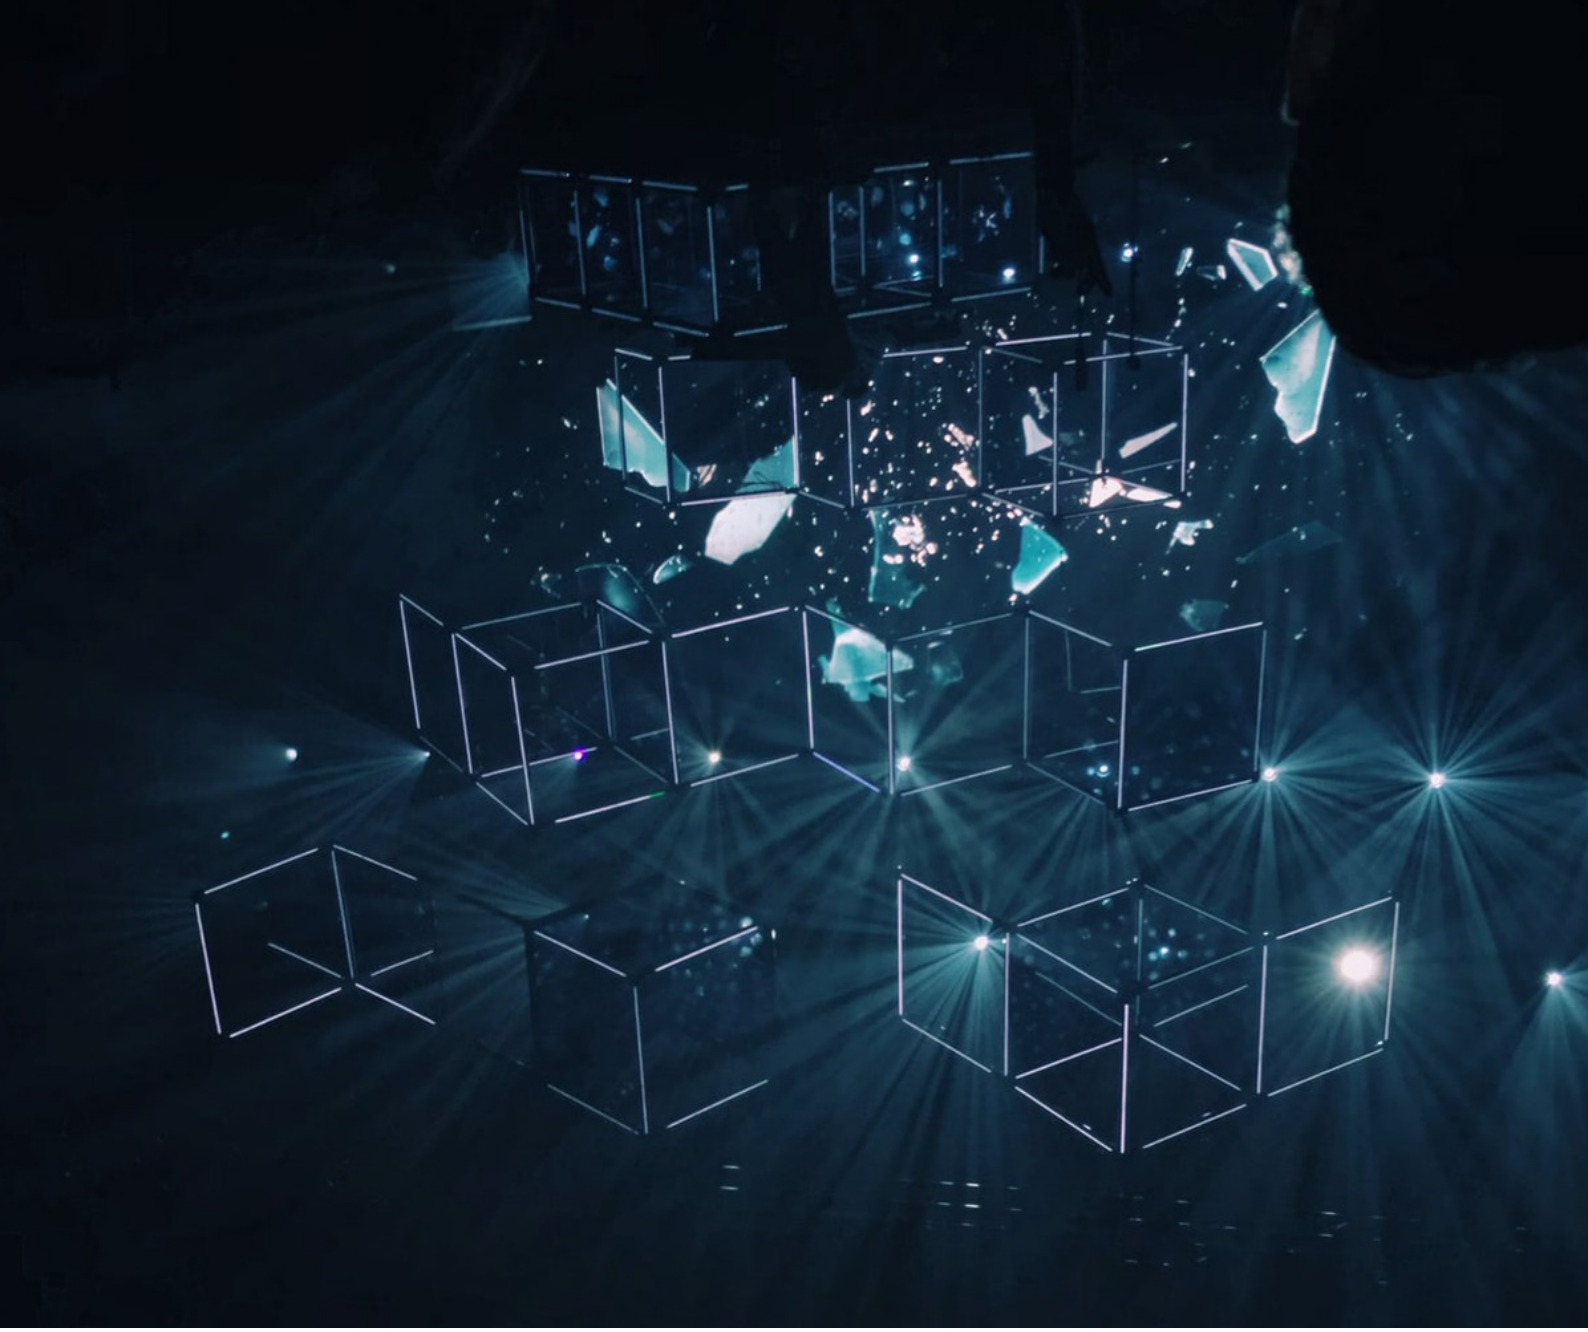

AI in Government

— Gavin Whyte

Chief Artificial Intelligence Officer,
Public Sector Network

Publicsectornetwork.co (02) 9057 9070

PUBLIC SECTOR NETWORK

Table of contents

1. Executive Summary

1. Executive Summary

AI in Government

In the midst of the COVID-19 pandemic the strategic importance of Artificial Intelligence (AI) in Government has become more evident than ever before. This has been clearly identified through the use of AI with mobile phone and geolocation in contact tracing. While COVID-19 still rages, there are multiple examples where we could implement AI in Government.

AI could improve the delivery of public services in areas like data quality, healthcare, transport and procurement.

But how do Governments position themselves to take advantage of AI powered transformations?

This white paper looks at key findings that came out of conducting interviews with 41 participants from various Departments of Government. We thereafter look at what an Ethical Framework in production will look like, and next steps for Government, in terms of implementation and adoption of an AI framework. Note this document is a collation of various descriptions of information used by AI.

2. Key Findings

2. Key Findings

Below is a summary of the key findings that were discussed amongst the interview participants.

2.1 Adoption of AI

One of the key aspects of a successful AI deployment is data, and it is quite evident that Government generates large volumes of data. While the adoption of AI varies in many divisions of Government, starting from AI thinking all the way to production of AI models, certainly the thinking and framework around the adoption of AI is becoming a key aspect across many government divisions. Considerations have been taken into account around the impacts of AI on current workloads. Lots of interesting thought processes and strategic thinking have gone into the application of AI, from start-up accelerators to proof of concepts. The use of automation and AI from basic reporting through to deep analysis, were key features during the discussions.

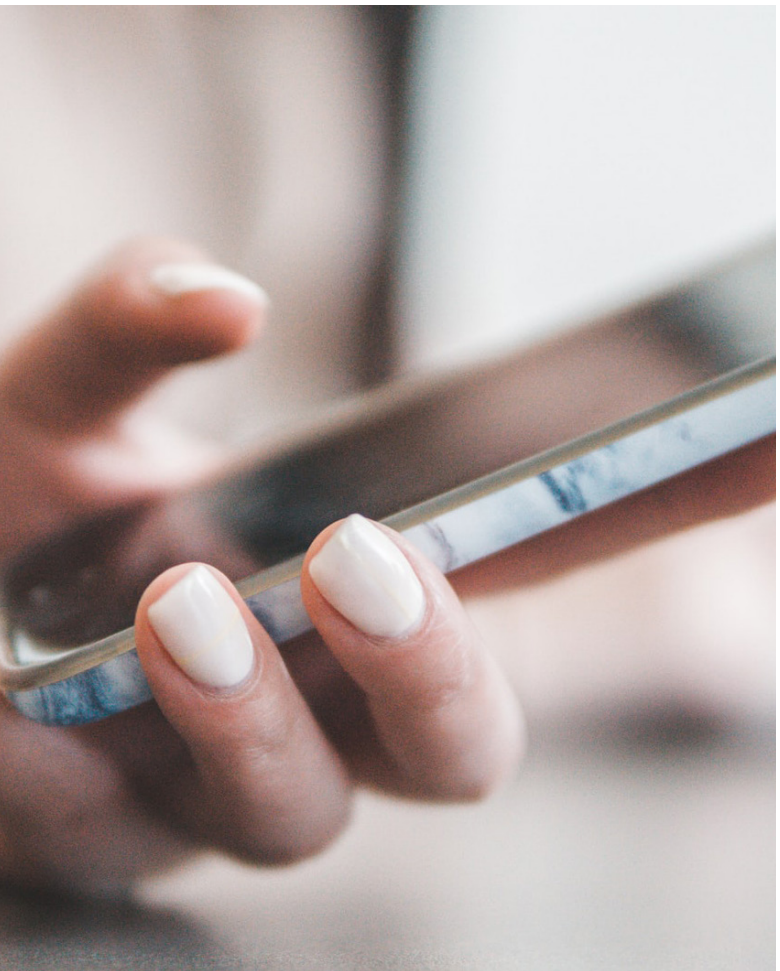

2.2 AI Technologies

When presented with the importance of various AI options like Cloud Intelligence, Augmented AI assistance and Autonomous AI, a vast majority of the participants favoured Cloud and Augmented AI rather than Autonomous AI.

While participants indicated its beneficial use and thought leadership, a preference was given to the adoption of AI today in Government. This was based on the maturity of their data policies and the gentle introduction of AI into their environment in a controlled manner, via an Augmented AI methodology. Traceability and auditability along with explainability and data privacy were key attributes that were emphasised. The importance of machine assisted behaviours augmenting the workforce was quite a clear winner in terms of usage, along with the implications and impacts to the workforce.

Organisational strategy moving towards a sophisticated cloud-based platform to scale up due to the large volumes of data was another major talking point. Data ethics and regulatory functions were also key features during these discussions. While we can explore AI as an umbrella, and how it fits in the toolkit. AI as an umbrella has far reaching benefits for Government, which needs to be explored further as strategic sessions for each government division.

The ability of leadership thinking around freeing up talent, with the augmentation piece of AI along with RPA, that allowed for the automation of current business processes, especially during these unprecedented times, would change the way government employees work and provide better insights to run an efficient division. At times participants spoke about AI confidence levels needed to be attained in terms of false positives and improving accuracy for those that have implemented the AI process.

2.3 Spending budgets

While a lot of government divisions have used business-as-usual (BAU) budgets to develop proof of concept projects, it was quite interesting to observe those that utilised BAU for simple experiments, and applied for project funding, were successful in their AI deployments. Whether to use BAU or CAPEX, is always the question when deploying from experimentation all the way to AI OPS production grade deployments. But many divisions did get the necessary CAPEX project budgets to fund AI projects.

During the discussions, a very important aspect around the automation of the existing workforce to free up resources to take on multiple tasks during the COVID-19 period was widely discussed. A detailed business plan was essential to describe the opportunities that AI presented, along with the efficiencies. This helped tremendously with CAPEX funding.

Senior executive buy-in was key around project funding and development of business use cases. Divisions have already acquired their funding aligned to organisation strategy, so a well-defined strategy is also key to delivery.

2.4 Privacy Impacts

While a large number of participants indicated that data was important, privacy was not really seen as a problem, as long as the divisions are being transparent. Adopting a privacy framework, as indicated by participants, was nonetheless going to be a key factor in establishing trust, transparency, and accountability.

Protecting citizen's information and developing deeper trust in Government is critical to maintaining and protecting their data. When data is shared, the ability of de-anonymisation should not be conducted only where permission is acquired for internal purposes. Good organisational thinking along with the establishment and activities of ethics committees and a focus on ethics, are key.

Having an audit trail is also an important activity during this process of usage. The issue of unstructured data anonymisation did arise during the discussions as an important activity, along with ensuring data governance and privacy policies are always up to date.

Protecting citizen's information and allowing citizens to develop trust in government, on the back of the pandemic, is important. By further extending this trust factor on the processes of how we maintain and protect data can either build trust or quickly erode public opinions. It's the former that is critical.

From a privacy and regulatory perspective, the single most important issue that was discussed was explainability and accountability. In all cases, data is sovereign to Australia and each department comes with responsibility and accountability for decisions. In any situation, the citizen is entitled to explanations through the decision-making process.

Deep learning was discussed with a few participants in detail, specifically around the processing of large volumes of data and the benefits that deep learning brings to Government. While powerful, it has a level of complexity. Clear broad guidance on data assets and how to implement proper deep learning along with the explainability, was critical.

RoboDebt also came up several times during this conversations, specifically around what not to do, although this provides a bit of insight into automated decision-making, and how having an ethical decision-making framework would have played a critical role. The ability to codify decisions is critical. The concept of privacy by design as a foundation, with built-in frameworks that protects human rights and implements ethics, was also critical. It is crticial we don't conflate privacy and security. Trust vs Surveillance e.g. RoboDebt (risk).

The issue of having a clear understanding of data privacy was important in every strategy and framework design. The kinds of privacy protections that are in place, and working with teams specifically that are authorised and use data, is key. Using technology and analytics to track access, in terms of privacy is critical. Options to build-in compliance or compliance by design, were also seen as very important.

Data-sharing also came up for discussion. Data stewards are key to managing the application of policies to data. Whether data is being made for internal purposes or for public use, it should be managed and treated as an asset. Privacy by design is not a constraint.

While participants indicated more work was required under information sharing legislation, the most challenging tasks are data tagging and classification, and can this process be automated given the vast amount of data that Governments hold.

Defining good data management practices and the ability to easily discover data, privacy provisioning through privacy by design, and security by design, were also seen as critical. The circumstances of deployment are a very important consideration, eg Health vs Security / Regulator vs ATO. It is important from a health perspective that requirements on certain data points will be required, like age, sex, ethnicity vs security and the implications to various government divisions needs to be taken into account.

The contextualisation of a federated model needs to be carefully considered. Upholding responsibility in the use the AI tooling is important. Explainability is critical i.e. chat bot informing decisions. Citizens have the right to understand the tech being used.

2.5 Cloud Computing and the use of big data

The use of big data is critical in processing large volumes of data generated by Government.

As we may have noticed in the news, NSW agencies will be expected to adopt a "public cloud by default" approach for all future IT procurements under a shake-up of the state government's cloud computing policy. Many divisions are currently in the process of applying a cloud first policy. While this is certainly an emerging space in terms of technology, participants indicated that they are currently in the process of migrating to the cloud, or more importantly, are on that journey. This is quite evident with a number of divisions utilising software like Azure and AWS cloud services to run loads and utilise the cloud for storage.

There are still a large number of divisions that run critical loads on premises, or are still looking at migrating to the cloud. Data governance and protection will play important roles in the migration of data to the cloud, along with data access from on premises or remote, with the necessary security mechanisms in play. While maturity is low in the use of cloud computing technology, talent plays an important part in the balance of business, technology, costs, data security and privacy. Clear guidelines are required in the use of Cloud computing and data protecting.

While cloud can be used for computationally intensive work, data sharing and migration from legacy systems are important considerations during the migration process.

Participants spoke about the implications of cloud first policy. They were concerned with future-proofing and having the ability to scale up without buying more software. The ability to re-architect is also going to be critical. For instance, how can SAS-type applications be moved to the cloud, or augment existing application to the cloud?
The following scenarios, will need to be taken into consideration:

1. Hybrid Cloud Scenarios
2. Ability to be aware of risks and breeches
3. Cloud Sandbox Environment
4. Departmental culture around cloud testing and implementation of cloud use cases and the culture in practise
5. The use of visualisations like Tableau, Qlik Sense and Power BI were all important aspects in the migration to cloud.

2.6 AI Strategy

To develop a good AI Strategy, there are several components to look at. With every AI strategy there is always a successful data strategy that needs to be in place, along with data governance, data access, data security and data privacy. Participants clearly pointed out that having the data strategy was a critical component in developing a good AI strategy, and this provides a better understanding and roadmap.

Some of the outputs that were discussed include:

- Use case identification.
- AI environments – Where do we start?
- Development of an ethical AI framework was absolutely key.
- How is the AI strategy relevant to the broader business strategic objectives?
- The ability to conduct education and strategy sessions for key stakeholders to become AI ready, bringing leadership teams on the journey.
- Identify gaps and provide training where necessary.
- Develop an AI framework to provide guidance and outputs.
- Look at identifying the low hanging fruit in terms of use cases to prove out technology and developing a business case for CAPEX funding.
- Identify tools necessary to deliver AI.
- Developing governance frameworks and models.
- Addressing the current operating model.
- Developing or updating a master data management framework.
- Developing a quality assurance framework.
- Developing a uniform architecture and data dictionary.
- Development of an open data science platform, along with guidelines.
- Work with broader strategic divisional objections.
- List several automation use cases that will free up individuals from menial tasks.
- Provide details around linkage to enterprise data platforms.
- Defining project key activities.
- Developing automated human decision-making frameworks.
- Review of the current business operating model in parallel to developing a strategy.
- Develop proof of concepts with various tooling and algorithms.
- AI(precision) + Automation(scale) + Ethics(process) will be the key recipe in achieving success.

2.7 Benefits of AI

While the benefits of AI are significant in government, some of these benefits were articulated by the participants in the following ways:

- Augment human effort to create value at scale.
- Improve the quality of services.
- Reduce paperwork.
- Improve citizen enquiries.
- Provide personalised services in government.
- Improving precision in analytics.
- Working better, smarter and faster, and improving effectiveness of government services.
- Delivery of ethical services for the community.
- Prediction is a key feature of AI to deliver successful outcomes.
- Doing more work through automation and utilising existing staff for high value tasks.
- Demonstration of exactly what the government does, through audit trails.
- Development of new skills and harnessing insights.
- Providing richer data fields and better insights.
- Spotting patterns of effectiveness and propensity matching.
- Operational efficiencies
- Understanding the current state, and targeting low hanging fruit.
- Developing AI strategy to scale workforce.
- Developing an AI playground for producing the best outcomes that are repeatable.
- Supporting the workforce, not replacing the workforce.
- Improving data literacy.
- Providing specialised AI training to different roles.
- The ability to explore data and realising that the benefits can be explored further
- Preventing work safe injuries with better predictive techniques, and optimised pathways.
- From a policy perspective, providing an ethical AI framework that is repeatable.

2.8 Risks of AI

A number of risks were also articulated by the participants:

- Getting the government workforce ready for the AI era.
- The growing complexity of AI technologies.
- Funding AI technologies.
- Growing concerns over algorithmic risk, black box, and model bias.
- Implementation of ethical AI.
- Discrimination of social groups, during model learning.
- Wrapping ethical guidelines around socially acceptable protection.
- Business readiness – need to approach AI not just in the context of analytics but from a business context and how is will be integrated into business flows.
- Appointment of responsibility and accountability of what the machine tells us.
- Know what is going on with explainability.
- The ability to understand whether the current environments can support AI driven capabilities.
- Getting senior stakeholder buy-in.
- Assessing a skills gap.
- Evaluating security and privacy risks.
- Ethical governance.
- Human sense and loss of control.
- Losing jobs.
- Integration and system risks.
- Funding for AI projects.
- Data security is a critical part of success.
- Capturing accurate information.
- Legality of AI and the ownership of AI.Legality of AI
- Risks in commercial readiness

2.9 Training, Tooling and Upskilling of Staff

Training, tooling and upskilling were seen as critical components to the success of the implementation of AI.

While several agencies encourage and deliver programs of work to discuss AI and its implementation, there is an inherent skills gap in both the strategy and delivery of successful AI platforms for government. There will be effort required in reaching a higher maturity for AI readiness. Teams and the HR divisional structure are key to evolving into a successful AI division, along with key stakeholder support.

Upskilling is certainly a key agenda item for most divisions, from training and understanding more about AI. Talent introduction from industry into government was discussed in several divisions. Overall, a lack of talent was clearly articulated amongst many divisions, along with the limitations of skill sets.

Many participants voiced their opinions on the establishment of a community of practice, saying it will be critical to drive a successful AI strategy in Government.

From a tooling perspective, repeatability, sandbox environments, framework and AI guidance, were seen as critical to successfully deploy AI into production.

3. Ethics Research Framework

3. Ethics Research Framework

The following is a summary of the thinking that goes into an ethical AI framework. It will assist departments and divisions to conduct an actual AI use case.

- *Algorithmic impact assessment of predictive systems*

When implementing an algorithmic impact assessment of predictive systems, the system should be able to make the existing prioritisation process faster and more accurate by automatically weighing the information provided by the departments in the context of the overall population, when assessing people's levels of vulnerability.

- *What are the typical models required to build an ethical AI framework in Government?*

The impact of each data point is a predictive modelling exercise that is learned and optimised in the model training phase of the predictive risk modelling. Each model has to be trained to recognise the optimum mapping between sets of predictors (data points used in the prediction process) and the corresponding outcome (actual prediction known as targets).

Each prediction can be utilised in many areas in government, from data quality to improving citizen's license renewal processes, or trying to book an appointment at the public hospital for surgery.

In terms of the patient optimisation journey at the hospital, currently this process is done by an individual. That individual has to take the criticality of the surgery, doctor's schedules, theatre bookings and many other data points into consideration around the patient's vital health statistics. This process could be developed with an ethical AI framework, but will need to be fair and regulated.

- *The establishment of an ethical framework*

To continue with the patient optimisation example, the risk assessment provides medical staff with relevant information about how to classify individuals at risk of experiencing an adverse outcome if not attended to in a timely manner.

On this basis, a set of the highest score individuals are provided with public patient bookings for surgery, including an optimised designated time and date for patients.

Along these lines, the main aim of the algorithmic system is prioritisation, which is conducted partially based on the individual medical information. The score resulting from the processing of these data points, together with data already held by the hospital, is used as part of the decision-making process. In particular, the risk score associated with each patient is used by medical administrative workers to conduct the final assessment and evaluate eligibility concerning the selection and provision of booking appointment.

The system is then designed to automate the booking of the recipient of the service automatically, but it assists in decision-making by providing human staff with a risk score which helps to rank patients for surgery.

- *Aims of the audit and methodology.*

The methodology for this analysis is mainly oriented towards determining both the accuracy of the algorithm in regard to its expected goals, and to identifying potential discrimination derived from its decisions. To achieve this, before the final phase of the analysis, three main steps were followed.

Firstly, the model design was reconstructed and the theoretical and methodological basis established for the algorithmic model. In this context, literature was reviewed on patients requiring surgery to have a basis for valuing the desirability of the model. The developing team provided valuable inputs and answers to the information requests to complete this process.

Secondly, an analysis of the situation was developed with optimising patient journeys that required surgery. This was aimed at framing the main factual information leading to optimised booking, where the system is implemented, as well as constructing hypotheses about ways of accurately measuring the risk of optimised patient booking. The above activities were also oriented towards establishing a set of hypotheses about algorithmic fairness, accuracy, and discrimination.

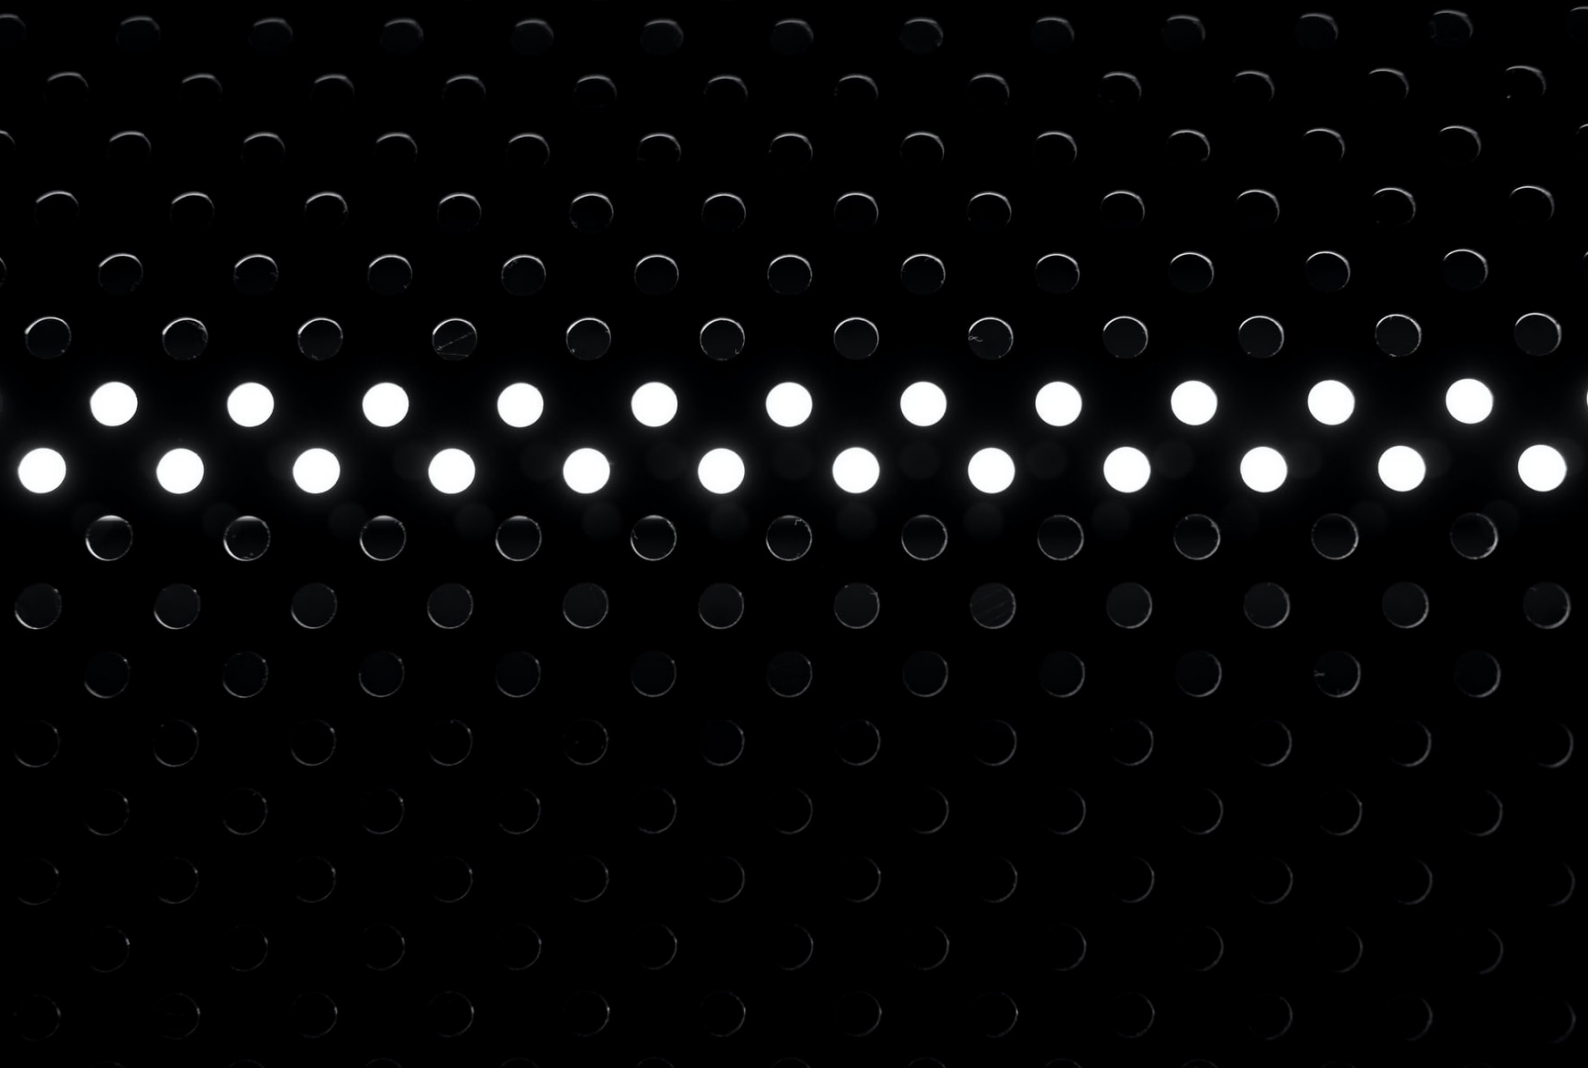

Thirdly, on this basis, these steps were followed:

a. define an assignment of elements in the data to patients,
b. define critical patient groups,
c. determine a set of metrics aimed at measuring bias, and
d. measure and compare across patient groups.

- *Algorithmic discrimination*

There are different approaches to defining algorithmic bias. From a research perspective there are those definitions that stress the unexpected character of algorithmic processing outcomes according to its predefined aims. "Algorithmic bias occurs when an AI model, trained on a given data set, produces results that may be completely unintended by the model creators."

There is also the perspective that stresses the incompatibility between the outcomes of algorithmic processing and its expected goals, and is greatly defined by contextual social factors beyond the creators' expectations, determining what can be considered as fair or ethical.

Algorithmic bias happens when "the data used to develop and refine algorithms reflect implicit values of the society in ways that are judged as irrational or unfavourable."

"Algorithmic bias is added by the algorithm itself and not present in the input data."

Beyond these conceptual perspectives, the literature increasingly recognises the importance of social or medical factors and the prominent role of technological developers in detecting and mitigating algorithmic bias.

The importance of implicit values in data collection, selection and use in this process has been stressed. This is one of the reasons why research has put particular emphasis on the ethical issues related to AI systems.

In order to frame algorithmic bias, different forms of discrimination should first be distinguished in terms of medical concepts. For instance, the vital signs that will improve the patient's chance to get a surgery booking at an optimal time.

There is a significant amount of thinking that goes into developing ethical frameworks for each Government Department.

4. Where to begin - Next Steps

4. Where to begin – Next Steps

- What are some of the typical frameworks used to develop AI strategy, business use cases and to track success?

At Public Sector Consulting, we have developed frameworks, along with the use of the BrewAI automation framework for creating sand boxes to prototype AI use cases. This will help departments to develop business use cases for CAPEX funding, or to utilise a sandbox for proof of concept, or to develop a project road to successful AI implementation.

- Typical AI project delivery lifecycle

With the application of a typical project delivery lifecycle, it is important to look at the following:

Qualify

- What are some of the opportunities where you will derive benefits from AI, endorsed by business owners?

Execute

- Qualification of Opportunity Backlog and define approval for business case
- Opportunity endorsed along with delivery and road map

Embed

- Solution stabilised in production
- Solution hand over to business operation team
- Solution measurement with determine solution trajectory to achieve business case benefits.

Solution

This is a tried and tested delivery lifecycle which includes operational cadences and check-points, as well as commercial models, customised to the organisation's preferred way of working. Along the journey, Public Sector Consulting and BrewAI tooling can make recommendations towards a strategic simplification of an organisation's objectives more broadly.

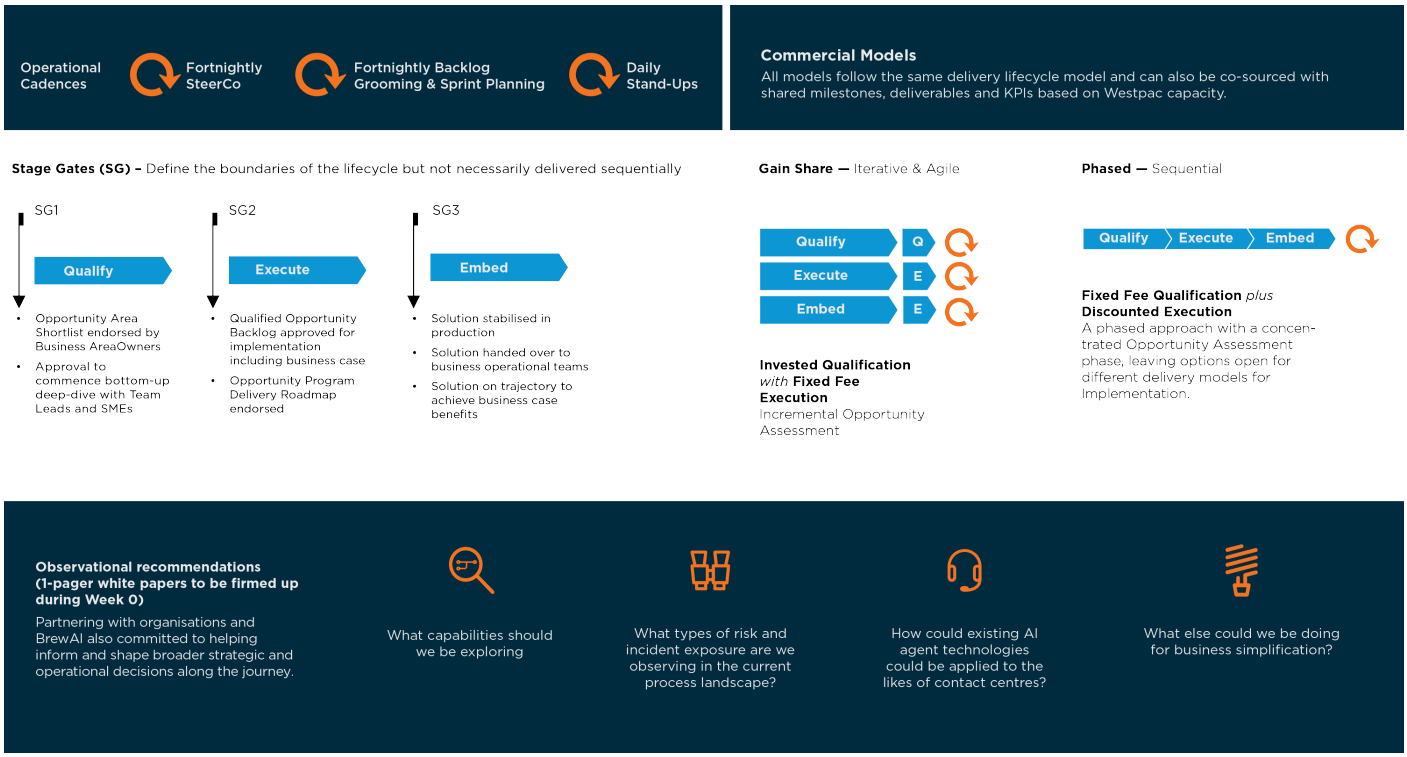

5. Conclusion

5. Conclusion

Developing an ethics framework and an AI framework for a successful AI implementation requires a number of factors to be taken into account. Understanding the business process; automation of menial tasks; the development of an AI Strategy; the application of AI to use case to achieve automation; development of an ethical AI framework with accountability; the proof of concepts and final production and business operation handover, are no easy tasks. This is on top of having a master data strategy. Having the right support, framework and process will help tremendously. National and State, strategical and ethical guidance is critical from an AI ethics point of view, currently ethical guidance is under review. Privacy by design is a critical factor within AI methods and this must be clearly distinguished from security.

The information presented in this white paper will allow the public sector to think and help define the processes required for a successful AI implementation.

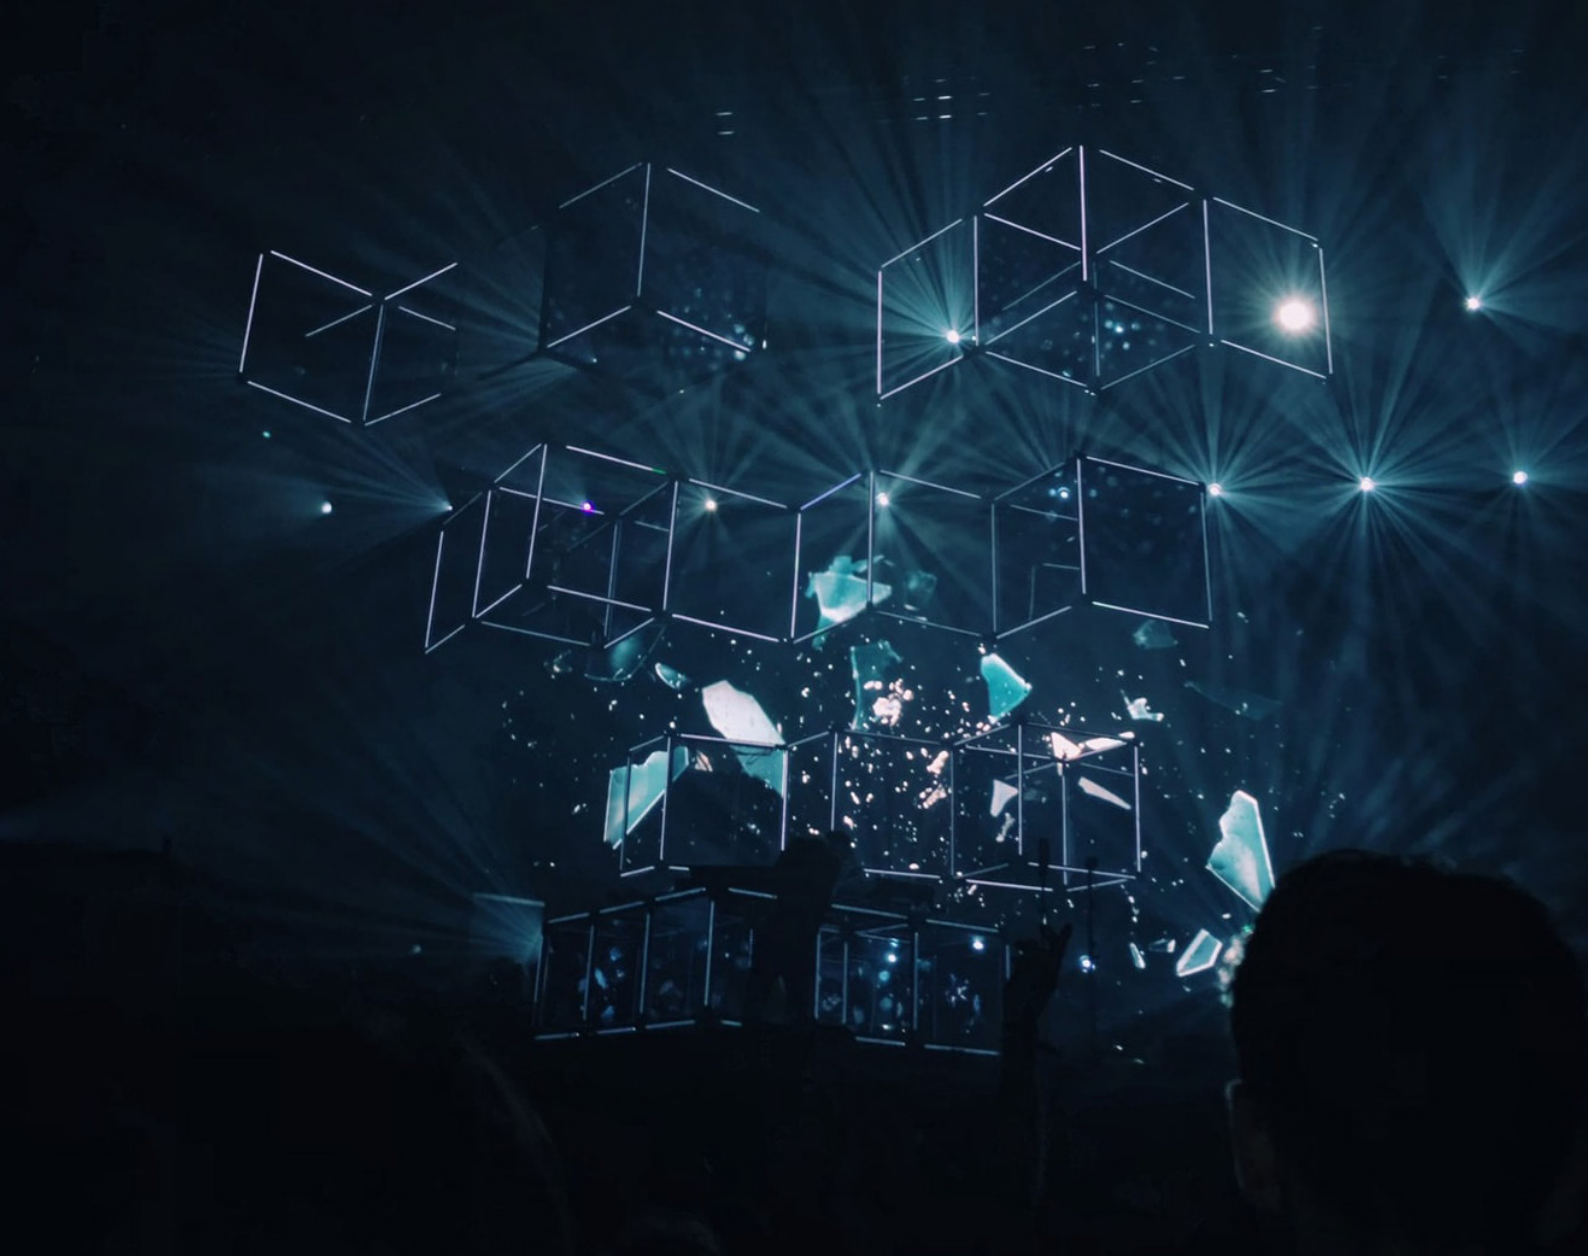

For further information please contact:

Gavin Whyte
Chief Artificial Intelligence Officer
Public Sector Network
E: Gavin@publicsectornetwork.co