Developing an Organization's AI Policy

Published February 2024

This guide has been developed by the Healthcare Communications Association's (HCA) Standards & Best Practice and Foresight Committees. It aims to provide support to members in developing their organization's AI policy, or as a review checklist for those who already have a policy in place. Due to the international make-up of the HCA's membership, this guide is written generically to be applicable across geographies. Consideration is therefore needed around specific local regulations, laws or frameworks that may be relevant in the areas outlined for inclusion in an organization's AI policy.

Developing an Organization's AI Policy

In October 2023, the Healthcare Communications Association published *The AI Roadmap – charting our future in healthcare communications (see Matt Lewis & Elizabeth Mercer, Current Medical Research and Opinion, 40:1, 5-7, DOI: 10.1080/03007995.2023.2261737)* The main tenets of that publication were that generative AI has a huge opportunity to transform our work at multiple levels; that we need to embrace AI now, as waiting is not a viable strategy; and that a clear AI policy is essential to provide appropriate guardrails to ensure responsible and ethical use within organizations.

As a robust AI policy is critical to help guide an organization's journey, this guide aims to complement *The AI Roadmap*, providing further insights for those developing an AI policy for their organization. It has been developed having reviewed existing policies which have been established in our sector; policy guides from across related business sectors; and finally by asking AI models themselves.

In producing an AI policy, all relevant stakeholders within an organization need to be consulted for approval. It is also important to ensure that any policy will help fulfil any requirements set by those to whom you supply goods and services. This guide provides the template structure, and suggested content to incorporate for an initial draft and should not usurp the internal approval process. Relevant stakeholders may include, but not be limited to: IT / AI leads, Legal, Regulatory, Board / Leadership Team/ Department heads, Financial Director, CEO

The organization needs to undertake the comprehensive communication and training required to ensure all employees and relevant stakeholders have access to, understand and commit to the requirements of the policy.

An organization may wish to ask employees to sign a copy of the policy to confirm understanding of the requirements it puts on them.

Artificial Intelligence (AI) is a broad term that encompasses a range of techniques and technologies that enable computers to perform tasks that would previously have required human intelligence for example, speech recognition, decision-making and problem-solving.

Generative AI is a specific subset of AI that centers on generating new content that has not previously been seen or created. Generative AI uses algorithms that learn patterns from available data and then generates new content, offering creative variations. It is the more widespread accessibility of Generative AI that has catalyzed the importance of organizations having policies for the use of AI, due to the additional risks this involves in various areas.

For the purpose of this guide the generic term of AI is used, however, a policy may want to be more specific in the terminology if this is relevant for the required understanding.

Policy Structure and Content Considerations

As a minimum, a robust AI policy needs to consider, and document, the following:

Purpose of the policy

This should explain why an AI policy is essential to ensure the guardrails are in place for the responsible and ethical use of AI within the organization. It could include the following:

- Recognition AI has a growing importance for the work the organization undertakes and how it can provide a step-change in how we are able to deliver innovative communications and their impact

- Highlight there are potential ethical concerns e.g. threat of redundancy and practical risks e.g. breach of confidentiality that need to be mitigated, which the policy aims to address

- Emphasize that AI is a tool to support our work and not a replacement for human expertise or decision-making

- Position the policy as the mandatory rules under which all use of AI needs to be benchmarked (this could also include the need for all employees to formally sign to confirm they are familiar and comfortable with the policy and its requirements)

- Recognize that alongside this policy, any relevant clauses in third-party service agreements or similar, need to also be considered

- Emphasize that this is a rapidly changing environment and employees need to keep up-to-date with changes to the policy and the implications that this may have on their use of AI

Which AI tools can be used and the process for gaining approval to add new tools to those authorized

Only the use of approved AI tools should be allowed.

Details on the process of evaluation and who undertakes the evaluation should be included as well as how a request for an evaluation can be made.

A full list of those AI tools approved for use should be incorporated or linked to. This should be regularly updated.

Similarly, any AI tools specifically not approved should be documented in the same way.

Ethical and legal considerations

- **Contractual**
  Employees must not upload or share on AI any confidential (client or company), proprietary data or content protected by regulation, unless explicit written approval has been obtained from the data owner. Data on customers, employees, employers or business partners would be included in this prohibited use, as would survey or clinical data.

  If the organization has any closed AI tools (i.e. proprietary and limited in its accessibility and usage) where this would be permissible, these should be listed as specific exceptions to this rule.

- **Privacy**
  The relevant data protection laws followed by the organization globally or relevant for the geography where the organization is based, remain applicable for AI. These should be highlighted in the policy emphasizing that uploading or sharing personal data on AI without permission, would be a breach of those data protection laws.

  If the organization has any closed AI tools where this would be permissible, these should be listed as specific exceptions to this rule.

In addition to both of the above policy statements, employees should also always consider when uploading or sharing information with AI, if they would be comfortable with that information being available in the public domain. Any uncertainty and the information should not be shared.

- **Intellectual property (IP)**
  This is a complex area when considering the use of public generative AI. IP includes areas such as copyright, patents, trademarks and trade secrets. It is important to ensure that the organization's or client's IP is protected and that we do not infringe others IP. There is the risk of the latter when creating content or images from public generative AI tools. These models use content and data to train the tools, some of which may have IP rights. It is important to be certain no IP rights have been infringed if you are going to commercially use something generated using these AI tools.

- **Access**
  Employees should not give access to closed AI tools to anybody outside of the organization or to others within the organization who do not have prior approval for access.

  The sharing of login details or other sensitive information should be prohibited. Details should also be added here outlining how access for an employee or external third party is applied for and granted.

- **Transparency**

  Ethically we should always be open about our use of AI. Embracing AI in a safe and responsible way should be seen as a real positive for our communications activity. AI presents immense possibilities, and we should want to maximize the potential it can offer.

  We should commit to declare our use of AI where it has had a significant role in the final communications output. Alongside this we should outline the processes we undertake as outlined in this policy, as to how we ensure our use is responsible and ethical and any outputs are appropriate and accurate.

Human oversight / Mitigating risks around accuracy and bias

All AI output needs to be fully reviewed using traditional human assessment and verified and edited as required, to ensure it is ethical, accurate, comprehensive and that there is no inherent bias or discrimination. It should also be approved as compliant for all industry codes, regulatory or legal frameworks that are relevant. This process should also ensure all the other sections of this policy have been applied appropriately and that this is evidenced.

All references generated through AI should be similarly assessed and checked to confirm their applicability and accuracy of citation.

This section may require the addition of details for an organization's formal approval process, should this be in place.

Diversity, equity and inclusion

Diversity, equity and inclusion should be the foundation of all communications work. It is important to ensure the use of AI does not exclude or underrepresent any individuals or communities.

Conversely, AI may also be used to support these aims by helping remove human bias.

In all the uses of AI as an organization, the question needs to be asked, are we being true to our DE&I aspirations?
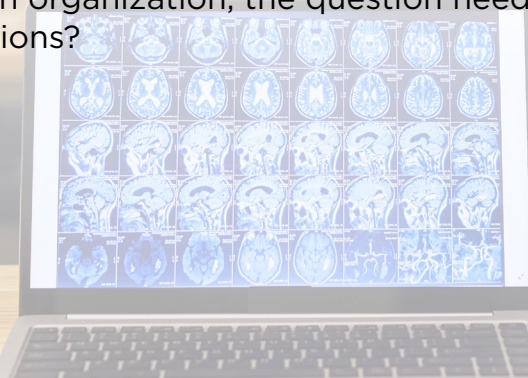
Employee training and support

Like any new technology or development, organizations need to be committed to the appropriate training and support to facilitate appropriate and ethical use. This will include employee understanding of their responsibilities surrounding the use of AI and specific practical training, as required, on the use of the AI tools available.

The mandated skills and knowledge training requirements of employees before using AI should be outlined here. As a minimum, full knowledge and understanding of the organization's AI policy should be required.

Suppliers' compliance

Organizations should ask at least the same standards of their suppliers, with regards the use of AI, as they practice themselves.

Formally request for all suppliers to agree to adhere to the organization's AI policy.

Consider if there are training requirements for suppliers, either by attending the organization's programme, or allowing their own equivalents.

Ongoing improvement and review process

In such a rapidly moving environment it is essential an AI policy remains current and effective. The organization should commit to monitor and evaluate AI use and review, update and communicate the AI policy regularly (suggest a maximum allowed period between reviews is documented, perhaps every 3 – 6 months).

The associated list of approved and unapproved AI tools should be a live list and not need to wait for a full policy review.

Legal and regulatory

The use of AI within the organization should always comply with the appropriate codes, regulatory and legal requirements applicable, irrespective of any requirements within the AI policy. Should a perceived conflict exist, this should be notified immediately to the Policy Officer.

Policy Officer

This should outline and provide full 24-hour contact details for the individual(s) who are responsible within the organization for the implementation and update of the policy, and as the point of contact for questions or to report suspected breaches.