




AI in the public services in the United Kingdom – some principles for doing it well

Tom Symons
Head of Government Innovation Research

nesta



Nesta is an innovation
foundation.

We back new ideas to tackle the
big challenges of our time.

Why AI and ML have an important role in public services

AI has the potential to:

- Codify best practice and roll it out at scale
- Remove human bias
- Enable evidence-based decision making in the field
- Spot patterns that humans can't see
- Optimise systems too complex for humans to model
- Quickly digest and interpret vast quantities of data
- Automate cognitive activities that require significant human effort

Specific existing use cases in the UK

AI is already being used by governments and public sector organisations for specific activities such as:

1. Analysing case notes to determine whether a child is likely to be taken into care (Xantura)
2. Spotting tumours in medical scans (Google Deepmind)
3. Detecting fraudulent benefits / tax claims (e.g. DWP (Universal Credit))
4. Enabling smart chatbots to answer citizen questions (e.g. Pepper in LB Enfield)
5. Initial diagnosis based on symptoms reported at GP surgeries (Babylon Health)
6. Identifying unlicensed rental properties (LB Barking and Dagenham/Greater London Authority)
7. Providing recommendations for health and social care plans for individuals to choose from (LB Harrow w/ IBM Watson)

Specific existing use cases in the UK

AI is already being used by governments and public sector organisations for specific activities such as:

Providing recommendations for health and social care plans for individuals to choose from (LB Harrow w/ IBM Watson)



Analysing case notes to determine whether a child is likely to be taken into care (Xantura)



Spotting tumours in medical scans (Google Deepmind)



Initial diagnosis based on symptoms reported at GP surgeries (Babylon Health)



Identifying unlicensed rental properties (LB Barking and Dagenham/Greater London Authority)



Detecting fraudulent benefits / tax claims (e.g. DWP (Universal Credit))



Enabling smart to answer citizen questions (e.g. Pepper in LB Enfield)





Ethical concerns

Critics have raised ethical concerns that AI could be used by governments and the public sector in ways that **invade privacy; or cause harm, unfairness and moral wrongs.**



Calls have been made for new codes, standards and principles to be created.

AI can yield great power, but with great power comes great responsibility

Examples of concerns raised about the use of AI

These ethical concerns - and many proposed solutions to them - tend to focus on one or more of the following **three stages** of deploying an AI.

0101
1001



CREATION

FUNCTION

OUTCOME

How the AI is created

How the AI works

What the AI is used to do

Examples of ethical concerns raised about the use of AI

CREATION	FUNCTION	OUTCOME
Does the AI use data that invades individuals' right to privacy?	Are the assumptions used by the AI correct?	Is the AI being used to do something unethical?
Is the training data accurate and truly representative?	Are the factors used by the AI to make a decision reasonable and fair?	Is anyone responsible / accountable if a negative outcome is produced by an AI?
Does the training data contain historic biases that could be perpetuated?	Can anyone see and understand how the AI works and audit how a given output was created?	Will people know if a decision affecting them was made by an algorithm?
What happens when the use of the AI renders the training data out of date?	Can we be sure the AI is protected against hacking and manipulation?	What recourse will people have if an AI discriminates against them or causes them harm?

A number of these have already provoked controversy

Providing recommendations for health and social care plans for individuals to choose from (LB Harrow w/ IBM Watson)



Analysing case notes to

Fears algos will remove children from their parents



Spotting tumours in medical scans (Google)

Data privacy



Initial diagnosis based on symptoms reported at GP

Questions about accuracy



Identifying unlicensed rental properties (LB Barking and Dagenham/Greater London Authority)



Detecting fraudulent benefits / tax claims (e.g.)

Questions about accuracy



Enabling smart to

Discomfort with automation of human roles and dealing with a robot



Public sector as a special case?

- Monopoly provider of the services it offers
- Interacts with very vulnerable people
- Decisions may have significant consequences on a person's life
- Democratically elected governments have special duties of accountability

How should public sector
organisations approach
the ethics of AI?

Google AI Principles



1. **Be socially beneficial**
2. **Avoid creating or reinforcing unfair bias**
3. **Be built and tested for safety**
4. **Be accountable to people**
5. **Incorporate privacy design principles**
6. **Uphold high standards of scientific excellence**
7. **Be made available for uses that accord with these principles**

See: <https://www.blog.google/technology/ai/ai-principles>



1. **Fairness** - AI systems should treat all people fairly
2. **Inclusiveness** - AI systems should empower everyone and engage people
3. **Reliability & Safety** - AI systems should perform reliably and safely
4. **Transparency** - AI systems should be understandable
5. **Privacy & Security** - AI systems should be secure and respect privacy
6. **Accountability** - AI systems should have algorithmic accountability

See: <https://www.microsoft.com/en-us/ai/our-approach-to-ai>

UK Government initial code of conduct for data-driven health and care technology



Department
of Health &
Social Care

1. Define the user
2. Define the value proposition
3. Be fair, transparent and accountable about what data you are using
4. Use data that is proportionate to the identified user need (data minimisation principle of GDPR)
5. Make use of open standards
6. Be transparent to the limitations of the data used and algorithms deployed
7. Make security integral to the design
8. Define the commercial strategy
9. Show evidence of effectiveness for the intended use
10. Show what type of algorithm you are building, the evidence base for choosing that algorithm, how you plan to monitor its performance on an ongoing basis and how you are validating performance of the algorithm.

See: <https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology>

European Commission's High Level Expert Group on Artificial Intelligence



“Trustworthy AI” founded on fundamental principles of individuals' rights

1. **Accountability**
2. **Data Governance**
3. **Design for all (by all - include diversity)**
4. **Governance of AI Autonomy (Human oversight)**
5. **NonDiscrimination**
6. **Respect for Human Autonomy**
7. **Respect for Privacy**
8. **Robustness**
9. **Safety**
10. **Transparency**

See: <https://ec.europa.eu/digital-single-market/en/news/draft-ethics-guidelines-trustworthy-ai>

Smart Dubai - AI Ethics



Fair

- Demographic fairness
- Fairness in design
- Fairness in data
- Fairness in algorithms
- Fairness in outcomes



Accountable

- Apportionment of accountabilities
- Accountable measures for mitigating risks
- Appeals procedures and contingency plans



Transparent

- Identifiable by humans
- Traceability of cause of harm
- Auditability by public



Explainable

- Process explainability
- Outcomes explainability
- Explainability in non-technical terms
- Channels of explanation

See: <https://smartdubai.ae/initiatives/ai-principles-ethics>

Nesta principles for public sector use of AI

- 1 - Every algorithm should be accompanied with a description of its function, objectives and intended impact, made available to those who use it.**
- 2 - A description of the data on which an algorithm was trained and the assumptions used in its creation should be published, together with a risk assessment for mitigating potential biases.**
- 3 - A list of all the inputs used by an algorithm to make a decision should be published.**
- 4 - Citizens must be informed when their treatment has been informed wholly or in part by an algorithm.**
- 5 - Every algorithm should have an identical sandbox version for auditors to test the impact of different input conditions.**

nesta



- 6 - When using third parties to create or run algorithms on their behalf, public sector organisations should only procure from organisations able to meet Principles 1-5.**
- 7 - A named member of senior staff (or their job role) should be held formally responsible for any actions taken as a result of an algorithmic decision.**
- 8 - Public sector organisations should commit to evaluating the impact of the algorithms they use in decision making, and publishing the results.**

What do we make of
these different codes,
standards and principles?

There's lots of overlap in the recommendations from codes, standards and principles:

CREATION	FUNCTION	OUTCOME
Reveal / publish the training data	Make the code of the AI transparent and open for inspection	Ensure intended and actual outcomes are fair, transparent and aligned with human values
Identify and minimise bias in the training data	Do not create or procure black box AIs	Ensure outcomes can be explained
Respect privacy / don't use data in ways that are creepy	Identify and minimise bias and limitations in the AI's assumptions	Ensure there is a process of oversight and evaluation
Do not use data on sensitive factors such as race and religion	Ensure the factors and function of the AI can be explained	Ensure a person is accountable for decisions made using the AI
Use personal data in compliance with GDPR	Offer identical sandbox versions of the AI to enable testing	Ensure outcomes are fair, inclusive and respect dignity and rights
	Ensure fairness by design	Ensure use of AI is known and there is a process of appeal
	Protect from manipulation and hacking by design	Mitigate against harms



But it's not that simple...

Different levels of complexity of AI have different consequences for ethics

Complexity	3	Unlimited quantities of unstructured training data such as video, photos, sound, free text	Dynamic model constantly evolving based on live data	Model used as one part of long and complex decision-producing chain
	2	Defined quantity of structured or unstructured training data used for a one-time creation of model	Static model created using one-time machine learning process	Model used to cover simple and clearly defined point of decision making process
	1	Defined number of structured datasets used for one-time weighting of model	Static model using human inputted rules weighted by machine learning	Model used to cover simple and clearly defined point of decision making process
		CREATION	FUNCTION	OUTCOME

nesta

0101
1001





Challenges for ethical approaches

Is it really possible to be transparent about the training data and assess for bias if it's unlimited and unstructured - e.g. thousands of hours of CCTV footage?

And can we meaningfully talk about 'explainability' if not even the developers of an AI know how it reasons? See example of AlphaGo AI, which played Go against itself to learn the optimal strategy.

Challenges for ethical approaches to Levels 2 and 3

In the tables that follow, for each recommendation:

✓ indicates it's straightforward / possible

X indicates it's extremely hard / impossible

~ indicates it's only possible in some circumstances

Viability of ethical recommendations for the 3 levels of AI: Creation

CREATION	Level 1	Level 2	Level 3
Recommendations	Defined number of structured datasets used for one-time weighting of algorithm	Defined quantity of structured or unstructured training data used for a one-time creation of algorithm	Unlimited quantities of unstructured training data such as video, photos, sound, free text
Reveal / publish the training data			~
Identify and minimise bias in the training data	~	~	~
Respect privacy / don't use data in ways that are creepy			~
Do not use data on sensitive factors such as race and religion		X	X
Use personal data in compliance with GDPR			~

Viability of ethical recommendations for the 3 levels of AI: Function

FUNCTION	Level 1	Level 2	Level 3
Recommendations	Static algorithm using human inputted rules weighted by machine learning	Static algorithm created using one-time machine learning process	Dynamic algorithms constantly evolving based on live data
Make the code of the AI transparent and open for inspection			(but meaningless)
Do not create or procure black box AIs			
Identify and minimise bias and limitations in the AI's assumptions		~	X
Ensure the factors and function of the AI can be explained		~	~
Offer identical sandbox versions of the AI to enable testing			X
Ensure fairness by design		~	~

Viability of ethical recommendations for the 3 levels of AI: Outcome

OUTCOME	Level 1	Level 2	Level 3
Recommendations	Algorithms used to cover simple and clearly defined point of decision making process	Algorithm used as one part of long and complex decision-producing chain	Algorithm used as one part of long and complex decision-producing chain
Ensure intended and actual outcomes are fair and transparent			
Ensure outcomes can be explained		~	~
Ensure there is a process of oversight and evaluation			
Ensure a person is accountable for decisions made using the AI			~
Ensure outcomes are fair, inclusive and respect dignity and rights			~
Ensure use of AI is known and there is a process of appeal		~	~
Mitigate against harms			

Where does this leave us?

A one size-fits-all approach that covers all instances of AI is unlikely to work, unless it's so high-level as to offer little practical guidance... But that might be ok

Any code that is created needs
to cover private sector partners
providing services to
government

Most public sector applications of AI are closer to level 1 than level 3, so implementing a code is possible... for now

The biggest unresolved issue is “explainability” - should we choose to avoid uses of AI we cannot adequately explain?

We need a diverse set of people involved at every stage of design, oversight and evaluation of AI

Can we have more emphasis on
- and faith in - professional
judgement?

But at the same time, recognise
the human dimension in the use
of these tools and the need for
education, training and care

Nesta's 10 questions for public sector use of AI

10 QUESTIONS TO ANSWER BEFORE USING AI IN PUBLIC SECTOR ALGORITHMIC DECISION MAKING


OBJECTIVE

Why is the algorithm needed and what outcomes is it intended to enable?



USE

In what processes and circumstances is the algorithm appropriate to be used?



IMPACTS

What impacts - good and bad - could the use of the algorithm have on people?



ASSUMPTIONS

What assumptions is the algorithm based on and what are their limitations and potential biases?



DATA

What datasets is / was the algorithm trained on and what are their limitations and potential biases?




INPUTS

What new data does the algorithm use when making decisions?



MITIGATION

What actions have been taken to mitigate the negative impacts that could result from the algorithm's limitations and potential biases?




ETHICS

What assessment has been made of the ethics of using this algorithm?



OVERSIGHT

What human judgement is needed before acting on the algorithm's output and who is responsible for ensuring its proper use?



EVALUATION

How, and by what criteria, will the effectiveness of the algorithm be assessed, and by whom?



@EddieACopeland #IdeaOnAPage

@nesta_uk

Get this infographic at
<http://bit.ly/IdeaOnAPage>