

Artificial Intelligence Best Practice Approach at Tendring District Council

A practical guide on the effective and ethical use of artificial intelligence.

Version Control

Version 1.0 – 20 December 2024

Summary

While AI will help us significantly in our daily tasks in the future, human oversight is essential to ensure accuracy, ethical use and decision making; and humans are accountable for both the correct process and the end product.

This will include verifying information, editing content, and ensuring there is no wider risk to the organisation.

If you follow the guidance below – and work with the wider organisation to develop and use best practice – then AI can be a real force for good for Tendring District Council (TDC).

Contents

1. Introduction
2. About this guide
3. Benefits of AI
4. Keeping up-to-date
5. Ethical and Safe Use Considerations
 - Accountability
 - Information Governance
 - Cyber Security
 - Fairness and bias in AI
 - Avoiding misinformation
 - Models to use
 - Copyright, Public Sector Information and Intellectual Property Rights
 - Decision-making
 - Transparency
 - Environmental sustainability
 - De-skilling
 - Accessibility and web content
 - When not to use AI
6. Best practice use of AI
 - Information
 - Effective prompts
 - Iterative process
 - Knowing when to use AI
7. Glossary
8. Acknowledgements

Introduction

Artificial Intelligence (AI) is one of the fastest growing fields of technological innovation, which has seen an explosion of readily-available tools in the past few years.

It is becoming increasingly prevalent, and already has the ability to improve the efficiency and quality of work across the authority; though it is perhaps some way off maturity in terms of realising high expectations for medium or long-term automation.

AI will not replace humans; but staff who can effectively leverage and use AI will unlock additional capacity, especially to carry out tasks which can only be done by a person.

About this guide

This approach is not a catch-all policy for AI use at Tendring District Council (TDC); with the speed of technical innovation in the field best practice and the abilities of AI will evolve quickly.

Rather, this approach is designed to steer staff and Members through the necessary ethical considerations required when using AI as it exists now, while a full policy can be carefully considered and adopted. It focuses predominantly on the use of 'large language models' (LLMs), examples of which include CoPilot and Chat GPT that are increasingly becoming commonplace names (though many others exist) – this approach does not explore more specific or specialist uses of AI which may be available to various services across the council. Likewise, this approach aims to assist with considered use of both conversational AI – things such as chatbots, producing human-like interaction – and generative AI, which creates content, such as a report or piece of text, as well as potential use of machine learning using large data sets.

Crucially this approach does not replace a range of existing council policies which all factor into AI use – but instead seeks to guide users through the relevant considerations. The introduction of an AI process should be approved through usual council procedures, including governance and financial considerations – not all AI models are free.

This document focuses on the deliberate and specific use of AI, but it is worth noting that AI is embedded into much of the day-to-day software that we use, from word processing, spell checking and graphic design, to communications and photo editing tools. This document does not address TDC's approach to identifying or receiving AI, such as checking job applicants' use. It should be noted that it can be very hard to detect (good) AI-generated content, despite well-publicised examples.

If you are considering using AI then please consult with the IT Service Desk – after having read this document – to discuss your proposals; this may also help with wider-use considerations and help to share best practice.

Benefits of AI

There are multiple applications of AI that can be beneficial for teams; the limit is your creativity (or even how you prompt AI's creativity). It can be used to draft all forms of written content – from reports, to social media posts – brainstorm ideas, write code, analyse data, automate routine tasks (such as taking minutes from meetings), spark creative ideas, improve productivity, and provide valuable insights.

It is important to remember there is more to AI than content generation, and AI is not the end product, but a shortcut or assistant along the way – which then gives us as humans more time to add extra value from our lived experience and expertise. Using AI for simple tasks allows humans to tackle more complex issues.

Departments are encouraged to consider how AI could resolve challenges within their services – and, using this approach as a basis, engage with relevant council teams to enable such solutions.

Keeping up-to-date

AI technology is continually evolving, and it is therefore crucial for those using it to stay current with these advancements to maximise the benefits and minimise the risks associated with AI use. This can be done in a number of ways, including:

- Reviewing AI guidelines and policies: Regularly review the latest guidelines from credible sources, such as the Local Government Association, the UK Government or the Information Commissioner's Office (ICO). In particular, reference [the Government Communication Service Generative AI Policy](#).
- Attending training: There are lots of opportunities to engage with workshops and webinars that can help you stay abreast of AI advancements. This might include training on how to use AI tools more effectively, updates on AI policy, or discussions on ethical AI use.
- Following AI news and trends: Keep up-to-date with the latest news and trends in AI. This could be through AI-focused websites, blogs, podcasts, or newsletters.
- Seek regular feedback and reflect: Use your experiences, and those of your colleagues, to continuously learn and improve how you use AI tools. Regular reflection and feedback sessions can help identify areas for improvement or new opportunities for AI use within your communications work.

As we move into the (not-too-distant) future, keeping up with AI's rapid development will be an integral part of how we all work. Staying informed will enable you to harness its power effectively, ethically, and safely.

In addition, this approach will be regularly reviewed to ensure it remains up-to-date. At any point, use of AI will be governed by relevant (and perhaps new) legislation and regulation – such legal frameworks will always take primacy over the guidance outlined in this document.

Ethical and Safe Use Considerations

Accountability

AI is a tool to be used by people; the person who uses it is ultimately accountable for the end product.

Therefore the onus rests on you to ensure that whatever outputs you use from AI are accurate, ethical and otherwise acceptable. You will be responsible for any failings.

Apply healthy scepticism to the AI outputs and remember to check, check and check again.

Consider the analogy of a driver using a sat nav; the driver is still responsible for not driving into a lake.

Information Governance

According to the UK National Cyber Security Centre, many LLMs do not automatically add information from queries to its model for others to query. This means that your chat history should not be visible or searchable by other users.

Whilst this provides some level of data protection, it is important to note that every time you enter information into an LLM, you are giving that data to the creator or host organisation. This means that you could be breaking the law or in breach of organisational policies by sharing personal or sensitive information.

Therefore you should be conscious about data privacy and security: be vigilant about the types of data shared. You must not input Person Identifiable Data (PID), Personal Confidential Data (PCD), sensitive personal data, or proprietary information. Keep in mind that data breaches can happen, and data transmitted could potentially be exposed.

You may also not have control over where any data is stored or processed – for example, are the AI provider's data centres based in the UK/EU and therefore covered by GDPR.

To help you in this, ensure your mandatory training on information governance best practices is up-to-date, and that you have read the organisational Information Governance Framework and Policy.

It should always be remembered that breaches of our Information Governance policy can have serious consequences for the organisation, both in terms of financial penalty, business impact and reputation.

Cyber Security

If using developed AI models – that is to say, those beyond basic ‘off the shelf’ or online platforms – or creating connections to internal IT systems, then consideration must be given to cyber security implications.

Consult with the IT Department before doing so, as is part of IT policy for any such software use.

Fairness and bias in AI

AI models may unintentionally generate biased or unfair content due to the data they were trained on. This is because they learn from a vast amount of internet text, inheriting both the beneficial knowledge and the biases present in this data.

Be mindful of this and critically review the output, ensuring it aligns with our principles of fairness and inclusivity. Do this yourself, but you can also ask the AI model to identify its own biases too.

Remember that an Equality Impact Assessment may be required for the task or service you are conducting; or that you may need to review an existing assessment if you are changing how you perform that service through the introduction of AI.

Avoiding misinformation

AI LLMs generate responses based on patterns it learned during training. They do not necessarily have access to real-time information or the ability to independently verify facts. Always cross-check important information from external, reliable sources.

A unique challenge with AI is the potential for “hallucinations” — instances where the model generates information that is not just biased, but outright incorrect or nonsensical. These “hallucinations” can occur if the AI misinterprets the input, or when it makes unfounded assumptions based on its training data.

To mitigate this, always verify the information generated, especially when it comes to statistical data, factual claims, or sensitive topics. If something seems off or incorrect, cross-verify with reliable sources or rephrase your input for a more accurate output.

You can also ask the AI to provide citations (its sources) for information, to make this cross-referencing easier and guard against hallucinations.

You should never use AI to create something which could mislead others, whether this is entirely AI generated, or using AI to alter or change existing content in a way that removes or replaces its original meaning or message.

Models to use

Using reputable AI models, particularly those which are part of wider TDC software or licence packages (for example, CoPilot which is a Microsoft product) can help to mitigate some of the risks outline above.

CoPilot is currently preferred by the Council as it is part of our wider IT licences; while Chat GPT is blocked as it does not conform to GDPR. These considerations should be applied to any AI model.

LLMs such as CoPilot can be good 'jack of all trades' AI models, but remember there may be specific models for specialist areas of work. If using them, ensure they meet the necessary legal and ethical requirements as set out here.

Copyright, Public Sector Information and Intellectual Property Rights

At the time of developing this approach the issue of copyright and LLMs is a complex one, with ongoing legal cases between LLM developers and those whose content was used to teach them.

Similarly it is not yet legally defined or tested as to who owns the copyright on content produced from LLMs – the AI creator/developer, or the person who inputs information and prompts.

Purely AI-generated content cannot be copyrighted.

As a public sector organisation much of whose work is in the public interest and available as Public Sector Information, this is unlikely to have an impact upon TDC use of AI LLMs – but it should be considered if you are looking to produce something which will be required to be held under TDC copyright, or when commissioning suppliers or contractors. When procuring work you may want to check the supplier's position on AI use.

It does not impact upon the established principle of copyright law that work carried out in the course of your employer is under the copyright of the employer, not you as an individual.

You should not use copyright-protected content as inputs into LLMs.

Decision-making

There are strict rules around automated decision-making (which may or may not involve AI) as set out in UK GDPR [explained by the Information Commissioners' Office](#).

Should any tasks implement automated decision-making – where there is no involvement by a human – then this will need carefully documenting, adding to a central register, and relevant privacy notices amended. People also have the right not to be subject to a decision based solely on automated processing; and so the legal effects of doing this could be significant to the council.

These rules must be adhered to if using AI in decision-making, along with consideration of the other ethical points within this approach – particularly around unconscious and data bias.

Transparency

Whether use of AI should be declared on the finished product is a matter of ongoing debate, and is not a black-and-white topic.

As a pragmatic approach, minor or background use of AI tools – such as spell-checking, basic photo enhancement, or brainstorming – should not require a declaration of AI use.

Where AI tools have been used as the substantive piece of output, with minimal human revision; or, for example, if used not just to edit or enhance an image or piece of content but to change its meaning and purpose; then a declaration that AI has been used may be appropriate.

The best solution is to acknowledge organisational use of AI – being transparent about this use, and the framework within which it is used – and have this as a general disclaimer (akin to a privacy notice or policy which describes general processes rather than necessarily specific uses of data).

TDC will publish this approach on its website and intranet as part of this transparency.

Environmental sustainability

Like many digital tools, AI is not carbon neutral as it relies on electricity – in the case of LLMs, a lot of data causing a heavy electrical usage. While not a specific consideration for day-to-day use, this should be balanced when considering strategic use. AI can, though, also be used to look at ways of reducing environmental impacts in other areas, so it has the potential to balance out its positive and negative impacts.

De-skilling

While there are many potential benefits to using AI, it is important that you do not become over-reliant on it, and lose the core skills and expertise – the human touch – where you can add value to and evaluate the outputs.

Accessibility and web content

Ensure that any content you produce meets best practice and regulatory requirements for accessibility – especially around Plain English. You can ask the AI to do this.

Be aware that for web content, search engine optimisation can penalise poor quality content – so ensure you are using AI effectively.

When not to use AI

There are certain scenarios and topics where the use of AI should be approached with extreme caution or avoided altogether. Given that AI lacks human judgement and sensitivity, it is not well-suited for delicate contexts where tact, empathy, and deep understanding are essential. Here are some examples:

- **Highly sensitive topics:** Matters such as bereavement, mental health, or other personal struggles necessitate a human touch. In such cases, it is advisable to rely on human judgement and compassion, rather than using AI-generated content that could unintentionally come across as cold or insensitive. Audiences may also feel angry or upset if they discover that content addressing sensitive matters was generated by an AI.
- **Cultural sensitivities:** When dealing with content that involves cultural contexts or sensitivities, it is best to have individuals who are well-versed in those cultures craft the communication. AI models might not have the depth of understanding required to address cultural nuances, and there is a risk of generating content that is unintentionally offensive or inappropriate.
- **Communications involving children or vulnerable populations:** When creating content that is aimed at or involves children, the elderly, or vulnerable populations, it is important to exercise an extra layer of caution and thoughtfulness. Human empathy and understanding of the nuances involved in communicating with these groups are essential.
- **Imagery and video:** At the time of development of this approach, AI tools to create images and video from scratch are extremely specialised, and should only be used by someone trained in the field. However, simpler AI tools can be used to edit and enhance existing material (provided all other considerations are met); while AI can also be useful in creating conceptual or storyboard content. As above, though, AI should not be used to alter existing content in a way that removes or changes its context or message. Using AI to replicate someone's likeness must be extremely carefully considered, and only done with the written (and recorded) consent of the individual(s) being replicated.
- **Specialised work:** Councils often deal with highly specialised and technical areas of expertise. While AI can be a useful research, drafting and editing tool, it is not a replacement for the knowledge of subject matter experts.

This list is not exhaustive. If in doubt about whether it is appropriate to use an AI tool, please discuss this with your line manager, IT Service Desk, or a senior communications lead.

While AI can be an invaluable tool for streamlining and enhancing communications, it is important to recognise its limitations. Rely on human expertise and judgement in situations that require sensitivity, precise legal language, cultural competence, or careful handling of crises.

Consider the reputational damage to TDC of poor, misguided or unethical use of AI.

Best practice use of AI

Collaboration and best practice

As AI is an evolving field, best practice will change and develop. TDC is keen to share learning and best-use cases that make AI useful, effective and efficient.

In order to do this, please share your successes with the IT Service Desk, who are researching and curating this information into a resource which can be shared with TDC staff.

Information

While trained on substantial amounts of data, AI is only as good as the information you put into it.

It is always useful, therefore, to provide AI with relevant information, such as: examples of similar work or templates you want it to use; previous reports, press releases or communications on a topic; a style guide.

This list is not exhaustive. While various AI model interfaces are different, most will give you the option to upload documents (considering the points above about information governance) – often useful for this background information – as well as text for your prompt inputs. It can also consider web information if you provide a web link.

Effective prompts

As well as providing enough information, being clear about what you are asking AI to do is critical – this is where effective prompts (the instructions you give AI) are vital. Some top tips are:

- **Define the AI's role:** Specifying the role in your prompt can help get the kind of response you want. For instance, you might want it to act as a brainstorming partner, an editor, a summariser, or data analyst. Start your prompts with this, such as “you are a public service communications professional”.
- **Be specific:** If you want a particular type of response, make sure to specify it in your prompt. For example, if you want a brief answer, you could start your query with “in one sentence...”, or “in less than 50 words”. If you want a detailed answer, ask for “a detailed explanation” or “a step-by-step guide”. If you don't have a pre-defined language setting, specify you want it in UK English (or in a later prompt, ask it to re-write it in UK English).
- **Provide context:** Providing context can help the model give a more relevant response. Instead of asking it to “write an article about the launch of a new service”, provide some background information about who the article is for and key information which is needed to generate an informative and engaging piece. For example, are you drafting something for other staff, a report to Management Team, or an email to Members?

- **Give an example:** AI often responds well to examples. For instance, if you're seeking to draft an announcement about a new sustainability initiative for local councils, you could prompt the AI with: "Write an announcement similar to the one below, but this time focusing on our new recycling programme for community groups." This provides the AI with a clear reference point and helps it understand the specific task you're asking it to perform.
- **Set constraints or limitations:** When you are dealing with broad topics, it is helpful to provide the AI with specific parameters. For example, if you're summarising meeting notes you could ask the model to "list five key points" or "capture specific actions given". This provides a specific focus for the response.
- **Guide the format:** If you need the response in a particular format, specify it in your prompt. For example, you could ask it to "outline a report using the attached Cabinet report template for promoting our local council's sustainability initiatives", or specify that you want responses in bullet points, a table (with specified columns), or something else. This not only specifies the topic, but also the structure of the response, ensuring alignment with recognised objectives.
- **Set the tone:** Tone can be crucial, especially when addressing sensitive topics. You can direct the AI to adopt a specific tone in its responses. For instance, if encouraging residents to take a certain action, say in your prompt "the tone should be encouraging and informative, emphasising the importance of community involvement and action".

Building on the above, here are eight elements every effective prompt should have. (Please note – for smaller tasks, you may not need each element, but do consider it.)

1. **Persona:** the role in responding: e.g. "You are an IT professional"
2. **Context:** what is the situation or topic?
3. **Task:** the main instruction. Split up complex tasks
4. **Exemplar:** provide an example of what you are looking to achieve
5. **Format:** specify the presentation or layout format you want. Do you want a table (if so, what columns), bullet points, image. (This may be provided within your exemplar.)
6. **Tone:** emotion, what impact do you want it to have?
7. **Framework:** what structure do you need, is there an industry-standard set of headings for example. (This may be provided within your exemplar.)
8. **Style:** are there any brand guidelines, style guide, tone of voice settings or instructions?

Iterative process

AI, in particular LLMs, are best used in an iterative process – that is to say, over a number of steps. AI models remember the context of the conversation to assist with this process.

Do not simply input one prompt, with multiple asks, and expect a solid answer. Break up the work into different prompts; ask it to use previous answers to then guide the next steps.

Ask AI to refine answers – it is unlikely the first response it gives will be the best one it can provide. If your first prompt does not give the required outcome, how can you vary it, expand on it. It is often handy to ask AI to number the paragraphs it gives you, so can then ask them to juggle around the order or refine specific sections.

Be strategic in your use of AI – ask AI what would be the best prompt to achieve a task, and what the best process would be to accomplish this iterative process.

Think about these top tips for the iterative process:

- **Link it back to the local:** When using AI-generated content, it is essential to customise the draft content with the Tendring/local context in mind. This could include referencing local services, organisations, or initiatives. This can also include statistics, trends, or news relevant to the area. Integrating these elements will make outputs more relatable and aligned to local strategies.
- **Rephrase and refine your prompt:** If the initial response does not meet your needs, consider altering your prompt. A slight change in phrasing, adding specific details, or asking the question from a different angle can lead to a more useful answer. For example, if the response to “what should be included in a tenancy management strategy?” isn’t sufficient, you could ask, “what are the top five key messages for a strategy about tenancy management?”
- **Ask for revisions:** The AI can refine its responses based on your feedback. If the initial response is not what you wanted, you can ask for changes. This might include adjusting the length, altering the tone, adding, or removing specific points, or shifting the emphasis. For example, if a press release draft is too long, you could ask, “can you condense this press release to fit one page without losing key points?”
- **Seek clarification:** If a response is unclear or too technical, you can ask the AI for clarification or a simpler explanation. For instance, if a description of a new health policy is laden with jargon, you might request, “can you explain this policy in plain English for the public?”
- **Verify information:** Cross-check important information from external, reliable sources, especially when the AI generates data, factual claims, or sensitive information.
- **Be sure to edit:** AI is a fantastic tool, but it does not replace the professional experience and skills you have. Responses should always be reviewed and refined by human judgment to ensure they meet plain English standards and effectively communicate your intended message.

Knowing when to use AI

AI can be a valuable tool for a variety of tasks across TDC. Three broad themes are:

- Getting you started (overcoming a blank page, or research)
- Refining or repurposing existing content (spell checking, re-writing)
- Automating a process or removing drudge tasks (coding, summarising meetings, making multiple versions of content)

The below list is not exhaustive, but hopefully it gives you some ideas to get started.

- **Copy editing:** While AI is not a professional copy editor, it can help in spotting glaring typos, grammatical errors, or awkward phrasing in written content. This can help streamline the editing process, though a human eye will still be necessary for nuanced editing work. A simple prompt of “edit and improve this:” could save you hours. NB – this cannot replace a human proof-read.
- **Devising strategies:** AI can be a helpful idea-generating partner when devising strategies or plans, especially using industry-recognised frameworks. You might ask it to “Outline a communication strategy for promoting a new mental health initiative,” providing a starting point that can then be refined and detailed.
- **Drafting communications:** AI can assist in drafting initial versions of messaging, saving time, and facilitating efficient communication. Remember, these drafts should always be reviewed by the appropriate expert and undergo the standard sign-off processes. It can also help if you require multiple versions of essentially the same messaging but for targeted audiences.
- **Creating content for social media:** AI can aid in generating content for social media, from posts raising awareness about service issues to those promoting behaviour change. Please note, the style and format of suggested copy may not be in line with today’s standards due to the pace at which social media best practice changes. Giving examples of successful posts can help it to generate effective copy for you, and use the Communications team to support with best practice on this.
- **Summarising text:** AI can also summarise information – whether that is to provide a summary of a report, a precis to include within newsletters and digests, or generating meeting summaries or minutes from a transcript or notes. Try using the prompt “summarise the below in less than 50 words to use as an executive summary”.
- **Brainstorming ideas:** AI can be a great way of getting over ‘writers’ block’, or as a research tool (if citations checked). It can also offer fresh perspectives on problems, such as identifying risks you may not have considered and support research (with properly cited sources).
- **Writing code:** AI can help with writing code for web development, managing Excel spreadsheets, and other software too.
- **Creating large amounts of variations:** Perhaps you have a presentation for numerous speakers, and want a script for each person with their lines in bold; or need multiple prompts based around a single scenario. Use AI to work through each version.
- **Supporting accessibility:** Such as generating subtitles for video or translations (subject to the proper checks).
- **Generating reports:** Creating reports based on templates, and previously provided information; for example, converting a Management Team Briefing Note into a Cabinet report, or a summary email into a briefing.
- **Identifying fraud:** There is the potential to use machine learning AI to identify patterns in data sets and flag potential fraud.

AI is not just about text. Think other media too, such as audio. Using in-built tools to get a transcript from an audio recording can then be used to provide summaries of

meetings, or provide concept pictures (noting the considerations around best-practice for AI image use).

Glossary

- **Large Language Model (LLM):** An artificial intelligence model trained on a vast amount of text data to understand and generate human-like text.
- **Prompt:** The words provided by the user to the AI, specifying the desired response or task.
- **Response:** The output generated by the AI in reaction to a prompt.
- **Interface:** The graphical or textual environment where users interact with an AI LLM, including input fields, response areas, and chat history.
- **Misinformation:** False or misleading information generated by AI, often unintentionally, due to the limitations of its training data.
- **Hallucinations:** Instances where AI generates information that is incorrect, nonsensical, or biased, often due to misinterpretation of input or unfounded assumptions based on training data.
- **Rephrase and Refine:** Altering the prompt given or adjusting its parameters to elicit a more useful or relevant response.
- **Ask for Revisions:** Requesting AI to refine its responses based on feedback, including adjusting length, tone, or specific points.
- **Seek Clarification:** Requesting AI to provide simpler explanations or clarification on unclear or technical responses.
- **Verify Information:** Cross-checking information provided by AI with external, reliable sources to ensure accuracy.
- **Feedback and Reflection:** Using experiences and feedback from colleagues to continuously improve AI usage, identifying areas for enhancement or new opportunities for AI integration.
- **Conversational AI:** AI which produces human-like interaction, such as chatbots
- **Generative AI:** AI which creates content, such as a report or piece of text.
- **Machine learning:** a type of AI that allows computers to learn and improve from data without being explicitly programmed. It uses algorithms to analyse large amounts of data, identify patterns, and make predictions.

Acknowledgements

This approach has, with permission, drawn heavily upon work done by the Essex Communications Group (of which TDC is a member), which itself was based upon work developed by the Mid and South Essex Integrated Care Board. The original document was developed in collaboration with AI.