Artificial Intelligence Policy

Table of Contents

1 Introduction

1.1 Overview

This policy will set out our organisational expectations when it comes to the use of Artificial Intelligence (AI) when working for Babergh and Mid Suffolk District Councils (BMSDC). It will provide necessary context and definitions, as well as the requirements that we expect all staff to follow when using AI in the context of their jobs.

Alongside this policy, we will offer Usage Guidance, showing worked examples of AI in action in the context of the local government environment.

1.2 Context

The Councils have used AI-styled tools and software for years. However, the recent significant proliferation and accelerating capabilities of AI-based tools and integration of AI features, particularly Generative Artificial Intelligence, into existing software, programs and processes, means that our governance and regulation thereof must also evolve. This policy is just one part of that.

We are not alone in introducing robust policies and guidance. Across the public sector, various organisations have worked to adapt to this ever evolving and ever-changing world. Some of the work done by other local authorities and public sector bodies has informed our own work. Below, you can find some of the documents we have used. If you would like more information about the best practice available, please contact ICT and Programmes.

- AI Playbook for the UK Government: This playbook offers guidance on using AI safely and effectively in government organisations.
- Guidance to Civil Servants on Use of Generative AI : This document provides general principles on the use of AI, including the importance of data privacy, accuracy and ethical considerations.
- A Guide to Using Artificial Intelligence in the Public Sector : Published by the Government Digital Services and the Office for Artificial Intelligence, this guide covers how to assess if AI will help meet user needs, how to implement

2 Policy Statements

2.1 Scope of Policy

This policy applies to all employees, councillors, or contractors who use AI tools in the course of their work for the council. We also refer to the duties of these parties when interacting with AI-generated content from third parties.

2.2 Statement

Several of the pieces of software that the Councils use as part of their day-to-day operations are AI based or have AI elements. However, other tools, particularly Microsoft Copilot, are made available to staff as an additional support. Use of Microsoft AI is not mandatory or required.

2.3 Responsibility for Output.

Regardless of whether an AI tool generates new content -or simply analyses existing information, you are responsible for use and output. None of the AI that we currently use in the Councils make independent decisions. As such, the user is entirely responsible for the output, and you **must** check and verify the product of AI you use.

3 Necessary Definitions

- Generative Artificial Intelligence (GenAI) - a form of artificial intelligence (AI) that generates text, images, or other media in response to prompts.
- Large Language Models (LLM) - A type of deep learning AI that processes language by identifying patters and consistencies in large training sets to simulate human language processing (e.g Copilot, Chat GPT, Claude).
- GenAI LLM - is a combination of GenAI and LLM presented in a single solution.
- Author - the individual who has requested output from a solution which provides GenAI content, analysis, or processing.
- Application Programming Interface (API) - a set of definitions and protocols that enables data transmission between separate software solutions.
- Plugins - a software component that adds a specific feature, often a connection to an external application or dataset, to an existing solution.
- Data Protection Impact Assessment (DPIA) - the process carried out on projects / changes involving personal data to help ensure compliance with data protection legislation and embed 'privacy by design'.

4 Data Protection

Using AI increases our risk when it comes to data security. When using AI, or considering acquiring any new system or software, keep the following principles

- Data Control: We ***must*** have full control over our data. This includes detailed knowledge of who has access to our data, the purpose of this data sharing and the expectations of how that data is stored/destroyed.

  - Where plugins or third-party applications are used data may be shared with the plugin provider and their data privacy policies will need to be reviewed. Any such requirement would potentially necessitate a separate and distinct project.

- Data Usage: Our data is used only to provide the service or purpose originally intended. ***It is not used to train AI models.***
- Security Measures: We ***must*** understand and comply with all of the security measures in place so that the security of our data is ensured.
- Compliance: Our data ***must*** be demonstrably handled according to regulations like GDPR.
- Monitoring: The system monitors for any abusive or harmful use of the service but ***must not*** store or review our data for these purposes.

5 Ethical Considerations

1. Job Context Usage*: **Only use the AI tool within your specific job context or for the purpose for which it is intended***. For example, a member of finance should only use AI in the context of their finance role, so that they may compare the outputs with their own knowledge to prevent inaccuracies.
2. Data Sensitivity: Never input sensitive or personal data into any AI tool, unless its safety is ensured. This includes classified information, sensitive data and personal identifiable data.
3. GDPR Compliance: Ensure that your use of AI complies with GDPR principles, respecting privacy and data protection laws.

4. Verify and Cite: Be aware that outputs from AI, particularly GenAI such as Copilot can be biased, experience hallucinations or contain misinformation. ***Always verify and appropriately cite any information generated by AI.***

5. Ethical and Responsible Use*:* Use AI ethically and responsibly, making sure it aligns with organisational values. ***Be transparent about AI use, ensure accountability, and avoid harm.***

6. Accountability*: **The end user is ultimately accountable for any use of outputs from any AI tool.*** As such verification, citing and transparency are vital.

7.Discrimination: GenAI must not be used to generate and publish content that is discriminatory, offensive, or inappropriate and careful checking of any content created must be undertaken.

6 Usage Policies

Below are more detailed policies around the use of AI. All users of AI in BMSDC will adhere to the Policy. In some instances, components of the Policy already exist elsewhere (Data Protection for instance) but they are repeated here for the sake of completeness.

Where possible a relevant example is provided. The example is provided solely to aid understanding and is in no way meant to be the only example of when or where the policy should be used.

6.1 Approved genAI Tools

- Staff are only permitted to use approved genAI tools for Work Purposes.
- Microsoft CoPilot is the only approved genAI tool within Babergh and Mid Suffolk District Councils.
- Should staff wish to adopt other genAI tools (or any other AI) they need to adhere to the principles and processes set out in the Procurement section of this document.
- The Artificial Intelligence Governance Boards directives must be followed in regards to any adoption AI.

6.2 Job Context Usage

- Use AI only for tasks related to your job role and area of professionalism.
    - Example: If you're a planning officer, use AI to support tasks like drafting planning reports, not for financial auditing.
    - Why: This ensures you can verify the accuracy of AI output and understand its limitations

6.3 Data Protection

- Never enter sensitive information or personal information into AI, unless its security is ensured. This includes classified or commercially sensitive information.
    - Example: Do not input residents' personal details like names, addresses, or council tax information.
    - Why: To protect privacy and comply with data protection laws.
- Follow GDPR rules to protect privacy. Ensure data minimisation, purpose limitation and respect data subject rights.
    - Example: You may ask AI to help you design a feedback survey on Waste Collection Services with a view to improving efficiency and satisfaction. Make sure you only ask for essential information (type of waste service used, frequency of collection, overall

satisfaction) avoiding the collection of potentially unnecessary personal details like full names, addresses or phone numbers. When the survey is launched ensure residents are aware of how their data will be used ensuring they have the option to opt out. Upon receiving the survey responses you may ask AI to help you analyse them – ensure you are clear about the purpose of that analysis – to provide efficiency and customer satisfaction improvement measures.

  - Why: To ensure compliance with legal requirements and protect individuals' rights.
- Ensure data accuracy before making decisions. Always verify the information provided by AI.
    - *Example:* Double-check AIs analysis before using it. You may ask AI to provide help building you a business case. If you don't provide it with costs it will likely make them up. Even if you do provide it with costs it might get its maths wrong.
  - Why: To make informed and accurate decisions.

6.4 Ethical and Responsible use.

- Use AI ethically and responsibly. Be transparent about its use to ensure accountability. Avoid causing harm.
    - *Example:* I have used Copilot to draft an email. I would add, as a footer to the email the following phrase: "I have used Microsoft Copilot to assist in drafting this email. AI helped draft the initial version of the email, and I made necessary adjustments to ensure clarity, accuracy and tone."
  - Why: To maintain trust, transparency and uphold ethical standards and accountability.
    - For more information on usage statements see section 9.
- Always verify and cite information as AI outputs can be biased or incorrect.
    - *Example:* Verify statistics generated by AI before including them in a meeting presentation. Ensure the relevant how statement is included in the Usage statement.
  - Why: To ensure the information is accurate and reliable and uphold ethical standards
- Take responsibility for the outcomes of using AI in your tasks. AI is a tool, but it is not operable without
  - Why: To ensure accountability and trust.
- Identify and mitigate any risks associated with using AI.
    - *Example:* Assess the risk of relying on AI for critical decision-making processes.

- Why: To prevent potential issues and ensure safe usage.

6.5 Reporting and Feedback

- Report any issues immediately. If you find information you shouldn't have access to, report it straight away to your line manager or to ICT and Programmes.
    - Example: If in using AI, the output from AI includes confidential HR information, and you are not routinely given access to this information, please report this as a matter of urgency
  - Why: To improve our information security.

6.6 Training and Support

- Complete all provided training sessions on how to use AI effectively.
    - Why: To ensure you are using AI correctly and efficiently
- Refresh your GDPR Training
  - Why: To ensure a relevant level of readiness in understanding your role in Data Protection.
- Follow all security policies to protect sensitive information.
  - Why: To minimise risk of data breaches and protect sensitive information.
- Use available user support resources if you encounter difficulties
  - Why: To resolve issues quickly and effectively. To aid in the identification of issues early.
- Stay informed about any updates or changes to AI.
  - Why: To make the most of new features and improvements.

6.7 Collaboration and Innovation

- Collaborate with other users to share insights and best practices.
    - Example: Work with colleagues in your wider department to explore new ways to use AI. For instance, if you are in the Housing Department, you may work with colleagues from Housing to explore new ways to use AI for tenant communications. Please remember that this is a use case and must be registered as such for review by the AIGB.
  - Why: To enhance the effectiveness and innovation of AI usage. To ensure best value of our use of AI.
- Inform the product team of new use cases.
    - Example: As an AI user in the Housing Department, you may work with colleagues from Housing to use AI to generate tenant

communications – this use case needs to be registered with the product team for review.
    - Why: To ensure transparency of use, assist with benefits identification and management of risk.
- Be open to exploring innovative ways to use AI within your job context.
        - Example: As an Exec Assistant, experiment with using AI to streamline the production of meeting minutes and actions.
    - Why: To find new and efficient ways to perform tasks.

6.8 Organisational Improvement Participation

Users of AI tools within BMSDC, especially AI license holders and product owners, agree that they will support the ICT team in maturing the wider improvement of digital maturity at the organisation.

- Support the governance board in making informed decisions about AI.
        - Why: To ensure the tool is used effectively and responsibly.
- Adhere to any decisions made by the governance board.
        - Why: To maintain consistency and accountability.
- Participate in the evaluation of to provide comprehensive feedback.
    - Why: To improve the tool and its implementation.
- Keep the organisational objectives in mind and work towards achieving them.
    - Why: To ensure our use of AI is aligned to organisational.
- Track and report on performance metrics related to your use of AI.
    - Why: To measure the impact and effectiveness of the tool.
- Report any access issues immediately. If you find information you shouldn't have access to, report it away. All users of AI in BMSDC must report any access issues to their line manager.
        - Example: If in using AI, the output from AI includes confidential HR information, please report this to via ….
    - Why: To improve our information security.
- AI, as an immature technology, will fail. If a prompt fails, please retry a few times. Issues related to non-performance should be raised with the product team.
    - Why: To monitor the stability of the tool and to capture core issues during the trial.
- Provide feedback on your use of AI and support your team in its adoption
        - Example: Share any changes to AI prompts that you have made that have improved the outputs AI has produced.
    - Why: To improve the tool and its usage.
- Engage actively in discussions related to AI.

        - Example: Participate in project updates to the wider organisation to discuss the benefits and challenges of using AI.
    - Why: To share knowledge and improve the tool's usage.
- Share your knowledge and experiences with new users of AI.
    - Why: To help others use the tool effectively.

7 AI Usage Policy specifically for Coding, SQL, and Reporting outside of Corporate ICT.

This section of the policy outlines the acceptable use of AI in software development and data workflows, ensuring that only enterprise-approved tools are used by staff outside of Corporate ICT.

All consumer-facing or unapproved AI tools are prohibited.

Using unapproved AI tools in any Coding, SQL or Reporting context including coding, SQL generation, or codified data analysis (e.g. Python, Dax etc) is a violation of this policy.

Only Microsoft Copilot is approved for use.

- Approved AI tools may be used for:
    - Drafting or reviewing code
    - Generating SQL, DAX, or reporting logic

But only if:

- No sensitive, proprietary (e.g. third party owned data schemas), or personal data is shared in prompts
- Output is manually reviewed by qualified Corporate ICT staff and successfully tested in non-production environments with the knowledge and support of Corporate ICT.
- AI-generated logic is clearly attributed and documented in code comments or version control
- Any code or queries assisted by AI must include a comment, e.g.:
  // Generated with Microsoft Copilot, reviewed by (*Developer Name)*
- Maintain audit trails for AI-assisted reporting, especially in compliance-sensitive contexts (e.g., financial reporting, compliance reporting).

Violations of this policy may result in disciplinary action.

To report concerns or misuse, contact:

- Your team lead or manager and the BMSDC ICT and Information Governance Team.

8 Statements of Usage

Whenever you use AI to produce anything you must include a statement to describe that AI was used to produce it and more specifically how it was used to produce it.

This will look like {Usage Statement}. {How Statement(s)}.

See below for the Usage Statement and How Statements we expect you to use.

8.1 Usage Statement

Below is an example of a Usage statement. Usage statements must be prominently displayed in the output created and shared with others.

“I have used <name of AI tool> to assist in <insert basic description of Output Type>.”

Output Types will be many and varied, such as Email, Report, Summary, Minutes, Data Analysis. Please use the most relevant.

For any reports/documents/presentations the Usage Statement must be included on the title page, for emails the footer and for all other potential artefacts the most relevant and prominent position.

This Usage Statement must be supplemented with all the relevant How Statements.

8.2 How Statements

As mentioned elsewhere in this document verification of AI output is vital. The below statements **must be added** to the usage statement to show how you have physically verified the content of the output.

- **Content Generation**: “Artificial Intelligence (insert name here) was used to generate the initial draft of the content, which I then reviewed and refined.”
- **Data Analysis**: “(tool name here) assisted in the preliminary analysis of the data, and I verified the results.”
- **Research**: “(Tool name) was employed to gather and summarise research findings, which I cross-checked for accuracy.”
- **Email Drafting**: “(tool name) helped draft the initial version of the email, and I made necessary adjustments to ensure clarity, accuracy and tone.”

More examples will be available in the Usage Guidance document.

9 Procurement and Introduction of Artificial Intelligence Tools

9.1 Scope

This section applies to any system, software, service or functionality that incorporates AI or is AI-based, whether:

- It is newly procured
- Already in use by the Councils
- Introduced through an upgrade, update, configuration change or feature enablement by an existing supplier

This includes but is not limited to:

- Stand-alone AI tools
- AI functionality embedded within existing systems
- “Optional”, “beta”, “preview” or “included at no additional cost” AI features
- Supplier-managed AI processing that acts on Council data

This section applies regardless of whether a formal procurement exercise is taking place.

9.2 Relationship with the Good Start Board

Where a proposal involves new or materially changed AI capability, it should be presented to the Good Start Board at the earliest opportunity.

The Good Start Board acts as an early-stage governance and assurance checkpoint to:

- Identify where AI is involved, including where it is embedded in existing supplier solutions
- Surface potential data protection, ethical, operational, financial, or reputational risks
- Ensure proposals are routed through the appropriate governance, assurance and approval processes, improving organisational awareness of proposed projects.

9.3 Introduction of AI via Existing Suppliers

Where an existing supplier introduces, enables, or proposes to enable AI-based functionality within a system already used by the Councils, this must be treated as the

introduction of AI, not business-as-usual change, and therefore must also be brought through the AI governance Board.

Examples include (but are not limited to):

- AI-assisted summarisation, classification, recommendation, or decision-support
- Generative AI used to draft content, responses, reports, or communications
- AI-based analytics, prediction, or automated processing
- AI functionality acting on resident, staff, partner, or Council data

Such functionality must not be enabled, piloted, or used until it has been reviewed in accordance with this policy.

9.4 Assessment and Approval

Before any AI-enabled system or functionality is procured, enabled, piloted or trialled, or is expanded beyond an existing approved use please contact ICT, Information Governance or Procurement for further information on completing the following:

- An AI Evaluation (or equivalent agreed assessment)
- Confirmation of data flows, including processing, storage, and retention
- Consideration of data protection requirements, including whether a DPIA is required
- Assessment of ethical, operational, and reputational risks
- Confirmation of human accountability for outputs and decisions

The outcome of this assessment must be recorded on the Councils' AI Register.

9.5 Supplier Responsibilities and Assurances

Where AI functionality is introduced via a supplier, the supplier must be able to clearly confirm:

- Whether Council data is used to train, fine-tune, or improve AI models
- Where AI processing takes place, including whether data is processed outside the UK and who, if anyone, it is shared with as part of processing.
- What safeguards exist to prevent unauthorised access, data leakage, or misuse

- How AI outputs can be reviewed, challenged, or overridden by Council staff

AI functionality must not be enabled by default without explicit agreement from the Councils.

9.6 Ongoing Change and Monitoring

AI governance does not end at initial approval.

Service owners, contract managers, and project leads must:

- Remain alert to supplier updates that introduce or expand AI capability
- Re-assess AI functionality where scope, data usage, or risk materially changes
- Escalate concerns or uncertainty through appropriate governance routes, including ICT and Programmes

Failure to appropriately identify, assess, or govern AI functionality may result in suspension of use or contractual review.

10 Acceptance of AI-Generated Submissions

As an organisation, we recognise the growing use of artificial intelligence (AI) tools to generate content for public and stakeholder submissions. Generally, though there may be service exceptions, we accept AI-generated submissions, provided they meet the standards of accuracy, legality, and accountability as per submission guidelines.

10.1 Transparency and Accountability

- Submitters are responsible for the content of their submissions, regardless of whether AI tools were used in their creation.
- We encourage all external contributors to clearly disclose when AI has been used to generate or alter content, in line with emerging best practice and sector guidance.
- Where required by law or guidance from relevant authorities, such disclosure may be mandatory.

10.2 Service Specific Requirements

- Individual services within the organisation may be subject to additional requirements or guidance from governing bodies, regulators, or professional

associations which may take precedence over the submission guidance in this document.
- Services should regularly review and, where necessary, update their procedures to ensure compliance with the latest advice, guidance, and legislation relevant to their domain.
- Services should seek advice from BMSDC ICT and Information Governance prior to implementing their own service specific submission guidance.

10.3 Review and Oversight

- Submissions identified as AI-generated should not be dismissed outright unless service specific or overriding policy takes precedence. Where service specific policy does not take precedence, staff should apply appropriate checks for accuracy, originality, and compliance, especially where content appears unusually polished or raises concerns about authenticity.
- Where the use of AI could compromise the integrity of a process, the relevant escalation and review procedures should be followed.

10.4 Education and Internal Guidance

- Service areas will be responsible for ensuring staff are briefed on how to handle AI-generated submissions, including recognising potential risks and maintaining a consistent, fair approach.
- The organisation will provide online materials that inform the public and stakeholders about the responsible use of AI in submissions.

11 Appendix 1. Artificial Intelligence: Technologies, Risks and Mitigations

This appendix provides an overview of common Artificial Intelligence (AI) technologies, the key risks associated with their use, and the policy controls that mitigate those risks within Babergh and Mid Suffolk District Councils.

This appendix should be read alongside the main policy sections on accountability, data protection, ethical use, governance, and approved tools.

11.1 Generative AI

Generative AI uses Large Language Models and natural language prompts to produce new content such as text, summaries, images, or drafts.

Within the Councils, Microsoft Copilot is the only approved Generative AI tool for work use. Copilot uses Natural Language Programming and Large Language Models to generate novel content.

11.1.1 Key Risks of Generative AI

- Generation of inaccurate, misleading, or inappropriate content
- Surfacing of information the user should not have access to
- Inappropriate reliance on AI output without professional judgement
- Reputational harm through publication of unverified AI-generated content

11.1.2 Mitigating Policy Elements

- AI outputs do not replace human judgement; users remain fully accountable for verification and use
- AI use is restricted to job-context-appropriate tasks
- Sensitive or personal data must not be entered into AI tools
- Usage Statements and How Statements ensure transparency and auditability
- Use is limited to approved AI tools only
- Information incorrectly surfaced should be reported as a data breach as per the Acceptable Use Policy

11.2 Large Language Models

Large Language Models emulate and respond to human language, generating responses by predicting likely word sequences based on training data. They underpin Generative AI systems such as Copilot.

11.2.1 Key Risks of LLMs

- Confident presentation of incorrect or outdated information
- Embedded bias within training data influencing outputs

- Limited explainability of how outputs are generated
- Inconsistent outputs to similar prompts.

11.2.2 Mitigating Policy Elements

- Mandatory verification, validation, and citation of AI-generated information
- Explicit requirement for ethical, responsible, and non-discriminatory use
- Clear policy position that AI cannot make independent decisions
- Human accountability for all outputs and decisions

11.3 Natural Language Processing

Natural Language Processing enables systems such as chatbots to interpret and respond to written language by identifying linguistic patterns.

11.3.1 Key Risks of NLP

- Misinterpretation of user intent or context
- Inappropriate automated responses
- Risk of personal data exposure if misused
- Over-reliance on automated responses

11.3.2 Mitigating Policy Elements

- Strict data minimisation and data protection requirements
- Prohibition on entering personal or sensitive information
- Requirement for human oversight and review of outputs
- Clear reporting and escalation routes for unexpected behaviour or access issues

11.4 Optical Character Recognition

OCR systems extract text or identifiers from images, commonly used in areas such as parking enforcement and document processing.

11.4.1 Key Risks of OCR

- False positives or recognition errors
- Financial or reputational impact from incorrect enforcement
- Data protection risks relating to image capture and transmission
- Over-reliance on OCR outputs as definitive evidence

11.4.2 Mitigating Policy Elements

- GDPR-aligned data handling, including purpose limitation and security controls
- Governance oversight where OCR is introduced or enabled via suppliers
- Clear human accountability for enforcement and decision-making

11.5 Deep Learning

Deep Learning systems use layered neural networks to identify complex patterns and make predictions. These systems are not currently used by the Councils.

11.6 Computer Vision

Computer Vision is an exceptionally limited version of AI, which informs the facial recognition technology on some Council laptops.

11.7 Robotic AI

Some drones incorporate minor directional AI into their internal gyroscopes in order to maintain their direction,

11.8 Swarm AI

The Councils do not currently deploy Swarm AI, mainly seen in autonomous drones.