suitability.** There has been a huge rise in 'democratisation' of data science with many companies offering off the shelf algorithms that can be quickly deployed rather than these being developed in house. There is a risk that they have not been designed appropriately or that they are used in ways they were not designed for.

[2] ICO. Big data, artificial intelligence, machine learning and data protection. September 2017. https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf

Stage	Example risks
3: Live Data Data is input into the model to obtain an output from the model.	▸ **Transparency of the use of machine learning.** The use of a machine-learning model may not be visible to the end user. Given increasingly important decisions and predictions are being made using machine learning, arguably from an ethical perspective there should be more transparency and failure to provide this could impact brand and trust. Even where a user can see their data is processed by a system (for example to choose whether to process their card payment) they may believe this is being done based on pre-defined rules rather than machine learning (i.e. deterministically not stochastically). ▸ **Accountability.** There is a risk that complementary or compensating procedures around the model are not put in place to identify, challenge and correct mistakes, bias and potential violations of law. Additionally, there is insufficient ownership for responsibility of the model (between those who designed, trained, approved and implemented it). ▸ **Data ownership.** There is a risk of a lack of clarity as to who owns data input into the model in terms of customising feature selection (e.g. if a user creates a bespoke recipe who owns that combination?) or labelling data (e.g. a user uploading additional data that might be used for future training). ▸ **Dealing with inaccuracy.** Machine learning solutions will be accurate to a certain degree of confidence but there is a risk that this isn't clear to users. Users may not be aware that an exact solution hasn't been achieved and they are being presented with a default or nearest match solution. Consideration may not be given to the relative impact of false positives and false negatives, for example in healthcare diagnosis.
4: Feedback Data is fed back to further train and refine the model.	▸ **Mutations.** Models can be 'broken' with sufficient malicious input (for example trained to classify a benign file as a virus or vice versa). ▸ **Unintended Feedback Loops.** Feedback loops themselves can reinforce bias and enact self-fulfilling prophecies. For example if a machine learning model provides an inappropriately low ranking to a university, this affects student applications, faculty recruitment and funding sending it further down the ranking. ▸ **Model selection.** There is a risk of sub-optimal decisions being taken as to whether to use historic models or update them for new data and how frequently to do this. ▸ **Missing feedback.** Where the model is used to exclude data there is a risk that there is no feedback as to how this data goes on to perform against model predictions (e.g. approval of credit).

If auditability is not designed in from the start it will be almost impossible to assure these models and the decisions they make after the event.

The key impacts of emerging technology on existing assurance approaches

Whilst developing approaches to each emerging technology in turn can provide useful guidelines for teams (where they land in isolation and this can be done quickly enough) we believe there are three more fundamental shifts in assurance approaches that need to be considered by assurance leaders:

From post to pre-assurance

Assurance after the event is increasingly irrelevant. Whether its machine learning models that can't be retrospectively audited, the risk of almost instantaneously processing millions of items incorrectly (but consistently) with RPA or the immutability of Blockchain. The impact of not assuring emerging technologies before the event will increase in line with the increase of the power and responsibility being entrusted to them as they are embedded into safety critical, or decision making, systems. Perhaps the most quoted example of this is a model used to support criminal sentencing in the US by looking at the likelihood of reoffending. This significantly under-predicted white males re-offending and over predicted black males based on questions which introduced bias into the algorithm[3]. Considering the impact of this example then merely detecting discriminatory decisions after the event will not be sufficient. Under the accountability provisions of legislation such as GDPR organisations will need to find ways to build discrimination detection into emerging technology to prevent such decisions being made in the first place.

Assurance after the event is increasingly irrelevant.

From timely to time limited assurance

Assurance teams spend a significant amount of effort in providing comfort over processes, profits and projects based on how well they are doing at a point in time and provide little comfort as to how long into the future the assurance will remain valid — what is the 'assurance decay'? If a continuously evolving model is working as expected now, what assurance do we have that it won't start producing erroneous decisions and predictions going forward? While this may be an implicit gap in how assurance is reported today, emerging technology will accelerate the need to address this. To achieve this, the scope of assurance plans and reporting need to evolve to address questions such as:

- What are the things that we have assumed remain constant for the assurance to be valid?
- What ongoing monitoring controls are there that the assurance and these assumptions remain valid?
- Are there any specific triggers which would cause us to revisit or revise this assurance as it would not be valid?
- What assurance is there over controls which cover ongoing change management and evolution of systems?

From data *analytics* to data *dialectics*

Dialectics is a discourse between two or more people holding **different points of view** about a subject but wishing to establish the truth through reasoned arguments.

Over the last decade assurance teams have increasingly attempted to use data analytics to improve the way they scope, risk assess and deliver their work. Even basic analytics have driven additional insight and comfort in areas ranging from fraud (e.g. ghost employees) to commercial benefits (e.g. duplicate payments). While many aspire to move towards more advanced analytics such as continuous controls monitoring, emerging technology significantly increases a challenge that has already slowed progress for teams in this area. Simply put:

- the 'black boxes' are getting darker. As we move into areas such as AI it is becoming harder to understand how systems are processing things; and
- the 'data exhausts' are getting bigger. Exponentially more data is being generated by technologies such as IoT.

While there will no doubt continue to be a role for traditional analytics moving forward (including over emerging technologies such as RPA), we believe that assurance teams should also develop a data dialectics approach — focusing less on testing what the system has done and more on what it could and should have done. A data dialectics approach involves both generating an appropriately granular independent expectation and using this and appropriate questioning to challenge and assure the output. To bring this to life:

- A simple example of generating an independent expectation in practice has been to predict store level revenue based on weather, footfall and advertising campaigns and using this to highlight stores reporting revenue out of line with central expectations.
- A simple example of using an appropriate questioning approach is querying a machine learning model to understand its sensitivity to changes in training data and for specific outcomes understand which features are most heavily driving this outcome and what would have to change to change the outcome. Even where the underlying model is inscrutable a data dialectics approach provides a step towards better algorithmic assurance.

[3] Angwin, Julia. Make Algorithms Accountable. The New York Times, 1 August 2016. http://www.nytimes.com/2016/08/01/opinion/make-algorithms-accountable.html?_r=1

Ethics — moving from an ethical dilemma to an ethical diorama

Our experience with assuring emerging technology so far has convinced us it is not enough to merely evolve the way that we think about existing assurance approaches but that we need to develop approaches to assure new areas — fundamentally we need to move beyond asking whether systems are doing things right to asking whether they are doing the right things.

There are a number of reasons for the increased relevance of ethics — for example the increased impact of emerging technology in areas such as health, privacy and government as well as the necessity to have principles that can guide us when rules and regulations lag behind technology and have not yet been codified to adequately assure and address risks. Increased interconnectivity has also meant that systems have the potential to impact a broader set of stakeholder groups than imagined or intended. For example, in the US the Federal Trade Commission found evidence of the credit limits of people being lowered based on the poor repayment histories of other people who happened to shop at the same stores as them[4].

As already stated, today digital trust and embedding ethics can be a key differentiator in the digital economy providing greater reach and competitive advantage. However, we believe that, in future, it will be a baseline consumer expectation as attitudes towards ethics and data continue to evolve- an 'order qualifier' rather than an 'order winner'. We believe stakeholders will also increasingly demand to be identified in a less homogenous way — both so they are confident that systems are making decisions that are tailored to them and so that they can be clear whether the system is treating them equally or equitably.

Assurance of ethics itself creates a dilemma — it's an area which has historically been uncomfortable for many assurance professionals (who have not felt they can hold themselves out as the bulwark of ethics) and which in its own right hasn't had a well-defined assurance approach. Ethics themselves add to this dilemma as they are emergent and approaches to embedding them into systems can fall on either the horn of hard coding rules which will necessarily evolve (the prospect of dealing with a technology hardcoded with ethics from the 1800s is hardly appealing) or asking the technology to infer ethics from data which may itself be inherently biased.

So how do we start to shift assurance from an ethical dilemma to an ethical diorama (i.e. a model representation of ethical assurance)? We believe that increasing transparency is key. In examining how ethics can be embedded in assurance processes through increased transparency, ironically a new three lines of defence emerges — firstly assurance that those people optimising and implementing these technologies are accountable and enabled to consider their end use, secondly that ethical matrices and other tools are embedded at the correct stage gates in design and finally that on an ongoing basis there are appropriate ethical processes and procedures put in place to ensure continued accountability between both the technologies and those that run them.

Ethics in business is by no means ground-breaking but with the proliferation of emerging technologies the stakes are much higher with a lapse in ethical behaviour whether intentional or unintentional having a much greater impact. What is considered ethical as mentioned above can change with time, location, legal and regulatory background and sociological changes and there is a lot of room for manoeuvre and potential risk. Ethical assurance over emerging technologies from conception to use, can help organisations demonstrate integrity, gain trust and reduce their exposure to risk.

Fundamentally we need to move beyond asking whether systems are doing things right to asking whether they are doing the right things.

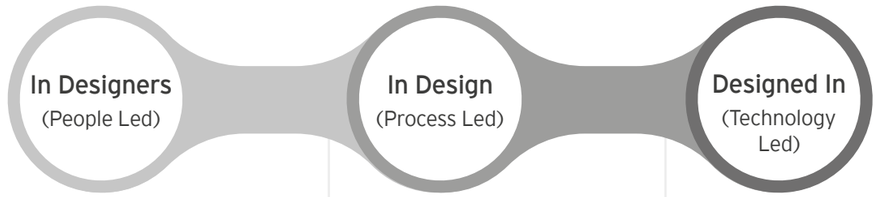

e.g. :

- Ethical training and standards for system designers and developers.
- Escalation and whistleblowing procedures where potential unethical use scenario detected.
- Increasing diversity of stakeholder views represented/ understood by designers.

e.g. :

- Ethical reviews considered at design/release stage gates.
- Development of ethical design principles/templates (e.g Open Data Institute (ODI)[5]).
- Independent 'red team' assessment of technology from broader group of impacted stakeholders.

e.g. :

- Hard coding of non-negotiable parameters.
- Design of compensating controls to allow issues to be monitored and dealt with outside of the system.
- Transparency to users of where system isn't 100% accurate and how it has made its decisions.
- Adequate accountability provisions.

[4] Federal Trade Commission. Big data: a tool for inclusion or exclusion. FTC, January 2016 https://www.ftc.gov/reports/big-data-tool-inclusion-or-exclusion-understanding-issues-ftc-report

[5] https://theodi.org/the-data-ethics-canvas

Three calls to arms for assurance leaders

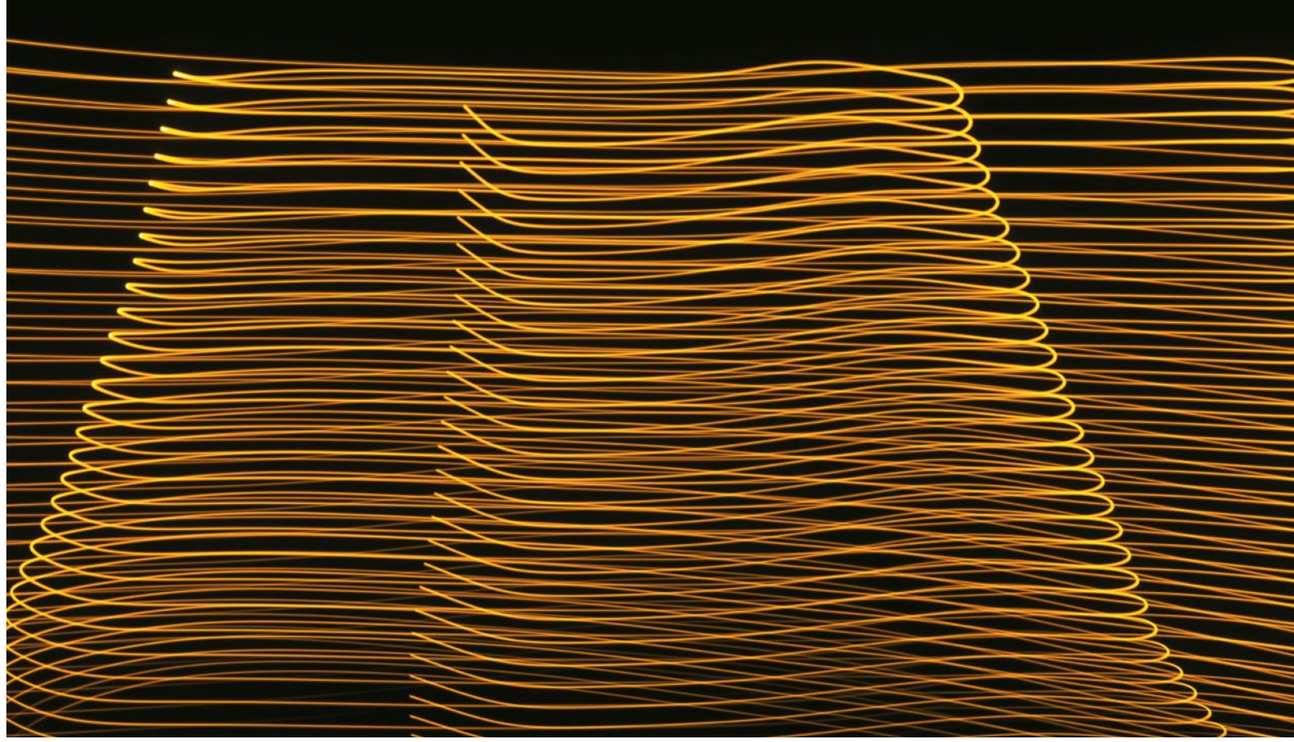

To the extent we accept that knowledge can't be codified quickly enough, regulation and accepted assurance approaches will continue to lag behind the rate at which emerging technology impacts our organisations.

This technology is already impacting our organisations and this will only increase — we need to quickly develop a plan that navigates a path between waiting (and potentially being too late) or over focusing on this at the cost of other areas that require attention. The reality is we have neither the luxury of doing nothing nor doing everything we would want to. We suggest three steps to consider in developing a practical response to assuring emerging technology risks.

Develop a rough map and start skirmishes

Starting work in this area is important both to address existing emerging technology risks as well as developing capability and confidence to deal with this as it increases in the future. In our work in this area we have found there are four key corners to consider in developing a rough map:

- **Verifiability:** What are the consequences of doing nothing now on our ability to assure but more importantly control this area in the future — will the horse already have bolted?
- **Visibility:** To what extent is the technology already understood with robust guidelines in place as to how it can be assured and controlled?
- **Value at risk:** What is the likely impact in the future of risks not being addressed in this area including the current direction of regulation (e.g. privacy)?
- **Velocity:** What is the speed of likely adoption and impact of this technology in the organisation in the future?

Having developed a view of where we should focus our efforts, it is important to start skirmishes early when we believe there will be an issue rather when there is one– agile assurance approaches such as white papers, project assurance, hypothesis testing and so forth provide one way of doing this while limiting resources committed.

Train the troops

From our own experience in developing approaches to assuring emerging technology we suggest three areas of focus to enable our teams to build the right skills to remain relevant to their organisations:

- **Give them first-hand experience:** *'The map is not the territory'* — teams can't prepare to deal with emerging technologies just by reading whitepapers (however well written and informative they might be…), attending breakfast briefings or webcasts. Training your entire team in becoming technical experts in data science isn't realistic either. To truly understand and be able to assure emerging technologies the team needs to get hands-on with them — this means seeing it in action, playing with it and gaining more than a superficial knowledge.
- **Develop effective communication and relationship skills:** The shift to pre-assurance may seem like a sensible step but for it to work in practice assurance professionals need to be involved up front. To do this they need more than ever to be able to build the relationships that will allow them to be invited to the table at the right time to stand shoulder-to-shoulder with the rest of the business — relying on assurance dictates and stage gates alone won't be enough to achieve this. Therefore as the deployment of emerging technologies increases so does the need for effective communication and relationship building skills in assurance teams.

Relying on assurance dictates and stage gates alone won't be enough to achieve this.

- **Train for higher order skills — the need to become more 'human':** Ethics is an area where we have clearly stated we need to collectively raise our game as an assurance profession in terms of embedding this into our assurance plans and therefore also in how we train our teams to understand and deal with this. However, we believe developing other higher-order skills will enhance the team's capability for dealing with emerging technology — whether that's in creativity (to help them find new approaches) or perhaps most importantly in how to deal with complexity. Even with today's technology, complexity is a key area where assurance often fails, for example gaps often occur in considering technologies' inter-relationship with other risks (e.g. master data, reports, application controls, interfaces). This will accelerate in the future and as *'simplicity does not precede complexity but follows it'* before our teams deliver off the shelf work programs we need to encourage them to stand back and to consider things such as these inter-relationships (between technologies, suppliers, risks, data to name a few). Therefore training teams to deal with and manage complexity (for example by training them in techniques such as problem-structuring methods) in order to design appropriate assurance will perhaps be the other key skill that makes a difference in the future.

We find it interesting that not all of the skills that assurance teams will need to further develop in response to emerging technology are technology orientated but believe that these skills will make assurance teams both more relevant to their organisations and better skilled for their future careers. If training and development plans can be clearly explained as such, training in these areas has an additional potential benefit of motivating and engaging people if they feel they are developing skills which will keep them relevant in an increasingly digital economy.

Adapt

As technology and organisations adapt we believe assurance functions must also move beyond the 'iteration' of the continuous improvement driven by measures such as effectiveness reviews and audit committee demands if they are to appropriately adapt. An approach we have applied to help assurance functions do this in practice considers adaptation across an additional two dimensions:

Iteration: This is an area most assurance departments already focus on to drive ongoing continuous improvement in existing processes by making them more efficient and effective.

Innovation: Choosing a limited number of 'big bets' where assurance teams can evolve or add value by doing something totally different. For example emerging technologies such as robotics have the potential for some more repetitive controls in frameworks such as SOX to be automated to allow more focus on other areas which require more judgement or are more complex.

Integration: It is difficult for assurance teams to have the resources to adapt alone and collaboration is another dimension which can allow them to do this more effectively. Working across the organisation and beyond (e.g. suppliers, peers) to keep up to date and where appropriate to collaborate with other initiatives and innovations can allow additional capabilities to be more quickly and cheaply developed and delivered.

In summary by starting skirmishes and focusing on those areas that matter the most first, by training for new skills and by adapting our assurance functions to remain relevant we believe assurance leaders will be better capable to move forward in an area where the map of how to do so is rarely fully documented.

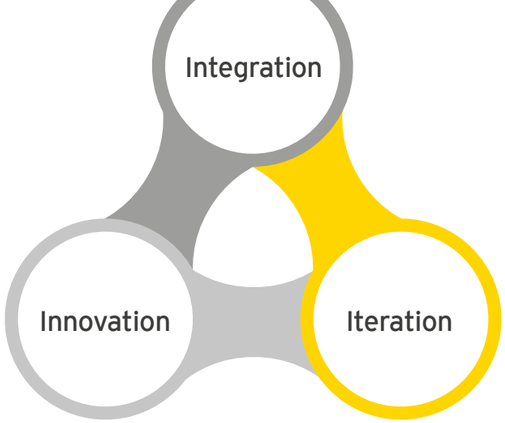

As a final thought we also believe assurance leaders will be more able to move forward through uncertainty if they have purpose – in this respect we believe that moving to consider ethical assurance provides both an additional opportunity as well as a challenge for leaders. In our experience for many who have chosen to work in the field of assurance it is personally important for them to be doing the right thing and making a difference. We believe that articulating the risks of emerging technology (perhaps especially the ethical ones) and the impact that assurance teams can have in protecting both organisations and individuals through assuring these will help motivate assurance teams to push through uncertainty and develop and deliver more robust approaches to assuring emerging technology in their organisations.

Conclusion

The pace of technological change is bringing with it unparalleled opportunities for companies to disrupt themselves and enter new markets. The promise of greater productivity, efficiencies and the elimination of human error is well documented. Less well documented are the new risks that emerging technologies are creating for organisations. The speed of adoption, complexity and ubiquity of these technologies means that these risks are rapidly increasing in both likelihood and impact and moreover often going unnoticed.

Current assurance approaches alone are insufficient to address these risks. Assurance leaders urgently need to engage with their stakeholders and the rest of the organisation to understand how emerging technologies impact their organisation now, and in the future. Resulting changes to assurance scopes and approaches require new skills and capabilities that assurance teams need to start developing today to remain relevant for the future. As part of this, ethical assurance will be key to help ensure that in embracing these new technologies organisations are confident that the way in which they are doing so is consistent with their brand and culture allowing them to demonstrate integrity and build essential digital trust.

Contacts

This paper has provided a few thoughts gathered at the beginning of 2018 as to how assurance approaches need to adapt to emerging technologies and the steps assurance leaders should take to achieve this in their organisations – like any perspective it will over time become outdated, less relevant and unfit for future assurance needs. We look forward to continuing to contribute to and help shape the debate. To discuss any of the topics in this paper (and how they have evolved since it was published) please contact one of the team below.

Richard Brown

UKI IT Risk and Assurance Leader
rbrown@uk.ey.com

Pragasen Morgan

Data Privacy
pmorgan@uk.ey.com

Piers Clinton-Tarestad

UKI IT Risk and Assurance Innovation Leader
pclintontarestad@uk.ey.com

Mike Rudberg

Assurance Transformation
mrudberg@uk.ey.com

Sofia Ihsan

Emerging Technology Assurance
sofia.ihsan@uk.ey.com

Chris Voogd

Assurance Analytics
cvoogd@uk.ey.com

Stuart McMeechan

Risk Analytics & Machine Learning
smcmeechan@uk.ey.com

Kevin Duthie

Global IT Risk Quality Leader
kduthie@uk.ey.com

EY | Assurance | Tax | Transactions | Advisory

About EY

EY is a global leader in assurance, tax, transaction and advisory services. The insights and quality services we deliver help build trust and confidence in the capital markets and in economies the world over. We develop outstanding leaders who team to deliver on our promises to all of our stakeholders. In so doing, we play a critical role in building a better working world for our people, for our clients and for our communities.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. For more information about our organization, please visit ey.com.

Ernst & Young LLP

The UK firm Ernst & Young LLP is a limited liability partnership registered in England and Wales with registered number OC300001 and is a member firm of Ernst & Young Global Limited.

Ernst & Young LLP, 1 More London Place, London, SE1 2AF.

EYG No.
ED None

Artwork by JDJ Creative Ltd.

In line with EY's commitment to minimise its impact on the environment, this document has been printed on paper with a high recycled content.

ey.com/uk