



AI GETS REAL

AT SINGAPORE'S CHANGI AIRPORT

A learning journey.

By Steve Lee and Steven Miller

Ranked as the best airport for seven consecutive years, Singapore's Changi Airport is lauded the world over for the efficient, safe, pleasurable and seamless service it offers the millions of passengers that pass through its facilities annually.¹ Much of Changi Airport's success can be attributed to the organisation's customer-oriented business focus and deeply embedded culture of service excellence, combined with a host of advanced technologies operating invisibly in the background. The framework for this technology enablement is Changi Airport Group's (CAC's) SMART Airport Vision—an enterprise-wide approach to connective technologies that leverages sensors, data fusion, data analytics, and artificial intelligence (AI), orchestrating these systems and capabilities into feedback loops, and deploying them with user-centric design to enhancing customer experience and operating efficiency.

Over the past decade, the number of passengers and aircraft moving through Changi Airport has increased substantially. Passenger movements through the airport have gone up from 37 million in 2009, to 56 million in 2015, and 66 million in 2018. The airport handles about 7,400 flights every week, or about one flight every 80 seconds.

These large increases in passenger volume, and corresponding increases in aircraft movements (Figure 1), created an imperative for it to continue excelling in spite of increasing business and operational complexity.



FIGURE 1

Changi Airport’s use of connective technologies to move towards its SMART Airport Vision has been developing since 2017 (Figure 2). At CAG, we initiated a growing portfolio of pilot projects and follow-on production deployments to validate how best to implement the sensing and network infrastructure, data management infrastructure, data analytics, and AI and Machine Learning (ML) capacity. These

efforts help us in continuously building out and improving our enterprise-wide intelligence pipeline at scale (Figure 3), and have resulted in a first wave of deployment of AI and ML-enabled applications across various functions that can sense better, analyse better, predict better, and also interact with people better.²

Predicting flight arrival times more accurately

A large-scale data fusion effort has been a key enabler for our efforts to substantially improve our ability to predict flight arrival times.³ The heavy inflow of arriving aircraft into Changi Airport not only affects runway traffic, but impacts the entire airport arrival flow process—from the landing gate, to immigration, baggage claim, right up to the taxi queue. Combined, this creates the full passenger arrival experience.

In a joint collaboration with SITA Lab, a technology provider specialising in aviation and airport technology services, we built a prototype system for predicting flight arrival times for long-haul flights (over four hours’ flight time) with greater accuracy than was previously obtainable. Historical data and real-time data on flight arrivals were used as inputs into a suite of ML models to accurately predict the arrival time from the time the aircraft was airborne. It took us six months to come up with the first implementation of a workable model that was good enough to validate the approach. The project team then focused on progressively making refinements and enhancements to the prediction model as well as to the data sources. As a result, prediction accuracy has steadily improved—we are currently able to predict arrival times for long-haul flights with nearly 95 percent accuracy.

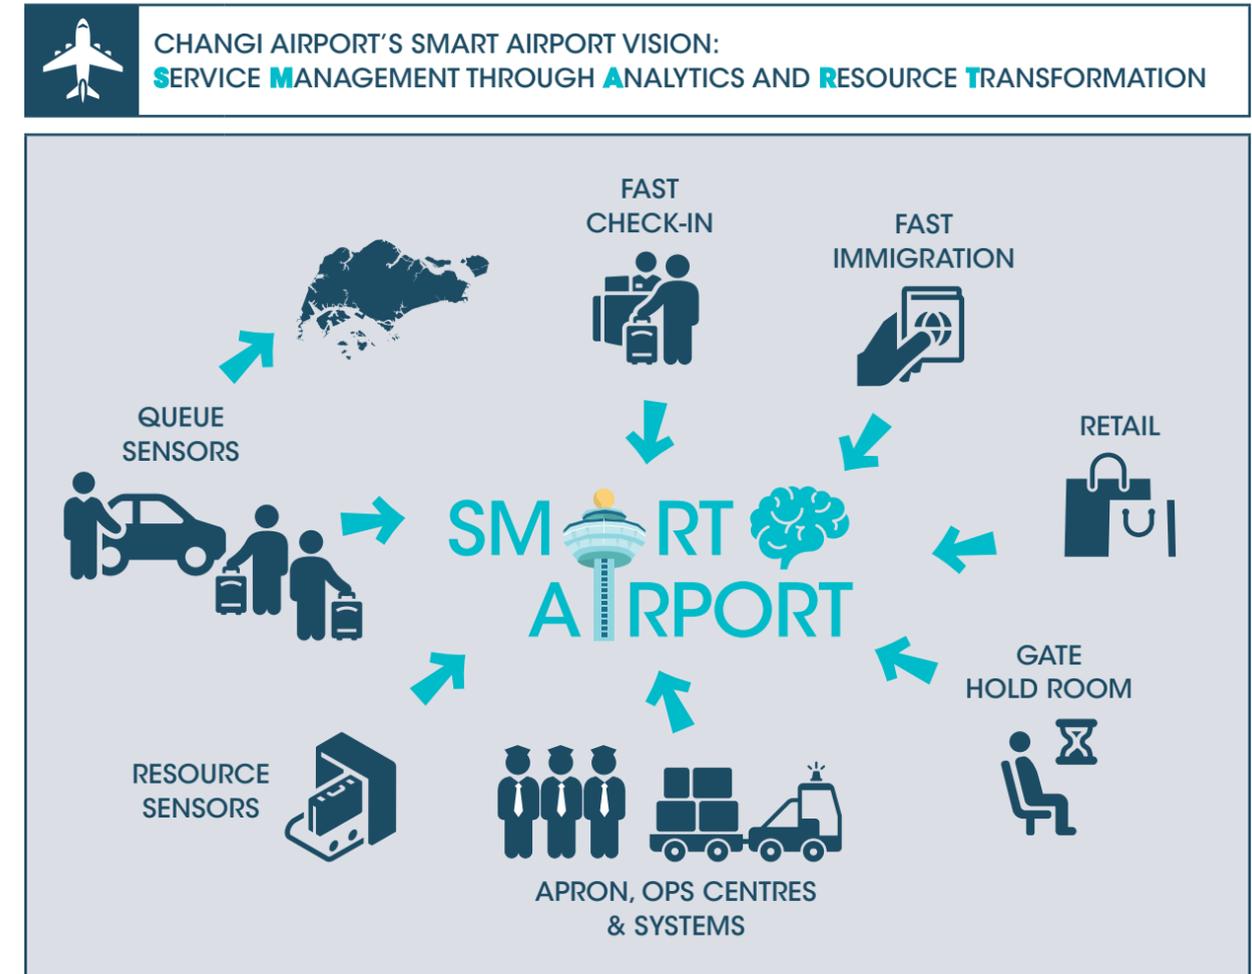


FIGURE 2

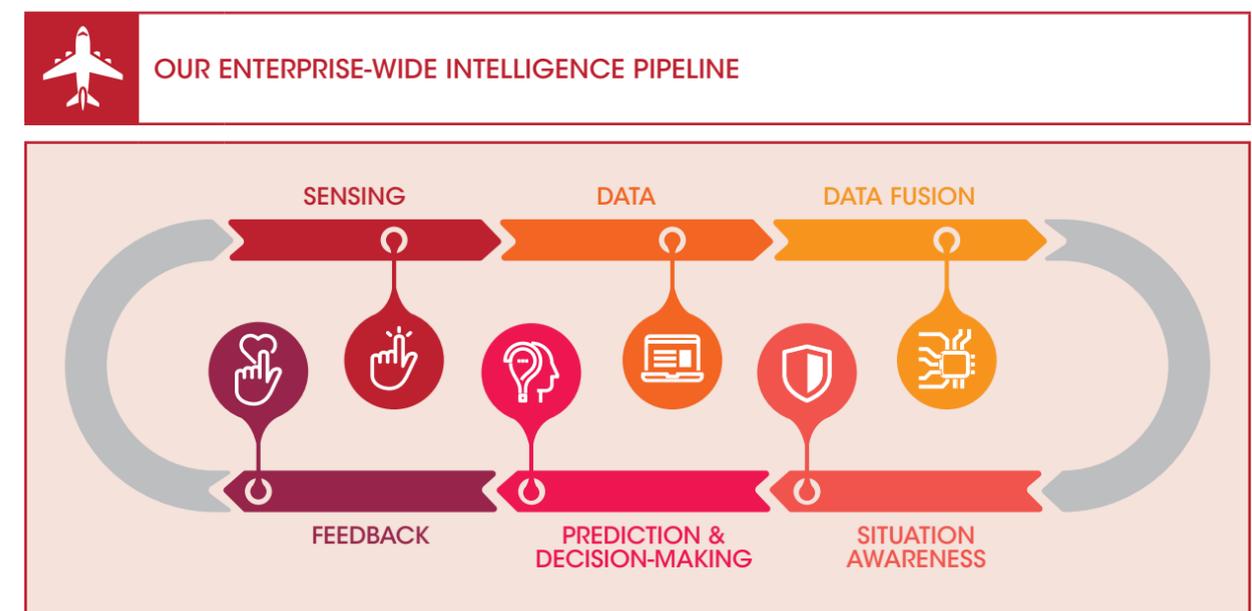


FIGURE 3

This more accurate prediction of flight arrival time is extremely valuable for operations supervisors who have to plan and deploy ground staff at the gate, as well as at various downstream processing points. It improves the airport's ability to support operations at the arrival gate, such as wheelchair assistance and other types of special needs. It also helps to improve service delivery planning and execution for immigration, baggage handling, and ground transport services. Overall, the mismatch between resource supply and demand throughout the various processing points of the passenger arrival journey is better managed. This helps to reduce the occurrence and severity of long passenger queues and to improve service staff utilisation and productivity as a result of less idle time.

Video analytics to improve airport safety and operational efficiency

Computer vision, video content analysis and video analytics are closely related terms used to describe technologies and applications that have been deployed for decades in a wide variety of commercial settings. Computer vision has been one of the main application areas for AI research, as well as for AI industry applications. Despite the availability of commercial solutions, our experience with video analytics illustrates that, while there are many general-purpose

AI techniques and solutions in the market, these off-the-shelf solutions need to be contextualised and customised in substantial ways to meet the requirements of our particular application setting and business needs.

DANGEROUS LEFT LUGGAGE

At CAG, we are developing video analytics capability to enable us to identify whether a piece of luggage left on the ground represents a dangerous situation. General-purpose automated tools for still-image analysis, as well as for video-image analysis, have improved dramatically in recent years, notably since the 2012 ImageNet competition that demonstrated the superiority of Deep Learning neural networks for tasks where very large data sets can be used to train the neural networks.⁴ The spectacular success of Deep Learning neural networks at the 2012 ImageNet competition, as well as in the ImageNet competitions that followed (2013-2017), are what launched the current 'Deep Learning revolution' across a wide-range of application settings that we are currently experiencing. Using Google's open source machine learning software, TensorFlow, and open source object libraries, we could easily train our software application to recognise luggage. Relatively easily and quickly, we could identify if there was a suitcase or other types of baggage left behind or left alone. What we really wanted to know, however,

While there are many general-purpose AI techniques and solutions in the market, these off-the-shelf solutions need to be contextualised and customised in substantial ways to meet our requirements.

was whether the left baggage represented a dangerous situation or not. It could be that a passenger sets aside a bag for a few moments, or even longer, to attend to something; or forgets to take with them one of the items they are carrying. The point is that, in most cases, bags left alone are not dangerous situations.

But trying to determine those very special cases, where a bag left alone has a high probability of being a dangerous situation, turned out to be very difficult to do, as there is hardly any training data available, and the many circumstances and scenarios involved made it nearly impossible to specify this type of detection using predefined rules. This example shows why one has to work with the technology, the vendors, and with internal staff to customise the context for specific circumstances.

FOREIGN OBJECT DEBRIS

We have also used visual analytics to inspect runways for 'foreign objects' in what the aviation industry refers to as Foreign Object Debris (FOD) detection. FOD are usually small items like keys, pens, batteries, coins, and even a workman's glove that inadvertently fall out of pockets. These seemingly trivial items can cause major damage to the aircraft and can even result in major accidents. Keeping runways free of FOD requires constant inspection and monitoring, as well as debris clearing, to ensure that the runway is perfectly safe for take-off or landing. Changi currently has two runways, each 4 km long, which require 24/7 inspection for FOD.

Prior to implementing our first visual analytics system for FOD detection in 2009, we asked ourselves if there is a way to use video analytics to support the staff involved in this extremely tedious and error-prone inspection work that is obviously not well-suited to humans. We went ahead and implemented a commercially available system specialised for this purpose, and it proved to be helpful. At the same time, it had a relatively high false alarm rate, as small puddles of water, pieces of paper, and other visual irregularities due to unusual lighting or shading, would trigger a FOD alert. False alarms are costly. For debris removal operations to

start, the runway needs to be cleared of all flights and the staff have to go on the runway and confirm whether or not there is FOD. This causes departure delays and disruptions, as flights taking off have to wait at the gate and arriving flights are put in a holding position.

Over time, we have been able to develop a good dataset that has helped us move to the next generation of video analytics capability for FOD detection. We collected large numbers of FOD image examples which showed true positives and true negatives, as well as false positives and false negatives. We then worked with a vendor who had strong general purpose capability with image and video analytics based on recent developments in ML (Deep Learning neural networks) for image recognition and classification. This newer approach, in combination with the training data set, reduced the false positive rate by 35 percent and even reduced the false negative rate by a small but practically useful amount.

These technological improvements substantially reduced the burden on operations personnel (who had to monitor the video analytics system) and unnecessary runway closures and their related disruptions. The key to realising this improvement was our historical data set. Without that, it would not have been possible to harness the new generation of higher performing neural network-based models for this application of video analytics. This is the power of data!

QUEUES AT SECURITY SCREENING

In three of the four terminals at Changi Airport, security screening is done at the boarding gate rather than at one central location. The passengers are delighted that they do not have to endure the arduous queues that are almost always associated with a centralised security screening area in most large, crowded airports. However, the decentralised approach to security screening, where screening machines and staff are located at each gate, poses a major operational challenge. This approach requires more security screening equipment to be purchased. It also creates a real-time need to know how many operations staff are required at each gate in order to minimise the build-up of queues of boarding passengers at that gate. This situation has to be managed over three terminals, each with approximately 60 gates.

We know the schedules of flight departures at each gate, which gates are going to be in use at which time, and if there are departure delays. But until recently, there has not been a practical, economical and reliable way to know the real-time situation of the queues of departing passengers at each gate. The arrival pattern of passengers at the gates



is probabilistic, and so is the amount of carry-on luggage per person. And while we may know the historical patterns of these variables, one never quite knows what will happen in real-time at any particular moment. We used to depend on our security operators to let us know when the queue starts forming. We also had operators manually looking at the CCTV screens to inform us about long queues. But there had to be a better way to monitor, and even predict, queue build-up at boarding gates using intelligent video analytics (IVA).

We created a large dataset using historical video image data from our existing cameras, and used this as a means of benchmarking the extent to which commercial vendors could meet our requirements. A competition was organised, where we screened the international vendor community, as well as local solution providers. Based on this, CAG invited 15 companies to participate in a competitive benchmarking exercise. This included international/multinational technology giants with video analytics products, highly specialised firms with deep expertise in this area, start-ups that were specialising in this area, and applied, translation-focused research institutes with these capabilities. To our surprise, none of the existing solutions from this very capable, pre-selected sample of 15 vendors could perform our test task to our required levels of performance. Their general-purpose capabilities were insufficient even with the training data we had given them, because the methods had not been tuned and customised to the context of our specific usage setting and needs. And that gap was large enough to make the off-the-shelf solution impractical to use in our everyday production setting.

Why was it so difficult to count people (as video analytics systems are usually very good at automatically counting people) and form an estimate of the crowd at the gates? Changi's departure gates have glass walls designed to create a spacious feeling for passengers. However, this created havoc for automated video analytics systems that count people, as there was no way of determining from the image if a person was inside or outside of the glass wall! The queues could also form in less distinct ways and sometimes mix with regular passenger traffic walking past that particular gate. Was the automated video analytics system counting the actual number of people waiting in the queue to go through security screening? Or was it also counting people that have already passed screening and are inside the boarding gate? Or people just walking by? This is just one

vivid example of why something that seems easy to do with today's AI technology could turn out to be not so easy in a specific setting.

We eventually proceeded with one video analytics vendor, and worked with them to contextualise and customise their solution for our required level of performance. In addition, we pursued a data fusion strategy to provide more context and hence more intelligence to our overall analysis. For example, we identified the type of aircraft that would be departing, as the capacity of the aircraft gives us an upper limit on the number of people who would be arriving at the gate. We found ways to use video analytics to create profiles of flights, to provide statistical information on the patterns of passenger arrivals at the gate by flight and by time of day/month/year. All of this combined enabled us to improve our ability to estimate, in real time, how many passengers have boarded, how many more will need to board, what the projection for the queue will be, as well as the manpower requirements at the gate to handle the predicted number of boarding passengers. This AI-enabled video analytics effort is enabling us to reduce queues at those gates with decentralised security screening, and to increase the productivity and deployment efficiency of the operations personnel who operate the screening checkpoints.

Customer service chatbots

Passenger volume levels at Changi are expected to keep increasing, and when our new Terminal 5 opens in 2030, passenger volume for all five terminals is expected to cross 130 million people annually. Unsurprisingly, the airport receives huge volumes of customer service calls each year, mostly about basic information like flight times, gates, transit times between terminals, baggage allowances, immigration procedures, lost-and-found requests, and restaurants and shopping. While peak times for asking these types of questions occur during the day and up until the evening, there are still quite a few who contact the call centre between midnight and 5 am. People like to ask questions, even though much of this information is available on the company website, and they want the assurance of a well-informed response.

Given that so many of these questions are repeats of the same inquiries, and given the heavy volumes, the employee turnover rates at Changi's call centres indicate that this is not a desirable job—at least not in its current form. Chatbots are the obvious strategy for improving this situation, but we all know the frustrations of using chatbots



that don't really understand the intent or specifics of what is being asked, and don't know when and how to transfer the call to a human if the interaction is not working out well.

CAG started live deployment of customer-facing chatbots at the beginning of this year. Prior to launching our own, we studied how other airports, airlines, and a variety of other industries used chatbots. We studied the types of questions that come through the airport call centre. We also studied the various sequence and bundling of questions, to see if there were ways to anticipate the customer's next question and provide them with what they might want even before they ask. We went through numerous 'design thinking' sessions and considered many alternatives for how to create easy-to-use interfaces and navigation sequences, and how we might combine text responses with other types of structured information support. We did experiments to see if people preferred to receive streamlined responses (with carefully curated and minimal amounts of content to answer their question) or more content-rich responses (with more supporting information). We worked with language technology specialists to choose which language technology platform to use given the types

of questions our customers ask, the types of dialogues they engage in, and other technical needs such as requirements for enterprise security and data protection.

In other words, the chatbot project involves a whole lot more than just chatbots. It is an entire user experience design effort, where the chatbot aspect of answering text questions is just one part of an AI-enabled, increasingly more context-aware (and intelligent) approach to understanding the customer's intent, and providing them with the support information they need.

Chatbots will enable us to reduce the pressures of volume and tedious repetition on our human call centre operators. Also, they will allow us to digitise more easily what people are asking, even with very large quantities of questions. This helps us to learn on a continuous basis, and adapt our services accordingly, e.g. when a specific inquiry or comment comes up frequently, we recognise that as a problem and fix it accordingly. Similarly, customer questions also give us insights into market demand for new service offerings.

We have learnt that the development of chatbots is an iterative process. It continuously improves through ongoing

usage in testing and production, through ongoing monitoring of responses for adequacy, errors and failures, ongoing refinement of the language technology models, and parallel refinements to the related models for guiding navigation, presenting supporting information, and user experience. We also put a lot of attention on how to achieve a smooth handover from the chatbot to the human operator when the chatbot is unable to handle a query. Over time, we hope to rely as much as we can on the chatbot while keeping service standards high. We will also continue to refine the way that the chatbots and their human supervisors work together to provide a better experience than could otherwise be provided with only chatbots or only people.

Lessons learned and managerial insights

Technology in itself is an enabler. Even with today's more advanced and rapidly improving AI capabilities, it is still true that technology is necessary but not wholly sufficient for delivering world-class levels of service excellence across all aspects of our customer interactions and operations. And we do not see this situation changing any time soon.

The highly iterative nature of all of our AI projects is a consistent theme—with each of these project iterations, our project teams learn and grow. They develop and redevelop the analytics and prediction engines, and refine and expand the data sets,

and learn and grow again. We keep cycling through the OODA loop: observe, orient, decide, and act. In the process, we are not just growing the technology, we are growing our operating practices and processes, as well as growing the learning capacity of our entire organisation.⁵

Our conceptual roadmap for realising the SMART Airport Vision provides a clear, simple visual representation to convey the big picture and strategic purpose for growing our portfolio of AI projects (Figure 4). The roadmap makes it easier for our cross-functional teams to communicate about which type of key outcome they are targeting for capability enhancement, and which level of key enabler they are contributing to and/or making use of. It also helps

our cross-functional teams clarify where and how the project fits into the overall management cycle.

From what we see on the horizon, we believe a lot of the AI technologies and vendor solutions that are in use today will be commodified and simplified. Organisations that understand this will learn how to move through successive iteration cycles faster, and be in a better position to leverage the OODA loop. For Changi Airport, the bulk of our AI-related investments in the near- and mid-term will go into contextualising the vendor-provided (or open sourced) AI models such that the analytics, prediction, optimisation and decision-making methods derived from the data work for our specific context and needs.

Based on our project experiences thus far, we summarise six managerial insights emerging from our journey with learnings on how to make AI real at Changi Airport. We believe these observations are applicable to a wide range of medium- and large-sized organisations that deliver services or products in the physical world—as in, those not 'born' as fully digital companies—and currently in the midst of various digital transformation efforts.⁶

1. The importance of a conceptual roadmap for guiding our growing portfolio of AI projects.
2. Rethinking our strategies for outsourcing, insourcing and co-sourcing our technology product and service providers to move with greater agility and speed.
3. The criticality and value of our internal knowledge of our specific business domain and business needs.
4. The criticality and value of iteration for addressing key management issues as well as for improving system and process performance.
5. The availability of new data sources, including open data sources, is a game changer.
6. The power of data fusion, through understanding what data is needed, what is available, and how to put the available data together for meaningful use.

At Changi Airport Group, connective and intelligent technologies are here to stay. But the question for us is: How can these new capabilities change the way we think, the way we work, and the way we serve our customers?

Once we understand how to get on the pathway of making these changes in mindset, practices, and engagement, we humans can do remarkable things working symbiotically with AI-enabled technologies and machines.⁷

Steve Lee

is the CIO and Group SVP for Technology at Changi Airport Group

Steven Miller

is Vice Provost (Research) and Professor of Information Systems (Practice) at Singapore Management University. He serves as an advisor to Changi Airport Group

The authors gratefully acknowledge the substantial efforts of the many CAG employees and partnering vendors who worked to realise these AI initiatives.

References

- 1 Skytrax, World Airport Awards, "The World's Best Airports of 2019".
- 2 Thomas H. Davenport, "The AI Advantage: How to Put the Artificial Intelligence Revolution to Work", MIT Press, 2018, provides additional information on the recent wave of deployment of AI (cognitive automation) applications in large corporate enterprises to improve business processes, analytics capabilities, and support tools for customer and staff engagement.
- 3 Previously appeared in Steven M. Miller, "AI: Augmentation More So than Automation", Special Supplement, Asian Management Insights, Vol. 5, Issue 1 (May 2018), Singapore Management University.
- 4 Dave Gershgorin, "The data that transformed AI research—and possibly the world", Quartz, July 26, 2017. A more technical talk that summarises the history and progress of ImageNet given by the creator of the effort, is Li Fei-Fei and Jia Deng, "ImageNet: Where are we and where are we going", 2017. Both these documents highlight the dramatic increase in automated image classification and recognition abilities that occurred as a result of training multi-layered (Deep Learning) neural networks on the very large ImageNet dataset, and applying those trained neural networks to the Image classification task. The spectacular success of Deep Learning neural networks at the 2012 ImageNet competition, as well as in the ImageNet competitions that followed (2013-2017) are what launched the current "Deep Learning" revolution across a wide-range of application settings that we are currently experiencing.
- 5 Nancy M. Dixon, "The Organizational Learning Cycle: How We Can Learn Collectively, 2nd edition", 2017, Routledge, an imprint of Taylor & Francis Group; and Ryan Smerek, "Organizational Learning & Performance: The Science and Practice of Building a Learning Culture", 2018, Oxford University Press.
- 6 More details available in Steve Lee and Steven M. Miller, "Lessons Learned and Managerial Insights as AI Gets Real at Changi", Special Supplement, Asian Management Insights, Vol 6, No. 1 (May 2019), Singapore Management University.
- 7 For more insight into the issues and possibilities associated with symbiotically combining human intelligence with machine intelligence, see Paul R. Daugherty and H. James Wilson, "Human + Machine: Reimagining Work in the Age of AI", 2018, Harvard Business School Press; and Steven M. Miller, "AI: Augmentation More So Than Automation", Special Supplement, Asian Management Insights, Vol. 5, Issue 1 (May 2018), Singapore Management University.

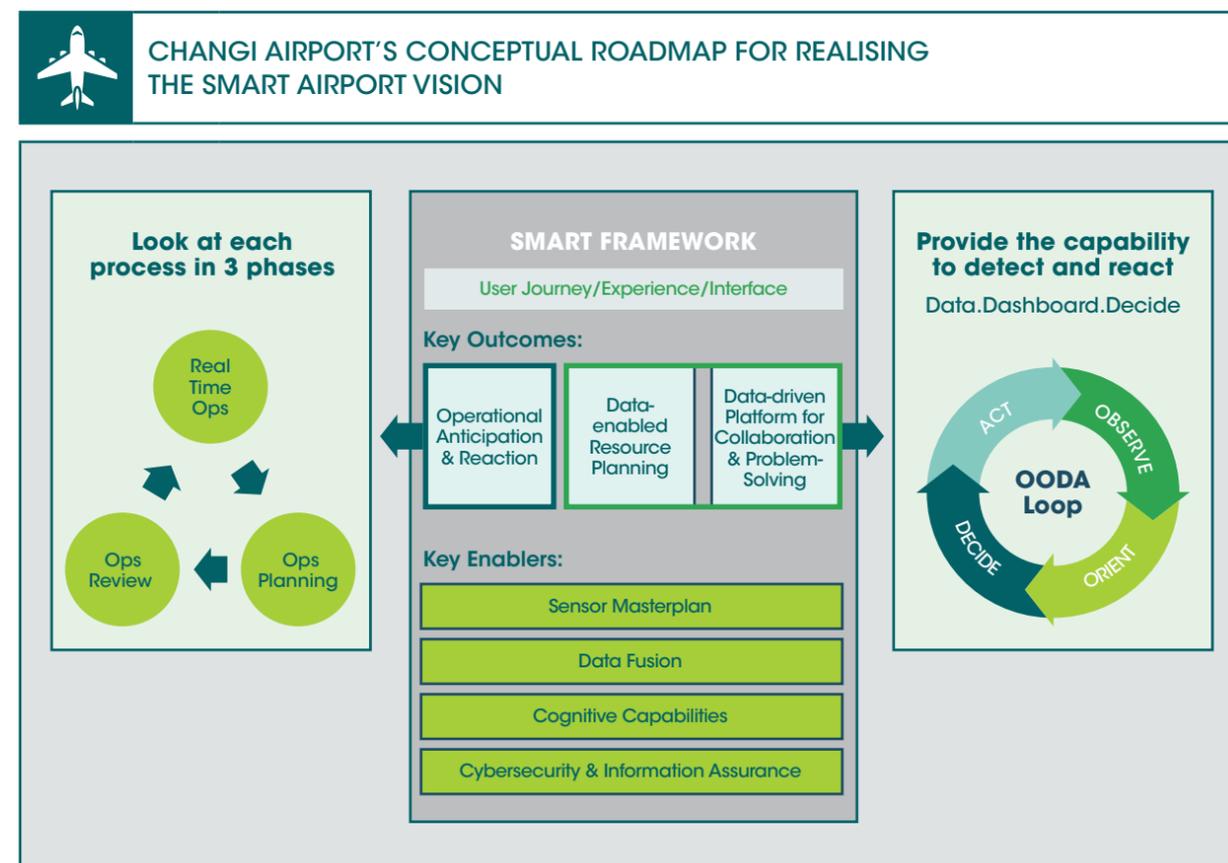


FIGURE 4