SAP Institute for Digital Government

SAP Thought Leadership Paper | PUBLIC
Public Sector

Redesigning Work with Artificial Intelligence

Accelerating the Adoption of Artificial Intelligence in Government Organizations

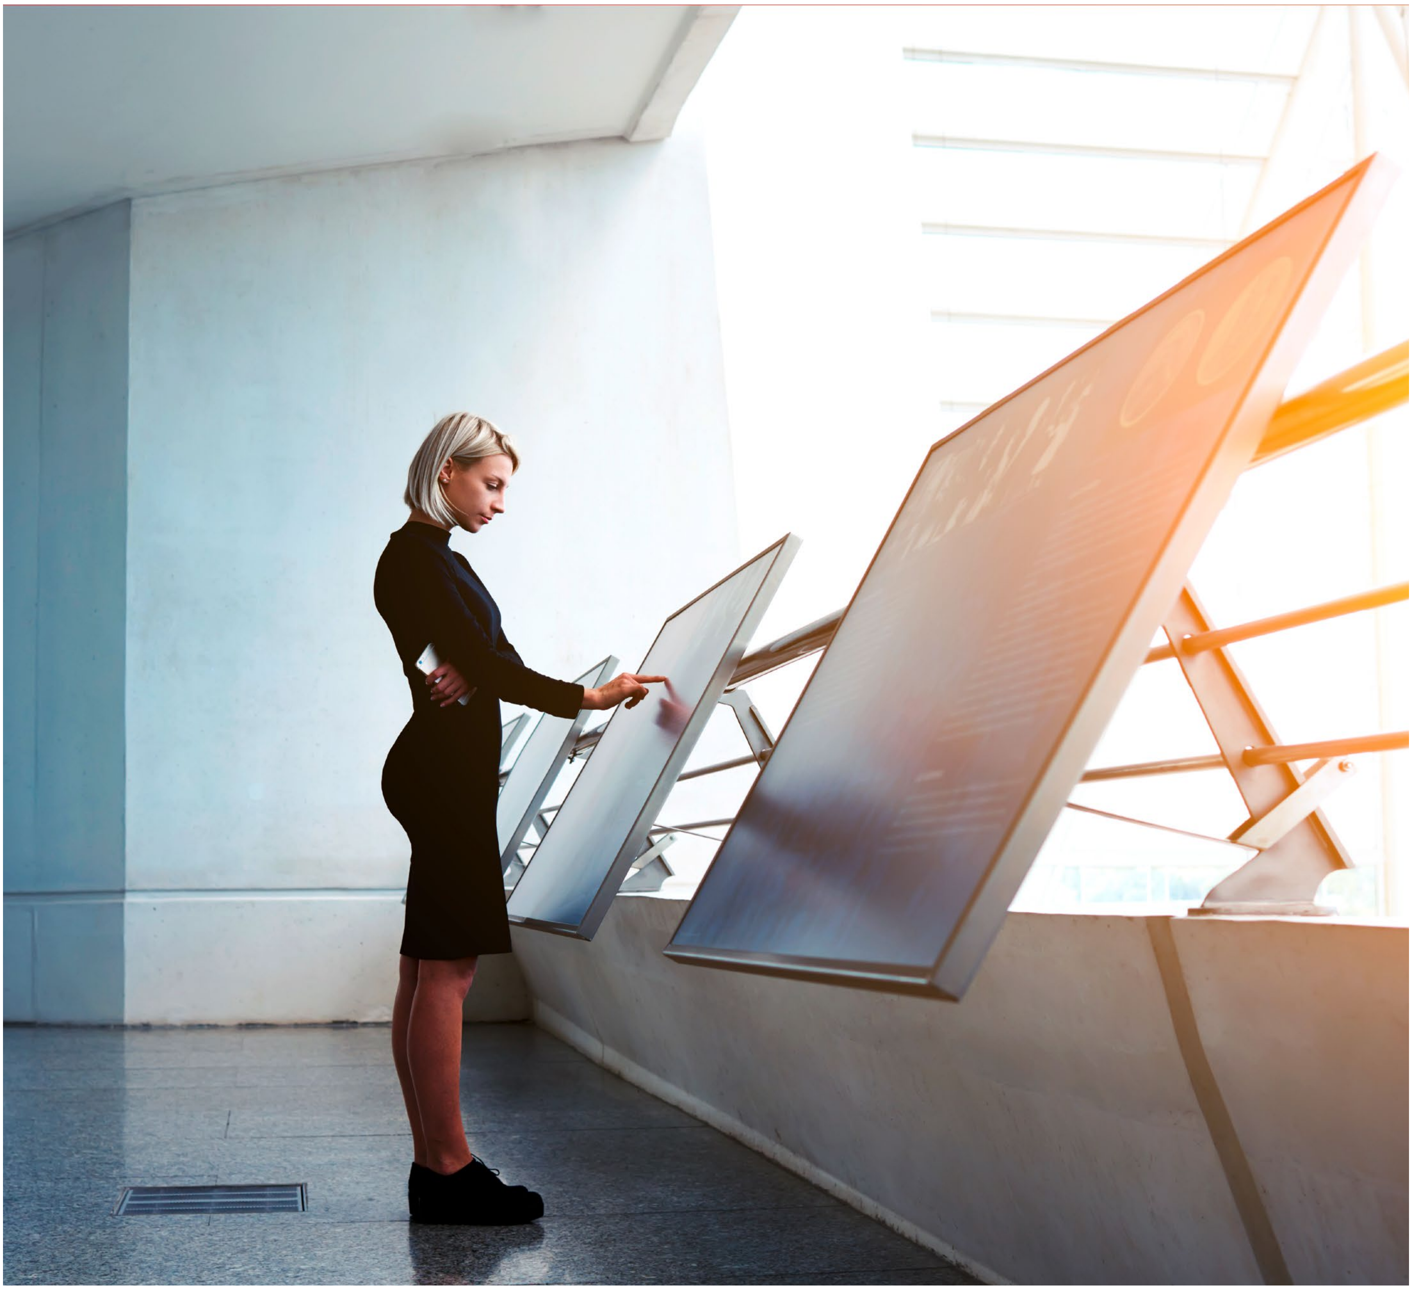

The University of Queensland Australia

Table of Contents

Authors
Doctor Reihaneh Bidar
Doctor Ida Someh
Doctor Tapani Rinta-Kahila
Doctor Katie Williams
Professor Marta Indulska
Ian Ryan

As government bodies become more agile, results-oriented organizations, they are turning to artificial intelligence (AI) to change traditional ways of working. This technology can revolutionize the efficiency and flexibility of public services and deliver outcomes that are not currently achievable for them.

Yet, the risks that AI poses to stakeholders, together with the sector's rules, policies, and legislation, constrain the ability to push AI-driven change.

So how can AI be successfully leveraged with minimal risks? What are the key design choices for developing effective and responsible processes in the public sector that maximise value for stakeholders?

The Move to Leveraging AI by Governments Around the World

Governments around the world have started to implement AI programs with a view to improving their capabilities and optimizing service delivery for citizens. This is a considerable change, with advancements in AI and digital technologies are forecast to contribute $315 billion to the Australian economy by 2028[1]. The Australian Public Service Commission reports that 18% of tasks and up to 40% of time used to perform some roles in public sector could be automated by 2030[2]. Such predictions are driving public sector agencies to explore the next horizon for work: its transformation through the use of AI. The number of pilot AI projects in the public sector has significantly increased. While there has been some success with deploying AI (for example, in medical diagnosis, assessing tax claims, and optimising transport), relatively few of the pilot projects progress to full implementation[3], let alone enable sustained value creation. A key reason for this is that best practice for effective design and implementation of AI-based work processes has only started to emerge.

1 Hajkowicz, S. A., Karimi. S., Wark, T., Chen, C., Evans, M., Rens, N., Dawson, D., Charlton, A., Brennan, T., Moffatt, C., Srikumar,S., Tong, K. J. (2019) Artificial intelligence: Solving problems, growing the economy and improving our quality of life. CSIRO Data61, Australia

2 Commonwealth of Australia. (2019). Our Public Service, Our Future. Independent Review of the Australian Public Service.

3 Benbya, Hind; Davenport, Thomas H.; and Pachidi, Stella (2020) "Special Issue Editorial. Artificial Intelligence in Organizations: Current State and Future Opportunities," MIS Quarterly Executive: Vol. 19 : Iss. 4 , Article 4. Available at: https://aisel.aisnet.org/misqe/vol19/iss4/4

AI differs fundamentally from previous, simpler technologies, such as rule-based automation and traditional reporting and analytics techniques. The complex characteristics of AI increase the risks to stakeholders when it is incorporated into work processes. These characteristics include:

- AI's outputs are probabilistic, so applying them to real-world cases as if they were definitive can create negative and unintended consequences
- AI operating mindlessly by only applying learning from the training data made available to it, can overlook important aspects of citizens' current circumstances when making decisions
- AI-based models 'drift' and the outputs can become biased over time, with relevance and accuracy suffering if the systems are not appropriately maintained

Because of these unique challenges, integrating AI into work processes introduces risks to citizens, and so it is crucial for governments to manage these risks effectively. In the private sector, citizens can often vote with their feet, whereas opting out of public services is often not feasible. These risks are inherent to AI and should be managed effectively through appropriate design choices that minimize any potential harm to citizens.

This paper offers the first practical steps to guide the responsible design of AI-based work processes in the public sector, by offering a risk-informed approach to classify tasks and design work-related elements.

Design Elements for AI-Based Work Processes

Successful deployment of AI requires attention to four process design elements. These are people (workers, their managers, and IT developers) who use technology (here, an AI system) to handle a range of tasks (various decisions on citizen-related services) within a specific structure (for example, hierarchies, lines of responsibility, and incentive systems). These four design elements are interrelated and can be arranged in various ways to ensure responsible and effective services.

The revolutionary nature of AI technology entails fundamental change to how work processes can be designed and executed. Work tasks, technology management capacity, people's skills and roles, and the structure governing the work all must change in the pursuit of effective design that minimizes risks.

So, how should public sector organizations approach the redesign of their work processes for the AI age in practice? We suggest the journey begins by understanding the types of tasks influenced by AI and the risks inherent in those tasks.

Process design elements	Definition
Task	Activity or set of activities undertaken to deliver an outcome (for example, operations or services to citizens)
Technology	Use of technology to support and transform everyday operations through automation or augmentation of work processes
People	Human actors in the organization who work jointly with the AI system to carry out tasks, such as managers, domain experts, and IT developers
Control structure	Organizational arrangements (for example, governance, risk-control, decision-making) to manage the dynamics that AI brings to the environment

AI calls for fundamental changes to the core process design elements of **tasks, technologies, people,** and **structure.** However, a responsible work re-design approach would proactively consider risks to stakeholders.

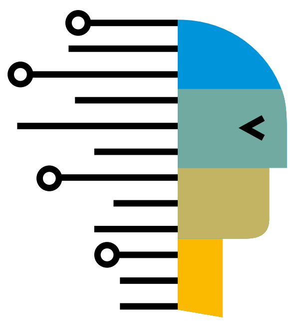

A Risk-Informed Task Classification for AI

Recent applications of AI in the public sector have brought significant unintended consequences for citizens and created widespread mistrust in the technology. For example, the Dutch government's use of AI caused controversy by systematically judging certain groups of recipients guilty of welfare fraud and led to the government resigning.
In the US, the use of predictive policing systems directed officers to areas more heavily populated by racial minorities, resulting in those minorities likely being unfairly targeted by law enforcement.

This set of examples illustrates why risk needs to be considered up-front in the design of AI-driven processes and carefully managed through various design choices. The risks that AI poses to citizens are context-specific and range from a loss of privacy and exercise of governmental control over citizens, to biases resulting in unfair treatment of the citizens, to significant adverse financial or health events. Identifying the risks early on can help to determine the right process design choices including defining appropriate roles for humans, effective governance mechanisms, and so on.

To enable a risk-informed approach that can guide the integration of AI technology, we classify tasks for AI along two dimensions: the level of risk and degree of specialization for the task:
The level of risk is the potential harm to citizens associated with the use of AI in public services.

- **Limited risks:** tasks with a low degree of potential harm to citizens

Example: government services related to establishing new companies or collecting property tax could be characterized as low risk tasks, because there is low likelihood of immediate harm or threat to citizens. AI operating incorrectly would mainly cause harm in terms of financial or other resources, but any damage would typically be somewhat revocable (for example, AI misidentifying a company as potentially fraudulent).

- **Significant risks:** tasks with extensive potential for bringing harm to citizens

Example: government services around healthcare, policing, and social services, for instance, tend to incorporate work tasks involving significant risk. This is because these services often have direct implications on citizens' wellbeing and safety, and AI operating incorrectly can potentially exert significant and potentially irrevocable damage (for example, an

AI diagnostic tool recommending an incorrect treatment for a patient).

The degree of specialization is the depth of knowledge required for performing tasks.

- **Low degree of specialization:** tasks that require only common knowledge and skills – they involve little or no complex decision-making
- **High degree of specialization:** tasks that require sophisticated knowledge and skills that can only be gained through special training – they involve a medium to high degree of complex decision-making

Example: the work of a government tax accountant who must understand the complexities of the tax law and use their knowledge and experience to interpret cases for which no clear guidelines exist.

When juxtaposing the levels of risk against the degree of specialization, four different classes of AI tasks emerge (see Figure 1 for definitions). We suggest that each class of AI tasks imposes its own requirements for optimal process design and minimization of risks. The following case studies help us to illustrate the concrete differences in these sets of design choices.

	Low degree of specialization	**High degree of specialization**
Significant risk	**High-reliability tasks** Tasks that require no or low levels of specialization but pose a high risk of potential harm to citizens and demand significant risk control and oversight	**Wicked tasks** Tasks that require high levels of specialization and pose a medium to high risk of potential harm for citizens, demanding significant risk control and oversight
Limited risk	**Mundane tasks** Tasks that require no or low levels of specialization and represent a low degree of potential harm for citizens, not demanding extensive risk control or oversight	**Expert tasks** Tasks that require medium to high levels of specialization but represent a low degree of potential harm to citizens, not demanding extensive risk control and oversight

Figure 1: The Four Classes of AI tasks

Redesigning Work Processes for AI

Here are four case studies to illustrate the decisions involved in designing AI-based work processes. Each of the cases represents one specific task type – mundane, high-reliability, expert, and wicked – and the associated design choices.

WORK REDESIGN FOR MUNDANE TASKS

Case study 1: Rego AI

Context: a European business authority tasked with boosting business and investment opportunities while ensuring key regulatory obligations were managed effectively. A key aspect of expanding the opportunities nationwide was to create an efficient, hassle-free process to handle the growing number of new company registrations. In 2017, the agency established a machine learning (ML) lab to develop prototype applications that handle real use cases across the agency, including new company registrations.

The work's original design: for registering new companies, human workers reviewed company-establishment documents sent by citizens, verified their validity, and examined them for any red flags before further processing.

The problem: for various reasons, many of the incoming documents had been left unsigned by the people who submitted them. Because these could not be processed and had to be returned to citizens, significant manual labour and unnecessary inefficiencies resulted.

Redesigned work for Rego AI: to improve the efficiency of the document review process, the agency divided the original task into smaller parts, to enable distribution of the work between AI and humans. Human workers were tasked with verifying the document information, which required specialist knowledge (such as judging whether the prospective company is legitimate), whereas the mundane work of checking whether a given document had been signed or not was assigned to AI.

4 Asatiani, Aleksandre; Malo, Pekka; Nagbøl, Per Rådberg; Penttinen, Esko; Rinta-Kahila, Tapani; and Salovaara, Antti (2020) "Challenges of Explaining the Behavior of Black-Box AI Systems," MIS Quarterly Executive: Vol. 19: Iss. 4, Article 7..

DESIGN CHOICES FOR MUNDANE TASKS (REGO AI)

Let us look at how AI enabled redesigning of the process and addressing of the agency's problem including the changes made across the four key work design elements.

The task for AI

AI was given the task of detecting whether a document was signed or not. Signature detection was classified as an example of a mundane task because:

- Identifying the presence or absence of a signature is simple and requires only a low level of specialization (i.e., one need only apply rudimentary knowledge and skill)
- It would pose only limited risk with minimal possible harm to citizens from errors (i.e., from false positives or false negatives such as the system claiming that the document is not signed when it is signed

Technology

A deep-learning AI model was developed to detect signatures in scanned images of the documents, in conjunction with a document filter that sends the documents for processing (if a signature was detected) or back to the submitter (if one was not detected). The model assesses the probability of the document having been signed and returns three figures: the likelihood of it being physically signed, the presence of a digital signature, or a signature being missing. Because of the low-risk nature of the task, high accuracy was prioritized over transparency, and the model was allowed to perform automatic actions. The AI model was a complex, black-boxed system whose actions were not easily explainable to humans.
This was not problematic, though, because human workers did not need explanations for why an image was classified as signed or not.

People

The burden on case workers was alleviated through handing off a 'mindless' part of the task of processing companies' establishment documents (checking for signatures) to AI. One of the interviewees described this as *"trying to lessen the manual workload and reserving the human decision-making for the more creative or advanced tasks"*. This freed them to focus their energies on the mindful portion of the task – evaluating the documents' content.

Control structure

Boundaries were established for the AI's operation to address the possibility of it misjudging a signature to a document: Human case workers would still check the AI's outputs when cases of potential error were escalated by citizens. They could do this easily by visually inspecting any documents for which queries were raised.

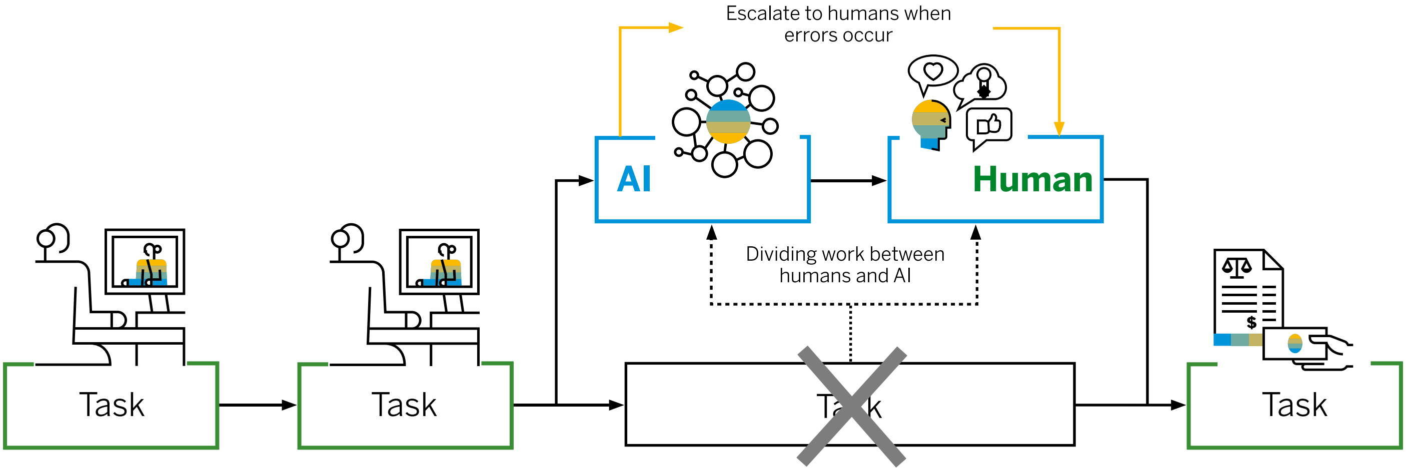

Figure 2: Design Choices for Mundane Tasks

People
The burden on case workers was alleviated through handing off a 'mindless' part of the task of processing companies' establishment documents (checking for signatures) to AI. One of the interviewees described this as *"trying to lessen the manual workload and reserving the human decision-making for the more creative or advanced tasks"*. This freed them to focus their energies on the mindful portion of the task – evaluating the documents' content.

Control structure
Boundaries were established for the AI's operation to address the possibility of it misjudging a signature to a document: Human case workers would still check the AI's outputs when cases of potential error were escalated by citizens. They could do this easily by visually inspecting any documents for which queries were raised.

- **The task:** splitting the original work process into parts, with the mindless tasks that do not require much specialization separated out for delegation to AI
- **Technology:** AI performing autonomous actions, prioritizing performance over transparency
- **People:** limited human intervention, with humans reassigned to new roles where mindful and specialized work is required
- **Control structure:** escalation to humans as a fallback mechanism if the AI's decision is challenged

Work Redesign for High-Reliability Tasks

Case study 2: Border AI

Context: Australian border-protection agency responsible for accurate identity verification and for clearance processes for passengers entering and leaving the country at an international airport. In 2021, the agency deployed the SmartGate system to perform identity verification via biometric matching across selected international airports, which afforded optimizing Australian citizens' movements across the nation's borders.

The work's original design: border-protection staff manually checked the passenger's travel documents and compared the person's face with the passport photo. If the documents were in order and the two matched, the passenger was admitted through border control.

The problem: manual verification of passengers' travel documents and identity is subject to human error and is resource-intensive. It created extensive queues, which imposed a burden on both border workers and passengers.

Redesigned work for Border AI: the agency leveraged technology to automate the passenger biometrics matching process, for greater precision, a better passenger experience, and reduced burden on workers. The process of identity verification was split into two closely connected subtasks, each of which was then assigned to suitable technology: verifying travel documents (i.e., comparing ePassport information against the customs database), and matching the passenger's face with the passport image. A digital process was implemented for the former, and the matching was delegated to AI.

DESIGN CHOICES FOR HIGH-RELIABILITY TASKS (BORDER AI)

Let us look at how AI functioned in the redesign of the work process to address this problem. Again, the changes can be considered in terms of the four process design elements.

The task for AI

AI was put in charge of performing the matching of a passenger's face with the person's passport image, which complemented the passenger verification process using ePassport data. Facial matching was classified as an example of a high-reliability task because:

- Comparing faces with passport photos involves a low degree of specialization, in that this task can be performed by means of common knowledge and skills
- The possibility of allowing a person with fabricated travel documents to enter/leave the country poses a significant risk with the possibility of harm to citizens, since it could jeopardize passenger safety and airport security (consider terrorism, for example)

Technology

SmartGate employed a camera and a deep-learning model designed for comparing a passenger's face with the passport image presented, in conjunction with an ePassport reader that compares passport information with details in the customs database. If the passenger is cleared, automatic gates open and the person proceeds through. Hence, the entire process is automated with self-service technology that applies several components. In the model's implementation, high performance was prioritized over explainability, to minimize congestion at the airport.

People

The AI system significantly decreased the repetitive routine work of the border officers, alleviating workload. However, the automation did not entirely eliminate this human role. Because of the high risk associated with border security and the fact that many passengers do not have ePassports, a parallel manual process performed by human officers was maintained. Also, some personnel were reskilled to assist passengers with the use of the SmartGate system (for example, instructing them in how to stand in front of the camera).

Control structure

SmartGate needed appropriate intervention points to facilitate secure passenger movements. For example, the system is designed such that any possible failure (such as photo mismatch) causes the passenger to be referred for manual processing to one of the officers. Human workers also served as active supervisors of the AI in the field.

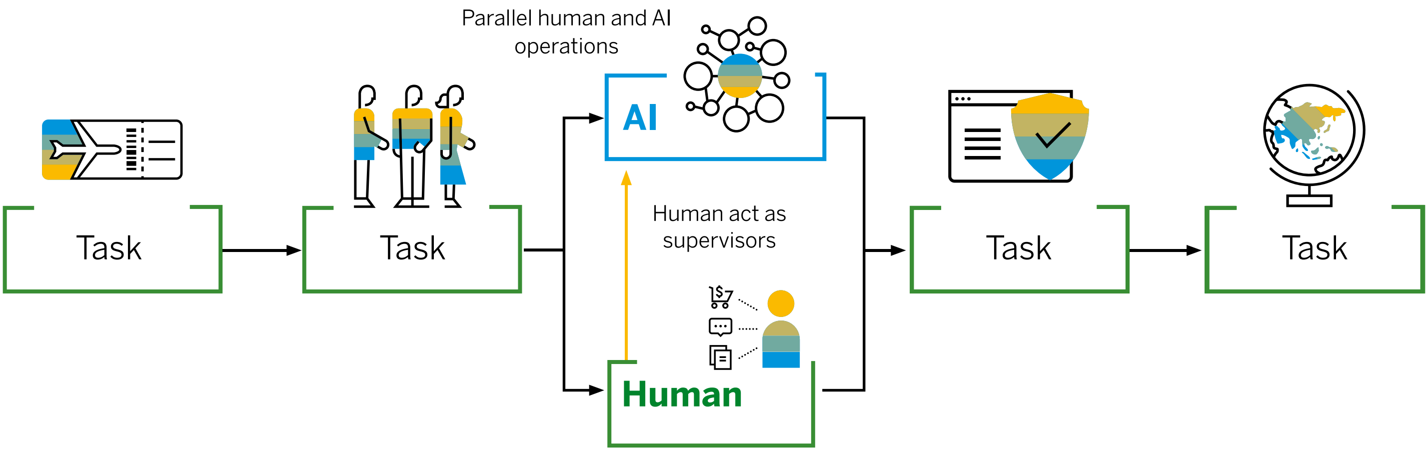

Figure 3: Design Choices for High-Reliability Tasks

- **The task:** dividing the original work into parts and creating completely digital operations to increase throughput and handing over low-specialization tasks to AI
- **Technology:** supervised automation with model performance prioritized over transparency
- **People:** parallel processes run by humans, with humans reskilled to support AI operations and some employee roles redefined
- **Control structure:** deploying active human supervisors to minimize risks

Work Redesign for Expert Tasks

Case study 3: Tax AI

Context: an Australian state revenue office responsible for collecting land tax and enforcing tax compliance. The agency deployed AI to enable efficient and effective collection of land tax, integrated with programs for managing related debt. The goal was a better tax-collection process and special support for citizens who were likely to become debtors.

The work's original design: human experts manually assessed citizens' tax profiles and past payment behaviour to determine who was at high risk of becoming a debtor. Experts would then reach out to those citizens and offer potential avenues for support.

The problem: the original work design relied on human expert applying 'rules of thumb' and experience to identify citizens at risk. The tasks were not performed consistently from officer to officer because workers with less experience at the agency could not accurately identify individuals at risk and their behaviour patterns. Also, even seasoned professionals were subject to biases. The revenue office wanted to decrease the number of citizens becoming tax debtors.

Redesigned work for Tax AI: to improve the accuracy of predicting who is likely to become a debtor, the agency delegated the information processing to AI designed to make predictions and offer a human agent corresponding behavior insights and recommendations. After reviewing the recommendations, the human debtmanagement officer decided whether the citizen seems to possess capacity to pay or, instead, requires additional financial support.

DESIGN CHOICES FOR EXPERT TASKS (TAX AI)

AI enabled redesigning this work process to address the problem. Again, let us look at the changes in terms of the key work design elements.

The task for AI

The existing task was augmented with an AI model charged with the task of predicting which citizens are at risk of insolvency. The goal was to prevent taxpayers from falling into debt. Predicting potential debtors was classified as an expert task because it:

- Requires a high degree of specialization, with officers needing detailed-level understanding of tax law and customer payment behavior, which demands special training
- Poses limited risk with a low possibility of harm to citizens – the task was just to identify citizens 'at risk' of land tax default and help with possible cases of hardship by probing their situation and offering payment plans

Technology

Two concurrent models (a deep neural network, and random-forest model) were trained to predict payment behaviour and, on that basis, provide guidance and recommendations for debt-management officers. The task's specialist nature required some degree of explainability so that the tax officers could justify the decision to contact a citizen. The ensemble of two models enabled striking a balance between performance and explainability: while the inscrutable deep-learning model facilitates solid performance, the random-forest model offers some explainability by depicting decision trees.

Also, developers devised a user interface that listed recommended actions and presented the AI model's work to debt-management officers graphically: these customer journey maps pinpointed specific behaviours (over a three-year span) that officers could click on to access specific details.

People

Debt-management officers used the models' predictions with the aid of the user interface, and made the final call on the citizen's situation (for example, offering payment plans or reminders of any eligibility for deductions). Augmenting the existing task with AI-based insights called for upskilling these officers. The agency invested in training officers to use the outputs and the interface effectively. Data scientists began educating tax experts on the principles of the AI system's design, including its input data, functionality, limitations, and output decisions. These customer-service staff had to understand the factors that the models consider and the outputs' probabilistic nature. The work of the tax officers changed slightly as they needed to start collecting more detailed textual data from customer interactions that could be used for improving the model's sensitivity. Thus, technical elements were introduced to the officers' previously non-technical role.

Control structure

The agency maintained a human-in-the-loop design as a way to control, verify, and sometimes challenge AI suggestions. Hence, the AI model was implemented to augment debt-management officers' work. The rationale for this choice is that each citizen's circumstances are different and AI 'is not 100% correct' – it can act mindlessly. Officers were trained to work with the AI interface and to make the final decisions recognising individuals' unique circumstances.

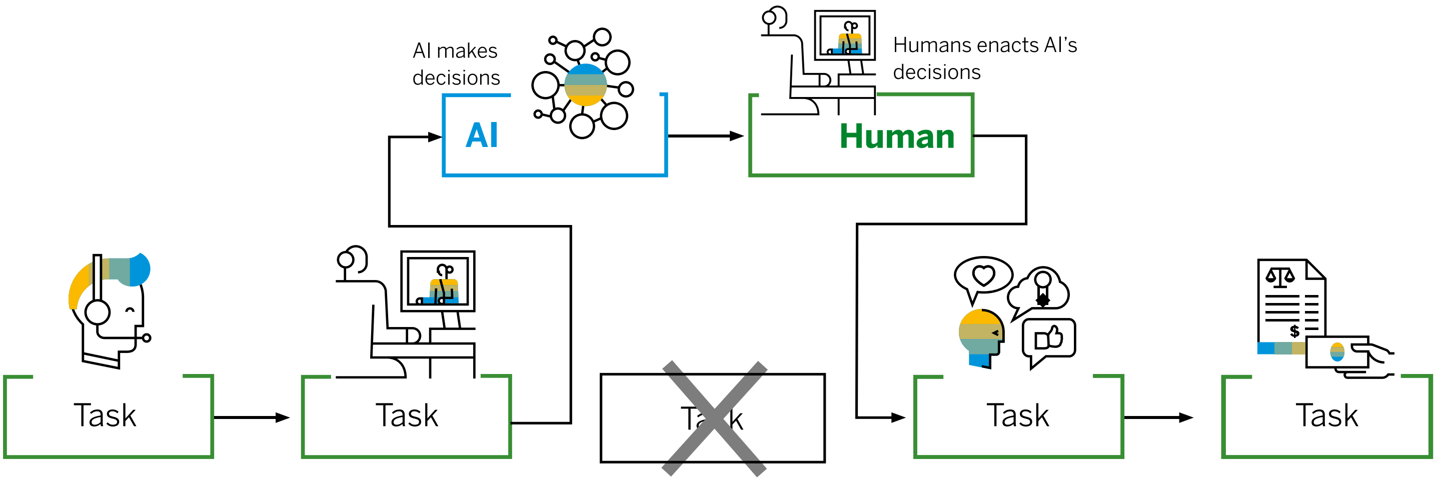

Figure 4: Design Choices for Expert Tasks

- **The task:** maintaining the core task and enriching it by adding data-driven insights to guide decisions
- **Technology:** augmenting work through guidance for human action, with balance between performance and transparency supported by a user interface explaining AI output
- **People:** reskilling such that human workers understand the logic behind the AI decisions and to effectively enact the AI's decisions. Additionally, the AI contributes to professional development by improving understanding of the tax domain, new roles added for humans to capture data and improve models
- **Control structure:** keeping the locus of decision-making with the humans, limiting AI's actions guarantees that decisions and actions are justified

Work Redesign for Wicked Tasks

Case study 4: Health AI

Context: an Australian public hospital's emergency ward waiting room, which typically faced surges of patients seeking clinical care, needed to scan patients for sepsis to ensure timely treatment.

The work's original design: the emergency department's triage nurses scanned incoming patients to identify those at risk of sepsis and prioritized their treatment such that those in need received timely intervention. Nurses would collect health data from patients, using a paper-based survey, and check for any values beyond sepsis thresholds. They would also visually observe patients in the waiting room to detect any signs of sepsis. Nurses were responsible for prioritizing the patients for clinical care, and the final decision for detecting sepsis relied on doctors.

The problem: while sepsis is treatable if diagnosed in time, many patients die because of a lack of timely diagnosis, due to delays in the prescription of antibiotics. Long queues in emergency departments, and significant clinical workloads, presented challenges for timely treatment. On the other hand, giving antibiotics to someone who does not have sepsis is not desirable either. The hospital wanted to reduce the number of people at risk of dying of sepsis whilst not prescribing antibiotics needlessly.

Redesigned work for Health AI: The data science team associated with the hospital trained an algorithm to identify at-risk patients in the emergency ward. As patients were triaged a new visual interface showed a real-time ranking of patients based on their level of sepsis risk. Nurses would monitor both the interface and patients in the waiting room, enacting suitable clinical treatment pathways for those identified as high-sepsis-risk patients.

DESIGN CHOICES FOR WICKED TASKS (HEALTH AI)

For this type of task, AI enabled redesigning the process, addressing the problem at hand.

Let us look at the changes made across the four key work design elements.

The task for AI

The existing task was augmented with an AI system designed for early detection of risk of sepsis in the hospital's emergency department. Predicting potential patients at risk of sepsis was classified as a wicked task because:

- Correctly detecting and diagnosing sepsis requires a high degree of specialization in the form of medical expertise
- The task presents significant risk in that a) lack of timely intervention could result in patient mortality in the event of false negatives and b) prescribing antibiotics to someone who does not have sepsis is costly and may also cause harm to the patient

Technology

An AI model was trained by means of various data sources, including electronic medical records, clinical notes, biomedical patterns, and patients' vital signs checked in the emergency room. Both explainability and performance were highly important, in light of the critical nature of the task. However, explainability was prioritized due to trust and accountability that the health context demands. After experiments involving several algorithms, with varying levels of explainability, the implementation team chose a model that takes advantage of logistic regression and gradient boosting techniques. This model performed as well as black-box ones – sometimes better – and its transparency afforded building a dashboard that not only shows the risk ranking but also explains why each patient was identified as having a high or low risk level. The model's decision-making was limited to presenting risk levels, and it made no final decisions. Its purpose was to assist the nurses with what they were already doing.

People

Nurses had been relying on simple vital-sign cut-off thresholds to identify whether a patient was at risk of sepsis; the clinical director gave the example that "people would look at patients and [find that since] a pulse of more than 120 is sepsis-positive, pulse 119 is not sepsis-positive". Augmentation with AI highlighted a need to reskill nurses to move away from the old paradigm and build a more nuanced understanding of the complexity of sepsis, beyond relying on simple rules of thumb. Moreover, discovering that nurses were too busy to scan patients for signs of sepsis prompted the creation of a new role: an AI-augmented triage nurse specifically responsible for monitoring deterioration in the emergency room's status, by referring to the AI system and observing patients in the waiting area.

Control structure

To design the system, the data science team sourced guidelines from their central AI ethics committee that had developed detailed standards for AI design and implementation in the public sector (for example, for reducing biases, increasing transparency, governing models), and actively engaged with them to ensure ethical design and implementation. In the emergency room, AI was used as a monitoring tool with a new hybrid role to review the tool. Other than AI monitoring, all patients were still reviewed by triage nurses to ensure work burden did not introduce mistakes into nurses' triage process.

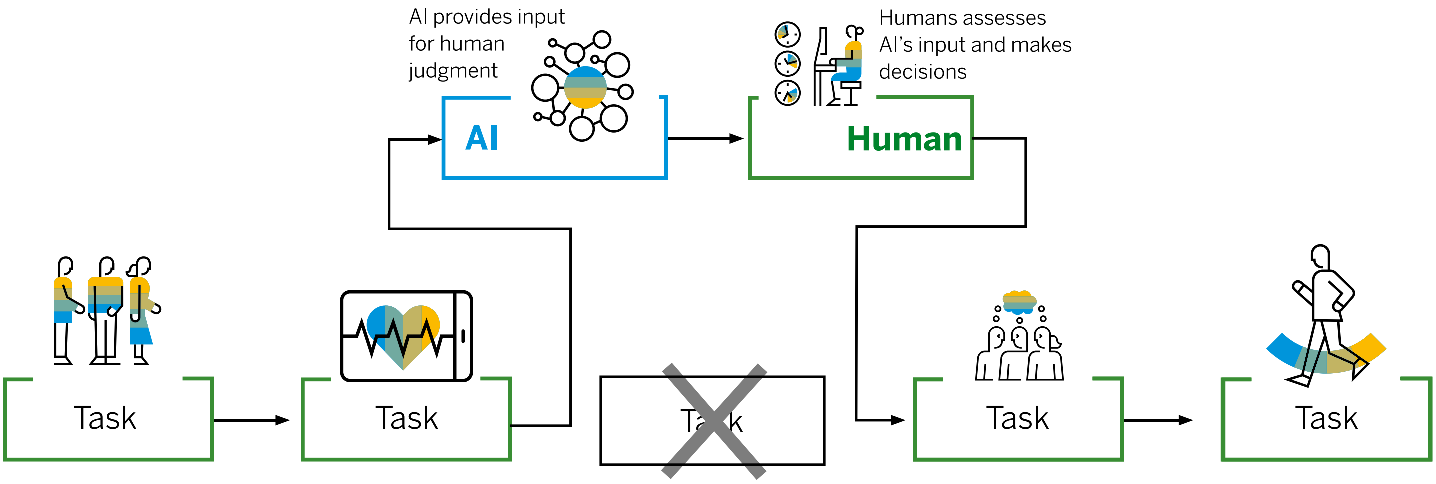

Figure 5: Design Choices for Wicked Tasks

- **The task:** retaining the core task and enriching it by adding a data-driven component to inform decisions
- **Technology:** augmenting work through informing human decisions, with transparency assigned precedence over performance, including a user interface to explain AI outputs
- **People:** reskilling humans to work with AI, creating a new role for AI-augmented sepsis detection
- **Control structure:** limiting AI's autonomy, integrating extensive ethics oversight and preventing alert fatigue

Work Design Choices in Summary

Synthesising key conclusions from the case studies, the final table presents a condensed characterization of appropriate design decisions for each of the four task types.

Recommendations stemming from this summary can help take work with AI in the right direction, toward mature design and responsibility.

	Low degree of specialization	High degree of specialization
Significant risk	**Task type:** high-reliability tasks **Technology:** supervised automation; model performance prioritized over transparency **People:** parallel processes run by humans; reskill humans to support AI operations **Control structure:** humans as active supervisors in the field	**Task type:** wicked tasks **Technology:** augmentation through informing human decisions, prioritize transparency over performance, build an explanatory user interface **People:** reskill humans and define new roles dedicated to AI work **Control structure:** limit AI's autonomy, extensive ethics oversight, prevent alert fatigue
Limited risk	**Task type:** mundane tasks **Technology:** full automation; priority for performance over transparency **People:** limited human intervention, reassign humans to new roles **Control structure:** escalate to humans when errors occur	**Task type:** expert tasks **Technology:** augmentation through guiding human action, balance between performance and transparency, build an explanatory user interface **People:** reskill humans, humans and machine both engaging in learning and development **Control structure:** limit AI's autonomy, humans justify AI decisions, prevent automation complacency

Assessing the potential risk of integrating AI within a process, creates the **optimal basis** for assuring the overall outcome of the program.

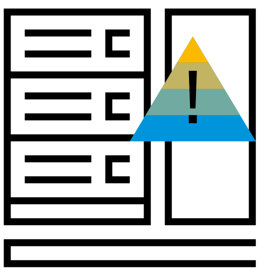

Recommendations

Integration of AI into work processes presents a unique set of challenges and opportunities. Within the public sector, a constant tension exists between demonstrating increased government productivity and driving social outcomes with limited resources, whilst ensuring the safety of its citizens. AI provides the sector with the ability to lessen this tension, as we have seen in the cases above, as long as the necessary guardrails are in place.

To further elevate responsible implementation of AI across the public sector we make the following recommendations:

1. Consider risks associated with using AI to automate or augment work upfront: best practice process design for AI needs to bring risk control and management to the forefront of AI implementation decisions, and utilize various configurations of design elements (task, technology, people, and structure) to enable and maximize responsible and ethical design.

2. Unpack your work processes and develop clear criteria on what tasks to delegate to AI and how: based on the outcomes of the risk assessment, carefully examine existing work processes and break them into subprocesses or tasks that can be delegated, with confidence, to humans or algorithms as appropriate. Separate the mindless, routine parts of the work process from those that need mindful judgement. Consider automation of the mindless parts, with appropriate level of human supervision depending on the risks involved. For tasks with high specialization requirements, use AI to augment work, and actively increase user engagement with AI decisions.

3. Avoid human skills erosion and AI model deterioration over time: AI's integration into work processes introduces opportunities to embrace continuous learning, for both humans and machines. Machines constantly learn from flows of new process data, and humans learn from the machine's new insights. For humans, enabling, absorbing, and acting on new knowledge requires agility and a desire to maintain human skills and expertise in an attempt to avoid automation complacency and skills erosion. The learning cycle equally applies to the AI. AI models need to be trained and retrained after implementation and enhancement, to avoid model drift and maintain high performance standards.

4. Implement control and governance structures to actively avoid citizen harm: rethink your existing control structures in new, innovative ways, to ensure work performance is appropriately monitored and risks are effectively managed over time. We have identified several ways to implement new control structures specific for AI, but more work and new approaches are required to ensure adequate guardrails and organizational levers are effectively used to manage and mitigate AI risks. This requires organizations to develop new instruments for measuring the impact of AI systems on citizens and other stakeholders. Such metrics need to go beyond the traditional financial realm to capture issues around bias and realization of human potential. They need to be integrated across reporting practices so as to enable a genuinely 360-degree view on the consequences of AI.

5. AI is not a set-and-forget technology - Be prepared to evolve your work processes in response to stakeholder feedback: don't cement your work processes, as using AI requires commitment to change, which could be triggered by the changing stakeholder needs and expectations, regulatory environment or the inadvertent outcomes AI creates for stakeholders. Such continuous change is not easy and can only be successful with strong leadership and change management skills.

We advise public managers to ensure there is **organizational commitment to, as well as investment in,** AI to develop systems that deliver benefits while minimising risks.

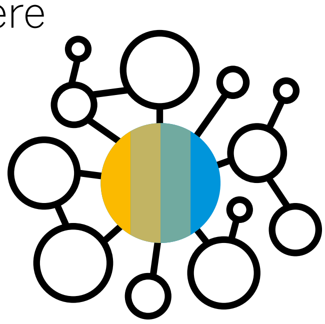

About the SAP Institute for Digital Government

We live in an increasingly disrupted world and are witnessing an unprecedented transformation of how governments, businesses, and citizens operate and interact.
This transformation is readily evident in the changing role of government as it addresses this disruption: increasing expectations of citizens in how they engage with government services; the ability of government operations to effectively and safely utilize the valuable data within and across the ministries; and creating secure and economically sustainable environments and delivering the mission of government in helping drive nation-building.

SAP has been a key enabler of government services and processes for over 30 years. As a global company, we have first-hand experience partnering with governments. In 2014, along with several academic and government institutions, SAP created the SAP Institute for Digital Government to support governments in responding to these challenges. The institute facilitates a forum for exchange of ideas and thought-leadership demonstrating the public value of digital government to tackle real-world, complex issues.

To learn more about the SAP Institute for Digital Government, please visit https://discover.sap.com/sap-institute-digital-gov/en-us/index.html.

The University of Queensland Researchers
Dr Reihaneh Bidar, Dr Ida Someh, Dr Tapani Rinta-Kahila, Dr Katie Williams, Professor Marta Indulska

SAP advisors
Ian Ryan, Ryan van Leent

Acknowledgements
We are grateful for the anonymous case organizations for their collaboration in data collection.

Citation:
Bidar, R., Someh, I., Rinta-Kahila, T., Williams, K., Indulska, M., and Ryan, I., (2023).
Redesigning Work with Artificial Intelligence: Accelerating the Adoption of Artificial Intelligence in Government Organisations
The University of Queensland and SAP SE.

87409enUS (23/03)