

AI Preferences  
Internet-Draft  
Intended status: Standards Track  
Expires: 22 January 2026

P. Keller  
Open Future  
M. Thomson, Ed.  
Mozilla  
21 July 2025

A Vocabulary For Expressing AI Usage Preferences  
draft-ietf-aipref-vocab-02

## Abstract

This document proposes a standardized vocabulary for expressing preferences related to how digital assets are used by automated processing systems. This vocabulary allows for the creation of structured declarations about restrictions or permissions for use of digital assets by such systems.

## 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://ietf-wg-aipref.github.io/drafts/draft-ietf-aipref-vocab.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-aipref-vocab/>.

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/>.

Source for this draft and an issue tracker can be found at <https://github.com/ietf-wg-aipref/drafts>.

## 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 22 January 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. Introduction . . . . .	3
2. Conventions and Definitions . . . . .	3
3. Statements of Preference . . . . .	4
3.1. Conformance . . . . .	5
4. Vocabulary Definition . . . . .	5
4.1. Automated Processing Category . . . . .	6
4.2. AI Training Category . . . . .	6
4.3. Generative AI Training Category . . . . .	7
4.4. AI Use Category . . . . .	7
4.5. Search Category . . . . .	7
5. Usage . . . . .	7
5.1. More Specific Instructions . . . . .	7
5.2. Vocabulary Extensions . . . . .	8
6. Exemplary Serialization Format . . . . .	8
6.1. Usage Category Labels . . . . .	8
6.2. Preference Labels . . . . .	9
6.3. Text Encoding . . . . .	9
6.4. Syntax Extensions . . . . .	9
6.5. Processing Algorithm . . . . .	10
6.6. Alternative Formats . . . . .	11
7. Consulting a Preference Expression . . . . .	11
7.1. Combining Preferences . . . . .	12
8. Applicability and Legal Effect . . . . .	12
9. Security Considerations . . . . .	13
10. IANA Considerations . . . . .	13

11. References . . . . .	13
11.1. Normative References . . . . .	13
11.2. Informative References . . . . .	14
Acknowledgments . . . . .	14
Authors' Addresses . . . . .	14

## 1. Introduction

This document defines a common vocabulary of terms for automated systems that process digital assets. The primary purpose of this vocabulary is to enable machine-readable expressions of preferences about how digital assets are used by automated processing systems in the context of training AI models and other forms of automated processing.

The terms defined by the vocabulary can be used to describe, in a standardized way, the types of uses that a declaring party may wish to explicitly restrict or allow. Preferences are then expressed as a grant or denial of permission concerning each of the types of use defined in the vocabulary. This ensures that preferences can be communicated, processed, and stored in a consistent and interoperable manner.

The vocabulary or the preferences that might be expressed do not proscribe how automated processing systems obtain or act on preferences. Separate documents will describe how preferences might be associated with assets. It is designed to ensure that preference information can be exchanged between different systems and consistently understood.

The vocabulary is intended to be usable both where expressing preferences results in legal obligations and where there are no associated legal protections. That is, preferences can be expressed to invoke specific protections, or they can be made without any presumption of specific legal consequences. Potential legal obligations include rights reservations made by rightholders in jurisdictions with conditional exceptions on copyright protections. Expressing preferences is without prejudice to applicable laws, including the applicability of exceptions and limitations to copyright.

## 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.

This document uses the following terms:

AI:

Artificial intelligence or machine learning, which are used interchangeably in this document, refer to computer systems or algorithms that are trained to accomplish a task.

AI Training:

Processing input data to identify statistical trends in order to produce an AI system.

Asset:

A digital file or stream of data, usually with associated metadata.

Declaring party:

The entity that expresses a preference with regards to an Asset.

### 3. Statements of Preference

The vocabulary is a set of categories, each of which is defined to cover a class of usage for assets. Section 4 defines these categories in more detail.

A statement of preference is made about an asset. Statements of preferences can assign preferences to each of the categories of use in the vocabulary. Preferences regarding each category can be expressed either to allow or disallow the usage associated with the category.

A statement of preferences can express preferences about some, all, or none of the categories from the vocabulary. This can mean that no preference is expressed for a given usage category.

Some categories describe a proper subset of the usages of other categories. A preference that is expressed for the more general category applies if no preference is expressed for the more specific category.

For example, the Automated Processing category might be assigned a preference that allows the associated usage. In the absence of any statement of preference regarding the AI Training category, that usage would also be allowed, as AI Training is a subset of the Automated Processing category. In comparison, an explicit preference regarding AI Training might disallow that usage, while permitting other usage within the Automated Processing category.

After processing a statement of preferences the recipient assigns each category of use one of three preference values: "allowed", "disallowed", or "unknown".

In the absence of a statement of preference, all usage categories are assigned a preference value of "unknown".

### 3.1. Conformance

This document -- and those documents that define concrete uses of this vocabulary -- describe how usage preferences are associated with assets. Conformance to the specification means following the normative language that defines the construction and interpretation of usage preferences. The process of obtaining preferences has very limited scope for variation between implementations.

Variation in what usage preferences might be obtained are limited to:

- \* the choice of which methods for expressing preferences -- such as those defined in [ATTACH] -- are implemented and used, and
- \* the terms that are in the vocabulary, as might be added by updates to this document; see Section 6.4.

There is considerably more discretion involved in respecting preferences. An entity MAY choose to respect these preferences when processing assets. This is done according to both:

- \* an understanding of the nature of that processing and how it corresponds to the usage categories where preferences have been expressed, and
- \* the legal context that applies; see Section 8.

Usage preferences can be overridden through express agreements between relevant parties. There are also many situations where other priorities could override any usage preferences. For example, people with accessibility needs might override a preference to disallow AI Use (Section 4.4) so that they might access automated captions or summaries. Another case might involve the use of assets for research. Such overrides might be explicitly permitted in law, or could be based on the judgment of individual system users.

## 4. Vocabulary Definition

This section defines the categories of use in the vocabulary.

Figure 1 shows the relationship between these categories:

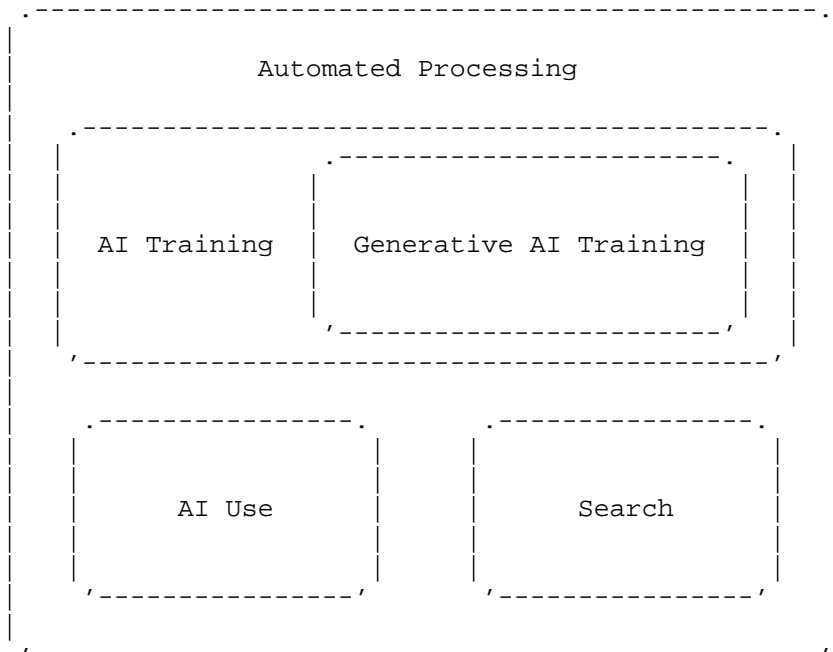


Figure 1: Relationship Between Categories of Use

#### 4.1. Automated Processing Category

The act of using one or more assets in the context of automated processing aimed at analyzing text and data in order to generate information which includes but is not limited to patterns, trends and correlations.

The use of assets for automated processing encompasses all the subsequent categories.

#### 4.2. AI Training Category

The act of training machine learning models or artificial intelligence (AI).

The use of assets for AI Training is a proper subset of Automated Processing usage.

#### 4.3. Generative AI Training Category

The act of training general purpose AI models that have the capacity to generate text, images or other forms of synthetic content, or the act of training more specialized AI models that have the purpose of generating text, images or other forms of synthetic content.

The use of assets for Generative AI Training is a proper subset of AI Training usage.

#### 4.4. AI Use Category

The act of using one or more assets as input to a trained AI/ML model as part of the operation of that model (as opposed to the training of the model).

The use of assets for AI Use is a proper subset of Automated Processing usage.

#### 4.5. Search Category

Using one or more assets in a search application that directs users to the location from which the assets were retrieved.

The purpose of defining a distinct Search category is to allow preferences to be expressed about search applications, independent of other categories of use. A distinct Search category allows for preferences specific to search applications, even if the use of AI is involved in their implementation.

The use of assets for Search is a proper subset of Automated Processing usage.

### 5. Usage

The vocabulary is used by referencing the terms defined in Section 4, directly or via mappings, in accordance with how they are defined in this document.

#### 5.1. More Specific Instructions

A recipient of a statement of preferences that follows this model might receive more specific instructions in two ways:

- \* Extensions to the vocabulary might define more specific categories of usage. Preferences about more specific categories override those of any more general category.

- \* Statements of preferences are general purpose, machine-readable statements that cannot override contractual agreements or more specific statements.

For instance, a statement of preferences might indicate that the use of an asset is disallowed for AI Training. If arrangements, such as legal agreements, exist that explicitly permit the use of that asset, those arrangements likely apply, unless the terms of the arrangement explicitly say otherwise.

## 5.2. Vocabulary Extensions

Systems referencing the vocabulary MUST NOT introduce additional categories that include existing categories defined in the vocabulary or otherwise include additional hierarchical relationships.

## 6. Exemplary Serialization Format

This section defines an exemplary serialization format for preferences. The format describes how the abstract model could be turned into Unicode text or sequence of bytes.

The format relies on the Dictionary type defined in Section 3.2 of [FIELDS]. The dictionary keys correspond to usage categories and the dictionary values correspond to explicit preferences, which can be either y or n; see Section 6.2.

For example, the following expresses a preference to allow AI training (Section 4.2), disallow generative AI training (Section 4.3), and expresses no preference for other categories other than subsets of these categories:

```
train-ai=y, train-genai=n
```

### 6.1. Usage Category Labels

Each usage category in the vocabulary (Section 4) is mapped to a short textual label. Table 1 tabulates this mapping.



Category	Label	Reference
Automated Processing	all	Section 4.1
AI Training	train-ai	Section 4.2
Generative AI Training	train-genai	Section 4.3
AI Use	ai-use	Section 4.4
Search	search	Section 4.5

Table 1: Mappings for Categories

Any mapping for a new usage category can only use lowercase latin characters (a-z), digits (0-9), "\_", "-", ".", or "\*". These are encoded using the mappings in [ASCII].

## 6.2. Preference Labels

The abstract model used has two options for preferences associated with each category: allow and disallow. These are mapped to single byte Tokens (Section 3.3.4 of [FIELDS]) of y and n, respectively.

## 6.3. Text Encoding

Structured Fields [FIELDS] describes a byte-level encoding of information, not a text encoding. This makes this format suitable for inclusion in any protocol or format that carries bytes.

Some formats are defined in terms of strings rather than bytes. These formats might need to decode the bytes of this format to obtain a string. As the syntax is limited to ASCII [ASCII], an ASCII decoder or UTF-8 decoder [UTF8] can be used. This results in the strings that this document uses.

Processing (see Section 6.5) requires a sequence of bytes, so any format that uses strings needs to encode strings first. Again, this process can use ASCII or UTF-8.

## 6.4. Syntax Extensions

There are two ways by which this syntax might be extended: the addition of new labels and the addition of parameters.

New labels might be defined to correspond to new usage categories. Section 5.2 addresses the considerations for defining new categories. New labels might also be defined for other types of extension that do not assign a preference to a usage category. In either case, when processing a parsed Dictionary to obtain preferences, any unknown labels MUST be ignored.

The Dictionary syntax (Section 3.2 of [FIELDS]) can associate parameters with each key-value pair. This document does not define any semantics for any parameters that might be included. When processing a parsed Dictionary to obtain preferences, any unknown parameters MUST be ignored.

In either case, new extensions need to be defined in an RFC that updates this document.

## 6.5. Processing Algorithm

To process a series of bytes to recover the expressed preferences, those bytes are parsed into a Dictionary (Section 4.2.2 of [FIELDS]), then preferences are assigned to each usage category in the vocabulary.

The parsing algorithm for a Dictionary produces a keyed collection of values, each with a possibly-empty set of parameters. The parsing process guarantees that each key has at most one value and parameters.

To obtain preferences for each of the categories in the vocabulary, iterate through the categories. For the label that corresponds to that category (see Table 1), obtain the corresponding value from the collection, disregarding any parameters. A preference is assigned as follows:

- \* If the value is a Token with a value of *y*, the associated preference is to allow that category of use.
- \* If the value is a Token with a value of *n*, the associated preference is to disallow that category of use.
- \* Otherwise, a preference is not expressed for that category of use.

Note that this last alternative includes the key being absent from the collection, values that are not Tokens, and Token values that are other than *y* or *n*. All of these are not errors, they only result in no preference being inferred.

An important note about this process and format is that, if the same key appears multiple times, only the last value is taken. This means that duplicating the same key could result in unexpected outcomes. For example, the following expresses no preferences:

```
train-ai=y, train-ai="n", train-genai=n, train-genai, all=n, all=()
```

If the parsing of the Dictionary fails, no preferences are expressed. This includes where keys include uppercase characters, as this format is case sensitive (more correctly, it operates on bytes, not strings).

This process produces an abstract data structure that assigns a preference to each usage category as described in Section 3.

#### 6.6. Alternative Formats

This format is only an exemplary way to represent preferences. The model described in Section 3, can be used without this serialization.

Any alternative format needs to define the mapping both from that format to the model used in this document and from the model to the alternative format. This includes any potential for extensions (Section 6.4).

The mapping between the model and the alternative format does not need to be complete, it only needs to be clear and unambiguous.

For example, an alternative format might only provide the ability to convey preferences for a subset of the categories of use. A mapping might then define that no preference is associated with other categories.

#### 7. Consulting a Preference Expression

After processing a preference expression (Section 6.5), an application can request the status of a specific usage category.

A single preference expression can be evaluated for a usage category as follows:

1. If the expression contains an explicit preference (either to allow or disallow), that is the result.
2. Otherwise, if the usage category is a proper subset of another usage category, recursively apply this process to that category and use the result of that process.

3. Otherwise, no preference is expressed.

This process results in three potential answers: allow, disallow, and no preference. Applications can use the answer to guide their behavior.

One approach for dealing with an "unknown" or "no preference" answer is to assign a default. This document takes no position on what default might be chosen as that will depend on policy constraints beyond the scope of this specification.

#### 7.1. Combining Preferences

The application might have multiple preference expressions, obtained using different methods.

If multiple preference expressions are active, all preference expressions are consulted (Section 7). This might result in conflicting answers.

Absent some other means of resolving conflicts, the following process applies to each usage category:

- \* If any preference expression indicates that the usage is disallowed, the result is that the usage is disallowed.
- \* Otherwise, if any preference preference allows the usage, the result is that the usage is allowed.
- \* Otherwise, no preference is expressed.

This process ensures that the most restrictive preference applies.

#### 8. Applicability and Legal Effect

This document provides a set of definitions for different categories of use, plus a system for associating simple preferences to each (allow, disallow, or no preference; see Section 3).

The categories of use that are defined as part of the vocabulary are not always clearly applicable or inapplicable to a particular system or application. The universe of possible systems is far more complex than any simple vocabulary is capable of describing. That means that some discretion could be involved in deciding whether a preference applies.

The expression of preferences might activate regulatory or legal consequences, which has implications for entities that consume those preferences. Their interpretation of the meaning of different terms could have legal ramifications. Different jurisdictions could reach subtly different conclusions about the applicability of each category of use to specific applications.

It is the responsibility of those that process affected assets to understand the legal implications of their use of digital assets.

This includes understanding:

- \* obligations regarding how preferences are obtained (in particular, which methods of associating preferences with content are expected to be understood),
- \* the specific uses to which assets are put,
- \* how preferences apply to the those uses, and
- \* how relevant jurisdictions might interpret those preferences.

These considerations will depend on jurisdiction and the details of the system.

## 9. Security Considerations

TODO Security

## 10. IANA Considerations

This document has no IANA actions.

## 11. References

### 11.1. Normative References

- [ASCII] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969, <<https://www.rfc-editor.org/rfc/rfc20>>.
- [FIELDS] Nottingham, M. and P. Kamp, "Structured Field Values for HTTP", RFC 9651, DOI 10.17487/RFC9651, September 2024, <<https://www.rfc-editor.org/rfc/rfc9651>>.

- [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>>.

## 11.2. Informative References

- [ATTACH] Illyes, G. and M. Thomson, Ed., "A Vocabulary For Expressing AI Usage Preferences", Work in Progress, Internet-Draft, draft-ietf-aipref-attach-02, 21 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-aipref-attach-02>>.
- [UTF8] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/rfc/rfc3629>>.

## Acknowledgments

The following individuals made significant contributions to this document:

- \* Cullen Miller
- \* Laurent Le Meur
- \* Leonard Rosenthol
- \* Sebastian Posth
- \* Timid Robot Zehta

## Authors' Addresses

Paul Keller  
Open Future  
Email: [paul@openfuture.eu](mailto:paul@openfuture.eu)

Martin Thomson (editor)  
Mozilla  
Email: [mt@lowentropy.net](mailto:mt@lowentropy.net)