The DE-BIAS API

Technical documentation

Image: Worldwide Wall - Associazione Culturale GoTellGo, Italy - CC BY

Document information

Current version	1.0
Contributors	Jason Liartis, Orfeas Menis (ThinkCode); Antoine Isaac, Hugo Manguinhas, Jochen Vermeulen (Europeana Foundation); Kristina Rose (DFF - Deutsches Filminstitut & Filmmuseum); Vassilis Tzouvaras (Datoptron); Kerstin Arnold (Archives Portal Europe)
History of changes		
Version	Change(s) applied	Contributors

Target audience

- Software developers and technical staff at cultural heritage institutions, aggregators, etc. interested in integrating the DE-BIAS API into their existing workflows of data processing.

Learning goals

- Understand the considerations behind the development of the DE-BIAS API
- Understand the requirements that the DE-BIAS API is meant to fulfil
- Understand the structure and parameters of a request to the DE-BIAS API (simple and detailed version)
- Apply that structure and parameters to one's own data
- Understand how to read the response of the DE-BIAS API (simple and detailed version)
- Apply the response of the DE-BIAS API to one's own data

Content note: Readers are advised that this document contains distressing words as part of the examples.

Table of contents

API request and response samples

This chapter illustrates the current status of the DE-BIAS API using a sample request for 'debiasing' multiple records and a series of sample responses covering the more detailed as well as the simpler version of the API. The DE-BIAS API can be accessed via the following URLs:

- Simple version at https://debias-api.ails.ece.ntua.gr/simple
- Detailed version at https://debias-api.ails.ece.ntua.gr

Simplified version of the API

The simplified version of the DE-BIAS API is not record-dependent and can be used to annotate/enrich collections/datasets where multiple records share common literals/values.

Request structure

```
HTTP: POST
```

```
{
  "language": <LANGUAGE>,
  "values": <VALUES>,
  "useNER": <USE_NER>,  // optional, can be omitted
  "useLLM": <USE_LLM>  // optional, can be omitted
}
```

The request structure for the simplified version of the DE-BIAS API consists of only two required parameters and two optional ones:

- <LANGUAGE>: A two-letter code according to ISO standard 639-1 indicating the language of the metadata values to be "debiased" (e.g. "en", "de", "fr", "it", "nl", ...)
- <VALUES>: A list of literals, e.g. descriptive metadata, to be processed by the API
- <USE_NER>: A boolean indicating whether to use Named Entity Recognition (NER) to filter out detected terms; the default value is "true", and can be set to "false"
- <USE_LLM>: A boolean indicating whether to filter terms that are only contextually contentious, by utilising a Large Language Model (LLM); the default value is "true", and can be set to "false"

3

Sample request

```
{
  "language": "en",
  "values": [ "A gouty man on a velocipede out shooting with his dogs, a
negro sniggers behind him. Engraving.", "This African rickshaw was invented
by a Swedish missionary and has been of great value in travelling, through
the continent." ],
  "useNER": false
}
```

Response structure

```
{
  "metadata": {
    "annotator": <ANNOTATOR>
    "thesaurus": <THESAURUS> //only one thesaurus, the DE-BIAS vocabulary,
     is supported at the moment so the value is always "null"
    "date": <DATE_TIME> }
  "results": [
    {
      "language": <LANGUAGE>,
      "literal": <LITERAL_VALUE>,
      "tags": [

        {
          "start": <STARTING_POSITION>,
          "end": <ENDING_POSITION>,
          "length": <LENGTH>,
          "uri": <VOCABULARY_URI>,
          "literal": <TERM_LITERAL>,
          "issue": <TERM_ISSUE>
        },
        {
          "start": ...,
          "end": ...,
          "length": ...,
          "uri": ...,
          "literal": ...,
          "issue": ...
        }
      ] //tags
    } // one result
  ] // results
} // outer object
```

The response structure for the simplified version of the DE-BIAS API includes the following parameters:

- <ANNOTATOR>: An identifier of the tool that was used to create this annotation
- <THESAURUS>: An identifier of the thesaurus that was utilised by the tool to create the annotation. Since the tool only supports the [DE-BIAS vocabulary](#) in its current version, this value is always null and should be ignored
- <DATE_TIME>: A timestamp indicating the time of the creation of the annotation
- <LANGUAGE>: The two-letter code according to ISO standard 639-1 indicating the language of the metadata values that have been 'debiased' (e.g. "en", "de", "fr", "it", "nl", …)
- <LITERAL_VALUE>: The text, i.e. descriptive metadata, that was processed by the API and to which the following tags, i.e. annotations, refer
- <STARTING_POSITION>: The starting position, i.e. character, in the text of the first detected term to which the specific annotation refers
- <ENDING_POSITION>: The ending position, i.e. character, in the text of the first detected term to which the specific annotation refers
- <LENGTH>: The length of the detected bias term
- <VOCABULARY_URI>: The vocabulary URI of the detected bias term
- <TERM_LITERAL>: The term literal as it appears in the vocabulary (encoded as skosxl:literalForm)
- <ISSUE>: The description of the issue associated with the term in the vocabulary (encoded as debias-o:hasContentiousIssue in the term's data respectively as dcterms:description in the description's data).

Sample response

```
{
  "metadata": {
    "annotator": "debias"
    "thesaurus": "null"
    "date": 2024-12-17T12:00:00Z }
  "results": [
    {
      "language": "en",
      "literal": "A gouty man on a velocipede out shooting with his dogs,
       a negro sniggers behind him. Engraving.",
```

```
  "type": "Request",
  "params": {
    "limitPerPredicate": <LIMIT>, // optional, can be omitted
    "language": <LANGUAGE>,
    "provenance": <PROVENANCE>,
    "useNER": <USE_NER>,  // optional, can be omitted
    "useLLM": <USE_LLM>   // optional, can be omitted
  },
  "totalItems": <TOTAL>,
  "items" : [
    {
    "id": <RECORD_ID>,
    "<FIELD_NAME>": [ <FIELD_VALUE>, ... ]
      ...
    },
    ...
  ]
}
```

The request structure for the detailed version of the DE-BIAS API consists of the following parameters:

- <CONTEXT_URI>: The URI of the JSON-LD context
- <LIMIT>: The number of mentions of the same bias term to be returned per metadata field (e.g. EDM property) per record (optional, can be omitted to get all occurrences of all detected terms)
- <LANGUAGE>: A two-letter code according to ISO standard 639-1 indicating the language of the metadata values to be 'debiased' (e.g. "en", "de", "fr", "it", "nl", …)
- <PROVENANCE>: A boolean value to inform the tool to return or not the provenance information associated with the annotation, meaning "created" and "creator" fields
- <USE_NER>: A boolean indicating whether to use Named Entity Recognition (NER) to filter out detected terms; the default value is "true", and can be set to "false"
- <USE_LLM>: A boolean indicating whether to filter terms that are only contextually contentious, by utilising a Large Language Mode (LLM); the default value is "true", and can be set to "false"
- <TOTAL>: An optional field to indicate the total number of items, i.e. in most cases metadata records, that are part of the request
- <RECORD_ID>: The local identifier of the metadata record, in relation to the "@base", or otherwise, the complete URI of the record within the relevant namespace (e.g. the "data.europeana.eu" namespace)

- <FIELD_NAME>: The namespace prefixed name of the metadata field (e.g. "dc:title" when operating with EDM properties)
- <FIELD_VALUE>: The literal value that was associated with the metadata field; more than one value may be supplied as an array

Sample request

```
{
  "@context": [
    "http://www.w3.org/ns/anno.jsonld",
    "https://www.europeana.eu/schemas/context/edm.jsonld",
    { "@base": "http://data.europeana.eu/item/" }
  ],
  "type": "Request",
  "params": {
    "limitPerPredicate": 1,
    "language": "en",
    "provenance": false,
  },
  "totalItems": 2,
  "items" : [
    {
      "id": "12345/XPTO",
      "dc:title": [ "a sample title", "a second sample title" ],
      "dc:description": [ "a sample description", "a second sample
       description" ]
      ...
    },
    {
      "id": "12345/XPTO_2",
      "dc:title": [ "another sample title", "another second sample title"
       ],
      "dc:description": [ "another sample description", "another second
       sample description" ]
      ...
    },
    ...
  ]
}
```

JSON-LD context (compatible with the W3C Annotation Model)

```
"as": "https://www.w3.org/ns/activitystreams#",
    "dc": "http://purl.org/dc/terms/",
    "dce": "http://purl.org/dc/elements/1.1/",
    "foaf": "http://xmlns.com/foaf/0.1/",
    "oa": "http://www.w3.org/ns/oa#",
    "rdf": "http://www.w3.org/1999/02/22-rdf-syntax-ns#",
    "xsd": "http://www.w3.org/2001/XMLSchema#",
    "soa": "http://sw.islab.ntua.gr/annotation/",
    "id": {
      "@id": "@id",
      "@type": "@id"
    },
    "type": {
      "@id": "@type",
      "@type": "@id"
    },
    "value": "rdf:value",
    "created": {
      "@id": "dc:created",
      "@type": "xsd:dateTime"
    },
    "creator": {
      "@id": "dc:creator",
      "@type": "@id"
    },
    "language": "dc:language",
    "Software": "as:Application",
    "name": "foaf:name",
    "Annotation": "oa:Annotation",
    "TextPositionSelector": "oa:TextPositionSelector",
    "RDFPropertySelector": "soa:RDFPropertySelector",
    "TextualBody": "oa:TextualBody",
    "body": {
      "@id": "oa:hasBody",
      "@type": "@id"
    },
    "scope": {
      "@id": "oa:hasScope",
      "@type": "@id"
    },
```

```
    "selector": {
      "@id": "oa:hasSelector",
      "@type": "@id"
    },
    "source": {
      "@id": "oa:hasSource",
      "@type": "@id"
    },
    "target": {
      "@id": "oa:hasTarget",
      "@type": "@id"
    },
    "Literal": "soa: Literal"
```

```
Content type: application/json
```

Response structure

```
{
  "@context": [
    "http://www.w3.org/ns/anno.jsonld",
    "https://www.europeana.eu/schemas/context/edm.jsonld",
    "@base": "http://data.europeana.eu/item/",
  },
  "type": "AnnotationPage",
  "partOf": {
    "type": "AnnotationCollection",
    "modified": <DATE_TIME>
  },
  "items" : [
    <ANNOTATION>,
    ...
  ]
}
```

The response structure for the detailed version of the DE-BIAS API includes the following parameters:

- <DATE_TIME>: A timestamp indicating the time of the creation of the annotation
- <ANNOTATION>: The detailed annotation following the structure shown below

```
{
  "id": "http://example.org/anno29",
  "type": "Annotation",
  "motivation": "highlighting",
  "body": <VOCABULARY_URI>,
  "target": [
    {
      "source": <RECORD_ID>,
      "selector": {
        "type" : "RDFStatementSelector",
        "hasPredicate": "<FIELD_NAME>",
        "refinedBy" : {
          "type" : "TextQuoteSelector",
          "exact": {
            "@value": <BIASED_TERM>,
            "@language": <LANGUAGE>
          },
          "prefix": <PREFIX>
          "suffix": <SUFFIX>
        }
      }
    },
    {
      "source": ...,
      "selector": {
        "type" : "RDFStatementSelector",
        "hasPredicate": ...,
        "refinedBy" : {
          "type" : "TextQuoteSelector",
          "exact": {
            "@value": ...,
            "@language": ...
          },
          "prefix": ...,
          "suffix": ...
        }
      }
    },
    ...
  ]
  ...
}
```

The "prefix" and "suffix", containing the text before and after the detected bias term, should also be stated unless the text fragment including the contentious term is at the start or at the end of the value. The character length of both prefix and suffix (separately) should be a minimum of 50 characters, as applicable. If the 50th character happens to be in the

middle of a word, then the complete word should be included respecting this way the 50 characters minimum.

Sample response

```json
{
  "id": "http://example.org/anno29",
  "type": "Annotation",
  "motivation": "highlighting",
  "body": "http://debias.example.org/term/xpto",
  "target": [
    {
      "source": "http://data.europeana.eu/item/some_cho",
      "selector": {
        "type" : "RDFStatementSelector",
        "hasPredicate": "dc:description",
        "refinedBy" : {
          "type" : "TextQuoteSelector",
          "exact": {
              "@value": "bias term",
              "@language": "en"
          },
          "prefix": "A description with a "
          "suffix": "Also, the "
        }
      }
    },
    {
      "source": "http://data.europeana.eu/item/some_cho",
      "selector": {
        "type" : "RDFStatementSelector",
        "predicate": "dc:title",
        "refinedBy" : {
          "type" : "TextQuoteSelector",
          "exact": {
            "@value": "bias term",
            "@language": "en"
          },
          "prefix": "The title contains also the same ",
        }
      }
    },
    ...
  ]
  ...
}
```

Background on the development of the API endpoint

Operational decisions taken

Based on the designs for the display of detected bias terms on Europeana.eu and a preliminary review of the DE-BIAS vocabulary, the following general decisions were taken for the development of the API to ensure a smooth integration of the DE-BIAS API with the core service platform of Europeana and to be certain about the resulting technical implications for the processing activity defined as part of the DE-BIAS project:

- **Representing bias terms as annotations (general)**
  - Due to the complexity and nature of the bias use case, which goes beyond traditional metadata, it was decided that the detected bias terms will be represented as annotations, both on the DE-BIAS tool and in Europeana.
  - As a consequence of this, a model representation for the annotations will be used that both systems can implement; this model representation will be added as a new application scenario for the Europeana Data Model (EDM) Annotations Profile.
- **Serving bias annotations for display on Europeana.eu (only for application to the core service platform of Europeana)**
  - The Europeana Annotations API will be used to store the bias annotations and serve them to the Europeana.eu website for display; the Annotations API will hence be extended to support bias annotations.
  - To guarantee that the data for bias terms are always fresh, the bias annotations that are returned by the DE-BIAS tool and stored in the Annotations API will not retain the metadata that is associated with a bias term and its contentious issue description, but only its URI.
- **Using the Metis dereference service (only for application to the core service platform of Europeana)**
  - When serving the data from the Annotations API to Europeana.eu, the API will use the Metis dereference service to dynamically obtain the data from the source DE-BIAS vocabulary (using its URI).
  - The Metis dereference service will be configured to support the DE-BIAS vocabulary (and implement a mapping, if needed, from the source format to skos:Concept).

- **Processing of selected data sets for validation and evaluation (only during the project lifetime and for application to the core service platform of Europeana)**
  - The processing tool developed by Europeana Foundation for the project must copy over the annotations to the Annotations API. In the context of the new post-publication pipeline, the DE-BIAS tool will then be applied to the minimum 4.5 million records that constitute the project's KPI for validation and evaluation.
  - As the bias terms will live separately from the metadata, there will be no need to do an official reindex in order to import the bias terms to the data.

Requirements for display and integration

Following up from the operational decisions listed above, a set of functional and non-functional requirements was defined for the model representation for the annotations and the API endpoint. They reflect design as well as technical considerations necessary for the integration into the common European data space for cultural heritage.

Functional requirements

The following functional requirements apply to each metadata record and the creation of annotations:

R1 - Term in a specific language: These are the word(s) that make up the bias term itself as defined in the DE-BIAS vocabulary.

R2 - Accompanying text: This is the contentious issue description of the bias term as detailed in the DE-BIAS vocabulary.

R3 - The fields and location within the text where the bias term was found: Given that there may be more than one mention of the same bias term for a specific metadata field (e.g. the description of the cultural heritage object), only the first mention of the bias term should be returned (even if the field has more than one literal value, e.g. a description with several paragraphs). This is to not overload the display with highlights.
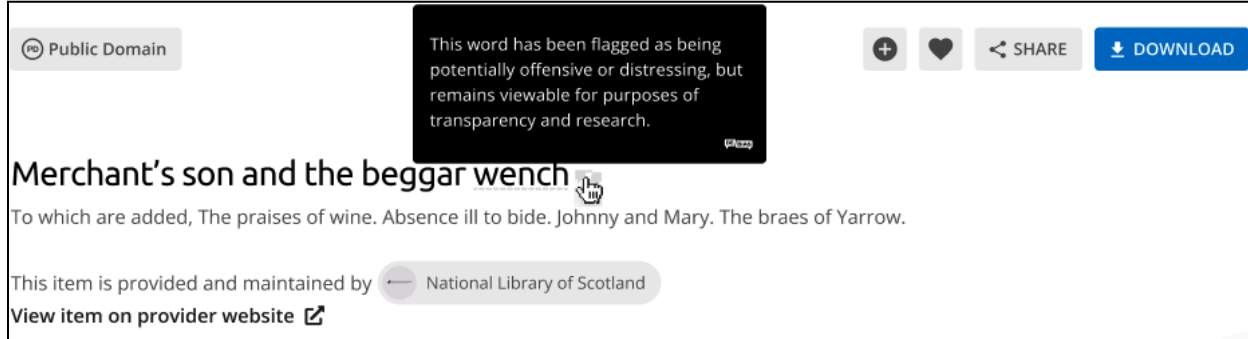
Fig. 1 - Wireframe screenshot for the display of bias terms on Europeana.eu

Non-functional requirements

The following requirements apply to the execution of the DE-BIAS API:

R4 - Efficiency: The data that is sent in the request and returned in the response must be optimised to take the least bandwidth for the use case at hand.

R5 - Performance: The tool must exceed a minimum throughput of 105 rpm (records per minute). The minimum rpm was calculated on the basis of meeting the KPI of the DE-BIAS project of at least 4.5 million records and an overall processing duration of 30 days. The rpm estimate is an absolute minimum considering that there may be downtime (even if short) and that the records to be analysed might vary in size. Finally, in order for the tool to be approved for formal integration, a much higher rpm is required as otherwise it would make it unsustainable to be applied on a regular basis to metadata records published on Europeana.eu.

R6 - Availability: The tool must meet an uptime of 99.8%. This requirement goes hand in hand with the previous requirement as otherwise it won't allow the DE-BIAS project meeting the processing timeline of 30 days. Additionally, any tool integrated with the data space must comply with the non-functional requirements of the data space set out by the European Commission[1].

---

[1] See the Tender specifications for the Deployment of a common European data space for cultural heritage: https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/tender-details/docs/etender/9835/9835_129416_EN-CNECT_LUX_2021_OP_0070%20Tender_specifications_Common%20European%20data%20space_001_ENG_V2.pdf and the Europeana Foundation's Playbook for software development and integration: https://drive.google.com/file/d/12h1O6OLDawoua1pWHlsq5_EI0D12k7T5/view (both last accessed: 7 February 2025).

Model representation for annotations

This section describes the model used to represent annotations as part of the output of the DE-BIAS API. The annotation intends to highlight the mention of the detected bias term in the metadata and associate it to a term from the DE-BIAS vocabulary (in its form of a knowledge graph). The same term or variations of it may occur in multiple places in the same metadata record.

DE-BIAS reuses an application of the W3C Annotation Model that was originally created in the scope of the [CRAFTED project](#) (funded under the Connecting Europe Facility of the European Union between September 2021 and June 2023 as a [Europeana Generic Services](#) project), which also serves DE-BIAS' purpose.

Guidelines

The following guidelines apply:

- An annotation should represent all the occurrences of exactly one bias term in a metadata record. In the specific context of using the DE-BIAS API on data prepared for publication on Europeana.eu, this means the edm:ProvidedCHO. In accordance with R3 (see chapter on [Functional requirements](#) above), the API may limit the number of occurrences represented.
- The "motivation" of the annotation must be "highlighting".
- The "body" of the annotation must be the URI of the detected bias term.
- The annotation must have one or more targets. Each target points to one occurrence of the bias term in the record metadata (e.g. the edm:ProvidedCHO). Each target:
    - Must be a specific resource where the "source" contains the URI of the metadata record (e.g. the edm:ProvidedCHO).
    - It must have one "selector" of "type" "RDFStatementSelector".
    - The selector must give the URI of the metadata field where the biassed term has been detected as "hasPredicate".
    - The selector must be "refinedBy" a "TextQuoteSelector" as follows:
        - It must have one "exact" property, containing the term as it appears in the metadata I.e. this might be the exact term as specified in the DE-BIAS vocabulary or any variation of the term as detected by the DE-BIAS API. If the metadata field where the biassed term was found contains a language tag (e.g. @xml:lang) then the value of the "exact" property must also have that same language tag.

> - The "prefix" and "suffix", containing text before and after the detected bias term, should also be stated unless the text fragment including the contentious term is at the start or at the end of the value. The character length of both prefix and suffix (separately) should be a minimum of 50 characters, as applicable. If the 50th character happens to be in the middle of a word, then the complete word should be included respecting this way the 50 characters minimum.

Rationale

- It was preferred to cluster all mentions of the same bias term in a single annotation in order to optimise the representation without losing semantics and compliance with the Web Annotation data model.
- The TextQuoteSelector was chosen to reduce the verbosity of the annotation while, at the same time, giving a relatively good perception of where the term was found.

Relevant specifications for the DE-BIAS use case can be found here:

- [RDFStatementSelector](#)
- [TextQuoteSelector](#)

Examples

Example 1

The following example is an edm:ProvidedCHO to be processed for 'debiasing'. Note that this example does not follow the syntax to be used in requests to the DE-BIAS API, but is mainly used for explanation purposes to illustrate the model representation.

```
{
  "id": "http://data.europeana.eu/item/some_cho",
  "type": "ProvidedCHO",
  "dc:description": {
      "@value": "A description with a bias term. Also, the bias term appears
twice.",
      "@language": "en"
  },
  "dc:title": {
      "@value": "The title contains the same bias term.",
      "@language": "en"
  },
  "dc:description": {
      "@value": "The second description also contains the same bias term.",
      "@language": "en"
  },
  ...
}
```

Example 2

An example of an annotation returned by the DE-BIAS API on a single edm:ProvidedCHO that contained a bias term in two fields (shown in Example 1 above). The annotation only highlights the first occurrence of the bias term per field, i.e. it omits the repeated mention of the bias term in the first dc:description as well as the mention of the same bias term in the second dc:description.

```
{
  "@context": [
    "http://www.w3.org/ns/anno.jsonld",
    "https://www.europeana.eu/schemas/context/edm.jsonld"
  ],
  "id": "http://example.org/anno29",
  "type": "Annotation",
  "motivation": "highlighting",
  "body": "http://debias.example.org/term/xpto",
  "target": [
    {
      "source": "http://data.europeana.eu/item/some_cho",
      "selector": {
        "type" : "RDFStatementSelector",
        "hasPredicate": "dc:description",
        "refinedBy" : {
          "type" : "TextQuoteSelector",
          "exact": {
            "@value": "bias term",
            "@language": "en"
          },
          "prefix": "A description with a ",
          "suffix": ". Also, the "
        }
      }
    },
    {
      "source": "http://data.europeana.eu/item/some_cho",
      "selector": {
        "type" : "RDFStatementSelector",
        "hasPredicate": "dc:title",
        "refinedBy" : {
          "type" : "TextQuoteSelector",
          "exact": {
            "@value": "bias term",
            "@language": "en"
          },
          "prefix": "The title contains the same "
        }
      }
    },
    ...
  ]
  ...
}
```

Requirements for the API

In combining all of the above, the following requirements have been defined for the API itself (non-functional) and to requests handled by and responses sent from the API.

19

Non-functional requirements for the API

The following should be considered in order to meet non-functional requirements:

- Support sending (at least) 100 records at a time
- Support HTTP compression
- Support Keep-Alive header or otherwise support HTTP/2 or /3
- Scale-up server instances to increase the overall rpm rate

Requirements for requests handled by the API

The following requirements apply to requests to be handled by the API:

- The API should be able to receive a list of records to be subject to 'debiasing'.
- The API should be able to receive a subset of the metadata of a record in one language. Only literal values are relevant to be 'debiased'.
- A single language must be indicated in the request that applies to all literal values of all records in the set that is being sent.
- By default, one mention per bias term should be returned per metadata field. To make it flexible enough to be applied in other situations within the data space than for Europeana.eu, a limit for the number of mentions could be indicated in the request.

The following snippet represents the syntax for requests:

```
{
  "@context": {
    <CONTEXT_URI>,
    "@base": "http://data.europeana.eu/item/"
  },
  "type": "Request",
  "params": {
    "limitPerPredicate": <LIMIT>,  // optional, can be omitted
    "language": <LANGUAGE>,
    "provenance": <PROVENANCE>
    "useNER": <USE_NER>,  // optional, can be omitted
    "useLLM": <USE_LLM>   // optional, can be omitted
```

```
    },
"totalItems": <TOTAL>,
    "items" : [
        {
            "id": <RECORD_ID>,
            "<FIELD_NAME>": [ <FIELD_VALUE>, ... ]
            ...
        },
        ...
    ]
}
```

The keywords in the snippet above mean the following:

- <CONTEXT_URI>: The URI of the JSON-LD context
- <LIMIT>: The number of mentions of the same bias term to be returned per metadata field (e.g. EDM property) per record (optional, can be omitted to get all occurrences of all detected terms)
- <LANGUAGE>: A two-letter code according to ISO standard 639-1 indicating the language of the metadata values to be 'debiased' (e.g. "en", "de", "fr", "it", "nl", ...)
- <PROVENANCE>: A boolean value to inform the tool to return or not the provenance information associated with the annotation, meaning "created" and "creator" fields
- <USE_NER>: A boolean indicating whether to use Named Entity Recognition (NER) to filter out detected terms; the default value is "true", and can be set to "false"
- <USE_LLM>: A boolean indicating whether to filter terms that are only contextually contentious, by utilising a Large Language Mode (LLM); the default value is "true", and can be set to "false"
- <TOTAL>: An optional field to indicate the total number of items, i.e. in most cases metadata records, that are part of the request
- <RECORD_ID>: The local identifier of the metadata record, in relation to the "@base", or otherwise, the complete URI of the record within the relevant namespace (e.g. the "data.europeana.eu" namespace)
- <FIELD_NAME>: The namespace prefixed name of the metadata field (e.g. "dc:title" when operating with EDM properties)
- <FIELD_VALUE>: The literal value that was associated with the metadata field; more than one value may be supplied as an array

Example

```
{
  "@context": {
    "http://www.w3.org/ns/anno.jsonld",
    "https://www.europeana.eu/schemas/context/edm.jsonld",
    "@base": "http://data.europeana.eu/item/"
  },
  "type": "Request",
  "params": {
    "limitPerPredicate": 1,
    "language": "en",
    "provenance": false
  },
  "totalItems": 2,
  "items" : [
    {
      "id": "12345/XPTO",
      "dc:title": [ "a sample title", "a second sample title" ],
      "dc:description": [ "a sample description", "a second sample description" ]
      ...
    },
    {
      "id": "12345/XPTO_2",
      "dc:title": [ "another sample title", "another second sample title" ],
      "dc:description": [ "another sample description", "another second sample
       description" ]
      ...
    },
    ...
  ]
}
```

Requirements for responses sent by the API

The following requirements apply to responses sent by the API:

- List of highlighted bias terms for all records. Each item must be an annotation that concerns the highlighting of a bias term within a metadata record and contain at least one (or the limit specified in the request parameter) highlight of the term per metadata field.
- Only the term URI needs to be exposed, as the additional metadata for the contentious term - such as the contentious issue description, the suggestion description, the suggested alternative terms if available, etc. - is maintained separately and can be gathered once connected to the term URI.
- The mention(s) of where the bias term was found (across all metadata fields) should be listed together within the same annotation.

- The response must be optimised to reduce bandwidth.

The following snippet represents the syntax for responses:

```
{
  "@context": [
    "http://www.w3.org/ns/anno.jsonld",
    "https://www.europeana.eu/schemas/context/edm.jsonld",
    "@base": "http://data.europeana.eu/item/",
  },
  "type": "AnnotationPage",
  "partOf": {
    "type": "AnnotationCollection",
    "total": 100,
    "modified": "2024-12-19T12:00:00Z"
  },
  "items" : [
    <ANNOTATION>,
    ...
  ]
}
```

The keyword <ANNOTATION> in the snippet above means the annotation as defined in the previous chapter about the Model representation for annotations.

Rationale

- The Annotation Page from the Web Annotation Data Model was used to group / return the list of annotations that are produced by the DE-BIAS API.
- If necessary an artificial Annotation Collection could be added to the Annotation Page using the "partOf" property, in order to indicate properties that apply to the whole collection such as the date that the annotations were generated (i.e. "modified") or provenance (when applicable to all annotations).