Simple Knowledge Organization System (SKOS)†

Arthur Smith

American Physical Society, 1 Research Road, Ridge, New York 11961, USA,
<apsmith@aps.org>

Arthur Smith is lead data analyst at the American Physical Society. He has worked in information technology roles for the *Physical Review* journals since 1995. He received a PhD in physics from Cornell University in 1991. He has given talks at Taxonomy Boot Camp (associated with the KM World conference), at the PIDapalooza meetings on persistent identifiers, and at a variety of related conferences.

Smith, Arthur. 2022. "Simple Knowledge Organization System (SKOS)." *Knowledge Organization* 49(5): 371-384. 46 references. DOI:10.5771/0943-7444-2022-5-371.

Abstract: SKOS (Simple Knowledge Organization System) is a recommendation from the World Wide Web Consortium (W3C) for representing controlled vocabularies, taxonomies, thesauri, classifications, and similar systems for organizing and indexing information as linked data elements in the Semantic Web, using the Resource Description Framework (RDF). The SKOS data model is centered on "concepts", which can have preferred and alternate labels in any language as well as other metadata, and which are identified by addresses on the World Wide Web (URIs). Concepts are grouped into hierarchies through "broader" and "narrower" relations, with "top concepts" at the broadest conceptual level. Concepts are also organized into "concept schemes", also identified by URIs. Other relations, mappings, and groupings are also supported. This article discusses the history of the development of SKOS and provides notes on adoption, uses, and limitations.

Received: 19 July 2022; Accepted 26 July 2022

Keywords: SKOS, knowledge representation, Semantic Web technologies

† Derived from the article titled "Simple Knowledge Organization System (SKOS)" in the ISKO Encyclopedia of Knowledge Organization, Version 1.0 published 2022-07-18. Article category: Standards and formats for representing data.

1.0 Introduction

The Simple Knowledge Organization System or SKOS (Miles and Bechhofer 2009) is a recommended standard from the World Wide Web Consortium (W3C) for representing the structure and content of a knowledge organization system (KOS) (Mazzocchi 2018). Different types of KOS differ in their internal structure, and we will review those differences briefly as they factor into the design of the SKOS recommendation. This introduction will also cover some related standards relevant to understanding SKOS.

1.1 Types of knowledge organization systems

Specific types of KOS include keywords (Lardera and Hjørland 2020), "thesauri" (Dextre Clarke 2019), classification schemes, and ontologies (Biagetti 2021). At the most basic level (keywords) these KOS's are just a non-hierarchical list of selected terms from a language, a vocabulary, ordered alphabetically. A thesaurus or classification scheme as a KOS adds a hierarchical structure to the terms and fundamentally is concerned with the underlying concepts represented by those terms. The hierarchy relation in a thesaurus may be generic (the higher concept is a more general form of the lower concepts, "human settlement" > "city" for example), instantive (the lower concepts are specific instances of the higher concepts, "city" > "Paris" for example) or partitive (the lower concepts are contained within the higher one, "Paris" > "Eiffel Tower"). In addition to hierarchy, thesauri often also include alternate terms or labels for the concepts and additional structured notes, or other features attached to the concepts; they may also include associative relations between concepts to indicate relatedness aside from hierarchy.

Classification schemes mostly use a notation or code (as opposed to a verbal sign) for each concept, which is generally used to impose a specific (non-alphabetic) ordering and a hierarchy on the collection of concepts.[1] Having a code or identifier for the concept also allows a single concept to be labelled with terms in multiple different languages, to easily associate alternative terms, and to allow the specific terms used to change over time (for correction or clarification) while not having to change the underlying identifier for the concept or the indexing of content items with those concepts.

The word ontology has acquired several different definitions but in the context of knowledge organization, ontologies (Biagetti 2021) can be considered a generalization of all these knowledge organization systems. They may include additional types of relationships between concepts and additional attributes attached to individual concepts or terms. A hypothetical ontology of human settlements might include geographic coordinates for each concept representing a place and a special relation indicating a rail link between places, for example. These special properties are usually very specific to the subject domain, so that these relations and attributes themselves become a significant part of the KOS in addition to the concepts and terms. Ontologies using such special properties are the one type of KOS that is not largely or fully supported by the SKOS standard.

1.2 ISO standards

There have naturally been attempts to standardize the representation of a KOS with several international standards developing over the years. ISO 2788 for monolingual thesauri (International Organization for Standardization 1974) and ISO 5964 for multilingual thesauri (International Organization for Standardization 1985) were developed as guidelines for consistency within or between indexing agencies and established the basic distinction between concepts (the semantically meaningful ideas being indexed) and terms (the specific words used to label these concepts in one or several languages). Both of these ISO standards were withdrawn in 2011 when ISO 25964 part 1 (International Organization for Standardization 2011) was released on the subject of *Thesauri for Information Retrieval*. A second part of ISO 25964 was released in 2013 to cover interoperability between vocabularies.

ISO 25964 is covered in more detail in the article on thesauri by Dextre Clarke (2019). The data model it recommends is quite complex, with about two dozen different types of objects involved. It allows attaching notes and other attributes to each term (text label) in the thesaurus and grouping concepts in unordered or ordered lists which may also be included as part of the concept hierarchy. It distinguishes between preferred and non-preferred terms (one preferred term per concept and language), and between hierarchical and associative relationships between concepts. In principle ISO25964 is sufficiently complete to represent every type of KOS except the general class of ontologies with their domain-specific relations and properties. The relationship between SKOS and these ISO standards will be discussed in detail in section 3.3 below.

1.3 Resource Description Framework (RDF) and the Semantic Web

In a 1999 book chapter (Berners-Lee and Fischetti 1999) Tim Berners-Lee and Mark Fischetti proposed the next phase of the World Wide Web that Berners-Lee had started in 1993, which they called “the Semantic Web”. This was a little more substantively described by Berners-Lee, Hendler and Lassila (2001). The World Wide Web Consortium (W3C) has been steadily developing recommendations for this next phase since then. The Web itself defined Universal Resource Locators (URLs) or more generally a Uniform Resource Identifier (URI) as the syntax and semantics of formalized information for location and access of resources via the Internet. (Internationalised Resource Identifier (IRI) is even more general, but we will stick with URIs in the following). A URI can be considered to represent a concept or entity as long as the owner of the underlying web server wishes to preserve it, so at least a fraction of URIs can be considered persistent identifiers. On the other hand, this also means that one cannot always rely on URIs to be stable when linking to them. URIs also share the same semantic problem as any other sign: different people (or even the same person at different times) may attribute different concepts to it.[2]

The Semantic Web builds on the World Wide Web by defining a Resource Description Framework (RDF) (Lassila and Swick 1999), that allows use of URIs as identifiers both for entities and for their properties or relationships. The central concept in RDF is the “triple”, a statement containing a “subject”, “predicate”, and “object” where both subject and predicate generally are URIs, and the “object” or value of the statement can be a string of text, a date, a number or other type of literal value, or another URI.

RDF, the underlying format for the SKOS standard, has now been used for a wide variety of published Web standards, both from the W3C and from other organizations such as Dublin Core (International Organization for Standardization 2017)[3], “FOAF” (Brickley and Miller 2005)[4] and a variety of ontologies used in the biomedical sciences (Smith et al. 2007). Semantic Web technologies including RDF and its extensions have found significant use in representing complex models of real-world systems (Allemang, Hendler and Gandon 2020).

2.0 SKOS structure

The elements of SKOS are well summarized in the synopsis of the W3C recommendation (Miles and Bechhofer 2009):

> Using SKOS, concepts can be identified using URIs, labeled with lexical strings in one or more natural languages, assigned notations (lexical codes), documented with various types of note, linked to other concepts and

organized into informal hierarchies and association networks, aggregated into concept schemes, grouped into labeled and/or ordered collections, and mapped to concepts in other schemes.

The language of SKOS is RDF. This means that every subject or predicate is a URI, while objects or values may be URIs or literal values. In the following discussion the skos: prefix should be understood to represent the first part of a URI in the SKOS standard, namely 'http://www.w3.org/2004/02/*skos*/core#'. For example, skos:Concept actually represents the full URI 'http://www.w3.org/2004/02/*skos*/core#Concept'. Also rdf: similarly represents the standard URI prefix for that namespace, 'http://www.w3.org/1999/02/22-rdf-syntax-ns#'.

2.1 SKOS concepts and their labels

The fundamental entity in SKOS is skos:Concept. This is a unit of meaning that can then be attached to documents or other entities for indexing and classification purposes. Every skos:Concept has a Uniform Resource Identifier (URI) which acts as the identifier for the concept and within the RDF context takes the place of the subject (or sometimes object) of RDF statements. One statement attached to every such concept is to assert that it is an instance of the class skos:Concept:

<concept URI> rdf:type skos:Concept .

Attached to every concept are labels which are literal strings with an attached (optional) language tag; in RDF notation they would look something like label@en where the 'en' piece is the 2-letter ISO language code for English. Labels in SKOS are either preferred (skos:prefLabel property), alternate (skos:altLabel property), or hidden (skos:hiddenLabel). These would be encoded for example as:

<concept URI> skos:prefLabel "label"@en .
<concept URI> skos:altLabel "string"@en .
<concept URI> skos:hiddenLabel "lable"@en .

Preferred labels are usually displayed to represent the concept and there should be only one preferred label per language code for a given concept. Within a coherent concept scheme there should be only one concept with a given preferred label (including language tag), so that the preferred labels can uniquely identify concepts, even though they are not the actual identifiers for them (see further Section 2.3). Preferred labels may change over time, so they also cannot be used as identifiers for that reason. Labels can be provided in any number of different languages; SKOS is inherently multilingual. Whether preferred labels are provided for every concept in every supported language is up to the maintainers; often this will not be the case so that application software will need a fallback mechanism to display concepts in languages where labels are missing.

Alternate labels are different words or phrases, acronyms, or other terms that should mean the same thing as the preferred label and intended meaning of the concept. They might also be terms commonly searched for which should return this concept. Alternate labels may or may not be displayed to end users in normal usage. Hidden labels are similar, except they are only used for searching, and not displayed to end users – this is commonly used for mis-spelled forms of the other labels.

While the labels attached to a concept URI may change, the concept represented by the URI should have a single unambiguous meaning. For our hypothetical collection of concepts about cities, a URI with preferred English label "Calcutta" at one point in time might later change to preferred label "Kolkata", while still representing the same city ("Calcutta" would then be a natural value for an alternate label). The meaning of the concept is more than just this label. Whether this URI represented the city proper, the metropolitan area, the port, or the city in some sense in all those aspects without being more specific would be up to the organization responsible for the URI. This conceptual meaning may not always be correctly interpreted by those using it for indexing purposes, of course.

2.2 SKOS relations and hierarchy

SKOS has two direct hierarchy relations: skos:broader and skos:narrower, which are inverses of one another. To follow the "city" example mentioned earlier, this could be denoted by:

<URI for human settlement> skos:narrower <URI for city>

And similarly, the inverse

<URI for city> skos:broader <URI for human settlement>

SKOS does not define the type of broader/narrower relation being invoked in these relationships, so the meaning could be generic as in this case, or instantive or partitive. A narrower (child) concept may have multiple broader (parent) concepts since SKOS inherently supports polyhierarchy where needed. No matter under which parent concept a child is found it has the same meaning since it has the same identifier (URI) referring to a single unit of meaning.

SKOS also defines a reflexive associative relation, skos:related. This is intended to link concepts that are not connected by hierarchy but have some conceptual relationship so that a person using one concept for searching or indexing may also have an interest in using the other. That this relation is reflexive means that whenever <A> skos:related <B> we also have <B> skos:related <A>, so the relation can be seen from either direction.

2.3 SKOS concept schemes

In SKOS, every concept should belong to one or more skos:ConceptScheme entities. These are the top-level organizing structures within SKOS. A concept is indicated as belonging to a concept scheme through the skos:inScheme predicate. Continuing with our human settlements example, it could be part of a general "Geographic Entities" concept scheme, so we would have statements like:

<URI for city> skos:inScheme <URI for Geographic Entities>

Note that the URI for the concept scheme is intended to represent the scheme in general, not the concept denoted by its label as for regular concepts; in SKOS "concepts" and "concept schemes" are disjoint. In other words, "URI for Geographic Entites" here represents that collection of concepts and relations, not a particular concept that could otherwise be placed in a hierarchical relation to the rest.

Concept schemes contain top concepts, which should be the broadest concepts in the hierarchy of the scheme; these are the entry points from the concept scheme to the hierarchy of concepts. These relationships are indicated by the "skos:hasTopConcept" predicate. Concept schemes are not exclusive: a concept can belong to more than one scheme, and a concept that is a top concept in one scheme may be a lower-down concept in the hierarchy of another.

A SKOS concept scheme does not necessarily correspond with a data file that might be provided for the associated SKOS entities and relationships. There may be several such files provided with different representation formats or at different URL locations to represent the same scheme, and a single file (or its interpretation as a graph of RDF triples) may contain information about multiple concept schemes and all their associated concepts and relationships.

2.4 Documentation properties

SKOS includes a set of documentation properties, subproperties of "skos:note", which can be used to annotate concepts. The values for these documentation properties may be literal (a string with optional language tag, as for labels), or alternatively a URI which may have additional associated properties. "skos:definition" provides a way to clearly explain the meaning of a SKOS concept. "skos:scopeNote" indicates how the concept should be used in indexing, while "skos:example" can provide examples of entities that should be indexed with the concept. "skos:historyNote" and "skos:changeNote" are to provide information about previous versions of the same concept, and "skos:editorialNote" is intended for housekeeping purposes of those administering the concept scheme, for example to indicate that a future review may be needed.

SKOS includes a provision for notations attached to concepts with the skos:notation property; this is typically a numeric or alpha-numeric code like the UDC or Bliss codes, or for example the Mathematics Subject Classification (MSC) codes (Fraser 2020, Section 6.2). An application might use the natural ordering and hierarchy provided by a notation to display the concepts in a concept scheme; this is one alternative to alphabetic ordering by preferred label that is the usual default with SKOS. As literal values skos:notation entries can have a custom datatype but otherwise are simple strings and do not provide a mechanism to describe the internal structure of such classification codes.

2.5 Other features of SKOS

SKOS has some provision for collections of concepts with the skos:Collection and skos:OrderedCollection classes; however these appear to be very rarely used as they are not considered subclasses of skos:Concept and therefore cannot be placed into a concept hierarchy in any natural manner. Concepts within a skos:Collection are listed with the skos:member property, while concepts in a skos:OrderedCollection are grouped into an rdf:List that is linked with the skos:memberList property to the ordered collection URI.

SKOS also includes mapping properties to allow relationships similar to the within-scheme relationships discussed previously, but to link between concepts from different concept schemes. Because subjects and objects in RDF are identified by URIs which are universal and unambiguous (in principle uniquely defined independent of any context) it is a simple matter to include in a data file associated with one concept scheme these mapping relations linking that scheme to another one (or several) that may be under a completely different domain. For concepts that have the same meaning as far as document indexing is concerned the "skos:exactMatch" property is provided. Concepts in separately maintained vocabularies are unlikely to be precisely identical in meaning simply because they are maintained separately and, for example, are likely to have distinct and incompatible hierarchical contexts, histories, labels, and so forth. But if the meanings are sufficiently close to be interchangeable for use in applications, this exact match relation is the right choice. For concepts that are similar in meaning but not exactly the same SKOS provides "skos:closeMatch". Hierarchical cross-scheme relations can also be described by "skos:narrow-

Match" and "skos:broadMatch", while "skos:relatedMatch" can link concepts that have some associative relation.

3.0 The development of SKOS

3.1 Purpose and history

RDF itself was designed in the context of standardizing the problem of information retrieval across the burgeoning Internet, and it was recognized from the start that representing thesauri and other types of knowledge organization systems would be an important component of helping make resources more easily discoverable (Baker et al. 2013). The idea of representing a vocabulary using language-independent concepts with labels in multiple languages was initially developed across several European projects: DESIRE (Development of a European Service for Information on Research and Education, 1997-2000), LIMBER (Language Independent Metadata Browsing of European Resources, 1999–2001) and SWAD (Semantic Web Advanced Development, 2001–2004) which produced the original draft of a "Simple Knowledge Organization System" (SKOS) (Miles, Rogers and Beckett 2004). This was taken up by the World Wide Web Consortium (W3C) as a working draft, and in 2006 the Semantic Web Deployment Working Group (SWD) was chartered to conduct the systematic review required to become an official W3C Recommendation.

The SWD collected specific use cases (Isaac, Phipps and Rubin 2009) to inform the development of the new recommendation. These included examples from the arts and sciences and from potential commercial applications such as the Product Life Cycle Support vocabulary, and even general library classification systems like UDC. From these use cases they derived a collection of requirements that the recommendation could support, and then the group debated and decided on which ones would be accepted and implemented. Some candidate requirements such as the need for a way to coordinate concepts (combining two or more to create a new one) were not accepted or implemented.

In addition to SKOS, SWD (led by Tom Baker and Guus Schreiber) was also responsible for the RDFa standard for embedding RDF in web pages (XHTML format) and for a collection of "Best Practice Recipes for Publishing RDF Vocabularies".[5] All the (240) issues raised by the working group across these three responsibilities are publicly documented.[6] Much of the detailed effort of the working group was focused on formalizing the representation of the essential structural elements of a thesaurus using RDF, with definitions, constraints, and inference rules defined using the Web Ontology Language (OWL) in addition to the standard property and class definitions provided by the basic RDFS standard. These formal details and decisions are discussed in depth in Baker et al. (2013).

3.2 Considerations and limitations

Those involved quickly noticed that the relationships and nature of terms and concepts in thesauri are generally informal, designed for intuitive human use, not machine logic. Introducing formalizations of their meanings would be too restrictive and not general enough for wide application. The most practical approach was to provide a structure that was as non-specific as possible, a "minimal ontological commitment" (Baker et al. 2013) to the nature of concepts and their relationships. The standard does not say what a concept is, other than that it has labels and relations to other concepts and one or more concept schemes. One consequence of this minimum commitment was to drop some of the more complex properties that had been part of early drafts. This included refinements of the broader/narrower relations to specific kinds of hierarchical relations (partitive, instantive, or generic); these were dropped from the final SKOS standard.

The question of whether hierarchy relations should be transitive (i.e. does <A> <skos:broader> <B> and <B> <skos: broader> <C> imply <A> <skos:broader> <C>?) was settled by deciding the standard properties would be non-transitive, but transitive superproperties would be logically entailed that do not need to be used directly. In the final standard, <A> <skos:broader> <B> and <B> <skos:broader> <C> does NOT imply <A> <skos:broader> <C>, but it does imply that <A> <skos:broaderTransitive> <C>. The skos:narrowerTransitive property is an inverse to skos:broaderTransitive with the same implications.

The early SKOS-core draft[7] also included a skos:subject property for use in indexing; this was dropped from the final recommendation as being outside of the scope of SKOS itself, and able to be handled by established properties provided for example by the Dublin Core standard. However, a skos:subject might have been a useful addition to advise data users that the value is the URI of a SKOS concept rather than some arbitrary resource or string; it would also have brought the question of distinguishing indexing and non-indexing concepts within the scope of the recommendation (Baker et al. 2013). The draft skos:subject predicate was heavily used in DBPedia for several years.

Concept symbols (in addition to language-specific labels) were also part of the initial draft but dropped in the final recommendation as being duplicative and underspecified. Given that the full range of Unicode characters is available in RDF literal string values this does not seem to be missed.

A number of issues were raised during the development and review process but ultimately never addressed by the recommendation itself. The standard does not specify what the URI corresponding to a SKOS concept or concept scheme should resolve to, or even that it need necessarily resolve at all. Some *de facto* standards seem to be developing in this area but it is not part of the original or associated recommendations.

The issue of non-indexing terms and node labels was also never resolved. These are used in thesauri and classification schemes to group narrower concepts in a logical fashion under a heading that is not in itself a concept within the domain of the thesaurus. Our hypothetical thesaurus of human settlements might include a collection of "cities by population" with narrower concepts "over 10 million", "2 million to 10 million", "500,000 to 2 million", *etc*. The node label "cities by population" is there only to group the narrower concepts, not to be used for indexing in itself (i.e., an empty class). There is no way within SKOS to retain the hierarchy of the thesaurus while indicating that the concept "cities by population" is somehow different from other concepts. SKOS does include several options for collections of concepts; however here the ontological commitment of the recommendation is not minimal enough: Collection and Concept are regarded as disjoint, so that the relation definitions disallow a collection from being a mid-level node in a hierarchy of concepts. Additional application logic or non-SKOS properties are needed to use a collection as a hierarchy node, or to indicate that a concept is a node label and not intended for document indexing. (Panzer and Zeng 2009) suggest creating Assignable and NonAssignable subclasses of skos:Concept to address the issue of non-indexing terms or node labels, with additional subclasses as needed for more specific cases.

SKOS concept schemes are the top-level organizing entities for the concepts, and the inScheme property relates concepts to the scheme or schemes they belong to. An issue raised during review was whether the hierarchy relationships could also be considered to belong to particular schemes. For example, a set of concepts could be structured in one hierarchy when considered part of one scheme, but in a different hierarchy in another. To allow this would require attaching something like the skos:inScheme predicate to each skos:broader/narrower statement, but statements about statements require reification which is far from simple, so this additional level of complexity is not part of the SKOS recommendation. Different concept schemes can provide the RDF for their scheme in separate graphs (individual RDF files) that may partially solve this problem, but that could lead to significant confusion when several such graphs are combined; the result would be unwanted polyhierarchy that might even include cycles within the hierarchy graph. So, in the final SKOS standard concept URIs should be considered to be part of a single hierarchy that is independent of concept scheme.

Thesauri and classification schemes often include phrases: pre-coordinated terms – two or more concepts linked together for indexing purposes.[8] This is to indicate that indexed documents are about the concept expressed by the phrase, often with the second concept in the context of the first. Some complex classification schemes have mechanisms for free combination of concepts with syntactically meaningful notation, for example the *Integrative Levels Classification* (ILC) (Binding, Gnoli and Tudhope 2021). The working group considered adding a mechanism to allow for this sort of joining of concepts with Boolean logic, but it ended up being deferred and there is no specific support for pre-coordination nor for syntactically meaningful free faceting in the SKOS recommendation.

Finally, the issue of how to handle deprecation and revision of a vocabulary or individual concepts within it was raised but never addressed specifically by the SKOS team. There are later W3C recommendations, for example, the *Data Catalog Vocabulary* (DCAT)[9] that may be considered to cover this, and indeed it does seem out of the main scope of SKOS. Nevertheless, the issue of revision handling is a practical issue faced by just about every maintainer of a vocabulary or thesaurus. One mechanism for handling deprecation could be to include mapping relations between deprecated concepts and their replacements, perhaps treating different versions of a concept scheme as if they were distinct schemes, which might require adding version identifiers to concept or concept scheme URIs. Does a new version of a concept scheme create new concepts (with new URIs), or are these concepts the same as the ones with the same labels in earlier versions? The SKOS standard does not address this question.

While in principle any thesaurus or classification scheme can be readily converted to the SKOS format just as a file format change[10] some of these issues and the nature of linked data and RDF itself can necessitate some re-thinking of the classification structure. For example, a hierarchy with a lowest level node having a list of different topics and then an Other category (for everything not specifically listed) at least needs a better label than Other, since concepts are defined independently of their placement in a hierarchy and can be reused outside of that context. Migrating to SKOS means asking questions like: does the same label mean the same thing in different places (in which case it can be one single concept) or different things (so the labels will need to be modified to be distinct and clear about meaning). So even the "minimal ontological commitment" principle to which SKOS tries to adhere does require the maintainers of what were previously more informal vocabularies to better label and define their concepts for wider usability.

3.3 ISO 25964 and SKOS

The W3C recommendation Simple Knowledge Organization System (SKOS) (Miles and Bechhofer 2009) was released in 2009 after almost a decade of development. Since this was prior to ISO 25964 it was not directly informed by that standard but rather by the predecessor ISO standards. The SKOS primer (Isaac and Summers 2009) includes an appendix giving the correspondence between SKOS and those earlier standards (2788/5964). ISO and the SKOS contributors developed an updated correspondence between the new standard and SKOS in 2012 (ISO TC46/SC9/WG8 working group and Isaac, 2012). There is considerable overlap between the two and some areas of difference.

The Concept class in SKOS serves the same purpose as the ThesaurusConcept object in ISO 25964's UML model, and concept schemes in SKOS play essentially the same role as the Thesaurus object in ISO 25964 as the top-level organizing structures for the controlled vocabulary. The mechanism for labeling concepts in SKOS has some similarities but is fundamentally different (and simpler) than the model in ISO 25964, where each label is an object in itself, the ThesaurusTerm (which may be a PreferredTerm or other type), to which a variety of additional properties can be attached. An extension of SKOS known as SKOS-XL (see section 4.3 below) allows URIs to be assigned to labels so that they can have additional properties and so be more closely aligned with the ISO 25964 ThesaurusTerm approach.

The SKOS relational properties are similar to but again simpler than the ISO 25964 AssociativeRelationship and HierarchicalRelationship objects, where in both cases a role can be attached to those relations allowing for more precise description of the type of relation. With regard to the SKOS mapping relations, ISO 25964 part 2 (International Organization for Standardization 2013) covers essentially this topic of mapping between vocabularies. It discusses strategies for mapping and suggests several additional relations that would be useful, particularly compound mappings, where an entry in one vocabulary corresponds to a combination (Boolean AND or OR) of two or more entries in another. See also (Zeng 2019) on interoperability in Knowledge Organization Systems. Here the model SKOS uses is simpler than is perhaps ideal for the purpose of fully addressing relations between different vocabularies and thesauri.

There is a lot of similarity between the SKOS note properties and the Note objects defined by ISO 25964, though there are some small differences in details. For instance, ISO 25964 does not have an Example note, but it does have a CustomNote that can have an arbitrary noteType value. Because ISO 25964 distinguishes concepts and terms (corresponding to the label values in SKOS) as separate objects, it attaches some types of notes to concepts and some to terms (and some are allowed on both). Within ISO 25964 notation is an optional string property within the ThesaurusConcept object. The SKOS approach is slightly more flexible in this case as it allows a data type to be attached to the notation string, so that several different notations could be used within the same concept scheme without conflict.

In ISO 25964 the ConceptGroup and ThesaurusArray objects are similar to the skos:Collection classes; however ThesaurusArray in ISO 25964 can be placed within the hierarchy of the thesaurus, and this sort of construct with node labels is often used in thesauri. This is an area where SKOS is missing support for a commonly used feature.

4.0 SKOS Usage

4.1 Using SKOS

SKOS itself does not have a property to relate a document being indexed with the specific concepts from a SKOS vocabulary it is about, although such a property was party of an early version of the SKOS proposal (Baker et al. 2013) (see section 3.2). The SKOS primer (Isaac and Summers 2009) recommends using the Dublin Core subject predicate: http://purl.org/dc/terms/subject, i.e. an RDF statement of the form:

<document URI> dct:subject <concept URI>

Libraries and museums have published linked open data with these sorts of relations between their collections and thesauri, but often they use their own custom properties instead of the Dublin Core one for this indexing relation. For example, the US Library of Congress developed the BIBFRAME ontology (Hawkins 2015)[11] to provide full bibliographic metadata for items in their collection, and BIBFRAME has its own subject property[12] which the library uses in their linked data to index works with concepts from their vocabularies. A system that indexes documents with a SKOS vocabulary for internal use can alternatively just record the document-URI relations in an internal database rather than creating a public RDF graph.

Naturally, concepts used in such a context need not be limited to a single SKOS concept scheme. Indexers can link their resources to any SKOS scheme with persistent URIs (which in principle they should all have). Applications making use of indexed resources of this sort need to be able to resolve the URIs or find some other lookup mechanism to know what the labels are, if possible in a specific language (i.e. the skos:prefLabel and skos:altLabel properties). Search engines and similar applications will be more useful if they have some understanding of the hierarchy (the skos:broader and skos:narrower properties) so that for example a search on a particular term could locate all documents indexed not just with that specific term but also any children of the concept represented by that term. Broader filtering could be provided via the skos:ConceptScheme so that only resources indexed within a particular scheme may be considered in a search.

Raw RDF is not particularly easy to understand for end-users; a minimal layer on top is the SPARQL endpoint,[13] which provides a somewhat standardized interface for querying a graph database using the SPARQL language (Prud'hommeaux and Seaborne 2008). Users need to know exactly what properties are used in the models in the graph (dct:subject or bibframe:subject or Wikidata's wdt:P921 for example) but for a given RDF data model this can be a useful and efficient approach. Again, in practice most presentation of items indexed with SKOS vocabularies involves search engines or more customized applications at the top-most user-interface level.

4.2 SKOS tools

A wide variety of both open-source and proprietary tools are available to work with SKOS vocabularies, some of them specifically designed or adapted for the SKOS format. These include tools for developing vocabularies, for validating them, for searching and browsing them and providing a web interface for each concept, and for making use of SKOS in indexing and searching. The following discussion gives a few examples without attempting a comprehensive list.

PoolParty (Schandl and Blumauer 2010) was designed to support SKOS thesaurus creation from the start and now includes additional components for validation and making use of SKOS vocabularies. The FAO AGROVOC vocabulary maintainers developed the general open-source SKOS development tool "VocBench" (Stellato 2014).[14] Many older thesaurus development tools also now include support for export in the SKOS format. Tools to take an existing vocabulary in, for example, a spreadsheet format and convert it to SKOS include SKOSify (Suominen and Hyvönen 2012) and SKOS Play (https://skos-play.sparna.fr/).

The qSKOS tool (Mader, Haslhofer and Isaac 2012)[15] is commonly used for validation and quality checking of SKOS vocabularies; it has been combined with SKOSify (Suominen and Mader 2013) and is also used by PoolParty, although both of those tools began with their own quality-checking components. The quality checks ensure for example that skos:prefLabel values are unique, that all label values have proper language tags, and that hierarchy is logically structured with no orphan concepts or cycles. Some of these quality checks ensure compliance with the SKOS standard while others are advisory and a matter for vocabulary maintainers to determine compliance. Links for this and other SKOS quality-checking tools are available at https://www.w3.org/2001/sw/wiki/SKOS/Validation.

SKOS Play is designed to render and visualize SKOS vocabularies and provides several output formats including a browsable website or interactive visualization. Skosmos (Suominen et al. 2015) is user-friendly software for a SKOS vocabulary allowing searching, browsing, and viewing of each concept. Skosmos is used as the front-end for many SKOS concept schemes available online including the AGROVOC vocabulary and the SKOS vocabularies provided by the Basel Register of Thesauri, Ontologies & Classifications (BARTOC).

4.3 SKOS extensions: SKOS-XL

The difference in treatment of term labels from ISO 25964 was noted above; standard SKOS does treats labels as literal values, not as objects with URIs that can have further properties attached. The SKOS standard includes (Miles and Bechhofer 2009. Appendix B) a SKOS-XL extension (SKOS eXtension for Labels), where each label does become such an object, allowing closer alignment with ISO 25964. With this optional extension every literal string value and language combination becomes a separate member of the skosxl:Label class, with value given by a (single) skosxl:literalForm predicate. SKOS concepts can have skosxl:Label values related through the skosxl:prefLabel, skosxl:altLabel, and skosxl:hiddenLabel predicates, which are the SKOS-XL analogues of the literal-valued skos:prefLabel, etc. predicates discussed earlier.

In addition SKOS-XL defines a predicate to show two labels are related: skosxl:labelRelation. The standard encourages practical use of this through subproperties; for example defining an acronym property that allows linking the SKOS-XL label for an organization name to the label for its acronym.

This extension illustrates a general characteristic of knowledge graphs using RDF: they rarely rely on only a single standard to define their classes and properties. Since everything is either a URI or a literal value, and URIs are indeed universal, it is a simple matter to make use of URIs from a wide variety of different sources in creating an information entity. This may present a bit of a problem for users of such RDF graphs, however, in that they need to be prepared to handle much more than just a single standard for entities and their relationships. Software that is designed to handle a SKOS thesaurus may not understand these additional features, so designers of such vocabularies need to be mindful of their tools and end-users in making use of such extensions. Most of the tools mentioned in the previous section now support SKOS-XL in some form, although that support was often added later rather than available from the start.

5.0 Impact of SKOS

5.1 Vocabularies and thesauri

Thousands of vocabularies and thesauri have been made available as linked open data (LOD) (Allemang, Hendler and Gandon 2020) and a large fraction of them use SKOS classes and properties to describe their structure. A 2012 survey found 478 SKOS vocabularies available (Manaf, Bechhofer, and Stevens 2012) and the number has grown significantly since then; 1214 SKOS vocabularies are provided just on the BARTOC Skosmos server (https://bartoc-skosmos.unibas.ch/) at the time of this writing. Specific examples include the US *Library of Congress Subject Headings* (https://id.loc.gov/authorities/subjects.html), the UN Food and Agriculture Organization *AGROVOC* vocabulary (https://www.fao.org/agrovoc/linked-data), the Getty *Art and Architecture Thesaurus* (AAT) (http://www.getty.edu/research/tools/vocabularies/lod/index.html), the *Unified Astronomy Thesaurus* (https://astrothesaurus.org), PhySH (Smith 2020), and over four hundred vocabularies provided as *Research Vocabularies Australia* by the Australian Research Data Commons (https://vocabs.ardc.edu.au/). Library and information professionals have been encouraged to use SKOS and convert their existing controlled vocabularies to the format (Frazier 2015). Some traditional classification schemes have provided both their standard format and an official SKOS version for use as Linked Data; an example here is recent releases of the MSC (Arndt et al. 2021).

In some instances, the vocabularies or thesauri use additional classes and properties to describe their structure, beyond what is available with SKOS (and SKOS-XL). For example the Getty vocabularies use an RDF representation of the ISO 25964 standard to get around the limitations mentioned above regarding the skos:Collection class; their iso:ThesaurusArray is a subclass of skos:Collection but can be placed within the hierarchy using an iso:superOrdinate property. See http://vocab.getty.edu/doc/ for more details.

5.2 Mappings between vocabularies

While each URI for a SKOS concept is in principle unambiguous in meaning, the ease with which new URIs can be independently created leads to a natural disorder and incompatibility when such a wide variety of indexing is applied. This is not a new problem for controlled vocabularies and has led some to doubt whether they provide any advantage over natural language search (Maniez 1997). Even within a single conceptual framework meaning may change as new knowledge is gained – biological taxa and even the definition of basic terms like species have changed over time, particularly with the advent of genetic analysis with DNA (Minelli 2022). At least some of this can be alleviated by publishing mappings between concept schemes. While the SKOS matching properties are not widely used in most published concept schemes as those schemes generally confine their attention just to relating the concepts they contain, a number of mappings between vocabularies that make use of these match properties have been published in recent years, an encouraging sign for mitigating this issue of incompatibility between vocabularies.

The STW Thesaurus for Economics (Kempf and Neubert 2016) provides an example of a pre-existing thesaurus migrated to SKOS and taking advantage of almost all the features SKOS provides, including the matching properties, which STW uses extensively to map their concepts to equivalent or related concepts in several other online resources, including the German Integrated Authority File (GND), DBPedia, and AGROVOC.

The *Global Agricultural Concept Space* (Baker et al. 2019) is planned to provide a common namespace (i.e., a common URI prefix) for concepts in food and agriculture. It includes a central SKOS concept scheme, GACS Core, and SKOS mapping relations to AGROVOC and other similar concept schemes which are widely used to index bibliographic records and agriculture-related organizations around the world. GACS Core extends its SKOS concepts with a collection of concept types and some special relations (such as gacs:hasProduct to relate an organism to the associated food product), but this is deliberately lightweight and designed for ease of maintainability, rather than providing a comprehensive ontology. A related project with mappings to other vocabularies is maintained by the US Department of Agriculture as the *National Agricultural Library Thesaurus Concept Space* (https://agclass.nal.usda.gov).

The ARIADNE project engaged in a mapping exercise (Binding and Tudhope 2016) between a variety of controlled vocabularies used in archeology datasets and the Getty AAT. Many of the concepts were interchangeable or close and could be mapped with skos:exactMatch or skos:closeMatch; others were more precise than the AAT concepts and were mapped with skos:broadMatch. With the Getty AAT as a central linking hub this allowed all these vocabularies to be unified and allowed items indexed with them to be searched as a coherent whole.

5.3 Non-vocabulary contexts

The SKOS label properties can be used in non-SKOS contexts because their definitions are particularly free of ontological commitment (unlike the relation properties which require skos:Concept instances as subject and object). For example Wikidata uses the SKOS labeling properties skos:prefLabel and skos:altLabel in its RDF dumps to indicate the labels on entities in that database but it does not otherwise define its entities to be skos:Concepts (see https://www.mediawiki.org/wiki/Wikibase/Indexing/RDF_Dump_Format).

In some applications SKOS is used somewhat more thoroughly but as only a small part of a larger ontology. DBPedia (Bizer et al. 2009) has long used SKOS to represent Wikipedia categories. In this case each category is declared as a skos:Concept and the SKOS label and hierarchy properties are used. However, the vast majority of DBPedia entities are for regular Wikipedia pages and are not part of this SKOS subset; instead, they are represented by other aspects of the DBPedia ontology.

The *Quantities, Units, Dimensions and Types* (QUDT) ontology (http://www.qudt.org/) has a concept class of its own (qudt:Concept) which is defined as a subclass of skos:Concept. The ontology then uses skos:altLabel extensively to provide alternate labels for its entities, uses some of the SKOS matching properties to link them to DBPedia, and also uses some of the SKOS documentation properties. Nevertheless, QUDT is not generally seen as a SKOS vocabulary since the primary purpose of the ontology is in the other relations and attributes of its classes, to precisely define quantitative units for data.

More complex ontologies can often be automatically converted to a SKOS representation by reducing the classes of the ontology to a hierarchical list of labeled concepts. The Financial Industry Business Ontology (FIBO. https://spec.edmcouncil.org/fibo/) is provided by its curators both in its full form as an OWL ontology and in a simplified SKOS format as the *FIBO Vocabulary*. In other cases, the SKOS representation may be developed by a third party. To one degree or another SKOS has been found useful in a wide variety of additional contexts beyond the traditional thesaurus or controlled vocabulary systems it was initially designed for.

5.4 New and experimental applications

Indexing documents with a controlled vocabulary is often a labor intensive process: understanding is needed to determine what topics a piece of text is about. Some automation can be done through string-matching and associated rules but designing such automated indexing systems in a reliable way can be complex. Recent advances in natural language processing and machine learning have shown signs that this sort of automation can be done more routinely, although they still need to start with a manually indexed training corpus. The National Library of Finland has produced a freely available software toolkit, Annif (Suominen, Inkinen and Lehtinen 2022) (available at https://annif.org) that can create such automated indexing tools from SKOS concept schemes, and they and others are actively using it to integrate semi-automated subject indexing into their metadata workflows.

Integrating SKOS with other open interoperable standards is a route taken by Skohub (https://skohub.io). This experimental software turns each SKOS concept in a concept scheme into a hub for content indexed with that concept, to which social media or other client applications can subscribe. Indexing software publishes the relationships to SKOS concepts to the Skohub service, which then pushes links to newly indexed documents out to clients that have subscribed to the related concepts. There are other techniques for developing topical feeds of interest to readers but this approach via open Web standards holds promise.

6.0 Considerations for the future

The W3C seems to consider SKOS a finished product; there has been no working group assigned to review or update it in the more than a decade since the recommendation was published. Among Semantic Web technologies SKOS has been reasonably successful, with thousands of published vocabularies and thesauri making use of it and a healthy collection of tools for creating, validating, and viewing or using it. The Semantic Web itself however has followed a different path than the original Web 3.0 (Markoff 2006) vision for it. While there are vast and growing quantities of RDF data available online, mechanisms for creating and making use of it have changed. Instead of relying on SPARQL endpoints (though some still operate usefully) RDF data and other information is now collected into larger Knowledge Graphs (Hogan 2022) used by major online services such as Google.

But the LOD/Semantic Web ecosystem is not the only place where a SKOS vocabulary can be useful. Documents or other items can be indexed using SKOS concepts simply by entering the URI (and/or other metadata) for the SKOS concept into a relational database or search collection. The fact that the identifier is a URI does not really matter for such purposes; just having a unique and persistent ID for a concept allows conceptual grouping of items, and the SKOS relationships can be used to create faceted hierarchies of documents or products for browsing purposes. All of this currently requires custom software, so a more consistent standard for indexing with SKOS vocabularies might allow for more general-purpose software tools in this area.

The fact that SKOS identifiers are URIs presents another area where applications could potentially work better with additional standards. URIs in common use are almost always actually URLs, that is they point to a location on the Internet that can be retrieved, so what should the response to retrieving a SKOS concept or concept scheme URI look like? The SKOS recommendation and associated notes from the SWD don't say, or even require that the URI be resolvable. In practice SKOS URIs usually resolve to some sort of concept scheme browser (like Skosmos) that may provide both HTML and RDF (XML, JSON-LD, *etc*.) representations of the concept scheme or the specific concept. However, at least to this point software using such vocabularies cannot rely on a particular structure for responses from these URIs, and so needs to be tailored to each vocabulary (or at least to each concept space). This may be another area where a more rigid standard or expectation could be set that would be helpful for users of such vocabularies; perhaps the Skosmos design will become a *de facto* standard here.

As has been noted above, some aspects of SKOS have received only limited use. An update to the recommendation to either improve their usability or remove them would help to keep this standard as simple as possible. On the other hand, some additions to the recommendation could also be helpful, even while trying to avoid the full complexity of ISO 25964 or more complete bibliographic systems like BIBFRAME. The issues left unresolved in the development of SKOS (see Section 3.2) have been worked around by implementers in the intervening years but could still benefit from improvements to the standard if that were possible. A standard to indicate non-assignable or non-indexing concepts would be helpful (Panzer and Zeng 2009); some mechanism for combining concepts (Binding, Gnoli and Tudhope 2021) would also be welcome. And defining a structure for notes and notation, similar to the SKOS-XL extension for labels, would permit some classification schemes to have their meaning more fully mapped to a semantic web context (Panzer and Zeng 2009).

SKOS provides a simple data model for a knowledge organization system that could be considered conceptually independent of RDF and the Semantic Web technologies that fostered it. In the longer run no matter how the patterns of RDF usage change it may make sense to contextualize the SKOS data model for other types of computing systems such as search engines and relational databases.

7.0 Conclusions

SKOS has become a widely adopted "simple" standard in the field of knowledge organization, with its ability to capture and easily share in a standardized form the content of controlled vocabularies, thesauri, classification systems and subject headings, although some traditional features of thesauri and classification systems are not supported. As with most standards in the field of knowledge organization SKOS represents a compromise between different needs and interests, and it cannot solve the problem of incompatibility between different KOS.

For a KOS needing additional capabilities (particularly for some types of ontology) SKOS can act as a basis with custom extensions used to address those missing features. Organizations with standardized thesauri in the arts and sciences and even financial and business domains have adopted and published openly available SKOS concept schemes as *linked open data* (LOD), available for anyone to use.

The relation of SKOS to *the Semantic Web* and LOD comes from its definition in terms of the Resource Description Framework (RDF); every entity in a SKOS thesaurus is a URI that could in principle be downloaded or viewed in a web browser. A SKOS concept scheme with its concepts and relations is a graph that can be queried with the SPARQL language. But SKOS can also be used in any context where an identifier is needed for the concepts in a vocabulary or thesaurus, the identifiers just happen to be URIs. And there are a wide range of tools for creating and making use of SKOS vocabularies.

SKOS (with the first word in its full name being Simple) does not support every feature needed for thesauri as described in the ISO 25964 standard. The SKOS-XL extension addresses some of those missing pieces, and third-party extensions are available to cover other features if needed.

The benefits of SKOS as a basic standard for representing and publishing vocabularies are clear, and it deserves to be widely understood by those involved in knowledge organization.

Notes

1. In the literature of knowledge organization, a distinction is often made between "classification systems" versus "verbal indexing languages" (see https://www.isko.org/cyclo/indexing#4.1), where the first group contains non-verbal codes, while the second contains verbal codes. There are, however, examples of purely verbal classification systems (see https://www.isko.org/cyclo/ir#4.1) just as there are examples of thesauri using non-verbal codes in addition to the verbal ones (e.g., the NCI-thesaurus of the National Cancer Institute). It should be said, as just as the notations in classification systems can be labeled in different languages, the preferred terms in thesauri can be multilingual and also assigned synonymous to each descriptor.
2. An anonymous reviewer wrote: "In my view, SKOS URIs almost certainly do not represent exactly the same meaning since, in addition to the implicit meanings given by (different) hierarchical structures, different KOS carry different (community and cultural) perspectives and scope notes of what appears at first glance to be the same terms denoting the same concept. Hence the preferred use in the KOS community of the looser semantics of the SKOS matching properties."
3. See also https://dublincore.org/
4. FOAF is an ontology describing persons, their activities, and their relations to other people and objects.
5. "Best Practice Recipes for Publishing RDF Vocabularies" see https://www.w3.org/2006/07/SWD/ and https://www.w3.org/TR/swbp-vocab-pub/
6. SWD Issue Tracking at https://www.w3.org/2006/07/SWD/track/issues/
7. The early SKOS-core draft is available at https://www.w3.org/TR/swbp-skos-core-spec/
8. ISO 25964-1: 2011(E), §6.3.1 presents two kinds of phrases: adjective phrases (e.g., "cold fusion") and prepositional phrases (e.g., hospitals for children).

9. About the *Data Catalog Vocabulary* (DCAT) see https://www.w3.org/TR/vocab-dcat-2/
10. SKOSify (Suominen and Hyvönen 2012), for example, is designed to do this, see also https://skosify.readthedocs.io/
11. See also https://www.loc.gov/bibframe/docs/index.html
12. BIBFRAME's "subject" property is available at: http://'id.loc.gov/ontologies/bibframe/subject
13. SPARQL is a recursive acronym for "SPARQL Protocol and RDF Query Language", a language for querying, retrieving, and updating data in an RDF graph.
14. See also http://vocbench.uniroma2.it
15. See also https://github.com/cmader/qSKOS

References

Allemang, Dean, James A. Hendler and Fabien Gandon. 2020. *Semantic Web for the Working Ontologist: Effective Modeling for Linked Data, RDFS, and OWL*. 3rd ed. New York: ACM Books.

Arndt, Susanne, Patrick Ion, Mila Runnwerth, Moritz Schubotz and Olaf Teschke. 2021. "10 Years Later: The Mathematics Subject Classification and Linked Open Data." In: *Intelligent Computer Mathematics. CICM 2021,* edited by Fairouz Kamareddine and Claudio Sacerdoti Coen. *Lecture Notes in Computer Science*, vol 12833. Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-030-81097-9_12

Baker, Thomas, Sean Bechhofer, Antoine Isaac, Alistair Miles, Guus Schreiber and Ed Summers, 2013. "Key Choices in the Design of Simple Knowledge Organization System (SKOS)." *Journal of Web Semantics* 20:35-49. https://doi.org/10.1016/j.websem.2013.05.001

Baker, Thomas, Brandon Whitehead, Ruthie Musker and Johannes Keizer. 2019. "Global Agricultural Concept Space: Lightweight Semantics for Pragmatic Interoperability". *npj Science of Food* 3:16. https://doi.org/10.1038/s41538-019-0048-6

Berners-Lee, Tim and Mark Fischetti, 1999. "Mind to Mind." In *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor. San Francisco, CA: Harper, 157-76.*

Berners-Lee, Tim, James Hendler and Ora Lassila, 2001. "The Semantic Web." *Scientific American* 284, no. 5:34-43

Biagetti, Maria Teresa. 2021. "Ontologies as Knowledge Organization Systems". *Knowledge Organization* 48, no. 2:152-76. Also available in *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli. https://www.isko.org/cyclo/ontologies

Binding, Ceri and Douglas Tudhope. 2016. "Improving Interoperability Using Vocabulary Linked Data." *International Journal on Digital Libraries* 17:5–21. https://doi.org/10.1007/s00799-015-0166-y

Binding, Ceri, Claudio Gnoli and Douglas Tudhope. 2021. "Migrating a Complex Classification Scheme to the Semantic Web: Expressing the Integrative Levels Classification using SKOS RDF." *Journal of Documentation* 77: 926-45. https://doi.org/10.1108/JD-10-2020-0167

Bizer, Christian, Jens Lehmann, Georgi Kobilarov, Sören Auer, Christian Becker, Richard Cyganiak and Sebastian Hellmann. 2009. "Dbpedia: A Crystallization Point for the Web of Data." *Journal of Web Semantics* 7: 154-65. https://doi.org/10.1016/j.websem.2009.07.002

Brickley, Dan and Libby Miller. 2005. *FOAF Vocabulary Specification*, FOAF Project. http://xmlns.com/foaf/spec/

Dextre Clarke, Stella G. 2019. "The Information Retrieval Thesaurus." *Knowledge Organization* 46: 439-59. Also available in *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli. http://www.isko.org/cyclo/thesaurus

Fraser, Craig. 2020. "Mathematics in Library and Review Classification Systems: An Historical Overview." *Knowledge Organization* 47: 334-56. Also available in *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli. https://www.isko.org/cyclo/mathematics

Frazier, Priscilla Jane. 2015. "SKOS: A Guide for Information Professionals." American Library Association, May 18, 2015. http://www.ala.org/alcts/resources/z687/skos

Hawkins, Les. 2015. "The Semantic Web and the BIBFRAME Initiative." *Serials Review* 41, no. 2: 106-7.

Hogan, Aidan et al. 2022. "Knowledge Graphs." *ACM Computing Surveys* 54, no. 4: 71:1-37. https://doi.org/10.1145/3447772

International Organization for Standardization. 1974. *ISO 2788-1974. Documentation - Guidelines for the Establishment and Development of Monolingual Thesauri*. Geneva: International Organization for Standardization.

International Organization for Standardization. 1985. *ISO 5964-1985. Documentation - Guidelines for the Establishment and Development of Multilingual Thesauri*. Geneva: International Organization for Standardization.

International Organization for Standardization. 2011. *ISO 25964-1: Information and Documentation - Thesauri and Interoperability with other Vocabularies - Part 1: Thesauri for Information Retrieval*. Geneva: International Organization for Standardization.

International Organization for Standardization. 2013. *ISO 25964-2: Information and Documentation - Thesauri and Interoperability with other Vocabularies - Part 2: Interoperability with other Vocabularies*. Geneva: International Organization for Standardization.

International Organization for Standardization. 2017. *ISO 15836-1:2017 Information and Documentation - The Dublin Core Metadata Element Set - Part 1: Core Elements*. Geneva: International Organization for Standardization.

Isaac, Antoine, Jon Phipps and Daniel Rubin. 2009. *SKOS Use Cases and Requirements.* World Wide Web Consortium. http://www.w3.org/TR/skos-ucr/

Isaac, Antoine and Ed Summers (Eds.), 2009. *SKOS Simple Knowledge Organization System Primer.* World Wide Web Consortium. https://www.w3.org/TR/skos-primer/

ISO TC46/SC9/WG8 working group and Antoine Isaac, 2012. *Correspondence between ISO 25964 and SKOS/SKOS-XL Models*, International Organization for Standardization. https://groups.niso.org/apps/group_public/download.php/12351/Correspondence%20ISO25964-SKOSXL-MADS-2013-12-11.pdf

Kempf, A. and J. Neubert. 2016. "The Role of Thesauri in an Open Web: A Case Study of the STW Thesaurus for Economics." *Knowledge Organization* 43:160-73.

Lardera, Marco and Birger Hjørland. 2020. "Keyword." In *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli. http://www.isko.org/cyclo/keyword

Lassila, Ora and Ralph R. Swick (Eds.). 1999. *Resource Description Framework (RDF) Model and Syntax Specification*. World Wide Web Consortium. https://www.w3.org/TR/1999/REC-rdf-syntax-19990222/ Note this has been replaced by updates in 2004 and 2014 that now define the current RDF 1.1 standard.

Mader, Christian, Bernhard Haslhofer and Antoine Isaac. 2012. "Finding Quality Issues in SKOS Vocabularies." *Proceedings of Theory and Practice of Digital Libraries* (TPDL 2012). Lecture Notes in Computer Science 7489. Berlin: Springer-Verlag, 222–33.

Manaf, Nor Azlinayati, Sean Bechhofer and Robert Stevens. 2012. "The Current State of SKOS Vocabularies on the Web." *Lecture Notes in Computer Science* 7295. Berlin: Springer, 270–84. https://doi.org/10.1007/978-3-642-30284-8_25

Maniez, Jacques. 1997. "Database Merging and the Compatibility of Indexing Languages." *Knowledge Organization* 24: 213-24.

Markoff, John. 2006. "Entrepreneurs See a Web Guided by Common Sense." *New York Times* November 12 2006. https://www.nytimes.com/2006/11/12/business/12web.html

Mazzocchi, Fulvio. 2018. "Knowledge Organization System (KOS)". *Knowledge Organization* 45: 54-78. Also available in *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli. http://www.isko.org/cyclo/kos

Miles, Alistair, Nikki Rogers and Dave Beckett. 2004. *SKOS-Core 1.0 guide*. World Wide Web Consortium. http://www.w3.org/2001/sw/Europe/reports/thes/1.0/guide/20040504/

Miles, Alistair and Sean Bechhofer (Eds.). 2009. *SKOS Simple Knowledge Organization System Reference*. World Wide Web Consortium. https://www.w3.org/TR/2009/REC-skos-reference-20090818/

Minelli, Alessandro. 2022. "Species." In *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli. http://www.isko.org/cyclo/species

Panzer, Michael and Marcia Lei Zeng. 2009. "Modeling Classification Systems in SKOS: Some Challenges and Best-Practice Recommendations." *Proceedings of the International Conference on Dublin Core and Metadata Applications 2009.* DCMI, 3-14. https://dcpapers.dublincore.org/pubs/article/view/974

Prud'hommeaux, Eric and Andy Seaborne (Eds.). 2008. *SPARQL Query Language for RDF.* World Wide Web Consortium. https://www.w3.org/TR/rdf-sparql-query/

Schandl, Thomas and Andreas Blumauer. 2010. "PoolParty: SKOS Thesaurus Management Utilizing Linked Data." In *The Semantic Web: Research and Applications. ESWC 2010*, edited by Lora Aroyo et al. *Lecture Notes in Computer Science*, vol 6089. Berlin: Springer, 421-5. https://doi.org/10.1007/978-3-642-13489-0_36

Smith, Arthur. 2020. "Physical Subject Headings (PhySH)." *Knowledge Organization* 47: 257-66. Also available in *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli, https://www.isko.org/cyclo/physh

Smith, Barry, Michael Ashburner, Cornelius Rosse, Jonathan Bard, William Bug, Werner Ceusters, Louis J Goldberg, Karen Eilbeck, Amelia Ireland, Christopher J Mungall, OBI Consortium; Neocles Leontis, Philippe Rocca-Serra, Alan Ruttenberg, Susanna-Assunta Sansone, Richard H Scheuermann, Nigam Shah, Patricia L Whetzel and Suzanna Lewis. 2007. "The OBO Foundry: Coordinated Evolution of Ontologies to Support Biomedical Data Integration." *Nature Biotechnology* 25: 1251-5. https://doi.org/10.1038/nbt1346

Stellato, Armando. 2014. "The History of VocBench, What We have been Doing during the last 15 Months and What the Future Holds." *Agricultural Information Management Standards Portal*. Food and Agricultural Organization of the United Nations. http://aims.fao.org/in-action/interviews/vocbench

Suominen, Osma and Eero Hyvönen. 2012. "Improving the Quality of SKOS Vocabularies with Skosify." In *Proceedings of the 18th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2012) Lecture Notes in Computer Science*, vol 7603, Berlin: Springer, 383-97.

Suominen, Osma and Christian Mader. 2013. "Assessing and Improving the Quality of SKOS Vocabularies." *Journal on Data Semantics* 3: 47–73. https://doi.org/10.1007/s13740-013-0026-0

Suominen, Osma, Henri Ylikotila, Sini Pessala, Mikko Lappalainen, Matias Frosterus, Jouni Tuominen, Thomas Baker, Caterina Caracciolo and Armin Retterath. 2015. "Publishing SKOS vocabularies with Skosmos." https://skosmos.org

Suominen, Osma, Juho Inkinen and Mona Lehtinen. 2022. "Annif and Finto AI: Developing and Implementing Automated Subject Indexing." *JLIS.It* 13: 265–82. https://doi.org/10.4403/jlis.it-12740

Zeng, Marcia Lei. 2019. "Interoperability." *Knowledge Organization* 46: 122-146. Also available in *ISKO Encyclopedia of Knowledge Organization*, edited by Birger Hjørland and Claudio Gnoli. http://www.isko.org/cyclo/interoperability