

On the composition of ISO 25964 hierarchical relations (BTG, BTP, BTI)

Vladimir Alexiev¹ · Antoine Isaac² · Jutta Lindenthal³

Received: 12 January 2015 / Revised: 29 July 2015 / Accepted: 4 August 2015 / Published online: 20 August 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract Knowledge organization systems (KOS) can use different types of hierarchical relations: *broader generic* (BTG), *broader partitive* (BTP), and *broader instantial* (BTI). The latest ISO standard on thesauri (ISO 25964) has formalized these relations in a corresponding OWL ontology (De Smedt et al., ISO 25964 part 1: thesauri for information retrieval: RDF/OWL vocabulary, extension of SKOS and SKOS-XL. <http://purl.org/iso25964/skos-thes,2013>) and expressed them as properties: *broaderGeneric*, *broaderPartitive*, and *broaderInstantial*, respectively. These relations are used in actual thesaurus data. The compositionality of these types of hierarchical relations has not been investigated systematically yet. They all contribute to the general *broader* (BT) thesaurus relation and its transitive generalization *broader transitive* defined in the SKOS model for representing KOS. But specialized relationship types cannot be arbitrarily combined to produce new statements that have the same semantic precision, leading to cases where inference of *broader transitive* relationships may be misleading. We define Extended properties (BTGE, BTPE, BTIE) and analyze which compositions of the original “one-step” properties and the Extended properties are appropriate. This enables providing the new properties with valuable semantics usable, e.g., for fine-grained information retrieval purposes.

In addition, we relax some of the constraints assigned to the ISO properties, namely the fact that hierarchical relationships apply to SKOS concepts only. This allows us to apply them to the Getty Art and Architecture Thesaurus (AAT), where they are also used for non-concepts (facets, hierarchy names, guide terms). In this paper, we present extensive examples derived from the recent publication of AAT as linked open data.

Keywords Thesauri · ISO 25964 · BTG · BTP · BTI · Broader generic · Broader partitive · Broader instantial · AAT

1 Introduction

One of the most important relations in knowledge organization systems (KOS) are the hierarchical relations. The main, general type of hierarchical relation, referred to as *broader (term)* (BT) in thesaurus standards, denotes a relationship from a specific entity to a more generic one—its direct “parent” in a conceptual hierarchy. The SKOS model for publishing KOS as semantic web resources [1] has formalized it as the *skos:broader* property. But thesauri may use various types of more specialized hierarchical relations (see [2, Sect. 3.2]). We reiterate their definition below, with examples from the Getty vocabularies: the Art and Architecture Thesaurus (AAT), the Union List of Artist Names (ULAN) and the Thesaurus of Geographic Names (TGN) [3]:

The latest ISO standard on thesauri ISO 25964 [4] has formalized these relations: its domain model includes a field *HierarchicalRelationship.role*. The corresponding OWL ontology [5] expresses them as properties in the ‘iso-thes.’ namespace: *broaderGeneric*, *broaderPartitive* and *broaderInstantial*.

✉ Antoine Isaac
aisaac@few.vu.nl
Vladimir Alexiev
vladimir.alexiev@ontotext.com
Jutta Lindenthal
jutta.lindenthal@gmail.com

¹ Ontotext Corp, Sofia, Bulgaria

² Europeana and VU Amsterdam, The Hague, The Netherlands

³ Independent consultant, Lübeck, Germany

Table 1 Hierarchical relation types

Relation	Abbrev.	Name	Example
Broader generic	BTG	Species/genus relation	Calcite is a <i>kind of</i> mineral (AAT)
Broader partitive	BTP	Part/whole relation	Tuscany is a <i>part of</i> Italy (TGN)
Broader instantial	BTI	Instance/kind relation	Rembrandt is an <i>instance of</i> (example of) a person (ULAN)

Similar relations have been defined and used in previous or parallel efforts to represent fine-grained thesaurus data:

- The German Gemeinsame Normdatei (GND) linked data service and the GND ontology [6]. E.g., the representation of *SG Dynamo Dresden*¹ has BTI relationship to *football clubs*:
(<http://d-nb.info/gnd/5055902-3>)
gnd:broaderTermInstantial
(<http://d-nb.info/gnd/4155742-6>).
- The FinnONTO SKOS extensions define properties broaderGeneric, broaderPartitive [7].
- Some vocabularies hosted by digiCULT-Verbund eG., maintained in the vocabulary management tool xTree, and represented in the exchange format vocnet [8].

Most recently, the Getty vocabulary program (GVP) thesauri (AAT, TGN, ULAN) have been published as linked open data. This raised the issue of which properties to reuse from these efforts, and which formal semantics are really suited for the data at hand in the Getty vocabularies [9, Sect. 2.3 Subject Hierarchy].

1.1 Hierarchical relations in Getty thesauri

Table 1 shows statistics of the three types of hierarchical relations for each GVP thesaurus as of May 2015. The count is from the following SPARQL² query in the GVP SPARQL endpoint³:

```
select ?thes ?rel ?typ (count(*) as ?c) {
  {select * {values ?rel
    {gvp : broaderGeneric gvp : broaderPartitive
     gvp : broaderInstantial}}}
  ?x ?rel ?y.
  ?x a ?typ; skos : inScheme ?thes.
```

¹ <http://d-nb.info/050559028/about/rdf>.

² The RDF query language, see <http://www.w3.org/TR/sparql11-query/>.

³ The query is quite expensive, please use <http://getty.ontotext.com/sparql>.

```
?typ rdfs : subclassOf gvp : Subject filter
  (?typ != gvp : Subject)
} group by ?thes ?rel ?typ
```

The query iterates over the three kinds of relations using VALUES. For each thesaurus entity ?x (a gvp:Subject), it fetches the thesaurus (skos:inScheme) and its “Record Type” ?typ (as a proper subclass of gvp:Subject). The types concept, administrative place, physical place, person, group (corporate body) and unknown person are the main thesaurus entities, intended to be used for indexing. The types facets, guide terms and hierarchy names are not used for indexing; they merely serve to structure the hierarchy.

In Table 2 below, the left column shows the thesaurus and entry type, the top row is the relationship having that type as its source (subject), and the cells show relationship counts.

A brief analysis of these numbers follows.

In AAT, most relations are BTG, but there are some BTP. There is a variety of situations, including

- Concept BTP concept: *calendars of relics* BTP *cabinets of relics*.
- Concept BTP guide term: *anvil components* BTP (*anvils and anvil accessories*).
- Guide term BTP concept: (*jewelry and accessory components*) BTP *jewelry*.
- Guide term BTP guide term: (*grinding and milling equipment components*) BTP (*grinding and milling equipment*).
- Concept BTP hierarchy name: *building divisions* BTP *single built works*.

In TGN, most relations are BTP (a place is part of another place). TGN place types (e.g., *inhabited place*, *seaport*, or *commercial center*) are AAT concepts, and we considered mapping the *place type* relation to a sub-property of BTI (e.g., *Sofia* BTI *inhabited place*). This would support use cases like the one described in Sect. 3.1. However, this has been postponed, pending a better understanding of place type hierarchies in AAT.

In ULAN, most relations are BTI, e.g., *Rembrandt* BTI *persons, artists*; *J. Paul Getty Trust* BTI *corporate bodies*. There are some BTP, e.g., *Getty Research Institute* BTP *J. Paul Getty Trust*.

Table 2 Breakdown of relation counts in GVP vocabularies

Thes\type	BTG	BTI	BTP	Total
AAT	44,953		456	45,409
Concept	43,350		450	43,800
GuideTerm	1564		6	1570
HierarchyName	38			38
TGN	1		1,483,515	1,483,516
GuideTerm			62	62
AdminPlaceConcept	1		792,089	792,090
PhysPlaceConcept			691,307	691,307
PhysAdminPlaceConcept			57	57
ULAN	2	217,155	17,754	234,911
GuideTerm	2			2
GroupConcept		23,227	16,694	39,921
PersonConcept		191,705	1057	192,762
UnknownPersonConcept		2223	3	2226
Total	44,956	217,155	1,501,725	1,763,836

1.2 The problem

As far as we know, there are very few datasets yet that use the ISO hierarchical relation properties. One reason is the relative novelty of this ontology (created 2013-12-09), but we see two other reasons:

Improper axiomatization for property composition

The ISO relations all are sub-properties of `skos:broader`, being a sub-property of `skos:broaderTransitive`, which is transitive. If we have these statements:

```
concept1 iso-thes:broaderGeneric concept2.
concept2 iso-thes:broaderPartitive concept3.
```

we can infer these statements:

```
concept1 skos:broader concept2.
concept2 skos:broader concept3.
concept1 skos:broaderTransitive concept3.
```

In general, we have the following inference for any chain of BTG, BTP, BTI:

```
(broaderGeneric|broaderPartitive|broaderInstantial)
→ broader → broaderTransitive.
```

But such composition of relationships is inappropriate in some cases: [10] argues that `broaderTransitive` is less error-prone if established separately for BTG and BTP, and neither for BTI, nor mixed paths of BTG and BTP.

In the SKOS context, the inference of `skos:broaderTransitive` statements should not be harmful. The SKOS specification assigns this property only extremely lightweight

semantics: it merely means that one concept is the ancestor of another in a KOS hierarchy, whatever semantic flavor the individual ‘parent’ relationships between intermediate concepts may have [11]. However, the KOS community generally knows that it is possible to combine hierarchical links in a finer-grained way, i.e., one that produces statements using the richer, specialized types of relations. As a result, the existence of `skos:broaderTransitive` has tended to raise expectations that its rather weak semantics cannot meet. There was a lively discussion about this subject on the SKOS mailing list from Nov 2013 to April 2014.⁴ ISO’s extension of SKOS could have solved the issue by defining `iso-thes:broaderGeneric`, `iso-thes:broaderPartitive` or `iso-thes:broaderInstantial` with appropriate compositional semantics (for example, by stating that `broaderPartitive` is transitive). But until now the ISO standard has postponed the definition of advanced formal semantics for these properties.

Note that our investigation relies on vocabularies where hierarchical relations follow the logical rules as described in the ISO thesaurus standard.

*Missing broader/narrower relationships for non-concepts*⁵ ISO 25964 defines the hierarchical properties to be applicable to concepts only. But important thesauri like AAT need to use them also for other resources (facets, guide terms and hierarchies as shown above) that take full part in defining hierarchies, while not qualifying as concepts according to the standard (to be used in indexing).

⁴ <http://lists.w3.org/Archives/Public/public-esw-thes/2013Nov/>.

⁵ The term “non-concept” is used in this paper to refer to AAT record types as well as vocabulary classes which are not of type `skos:Concept`.

Table 3 Summary of compositionality rules

First/second	BTG _x	BTP _x	BTI _x
BTG _x	BTGE	BTPE	n/a
BTP _x	BTPE	BTPE	No
BTI _x	BTIE	No	n/a

OWL ontology language, this is written as a sub-property axiom:

```
gvp:broaderGeneric rdfs:subPropertyOf
  gvp:broaderGenericExtended.
```

We then analyze which compositions of the original “one-step” properties and the Extended properties are appropriate. We use property chains (denoted by “/”) and analyze appropriate inferences (“→”) case by case. On the left side, “BT*x” means “BT* or BT*E”, i.e., the left member of the inference rule at hand can match either an original ‘one-step’ statement or a statement inferred from other statements, in a recursive fashion. Table 3 is a summary of our findings. First/second represent the first and second property in a chain. “n/a” means that we consider such a situation should not appear in a vocabulary, while “no” means no relation should be inferred.

BTG_x/BTG_x → BTGE

If *X* is a kind of *Y* and *Y* is a kind of *Z* then *X* is a kind of *Z*.

Example (AAT): *baking pans* BTG *bakeware* BTG (vessels for cooking food) implies *baking pans* BTGE (vessels for cooking food). In OWL, this is written as a property chain axiom:

```
gvp:broaderGenericExtended owl:propertyChainAxiom
  (gvp:broaderGenericExtended gvp:broaderGeneric
  Extended).
```

BTG_x/BTP_x → BTPE

If *X* is a kind of *Y*, which is part of *Z* then *X* is part of *Z* (since *X* can play the role of *Y*).

Example (AAT): *beak irons* BTG *anvil components* BTP (anvils and anvil accessories) implies *beak irons* BTPE (anvils and anvil accessories).

BTG_x/BTI_x → n/a

X BTG *Y* means that *X* is a subclass of *Y*, so *Y* is a generic concept. But *Y* BTI *Z* means that *Y* is an individual (named entity) of type *Z*. In the context of one KOS it is not appropriate to have *Y* be both a class and an individual. (Perhaps *Y* can be treated as an individual in one KOS and a class in another KOS). Indeed, in the GVP thesauri, such a situation does not appear, as can be checked with this query:

```
select * {?x gvp:broaderGeneric ?y. ?y
  gvp:broaderInstantial ?z}
```

BTP_x/BTG_x → BTPE

If *X* is part of *Y*, which is kind of *Z* then *X* is part of *Z* (since *Y* can play the role of *Z*).

Example (AAT): *anvil components* BTP (anvils and anvil accessories) BTG (forging and metal-shaping tools) implies *anvil components* BTPE (forging and metal-shaping tools).

BTP_x/BTP_x → BTPE

If *X* is part of *Y* and *Y* is part of *Z* then *X* is part of *Z*.

Example (TGN): *Sofia* BTP *Bulgaria* BTP *Europe* implies *Sofia* BTPE *Europe*.

Note that mereological relationships are not a semantically homogenous class, and unrestricted use of part-of relationships may break transitivity. There are many counterexamples for transitive part-whole relationships:

- *Mick Jagger's thumb* is part of *Mick Jagger*, who is part of *The Rolling Stones*. Is *Mick Jagger's thumb* part of *The Rolling Stones*? [15, p. 133]. It depends on the point of view, but likely it will not be true, since here we are mixing different mereological relations (extrinsic property “membership” vs. intrinsic property “part-of”), and different categorial entities (*Mick Jagger* is a person, and the *Rolling Stones* are a group).
- *Istanbul* is part of *Turkey*, which is part of *Asia* as well as *Europe*; yet without taking further attributes into account, it is not decidable if *Istanbul* is part of *Asia*, or *Europe*, or both. (Actually it is the only metropolis to date extending over two continents.) Here we have to take care not to mistake partial inclusion for full inclusion.
- *Netherlands Antilles* BTP *Netherlands* BTP *Europe* (this was true until 1954 and is in TGN with historic date qualification). Yet, *Netherlands Antilles* BTPE *Europe* is not true. Here we are mixing administrative inclusion with geographic (physical) inclusion.
- *Chicken feet are part of chicken*, which is part of *chicken soup*, yet you would not normally put chicken feet in soup (chicken feet are considered a delicacy in some cuisines, but not in soup). Here, we are mixing member meronyms with substance meronyms.

References [12,13] give a lot more examples of variegated part-of relations. Reference [13] defines six kinds of part-of relations: component/integral object, member/collection, portion/mass, stuff/object, feature/activity, place/area. Wordnet also distinguishes between three kinds of meronym: member, part, substance. See [14] for an excellent introduction to the topic.

Reference [16] considers the ISO 2788 subtypes of BTP (systems and organs of the body, geographical locations,

Table 4 Sample searches using type and parent place

Name	Ancestor	Ancestor	Results
Sofia	Mexico	Chihuahua	Places called “Sofia” in Mexico, Chihuahua state
Sofia	Ranch	Mexico	Ranchos called “Sofia” in Mexico
	Religious	China	Religious centers in China
	Orthodox	Europe	Orthodox Eastern (Christian), Orthodox Jewish, or Orthodox Caliphate centers in Europe
	Undersea	Ocean	Undersea features (e.g., caves, ridges) in oceans

disciplines/fields of study, social structures). Further, it considers a sub-property hierarchy for RT (related).

But ISO 25964 defines only one variety of BTP, so these nuances are lost. Further, our purpose in this paper is simpler, to analyze the interactions of BTG, BTP, BTI.

Note: the standard restricts part-whole relationships to situations of unique part of the specific whole, e.g., a *car wheel* is uniquely a part of a *car*, not a part of a *bicycle* (at least not usually). Since *chicken*, including *feet*, is not uniquely and exclusively part of *chicken soup*, this relationship would not qualify as BTP according to the ISO standard rules. (Rather, the relationship should be represented by an associative relationship). But as AAT examples in this paper show, this restriction is often not followed in actual thesauri, even well-organized ones.

$BTP_x/BTI_x \rightarrow no$

Counterexample: *Sofia* BTP *Bulgaria* BTI *country*. But *Sofia* BTI *city* (or *inhabited place*), and there is no broader relation between *city* and *country* at all.

$BTI_x/BTG_x \rightarrow BTIE$

Example: *Mt Athos* BTI *orthodox religious center* BTG *Christian religious center* implies *Mt Athos* BTIE *Christian religious center*.

$BTI_x/BTP_x \rightarrow no$

Counterexample: *Statue of Liberty pedestal* BTI *pedestals*, some of which are part of some *statues*. But that particular *pedestal* is neither BTI nor BTP *statues* in general. In the figure below, we could infer the dashed relation (generalization of BTP from one instance thereof): i.e., $X \text{ necessary BTP } Y$ and $Z \text{ BTI } X$ and $T \text{ BTI } Y \rightarrow Z \text{ BTP } T$. But in the case under consideration, we have only three nodes in sequence, not four nodes.

$BTI_x/BTI_x \rightarrow n/a$

If X is an instance of Y , which is an instance of Z , then Z is a metaclass.⁸ Metaclasses have found very useful applications in programming. OWL also allows classes of classes, notably via a technique called punning [17]. However, we are

not aware of any useful examples in thesauri, so we disallow this inference.

Finally, we define the union of the Extended relations as $BTGE|BTPE|BTIE \rightarrow gvp:broaderExtended$. It allows a user to access the Extended hierarchies uniformly, yet consistently (unlike $skos:broaderTransitive$, whose grain is not so fine as it combines all flavors together to produce statements that reflect the—sometimes useful but essentially vague— notion of “ancestor”).

3 Using the new properties

In this section, we present examples of how the developed Extended relations can be used.

3.1 Query expansion in information retrieval

The main purpose of hierarchical relationships with fine-grained semantics is to ensure that query expansion will yield reliable results in information retrieval. For example:

- If *Sofia* BTP *Bulgaria* BTP *Europe* then *Sofia* BTPE *Europe*. This enables a search for places in *Europe* to also find *Sofia*.
- If *Mt Athos* BTI *orthodox religious centers* BTG *Christian religious centers* BTG *religious centers* then *Mt Athos* BTIE *religious centers*. This enables a search for *religious centers* to also find *Mt Athos*.

Note that if full query expansion over hierarchies is employed in such cases, the retrieval interface should facilitate the choice of whether or not to include individuals linked by an instance relationship in the query results. Otherwise the user might be overwhelmed with unwanted results.

Reference [18, Sect. 2.1] discusses using BTI to represent TGN “place type” and $broaderExtended$ to access a place’s “ancestors” (be that place’s types or parent places). The user can enter keywords for the place name or any of its ancestors, enabling searches like (Table 4):

⁸ <http://en.wikipedia.org/wiki/Metaclass>.

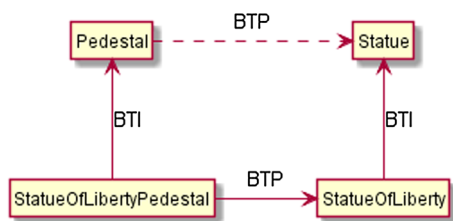


Fig. 2 Inference between BTI and BTP




Fig. 3 Inference between BTG and BTP

This is not implemented in TGN because a chain like *Sofia* BTP *Bulgaria* BTP *Europe* BTI *Continents* would infer *Sofia* skos:broaderTransitive *Continents*, as explained in Sect. 1.2. It was judged that such statements would be too confusing for users, even with warnings in the documentation.

Furthermore, the relevance of BTGE chains may decrease as one goes higher in a thesaurus hierarchy, towards increasingly abstract concepts. E.g., AAT includes this secondary BTG chain: *continents* BTG *landmasses* BTG *landforms* BTG *hypsographic features* BTG *earth sciences concepts* BTG *physical science concepts* BTG *scientific concepts*. It would be useful to type *Europe* as a *hypsographic feature* (as opposed to a *hydrographic* or *vegetal feature*), but going up to *scientific concepts* is hardly useful.

3.2 Beyond chain inferences

Figure 2 provided an example of non-chain inference. We can imagine many more such cases, e.g., X necessary BTP Y and Z BTG $Y \rightarrow X$ BTP Z as for the following case where keyboards would be a necessary part of keyboard instruments (Fig. 3):

3.3 Quality checking

Properly inferred Extended relations can be used to detect mistakes in the hierarchical structure of a vocabulary. Employing inference rules for quality checking is subject of future research, but here we provide a couple of examples.

3.3.1 Disjoint properties

There (used to be) a mistake in AAT: *swell boxes* BTG *organ components* (correct), and *swell boxes* BTG *organs (aerophones)* (incorrect). After the inference *swell boxes* BTG

organ components BTP *organs (aerophones)* \rightarrow *swell boxes* BTPE *organs (aerophones)*, we can catch this mistake if we declare that BTGE and BTPE are disjoint properties. Note that SKOS has similar disjointness declarations, e.g., between skos:broader and skos:related.

3.3.2 Disjoint children

The AAT guide term (*containers by function or context*) is a division by function. Such a division should be orthogonal, i.e., a disjoint union of its children. The orthogonality can be confirmed using the hierarchical view⁹ at the Getty site. Some of the children names allow non-disjoint interpretation, e.g., one could suppose that (*containers for personal use*) and (*containers for textiles and needlework*) can have a common concept. But closer examination of the former shows that their semantic coverage is indeed disjoint: it includes implements for health care, hygiene, tobacco use, and personal gear, none of which are related to textile/needlework.

Figure 4¹⁰ shows that AAT *palette cups*¹¹ is both a direct child of (*containers by function or context*) (correct), and a descendant (chain of length 7) through *cups (drinking vessels)* (incorrect).

Palette cups are “*Cups which attach to a painter’s palette and are used for holding small quantities of medium or solvent*”. So a *palette cup* is in fact not a *cup (drinking vessel)*: although the majority of cups are drinking vessels, this one is not (solvent-drinking creatures excluded).

After inferring the Extended relations, this semantic mistake can be detected using the disjointness of divisions by function. It cannot be detected by following the chain of isothes: properties only. After detecting potential issues, manual inspection is required to pinpoint the problem [*palette cups* are not *cups (drinking vessels)*] but inferencing can at least flag the issue.

3.4 Inferring ISO 25964 and SKOS relations in AAT

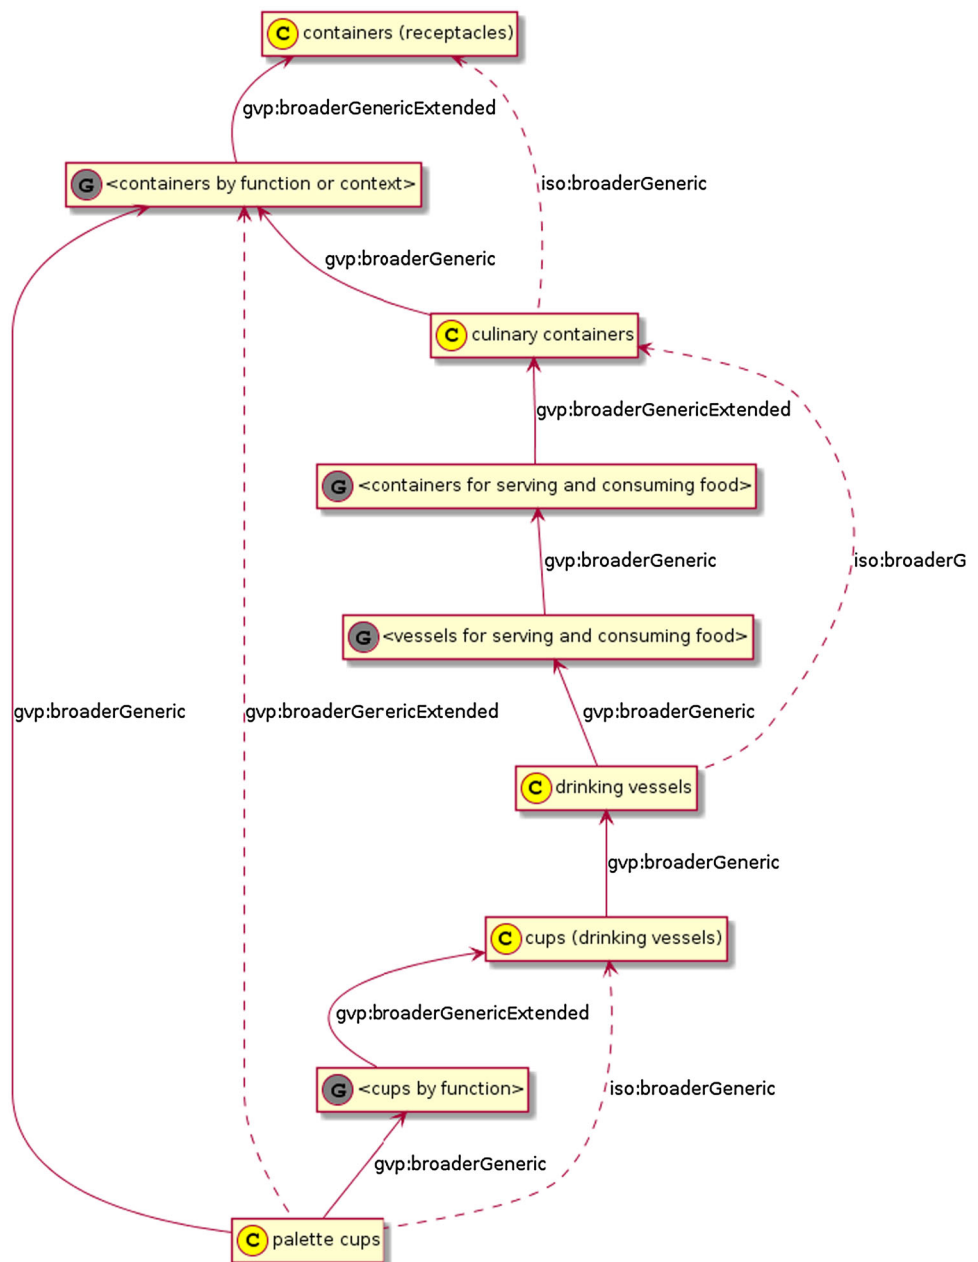
As explained earlier, we could not use directly the ISO 25964 relations in AAT since the ISO relations apply to skos:Concept only. AAT organizes its hierarchy of concepts using “non-concept” types of resources (facets, guide terms, hierarchies) that do not mirror the patterns of the ISO 25964 or SKOS ontologies. It may be possible to reconcile the approaches at hand, but this is not in the scope of this paper. The specific AAT relationships had to be expressed in the published data, as existing AAT usage may depend heavily

⁹ <http://vocab.getty.edu/hier/aat/300197200>.

¹⁰ In Figs. 4 and 5, *C* concept, *G* guide term.

¹¹ <http://vocab.getty.edu/aat/300215585>.

Fig. 4 The palette cups problem



on them. So for instance instead of `iso-thes:broaderPartitive`, we had to use our own property `gvp:broaderPartitive`. However, we would like to infer a SKOS/ISO relation between two concepts when they are connected directly by an appropriate Extended relation.

Even if the AAT model differs from ISO, we argue that the compositional semantics that one can associate to the generic, partitive and instantial “flavors” of hierarchical relations are essentially identical in both models. The Extended relations capture these “appropriate chains”, so we can fulfill our requirement by declaring inference rules that derive ISO/SKOS relationships from our Extended relations.

For example, consider the hierarchical context of *anvil components*¹² in Fig. 5. For brevity, we have omitted some nodes above *anvils* and *anvil accessories*, and additional relations `skos:member` and `iso-thes:superOrdinate` have been removed for clarity (examples of these relations are visible in Fig. 1):

We want to infer the dashed ISO properties from the `gvp:properties`. When two concepts are connected by GVP’s `BT*E` and there are no intervening concepts, we want to infer ISO `BT*` and `skos:broader`.

¹² <http://vocab.getty.edu/aat/300024872>.

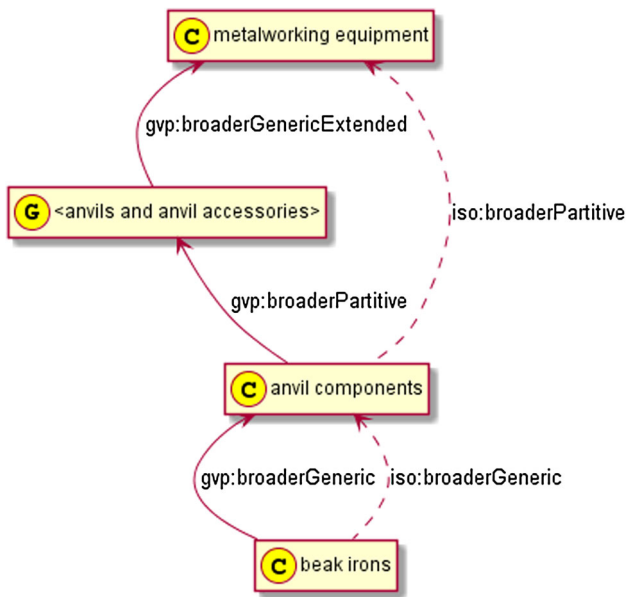


Fig. 5 Inferring ISO relations between concepts

The full reasoning details are provided in [9, Sect. 6.1] (in particular 6.1.10). Figure 6 below is a simplified version that sums up the formal semantics linking our Extended hierarchical properties to the ISO 25964 and the core SKOS properties. Hollow arrows indicate sub-property links, “concept–

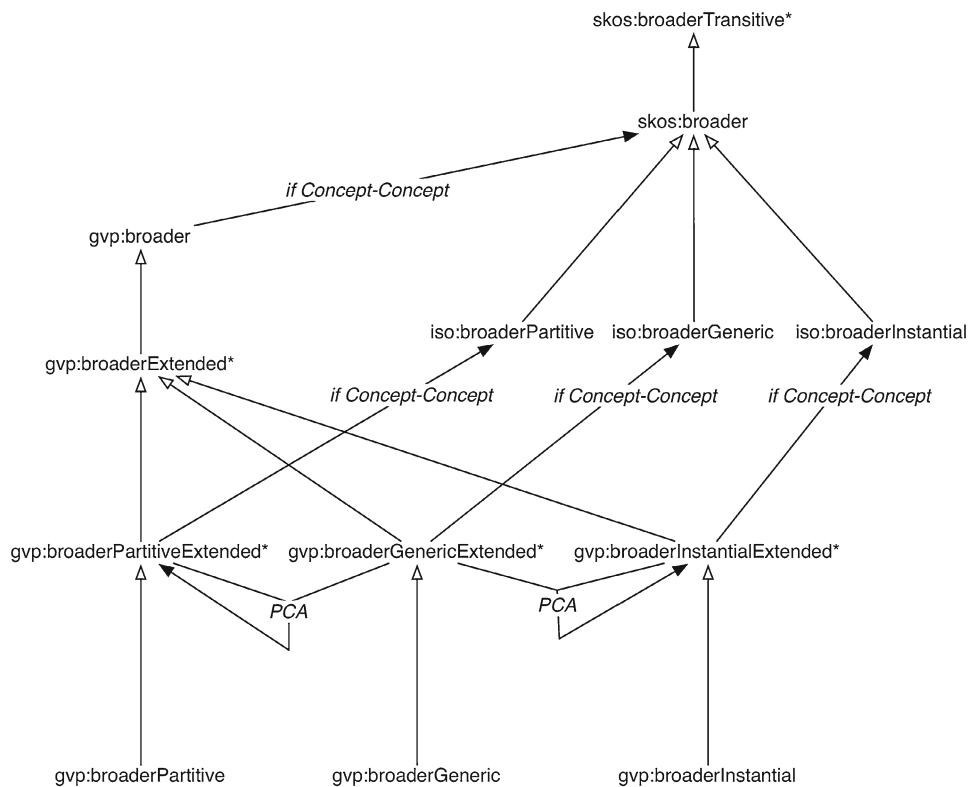
concept” plain arrows indicate inference of the target property when both subject and object of the considered statements are (ISO/SKOS) concepts, and “PCA” plain arrows represent deduction of a property by a property chain axiom.

4 Conclusions

Though qualified hierarchical relationships are suggested in the very first ISO thesaurus standard for thesauri (1974), this standard and its successors (e.g., ANSI/NISO Z39.19, ISO 25964) did not explicitly elaborate on properties of semantic relationships, such as transitivity, symmetry or reflexivity. However, most vocabularies did not make use of the possibility to explicitly distinguish between different types of hierarchies anyway. Instead, the thesaurus BT/NT relationship was and still is generally employed to cover all of these semantically different situations. Hierarchies in existing thesauri often are rather associations than sound hierarchical statements.

This is a severe impediment for fine-grained knowledge representation and query expansion. Since the hierarchies of existing thesauri often are an indiscriminate mixture of semantically highly different relations, automated inference of indiscernible high-level statements (such as with the skos:broaderTransitive property) leads to unsatisfactory retrieval results with respect to precision.

Fig. 6 Formal inference semantics unifying GVP, ISO 25964 and SKOS hierarchical properties



In the perspective of the linked data scenario and the creation of mappings across vocabularies, the logical grounding of relationships becomes increasingly important to ensure true interoperability and better “semantic” services. Further elaboration on logical conditions and restrictions of semantic relationships in KOS, especially concerning the part-whole relationship, is needed.

This paper presents solutions to tackle these issues. Our exploration of compositionality of BTG, BTP, BTI is the first step towards an agreement on formal semantics for these relationships, which can be re-used across thesauri to share finer-grained data. We hope that these will be adopted by the ongoing standardization work around the ISO 25964 ontology extension for SKOS [5], and such work has already started.

We also believe that our compositional semantics, as implemented in thesauri like AAT, opens new avenues for building innovative, helpful services for retrieval of objects described using these KOSs. The faceted browsing retrieval paradigm is an especially promising area. In spite of its popularity, many such browsing interfaces still rarely exploit the various semantic relationships within concept hierarchies at their full potential. Using the appropriate (inference) steps to traverse hierarchical graphs from mixed hierarchies, a system can present more relevant results for users while keeping a good precision level.

Acknowledgments This research was conducted as part of the semantic publication of the Getty vocabularies, <http://vocab.getty.edu>. The support of the J. Paul Getty Trust is gratefully acknowledged. This paper extends on an earlier presentation of this work [19].

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

All links were accessed on 2014-04-29 or later.

1. Baker, T., Bechhofer, S., Isaac, A., Miles, A., Schreiber, G., Summers, E.: Key choices in the design of simple knowledge organization system (SKOS). *J. Web Semant.* **20**, 35–49 (2013). doi:10.1016/j.websem.2013.05.001
2. Harpring, P.: Introduction to controlled vocabularies: terminology for art, architecture, and other cultural works. Getty Research Institute (2010). <http://www.getty.edu/publications/virtuallibrary/160606018X.html>
3. Getty vocabularies. <http://www.getty.edu/research/tools/vocabularies/>
4. ISO 25964 thesaurus schemas. <http://www.niso.org/schemas/iso25964>
5. De Smedt, J., Isaac, A., Clarke, S.D., Lindenthal, J., Zeng, M.L., Tudhope, D.S., Will, L., Alexiev, V.: ISO 25964 part 1: thesauri for information retrieval: RDF/OWL vocabulary, extension of SKOS and SKOS-XL (2013). <http://purl.org/iso25964/skos-thes>
6. Deutsche Nationalbibliothek—linked data service. <http://www.dnb.de/EN/lds>; GND ontology. <http://d-nb.info/standards/elementset/gnd>
7. Suominen, O.: FinnONTO SKOS extensions (2011). <http://www.ldf.fi/schema/skosexl/>
8. Vitzthum, A.: Standards und Schnittstellen für webbasierte Vokabularverwaltungssysteme. FG Dokumentation, Berlin (2012). http://www.museumsbund.de/fileadmin/fg_doku/termine/2012_10/Beitraege/Vitzthum-berlin_20121018.pdf
9. Alexiev, V., Joan C., Gregg G., Harpring, P.: Getty vocabularies linked open data: semantic representation, 3.0 edn. Getty Research Institute (2015). <http://vocab.getty.edu/doc/>
10. Lindenthal, J.: Ambiguities in representing thesauri using extended SKOS—examples from real life. NKOS 2012 Workshop, TPDL 2012, Paphos, Cyprus (2012). https://at-web1.comp.glam.ac.uk/pages/research/hypermedia/nkos/nkos2012/presentations/TPDL2012_NKOS_Ambiguities.pptx
11. Miles, A., Bechhofer, S.: SKOS simple knowledge organization system reference. W3C recommendation (2009). <http://www.w3.org/TR/skos-reference/>
12. Varzi, A.: Mereology. In: Zalta, E.N. (eds.) *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition). <http://plato.stanford.edu/archives/fall2014/entries/mereology/>
13. Winston, M., Chaffin, R., Hermann, D.: A taxonomy of part-whole relations. *Cogn. Sci.* **11**, 417–444 (1987)
14. Keet, C.M.: Introduction to part-whole relations: mereology, conceptual modelling and mathematical aspects. Free University of Bozen-Bolzano, KRDB Research Centre Technical Report KRDB06-3 (2006). <http://www.inf.unibz.it/krdp/pub/TR/KRDB06-3.pdf>
15. Allemang, D., Hendler, J.A.: *Semantic Web for the Working Ontologist: Modeling in RDF, RDFS and OWL*. Morgan Kaufmann Publishers/Elsevier, San Francisco (2008)
16. Tudhope, D., Alani, H., Jones, C.: Augmenting thesaurus relationships: possibilities for retrieval. *J. Digit. Inf.* **1**(8) (2001). <http://journals.tdl.org/jodi/article/view/181/160>
17. Motik, B.: On the Properties of Metamodeling in OWL, ISWC 2005. Galway, Ireland (2005)
18. Alexiev, V.: TGN place type relation (discussion document) (2014). <http://lists.w3.org/Archives/Public/public-esw-thes/2014Mar/att-0001/TGN-place-type.docx>
19. Alexiev, V., Lindenthal, J., Isaac, A.: On compositionality of ISO 25964 hierarchical relations (BTG, BTP, BTI). In: Presentation at 13th European Networked Knowledge Organization Systems (NKOS'14) Workshop at DL2014 Conference, London (2014). <http://vladimiralexiev.github.io/pres/20140912-NKOS-compositionality/>