

# SWITCH Innovation Lab

## “Image to Concept”

### Final Project Report

Beat Estermann, Sebastian Berlin, Andrea Bertino, André Costa, Florence Devouard, Eugene Egbe, Navino Evans, Matthias Ruediger, Sebastian Sigloch

Version 1.1

6 October 2023

In cooperation with:

**SWITCH**



### **Image on the cover page**

Homepage image of the ISA Tool by Islahaddow, using File:2017 06 Ali- 00213.jpg by user:Alimdaihl, [CC BY-SA 4.0 \(Wikimedia Commons\)](#)

### **Copyright license**

The present report is made available under the CC BY 4.0 ([Creative Commons Attribution 4.0 International](#)) copyright license.

# Executive Summary

By means of its “Research Data Connectome”, SWITCH, the Swiss national infrastructure service provider for higher education and research, seeks to connect open research using linked open data technologies. The goal is to make the data accessible, interoperable and valuable for research, education and innovation. In order to kick-start the development of new services, SWITCH carries out so-called “InnoLab” projects which have an experimental character and are geared towards generating quick learnings.

The present InnoLab project brought together researchers and software developers from SWITCH, Wikimedia Sverige, the University of Applied Sciences of the Grisons as well as the Bern University of Applied Sciences. The goal was to develop a microservice that supports the semi-automatic tagging of images in order to interlink them with concepts on Wikidata. It thus facilitates the search and discovery of relevant images by researchers and other interested parties. The microservice builds upon an existing crowd-sourcing tool, the ISA Tool, that has been deployed on Wikimedia Commons in 2019 where it is used to apply “depicts” statements describing the content of images stored in the free media repository.

The semi-automatic tagging functionality added to the ISA Tool in the course of the present project relies on two distinct algorithms: One of them is used to extract entities from the image itself. For this purpose, the Google Cloud Vision service available on Wikimedia Commons is used. The other one extracts entities from the image metadata, thus leveraging earlier efforts made to describe the content of the images. At the time of writing, the enhanced version of the ISA Tool is available in the test environment and can be used to add “depicts” statements to images on Wikimedia Commons. Plans to deploy it to production have been postponed due to several remaining bugs.

The key learnings gained in the course of the project can be summarized as follows:

- There are several issues that need to be tackled to allow for wider use and promotion of the ISA Tool: performance issues, reliability issues, improvement of multilingual support.
- Once these issues have been resolved, measures should be taken to increase the visibility and take-up of the tool among potential contributors. As an accompanying measure, it would be advisable to assess and monitor the relevance of the ISA Tool in comparison to other tools and methods employed to add Structured Data on Commons. Moreover, activities to further promote the tool among the volunteer community should be accompanied by a dialogue with various stakeholders on what constitutes “good” tagging of images.
- The algorithms used for semi-automatic tagging should be further improved and/or complemented; a variety of avenues to be pursued to this effect have been suggested.
- Research use cases in the context of the SWITCH Research Data Connectome should be facilitated by developing alternatives to the current requirement of uploading all media files to Wikimedia Commons. Some initial use cases have been identified in the areas of digital humanities, medical libraries etc.
- Requirements arising from research use cases making use of “depicts” statements beyond their current use for search and discovery should be further investigated.
- If the ISA Tool is to be used on a large scale in the context of the SWITCH Research Data Connectome, the conclusion of contractual agreements with service providers may be indicated. Roles and responsibilities with regard to deployment, operations and maintenance need to be clarified.

# Table of Contents

1	Introduction.....	5
2	Description of the Technical Solution .....	7
2.1	ISA Tool.....	10
2.2	Machine Vision Service (Google Cloud Vision).....	11
2.3	Metadata-to-Concept Module .....	11
2.4	Maintenance of the Tools .....	12
3	Expected Use .....	12
3.1	Usage of the ISA Tool .....	12
3.2	Current Applications Making Use of “depicts” Statements .....	13
3.3	Further Use Cases Making Use of “depicts” Statements .....	13
4	Added Value for the Connectome Project .....	14
4.1	Image and Context Metadata for the Linked Data Pipeline .....	14
4.2	Enrichment of Image Metadata from Swiss Projects in Relation to Wikimedia Commons .....	14
4.3	Enrichment of Image Metadata from Swiss Libraries, Repositories and Archives.....	15
4.4	Open Questions.....	15
5	Current Limitations and Possible Next Steps .....	16
5.1	ISA Tool.....	16
5.2	Metadata-to-Concept Module .....	16
5.3	Machine Vision Extension .....	18
5.4	Media Repository .....	19
5.5	Knowledge Graph.....	19
6	Monitoring and Evaluation.....	19
7	Conclusions and Outlook.....	21
7.1	Further Improvement and Promotion of the ISA Tool .....	22
7.2	Further Development of the Algorithms for Semi-Automatic Tagging.....	22
7.3	Facilitation of Research Use Cases .....	23
7.4	Contractual Aspects.....	24
8	List of Annexes.....	25

# 1 Introduction

The goal of the SWITCH Connectome project<sup>1</sup> is to connect open research data by means of linked open data technologies to make it accessible, interoperable and valuable for research, education and innovation. In order to improve the data enhancement and aggregation process, SWITCH supports the implementation of a series of so-called “InnoLab” projects the goal of which is to develop prototypical solutions of data enrichment services for the Research Data Connectome.

For the present InnoLab project, which was implemented between July 2022 and May 2023, the national infrastructure service provider for the academic sector in Switzerland teamed up with the Bern University of Applied Sciences, the University of Applied Sciences of the Grisons, and Wikimedia Sverige to develop an **Image-to-Concept Microservice** which supports the semi-automatic tagging of images on Wikimedia Commons by interlinking them to concepts on Wikidata, indicating what is “depicted” in the respective image. It thereby makes use of two distinct approaches to retrieve said concepts:

1. **Entity extraction from the image itself:** For this purpose, the Google Cloud Vision service available on Wikimedia Commons is used (Wikidata-based tag suggestions are provided through a publicly documented, open API).
2. **Entity extraction from the image metadata:** For this purpose, a “Metadata-to-Concept” module was developed as part of the InnoLab project.

The two approaches were combined in a single crowdsourcing tool – the **ISA Tool**<sup>2</sup> – a pre-existing tool which was enhanced in the context of the InnoLab project so that it can be combined with different entity extraction modules and is able to support machine learning approaches drawing on the information generated by the users of the tool. To facilitate the further development of the tool, the software architecture was designed in such a way as to allow for the replacement of the different entity extraction modules by other (open source) services. Data between the ISA Tool and the entity extraction modules are exchanged via publicly documented, open APIs.

The enhanced version of the ISA Tool has been made available in a test environment<sup>3</sup>, where the semi-automatic tagging feature has been available by default since May 2023. Its deployment to production has been postponed due to some remaining bugs and after initial deployment plans were stalled due to prolonged maintenance issues of the production version between May and August 2023.

SWITCH aims to utilize this “Image-to-Concept Microservice” as a Data Enrichment module as part of the Linked Data Pipeline, enriching image metadata from given data providers as part of the integration process for the Connectome Knowledge Graph. These metadata will be reused through the SWITCH Open Data Navigator<sup>4</sup> and the Connectome API<sup>5</sup>.

---

<sup>1</sup> <https://www.switch.ch/connectome/>

<sup>2</sup> [https://commons.wikimedia.org/wiki/Commons:ISA\\_Tool](https://commons.wikimedia.org/wiki/Commons:ISA_Tool)

<sup>3</sup> <https://isa-dev.toolforge.org/>

<sup>4</sup> <https://opendatanavigator.switch.ch>

<sup>5</sup> <https://opendatanavigator-test.switch.ch/api/graphql> (currently in development)

The present project report is structured as follows:

- Table 1 below gives an overview of the project team members, their affiliation and their role in the project.
- Section 2 contains a description of the technical solution of the Image-to-Concept Microservice and the enhancements that were made to the ISA Tool in the course of the InnoLab project.
- Section 3 describes the expected use of the Image-to-Concept Microservice in the form of short use case descriptions, complemented by current usage statistics of the ISA Tool in the context of Wikimedia Commons.
- Section 4 describes a possible test implementation of the Image-to-Concept Microservice in the Connectome Pipeline and highlights different options for the roll-out of the ISA Tool in various contexts.
- Section 5 lists the known issues and assesses possible avenues for the further development of the Image-to-Concept Microservice.
- Section 6 gives an overview of the project team’s reflections with regard to the continuous monitoring and the evaluation of the Image-to-Concept Microservice.
- Section 7 comprises some concluding remarks and an outlook on the next steps.

**Table 1:** Project Team (in alphabetical order)

Name	Organization	Role
Andrea Bertino	SWITCH	Integration with the Research Data Connectome
Sebastian Berlin	Wikimedia Sverige	Software Development
André Costa	Wikimedia Sverige	Coordination on the side of Wikimedia Sverige
Florence Devouard	Wikimedia Sverige (contractor)	Community Outreach
Eugene Egbe	Wikimedia Sverige (contractor)	Software Development
Beat Estermann	University of Applied Sciences of the Grisons	Project Coordination
Navino Evans	Wikimedia Sverige (contractor)	Software Project Manager / Software Architect
Matthias Ruediger	Bern University of Applied Sciences	Software Development
Sebastian Sigloch	SWITCH	Coordination on the side of SWITCH

## 2 Description of the Technical Solution

This section contains a short description of the Image-to-Concept Microservice and the enhancements that were made to the ISA Tool in the course of the InnoLab project. It also describes the current maintenance situation of the ISA Tool and the other components of the Microservice. Iteration reports detailing the software development tasks and the user tests carried out in the course of the project can be found in Annex 1. Further documentation of the ISA Tool can be found on the tool's homepage on Wikimedia Commons<sup>6</sup>. A description of known issues and prospects for the further development of the tool can be found in section 5.

The **Image-to-Concept Microservice** currently consists of three components that interact with Wikimedia Commons and Wikidata (see figure 1):

- The **ISA Tool** is a web application for the tagging of images. The tagging of the images by means of the ISA Tool typically takes place in the context of crowdsourcing campaigns. For this purpose, a set of images from Wikimedia Commons is assigned to a campaign defined in the ISA Tool by a campaign organizer, upon which users are invited to browse through the images and to apply tags that describe what is “depicted” on a given image. The ISA Tool is currently set up for the use in combination with **Wikimedia Commons** (media repository) and **Wikidata** (knowledge graph). This means that all the images to be processed by the tool previously need to be uploaded to Wikimedia Commons. Furthermore, all the tags that are applied to a given image need to have their corresponding Wikidata item. Thus, if a specific named entity or a concept is missing on Wikidata, it needs to be added to Wikidata before it can be used in the context of the ISA Tool.
- **Google Cloud Vision**<sup>7</sup> is an image analysis service based on machine learning. In the context of **Wikimedia Commons**, Google Cloud Vision has been set up to power the machine vision extension<sup>8</sup>. It automatically provides a probabilistic interpretation of what is depicted in a given image, based on an analysis of the image itself. Suggested labels from Google Cloud Vision are mapped to **Wikidata** items. The service can be used via an open API.
- The **Metadata-to-Concept Module**<sup>9</sup> is an entity extraction service (currently using an algorithm that does not rely on machine learning) which provides an interpretation of what is depicted in a given image, based on an analysis of the descriptive metadata. Suggested labels come in the form of **Wikidata** items. The service can be used via an open API.

The ISA Tool and the Metadata-to-Concept Module have been published under a free/libre open source software license, whereas the Google Cloud Vision service is a proprietary service owned by Google. The tag suggestions of the machine vision extension can be queried through a publicly documented, open API.

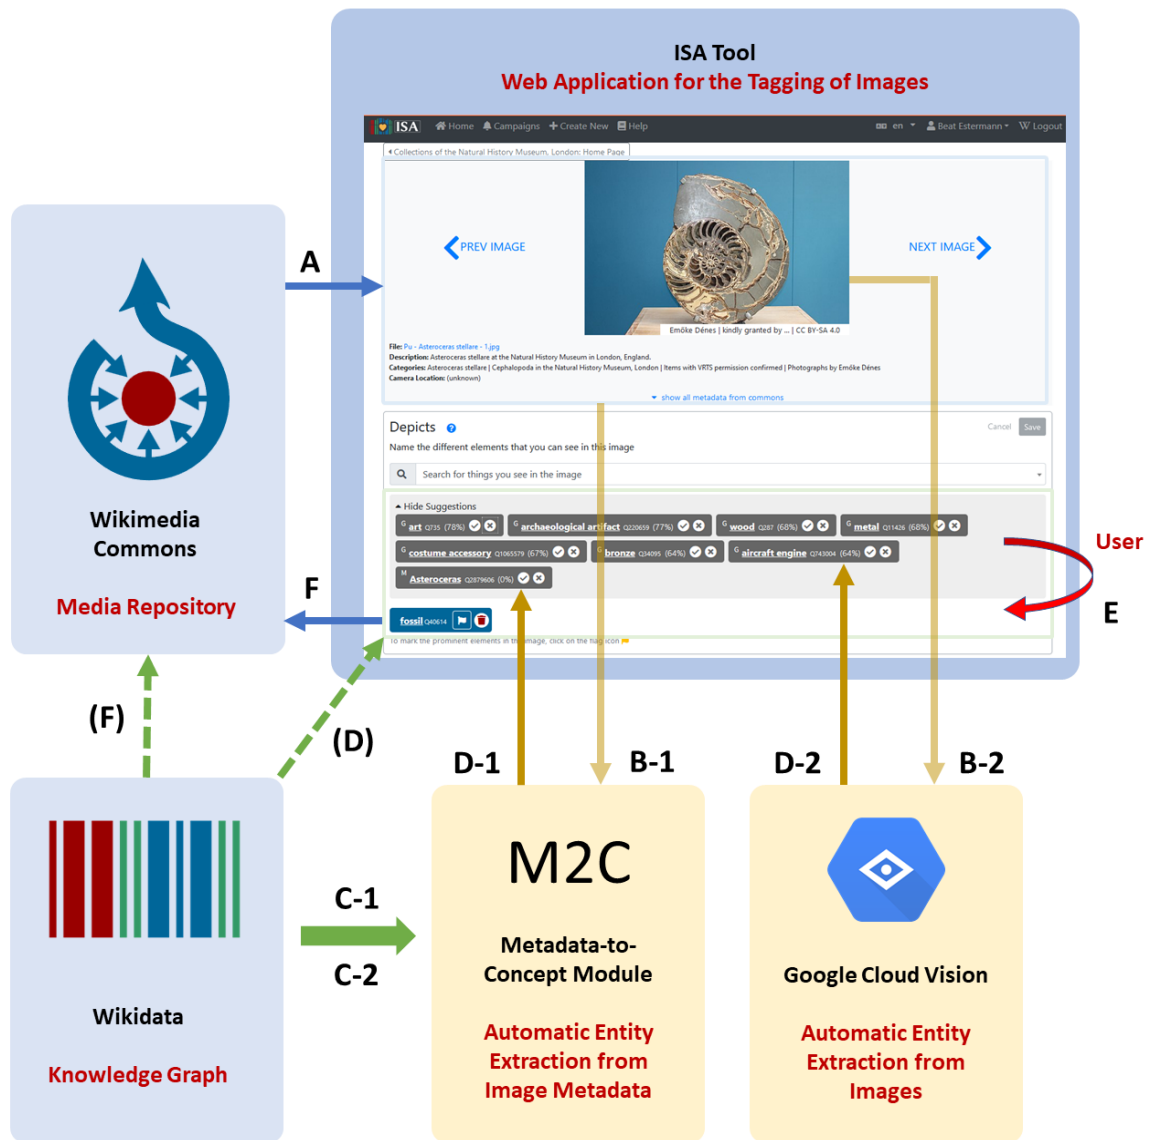
---

<sup>6</sup> [https://commons.wikimedia.org/wiki/Commons:ISA\\_Tool](https://commons.wikimedia.org/wiki/Commons:ISA_Tool)

<sup>7</sup> <https://cloud.google.com/vision/>

<sup>8</sup> <https://www.mediawiki.org/wiki/Extension:MachineVision/Developers>

<sup>9</sup> API: <https://m2c.wmcloud.org/>; code repository: <https://github.com/Wikimedia-Sverige/m2c>



**Figure 1:** Components of the Image-to-Metadata Microservice

Figure 1 illustrates the functioning of the Image-to-Metadata Microservice:

- Images from **Wikimedia Commons** that are part of a tagging campaign are displayed in the **ISA Tool**, along with their metadata (A).
- The **ISA Tool** sends out API calls to the **Metadata-to-Concept Module** (B-1) and to the **Machine Vision Extension** powered by **Google Cloud Vision** (B-2) in order to request tag suggestions. Both services express their suggestions in the form of **Wikidata** identifiers (C-1 / C-2) and send their response back to the **ISA Tool** (D-1 / D-2), which presents the suggested tags to the user in a dedicated section of the user interface.
- In the **ISA Tool**, users can use the free search field or the suggested tags to select the **Wikidata** items corresponding to the objects that are depicted in a given image. If a user decides that one of the suggestions is correct and clicks on the corresponding checkmark, the given tag disappears from the suggestions pane and is displayed instead in the section with the accepted tags (highlighted in blue) (E). If a user decides that a suggested tag is outright wrong, they have the

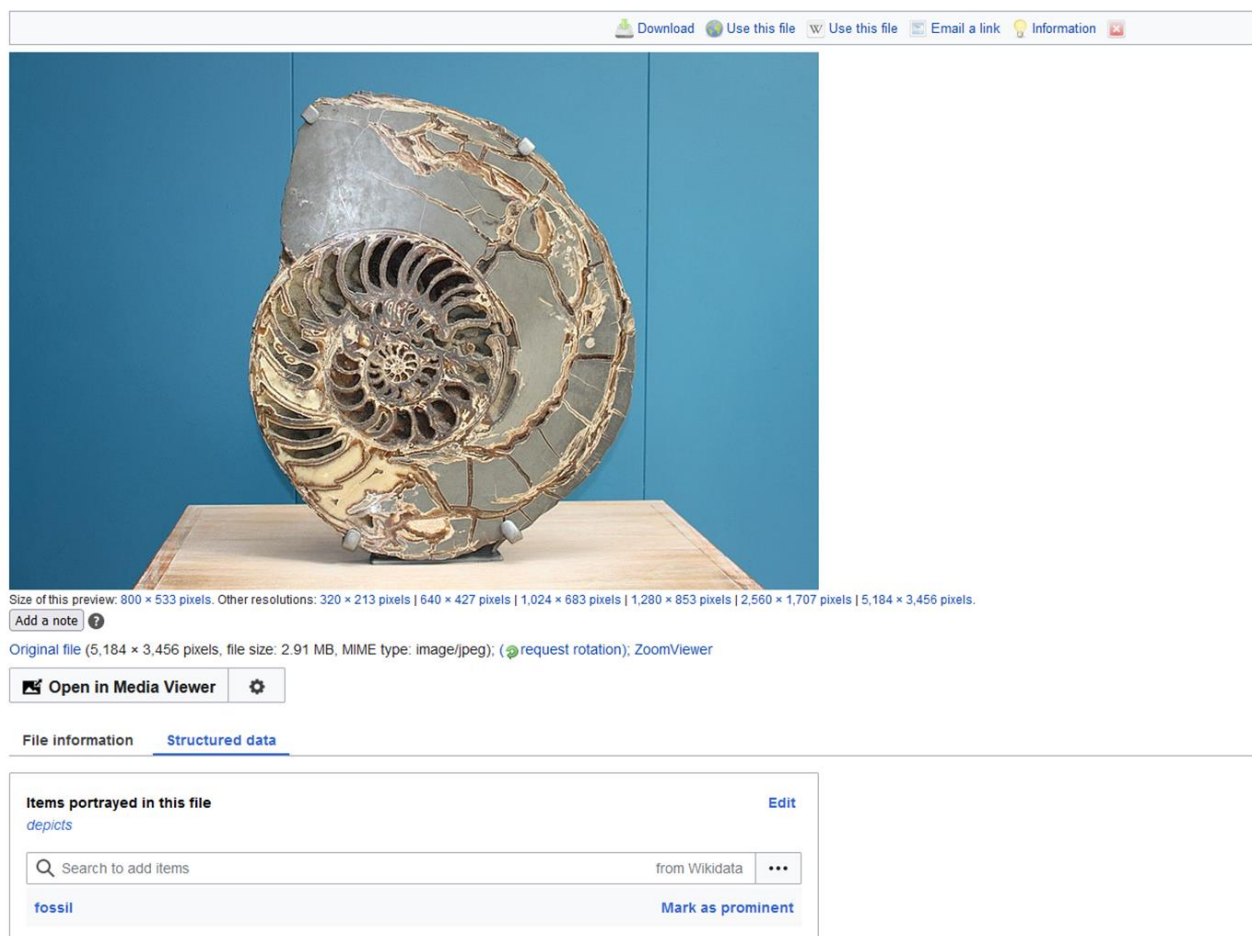


possibility to remove the tag from the suggestions pane. Both types of decision are stored by the **ISA Tool** for further processing.

- When a user of the **ISA Tool** is done selecting tags, they hit the “Save” button to write the selected tags to **Wikimedia Commons** in the form of “depicts” statements (F). Figure 2 shows the corresponding Wikimedia Commons page with a new “depicts” statement in the “Structured data” section. By means of “Structured data”, the image metadata is linked to the **Wikidata** knowledge graph, and the data can be queried through a SPARQL endpoint<sup>10</sup>.

## File:Pu - Asterocheras stellare - 1.jpg

From Wikimedia Commons, the free media repository



The screenshot shows the Wikimedia Commons page for the file 'File:Pu - Asterocheras stellare - 1.jpg'. At the top, there are navigation links: 'Download', 'Use this file', 'Use this file', 'Email a link', and 'Information'. Below these is a large image of a fossilized ammonite shell, cut open to reveal its internal structure, displayed on a wooden table against a blue background. Under the image, there is a note about the preview size (800 x 533 pixels) and other resolutions. Below that is an 'Add a note' button and a link to the 'Original file' (5,184 x 3,456 pixels, 2.91 MB). There is also an 'Open in Media Viewer' button. The page is divided into 'File information' and 'Structured data' tabs. The 'Structured data' tab is active, showing a section titled 'Items portrayed in this file' with an 'Edit' link. Below this is a search box for adding items from Wikidata, with a dropdown menu showing 'fossil' and a 'Mark as prominent' button.


**Figure 2:** Wikimedia Commons page with a “depicts” statement

In the following, the three components of the Image-to-Concept Microservice are described in more detail, with a special focus on the results of the software development activities carried out in the course of the InnoLab project.

<sup>10</sup> [https://commons.wikimedia.org/wiki/Commons:SPARQL\\_query\\_service](https://commons.wikimedia.org/wiki/Commons:SPARQL_query_service)

## 2.1 ISA Tool

At the outset of the project, the ISA Tool was a crowdsourcing tool that had been successfully used for the tagging of images since 2019. The goal of the InnoLab project was to add a semi-automatic tagging functionality that supports users in selecting appropriate tags to describe the content of the images. In contrast to the earlier version, the enhanced version of the ISA Tool now contains a collapsible suggestions pane (see figure 3).



**Figure 3:** Enhanced version of the ISA Tool with the suggestions pane

Apart from this change in the user interface, the following improvements were made to the software in the course of the InnoLab project:

- Support for TIFF images was added to the ISA Tool, which makes the tool more valuable for image collections provided by heritage institutions.
- The tool's database was extended so that it is now able to store tag suggestions obtained from various online services as well as the user decisions pertaining to these suggestions (acceptance / rejection of a given tag).
- The tool's functionality was extended so that it is now able to send API calls to different online services to solicit tag suggestions and to present them in the user interface. Furthermore, it is able to send recorded user decisions back to said online services to help them improve their algorithms. This feedback can be given on the fly, in a continuous machine-learning approach, or the data can be downloaded for analysis in the form of a data dump.
- A mechanism was implemented that allows for the monitoring of automatic tag suggestions and tags entered by users based on community guidelines. Currently, tags that would clearly violate

community guidelines are suppressed from the automatic tag suggestions, and users of the ISA Tool are prevented from saving such tags. For the moment, this deny list simply replicates the blocklist already in use in connection with the Machine Vision Extension.<sup>11</sup> It can be expanded in the future, to the extent that consensus is reached among the community regarding such guidelines.

## 2.2 Machine Vision Service (Google Cloud Vision)

No direct improvements were made to the Machine Vision Service. The ISA Tool was however set up to feed back the results of user decisions made on the basis of tag suggestions obtained from this service. It will thus provide valuable information for the continuous improvement of the underlying algorithm, using a machine learning approach.

## 2.3 Metadata-to-Concept Module

The Metadata-to-Concept Module was developed from scratch. It complements the tag suggestions obtained from the Machine Vision Service. In contrast to the former, the tag suggestions generated by this module are not based on the analysis of the image itself, but on the analysis of image metadata already present on Wikimedia Commons (name and description of the image, Commons categories, etc.). The Metadata-to-Concept Module is particularly useful in the case of images that have been uploaded along with valuable metadata. It also allows to leverage the volunteer effort that has already been made to add descriptions or to assign the images to relevant categories. It thus leverages earlier efforts both by data providers and by the Commons community in order to facilitate the generation of adequate “depicts” statements.

The current version of the module is a functional prototype in the form of a JSON-REST service.<sup>12</sup> It has been developed on a budget of 100 working hours and should be improved in the future (see section 5).

The Metadata-to-Concept Module receives a Wikimedia Commons file name through an API call, gets the metadata about the file via the Wikimedia Commons API<sup>13</sup> and evaluates this metadata, using a linguistic model (Natural Language Processing). For this purpose, the open-source software library SpaCy<sup>14</sup> is used. The module then returns tag suggestions in the form of Wikidata items (Q numbers). Before returning the results, unwanted tag suggestions (e.g. names, disambiguation pages, days of the week, calendar months, taxonomic ranks, etc.) are filtered out based on a blocklist. The number of suggestions is greatly dependent on the language and amount of metadata provided. The module currently works reasonably well for metadata in English, but provides only very few suggestions if the metadata is in another language. The functionality is provided by a Python program using Fastapi and SpaCy. Unicorn/Uvicorn in combination with nginx is used as the web server.

---

<sup>11</sup> <https://github.com/wikimedia/operations-mediawiki-config/blob/master/wmf-config/InitialiseSettings.php>

<sup>12</sup> Available at: <https://m2c.wmcloud.org/extract/>

<sup>13</sup> In the current setup in connection with the ISA Tool, it is the ISA Tool which contacts the Wikimedia Commons API and then passes along the metadata to the Metadata-to-Concept Module as it needs this data for other purposes as well.

<sup>14</sup> <https://spacy.io/>

## 2.4 Maintenance of the Tools

The ISA Tool and the Metadata-to-Concept Module are maintained by Wikimedia Sverige, in collaboration with volunteer developers. Wikimedia Sverige intends to continue supporting the tool with necessary software updates and to act on critical bug reports. Further development of the tools (see section 5) would require additional resources.

The Machine Vision Extension (based on Google Cloud Vision) is maintained by the Wikimedia Foundation.

## 3 Expected Use

This section contains several use case descriptions illustrating how the Image-to-Concept Microservice is expected to provide added value in the context of the data ecosystems it has primarily been designed for, i.e. Wikimedia Commons/Wikidata and the SWITCH Research Data Connectome.

The ISA Tool has been designed to support the addition of “depicts” statements and captions to images on Wikimedia Commons. The use cases described in this section therefore focus on the added value arising from this information under the current constraints. See section 5 for reflections on how some of the current limitations of the tool could be overcome in the future in order to expand the scope of the use cases.

### 3.1 Usage of the ISA Tool

The ISA Tool has been successfully used for the addition of structured description metadata (“depicts” statements and captions) to images on Wikimedia Commons since 2019. Many of such edits have been made in the course of dedicated crowdsourcing campaigns<sup>15</sup>. Such campaigns typically last for a few weeks and feature between a few hundred to several thousands of images. They typically attract up to a few dozen contributors, making between a few hundred up to tens of thousands of contributions. Thus, between 2019 and 2022, 674 users made 440’000 edits by means of the ISA Tool<sup>16</sup>.

The enhancements made to the tool in the course of the current project have been subject to repeated usability tests by real users of the tool, so that it is safe to assume that a) the overall usability of the tool has not been hampered by the changes made, and b) thanks to the tag suggestions, the effort for users required to add “depicts” statements has been reduced on average, with actual effort varying greatly depending on the image, the available metadata, prior knowledge of the user, as well as their assiduousness. Given that there is an ongoing debate within the community on Wikimedia Commons as to what constitutes “good” tagging<sup>17</sup>, no attempts have been made to systematically assess the quality of the tags added.

In view of the future usage of the ISA Tool, the organization of dedicated crowdsourcing campaigns or similar efforts to tag specific collections of images will be key.

---

<sup>15</sup> For an overview of some of these campaigns, see: [https://commons.wikimedia.org/wiki/Commons:ISA\\_Tool/Challenges](https://commons.wikimedia.org/wiki/Commons:ISA_Tool/Challenges)

<sup>16</sup> See the contribution statistics for the ISA Tool: [https://commons.wikimedia.org/wiki/Commons:ISA\\_Tool/Stats](https://commons.wikimedia.org/wiki/Commons:ISA_Tool/Stats)

<sup>17</sup> See [https://commons.wikimedia.org/wiki/Commons:Depiction\\_guidelines](https://commons.wikimedia.org/wiki/Commons:Depiction_guidelines) for the current interpretation of what constitutes rough consensus by the community.

Note that the use of the ISA Tool is not the only method used by contributors to Wikimedia Commons to add “depicts” statements.<sup>18</sup>

## 3.2 Current Applications Making Use of “depicts” Statements

There are currently at least three applications that make use of “depicts” statements:

- **MediaSearch**<sup>19</sup> is a new search function on Wikimedia Commons that was activated in May 2021<sup>20</sup>. In addition to using Commons categories and wikitext from templates, it uses structured data on Commons, including “depicts” statements, and information from Wikidata to find the most relevant and extensive results. Compared to the traditional search function on Wikimedia Commons, it offers a much better experience to users who use the search function in languages other than English.
- **ImageSuggestions** is a Wikipedia Extension/App used to provide image suggestions for Wikipedia editors, which was launched in February 2022. It comes in two flavors: as a MediaWiki extension for experienced editors, and in the form of suggested edits for beginning users. If available, the tool relies on P18 (image) and P373 (Commons category) statements on related Wikidata entries to suggest images. Where these are not present, it uses the MediaSearch, which in turn also relies on “depicts” statements, to generate suggestions.<sup>21</sup>
- **View it!** is a Wikipedia Extension/App used to display related images for Wikipedia articles, Wikidata entries and the like that was launched in January 2023.<sup>22</sup> It relies on Commons categories and on “depicts” statements to identify related images.<sup>23</sup>

## 3.3 Further Use Cases Making Use of “depicts” Statements

In the course of the InnoLab project, the following three use cases have been identified that are worthwhile pursuing. Like the Commons- and Wikipedia-related use cases, they are focused on facilitating discovery and promoting the re-use of images:

- **Integration of an image search on library discovery systems.** To achieve this, the structured data search<sup>24</sup> on national library discovery systems, such as Swisscovery<sup>25</sup>, would need to be extended by an image search that functions in a way similar to “View it!”. SWITCH will help establish the

---

<sup>18</sup> The feature for adding “depicts” statements on Wikimedia Commons was enabled in April 2019. At the time of writing, almost four years later, approx. 11.5 million (12.5%) of the 91 million images on Wikimedia Commons have at least one “depicts” statement.

<sup>19</sup> <https://commons.wikimedia.org/wiki/Special:MediaSearch>

<sup>20</sup> <https://diff.wikimedia.org/2021/05/19/media-search-is-available-for-all-users-on-wikimedia-commons/>

<sup>21</sup> <https://www.mediawiki.org/wiki/Extension:ImageSuggestions>

<sup>22</sup> [https://meta.wikimedia.org/wiki/View\\_it!\\_Tool](https://meta.wikimedia.org/wiki/View_it!_Tool)

<sup>23</sup> Example for the use of “depicts”: this [image](#) shows up in View it! on the article <https://de.wikipedia.org/wiki/Sandgrube> thanks to the “depicts” statement; the image is currently (as of 13.01.2023) not assigned to a Commons Category that would do the trick.

<sup>24</sup> For the description of a pilot implementation of structured data search on library discovery systems, see: [https://outreach.wikimedia.org/wiki/GLAM/Newsletter/January\\_2020/Contents/Switzerland\\_report](https://outreach.wikimedia.org/wiki/GLAM/Newsletter/January_2020/Contents/Switzerland_report)

<sup>25</sup> <https://swisscovery.slsp.ch>

contact to SLSP, the organization behind Swisscovery, with the goal to include an image search, including Wikimedia Commons, on their discovery portal.

- **Integration of an image search on discovery systems for researchers.** A first implementation will be considered in the context of the SWITCH Research Data Connectome Project. Such an integration would allow for metadata improvements when adding image resources to the Connectome Knowledge Graph.
- **Provisioning of the ISA Tool to researcher disciplines.** Next to the SWITCH Research Data Connectome Project, SWITCH may disseminate the ISA Tool to those research discipline multipliers that would strongly benefit from the use of such tools. First discussions took place with organisations that provide services for the disciplines of Digital Humanities and Medicine. SWITCH will include the ISA Tool in future communications.

In addition, it remains to be explored how “depicts” statements could be used in the context of **specific research projects**. A further avenue to be explored is to investigate to what extent the approach could be used to **allow heritage institutions to reach further insights into their collections**.

## 4 Added Value for the Connectome Project

The ISA Tool indicates the following added values for the Connectome project that could be materialized during follow-up activities:

- image and context metadata for the Linked Data Pipeline;
- enrichment of image metadata from Swiss projects in relation to Wikimedia Commons;
- enrichment of image metadata from Swiss libraries, repositories, and archives.

They are elaborated in more detail below, followed by a list of open issues that have been identified during the project and that would need to be addressed depending on the future usage scenarios.

### 4.1 Image and Context Metadata for the Linked Data Pipeline

The Linked Data Pipeline of the Connectome Project allows for the harvesting, pre-processing, mapping, validation and importing of metadata from open data resources using the rescs.org schema. Aggregated data is available for the Communities through the Beta Version of the Connectome API and the Open Data Navigator. Currently, no image metadata is harvested using the Linked Data Pipeline. However, the Rescs.org structure of the Connectome Knowledge Graph could be adapted for this purpose and the data structure could be extended to include “depicts” statements.

### 4.2 Enrichment of Image Metadata from Swiss Projects in Relation to Wikimedia Commons

So far, one project could be identified that does not present any issues related to rights clearance, as the photos concerned have already been uploaded to Wikimedia Commons: The Annemarie Schwarzenbach Photo Edition project, which is jointly led by the University of Zurich and the University of Geneva. The project proposal will be submitted in fall 2023. If it gets accepted, the project is expected to start in summer 2024. Further projects of that kind could be identified by talking to the Swiss heritage institutions that have already made content available on Wikimedia Commons.

## 4.3 Enrichment of Image Metadata from Swiss Libraries, Repositories and Archives

Next to the enrichment of image metadata that are already stored on Wikimedia Commons, there are projects and collections that are not allowed to store image data on Wikimedia Commons due to intellectual property and/or privacy issues. Projects and collections where this applies include: the poster collection of the Swiss National Library; the medical objects collection of the Lausanne University Hospital (CHUV); projects hosted by the Swiss National Data and Service Center for the Humanities (DaSCH); or the photograph collections referenced on the online platform for Swiss audio-visual heritage (Memobase).

Such projects and initiatives would highly profit from the deployment of the ISA Tool in an environment with access control. First discussions with representatives of these projects have taken place. Further dissemination efforts are needed prior to taking decisions on a larger deployment.

In addition, deploying the ISA Tool in the context of Europeana might be an option to be further investigated.

## 4.4 Open Questions

Depending on the future usage scenarios, the following questions should be addressed in a follow-up project, before engaging in further development and deployment activities:

- What feedback do users provide to the enhanced version of the ISA Tool, especially also users from within the research community? – Larger tests need to be carried out once the remaining issues have been resolved and after the enhanced version of the tool has been deployed to production.
- How reliable is the current service maintained by Wikimedia Sverige in cooperation with volunteers? To what extent is Wikimedia Sverige in a position and willing to take responsibility for the maintenance of the ISA Tool in the future? To what extent is the organization ready to act as the primary interlocutor and partner on the side of the Wikimedia Movement when it comes to further improving the tool (e.g. as a partner in an international R&D consortium)? What are the alternatives if no satisfactory solution can be found with Wikimedia Sverige?
- What are the requirements for operations when forking the ISA Tool for specific usage by the research community?
- What would the legal and contractual situation look like when using the ISA Tool in connection with Google Vision outside the context of Wikimedia Commons? What type of agreement could possibly be reached with Google?
- What are the competencies and efforts needed for deployment, operations and maintenance of a service targeted at Swiss research communities provided by SWITCH?

# 5 Current Limitations and Possible Next Steps

## 5.1 ISA Tool

Several issues have been identified that hamper the use and wider promotion of the ISA Tool:

- 1. Performance Issues:** The enhanced version of the ISA Tool regularly encounters performance issues resulting in pages loading slowly or error screens being displayed, requiring the user to reload the page. Thus, at the time of writing (beginning of September 2023), the version of the tool that would support semi-automatic tagging has not yet been deployed to production. Intermittently, performance issues have also affected the production version of the tool.
- 2. Maintenance Issues:** Between May and August 2023, the production version of the ISA Tool was broken for reasons that have not been entirely clarified. It took the software developers tasked with the maintenance of the tool about 3 months to recover a functioning version of the tool. During this period, all the running campaigns and any outreach efforts to promote the tool came to an unexpected halt. The deployment of the enhanced version of the tool has been delayed accordingly.
- 3. Reliability Issues:** The campaign statistics functionality seems to malfunction in the case of big campaigns. This issue is currently being worked on.<sup>26</sup>
- 4. Translation of the interface:** Currently, the user interface is available in 18 languages. Several major languages are still missing. The tool has been set up for translation; adding further languages therefore just requires a crowdsourcing effort.

Once the performance issues have been resolved, further efforts should be made to promote the tool among potential contributors – through visibility campaigns, live demos, conference presentations, etc.

## 5.2 Metadata-to-Concept Module

The Metadata-to-Concept Module currently uses the name and the description of the image, along with a language indication if available, as inputs for a syntactic analysis in a Natural Language Processing approach. The same goes for Commons categories present on the image description page, which are also analyzed as text strings. Based on this analysis, possible candidates of depicted concepts are extracted and matched against concepts already described on Wikidata. Whenever there is a good match, the corresponding Wikidata identifier is returned as a suggestion for a “depicts” statement. This procedure is applied to text strings in English as well as to text strings for which no language has been specified. In the latter case, the algorithm assumes that the text string presented is in English. Text strings in other languages are ignored. Automatic language recognition has been tested in the course of the development process; it was abandoned due to poor results given the fact that the text material fed to the algorithm consists of rather short strings that often do not contain grammatically well-formed sentences, while the number of languages to be distinguished is large. Furthermore, the language may vary across the different metadata fields for the same image.

There are several avenues that could be (further) explored for the improvement of the algorithm:

- 1. Improvement of the existing algorithm based on training data**

---

<sup>26</sup> See Phabricator ticket T328038.



Training data can be generated in two ways:

- a) The quality of current suggestions can be evaluated on the basis of an **analysis of data dumps from the ISA Tool** where proposed suggestions, accepted suggestions and rejected suggestions are continuously logged, along with further “depicts” statements that are applied to a given image. Such log data could eventually also serve as a basis for a machine learning approach, whereby the users of the ISA Tool are continually generating new training data that is used to further improve the algorithm.
- b) Training data can be **manually curated**. This would have the advantage of being able to take the same perspective as the algorithm, e.g. taking into account only the metadata and not the image itself in generating “depicts” statements. This task is however quite onerous, as this approach would require the manual description of at least 250 sample images per language. Thereby, it is not possible to just use an existing corpus of already curated images on Wikimedia Commons, as it is impossible to verify whether existing “depicts” statements have been added based on the image itself, based on its metadata, or based on both.

## 2. Adding Named Entity Recognition

Currently, the Metadata-to-Concept algorithm analyzes the text strings on a purely syntactic level, trying to determine a sentence structure and extracting all terms that are central to the sentence. In cases where this is not possible, all nouns are extracted, but this does not always work reliably, since a minimum amount of sentence structure must also be present to allow for the determination of the word type.

Given the fact that the metadata often contains named entities (names of people, organizations, or places, time indications, historical periods, titles of creative works, etc.), a Named Entity Recognition function could be added to the algorithm, which would be able to identify named entities through name pattern recognition and look-ups in existing registries, including Wikidata, independently from the analysis of sentence structures. This would most likely remedy a current shortcoming of the present algorithm consisting in the fact that it often extracts individual words from longer titles, individual words or names from composite names, single taxon names from composite taxons, and so on.

If external databases are used to identify named entities, this approach could also be used to suggest new Wikidata items to be created by the users of the ISA Tool before adding them in “depicts” statements.

## 3. Adding semantic analysis

While testing the current algorithm, it turned out that many “false positives” (which are currently filtered out by means of a blacklist) could be exploited to improve the recognition of compound names of people and places as well as composite biological taxa.

Semantic analysis could also be used to make more efficient use of the Wikimedia Commons categories already present on the image description pages, as their respective Wikidata item most often contains a pointer to a concept that should be suggested as a “depicts” tag.

Furthermore, a semantic approach could be used to remove “false positives” from tag suggestions generated by means of Natural Language Processing or Named Entity Recognition approaches. This would for example allow for the filtering out of erroneous suggestions for “depicts”

statements pointing to the painters, photographers, or institutions that have the depicted object in their holdings. In many cases, this information is present as structured data on the image description page, but may also be repeated in the file name or in the description, potentially leading to “false positives”. In the same vein, hidden categories present on the image description pages (which are currently ignored by the algorithm) could be exploited to root out “false positives” obtained from the analysis of name and description fields.

If such information is not yet present as structured data on the image description page, the algorithm could be trained to semantically distinguish the various information in the descriptive metadata to assign it to different types of statements (e.g. creator, holding institution, owner, viewed from, etc.) that users can add by means of the ISA Tool in addition to the “depicts” statements. In this way, the algorithm could step by step become semantically more sophisticated, avoiding some of the false positives and allowing users to assign others to the correct type of statement.

#### **4. Adding support for other languages than English**

As mentioned above, reliable language recognition is not possible due to the small amount of text. Nevertheless, for metadata fields where the language has been specified on the Wikimedia Commons page, the algorithm could be trained to process also other languages than English.

Furthermore, the algorithm could be programmed to use further default languages, based on their frequency of use on Wikimedia Commons and their distinguishability with regard to other languages used as default languages<sup>27</sup>.

#### **5. Combination of metadata analysis and machine vision approach in one algorithm**

Currently, the Metadata-to-Concept algorithm has no knowledge of the actual objects or scenes depicted, while the Machine Vision algorithm does not take into account the already existing metadata. In the future, algorithms combining the two approaches may be trained to yield better results than the current algorithms taken individually.

### **5.3 Machine Vision Extension**

The Machine Vision algorithm is not applied to all images on Wikimedia Commons by default. It therefore currently requires the manual inclusion of the respective Commons categories by Wikimedia Foundation staff upon request via a Phabricator ticket. This adds an extra manual step when setting up a new campaign in the ISA Tool and a waiting period of up to several weeks before Machine Vision suggestions are activated for the given campaign. Ideally, the Machine Vision Extension would be enhanced to allow for the inclusion of new categories by means of an API call from within the ISA Tool at the time of creation of a new campaign. In the case of “living” categories, there should also be a function to include newly added images in the analysis.

The Machine Vision Extension currently relies on a proprietary algorithm and is dependent on Google. As Google Cloud Vision does not natively support Wikidata, the extension relies on a (relatively old) mapping table to map tags between the Google Cloud Vision algorithm and Wikidata.

A possible way forward would consist in developing an open source algorithm to replace the current proprietary algorithm – either from scratch or based on an already existing solution. This algorithm could

---

<sup>27</sup> For example, while it may be difficult to tell Italian and Spanish or Czech and Slovak apart in an automatic fashion, it should not be a big deal to assign text strings to the hypothetical default languages English, Chinese, Arabic, and Hebrew.

be trained using Wikidata items directly, without the detour via a mapping table. Thereby, the existing “depicts” statements on Wikimedia Commons could be used as training data.

An alternative (or complementary) way forward would consist in developing open source algorithms for special types of images for which Google Cloud Vision does not perform well. Such algorithms could be added to the ISA Tool in addition to the two current ones.

## 5.4 Media Repository

Currently, the Image-to-Concept Microservice can be used only in connection with Wikimedia Commons. This means that the use of the service is restricted to images that can be published under a free copyright license on the Internet (no copyright issues, no data protection issues, etc.). From the point of view of research use cases that are of interest in the context of the SWITCH Research Data Connectome, this is a considerable limitation. In order to overcome it, the Image-to-Concept Microservice could be set up to be used in connection with other media repositories that may have more restrictive access and re-use policies than Wikimedia Commons. The tagging could still be done on the basis of Wikidata, or it could be extended to other knowledge graphs (see section 5.5. below). In any case, if separate instances of the service are to be run in the context of other media repositories, this would require fostering a separate user community that takes care of the tagging of the images.

While the ISA Tool and the Metadata-to-Concept Module could certainly be ported to be used in connection with another Media Repository<sup>28</sup>, it must be noted that the Machine Vision Extension relies on an agreement with Google for the use of its Cloud Vision algorithm in the context of Wikimedia Commons. Thus, a similar agreement would need to be reached in view of its use in connection with another media repository. Alternatively, another machine vision algorithm could be developed (see section 5.3).

## 5.5 Knowledge Graph

Currently, the Image-to-Concept Microservice can be used only in connection with Wikidata. It would be possible to use it also in connection with other, possibly complementary, knowledge graphs. To do so, algorithms creating pointers to these other knowledge graphs would need to be specifically trained for this purpose. The Knowledge Graph of the SWITCH Research Data Connectome Project could be one way to reuse the Image-to-Concept Microservice.

# 6 Monitoring and Evaluation

In the course of the InnoLab project, reflections were made with regard to relevant statistics for the monitoring, evaluation, and improvement of the ISA Tool and the different algorithms. Annex 3 contains an overview of possible statistics and their assessment. Thereby, four categories of statistics have been distinguished:

1. Monitoring of the functioning and usage of the tool;

---

<sup>28</sup> While connecting the ISA Tool to another MediaWiki-based image repository would likely be trivial, completely segregating the ISA Tool from the underlying repository might require a significant amount of work: First, the architecture would need to fully separate the frontend from the backend API calls, which currently may not entirely be the case. Second, the current implementation of the frontend makes a series of assumptions about the sort of data it gets from the image repository; thus an additional layer of abstraction would be required for the tool to be able to connect to various types of image repositories.

2. Campaign statistics as feedback to the users and campaign managers (basic statistics to be included in every campaign);
3. Improvement of the algorithms used in the context of the ISA Tool (manually or through machine learning);
4. Monitoring of the relevancy and effectiveness of the ISA Tool and the different algorithms used to suggest tags.

While the first three aspects have been addressed in the context of the InnoLab project, the fourth one reaches well beyond the scope of the project and should be addressed as part of a separate project that provides the basis for strategic decision making in view of the further development and promotion of tools facilitating the addition of structured data on Wikimedia Commons and eventually also directly in the context of projects related to the SWITCH Research Data Connectome. In the following are some reflections that should be taken into account in strategic decision making in this area:

- 1. There is disagreement among the community on Wikimedia Commons as to what constitutes “good tagging” of images.**

Apart from some early assessments by the Media Search team, there is little reflection and dialogue on possible use cases and their requirements regarding structured data on Wikimedia Commons.

**Rationale:** As long as there is no agreement on what constitutes “good tagging”, tools such as the ISA Tool and the algorithms employed to suggest tags, cannot be assessed against a shared quality standard, and possibilities to further develop them in a way to nudge contributors in accordance with shared community norms are limited. This aspect has a certain relevance, as criticism has been voiced in the past by community members pointing to potential biases in tagging behavior caused by the use of certain tools. It is thereby unclear whether alleged biases are induced by the tools themselves, by the fact that such tools are used by different types of users (e.g. new contributors), or by other factors. Also, there is no empirical evidence as to the type and extent of such alleged biases.

**Proposed approach:** For pragmatic purposes, it should be assumed that each tool or method used to add structured data to Wikimedia Commons comes with its biases. To the extent possible, these biases should be oriented in a way to nudge contributors in accordance with shared community norms. This goes for the design of the tools themselves as well as for the settings in which their use is promoted. To come to grips with what constitutes “good tagging” of images, it would be useful to create an overview of (possible) use cases with regard to Structured Data on Commons and to systematically document their requirements with regard to the data.

- 2. There are several tools that facilitate the addition of structured data on Wikimedia Commons. There is however neither an ongoing monitoring, nor a proper evaluation of these tools as to their relevance and effectiveness.**

In addition to the standard user interface and the Image Upload Wizard<sup>29</sup>, there are numerous tools that facilitate the addition of structured data on Wikimedia Commons. Examples include the

---

<sup>29</sup> <https://commons.wikimedia.org/wiki/Special:UploadWizard>

ISA Tool, Depictor<sup>30</sup>, OpenRefine<sup>31</sup>, Computer Aided Tagging<sup>32</sup>, ACDC<sup>33</sup>, SDC<sup>34</sup>, or ImageAnnotator<sup>35</sup>. Furthermore, contributors may use bots to generate “depicts” statements *en masse*, e.g. based on Commons categories or based on metadata from partner institutions<sup>36</sup>.

**Rationale:** Without knowing the relative relevance and effectiveness of the different tools, it is impossible to make informed decisions with regard to their further development and their promotion.

**Proposed approach:** In order to allow for the evaluation of the various tools, it would be useful to create an overview of the tools facilitating the addition of structured data on Wikimedia Commons and their perspectives for future development. Furthermore, growth statistics of structured data on Wikimedia Commons over time should be made available, along with an evaluation of the compared relevance and effectiveness of the various tools.

### 3. Structured Data on Commons can be seen as a replacement or a complement to the use of categories on Wikimedia Commons; there are currently no statistics as to their respective use and relative importance.

**Rationale:** Without knowing the extent and the relative importance of the use of structured data compared to the use of categories on Wikimedia Commons, it is difficult to make informed decisions with regard to the development and promotion of tools that support one or the other or that – as in the case of the Metadata-to-Concept Module – may assist users in the translation of one into the other.

**Proposed approach:** Complement the statistics to be provided according to point 2 by longitudinal comparative statistics regarding the use of structured data vs. the use of categories on Wikimedia Commons. Explore and assess ways to facilitate the translation of categories into structured data and vice versa.

## 7 Conclusions and Outlook

In the course of the present project, a semi-automatic tagging functionality was added to the ISA Tool. The new functionality relies on two distinct algorithms: One of them is used to extract entities from the image itself. The other one extracts entities from the image metadata, thus leveraging earlier efforts made to describe the content of the images.

The project had an experimental character: The new functionalities have been implemented on a test version of the tool. Deployment to production has been postponed due to some remaining bugs. During the project, a variety of learnings could be gathered, which are summarized in the following subsections.

---

<sup>30</sup> <https://commons.wikimedia.org/wiki/Commons:Depictor>

<sup>31</sup> <https://commons.wikimedia.org/wiki/Commons:OpenRefine>

<sup>32</sup> [https://commons.wikimedia.org/wiki/Commons:Structured\\_data/Computer-aided\\_tagging](https://commons.wikimedia.org/wiki/Commons:Structured_data/Computer-aided_tagging)

<sup>33</sup> <https://commons.wikimedia.org/wiki/Help:Gadget-ACDC>

<sup>34</sup> [https://commons.wikimedia.org/wiki/User:Magnus\\_Manske/sdc\\_tool.js](https://commons.wikimedia.org/wiki/User:Magnus_Manske/sdc_tool.js)

<sup>35</sup> <https://image-annotator.toolforge.org>

<sup>36</sup> See for example METbot: <https://commons.wikimedia.org/wiki/User:METbot>

## 7.1 Further Improvement and Promotion of the ISA Tool

In the course of the project, the ISA Tool was found to be fit for the intended purpose of adding new functionalities and to test them with real users. As the usage statistics show, the tool is regularly being used for several campaigns every year and attracts a considerable number of users, even though the tool's significance compared to other tools and methods used to add “depicts” statements to images on Wikimedia Commons remains unclear as comparative statistics are lacking.

The project has nevertheless brought to the fore a series of challenges and shortcomings that should be addressed in order to encourage the wider use and facilitate proactive promotion of the ISA Tool in the future:

1. **Performance Issues:** The tool regularly encounters performance issues resulting in pages loading slowly or error screens being displayed, requiring the user to reload the page.
2. **Maintenance Issues:** Between May and August 2023, the production version of the ISA Tool was broken for reasons that have not been entirely clarified.
3. **Reliability Issues:** The campaign statistics functionality seems to malfunction in the case of big campaigns.
4. **Translation of the interface:** Several major languages are still missing and should be added.

Once these issues have been resolved, measures should be taken to increase the visibility and take-up of the tool among potential contributors. To this end, it would be useful to assess and monitor the relevance of the ISA Tool in comparison to other tools and methods used to add Structured Data on Commons.

Furthermore, from discussions with the community, it appears that it is unclear what constitutes “good” tagging of images. It therefore seems important to engage in a further dialogue around this question with the community on Wikimedia Commons, with developers of search and discovery tools, as well as with potential further users of Structured Data on Commons. To inform this discussion, it would be useful to have comparative statistics at hand as regards the relative use of categories and “depicts” statements on image description pages over time.

## 7.2 Further Development of the Algorithms for Semi-Automatic Tagging

Currently two algorithms are being used by the ISA Tool to support the semi-automatic tagging of images:

1. The Machine Vision Extension, based on the proprietary Google Cloud Vision algorithm.
2. The Metadata-to-Concept Module.

There are various possibilities to improve or complement these algorithms that may be considered in view of the further development of the tool:

1. The ISA Tool could be set up to provide feedback regarding accepted and rejected tag suggestions coming from the Machine Vision Extension. From the point of view of the community, the improvement of the algorithm would most likely take place in a black-box fashion without further communication between Google and the community.
2. Inclusion of images in the Machine Vision Extension could be automatized on a category basis (supporting subcategories); the dynamic inclusion of new images in already included categories could be supported.

3. The Metadata-to-Concept Module could be further improved along the lines sketched out in section 5.2 of this report.
4. A free, open source alternative could be developed that could eventually replace the Google Cloud vision algorithm.
5. Specialized algorithms could be developed that excel at generating tag suggestions in cases where the current algorithms are failing. The ISA Tool has been designed in a way that allows for the addition of further algorithms. The main limiting factor is the number of false suggestions presented to the user (which greatly hamper the usability and effectiveness of the semi-automatic tagging functionality).

In a more general manner, it might be useful to further investigate the question under which conditions the users of the ISA Tool make active use of the suggestions, and under which conditions they simply ignore them and favor other methods to select their tags instead<sup>37</sup>.

### 7.3 Facilitation of Research Use Cases

As the first investigations for the identification of research use cases have shown, the main showstopper with regard to the use of the ISA Tool in the current setting is the requirement that all images be made available on the Internet under a free copyright license.

This challenge may be overcome by setting up an alternative media repository where copyright and/or access restrictions can be put in place. The porting of the software tools that are available under an open source license (ISA Tool and Metadata-to-Concept Module) would be possible. In this case it would make sense to clarify between SWITCH and Wikimedia Sverige to what extent it would be beneficial for both parties if the software code were further developed in common, with the perspective of supporting a variety of media repositories. Alternatively, the software code could be forked and adapted to the needs and requirements of the SWITCH Research Data Connectome Project.

Given the current setup, the porting of the Machine Vision algorithm would not be as straightforward: SWITCH could try to reach a separate agreement with Google in order to deploy Google Cloud Vision also on its own media repositories, or an alternative algorithm could be developed. In the meanwhile, the ISA Tool could be used on the alternative media repository without the tag suggestions from the Machine Vision Extension.

If the alternative media repository is expected to be used in a crowdsourcing setting, this would require fostering a contributor community, establishing and enforcing community norms, and fulfilling all the responsibilities that typically fall to the platform provider – this would require an extra effort for tasks that are currently covered by the Wikimedia Foundation and/or the Wikimedia Commons volunteer community. There are ways to avoid a part of this extra effort by implementing each image tagging project in a siloed approach, giving full control over the tagging process to the respective research project. This approach would have its pros and cons that should be pondered carefully before opting for one or the other approach.

---

<sup>37</sup> In the current implementation of the ISA Tool, users can simply hide the tag suggestions if they think that they are of not much help. They may use the existing metadata on the image description pages instead to identify appropriate tags (this metadata is also displayed from within the ISA Tool), or they may simply have a close look at the image and use the search functionality of the ISA Tool to find appropriate tags.

In any case, given the experience with the ISA Tool, reaching a high level of completeness of “depicts” statements on a specific set of images would always require the targeted mobilization of contributors by means of a crowdsourcing / expert sourcing campaign or by putting dedicated staff to the task.

In the longer run, it would be useful to monitor to what extent the addition of “depicts” statements gains traction on Wikimedia Commons over time, regardless of the tools and methods being used for adding them. If the trend further gains momentum, it would be advisable for the research community to try to leverage the community dynamic by making a maximum of their images available on Wikimedia Commons.

A further avenue that could be explored in the context of the SWITCH Research Data Connectome Project is the use of “depicts” statements in research projects beyond the search and discovery use case. If the tagging of images is used for analytical purposes, specific requirements in terms of the quality, consistency and completeness of the tags may arise that would require particular attention.

## 7.4 Contractual Aspects

If the ISA Tool is to be used on a larger scale in the context of the SWITCH Research Data Connectome Project, it might be advisable to verify whether the current contractual arrangements (or absence thereof) are sufficient for the use cases in question. Aspects that might be of relevance in this context are:

- The contract between Google and the Wikimedia Foundation regarding the use of the Google Cloud Vision algorithm in connection with Wikimedia Commons.
- The formalization of commitments regarding the operation and maintenance of the ISA Tool itself and the different modules and extensions employed to generate tag suggestions (Machine Vision Extension, Metadata-to-Concept Module).

If the tools are ported to be used in connection with a different media repository, the same applies *mutatis mutandis*.



## 8 List of Annexes

The following annexes are provided in form of separate documents:

- Annex 1: Iteration Reports (Software Development and User Tests)
- Annex 2: Sample Requests and Responses for the Metadata-to-Concept Module
- Annex 3: Baseline Statistics for the ISA Tool