

Chambers Ireland submission on the development of a National Artificial Intelligence Strategy

Executive Summary

Chambers Ireland is a business representative organisation, our members are the chambers of commerce in the cities and towns throughout the country. Each of our member chambers is central to their local business community and all seek to promote thriving local economies that can support sustainable cities and communities.

The current generation of Artificial Intelligence (AI) technologies will have profound effects on our economy and our society. Many good, some bad, most unavoidable.

Chambers Ireland wants the proposed National Artificial Intelligence Strategy to be sensible and grounded in reality. There are strong limitations to what people call AI and the hyperbolic discussions that surround the industry are rarely helpful in assessing those limitations and could unintentionally provoke the creation of regulatory environment which is hostile to the development of AI tools.

There are two bad regulatory environments for AI. Firstly, regulation could be used to protect incumbent industries which are at risk of automation from the forces of competition particularly in non-traded sectors. Support for such a regime could emerge where sectors and industries which have previously been protected from external competition suddenly have to accommodate new entrants to the market, whether they be home-grown competitors or competitors which originate from abroad. The temptation may be to erect barriers to competition which will be defended using fears about AI technologies. Ultimately this will result in our domestic economy's productivity diminishing relative to competitor nations.

Secondly, there is the risk of making categorical error regarding the nature of the AI technologies which are available to us and creating a regulatory regime that is based on a mischaracterisation of what these tools are, leading to potential local innovation in the field of AI being further incentivised to offshore.

Underlining this is that AI is both an opportunity and a risk to the Irish economy, but it is inevitable. If we are not fully engaged with the developments in the field of AI then we will find that economic and state actors will be ultimately end up using sub-optimal tools which have been developed for other populations, in other countries and we will be blind to the problems that they carry with them in their wake.

Recommendations:

- Regulation of AI should be proportionate, non-protectionist, non-industry specific, and grounded in the reality of these AI technologies
- Certain industries will be severely affected by automation, there should be a departmental programme that aims to identify such industries at the earliest opportunity
- A framework for cross-sectoral risk analysis and assessments must be developed
- Sectoral level recommendations should be made
- National Training Fund resources should be directed towards people who are currently working in such industries so that they can upskill and retrain
- Improvements need to be made regarding accessing training schemes for those who are currently in work
- The departments of Employment Affairs and Social Protection, Education, and Business need to collaborate to ensure that pathways through continued employment are supported for those in vulnerable sectors
- In-work training schemes and tailored transition educational options need development
- Introduction of a voucher model for funding future-proofed skills development courses targeted at SMEs
- Increase investment in entrepreneurship and innovation education for secondary students
- Increase investment in career guidance to ensure that young people are aware of the future risks and opportunities that are arising from the digital economy
- A cross-departmental framework for the publication of non-personal public data needs to be finalised
- A body, such as the CSO, should become the state body which holds all non-personalised public data, providing data services to all departments and state bodies, ensuring that local departmental data structure idiosyncrasies do not inadvertently create data silos
- Public data should be viewed as a public resource; therefore, state bodies should internalise the principle that sets of data should never be restricted to a single service provider

- The Department of Public Expenditure, and the Office of Government Procurement should develop a strict AI tool procurement policy to ensure that the considerable risks to the public purse that may be involved in poorly applied AI tools will be mitigated
- Every utilisation of an AI tool, product, or service by a state body should involve a risk assessment, a bias mitigation plan, and public, independent verification
- The SME favourable principle of think-small-first must be applied to all AI GovTech
- The Department of Enterprise, Business and Innovation should create an awareness campaign about the potential benefits, and risks, associated with AI tools, aware that the benefits are most likely to accrue to the non-traded sector
- More AI related in-job training for ICT specialists in SMEs should be supported by the National Training Fund
- For people with non-ICT backgrounds, more ICT skills and AI skills training programmes should be developed by the Irish third level sector, and the Institute for Public Affairs in particular.
- Our existing anti-discrimination law should form the basis of a regulatory response to AI-tools
- The existing legal framework for countering discrimination should be the foundation for any regulations which aim at anti-discrimination regulation of AI technologies
- Where someone chooses to use an AI tool, then an action against them under the ESA should remain possible
- Regulations should be formed in such a way that the individuals or organisations which use AI-tools remain responsible for the consequences of using these tools
- The Department of Enterprise, Business and Innovation should create tailored awareness campaigns about the implicit and explicit legal obligations associated with the use of AI-tools for business and for consumers.
- The user of an AI tool should be required to be able to demonstrate the efforts they took to compensate for unlawful bias, and the consequent mitigating efforts they took upon establishing that discrimination had occurred

Contents

Background	5
An overview of Artificial Intelligence (AI)	5
The limitations of AI	8
Regulatory Risks	9
The risk of protectionism	9
Bias and the risk of category error	10
The impact of AI on Ireland	12
Automation as a Threat/Opportunity	12
Irish Enterprise and AI	14
Top AI concerns for Irish business	16
GovTech AI opportunities	17
Priority actions for government and industry	19
Ethics of AI/governance issues	21
Cybersecurity	24

Background

An overview of Artificial Intelligence (AI)

In a discussion about Artificial Intelligence, the first element that needs consideration is what are we actually talking about, when we are talking about AI.

Artificial Intelligence contains many associated and related activities, including data science, deep learning, machine learning, autonomous vehicles, cybernetics, robotics, etc.

Most of these, and the current generation of 'Artificial Intelligence' products are forms of applied statistical learning: Given a sufficiently large data set, a parameterised model may be fitted to a set of data such that it can identify associations between different data elements, or rather, given a certain set of data it can, with a likelihood probability that arises from the model trained on that data, that model may classify new data according to 'known' patterns that originate from the primary, training dataset. The model 'learns' by altering parameters internal to it, so that its success at assigning the correct classification is maximised.

What it 'learns' or 'knows' depends on the approach taken. Some models are the consequences of datamining techniques.

Unsupervised learning is the most abstract of these and doesn't involve human labelling of the training data which are fishing expeditions that delve into a large pool of data to see what if anything interesting turns up. Network analysis is an area where this approach is useful, say if you were looking for patterns of company ownership, or were looking for data bottlenecks in your computer network, or perhaps unusual patterns of money transfers within financial data, this approach may throw up patterns which are distinct from the typical parts of the network.

This approach is useful for throwing up potential hypotheses, or suggested associations. Such associations would then need to be tested on independent real-world data to determine their validity.

Others involve supervised learning, where the data is labelled in advance, and the primary data is interrogated to determine if there are associations with defined outputs. Typical of this would be exploring the genetic data of tumour biopsies and to see what correlations there are with patient outcomes.

Other models are trained using reinforcement techniques. Autonomous vehicles are a good example here; the models are trained to avoid having the human driver intercede and take control of the vehicle. They create what amounts to a library of similar circumstances and then take the path of actions that is least likely to require driver intervention.

Most of the activity that falls under the heading of "AI" is:

- a classification problem, or
- some form of optimisation under a constraint problem, or
- a combination of these.

The degree to which a model is successful at this classification is often determined by the nature of the problem. Some problems are very amenable to solving.

For example: Spam filters - They have gone through a revolution where once they targeted particular key words, they now use as a data set, all the emails that the emailing system has seen tagged as spam, and predicts the likelihood that a new email fits the 'spam' pattern.

When you receive new spam and tag it as such, or untag something incorrectly tagged by the software as spam, it becomes part of the learning data set the next iteration of the spam detection software uses to learn from.

Some are very hard to solve.

For example: Fully automated vehicles – They are still not available. There are lots of circumstances where they do work, e.g. on highways - which are pretty simple, with one-way traffic, low relative speeds, well defined boundaries, etc. - they work. Put them in a shopping centre car park, with trolleys, kids, conflicting traffic flows, and they don't¹.

¹ <https://arstechnica.com/cars/2019/06/at-tesla-shareholder-meeting-musk-assures-there-is-not-a-demand-problem/>

And then there are the impossible problems:

For example: The weather – Met Eireann can improve their weather forecasting, but no number of neural networks is going to let us predict the weather more than a few days out, and even then, only in the simplest of circumstances.

Together, this places constraints on the utility of Artificial Intelligence. For narrow self-contained things where there is complete information – like the ads on Google – they can be incredibly useful, as Google ‘knows’ what you’ve clicked in the past, what you were searching for when you were presented with the ad, and what you have searched for in the past. Together Google can present you with the ads which optimise the money that Google receives. Google has knowledge of everything that happens within that space.

Then there are complicated systems with incomplete information, such as getting a robot to walk. There are many kinds of environments where the robot must walk where there is no data to work with. So, the robot uses generalised rules and the designers hope that they apply to new situations, slowly the system gathers more and more data, but where there will never be complete knowledge.

Then there are the complex environments such as the economy, where prediction efficacy is likely to be low, and even where it is useful is likely to be of limited use (in time and/or range).

Knowing when and where to apply the tools of AI requires both an expertise with AI and an expertise with the subject domain.

Without such individuals who hold this kind of joint expertise, then it’s unlikely AI solutions will be effective, either because the AI people will not be aware of the problems that need solving, or they will not understand the problems sufficiently to be able to produce something useful.

The limitations of AI

All AI solutions involve human intervention to some degree. People are involved in defining the problem that is to be solved, coding the data that is being used for classification, or assessing the ecological validity of the patterns that a model throws up.

Many of the problems that are associated with AI are a consequence of this human activity. There are biases in the initial collection of data, if we look at where AI tools have been applied to crime data in the [Great Britain](#)² we can see many of the problems that can arise when we ask AI to solve problems which the technologies are ill-suited for.

These tools are created with data that are sourced from an environment where there are data collection issues. Certain geographical areas are under-policed, and other over-policed. Certain crimes may be prosecuted in some areas and ignored in others. Racial factors can be implicit in geographical areas and may be used as just cause to defend 'random' searches that are far from random.

AI tools are probabilistic in their mechanics, not certain. And, unless someone is very, very careful they can be influenced in all kinds of non-obvious ways. Therefore, care needs to be taken before they are applied to a problem, and institutions/individuals who choose to use them need to monitor whether they are malfunctioning, or misapplied. We all have a responsibility to ensure that we do not discriminate against others in unlawful ways, regardless of how discrimination might occur. This holds true for AI as much as it does for any other means by which a person might be implicitly or explicitly biased.

Those who intend to use the tools ought to educate themselves about the limitations of these technologies so that they can protect both themselves and the public from these limitations.

Policy makers ought to be similarly cautious, not only in where they apply these tools, but also in how they decide to regulate them.

² https://rusi.org/sites/default/files/20190916_data_analytics_and_algorithmic_bias_in_policing_web.pdf

Regulatory Risks

The risk of protectionism

Certain industries are likely to experience a profound shock as these new AI tools are put on the market. Industries like underwriting, trading, banking, brokerage are probably going to see working practices transformed. Businesses with legacy liabilities and large overheads are probably underprepared for potential revenue declines. Just as Kodak was unprepared for the change in business model that digital cameras implied, some firms will struggle to meet the challenges of upstart competitors that will rely on the automation-driven efficiencies which AI tools will offer.

It is hard to see what government policy can do for such sectors, though it is likely that the relatively high-skilled workforce displaced by such technologies could find other work opportunities in the services sector. However, it may be necessary to review the educational and skills development structures which are available to those who are currently in work. Furthermore, attempts should be made to identify industries-at-risk-of-displacement at the earliest opportunity so that appropriate programmes for upskilling, and retraining can be developed for them in a proactive fashion.

Recommendations:

- Regulation of AI should be proportionate, non-protectionist, non-industry specific, and grounded in the reality of these AI technologies
- Certain industries will be severely affected by automation, there should be a departmental programme that aims to identify such industries at the earliest opportunity
- National Training Fund resources should be directed towards people who are currently working in such industries so that they can upskill and retrain
- Improvements need to be made regarding the accessibility of training scheme for those who are currently in work

Bias and the risk of category error

Another area where AI problems can be expected to emerge will be around bias and discrimination. This is because people are at the heart of AI. AI technologies are essentially ways of making tools that people will use.

While the process is complicated for the lay person, at a base level, AI tools are just matrices which take in numbers and output others. The mathematics involved are explainable with a Higher-Level Leaving Cert degree of complexity, and numerous platforms have developed that allow people to apply these tools without having to dig into the mathematics.

There are a couple of mechanisms which can allow for AI tools to be discriminatory. Explicit and implicit biases can appear in the data that feeds the development of these AI tools, or tools can be applied in circumstances which are outside of their envelope of validity.

In areas where bias already exists there can be structural reasons as to why the data that currently exists contains bias; consider the employment of women at higher levels of management, if your algorithm was optimised to select for candidates which are most likely to be employed in a certain industry, based on past success rates, then you might find the algorithm to be automatically excluding women from shortlisting. Similarly, different credit risks could be assigned to similar people based upon where they live rather than who they are.

The risk of bias is inherent to these tools, because that's how they are supposed to work: People who are of lower credit risk should find it easier to be approved for a loan. But unfortunately, some biases are spurious, and applying these tools in the right way needs expert knowledge of the domain.

This creates a profound problem for creating 'unbiased' algorithms; algorithms ought to be biased, just not in bad ways.

If an individual or institution, be it a state body or private business, chooses to use AI tools, then they should firstly be aware of the potential for unintended bias, and secondly test, monitor, and compensate for biases that emerge in the algorithms they are using.

This means that it is the outcomes that people experience which should be regulated, not necessarily the algorithms, or the AI tools, themselves. If the technologies are regulated then this will stifle local development, meaning that the tools which will ultimately be used in the Irish context will have been developed for other populations, and so may have undetermined, and uncorrectable biases in the Irish context.

The Equal Status Act already provides for protections against discrimination on a number of grounds, regardless of the cause of such discrimination, if such discrimination occurs as the result of a bad AI tool, or because of the misapplication of an AI tool then that discrimination is already unlawful.

Recommendations:

- The existing legal framework for countering discrimination should be the foundation for any regulations which aim at anti-discrimination regulation of AI technologies
- Where someone chooses to use an AI tool, then an action against them under the ESA should be possible
- The user of an AI tool should be required to be able to demonstrate the efforts they took to compensate for unlawful bias, and the consequent mitigating efforts they took upon establishing that discrimination had occurred

The impact of AI on Ireland

AI is to be considered in two ways, firstly there are the threats associated with automation, which is an old threat, but one that will hit new sectors.

Secondly, it can make untenable problems amenable to solution, thereby facilitating affordances previously unavailable.

As a strength, we have an IT base which is strong for an economy of our size, and many large multi-national companies use Ireland as their base.

As a weakness, our domestic supply of people with AI and IT skills is low, we rely on foreign talent which narrows the possibility of cross pollination from the multi-national sector to the domestic sector, furthermore it is hard for local SMEs to compete with multi-nationals for pay, working conditions, and professional opportunities.

Our small market also makes us niche in terms of software solutions. Even if a new product is created and is successful here, it's hard to globalise, as once the idea has been shown to work, it's easier for the Irish start up to be bought up by a large multi-national than it is for an Irish start up to compete with a multi-national that emulates the successful business model, but with an international marketing reach.

Often our highest potential start-ups migrate to the larger global centres of IT so that they have access to the capital and skills they need to scale. This means that while the innovation may arise here, the value is often accrued in other jurisdictions.

Automation as a Threat/Opportunity

Some great work was done by Frey and Osborne (2013)³ on "The future of employment: How susceptible are jobs to computerisation?" where they hypothesised that certain types of activities are less amenable to automation than others. They highlighted:

- Perception and Manipulation
- Creativity, and
- Social Intelligence

³ https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

as hard to reproduce and then ranked occupations about how dependent they were on these kinds of activities.

Having mapped these against US jobs data, they found that Office and Administrative support, is very likely to be wiped out through automation, Sales and Related, is very likely to disappear, much of the service industry will go, as to will the majority of Transport and Material Moving. Such a comprehensive mapping onto the Irish domain has yet to be conducted.




FIGURE I. A sketch of how the probability of computerisation might vary as a function of bottleneck variables.

Frey and Osborne (2013)

Obviously, AI related job redundancy will only be a sub-set of the total number of jobs that are vulnerable to automation, but they highlight some ways in which sectors will be affected.

With some companies starting to supply off-site manufactured modular building units we can see the start of this shift. Traditional joiners are having to upskill to produce pre-fabricated units in a factory setting, while the design work is becoming increasingly complex but standardised. School units can now be designed using drag and drop software, in the long run this will reduce the input of the higher skilled design professionals while artisan workers who construct the units to tight specifications are having to be trained in on new construction techniques and technologies.

Through initiating a shared framework of building standards, the Northern Ireland department of education reduced the degrees of design freedom so that pre-built units from different suppliers could be seamlessly integrated, to a degree it reduced the creative element and but has facilitated automation. In so doing they are creating stable, high-skilled, and well-paying offsite construction jobs that facilitate a good work/life balance.

Already in the Irish domain we can see companies like Vizlegal creating research products for the legal sector that will reduce the lower skilled work element within a highly skilled domain. No doubt there will be fewer low skilled legal jobs, but with the other elements of legal work involving creative argument, as a profession, it is not likely to be automated away, and is likely to see a productivity boom.

These developments suggest that there needs to be a sectoral level analysis of the kind of work that people do, so that we can understand what will happen for the workforces involved. IT induced automation has been fundamentally changing our working environments for decades, but the tendency has been for these changes to amplify the productivity of those who are working in the affected sectors.

IT efficiencies have not led to sector-wide job destruction, but tend to create new opportunities further up the value chain, and at a macro level new opportunities in the service sectors.

Recommendations:

- A framework for cross-sectoral risk analysis and assessments must be developed
- Sectoral level recommendations should be made
- The departments of Employment Affairs and Social Protection, Education, and Business need to collaborate to ensure that pathways through continued employment are supported for those in vulnerable sectors

Irish Enterprise and AI

Irish enterprises are likely to be underprepared for the AI shocks that will hit their sectors. We already know about the productivity crisis in medium sized businesses (50-250 employees). The EU Commission highlighted the “low percentage of the workforce with basic digital skills reflects the insufficient integration of digital skills in the education and training system” as a particular concern for Ireland in their “Country Report Ireland 2019”⁴ With basic skills they refer to knowing how to use email etc.

⁴ https://ec.europa.eu/info/sites/info/files/file_import/2019-european-semester-country-report-ireland_en.pdf

We have a large cohort of workers who have almost no digital skills, and those that have skills are in huge demand with companies which the Irish domestic sector will find it almost impossible to compete with.

There needs to be an effective way of upskilling those who are in the workforce, and those who are outside of it. There are good signs within the small and micro-industry businesses regarding productivity as they are outpacing their sister companies of medium size in terms of productivity and use of technology, but our indigenous non-traded services sector is likely to be affected very profoundly by AI.

There are three trails that need to be blazed in integrating AI with Irish industry.

- Firstly, there needs to be more home-grown data scientists.
- Secondly, there needs to be more home-grown data engineers.
- Thirdly, indigenous firms need to look towards ways in which these new technologies can complement their businesses.

As an economy however there may be issues arising in the non-tradeable services. While industries like accounting, financial services, insurance, and real estate are very likely to be transformed by new technologies over the coming years, others industries are not as easily mapped onto the Irish market (legal services, public administration, health services) and so are not as likely to benefit from the new technologies because on the one hand they will not have competitive reasons for applying the changes in work practices involved, and on the other there will be institutional incentives to resist change. This means that we are likely to lose out (as a whole economy) competitively as our cost base will likely remain high, relative to those country that do secure a deeper integration of AI technologies.

Recommendations:

- In-work training schemes and tailored transition educational options need development
- Introduction of a voucher model for funding future-proofed skills development courses targeted at SMEs
- Increase investment in entrepreneurship and innovation education for secondary students
- Increase investment in career guidance to ensure that young people are aware of the future risks and opportunities that are arising from the digital economy

Top AI concerns for Irish business

Certain sectors are likely to see a huge proportion of their work automated away.

Such as:

- Lower skilled professional work (bookkeeping, legal research, etc.)
- Office and Administrative support
- Sales and Related Services
- Transport and Material Moving

Within these industries, there will be much disruption, which is more likely to affect companies with legacy capital holding, staff who will not be appropriately skilled and retrainable. These firms are likely to lose out to smaller firms with lower overheads that can benefit from the near zero marginal costs of acquiring new customers.

However, while the small size of the Irish market will abate some of this. In the absence of a thriving indigenous industry this means at a macro level there is the risk that only AI products which have global market reach will be available to the domestic economy.

A side effect is that products tailored for British, or US markets, but dominating here could have sub-optimal outcomes in an Irish setting, both in terms of their applications, but also structurally as the outsourcing of AI tools could result in amplifying the productivity crisis in our medium sized firms, which are already lagging in terms of “Digital Skills”.

Where the scale of jobs in a sector is likely to decline this will most likely be felt by the lowest skilled members of the sector – new entrants for professions and sectors which are undergoing a technological transformation are likely to find it more difficult to get the entry level jobs that will give them a foothold in that industry.

It is therefore vital that careers advice in schools can be informed by sectoral level analysis so that young people can plan their career trajectories informed by where the labour opportunities will be, rather than where they have traditionally been. While not all of the ‘safe bet’ jobs will be affected by AI initiated automation, the sectors that have previously been seen as safe will be disproportionately affected by the innovations that do arise.

GovTech AI opportunities

There are huge productivity gains to be made from deeper implementation of AI product (in areas where sufficient data are available). Regarding these large data stores, typically they are held by large companies, the state, or are part of open data sets.

Government data, and open data sets will likely be the primary sources of indigenous industry opportunities. With this in mind it is vital that Irish government services fundamentally review their relationship with the public's data, and the data of the public services. Clearly GDPR is the framework which defines the public's data and such data needs to be protected, but with respect to public data, departments will have to become better at sharing the non-personalised, or aggregate, data and statistics that are at the core of the operations of our public and services, if they are to garner the benefits of innovations in GovTech.

In the absence of a healthy, local, AI and startup ecology state bodies will be forced to use AI tools and products which have been developed for other jurisdictions and trained on datasets which will not necessarily match the Irish population.

A government wide open-data strategy (for non-personal data) will also have the benefit of helping to ensure that the products and tools which state bodies use can be verified. An issue with AI products and tools is that they are 'black boxed'; it is often unclear why the matrices they use to channel input data towards outputs. Typically, they are not readily analysable at the functional level, but rather at the behavioural level – it is easy to interrogate their outputs, but not how the programmes reached those outcomes. As a consequence, and particularly in the monopsony context which state bodies often find themselves to be in, in the absence of an open-data context the state body can only trust that the AI tools they are using are working correctly, but there is no possibility of verifying that these tools are operating correctly.

Should state bodies continue to use the "walled garden" approach, where certain organisations have limited access to restricted datasets, not only will they not be able to fully gain from the potential benefits of GovTech they will also be ignoring the known risks of using AI tools which will open up state bodies to liability if it ultimately emerges that there was a bias in the tools that they have been using.

Irish state bodies can learn from the experience of other jurisdictions when it comes to AI GovTech. When they go well, AI technologies can have large upside benefits, but this also comes with large downside risks, particularly regarding uncertainty about:

- whether the final product can work

- whether the final product does what you think it is doing
- biased datasets
- unintended consequences

The classic error here was with Palantir developing a product with the Mayor's office in New Orleans⁵ which was offered to the city for free, with a service plan attached should they subsequently decide to use the tools (Palantir ultimately rolled out the service to police departments across the states, and internationally).

This scheme transformed from being a diversionary project aimed at allowing social services to involve themselves in the lives of people who are at risk of violence, into a pre-emptive policing project which justified search and seizure activities which sought to arrest people even when there was no complaint against them. This was due to the unintended consequence of using pre-existing bias from the police force, which ultimately became incorporated into the AI product.⁶

The law of unintended consequences should be keenly considered when thinking of where the state can collaborate better with business using AI, as any current bias implicit in the system will be amplified, and this will be statistically provable if there is a software package deployed, which would likely become a serious liability for any state body that utilised such a product.

Recommendations:

- A cross-departmental framework for the publication of non-personal public data needs to be finalised
- A body, such as the CSO, should become the state body which holds all non-personalised public data, providing data services to all departments and state bodies, ensuring that local departmental data structure idiosyncrasies do not inadvertently create data silos
- Public data should be viewed as a public resource; therefore, state bodies should internalise the principle that sets of data should never be restricted to a single service provider
- Every utilisation of an AI tool, product, or service by a state body should involve a risk assessment, a bias mitigation plan, and public, independent verification

⁵ <https://www.theverge.com/2018/2/27/17054740/palantir-predictive-policing-tool-new-orleans-nopd>

⁶ <https://www.theverge.com/2018/3/15/17126174/new-orleans-palantir-predictive-policing-program-end>

- The SME favourable principle of think-small-first must be applied to all AI GovTech

Priority actions for government and industry

Considerable attention, energy, interest and resources needs to be put to educating the public, the business community, and officials regarding the nature of AI, the limitations of AI tools, and the opportunities that they present.

IT projects are typically hard as they involve people, who are often expert within their domain, trying to communicate their needs with others who do not conversant with that domain. AI projects have the added issue where even to those who are working on a project, the functions of the final product can be entirely opaque.

Even honest actors can find themselves talking at cross purposes, not even realising that miscommunication is occurring. Therefore, even more than typical IT projects, the lions share of the work is done in the conceptualisation of the problem, not in the solutions phase.

The human element in creating AI tools must be highlighted at every opportunity. AI tools are domain specific and are limited by the abilities of those who define the problems they aim to address and are limited by the skills of those who create the solutions. AI is mostly applied statistics; it is not magic. And most people have a legitimate take their statistics with a pinch of salt.

Ideally the creators of AI tools would have both domain expertise, and also AI expertise, but there is a very limited set of individuals in any field which have both. In the absence of those people, data scientists will be working with domain experts who don't understand AI, and the inverse will likely be true. This means that businesses, and government, must be very wary of overpromising by the industry.

Ironically this overpromising is also a risk to the AI industry as threats which are not real tend to dominate the discussion (i.e. we see this with the spurious application of thought-experiment trolley problems to autonomous vehicles; an inappropriate application of a philosophical 'intuition pump' that relates to utilitarianism, to the problem of autonomous vehicles, which completely ignores how an autonomous vehicles choose their routes).

What will hold Irish government and industry back is the lack of people with ICT skills. AI specific needs are a subset of ICT related needs, they depend on Ireland having people with a deep understanding of IT, and also mathematics/statistics. The wider the number of people who have an understanding of IT, the wider the pool of people who can work within AI will be.

But to satisfy that demand we will need to be upskilling some people who are already very highly skilled individuals.

There are local skills shortages for higher level data scientists, there are few data engineers which significant skills, and most businesses are unaware of the opportunities, and the threats, of AI to their businesses. These skills gaps are not easily bridged by recruitment from outside of Ireland as many businesses, even large multinational firms, are finding it difficult to recruit due to the costs of housing.

The decision to teach coding (through python) to leaving cert level is good move in expanding the supply of those with ICT skills and is welcomed, but this year's 16 year olds will still be entering the ICT jobs market in seven or eight years' time, and when they do they will have a world of opportunity open to them. The recruitment challenges will not abate quickly.

Recommendations:

- The Department of Enterprise, Business and Innovation should create an awareness campaign about the potential benefits, and risks, associated with AI tools, aware that the benefits are most likely to accrue to the non-traded sector
- The Department of Public Expenditure, and the Office of Government Procurement should develop a strict AI tool procurement policy to ensure that the considerable risks to the public purse that may be involved in poorly applied AI tools will be mitigated
- More AI related in-job training for ICT specialists in SMEs should be supported by the National Training Fund
- For people with non-ICT backgrounds, more ICT skills and AI skills training programmes should be developed by the Irish third level sector, and the Institute for Public Affairs in particular.

Ethics of AI/governance issues

The primary issue with AI is that is really applied statistics, and statistics can be inaccurate or misrepresented. There are many reasons why a statistical result may be wrong, though usually these errors arise because you have collected the wrong data or are asking the wrong questions of your data. This can be because there is a fundamental misunderstanding of the problem that is under consideration, or because the data itself is not representative of the population you are interested in. More usually the problem is a combination of the two.

As with any other tool, AI can have its productive uses, or it can be used recklessly, or it can be wielded as a weapon.

We should try to resist the urge to treat AI products as a class apart from other IT tools. We must consider instead how they are being used. If, as a hypothetical, somebody created a product for the Irish market that constructed a credit point system that rated people based upon; credit history, court records, known addresses, social media data etc. this would likely (unless the creator took express efforts to the contrary) prejudice people from protected minority and ethnic backgrounds unless they had controlled for these issues (this is an issue which has occurred in many other countries).

The developer in this case would likely be reckless in their development process having not considered the biases in the data involved.

If someone used this product to deliberately avoid employing people from such backgrounds (as has occurred in other countries) then they would be in breach of our existing law.

Fundamentally, our National AI Framework should apply our existing Employment, Equal Protection, laws and regulations to AI-tools.

There may be value in supporting a voluntary industry organisation modelled on the Advertising Standards Authority, or the Association of Irish Market Research Organisations, which draws up guidelines for businesses in that field which would also highlight the risks associated for business, the public, and government regarding AI.

There also needs to be increased efforts within the International Standards Organisation ISO/IEC JTC 1/SC 42 Artificial Intelligence programme.

And inevitably EU regulations will apply.

Ethics and governance are important to consider when looking at how AI tools are applied, but there is a risk of moral panic reifying problems which don't really exist regarding what we are now calling AI. Data protection is obviously important, but that is covered by GDPR legislation, and it is vital that we do not allow for discrimination, but this is already covered by the equal status acts.

What is important regarding AI is a deep consideration of how these tools are used, and that the users/developers know the limitations of AI-tools. Subsequently, it will be important for users to monitor these tools for error, and then to adapt the tools to prevent these limitations becoming manifest (just as they should so with any other IT tool) - ultimately individuals, companies, and organisations will be liable if they use bad AI-tools, or misapply them in ways that are unfair to others - simply because the causal element is software will not protect them from liability.

If government chooses to regulate in such a way that it regulates the technology, rather than the application of that technology we not only risk stalling local efforts in the field of AI. This is not only a risk to business and industry, if AI tools are difficult to develop here because of regulatory requirements, the inevitable consequence is that Ireland will become dependent on AI tools that are trained on datasets that originate in other populations, and so will contain the biases associated with those populations without controlling for the biases that exist in our home population.

There needs to be a massive educational effort made to demystify what AI tools can do, and to enlighten people as to what the limitations of these tools are. This will be important for consumers, businesses and officials as all are likely to be targeted by unethical actors that will be promoting products that are founded on fictions, 'caveat emptor' is particularly relevant to AI-tools.

Just as with other IT-tools, AI-tools will be used to automate things which people already do, helping them focus on the important elements of their work. When people are choosing to apply these tools to their work it is essential that they are responsible for their choices, and it is important that a regulatory framework that looks towards AI does not absolve them of their responsibilities to apply these tools wisely and appropriately.

A real risk associated with over-regulating the sector, is that such regulations will become a perfunctory check list, and that meeting those requirements will be a sufficient defence to claim that due diligence was conducted. No check-list style approach will be suitable for commissioning or applying AI-tools as no regulation will be able to keep pace with the range of tools available to those working in the field. Furthermore, the appropriate use of AI-tools involves regularly monitoring their behaviour to identify their deficiencies, because they are essentially fallible tools which give probabilistic results not certain ones.

As AI-tools are probabilistic tools (much like with medical screening) they will always give a certain percentage of false positives, and another percentage of false negatives. Deciding what the tolerance for these errors will be requires a person judgement which is founded on the likelihood of that form of error, and the costs associated with getting it wrong. Where these costs or risks of making such an error cannot be borne (say at the individual child level with regard to child protection services), then [that is an area where it is inappropriate to use AI tools](#)⁷.

Individuals need to know these aspects to AI-tools before they choose to use them, and they need to bear responsibility for applying them in a given circumstance.

Recommendations:

- Our existing anti-discrimination law should form the basis of a regulatory response to AI-tools
- Regulations should be formed in such a way that the individuals or organisations which use AI-tools remain responsible for the consequences of using these tools
- The Department of Enterprise, Business and Innovation should create tailored awareness campaigns about the implicit and explicit legal obligations associated with the use of AI-tools for business and for consumers.

⁷ <https://www.communitycare.co.uk/2018/03/29/artificial-intelligence-childrens-services-ethical-practical-issues/>

Cybersecurity

The vulnerability of Ireland, as a target for criminal and political cybersecurity threats has become a business concern. Though included within the departmental headings for this consultation it is not primarily an AI concern. It will however become a concern to both those who are developing AI tools, and those who will use them.

AI tools will be most useful when it automates tasks which would otherwise absorb out time and energy, and so free us up for more useful tasks. But by removing the human from the function these functions will become vulnerable to cybersecurity threats, as these tools come to permeate deeper into our personal and professional lives, they will become more attractive targets for malefactors,

For those developing AI tools and services, considerable efforts will need to be developed to create these products and it is vital that Ireland has robust, secure digital and legal structures for these companies and entrepreneurs to work within.

The companies based here, who contribute a substantially to the Corporate Income Tax which the state receives need their Intellectual Property, AI tech and otherwise, to be secure.

Furthermore, many of these firms are also headquartered here, reporting to the Irish Data Protection Commissioner (DPC) regarding their data privacy and GDPR obligations. With so much of the world's data housed here it is paramount that we ensure that the data we are entrusted with remains secure. Irish security threats and vulnerabilities risk becoming security threats for other countries which could hurt Ireland's standing as a safe base for such countries. Ultimately, the long-term presence of these data intensive companies depends on the capacity of our public sector respond to cyberattack threats.

Ireland needs to ensure that there is a robust cross-departmental public sector strategy to pre-empt these threats. Such a strategy should see the development of a civilian cybersecurity agency which collaborates with Justice officials, the DPC, industry, security experts and our third level sector to identify, monitor, and respond to cybersecurity threats