2026

Best Practice Guide for the Responsible Use of Generative AI in Advertising

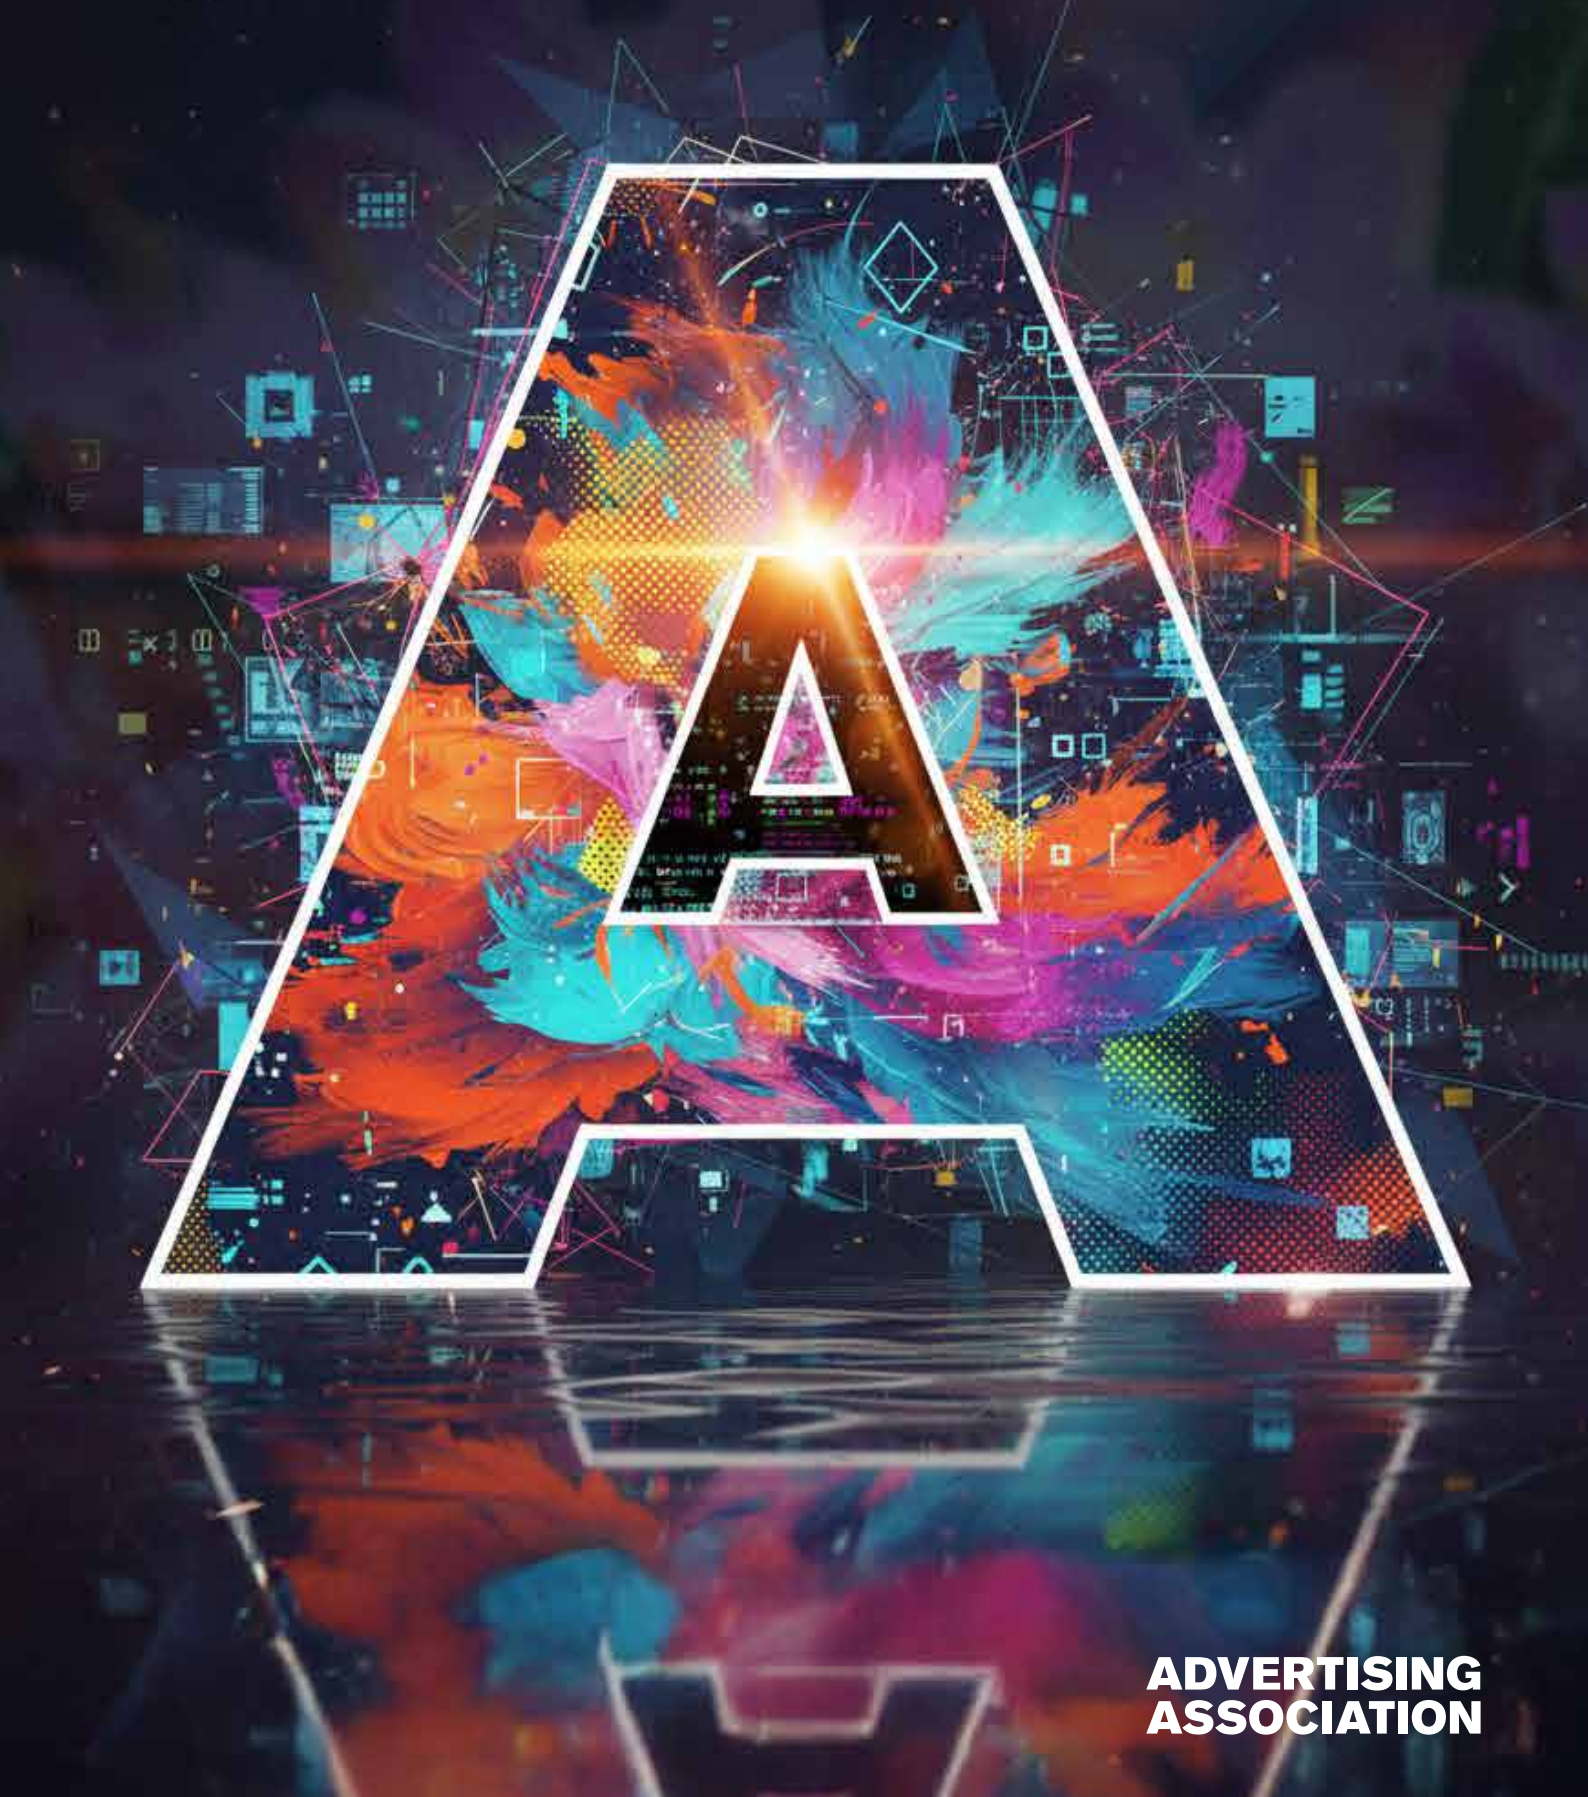

ADVERTISING ASSOCIATION

Foreword

Generative AI is transforming the UK's advertising landscape with remarkable speed. From enhancing creative development to streamlining campaign production, practitioners across the UK are adopting this technology to improve efficiency and effectiveness. As this transformation accelerates, advertisers and marketers are rightly focused on ensuring responsible implementation that maintains public trust.

That's why, as Deputy Chair of the Online Advertising Taskforce, I am delighted to welcome this Best Practice Guide for the Responsible Use of Generative AI in Advertising which addresses that need directly. Developed under the auspices of the Online Advertising Taskforce, it translates eight core principles for responsible generative AI use into actionable guidance for industry practitioners.

The collaborative development process, between industry, the ASA, and Government has ensured these recommendations align fully with both statutory requirements and the ASA's self-regulatory framework. This alignment provides clarity for advertisers, agencies and media whilst maintaining the flexibility needed for continued innovation.

By prioritising transparency, fairness, and robust governance, this accessible guide supports the UK advertising industry's commitment to responsible technology adoption. It provides practical tools to help practitioners navigate the opportunities and challenges of generative AI whilst upholding the standards that underpin consumer trust.

The principles and guidance outlined here represent the collective expertise of industry leaders committed to ensuring generative AI enhances rather than undermines advertising's role in supporting consumer choice and market competition.

Mark Lund OBE, Deputy Chair of the Online Advertising Taskforce

Executive Summary

AI is reshaping the advertising industry, and advertising practitioners are already using Generative AI (GenAI) to transform every facet of the process – from creative ideation to production. While offering significant benefits, GenAI can potentially introduce risks in areas such as deception, bias, and privacy erosion. Growing use of AI-generated content could impact the industry's long-term standing with consumers, with notable implications for trust in advertising.

That is why the Advertising Association's members: advertisers, agencies, and media, under the auspices of the joint Government-Industry Online Advertising Taskforce, have come together to develop this voluntary Best Practice Guide, building on and operationalising the IPA and ISBA principles for ethical AI use in advertising. We believe that a coherent, industry-led approach, which advocates consistent, high-standard practices, is the most effective means to pre-empt potential harms before they may occur. For those newer to AI technologies, we hope this guide sets out the key principles and considerations to build your confidence in appropriate use of AI tools.

The Guide is designed to enhance public trust and confidence in AI and ensure that responsible AI has a central role in growth and innovation, particularly in the advertising and marketing industry. It also tackles these challenges facing advertising practitioners by providing practical advice and recommendations which complement UK laws, like UK GDPR and the Equality Act 2010, as well as the UK's self-regulatory framework, such as CAP/BCAP codes and ASA standards, to foster ethical innovation while upholding integrity.

The Guide's principal aims are to deliver:

- Practical, high-level guidance for deploying GenAI responsibly in advertising
- A framework for implementation
- A framework for governance
- A document for shared standards and definitions

By following the Guide, practitioners can harness GenAI's creative and operational potential whilst upholding the ethical standards that maintain consumer trust and support sustainable industry growth.

The Guide's recommendations are anchored in eight principles which cover the following areas:

- Transparency
- Responsible Data Use
- Bias and Fairness
- Human Oversight and Accountability
- Societal Wellbeing
- Brand Safety and Reputation
- Environmental Stewardship
- Continuous Monitoring

Introduction

This voluntary Best Practice Guide ('Guide') supports UK advertising industry practitioners in using generative artificial intelligence ('GenAI') responsibly. It provides guidance on transparency, consumer protection, and maintaining advertising integrity whilst enabling innovation and growth.

GenAI is increasingly central to advertising – from ideation to content creation and personalisation. These systems include large language models, image generators, voice synthesis, video creation tools, and hybrid systems. They can create text, images, audio, video, or other media based on prompts, training data, or user inputs. Given this complexity, this Guide helps practitioners navigate these changes by distilling technical issues into accessible language. It does not create new rules or replace existing legal frameworks. Instead, it offers a common reference point for businesses making ethical decisions about AI deployment, complementing existing regulation with actionable best practices.

1. Objectives and Scope

1.1 Purpose

This Guide aims to:

- Establish high-level, voluntary best practice guidelines for GenAI use in online advertising
- Provide practical guidance for identifying and managing GenAI-related risks
- Support effective consumer protection and prevent misleading or harmful AI-generated content
- Create shared standards and common language across the advertising ecosystem
- Promote public trust in the advertising industry
- Enable responsible innovation and competitive advantage
- Complement existing legal, regulatory, and self-regulatory frameworks

1.2 Scope and Applicability

This Guide is designed to be relevant to the following organisations:

- Advertisers and brands
- Advertising agencies
- Media owners
- Technology providers and platforms serving the advertising industry

Whilst developed for the UK market, the principles are sufficiently flexible to accommodate international interpretations and applications.

To ensure consistency with existing obligations, the Guide operates within established legal and regulatory frameworks whilst providing practical interpretation of their application to GenAI contexts. These frameworks include:

UK Legislation:

- Data Protection Act 2018 and UK GDPR
- Digital Markets, Competition and Consumers Act (2024)
- Copyright, Designs and Patents Act (1988)

- Equality Act 2010

Self-Regulatory Standards:

- CAP/BCAP Codes and ASA standards (UK)
- EASA, ICAS, and ICC standards (international)

Practical Guidance:

- ICO AI and Data Protection Risk Toolkit
- ASA/CAP guidance on GenAI & advertising regulation
- ASA/CAP guidance on disclosure of AI in advertising

By working alongside these frameworks, the Guide provides a common reference point for businesses navigating the evolving AI landscape.

2. Governance Framework

2.1 General Guidance

Organisations should establish GenAI governance frameworks that prioritise ethical use, transparency, and accountability. When developing these frameworks, organisations should consider aligning with the UK Government's AI regulation principles, adapting them as appropriate to the advertising context.

Effective governance frameworks should include:

- Designated accountability for GenAI systems and outputs
- Clear processes for approving GenAI deployment
- Mechanisms for ongoing monitoring and risk management
- Alignment with organisational values and legal obligations

Responsibility for AI governance should be assigned to appropriate personnel or committees with appropriate technical understanding, decision-making authority, and accountability for managing AI-related risks.

2.2 Risk Assessment Guidance

All GenAI applications should undergo risk assessment before deployment. Risk levels should be determined by evaluating potential consumer harm, not solely by technical characteristics.

Recommended risk categories:

High Risk	Limited Risk	Minimal Risk
Requires senior executive approval and multiple review stages.	Requires managerial approval and expert review	Requires automated screening and periodic audits
Examples: ▪ *AI-generated content in regulated sectors (financial services, healthcare, gambling)* ▪ *AI-generated content targeting vulnerable populations* ▪ *AI-generated content which makes claims about health, safety, or product efficacy*	*Examples:* ▪ *Authorised and disclosed use of real people's likenesses*	*Examples:* ▪ *AI-generated creative content for general advertising (without likenesses)*

2.3 Ethics by Design

It is recommended that organisations developing or deploying GenAI systems should embed ethical considerations into system architecture from the design stage, rather than addressing ethics as an afterthought. This includes:

- Conducting ethical impact assessments before development begins
- Involving diverse stakeholders in design decisions
- Building transparency and explainability features into system architecture
- Implementing technical safeguards against bias and discrimination
- Designing audit and monitoring capabilities from the outset

For organisations using third-party GenAI tools, ethics by design means conducting due diligence on vendors' ethical safeguards, selecting systems that align with organisational values, and configuring tools to maintain ethical standards in their specific use context.

2.4 Training and Awareness

Effective AI governance recommends that personnel involved in AI deployment understand relevant ethical principles, legal obligations, and organisational policies. Training programmes should be proportionate to organisational size and employee roles:

For organisations with capacity for formal training:	**For smaller organisations:**
For example: ▪ Employee onboarding covering this Guide's principles and relevant legal frameworks (UK GDPR, Equality Act 2010, CAP/BCAP Codes) ▪ Develop role-specific training (e.g., creative teams on transparency and bias; data teams on privacy; management on risk assessment) ▪ Provide regular updates when regulations, guidance, or organisational policies change ▪ Provide access to resources such as ICO guidance, ASA rulings, and case studies	For example: ▪ Advise all staff using AI to read this Guide ▪ Maintain accessible resources (e.g., internal wiki, shared documents) with key requirements ▪ Direct staff at the ICO and ASA websites to access guidance resources. ▪ Designate a responsible person to stay informed about regulatory changes and brief the team ▪ Encourage participation in industry forums or working groups

Training effectiveness should be assessed through:

- Staff understanding of when to escalate AI-related concerns
- Compliance with internal approval and review processes
- Reduced incidents of non-compliant or ethically problematic AI outputs

2.5 Reporting

Organisations following this Guide should consider how they will demonstrate adherence to its principles. This may include internal audits, transparency reporting, or participation in industry forums.

To support ongoing development and refinement of the Guide, stakeholders are encouraged to:

- Share learnings and challenges in implementing these principles
- Participate in industry working groups or forums dedicated to responsible AI in advertising
- Contribute feedback on gaps, ambiguities, or areas requiring clarification

This Guide will be reviewed and updated regularly to reflect:

- Technological developments in AI systems and capabilities
- Evolving regulatory frameworks and legal precedents
- Emerging best practices and industry standards
- Feedback from stakeholders on practical implementation challenges

Review cycles should occur every 18-24 months, with ad-hoc updates as significant developments warrant. Updates should be developed by the Advertising Association in consultation with industry stakeholders, regulators, and consumer representatives.

3. Core Principles

Principle 1: Ensuring Transparency

Disclosure of AI-generated or AI-altered advertising content should be determined using a risk-based approach that prioritises prevention of consumer harm.

Transparency enables consumers to make informed choices and builds trust in individual brands and the advertising ecosystem. When consumers understand how content is created and why they are seeing particular advertisements, they can evaluate claims critically and exercise meaningful choice.

The decision to disclose GenAI use should be proportionate to the potential for consumer harm or misinterpretation. Content that is clearly deceptive or misleading should never be used, whether GenAI-use is disclosed or not. Content that could potentially confuse consumers about facts, product capabilities, endorsements, or the reality of depicted events would benefit from clear disclosure. Content that is obviously fictional, fantastical, or impossible does not typically require AI-specific labelling, though standard advertising disclosure rules would still apply.

When assessing disclosure requirements, organisations should consider whether:

- *A reasonable consumer might be confused about the nature or origin of the content*
- *The content depicts real people using AI-generated representations without clearly indicating their level of involvement.*
- *The content uses AI to synthesise or alter the voice of real individuals, particularly public figures, without clear indication of the artificial nature of the audio.*
- *The content presents AI-generated customer testimonials, reviews, or user-generated content that could be mistaken for genuine consumer experiences.*
- *The content features AI-generated spokespersons, personalities or brand ambassadors that could be perceived as real people rather than digital creations.*
- *AI-use affects the potential interpretation of claims about products or services. If it might materially affect a transactional decision, then it would be likely prohibited as misleading.*
- *The context and media in which the advertising content appears.*
- *The content targets vulnerable populations.*

For detailed guidance on implementing risk-based disclosure frameworks, organisations may refer to the report "Beyond Simple Labelling: A Framework for Labelling AI-Generated Content in Advertising" (ICAS, 2024).

Principle 2: Ensuring Responsible Use of Data

Personal data used for GenAI applications—including model training, algorithmic targeting, and personalisation—should comply with data protection law and respect individuals' privacy rights.

GenAI systems often rely on large volumes of personal data to function effectively. Responsible use ensures that individuals' privacy rights are respected and consumer trust is maintained, and regulatory penalties are prevented. Consumers are also increasingly aware of–and

concerned about – how their data is used. Responsible practices – such as transparency, consent, and clear opt-out mechanisms – help maintain public trust in digital advertising and prevent backlash against brands or platforms.

Organisations should:

- *Comply with UK GDPR and Data Protection Act 2018 requirements*
- *Obtain valid consent or establish another relevant lawful basis for processing personal data in GenAI systems*
- *Provide clear information about how personal data is used in training GenAI models*
- *Respect individuals' rights to access, rectification, erasure, and objection*
- *Apply data minimisation and purpose limitation principles*

Beyond personal data protection, organisations should respect intellectual property rights, including:

- *Avoiding the unauthorised use of copyrighted material in training data*
- *Ensuring AI-generated content does not infringe existing copyrights or trademarks*
- *Obtaining appropriate licences or permissions for commercial use of third-party content*

When using third-party AI models or platforms, organisations should verify that:

- *The provider's data handling practices align with UK data protection standards*
- *Terms and conditions clearly specify data usage, retention, and rights*
- *Training data sources are appropriately licensed*
- *Users' consent extends to third-party processing where applicable*

Organisations should consult ICO guidance, including the AI and Data Protection Risk Toolkit, for detailed implementation advice. They should also monitor developments in AI-specific regulation that may introduce additional data protection requirements.

Principle 3: Preventing Bias and Ensuring Fairness

GenAI systems should be designed, deployed, and monitored to prevent discrimination and ensure fair treatment of all individuals and groups.

AI systems can unintentionally replicate or amplify existing societal biases related to race, gender, age, disability, socioeconomic status, or other protected characteristics. In advertising, this manifests as unfair targeting, systematic exclusion of certain groups, or reinforcement of harmful stereotypes. Beyond being unethical, discriminatory AI practices may breach the Equality Act 2010 and undermine consumer trust.

Preventing bias requires proactive measures throughout the AI lifecycle—from initial design through ongoing deployment. Fair AI systems benefit everyone, align with legal obligations, and build long-term brand credibility.

Organisations are advised to:

- Assess AI systems for potential discriminatory outcomes before deployment
- Use diverse and representative training datasets
- Test AI systems for bias across different demographic groups

- Monitor deployed systems for discriminatory patterns in targeting or content generation
- Implement corrective measures when bias is detected
- Document bias mitigation efforts and outcomes
- Establish clear governance processes for responding to bias complaints or concerns from consumers or stakeholders

Fairness should encompass:

- **Impartiality**: AI decisions should not unfairly favour or disadvantage protected groups.
- **Inclusion**: Advertising should reach diverse audiences appropriately, not systematically exclude groups.
- **Representation**: AI-generated content should avoid stereotypical or harmful portrayals.
- **Equal opportunity**: All consumers should have fair access to relevant advertising, offers, and information, whilst respecting legitimate targeting criteria and consumer preferences.

Organisations should consult ICO guidance on fairness, bias and discrimination in AI systems, and ensure compliance with the Equality Act 2010.

Principle 4 – Ensuring Human Oversight and Accountability

AI-generated advertising content should be subject to appropriate human oversight before publication, with the level of oversight proportionate to potential consumer harm.

Whilst AI systems can produce efficient and scalable content, they can also generate misleading, inappropriate, or harmful material that damages brand reputation or deceives consumers. Human oversight ensures AI outputs align with brand values, legal standards, and ethical requirements including CAP/BCAP Codes and ASA rulings.

Effective oversight maintains the balance between AI's efficiency benefits and the need for accountability, particularly as consumers and regulators increasingly scrutinise AI-generated advertising.

Human oversight mechanisms should be proportionate to risk levels (as defined in Section 2.2):

High-risk applications	Medium-risk applications	Low-risk applications
Could include ▪ Multiple stages of expert human review before publication ▪ Senior-level approval for deployment ▪ Ongoing monitoring after publication ▪ Seeking advice from relevant third parties, like the ICO or CAP Copy Advice	Could include ▪ Single expert human review before publication ▪ Managerial-level approval ▪ Periodic audits of outputs	Could include ▪ Automated screening against brand guidelines and regulatory standards ▪ Spot-check reviews by trained personnel ▪ Clear escalation procedures for flagged content

Platforms may consider implementing safeguards to ensure that users – particularly small businesses without in-house expertise – have access to tools and guidance that facilitate appropriate human oversight.

Organisations should document oversight processes and maintain records of review decisions, particularly for high-risk applications.

Principle 5 – Promoting societal wellbeing

GenAI should not be used to create, distribute, or amplify harmful, misleading, or exploitative advertising content. Where possible, AI should be deployed to enhance consumer protection and advertising standards.

AI systems' scale and efficiency create both risks and opportunities. Without appropriate safeguards, AI can rapidly generate and distribute misleading claims, fraudulent offers, or content that exploits vulnerable consumers. Conversely, AI can be leveraged to improve content moderation, detect harmful patterns, and ensure compliance with advertising standards.

Preventing harm is not merely about avoiding regulatory penalties – it protects consumers, maintains industry credibility, and builds long-term brand trust.

Organisations are advised to:

- Implement content screening to prevent AI from generating misleading, fraudulent, or harmful advertising
- Establish safeguards against targeting vulnerable populations with exploitative or inappropriate content
- Monitor AI systems for patterns that could facilitate harmful outcomes (e.g., promoting unsafe products, encouraging addictive behaviours)
- Ensure AI-generated claims about products or services are substantiated and compliant with CAP/BCAP Codes.
- Block or flag content that violates advertising standards before publication.
- Coordinate with advertising platforms and partners to prevent harmful content from circumventing safeguards across different channels.

Where practical, organisations should use AI proactively to:

- Enhance automated detection of harmful or non-compliant content
- Improve targeting accuracy to prevent inappropriate ad placements
- Support positive social outcomes through public interest campaigns
- Strengthen brand safety and suitability mechanisms

Note: Preventing harm takes precedence over efficiency gains. If AI systems cannot reliably avoid harmful outputs, they should not be deployed.

Principle 6 – Driving Brand Safety and Suitability

Organisations should assess and mitigate brand safety and suitability risks from AI-generated content and AI-driven ad placement, ensuring GenAI systems align with brand values and safety standards.

GenAI systems can rapidly create and distribute content at scale, increasing both efficiency and risk. AI-generated content may inadvertently misalign with brand values, include offensive material, or appear in brand-unsafe contexts alongside harmful content. The speed and volume of AI operations mean reputation damage can occur quickly and at scale.

Proactive brand safety measures—including robust content screening, placement controls, and real-time monitoring—are essential to prevent GenAI use from undermining brand credibility.

Organisations should implement appropriate measures, which may include:

- Conducting pre-deployment risk assessments specifically evaluating brand reputation risks of AI-generated content
- Configuring AI systems with brand safety parameters, including prohibited topics, tone guidelines, and visual standards
- Implementing content review processes proportionate to risk (aligned with Section 2.2)
- Monitoring AI-generated content and ad placements in real-time to detect brand safety violations
- Establishing rapid response procedures for when AI creates reputation-damaging content
- Using AI-powered brand suitability tools to prevent ads appearing alongside inappropriate content
- Conducting regular audits of AI outputs for alignment with brand values and cultural sensitivities

Brand safety controls should address areas such as:

- Content appropriateness (avoiding offensive, controversial, or misaligned messaging).
- Placement safety (preventing ads appearing near harmful, illegal, or inappropriate content).
- Cultural and regional sensitivities (ensuring content respects diverse audiences)
- Consistency with brand voice, values, and positioning.
- Crisis communication preparedness (having plans for addressing AI-related brand safety incidents publicly).

Principle 7 – Promoting Environmental Stewardship

When selecting GenAI tools and approaches, organisations should consider environmental implications alongside business objectives, favouring energy-efficient options where practical.

AI systems, particularly large-scale model training and high-volume inference, consume significant energy. Whilst individual advertising applications may have modest environmental footprints, collective industry usage creates meaningful impact. Organisations can support sustainability by making informed choices about AI deployment, recognising that environmental responsibility increasingly influences brand reputation and aligns with UK Net Zero targets.

Practical considerations include:

- Choosing energy-efficient AI models where they meet business requirements (e.g., smaller models for simpler tasks)
- Evaluating whether AI is necessary or whether traditional methods are sufficient
- Considering lifecycle environmental impact when comparing AI solutions to alternatives (e.g., AI-generated imagery vs physical production and travel)
- Where practical, selecting AI providers committed to renewable energy and carbon neutrality
- Consider signing up to AdGreen (www.weareadgreen.org), which tracks AI usage during production and enables users to measure the emissions impact of their AI tools.

Note: Environmental considerations should be balanced with other principles. They should not override consumer protection, fairness, or transparency requirements.

Principle 8: Ensuring Continuous Monitoring and Evaluation

Deployed GenAI systems should be monitored continuously to detect performance degradation, bias drift, compliance failures, or other issues requiring intervention.

AI systems do not remain static after deployment. They can evolve, degrade, or produce unintended outcomes over time due to changes in training data, operational context, or user behaviour. Without ongoing monitoring, organisations may fail to detect biased targeting, non-compliant content, or misalignment with brand values until significant harm occurs. Continuous monitoring enables organisations to identify and address issues promptly, maintain compliance with evolving standards, and adapt to changing consumer expectations.

Monitoring mechanisms could include areas such as:

- **Performance tracking**: Regular assessment of AI outputs against quality and compliance standards
- **Bias detection**: Ongoing analysis to identify discriminatory patterns in targeting or content generation
- **Compliance checks**: Verification that AI-generated content adheres to CAP/BCAP Codes and legal requirements
- **Brand alignment reviews**: Ensuring AI outputs remain consistent with organisational values
- **Audit trails**: Documentation of AI decisions, outputs, and interventions for accountability

Monitoring frequency should be proportionate to risk:

- **High-risk applications**: Real-time monitoring with automated alerts and immediate human review capabilities
- **Medium-risk applications**: Weekly or monthly audits with documented review processes
- **Low-risk applications**: Quarterly reviews with spot-checking protocols

When monitoring detects issues, organisations should consider:

- Suspending problematic AI functions if consumer harm is imminent
- Investigating root causes (model drift, data quality, configuration errors)
- Implementing corrective measures before resuming operations
- Documenting incidents and responses for future reference

Organisations should ensure they have adequate technical and human resources to implement their chosen monitoring approach effectively.

Note: Organisations should periodically review and update their AI governance policies to address evolving risks, regulatory changes, and technological developments.

4. Conclusion

This Best Practice Guide provides the UK advertising industry with a practical framework for deploying GenAI responsibly in online advertising.

The eight core principles – transparency, responsible data use, fairness, human oversight, harm prevention, brand safety, environmental consideration, and continuous monitoring – address the key ethical challenges that AI presents whilst enabling innovation. Together, they help advertisers, agencies, media owners, and technology providers navigate AI's opportunities and risks.

Responsible GenAI use requires more than good intentions. It demands robust governance, proportionate risk assessment, and ongoing vigilance. Organisations that embed these principles into their operations will be better positioned to maintain consumer trust, ensure regulatory compliance, and build sustainable competitive advantage.

GenAI will continue to evolve, and this Guide will evolve with it. Organisations adopting these principles are encouraged to share their experiences, contribute to industry forums, and collaborate on updates as technology, regulation, and best practices develop.

By committing to responsible AI practices, the advertising industry can harness GenAI's creative and operational potential whilst upholding the ethical standards consumers expect and deserve.

ANNEX (A) - Definitions

GenAI is a broad and evolving term that can mean different things to different stakeholders– advertisers, agencies, tech providers, regulators, and consumers. Clear definitions ensure everyone speaks the same language when discussing risks, responsibilities, and opportunities, and help organisations understand when they are using AI in ways that may trigger legal, ethical, or operational responsibilities. These definitions are tailored to online advertising and aligned to the Guide's objectives.

AI-Generated Content: Any advertising material where GenAI has been used to create, substantially alter, or enhance content in ways that materially affect its meaning, appearance, or impact.

Bias (in AI systems): Systematic errors or skewed outcomes in AI systems that unfairly favour or discriminate against certain groups, individuals, or characteristics. Bias can arise from unrepresentative training data, flawed algorithms, or the replication of existing societal prejudices, and may result in discriminatory targeting, exclusion, or stereotyping in advertising.

Deepfake: AI-generated synthetic media that convincingly depicts real individuals saying or doing things they did not actually say or do, created by manipulating or fabricating audio, video, or images using AI techniques. In advertising, deepfakes pose particular risks when used without consent or adequate disclosure.

High-Risk Applications: AI applications that pose significant potential for consumer harm, including: unauthorised or undisclosed use of real people's likenesses, advertising in regulated sectors (financial services, healthcare, gambling), content targeting vulnerable populations, or unsubstantiated claims about product or service efficacy.

Human Oversight: The process by which qualified personnel review, validate, or intervene in AI-generated outputs or decisions to ensure accuracy, compliance with ethical standards and regulations, and alignment with brand values. The level and frequency of human oversight should be proportionate to the risk level of the AI application.

Personalisation (AI-driven): The use of AI to tailor advertising content, messaging, or delivery to individual users based on their characteristics, behaviour, or predicted preferences, going beyond traditional demographic segmentation to create unique or highly specific advertising experiences.

Synthetic Media: A subset of AI-generated content. Audio, visual, or audiovisual content that has been artificially generated or manipulated using AI to create realistic but fabricated representations of people, events, or scenarios.

Training Data: The datasets used to develop and teach AI systems, which may include text, images, videos, or other content. In advertising contexts, training data influences what AI systems can generate, how they make decisions, and what biases they may contain.

Vulnerable Populations: Groups requiring enhanced protection in advertising contexts due to factors that may increase susceptibility to harm or exploitation. This includes children (under 18), elderly consumers, individuals with cognitive impairments, and those experiencing financial distress, or other groups requiring enhanced protection.

ANNEX (B) – Recommended Cumulative Framework for Implementing Best Practice

Principle	Minimum	Better	Best
1. Ensuring Transparency	- Review B/CAP Codes, in particular Section 3 on Misleading advertising,[1] - Follow ASA Quick Guide on misleading advertising rules.[2] - Disclose AI use where the omission of that fact could mislead consumers	- Apply risk-based disclosure decisions for all AI-generated content - Document disclosure decisions and rationale - Review disclosure practices quarterly	- Adopt *Beyond Simple Labelling*[3] framework organisation wide - Implement consistent disclosure protocols across all campaigns - Conduct annual reviews aligned with evolving industry standards
2. Ensuring Responsible Use of Data	- Comply with UK GDPR and Data Protection Act 2018 - Follow ICO Guidance on AI and Data Protection[4] - Ensure PECR compliance for cookies	- Implement granular consent mechanisms for AI processing - Conduct annual data protection audits - Verify third-party AI provider compliance	- Publish annual transparency reports on AI data use - Commission independent data protection audits - Maintain comprehensive documentation of AI data flows - Conduct quarterly compliance reviews
3. Preventing Bias and Ensuring Fairness	- Ensure AI use complies with Equality Act 2010.[5] - Protect special category data per UK GDPR Article 9[6] - Conduct basic bias checks before deployment	- Conduct bias audits on AI models and datasets - Test outputs across diverse demographic groups - Review ICO Guidance on bias and discrimination[7] - Implement corrective measures when bias detected	- Train AI models on diverse, representative datasets - Document bias mitigation processes comprehensively - Conduct ongoing monitoring for discriminatory patterns - Publish fairness assessments and outcomes

[1] 03 Misleading advertising - ASA | CAPCAP Code and 03 Misleading advertising - BCAP Code
[2] A quick guide to changes to the misleading advertising rules - ASA | CAP
[3] Beyond Simple Labelling
[4] Guidance on AI and data protection | ICO
[5] Equality Act 2010: guidance - GOV.UK
[6] UK GDPR, Art.9
[7] What about fairness, bias and discrimination? | ICO

4. Ensuring Human Oversight and Accountability	- Maintain awareness of responsibilities under the B/CAP Codes - Implement human review for high-risk AI applications - Document approval processes and decisions	- Apply risk-proportionate oversight (Section 3.2) - Require managerial approval for medium-risk applications - Maintain review records for all AI-generated content	- Establish dedicated AI governance committee - Implement multi-stage review for high-risk applications - Align oversight processes with EASA and ICAS standards - Update governance framework annually
5. Promoting societal wellbeing	- Review B/CAP Codes in particular Section 4 on Harm and offence[8][9]. - Follow rule 1.3 of the CAP Code[10] and/or rule 1.2 of the BCAP Code[11] on responsibility to consumers - Screen AI outputs for misleading or harmful content - Block content violating advertising standards	- Implement automated content filtering for harmful patterns - Apply human review for high-risk categories (finance, health) - Monitor for exploitation of vulnerable groups	- Deploy AI-powered content analysis tools - Establish policies ensuring brand values alignment - Assess impact on vulnerable and marginalised groups - Balance commercial objectives with societal benefit - Proactively prevent harmful content generation
6. Driving Brand Safety and Reputation	- Maintain awareness of responsibilities under the B/CAP Codes - Configure AI with basic brand safety parameters - Monitor for brand-unsafe ad placements	- Conduct pre-launch brand reputation risk assessments - Test for culturally insensitive content - Implement real-time monitoring of AI placements - Establish rapid response procedures	- Deploy AI-powered brand suitability tools - Conduct real-time monitoring with automated alerts - Train custom models on brand-approved content - Maintain comprehensive brand safety documentation - Audit AI outputs regularly for brand alignment
7. Promoting Environmental Stewardship	- Review Ad Net Zero's AI Sustainability guidance[12] - Consider environmental impact in AI tool selection	- Compare energy consumption of AI tool options - Select lightweight models for simpler tasks	- Establish comprehensive environmental impact policies - Quantify AI-related energy consumption - Partner with carbon-neutral AI providers

[8] BCAP Code Section 4
[9] CAP Code Section 4
[10] CAP Code Section 1.3
[11] BCAP Code Section 1.2
[12] Ad Net Zero AI x Sustainability Exploratory Q1 2025 Series Summary, including Key Quotes and Learnings | Ad Net Zero Limited

	- Favour energy-efficient options where practical	- Assess AI vs non-AI environmental trade-offs - Document environmental considerations	- Promote AI use with positive environmental outcomes - Report on annual sustainability efforts annually
8. Ensuring continuous Monitoring and Evaluation	- Maintain awareness of responsibilities under the B/CAP Codes - Implement human monitoring of AI outputs - Conduct periodic reviews of AI performance	- Deploy real-time monitoring for high-risk applications - Conduct monthly audits of AI targeting and content - Document monitoring findings and corrective actions	- Implement continuous monitoring with automated alerts - Conduct risk-proportionate audits (daily/weekly/quarterly) - Explore novel AI applications proactively - Publish annual compliance reports - Maintain comprehensive audit trails

ADVERTISING ASSOCIATION

Find us at:
95 Aldwych, London,
WC2B 4JF

adassoc.org.uk
aa@adassoc.org.uk

@advertising-association
@ad_association
@ad_association