Data Protection Impact Assessment - template

Document Name	Machine Learning – Text classification- Data Protection Impact Assessment
Author/Owner (name and job title)	FOIA s.40(2) - Personal data that doesn't fall undi, Project Manager
Department/Team	Business Development Group
Document Status (draft, published or superseded)	
Version Number	1.0
Release Date	19/05/2021
Approver (if applicable)	
Review Date	
Distribution (internal or external)	Internal

Guidance for completing this template

Complete this Data Protection Impact Assessment (DPIA) template if your 'Screening Assessment - do I need to carry out a DPIA?' indicates a high risk to individuals. If you are unsure whether you need to complete a DPIA use the **screening assessment** first to help you decide.

Aim to complete your DPIA as early as possible as the outcome of the assessment could affect the viability of your plans. In extreme cases, you will not be able to continue with your plans without changing them, or at all.

Guidance notes are included within the template to help you with its completion- just hover your mouse over any blue text for further information.

The Information Management Service is also available for further advice and support. Please keep in mind our service standards if you require advice.

1. Process/system overview

1.1 Ownership

Project Title:	BDG 184b – Machine Learning – Text classification
Project Manager:	FOIA s.40(2) - Personal data that doesn't fa l unc
Information Asset Owner:	Mike Fitzgerald, Director of Digital, IT and Business Services
Data controller(s)	ICO
Data processor(s)	Microsoft *ICS.AI is the contracted developer of the text classification solution. They will be creating the solution within ICO managed environment.*

1.2 Describe your new service or process

In this document with reference to Customers it is defined as anyone who has sent an email into the dataprotectionfee@ico.org.uk mailbox which is made up of a majority of organisations although not exclusive. ICO will introduce a new service that will send an automated email to customers if they are making a change of address. Emails sent to the dataprotectionfee@ico.org.uk inbox will be classified using a machine learning service which will then automatically send an email.

Training Phase:
To develop the model ICS.AI will have to train the model to define the classification the emails will be measured against. This will require emails from ICO inbox dataprotectionfee@ico.org.uk.

During the training phase all data will remain within ICO managed services inside of UK region. Our privacy has been updated to ensure to include the use of machine learning. Once training has been completed no data will be left in the model and data will be deleted from the inbox dataprotectionfee@ico.org.uk as standard ICO email policy.

During the model training phase, no automated emails will be sent to customers.

Production:
When this service goes live all emails sent to dataprotectionfee@ico.org.uk will be classified and based on the model calculation they will receive a predefined automated response.

All emails should receive an automated response.

For a change of address classification, an email is sent to the customer containing information about how they can complete the change using the new change of address service online.
Other emails classified as a generic query will receive an automated email acknowledging receiving the request and an expected SLA of when the request will be completed.

Emails sent to dataprotectionfee@ico.org.uk typically contain information relating to a registration query i.e., name change or change of address. This information is publicly available on the data registration online.

However, there is no restriction on what information can be entered in the email, there is a risk of unintentional or intentional personal data being included by a customer and being processed.

The machine learning services will be within the ICO managed Azure subscription services and within the UK region. Data will not leave the UK during the processing or training.

A privacy notice will be updated to discourage users from entering personal data not related to registration queries and inform them of the email processing involved.

The dataprotectionfee@ico.org.uk retention is set as the standard ICO email policy.

Emails sent to the inbox will not be lost or altered by the machine learning process. The current design means that all emails sent to the inbox will be unaltered and will be processed as per normal by CRM.

All ICO registration agents with required permissions will be able to access and look at the inbox as normal.

Development for the machine learning text classification model has been contracted to a company called ICS.AI. They will be completing work within our managed ICO environment and will not hold or process any data outside of our enviroment. Contract is under the Gcloud framework.

1.3 Personal data inventory - explain what personal data is involved

Categories of data	Data subjects	Recipients	Overseas transfers	Retention period
Emails sent to the dataprotectionfee@ico.org.uk mailbox will have its contents classified. This will include the email address, subject title and the contents of the email which may contain information relating to registrations queries around address of trader/trading name/contacts for the registrations/payment categories or any other information they input in the email body. Email header information will be removed and is not processed by machine learning text classification service. No Cookies are used for this service.	Members of the public submitting email queries to dataprotectionfee@ico.org.uk	ICO Microsoft	None.	ICO standard email policy

This inbox is reserved for registration related queries and there are no restrictions preventing customers sending in personal information. A privacy notice will be updated on the website to inform users that any information sent to the above inbox will be processed using a text classification service provided by Microsoft Azure services which are managed within the ICO enviroment.				
Azure Database Raw text strings of email will be uploaded and through the text classification process data will be put into a format where statistics calculated from our training data is used to calculate probabilities for new data. Probabilities are calculated separately for each class. This means that we first calculate the probability that a new piece of data belongs to the first class, then calculate probabilities that it belongs to	Members of the public submitting email queries to dataprotectionfee@ico.org.uk	ICO Microsoft	None.	Data will be deleted after the training of the text classification model. This deletion of data will be done once the model has been created. The creation of the model and deletion of the data will be done by ICS.AI which is a contracted company for developing the text classification model. ICS.AI will be creating the model and deleting the data from the Azure database server within our managed ICO environment and will not hold or process any data outside of our enviroment.

the second class, and so on for all the classes				

1.4 Identify a lawful basis for your processing

ICO

The lawful basis for processing under GDPR article 6 is 6(1)(e) **Public task:** the processing is necessary for you to perform a task in the public interest or for your official functions, and the task or function has a clear basis in law.

1.5 Explain why it is necessary to process this personal data

The Machine learning project will provide a text classification service, that will allow ICO to response to customers based on the content of their email query.

Customers sending emails which are classed of type" change of address" will receive an automated email advising them of the new online service and further information required to process a change request in the form of a security pin.

This will efficiently direct customers to a new service and provide instructions on how to request a security pin. Reducing registration queries directed to the inbox.

Other customer queries will receive a generic response that acknowledges the receipt of the request and further information on expected response times. This will enhance our customer experience by reassuring their query has been received and reducing the number of follow up queries.

The text classification service will provide a quick and efficient way for customers to get the further information when requesting a change without contacting the helpline.

The project is necessary to reduce the number of staff to manually process incoming email. The current backlog is circa 4000 and the volume of emails is increasing daily. The backlog will increase further with expected growth of the registration. This back log and cannot be managed without hiring additional staff and working additional hours.

1.6 Outline your approach to completing this DPIA

We consulted our internal cyber security team to complete a security opinion reports.

Our internal information management team will be consulted to provide advise and guidance on data privacy.

Externally, ICS.AI have provided a proof of concept of the text classification service which was assessed before project was approved. ICS.AI have provided a number of similar services for the public sector and expert in their field.

ICS.AI are part of the UK government digital marketplace and approved for providing IT services.

Public consultation is not appropriate as internal processes are being updated to make ICO business processes more efficient. The original intent of the email is not being used for other purposes.

2.0 Data flows

2.1 Provide a systematic description of your processing, from the point that the data is first collected through to its destruction.

If your plans involve the use of new technology, you should explain how this technology works and outline any 'privacy friendly' features that are available.

High level Data Flow
Training process

FOIA s.31 - Law enforcement

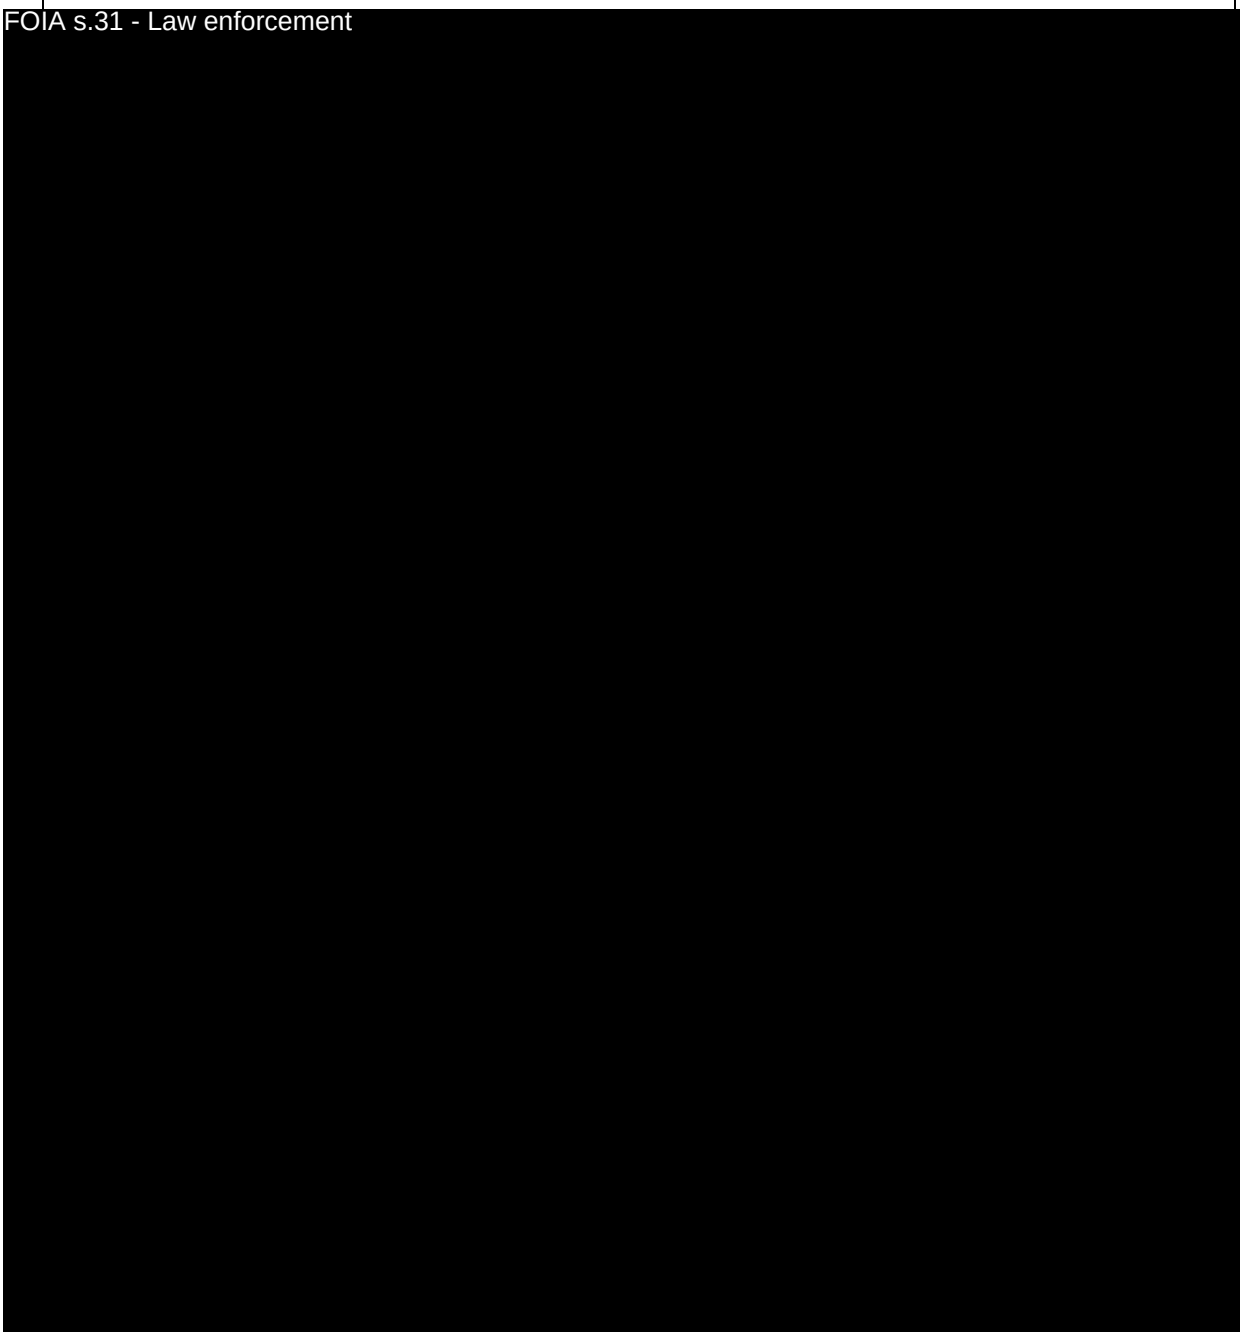

FOIA s.31 - Law enforcement

Production Dataflow.
High level Data Flow

FOIA s.31 - Law enforcement

3.0 Key principles and requirements

Purpose & Transparency

1. Will you need to update our privacy notices?

Yes ☒ No ☐

2. If you are not updating our privacy notices how do you intend to communicate information about your processing to the data subjects?

ICO AI guidance advises that the transparency information on training data is to be given prior of starting.

A privacy notice has been updated before the training phase to inform customers of the additional processing using machine learning services.

Although the service uses Machine learning technology for text classification no decision or judgement is made that impacts on the individual. All business processes remain unaltered and their request are not biased or altered in any way.

A privacy notice will be updated prior to the service going live.

3. If consent is your lawful basis for processing personal data are you maintaining appropriate records of the data subjects consent?

Yes ☐ No ☐ N/a ☒

4. If legitimate interests is your lawful basis for processing have you completed a legitimate interest assessment?

Yes ☐ No ☐ N/a ☒

If applicable, please provide a link to your completed assessment.

Accuracy

5. Are you satisfied the personal data you are processing is accurate?

Yes ☒ No ☐

6. How will you ensure the personal data remains accurate for the duration of your processing?

The data will not be altered, and original email will remain in the inbox.

7. If the personal data isn't being obtained directly from the data subject what steps will you take to verify accuracy?

N/A

Minimisation, Retention & Deletion

8. Have you done everything you can to minimise the personal data you are processing?

Emails headers will be removed and by the model as part of the text classification process.

Yes ☒ No ☐

9. How will you ensure the personal data are deleted at the end of the retention period?

A Retention policy of 12 months is applied on the mailbox. All training data will be deleted after the model has been created from the Azure Database while the original email will remain in the inbox for 12 months.

10. Will you need to update the retention and disposal schedule?

Yes ☐ No ☒

Integrity and confidentiality

11. Where will the personal data be stored?

Data will be stored in the ICO managed Azure subscription.

12. Are there appropriate access controls to keep the personal data secure?

Access to the training model will be restricted by role base access to the azure subscription. Access control is managed by ICO global admins who have access to the azure subscription which will be 2 users (Digital platform Architect and infrastructure Architect). The permitted users will be ICO IT admins and ICS.AI staff while creating the solution. Post training when the service is live access will be restricted to ICO IT admins and ICS.AI for support and maintenance of the service when required. ICS. AI may access ICO enviroment if the service is not working, they will not access the service as pat of Business-as-usual practices.

Yes ☒ No ☐

13. Have you contacted the cyber security team for a security assessment of your plans?

Yes ☒ No ☐ N/a ☐

If applicable, please provide a link to any assessment.

FOIA s.31 - Law enforcement

14. Please explain the policies, training or other instructions you intend to put in place to enable staff to operate the new system or process securely.

N/A service will only be accessible to specific users with elevated administration permissions.

Accountability

15. Who will be the Information Asset Owner for this personal data?

Director of Digital, IT and business services

16. Will you need to update our Article 30 record of processing activities?

Yes ☒ No ☐

17. If you are using a data processor have you agreed, or will you be agreeing, a written contract with them?

Yes ☒ No ☐ N/a ☐

Individual Rights

18. Is there a means of providing the data subjects with access to the personal data being processed?

Yes ☒ No ☐

19. Can inaccurate or incomplete personal data be updated on receipt of a request from a data subject?

Yes ☒ No ☐

20. Can we restrict our processing of the personal data on receipt of a request from a data subject?

Yes ☒ No ☐

21. Can we stop our processing of the personal data on receipt of a request from a data subject?

Yes ☒ No ☐ N/a ☐

22. Can we extract and transmit the personal data in a structured, commonly used and machine readable format if requested by the data subject?

Yes ☒ No ☐ N/a ☐

23. Can we erase the personal data on receipt of a request from the data subject?

Yes ☒ No ☐

Risk Description	Response to Risk	Risk Mitigation	Expected Risk Score		
			I	P	Total
			See Appendix 1 – Risk Assessment Criteria		
A customer receives an incorrect response as a result of the automated email response.	Accept	*Expected mitigation: The classification scope is limited to change of address and a generic response stating we have received the customers request and that which will be processed within an estimated timeframe. Customers incorrectly classified would receive the default response which is an acknowledgement. This will not have an impact on personal data. Only emails with an 80% certainty of a change of address request will be sent an email containing the link to change of address form.*	*1*	*1*	*1-Low*
Failure to provide transparency information prior to starting the Model training process	Avoid	*Expected mitigation:* *A privacy notice has been updated to inform customers the additional use for training purposes including machine learning.*	*1*	*1*	*1-Low*

User bias and discrimination.	Accept	*<u>Expected mitigation</u>:* *Users will NOT be categorized/labelled, and their requests will NOT be prioritised for preference.* *There is no decision made on the text classification that will alter the content and affect the processing of their email.* *By default, all current business processes remain, and the customer email will arrive in the inbox and be processed by the ICO registration department as current process. The Text classification will only send an acknowledgement email OR an email informing them of a change of address service.*	*1*	*1*	*1 - Low*

4.0 Risk assessment

5.0 Consult the DPO

Guidance: Submit your DPIA for consideration by the DPIA Forum. The process to follow is here.
Any recommendations from the DPOs team will be documented below and your DPIA will be returned to you. You should then record your response to each recommendation.

	Recommendation	Date and project stage	Project Team Response
1.	**See emails dated 11/03/2021 & 23/03/2021**	Planning 23.03.21	All recommendations accepted including Privacy notice updated on contact us page to include use of machine learning tools for training phase.
2	**Privacy notice before go live**	Planning 19.06.21	Privacy notice updated on contact us page to include statement we may send an automatic reply from insights from machine learning tools.

6.0 Integrate the outcomes back into your plans

Guidance: Identify who is responsible for integrating the DPIA outcomes. The outcomes include any expected mitigation you need to take as identified in your risk assessment and any further actions resulting from the DPOs recommendations.

Action	Date for completion	Responsibility for Action	Completed Date
Update Privacy Notice for training	03.03.21	Raymond Wong	03.03.2021

Delete Training Data	28.04.2021	ICS.AI	28.04.2021
Update Article 30 record of processing activities (3.0 Q18)	Before Go-Live	FOIA s.40(2) - Personal data that doesn't fall unde	19.05.2021
Update Privacy Notice before Go-Live	28.05.21	FOIA s.40(2) - Personal data that doesn't fall unde	19.05.2021

7.0 Expected residual risk and sign off

Guidance: Summarise the expected residual risk below. This is any remaining risk *after* you implement all of your mitigation measures and complete all actions.

It is never possible to remove all risk so this section shouldn't be omitted or blank. If the expected residual risk remains high (e.g. red on the traffic light scoring in the Appendix) then you will need to consult the ICO as the regulator by following the process used by external organisations.

The text classification and automated reply service is reliant on the Microsoft services being available and secure.

It is not possible to guarantee the service will always be available and that there is will be no security related incidents in future. The residual risk is low for Microsoft to be affected by either service interruption or security incident.

The overall residual risk is low , customers are not subject to bias/discrimination and the impact on users in an event of a false positive would be they receive an email with reference to a new change of address link.

7.1 IAO sign off

IAO (name and role)	Date	Project Stage
Mike Fitzgerald	19.05.21	Planning

8.0 Change history

Guidance: To be completed by the person responsible for completing the DPIA and delivering the system, service or process)

Version	Date	Author	Change description
V0.1		FOIA s.40(2) - Personal data th	First Draft
V0.2		Steven Johnston	Completed section 5.0 DPIA forum recommendations

Appendix 1: Risk Assessment Criteria

The following criteria are aligned with our corporate risk assessment criteria.

Impact

Impact is the consequence of a risk to the rights and freedoms of individuals being realised. Factors to consider include the financial harm or emotional distress that can be expected to occur.

Impact	Scoring criteria
Very low (1)	No discernible impact on individuals.
Low (2)	Individuals may encounter a few minor inconveniences, which they will overcome without any problem (time spent re-entering information, annoyances, irritations, etc).
Medium (3)	Individuals may encounter significant inconveniences, which they will overcome despite a few difficulties (extra costs, denial of access to business services, fear, lack of understanding, stress, minor physical ailments, etc)
High (4)	Individuals may encounter significant consequences, which they should be able to overcome albeit with serious difficulties (misappropriation of funds, blacklisting by financial institutions, property damage, loss of employment, subpoena, worsening of health, etc).
Very high (5)	Individuals which may encounter significant, or even irreversible consequences, which they may not overcome (inability to work, long-term psychological or physical ailments, death, etc.).

Probability

Probability is the likelihood of a risk to the rights and freedoms of individuals being realised. Factors to consider include the expected frequency of occurrence, and the motivation and capability of threat sources (eg does the threat require insider knowledge and/or significant technical resources to exploit any vulnerability?).

Probability	Scoring criteria
Very low (1)	0-5% - extremely unlikely or improbable For example, the risk has not occurred before or is not expected to occur within the next three years.
Low (2)	6-20% - low but not improbable

	For example, the risk is expected to occur once a year.
Medium (3)	21-50% - fairly likely to occur For example, the risk is expected to occur several times a year.
High (4)	51-80% - more likely to occur than not For example, the risk is expected to occur once a month.
Very high (5)	81-100% - almost certainly will occur For example, the risk is expected to occur once a week.

Risk level

Risk level is a function of impact and probability, and is represented by a RAG rating.

Probability / Impact	Very low (1)	Low (2)	Medium (3)	High (4)	Very high (5)
Very high (5)	Amber (5)	Amber (10)	Red (15)	Red (20)	Red (25)
High (4)	Green (4)	Amber (8)	Amber (12)	Red (16)	Red (20)
Medium (3)	Green (3)	Amber (6)	Amber (9)	Amber (12)	Red (15)
Low (2)	Green (2)	Green (4)	Amber (6)	Amber (8)	Amber (10)
Very low (1)	Green (1)	Green (2)	Green (3)	Green (4)	Amber (5)

Risk acceptance criteria

These criteria are guidelines only, and any risk treatment decisions should be made on a case-by-case basis. For example, it may be prudent to reduce a low risk because of legal and regulatory requirements.

Risk level	Acceptance criteria
Low (Green)	Within this range risks can be routinely accepted.
Medium (Amber)	Within this range risks can occasionally be accepted but shall be kept under regular review.
High (Red)	Within this range risks shall not be accepted and immediate action is required to reduce, avoid or transfer the risk.

Appendix 2: example risks to data subjects

Guidance: The following are examples of common risks associated with the processing of personal data to assist with your risk assessment. Not all of them will apply to your processing and the list is not exhaustive – you should consider specific risks that are relevant to your plans.

- Data is processed for unspecified / unlawful purposes/ not within expectations of data subjects
- Excessive data is processed
- Data is not kept up to date
- Data is kept for longer than is necessary by us
- Data is kept for longer than is necessary by data processor
- Data processed in contravention of data subject rights
- Data subjects unable to exercise their rights
- Data stolen or modified in transit
- Data stolen or modified at rest in our premises
- Data stolen or modified at rest in data processor premises
- Data transferred overseas to a jurisdiction that does not adequately protect data subject rights
- Re-identification of pseudonymised data by data processor or third party
- Unauthorised destruction or loss of data
- Data processor network / system / online portal not secure
- Data processor fails to process data in accordance with our instructions
- Personal data of children processed without appropriate safeguards / parental authority
- Consent of data subject not freely given (for example employer / employee processing)
- The data subject is particularly vulnerable (elderly or disabled) or is there a potential imbalance of power between the individual and the data controller (employee/employer)
- Source of data poses risks re accuracy (obtained from a unverified or old list)
- Risk to accuracy of data due to matching / combining data from different sources
- Use of new technology, e.g. fingerprinting, face recognition
- Monitoring or recording individuals
- Using profiling according to characteristics or behavior
- Non-compliance with DP principles

9.0 Template Change History (for Information Management Service only)

Version	Date	Author	Change description
v0.1	01/06/2020	Steven Johnston	First draft
v1.0	07/10/2020	Steven Johnston	First release
v1.1	07/01/2021	Iman Elmehdawy	Amendment to guidance note page 2.