CoRE-AI™
Coalition for Responsible Evolution of AI

SUBMISSION BY
THE COALITION FOR RESPONSIBLE EVOLUTION OF AI (CoRE-AI)

on

Draft Guidelines for Responsible Labelling of Synthetically Generated Content in Advertising

Issued by:
Advertising Standards Council of India (ASCI)

Date of Draft:
8 May 2026

Response Timeline:
13th June 2026

July 2026

Submission by The Coalition for Responsible Evolution of AI (CoRE-AI) on Draft Guidelines for Responsible Labelling of Synthetically Generated Content in Advertising

(C) 2026 The Dialogue®. All rights reserved.

About Coalition for Responsible Evolution of AI (CoRE-AI)
Established in July 2024, CoRE-AI (Coalition for Responsible Evolution of AI) is India's largest multi-stakeholder coalition focused on fostering the responsible and ethical development of AI technologies. By bringing together stakeholders from government, industry, academia, startups, and civil society, CoRE-AI aims to drive collaborative efforts that address the risks associated with AI while maximizing its societal benefits. Hosted by The Dialogue, the initiative seeks to guide India's AI journey, ensuring that technological advancements align with ethical standards to benefit the broader public.

For more information: Visit https://core-ai.in/

Publication date: July 6, 2026

Catalogue no.: CAI/WC/0726/01

Suggested Citation: CoRE-AI Secretariat. (2026). Submission on the Draft Guidelines for Responsible Labelling of Synthetically Generated Content in Advertising. Coalition for Responsible Evolution of AI (CoRE-AI).

Disclaimer: The facts and information in this report may be reproduced only after giving due attribution to CoRE-AI (Coalition for Responsible Evolution of AI)

Contents

Preamble

This response is submitted to the **Advertising Standards Council of India (ASCI)** in the context of its public consultation on the **Draft Guidelines for Responsible Labelling of Synthetically Generated Content in Advertising**, dated *8 May 2026*. At the **Coalition on Responsible Evolution of AI (CoRE-AI)**, we welcome the initiative taken by ASCI to proactively address the growing deployment of Artificial Intelligence and synthetically generated content across advertising ecosystems in India.

The proliferation of AI-generated content in advertising, spanning synthetic voices, deepfakes, virtual influencers, and AI-augmented visuals, poses genuine questions of consumer transparency and trust. ASCI's effort to develop a principled, risk-calibrated framework is therefore both timely and commendable. This response engages with the draft guidelines section by section, acknowledging its sound foundations while flagging areas where greater definitional clarity, operational guidance, or enforcement architecture would strengthen the final instrument.

This submission is structured to mirror the order of the draft guidelines and covers: (i) Objective and Scope; (ii) the Risk-Based Framework and Assessment; (iii) the three risk tiers ie. High, Medium, and Low; and (iv) broader observations including enforcement architecture, implementation timeline, consumer awareness and consistency with international practice.

I. Objective and Scope

A. What the Draft Gets Right

The foundational objective of the guidelines to ensure that advertisements remain "honest, decent, safe, and fair" even as brands increasingly harness AI-powered tools is sound and unambiguous. The deliberate decision to ground the framework in consumer outcomes rather than in the regulation of AI as a technology is particularly commendable. This is a sophisticated and forward-looking drafting choice, one that insulates the guidelines from rapid obsolescence as the underlying technology evolves. Regulating the effect rather than the tool is the appropriate policy posture.

The definition of synthetically generated content is suitably broad: "any advertisement that is artificially created, modified, or materially altered to appear authentic, including deepfakes, synthetic spokespersons, AI voices, and materially altered imagery." This definition captures the plurality of AI-driven advertising formats currently in use without attempting to enumerate every conceivable technique, thereby retaining adaptive relevance.

The dual assessment questions anchoring the risk framework are also well-framed: whether AI use materially influences purchase decisions and whether the absence of a disclosure would create a false or misleading impression of the product's claims/benefits. These are outcome-oriented, consumer-centric tests that provide an intelligible standard for advertisers and adjudicatory bodies alike.

B. Areas Requiring Greater Clarity

While the definition of synthetically generated content is broad in scope, the phrase "materially altered to appear authentic" requires further unpacking. Specifically:

i. What threshold **distinguishes "material" alteration from minor enhancement**, particularly for image-based advertising? While Section 3A (Low Risk) lists examples such as colour correction and blemish removal, these lists are illustrative rather than exhaustive, and disputes at the margin are foreseeable. A set of guiding principles, rather than just examples, would be more durable.

ii. The term **"to appear authentic" implies intentionality**. The guidelines **should clarify whether the test is subjective (what the advertiser intended) or objective (how a reasonable consumer would perceive the content).** The latter is the appropriate standard for a consumer protection framework.

iii. The **phrase "materially influences purchase decisions" appears both in the scope definition and as an assessment criterion** without a clear benchmark. Guidance on **how this is to be evaluated, for instance, by reference to the centrality of the AI-generated element to the product claim would reduce inconsistent application**.

iv. Additionally, **the guidelines do not address the role of advertising agencies and technology platforms in the compliance chain.** Advertisers may delegate production of AI-generated content to third-party vendors. **Clarity on primary responsibility whether it rests with the brand, the agency, or the platform is essential for workable enforcement.**

A more fundamental concern is the potential for regulatory duplication. **The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026 (SGI Amendments)**, notified on 10 February 2026, already establish comprehensive due diligence obligations pertaining to synthetically generated content, including specific exclusions to ensure that routine, good-faith uses of AI for editing, formatting, or improving content are not captured. **The Guidelines for Prevention of Misleading Advertisements and Endorsements, 2022,** separately impose duties on advertisers to ensure truthful and honest representations. The necessity for a parallel labelling regime, over and above these existing frameworks, requires clearer articulation to avoid duplicative compliance burdens, potential inconsistency between frameworks, and regulatory confusion for advertisers.

II. Risk-Based Framework and Assessment

A. What the Draft Gets Right

The adoption of a tiered, risk-based framework is the most strategically sound feature of these guidelines. A single, uniform disclosure requirement applied to all AI use in advertising would have been both disproportionate and counterproductive, risking label fatigue of the kind the draft itself expressly seeks to avoid. The three-tier structure of High Risk (Prohibited), Medium Risk (Labelling Required), and Low Risk (No Labelling Required) correctly calibrates regulatory burden to actual consumer harm potential.

The two-question assessment mechanism is practically useful. It asks: (i) whether AI use materially influences consumer decisions; and (ii) whether the absence of disclosure would create a false or misleading impression. These are questions that a competent advertiser can self-apply during the campaign design stage, and that a regulator can apply retrospectively. The framework therefore functions both as an ex ante compliance tool and an ex post adjudicatory standard.

B. Areas Requiring Greater Clarity

Notwithstanding the above, the risk-assessment mechanism would benefit from the following refinements:

i. The **two assessment questions function in combination but their interplay is not fully articulated**. For example: what is the outcome if an AI element materially influences purchase decisions but its absence would not create a false impression? **A decision tree or an illustrative matrix mapping question-responses to risk categories would materially enhance the framework's usability.**

ii. The **framework is self-assessed by the advertiser**. While this is the standard approach in industry self-regulation, there is **no indication of whether ASCI proposes to issue sector-specific guidance, maintain a look-up table of common ad formats, or establish a safe-harbour mechanism for good-faith self-assessment errors.** These additions would reduce compliance uncertainty, particularly for small and medium-sized advertisers.

III. Category 1 | High Risk (Prohibited Content)

A. What the Draft Gets Right

The enumeration of High Risk content as a category of outright prohibition, rather than as content requiring enhanced disclosure, is the correct approach. The examples provided are well-chosen and capture the most egregious forms of AI misuse in advertising:

i. The prohibition on fabricated endorsements and testimonials (Clause 1A) addresses a practice that directly undermines consumer trust and violates foundational advertising standards irrespective of AI involvement. The explicit mention of this category removes any ambiguity that AI-generated testimonials might be treated differently from conventional false endorsements.

ii. The prohibition on deepfakes and the use of a person's likeness without consent (Clause 1D) is essential and appropriately aligned with both the right to publicity and the emerging legal and regulatory discourse on non-consensual synthetic media. This provision correctly signals that consent, not merely disclosure, is the operative standard for likeness replication.

iii. The prohibition on AI-generated fictional authority figures with identifiable professional cues (Clause 1E) such as a synthetic doctor endorsing a health supplement is an important innovation. This addresses a growing and particularly harmful category of AI misuse that existing advertising codes did not explicitly contemplate.

iv. The prohibition on fake locations (Clause 1C), particularly in sectors such as real estate and travel where location authenticity is central to a consumer's decision, is well-targeted and practically significant.

B. Areas Requiring Greater Clarity

Several provisions in the High Risk category would benefit from definitional or operational refinement:

i. Clause 1B prohibits "exaggerating product results or features through claims or visual representations to create a misleading impression." This is a pre-existing principle under the ASCI Code and its inclusion here, without qualification, conflates ordinary advertising exaggeration (which may or may not involve AI) with AI-specific risk. **The provision should be redrafted to clarify that it is concerned specifically with AI-enabled exaggeration ie., where synthetic content creates a materially false impression of product performance that would not arise from conventional advertising.** Otherwise, the provision is redundant with existing ASCI rules rather than complementary to them.

ii. The phrase **"infringe on rights"** used in the High Risk category definition is over-broad, given that both illegal conduct and violations of the ASCI Code are already separately captured. **The phrase should be deleted to avoid inadvertently broadening the category beyond its intended scope.**

iii. Clause 1D prohibits use of "copyrighted work" without consent in the same breath as deepfakes and likeness rights. Copyright infringement is governed by separate intellectual property legislation and enforcement mechanisms. **While the principle is sound, bundling it within a consumer-facing advertising labelling framework may create confusion about the appropriate enforcement forum.**

iv. The High Risk category does not address AI-generated content targeting children or other vulnerable populations, despite the draft's stated concern with such populations in its Objective section. **Given the acute sensitivity of advertising directed at minors, an explicit High Risk sub-category for AI-generated content in child-directed advertising would close a significant gap.**

v. The draft is silent on the consequences of a High Risk violation beyond the statement that such content "will violate the ASCI Code even if an AI label is used." **Clarity on the remedial and enforcement implications including whether High Risk violations attract expedited complaint processing would be beneficial.**

Further, it must be noted that most **conduct enumerated under the High Risk category is already prohibited under existing law ie**. the ASCI Code and the Misleading Advertisements Guidelines, independent of any AI labelling requirement. The draft itself acknowledges that AI labels do not cure High Risk violations. This confirms that the labelling framework adds no remedial value to the High Risk category; it is **worth considering if this categorization is even needed and or should the same be enforced through existing mechanisms** without the need for a parallel AI-specific instrument in this tier.

IV. Category 2 | Medium Risk (Labelling Required)

A. What the Draft Gets Right

The Medium Risk category represents the operational heart of the guidelines ie. the tier where disclosure obligations are triggered and where the greatest volume of AI-assisted advertising will fall. The examples enumerated are well-calibrated:

i. The mandatory disclosure for virtual and synthetically generated influencers and ambassadors (Clause 2A) is appropriate and consistent with international best practice. Consumers interacting with AI-generated influencers have a legitimate interest in knowing that the personality, lifestyle, and endorsements they are engaging with are synthetic.

ii. The requirement to label replication of a real person's likeness or voice even with consent for personalised messaging (Clause 2B) is a nuanced and important provision. The addition of "even with consent" correctly recognises that consumer transparency is a distinct interest from the rights of the individual whose likeness is used. A person may consent to their likeness being used without that use being apparent to consumers; labelling serves the latter.

iii. Clause 2C, which requires disclosure for AI-generated visuals depicting product performance unless the visuals accurately replicate actual performance, draws a clear and principled line: AI may be used to illustrate, but not to misrepresent. This is a sound standard that reflects both consumer protection values and the legitimate creative uses of AI in advertising.

iv. The specific disclosure requirement for AI-generated paid or sponsored product suggestions (Clause 2G) to be labelled specifically as 'sponsored by' addresses a newly emerging category of AI-driven commercial recommendation at the intersection of native advertising and algorithmic suggestion.

B. Areas Requiring Greater Clarity

Notwithstanding its strengths, the Medium Risk category raises the most significant concerns in this submission, both in terms of scope and practical effect:

i. Industry data indicates that a significant majority of Indian media and advertising companies already use AI models in production at scale. The Medium Risk category, which encompasses broadly defined uses such as "creating realistic events, settings, or situations entirely with AI," AI-generated product visuals, and AI-generated sound effects, would in practical effect require labelling on the vast majority of contemporary advertisements. **A disclosure regime that labels nearly all advertisements provides no meaningful signal to consumers and would normalise labels to the point of ineffectiveness, precisely the label fatigue the draft seeks to avoid.**

ii. Several Medium Risk examples describe conduct that is functionally equivalent to traditional production techniques that have never required disclosure. Creating realistic settings or situations entirely with AI is analogous to constructed sets and cinematography. **Demonstrating a product that does not currently exist (Clause 2E) is a practice with decades of precedent through architectural renderings and CGI, neither of which has historically required consumer disclosure. The test should be whether the ad misleads about the product, not how it was made.**

iii. Clause 2B addresses replication of a real person's likeness "for personalised messaging." **The scope of "personalised messaging" is ambiguous. Does this extend to mass-broadcast advertisements dynamically customised using AI** (e.g., a single advertisement where the spokesperson's dialogue is individually tailored for each viewer through AI)? **Emerging dynamic ad-serving technologies make this a live question, and the guidelines should address it explicitly.**

iv. Clause 2F requires disclosure for AI-generated exaggerated sound effects "highly relevant to the product's core features." **The term "highly relevant" introduces a threshold that is difficult to assess consistently.** For a headphone brand, any audio in the advertisement is arguably highly relevant; for a food brand, enhanced sizzle effects could be relevant; for a car brand, engine sounds may be central. **Greater guidance on what constitutes relevance to core features or alternatively, adopting a broader disclosure default for AI-altered audio would reduce ambiguity.**

v. For advertisements that already carry clear, substantive creative disclaimers, such as "Sound Enhanced" or "3D Representation Only", imposing an **additional AI-specific platform label creates redundant double-labelling that degrades user experience without providing incremental consumer benefit.** ASCI should introduce an **exemption for advertisements that already feature standard, clear disclosures regarding creative enhancements.**

vi. The medium risk category includes "creating realistic events, settings, or situations entirely with AI" alongside "background and ambient elements" in the low risk category. This creates a direct internal contradiction. **A generic AI-generated meeting room or outdoor setting could fall in either category.** The **reference to 'settings' and 'situations' should be removed or substantially clarified with concrete examples of the kind of content that would actually mislead a consumer.**

vii. The **guidelines do not address whether the Medium Risk labelling obligation applies to advertisements on all media formats equally** like television, digital video, audio-only platforms, out-of-home, print **or whether platform-specific execution standards are contemplated.** A label that is legible on a digital video advertisement may be impractical on audio-only or outdoor formats. **The guidelines should either address multi-format execution or cross-refer to media-specific guidance to be developed.**

> **Additional Concerns**
>
> Recent academic research from NYU Stern and Emory University found that disclosure of AI involvement in advertising reduces click-through rates by an average of 31.5%. This evidence base should be squarely confronted in the final guidelines. Where mandatory labelling is retained for the Medium Risk category, ASCI should demonstrate with evidence that the consumer transparency benefit outweighs the established commercial and innovation cost, particularly for MSMEs and small advertisers who stand to be disproportionately affected.

V. Category 3 | Low Risk (No Labelling Required)

A. What the Draft Gets Right

The Low Risk carve-out is an essential component of a proportionate regulatory framework. The examples cited are well-chosen and reflect a practical understanding of how AI tools are routinely embedded in modern content production:

i. Routine post-production enhancements like colour correction, noise reduction, standard blemish removal, and minor lighting adjustments (Clause 3A) are correctly excluded from labelling requirements. Requiring disclosure for such commonplace production practices would be administratively burdensome and would contribute precisely to the label fatigue the draft seeks to avoid.

ii. The exclusion of purely decorative AI-generated backgrounds, abstract skylines, and ambient audio (Clause 3B) draws a principled distinction between elements that influence consumer perception of the product and those that merely provide aesthetic context. This distinction is sound and should be retained.

iii. The "fantastical elements" exception (Clause 3C) for obvious, unrealistic effects such as dragons or fairies correctly recognises that audiences apply contextual literacy to clearly fantastical content. No reasonable consumer is misled into believing that a dragon is a real product attribute.

iv. The administrative and text uses carve-out (Clause 3D) covering AI-generated advertising copy, accessibility descriptions, and document preparation reflects the ubiquitous deployment of AI in back-end content production and correctly reserves disclosure obligations for consumer-facing material that affects perception of the product.

B. Areas Requiring Greater Clarity

The Low Risk category, while well-conceived, raises some line-drawing questions:

i. The phrase "minor enhancements" (Clause 3A) does necessary work but is not precisely defined. In the context of beauty and personal care advertising, AI-enhanced skin or body imagery occupies a contested space: it may appear to be standard retouching but may create aspirational impressions that are materially misleading. **The guidelines should clarify whether AI-driven body-modification enhancements, beyond "standard blemish removal", fall within Low Risk or Medium Risk, particularly given the existing discourse around unrealistic body image standards in advertising.**

ii. The distinction between decorative ambient sound (Low Risk, Clause 3B) and exaggerated product-relevant sound effects (Medium Risk, Clause 2F) is sensible in principle but may be contested in practice for certain product categories. **The guidelines would benefit from one or two additional worked examples illustrating how this distinction operates across sectors such as food, automotive, and consumer electronics.**

iii. Purely fictional or animated synthetic characters presenting no realistic risk of deception should be formally accommodated within the Low Risk category alongside the fantastical elements exception.

> **Additional Concerns**
>
> ASCI should consider whether the Low Risk exemption rationale, that AI use has no material impact on a consumer's ability to make an informed choice, applies equally to a broader set of Medium Risk examples that are functionally indistinguishable from traditional production techniques. Where existing regulatory frameworks already adequately address the underlying consumer harm, an additional AI labelling requirement does not enhance protection.

VI. Broader Observations

A. Regulatory Overlap and the Necessity Test

As noted above, India's existing regulatory landscape, including the SGI Amendments (2026) and the Misleading Advertisements Guidelines (2022), already provides comprehensive protection against the core consumer harms that these guidelines seek to address. Traditional advertising has always used tools to construct, enhance, or

dramatise reality: airbrushing, CGI, set design, special effects, and stock photography compositing. None of these have required disclosure, and AI is, in functional terms, a new method to achieve the same results.

The necessity for a parallel AI-specific labelling regime, over and above existing frameworks, has not been evidenced by documented consumer complaints or studies demonstrating that Indian consumers have been misled or harmed specifically because an advertisement was produced using AI tools (as distinct from being misleading in substance). A fundamental regulatory principle is that intervention should be proportionate to demonstrated harm. ASCI should publish its evidence base for the Medium Risk labelling obligations in particular.

B. Implementation Timeline and Industry Readiness

ASCI is actively building its enforcement infrastructure, and it is reported that these guidelines may be endorsed by the Ministry of Information and Broadcasting. Rushed implementation risks disrupting the digital ecosystem before the industry has had sufficient opportunity to integrate technical provenance standards such as C2PA that would make AI identification and labelling technically reliable. A **minimum grace window of twelve months before enforcement commences**, coupled with a phased implementation approach for smaller advertisers, would be appropriate. **ASCI should also extend the public consultation period to sixty to ninety days given the complexity and breadth of the framework's impact**, and should specifically engage small and medium-sized businesses and entrepreneurs to assess the practical compliance burden.

C. Consumer Awareness

The effectiveness of any disclosure regime depends on consumer literacy. Disclosures are only meaningful if consumers understand what an AI label signifies and why it is relevant to their decision-making. Without a parallel public awareness initiative, disclosures risk being noticed but not understood. Labels and warnings are also well established to drive behaviour change primarily through negative affect rather than information transfer; mandatory labels on benign advertisements can therefore cause unwarranted reputational harm to brands without any corresponding consumer protection benefit.

D. Consistency with International Practice

The prevailing international approach to AI-generated content in advertising is risk-based, harm-focused, and calibrated rather than requiring blanket labelling of AI-assisted advertisements. In the UK, the Committee of Advertising Practice has confirmed there is no blanket legal requirement to disclose AI use in advertising, and that AI-specific disclosure is warranted only where failure to disclose renders an advertisement misleading. Australia's self-regulatory system similarly applies existing advertising codes to all content regardless of production method. It recognizes that existing advertising codes apply to all content, regardless of how it is created - AI or otherwise. India's draft guidelines should align with this international trajectory by focusing squarely on genuinely misleading or harmful AI uses, rather than adopting an approach that risks over-regulation relative to its global peers.

VII. Conclusion

CoRE-AI remains committed to constructive engagement with ASCI on these matters and welcomes the opportunity to participate in further consultations or technical working groups as the guidelines are finalised.

core-ai.in

secretariat@core-ai.in

@Core-Ai

@CoalitionforResponsibleEvolutionofAI