

AI Preferences
Internet-Draft
Intended status: Experimental
Expires: 12 March 2026

B. Silver
Advance
8 September 2025

Vocabulary For Expressing AI Substitutive Usage
draft-silver-aipref-vocab-substitutive-00

Abstract

This Internet Draft proposes a category entitled "AI Substitutive Use" which would enable parties to express a preference regarding how digital assets are used by automated processing systems, with a focus on post-training (inference-time) uses that are likely to result in the creation of AI-generated outputs that substitute for the original asset. The proposal is for this category to nest within the larger category of Automated Processing, currently envisaged in the working group draft [AIPREF-VOCAB] (21 July 2025).

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://datatracker.ietf.org/doc/draft-silver-aipref-vocab-substitutive/>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-silver-aipref-vocab-substitutive/>.

Discussion of this document takes place on the AI Preferences Working Group mailing list (<mailto:ai-control@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/ai-control/>. Subscribe at <https://www.ietf.org/mailman/listinfo/ai-control/>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 12 March 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Rationale	2
2. Conventions and Definitions	3
3. Vocabulary Definition	4
3.1. AI Substitutive Use Category (New)	4
4. IANA Considerations	4
5. Normative References	4
6. Informative References	5
Author's Address	5

1. Rationale

Existing mechanisms for expressing preferences, including those under consideration by the AI Preferences WG do not address concerns which have been strongly articulated about the practice of using digital assets as input to AI models to generate outputs which substitute for or undermine the value of the original assets. This gap leaves a broad group of stakeholders (including creators, journalists and publishers) without a means to express a preference regarding a type of use which is already having a material adverse impact on their rights. Developers and deployers of AI systems are also left without a clear, standardized preference signal regarding such uses, which results in "blunt" approaches to gathering such content - exposing them to legal risk. This proposal intends to define a tailored preference category to address the specific need, improve visibility across the board and support continued broad access to information

and content.

The use of digital assets for inferencing "in real time" is widespread as a means of improving the accuracy and contextual relevance of outputs, such as through the use of techniques such as Retrieval-Augmented Generation (RAG) [RAG2020]. The flipside of that value is that such outputs are inclined to substitute for or dilute the value of the original asset, which decreases user engagement with the original asset. This harms revenue opportunities and undermines the ability of the owner or distributor of the original asset to connect directly with their intended audience. For example, the use of journalistic material to create AI-generated summaries which have resulted in the substantial reduction of internet traffic to online publications. In the longer term, this jeopardises the sustainability of those enterprises and the underlying incentives to create and publish such material. To mitigate this, some are moving content behind paywalls and deploying other means of limiting open access - diminishing access to information and content.

Should incentives to create diminish, AI innovation will also suffer as a result of less quality content on which to build a distribution funnel. This would also undermine the sustainability and verifiability of news and information services relied upon by the public and government institutions. Where the AI model or platform takes on the role of information gatekeeper and shaper, connections between the public and original sources can be severed (or warped), which undermines the ability and willingness of internet users to ensure what they are reading, hearing or watching matches the original source(s), allowing factual misrepresentations to propagate and go unchecked.

Creators have also justifiably expressed the need for a preference that addresses the use of their assets to create derivative works "in the style of" such original assets. Creators are harmed by the unfettered use of their works as inputs to AI Models to create outputs which dilute the market for their works, adopting distinctive elements and styles established by the creators themselves - which also harms their moral rights and interests to protect the integrity of their works and ensure attribution.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

For the purposes of this document, the following terms are used:

- * ***Post-training (inference-time)***: Uses of an AI/ML model that occur after the model has been trained and frozen, typically when generating outputs in response to inputs at runtime.
- * ***Retrieval-Augmented Generation (RAG)***: A technique where external content is retrieved at query time and supplied to a model to condition the generated output. This document references RAG as a common mechanism by which substitutive outputs may be produced [RAG2020].

3. Vocabulary Definition

3.1. AI Substitutive Use Category (New)

The Act of using one or more assets as input to a trained AI/ML model (as opposed to the training of the model) which results in an output which incorporates, summarizes, aggregates or reproduces the assets, including stylistic elements thereof; provided, however that this category does not cover the use of a lawfully acquired digital asset where carried out directly by an end user (as opposed to a search application or bot) as input to a trained model to create a summary of such digital asset.

The use of assets for AI Substitutive Use is a proper subset of Automated Processing usage [AIPREF-VOCAB].

This category is distinct from AI Training or Generative AI Training, as it addresses uses that occur after a model has been trained, during inference. It is also distinct from Search, which covers uses that direct users back to the original asset. Substitutive Use, by contrast, describes outputs that replace, reduce the utility of, or make the source asset redundant to users by summarizing, reproducing, or restyling its contents.

Consistent with that objective, this category would not apply where end users are summarising digital assets which they have already acquired, outside of the context of search or retrieving such assets from online locations in summarized form.

4. IANA Considerations

This document has no IANA actions.

5. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.
- [RFC9309] Koster, M., Illyes, G., Zeller, H., and L. Sassman, "Robots Exclusion Protocol", RFC 9309, DOI 10.17487/RFC9309, September 2022, <<https://www.rfc-editor.org/info/rfc9309>>.

6. Informative References

- [AIPREF-VOCAB] IETF AI Preferences Working Group, "AI Preferences Vocabulary", Work in Progress, Internet-Draft, draft-ietf-aipref-vocab, 21 July 2025, <<https://datatracker.ietf.org/doc/draft-ietf-aipref-vocab/>>.
- [RAG2020] Facebook AI Research, University College London, and New York University, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks", NeurIPS 2020, 2020, <<https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>>.

Author's Address

Bradley Silver
Advance
United States of America
Email: bsilver@advance.com