Linked Data and IIIF

Integrating Taxonomy Management with Image Annotation

Gene Loh
Research and Development
Synaptica, LLC
Kuala Lumpur, Malaysia

Abstract—Linked Data **provides a natural platform for taxonomy management, and when used in concert with the *International Image Interoperability Framework* (IIIF) for accessing graphical image resources, provides an architecture that is suited to an image annotation application. The combination of Linked Data and IIIF specifications can be used to develop extensible discovery tools based on the *Semantic Web Stack*.**

Keywords—Linked Data; Semantic Web; IIIF; image annotation

I. INTRODUCTION

Taxonomy management systems are designed for knowledge organisation of formally structured concept schemes that describe collections of concepts such as subjects, people and places. Words that have multiple meanings and variant terms are disambiguated by linking them together with any number of relationships that can assert facts about how concepts and entities relate to one another.

The ability to view text and images, like JPEG, within a page, has been a key feature of the Internet since the creation of the web browser. Visual images are a powerful medium for communicating ideas and information, and they provide a valuable complement to textual content. A vast amount of information resides inside photographs, paintings, diagrams, and drawings, which is comprehensible to the human eye but relatively inaccessible to machine queries. Well established techniques exist to support searching and browsing images based on the metadata that has been applied to whole images, but search and browse access to specific features within images is a relatively immature field.

While a certain level of information access can be supported by image-level metadata, a much richer knowledge discovery experience can be provided if specific visual details are individually identified. An analogy with information access to physical books illustrates the value proposition for sub-image annotation. In a physical library, card catalogues and bibliographic databases provide a means to identify shelves of books and individual books. Tables of content and subject indexes then facilitate a deeper level of information access to specific pages, sections and paragraphs within a book.

Image-level metadata may be compared with card catalogues and bibliographic databases. They take the user to a discrete work but cannot take the user inside the work to discover its interior content. Sub-image annotation may be compared with tables of contents and subject indexes as they allow the user to search and to navigate inside images.

IIIF specifications are designed to provide structures for digital image repository and the delivery of collections of images and sub-image regional content and metadata, including annotations.

The Semantic Web provides a framework of component specifications and standards that enables taxonomy management using Linked Data to be integrated with IIIF resources. The underlying graph data structure of a Linked Data application enables relationship paths to be established between *resources* in taxonomies and corresponding resources within a collection of images through ubiquitous standards, protocols, data models and ontologies.

II. TAXONOMY MANAGEMENT AND LINKED DATA

Both whole image and sub-image annotations may be semantically indexed using controlled vocabulary terminology.

- *Semantic indexing* specifically means that semantically expressive relationships such as topic, location, and creator, may be used to index items of content or resources of images and sub-images to descriptive terminology.
- *Controlled vocabulary terminology* specifically means that indexing is performed using conceptual entities rather than plain text strings. Conceptual entities have associated descriptive metadata, and URI's and are therefore unique. Concepts are members of a concept scheme that may organise the concepts into a Knowledge Organisation System (KOS) using hierarchical, associative, and equivalency relationships.

A. Taxonomy Management System

A taxonomy management system is used for knowledge organisation, in particular, to create controlled vocabularies comprising *ontologies* that govern *concepts* which are often published as authoritative terms for any topic. Taxonomists, editors and subject matter experts define concepts that form vocabularies, and assign relationships that are defined by ontological rules, like association and hierarchy, as well as property constraints like data type, cardinality, and input field restrictions. Taxonomy management software typically adopt industry standards and specifications such as ISO 25964,

ANSI/NISO Z39.19, SKOS, SKOS-XL, and OWL and RDFS ontologies.

B. *Linked Data*

The application of Linked Data may be seen as an adjunct to taxonomy management, but plays an important role in supporting interoperability between an increasing number of knowledge organisation systems. Import and export of data, including live queries of endpoints can be achieved between Linked Data sources.

A Linked Data application may adopt a Resource Description Framework (RDF) triplestore as the underlying database, allowing inference rules to be embedded, simplifying queries and greatly improving graph pattern matching performance, which in turn facilitates visualisation and discovery of large datasets. Vocabularies in a Linked Data environment can be published in standard RDF notations, or *concrete RDF syntax*, like Turtle, RDFa, JSON-LD, and TRiG [1].

Core to interoperability is the adoption of Linked Data principles, both in the architecture of a software application, as well as the development of vocabularies within a taxonomy management system. Since Linked Data applications often employ RDF triplestores as their primary database, their resulting data models inherently use Uniform Resource Identifiers (URI). Every data row has three elements; a *subject*, *predicate*, and *object*, often referred to as a *triple*. Each row can further be assigned to a *named graph*, which can be used for data segmentation. Data represented as triples inherit graph model properties as relationships are established with each object element in a triple.

Subject	Predicate	Object
<http://id.loc.gov/authorities/names/n50031102>	prov:influenced	<http://id.loc.gov/authorities/names/n78030997>
<http://id.loc.gov/authorities/names/n50031102>	foaf:name	"Charles Babbage"
<http://id.loc.gov/authorities/names/n78030997>	foaf:name	"Ada Lovelace"

Fig. 1. RDF triples

Graph models are natural structures in which controlled vocabularies can be built upon, since concepts are defined not only by string literal data type properties like *label* and *description*, but also by relationships to other concepts and resources, internal or external to each vocabulary. These relationship links produce a connected graph which can be queried using the standard RDF query language, SPARQL.

The example in Figure 2 illustrates how concepts might be connected from three distinct Linked Data vocabularies; GeoNames (Place), Library of Congress Name Authority File (Person), and Library of Congress Subject Headings with relationship predicates from well established ontologies; Person (countryOfBirth), DBPedia (influenced), and the BBC (knownFor).

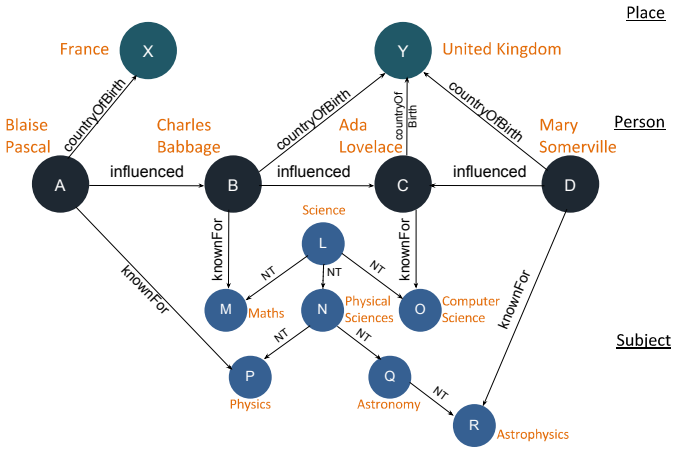

Fig. 2. Connecting controlled vocabularies in Linked Data

Interoperability is the cornerstone of a Linked Data application since the underlying *graph store protocol* is standard across all SPARQL endpoints, making it possible to perform live federated queries across disparate sources of data. Many providers of Linked Open Data vocabularies also make available their datasets in RDF notation that can be ingested, and thus cached locally in a server, for quicker access.

Examples of Linked Open Data vocabularies are Library of Congress Name Authority File, the Getty Research Institute Art and Architecture Thesaurus, and IconClass.

```
PREFIX bbccore: <http://www.bbc.co.uk/ontologies/coreconcepts/>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?personName {

    ?person bbccore:knownFor/skos:broader* <http://id.loc.gov/authorities/subjects/sh89005705> ;
        foaf:name ?personName .

}
```

Fig. 3. SPARQL query

III. International Image Interoperability Framework (IIIF)

IIIF is a community effort with similar objectives of interoperability to Linked Data in liberating access to image collections and their metadata from bespoke and proprietary applications [2]. IIIF digital image repositories augment standard image server features by extending traditional URI links to images over HTTP to include metadata encapsulated by the IIIF Presentation API data model. Additionally, IIIF images can be accessed or referenced in full, in part, or transformed using the Image API.

A. *IIIF Image API*

Very large and complex images are impractical to deliver over the web as discrete images files. The IIIF Image API provides the means for accessing renditions of an image from its full-sized original file to thumbnails or regions of the image of any zoom level, using standard URI's which can be interpreted by any web browser as shown in Figure 4.

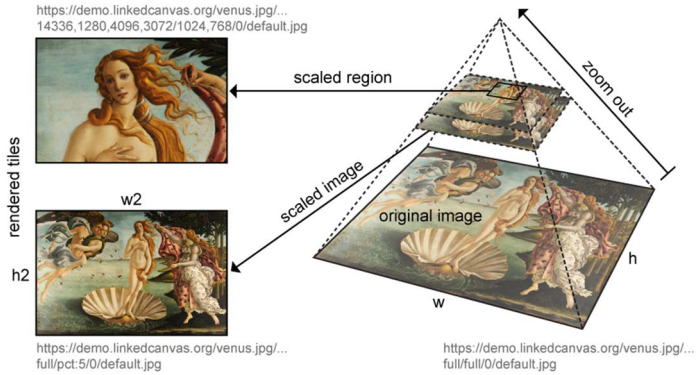

Fig. 4. Referencing images with IIIF Image API compliant URI's

Since the Image API is a web service that returns an image based on a standard HTTP request, the web browser does not distinguish a dynamically generated image from a static image file served from a web server. This provides a high level of utility as no additional front-end code need to be developed for the browser to retrieve images from a IIIF image server. An application that provides image pan-and-zoom features need only to provide regional coordinates and size parameters in the URI of a HTTP request in order for the IIIF server to transform the original image. The resulting rendered image tile is cached by the image server to reduce resource utilisation should the same URI be requested in future.

```
{scheme}://{server}{/prefix}/{identifier}/{region}/{size}/{rotation}/{quality}.{format}
```

Fig. 5. IIIF Image API URI format

Image pan-and-zoom applications often employ an algorithm that performs IIIF Image API requests on a limited number of sub-image tiles in order to reduce the number of connections to the webserver. The resulting pyramidal tile set of sub-images of different zoom levels can often be cached in its entirety on the webserver, which results in more efficient browser-server interaction.

B. *IIIF Presentation API*

The Presentation API provides the ability to add metadata to describe any image or region as an ordered collection, specified by basic primary resources comprising *manifest*, *sequence*, *canvas*, and *content* [3]. The Presentation API is completely independent of the Image API, and both API's are typically implemented in separate layers of an application's component architecture.

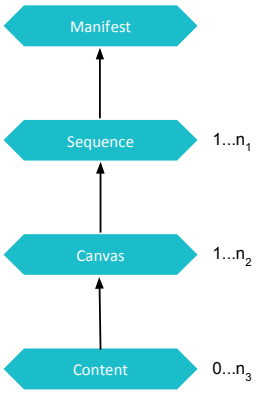

Fig. 6. IIIF Presentation API Primary Resource Types

An application may use one API and not another, or both API's together. The Presentation API provides both a data model for structuring collections of image and image-related resources, and a means for data interchange between IIIF-aware applications through *manifest* files.


```
{
   "@context": "http://iiif.io/api/presentation/2/context.json",
   "@id": "http://4ba5dee7-c88c-4569-ab5a-b6d6b0f29f07",
   "@type": "sc:Manifest",
   "label": "Sandro Botticelli, The Birth of Venus",
   "metadata": [],
   "description": "The painting was commissioned by Lorenzo di Pierfrancesco de'Medici, a cousin of
Lorenzo the Magnificent. The theme was probably suggested by the humanist Poliziano. It depicts
Venus born from the sea foam, blown by the west wind, Zephyr, and the nymph, Chloris, towards one of
the Horai, who prepares to dress her with a flowered mantle. The seascape, stunning for its
metaphysical tone and almost unreal quality, is illuminated by a very soft, delicate light. The real
meaning of this dreamlike vision is still under scholarly debate and investigation but is
undoubtedly linked with the Neo-Platonic philosophy, widely cultivated in the Medici court.",
   "license": "https://creativecommons.org/licenses/by/3.0/",
   "attribution": "N/A",
   "sequences": [
      {
         "@id": "http://18668650-63c6-426a-a2cf-1a4e2ecc53e1",
         "@type": "sc:Sequence",
         "label": [
            {
               "@value": "Normal Sequence",
               "@language": "en"
            }
         ],
         "canvases": [
            {
               "@id": "http://84aae598-febc-4b3a-8d13-be58e72bb469",
               "@type": "sc:Canvas",
               "label": "The Birth of Venus",
               "height": 18840,
               "width": 30000,
               "images": [
                  {
                     "@context": "http://iiif.io/api/presentation/2/context.json",
                     "@id": "http://43b2531b-87b8-499b-8974-4606208e15cb",
                     "@type": "oa:Annotation",
                     "motivation": "sc:painting",
                     "resource": {
                        "@id": "https://demo.linkedcanvas.org/odyd9x95.jpg/full/full/0/default.jpg",
                        "@type": "dctypes:Image",
                        "format": "image/jpeg",
                        "service": {
                           "@context": "http://iiif.io/api/image/2/context.json",
                           "@id": "https://demo.linkedcanvas.org/odyd9x95.jpg",
                           "profile": [
                              "http://iiif.io/api/image/2/level2.json",
                              {
                                 "supports": [
                                    "canonicalLinkHeader",
                                    "profileLinkHeader",
                                    "mirroring",
                                    "rotationArbitrary",
                                    "sizeAboveFull"
                                 ],
                                 "qualities": [
                                    "default",
                                    "color",
                                    "gray",
                                    "bitonal"
                                 ],
                                 "formats": [
                                    "jpg",
                                    "png",
                                    "gif",
                                    "webp"
                                 ]
                              }
                           ]
                        },
                        "height": 18840,
                        "width": 30000
                     },
                     "on": "http://84aae598-febc-4b3a-8d13-be58e72bb469"
                  }
               ]
            }
         ]
      }
   ],
   "structures": []
}
```


Fig. 7. IIIF Presentation API manifest [4]

Front-end applications like Mirador Viewer and Universal Viewer are able to ingest IIIF manifest files and provide their respective specialised feature set like side-by-side image comparison, and annotation, or as a general viewer of the manifest image resources.

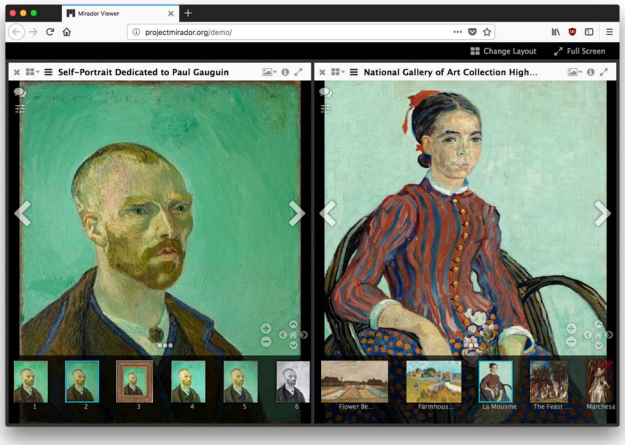

Fig. 8. Mirador Viewer [5]

IV. Linked Data and IIIF

Since IIIF resources are represented by URI's, they can be used either as a *subject* or *object* in a Linked Data context. A Linked Data application can adopt the Image API and define its own ontological properties independent of the Presentation API as depicted by the *:hasThumbnail* predicate in Figure 9.

```
<https://demo.linkedcanvas.org/venus.jpg/full/full/0/default.jpg>
    :hasThumbnail
        <https://demo.linkedcanvas.org/venus.jpg/full/pct:5/0/default.jpg> .
```

Fig. 9. Representation of IIIF Image API resource in an RDF triple

Similarly, resources in the Presentation API (Collection, Manifest, Sequence, Canvas, Annotation, AnnotationList, Range, Layer, and Content) are all referenced by URI's and may be represented the same way in a Linked Data model.

In order to use Linked Data controlled vocabularies in the Presentation API, we examine adopting the additional type, *Annotation*, defined in the IIIF specification. While Linked Data principles can be employed at any level of the IIIF Presentation API stack, when performing semantic tagging—connecting concepts in a Linked Data taxonomy to IIIF image resources—we may insert the primary resource subject URI from a controlled vocabulary into an *annotation* resource *on* a IIIF *canvas*, as illustrated in Figure 10.

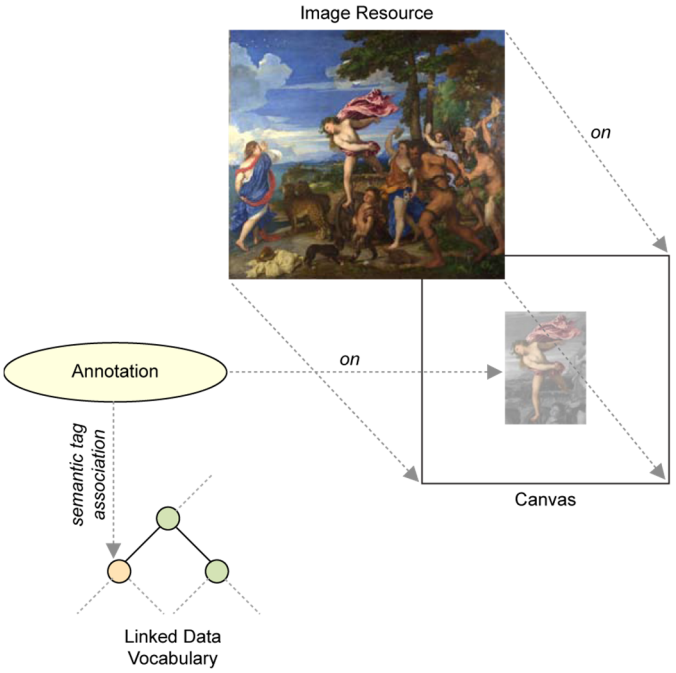

Fig. 10. Semantic tagging with IIIF (proposed)

At the time of writing, the IIIF Presentation API is at version 2.1.1, and does not directly address semantic tagging in its specification. There are a number of ways this might be adopted, either as an *advanced association feature*, or as a IIIF *Open Annotation* extension. For simplicity, we demonstrate this in Figure 11 by using the Open Annotation Semantic Tag type, and augmenting the IIIF *resource properties* to include the *rel* and *href* keys.

```
{
  "@context": "http://iiif.io/api/presentation/2/context.json",
  "@id": "https://demo.linkedcanvas.org/annotatedImage1/annotation/anno1",
  "@type": "oa:Annotation",
  "motivation": "sc:painting",
  "resource": {
    "@id": "https://demo.linkedcanvas.org/annotatedImage1/tag1",
    "@type": "oa:SemanticTag",
    "rel": "dcterms:subject",
    "href": "http://vocab.getty.edu/aat/300055165"
  },
  "on": "https://demo.linkedcanvas.org/annotatedImage1/canvas/p1#xywh=100,100,500,300"
}
```

Fig. 11. Proposed semantic tagging JSON-LD structure in IIIF

In the example in Figure 11, a region of a painting is tagged with the Dublin Core Metadata Initiative (DCMI) dcterms:subject associative relationship (*rel—predicate*) to a term in the Getty AAT vocabulary (*href—object*).

V. Potential Applications

Practical applications can be developed on a Linked Data taxonomy management system which leverages IIIF-compliant open source front-end manifest and deep-zoom image platforms to produce extensible and interoperable image annotation systems. Furthermore, such applications are able to define sub-image regions, or points of interest, of various shapes—like polygons, circles, and markers—to be indexed against concepts in controlled vocabularies.

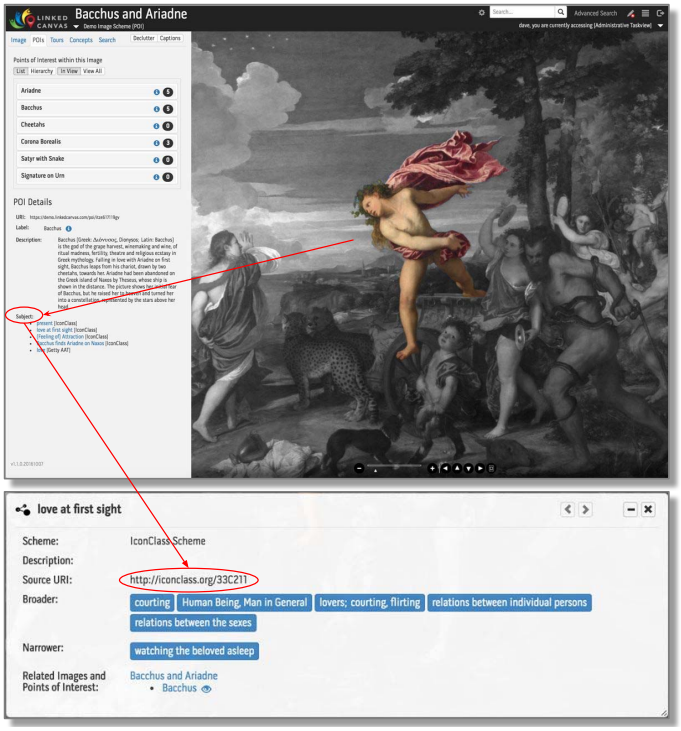

Fig. 12. Semantic tagging with IIIF

Bidirectional relationships between Linked Data and IIIF resources enrich the semantics of a dataset. Search and reporting features would enable visual discovery beyond results of immediate associative relationships, to also include inferred results following predefined ontological rules, by expanding relationship paths through walking of the interconnected graph between primary resource in Linked Data vocabularies and IIIF images.

Fig. 13. Connecting IIIF image resources with Linked Data vocabularies

Any application which includes semantic tagging features for viewing and manipulating digital images compliant with IIIF Image and Presentation API specifications would benefit from having an RDF triplestore backend that can be used to serve IIIF manifests and to provide ease of data migration as resources can be imported and exported in RDF notations. The IIIF Presentation API data model has the potential to be transformed into a Semantic Web compliant ontology. Applications may then adopt the *Linked Data Platform* or *Graph Store Protocol* to provide a RESTful API for CRUD (create, read, update, and delete) operations when annotating images.

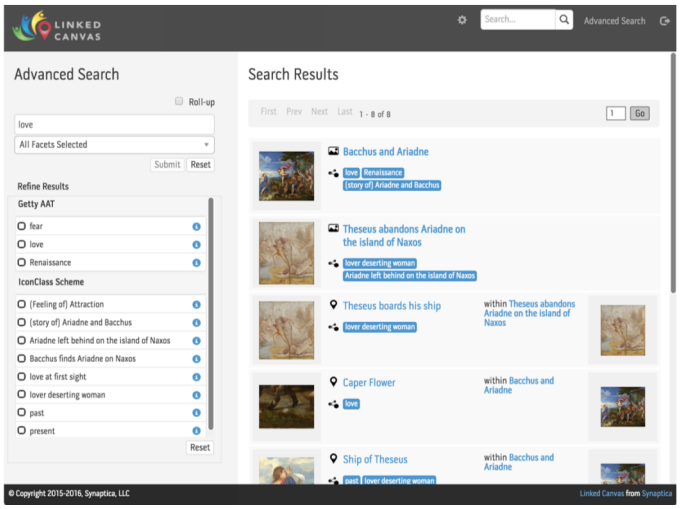

Fig. 14. Searching tagged and related concepts across whole and sub-images

When Linked Data and IIIF are implemented along with knowledge organisation processes and methods, there should be improvements in any or all of the following aspects of information access:

- *Search:* Start with an expression in the form of a search phrase, followed by iterations of disambiguation, search expansion, redirection, and filtering.
- *Refinement:* Also an iterative process for narrowing down of search results based on retrieved metadata.
- *Browse:* Begins with a presentation of organized lists or graphs of related things, and then provides pathways for users based on their interactions.
- *Discovery:* Interrupts the flow of a search or browse upon the surfacing of associated concepts or content of interest.

VI. Limitations of the Presentation API

While the IIIF Image API has potential ubiquity inherent in its design architecture (because every graphical web browser builds in support for handling images), the Presentation API has a more restrictive scope of use. Its design philosophies are borne out of initial requirements to support images of books, manuscripts, photographs, maps, paintings, and digitised renditions of other physical objects. As the adoption of IIIF specifications increases, there is a need to support more use cases, and therefore, more media formats and greater

extensibility in advanced association features like semantic tagging proposed in this paper.

Furthermore, the ancillary specification, Open Annotation, will be superseded by Web Annotation such that many aspects of the current Presentation API are earmarked for deprecation in the subsequent version of the specification. The Presentation API could also benefit from stricter adherence to Semantic Web standards and ontological models like RDFS and OWL to promote even greater interoperability and server-side adoption. To this end, the IIIF specification can be published as a coherent ontology.

VII. Conclusion

Linked Data serves as a bridge between controlled vocabularies and images on a IIIF platform. A concept in a taxonomy, defined as a primary resource in Linked Data, can reference other primary resources defined by the Presentation API; *manifest*, *sequence*, *canvas*, *content*, etc.; or any URI defined by the Image API. Conversely, with minor augmentation to the IIIF specification, *annotations* in the Presentation API can index primary resources in a taxonomy of concepts.

This level of interoperability enables collections of resources defined in the Presentation API to extend beyond textual annotations, and for resulting applications that adopt IIIF to also adhere to Linked Data principles.

The use of controlled vocabularies and semantic indexing plays a crucial role in making connections between related content that is inherently non-textual and thus non-searchable. Concepts used to describe images and sub-image annotations can themselves become gateways to connect to external datasets that are also related to the same concepts. Linked Data knowledge organisation systems and methods can greatly enhance the ability to annotate and provide access to image content served by IIIF Image and Presentation API specifications.

Acknowledgment

The author would like to thank Mr Dave Clarke, and Mr H L Tay for their help.

References

[1] W3C, “RDF 1.1 Concepts and Abstract Syntax”, 2014. [Online]. Available: https://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/. [Accessed: 13 October 2017].

[2] IIIF, “About IIIF”, 2017. [Online]. Available: http://iiif.io/about/. [Accessed: 26 September 2017].

[3] Appleby, Michael, Tom Crane, Robert Sanderson, Jon Stroop, and Simeon Warner, “IIIF Presentation API 2.1.1”, 2017. [Online]. Available: http://iiif.io/api/presentation/2.1/. [Accessed: 26 September 2017].

[4] Bodleian Libraries, University of Oxford, “IIIF Manifest Editor”, 2017. [Online]. Available: http://iiif.bodleian.ox.ac.uk/manifest-editor. [Accessed: 13 October 2017].