INTERNATIONAL STANDARD

ISO 21127

Third edition

Information and documentation — A reference ontology for the interchange of cultural heritage information

Information et documentation — Une ontologie de référence pour l'échange d'informations du patrimoine culturel

PROOF/ÉPREUVE

Reference number
ISO 21127:2023(E)

Contents

Page

Foreword iv
Introduction v
1 Scope 1
2 Normative references 1
3 Terms and definitions 1
4 Objectives 9
5 Compatibility 10
6 Applied form 10
6.1 General 10
6.2 Naming conventions 11
6.3 Inheritance and transitivity 12
6.4 Shortcuts 12
6.5 Logical expressions used in ISO 21127 12
6.6 Property quantifiers 14
7 Modelling principles 15
7.1 Reality, knowledge bases and ISO 21127 15
7.2 Authorship of knowledge base (KB) contents 17
7.3 Extensions 17
7.4 Minimality 19
7.5 Monotonicity 20
7.5.1 Open World principle 20
7.5.2 Monotonicity of the document 20
7.5.3 Monotonicity of the data 21
7.5.4 Monotonicity of the knowledge base 21
7.5.5 Monotonicity and time-dependent properties 21
7.6 Disjointness 22
8 Introduction to the basic concepts 22
8.1 General 22
8.2 Relations with events 24
8.2.1 General 24
8.2.2 Spatial relations 27
8.2.3 Temporal relations 29
8.2.4 Spatiotemporal relations 31
8.3 Specific modelling constructs 33
8.3.1 Types 33
8.3.2 Temporal relation primitives based on fuzzy boundaries 34
9 Class declarations 37
10 Property declarations 91
Annex A (informative) **Scope precision — Intended scope** 175
Annex B (informative) **Deprecated classes and properties** 176
Bibliography 179

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of ISO document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

ISO draws attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO takes no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents. ISO shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see www.iso.org/iso/foreword.html.

This document was prepared by Technical Committee ISO/TC 46, *Information and documentation*, Subcommittee SC 4, *Technical interoperability,* in collaboration with the International Committee for Documentation (CIDOC).

This third edition cancels and replaces the second edition (ISO 21127:2014), which has been technically revised.

The main changes are as follows:

- deprecated 13 overspecialised classes and 15 overspecialized properties;
- added 8 properties to replace 8 deprecated properties in order to support chronological reasoning;
- added 4 (sub)classes and 17 properties to align with OCG standards for geospatial data;
- added 4 (sub)classes and 12 properties for more detailed conceptualizations of existing concepts;
- provided further clarification of concepts through the addition real-life examples, references, and first order logic axioms;
- extended explanatory introductory sections to clarify the standard and its maintained scope.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html.

Introduction

This document is the culmination of more than a decade of standards development work by the International Committee for Documentation (CIDOC) of the International Council of Museums (ICOM). Work on this document began in 1996 under the auspices of the ICOM-CIDOC Documentation Standards Working Group. The document[1] provided by CIDOC formed the basis for ISO 21127 which was first published in 2006. While the initial impetus for the work came from the museum community, it has since spread to encompass other types of cultural heritage institution. This document has been appropriated and extended to meet the needs of other institutions dealing with cultural heritage.

The primary purpose of this document is to offer a conceptual basis for the integration, mediation, and exchange of information between cultural heritage organizations such as museums, libraries, and archives. This document aims to provide a common reference point against which divergent and incompatible sources of information can be compared and, ultimately, harmonized.

ISO 21127 is an ontology[2] for cultural heritage information: a formal representation of the conceptual scheme, or "world view", underlying the database applications and documentation systems that are used by cultural heritage institutions. It is important to note that this document aims to clarify the logic of what cultural heritage institutions do in fact document; it is not intended as a normative specification of what they should document. The primary role of this document is to enable information exchange and integration between heterogeneous sources of cultural heritage information. It aims to provide the semantic definitions and clarifications needed to transform disparate, localized information sources into a coherent global resource, be it within an institution, an intranet, or on the Internet.

The specific aims of this document are to

— serve as a common language for domain experts and IT developers when formulating requirements,

— serve as a formal language for the identification of common information contents in different data formats; in particular to support the implementation of automatic data transformation algorithms from local to global data structures without loss of meaning. These transformation algorithms are useful for data exchange, data migration from legacy systems, data information integration, and mediation of heterogeneous sources,

— support associative queries against integrated resources by providing a global model of the basic classes and their associations to formulate such queries, and

— provide developers of information systems with a guide to good practice in conceptual modelling.

The ISO 21127 ontology is expressed as a series of interrelated concepts with definitions. This presentation is similar to that used for a thesaurus. However, the ontology is not intended as a terminology standard and does not set out to define the terms that are typically used as data in cultural heritage documentation. Although the presentation provided here is complete, it is an intentionally compact and concise presentation of the ontology's 81 classes and 160 unique properties. It does not attempt to articulate the inheritance of properties by subclasses throughout the class hierarchy. However, this definition does contain all the information needed to infer and automatically generate a full declaration of all properties, including inherited properties.[3]

1) The CIDOC CRM Special Interest Group continues to maintain a version of this original document, usually known as the "CIDOC Conceptual Reference Model" or CIDOC CRM[15].

2) In the sense used in computer science, i.e. it describes in a formal language the relevant explicit and implicit concepts and the relationships between them.

3) A class and property reference hierarchy can be found in Reference [15].

Information and documentation — A reference ontology for the interchange of cultural heritage information

1 Scope

This document gives a curated, factual knowledge about the past at a human scale. It specifies all information required for the exchange and integration of heterogeneous scientific and scholarly documentation about the past at a human scale and the available documented and empirical evidence for this.[4)]

A more detailed and useful definition can be articulated by defining both the intended scope, a broad and maximally-inclusive definition of general application principles, and the practical scope, which is expressed by the overall scope of a growing reference set of specific, identifiable documentation standards and practices that this document aims to semantically describe, restricted, always, in its details to the limitations of the intended scope.

The practical scope[5)] of this document is expressed in terms of the set of reference standards and de facto standards for documenting factual knowledge. This document covers the same domain of discourse as the union of these reference standards; this means that for data correctly encoded according to these documentation formats there can be an ISO 21127-compatible expression that conveys the same meaning.

2 Normative references

There are no normative references in this document.

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

ISO and IEC maintain terminology databases for use in standardization at the following addresses:

— ISO Online browsing platform: available at https://www.iso.org/obp

— IEC Electropedia: available at https://www.electropedia.org/

4) Annex A elaborates on this definition.

5) The practical scope of ISO 21127, including a list of the relevant museum documentation standards, is discussed in more detail on the CIDOC CRM website at < https://cidoc-crm.org/scope.html >.

3.1
class
category of items that share one or more common traits serving as criteria to identify the items belonging to the class

Note 1 to entry: These properties need not be explicitly formulated in logical terms, but may be described in a text (here called a scope note) that refers to a common conceptualisation of domain experts. The sum of these traits is called the intension of the class. A class may be the domain or range of none, one or more properties formally defined in a model. The formally defined properties need not be part of the intension of their domains or ranges: such properties are optional. An item that belongs to a class is called an instance of this class. A class is associated with an open set of real-life instances, known as the extension of the class. Here “open” is used in the sense that it is generally beyond our capabilities to know all instances of a class in the world and indeed that the future may bring new instances about at any time (Open World). Therefore, a class cannot be defined by enumerating its instances. A class plays a role analogous to a grammatical noun, and can be completely defined without reference to any other construct (unlike properties, which shall have an unambiguously defined domain and range). In some contexts, the terms individual class, entity or node are used synonymously with class. For example, “Person” is a class. To be a “Person” may actually be determined by DNA characteristics, but everyone knows what a “Person” is. A “Person” may have the property of being a member of a “Group”, but it is not necessary to be member of a “Group” in order to be a “Person”. It is impossible to know all the "Persons" of the past. There will be more “Person” in the future.

3.2
complement
<of class A with respect to one of its superclasses B> set of all instances of B that are not instances of A

Note 1 to entry: Formally, it is the set-theoretic difference of the extension of B minus the extension of A. Compatible extensions of the CIDOC CRM should not declare any class with the intension of them being the complement of one or more other classes. To do so will normally violate the desire to describe an Open World. For example, for all possible cases of human gender, male should not be declared as the complement of female or vice versa. What if someone is both or even of another kind?

3.3
disjoint
having no common *instances* (3.8) in any possible world

Note 1 to entry: Classes are disjoint if the intersection of their extensions is an empty set.

3.4
domain
class (3.1) for which a property is formally defined

Note 1 to entry: This means that instances of the property are applicable to instances of its domain class. A property shall have exactly one domain, although the domain class may always contain instances for which the property is not instantiated. The domain class is analogous to the grammatical subject of the phrase for which the property is analogous to the verb. It is arbitrary which class is selected as the domain and which as the range, just as the choice between active and passive voice in grammar is arbitrary. Property names in this document are designed to be semantically meaningful and grammatically correct when read from domain to range. In addition, the inverse property name, normally given in parentheses, is also designed to be semantically meaningful and grammatically correct when read from range to domain.

3.5
endurant
entities which are wholly present at any time they are present

Note 1 to entry: See Reference [87], pp. 166-181.

3.6
extension
set of all real-life instances belonging to the class that fulfil the criteria of its *intension* (3.9)

Note 1 to entry: The extension of a class is “open” in the sense that it is generally beyond our capabilities to know all instances of a class in the world and indeed that the future may bring new instances about at any time (Open World). An information system may at any point in time refer to some instances of a class, which form a subset of its extension.

3.7
inheritance
duplication of properties from a class to its subclasses

Note 1 to entry: Inheritance of properties from superclasses to subclasses means that if an item *x* is an instance of a class A, then all properties that shall hold for the instances of any of the superclasses of A shall also hold for item x, and that all optional properties that may hold for the instances of any of the superclasses of A may also hold for item x.

3.8
instance
items having properties that meet the criteria of the *intension* (3.9) of the classes

Note 1 to entry: The number of instances declared for a class in an information system is typically less than the total in the real world. For example, the reader is an instance of Person, but they are not mentioned in all information systems describing Persons.

Note 2 to entry: For example, the painting known as the “The Mona Lisa” is an instance of the class E22 Human-Made Object. An instance of a property is a factual relation between an instance of the domain and an instance of the range of the property that matches the criteria of the intension of the property. For example, “The Mona Lisa” *has former or current owner.* The Louvre is an instance of the property *P51 has former or current owner (is former or current owner of).*

3.9
intension
intended meaning of a class

Note 1 to entry: The intension of a class consists of one or more common traits shared by all instances of the class or property. These traits need not be explicitly formulated in logical terms, but may just be described in a text (here called a scope note) that refers to a conceptualisation common to domain experts.

3.10
interoperability
capability of different information systems to communicate some of their contents

Note 1 to entry: Interoperability may mean that two systems can exchange information, and/or that multiple systems can be accessed with a single method.

Note 2 to entry: Generally, syntactic interoperability is distinguished from semantic interoperability. Syntactic interoperability means that the information encoding of the involved systems and the access protocols are compatible, so that information can be processed as described above without error. However, this does not mean that each system processes the data in a manner consistent with the intended meaning. For example, one system may use a table called “Actor” and another one called “Agent”. With syntactic interoperability, data from both tables may only be retrieved as distinct, even though they may have exactly the same meaning. To overcome this situation, semantic interoperability has to be added. This document relies on existing syntactic interoperability and is concerned only with adding *semantic interoperability* (3.24).

3.11
inverse of
reinterpretation of a property from range to domain without more general or more specific meaning

Note 1 to entry: The inverse of a property is similar to the choice between active and passive voice in some languages. In contrast to some knowledge representation languages, such as RDF and OWL, this document regards that the inverse of a property is not a property in its own right that needs an explicit declaration of being inverse of another, but an interpretation implicitly existing for any property. The inverse of the inverse of a property is identical to the property itself, i.e. its primary sense of direction. For example, “Entity *is depicted by* Physical Human-Made Thing” is the inverse of “Physical Human-Made Thing *depicts* Entity”.

3.12
knowledge creation event
organized transfer of knowledge of a group of domain experts into an information system under preservation of the relationship between those regarding the created content as representing their knowledge and the knowledge itself

Note 1 to entry: All knowledge contained in an information system must have been introduced into that system by some human agent, either directly or indirectly. Despite this fact, many, if not most, statements within such a system will lack specific attribution of authority. That being said, in the domain of cultural heritage, it is common practice that, for the processes of collection documentation and management, there are clearly and explicitly elaborated systems of responsibility outlining by whom and how knowledge can be added and or modified in the system. Ideally these systems are specified in institutional policy and protocol documents. Thus, it is reasonable to hold that all such statements that lack explicit authority attribution within the information system can, in fact, be read as the official view of the administrating institution of that system. Such a position does not mean to imply that an information system represents at any particular moment a completed phase of knowledge that the institution promotes. Rather, it means to underline that, in a CH context, a managed set of data, at any state of elaboration, will in fact embody an adherence to some explicit code of standards which guarantees the validity of that data within the scope of said standards and all practical limitations. So long as the information is under active management it remains continuously open to revision and improvement as further research reveals further understanding surrounding the objects of concern. A distinct exception to this rule is represented by information in the data set that carries with it an explicit statement of responsibility.

Note 2 to entry: In this document, such statements of responsibility are expressed through knowledge creation events such as E13 Attribute Assignment and its relevant subclasses. Any information in a model using this document that is based on an explicit creation event for that piece of information, where the creator’s identity has been given, is attributed to the authority and assigned to the responsibility of the actor identified as causal in that event. For any information in the system connected to knowledge creation events that do not explicitly reference their creator, as well as any information not connected to creation events, the responsibility falls back to the institution responsible for the database/knowledge graph. That means that for information only expressed through shortcuts such as *P2 has type*, where no knowledge creation event has been explicitly specified, the originating creation event cannot be deduced and the responsibility for the information can never be any other body than the institution responsible for the whole information system. In the case of an institution taking over stewardship of a database transferred into their custody, two relations of responsibility for the knowledge therein can be envisioned. If the institution accepts the dataset and undertakes to maintain and update it, then they take on responsibility for that information and become the default authority behind its statements as described above. If, on the other hand, the institution accepts the data set and stores it without change as a closed resource, then it can be considered that the default authority remains the original steward.

3.13
monotonic
<of a knowledge base> having a set of conclusions derived through inference rules that does not reduce, irrespective of whatever additional propositions can be inserted

Note 1 to entry: Monotonic reasoning is a term from knowledge representation. A reasoning form is monotonic if an addition to the set of propositions making up the knowledge base never determines a decrement in the set of conclusions that may be derived from the knowledge base via inference rules. In practical terms, if experts enter subsequently correct statements to an information system, the system should not regard any results from those statements as invalid, when a new one is entered. The ISO 21127 ontology is designed for monotonic reasoning and so enables conflict-free merging of huge stores of knowledge.

3.14
multiple inheritance
possibility for a class to have more than one immediate superclass

Note 1 to entry: The extension of a class with multiple immediate superclasses is a subset of the intersection of all extensions of its superclasses. The intension of a class with multiple immediate superclasses extends the intensions of all its superclasses, i.e. its traits are more restrictive than any of its superclasses. If multiple inheritance is used, the resulting "class hierarchy" is a directed graph and not a tree structure. If it is represented as an indented list, there are necessarily repetitions of the same class at different positions in the list. For example, Person is both an Actor and a Biological Object.

3.15
multiple instantiation
case that an instance of class A is also regarded as an instance of one or more other classes B1...n

Note 1 to entry: When multiple instantiation is used, it has the effect that the properties of all these classes become available to describe this instance. For instance, some particular cases of destruction may also be activities (e.g. Herostratos' deed), but not all destructions are activities (e.g. destruction of Herculaneum). In comparison, multiple inheritance describes the case that all instances of a class A are implicitly instances of all superclasses of A, by virtue of the definition of the class A, whereas the combination of classes used for multiple instantiation is a characteristic of particular instances only. It is important to note that multiple instantiation is not allowed using combinations of disjoint classes.

3.16
open world
assumption that the information stored in a knowledge base is incomplete with respect to the universe of discourse it aims to describe

Note 1 to entry: The "open world assumption" is a term from knowledge base systems. This incompleteness may be due to the inability of the maintainer to provide sufficient information or due to more fundamental problems of cognition in the system's domain. Such problems are characteristic of cultural information systems. Our records about the past are necessarily incomplete. In addition, there may be items that cannot be clearly assigned to a given class. In particular, absence of a certain property for an item described in the system does not mean that this item does not have this property. For example, if one item is described as Biological Object and another as Physical Object, this does not imply that the latter may not be a Biological Object as well. For example, one cannot list "all Physical Objects known to the system that are not Biological Objects in the real world", but one may of course list "all items known to the system as Physical Objects but that are not known to the system as Biological Objects". Therefore, complements of a class with respect to a superclass cannot be concluded in general from an information system using the open world assumption.

3.17
perdurant
entities which extend in time

Note 1 to entry: "The difference between enduring and perduring entities (referred to in this document as *endurants* (3.5) and *perdurants* (3.17)) is related to their behaviour in time. Endurants are wholly present (i.e. all their proper parts are present) at any time they are present. Perdurants, on the other hand, just extend in time by accumulating different temporal parts, so that, at any time they are present, they are only partially present, in the sense that some of their proper temporal parts (e.g. their previous or future phases) may be not present. For example, the piece of paper the reader is reading now is wholly present, while some temporal parts of their reading are not present any more. Philosophers say that endurants are entities that are in time, while lacking however temporal parts (so to speak, all their parts flow with them in time). Perdurants, on the other hand, are entities that happen in time, and can have temporal parts (all their parts are fixed in time)." (Reference [87], pp. 166-181).

3.18
property
named characteristic of a class to which values can be assigned

Note 1 to entry: A property serves to define a relationship of a specific kind between two classes. The property is characterized by an intension, which is conveyed by a scope note. A property plays a role analogous to a grammatical verb, in that it shall be defined with reference to both its domain and range, which are analogous to the subject and object in grammar (unlike classes, which can be defined independently). It is arbitrary, which class is selected as the domain, just as the choice between active and passive voice in grammar is arbitrary. In other words, a property can be interpreted in both directions, with two distinct, but related interpretations. Properties may themselves have properties that relate to other classes (This feature is used in this model only in order to describe dynamic subtyping of properties). Properties can also be specialized in the same manner as classes, resulting in IsA relationships between subproperties and their superproperties. In some contexts, the terms attribute, reference, link, role or slot are used synonymously with property. For example, “Physical Human-Made Thing *depicts* Entity” is equivalent to “Entity *is depicted by* Physical Human-Made Thing”.

3.19
property quantifier
declaration of the allowed number of *instances* (3.8) of a certain *property* (3.18) that can refer to a particular instance of the range class or the domain class of that property

Note 1 to entry: Property quantifier declarations are ontological, i.e. they refer to the nature of the real world described and not to our current knowledge. For example, each person has exactly one father, but collected knowledge may refer to none, one or many.

3.20
query
request for information from an information system expressed so that the response can be calculated automatically

3.21
range
class that comprises all potential values of a property

Note 1 to entry: That means that instances of the property can link only to instances of its range class. A property shall have exactly one range, although the range class may always contain instances that are not the value of the property. The range class is analogous to the grammatical object of a phrase for which the property is analogous to the verb. It is arbitrary which class is selected as domain and which as range, just as the choice between active and passive voice in grammar is arbitrary. Property names in ISO 21127 are designed to be semantically meaningful and grammatically correct when read from domain to range. In addition, the inverse property name, normally given in parentheses, is also designed to be semantically meaningful and grammatically correct when read from range to domain.

3.22
reflexivity
binary relation on a set X that relates every element of X to itself

Note 1 to entry: Reflexivity is defined in the standard way found in mathematics or logic: A property P is reflexive if the domain and range are the same class and for all instances x, of this class the following is the case: x is related by P to itself. The intention of a property as described in the scope note will decide whether a property is reflexive or not. An example of a reflexive property is E53 Place*. P89 falls within (contains):* E53 Place.

3.23
scope note
textual description of the *intension* (3.9) of a class or property

Note 1 to entry: Scope notes are not formal modelling constructs, but are provided to help explain the intended meaning and application of ISO 21127’s classes and properties. Basically, they refer to a conceptualisation common to domain experts and disambiguate between different possible interpretations. Illustrative example instances of classes and properties are also regularly provided in the scope notes for explanatory purposes.

3.24
semantic interoperability
capability of different information systems to communicate information consistent with the intended meaning

Note 1 to entry: In more detail, the intended meaning encompasses the data structure elements involved, the terminology appearing as data, and the identifiers used in the data for factual items such as places, people, objects, etc. Obviously, communication about data structure must be resolved first. In this case, consistent communication means that data can be transferred between data structure elements with the same intended meaning or that data from elements with the same intended meaning can be merged. In practice, the different levels of generalization in different systems do not allow the achievement of this ideal. Therefore, semantic interoperability is regarded as achieved if elements can be found that provide a reasonably close generalization for the transfer or merge. This document is only concerned with semantic interoperability on the level of data structure elements.

3.25
shortcut
formally defined single property that represents a deduction or join of a data path in the ontology

Note 1 to entry: The scope notes of all properties characterized as shortcuts describe in words the equivalent deduction. Shortcuts are introduced for the cases where common documentation practice refers only to the deduction rather than to the fully developed path. For example, museums often only record the dimension of an object without documenting the Measurement that observed it. This document declares shortcuts explicitly as single properties in order to allow the user to describe cases in which he has less detailed knowledge than the full data path would need to be described. For each shortcut, this document contains in its schema the properties of the full data path explaining the shortcut.

3.26
strict inheritance
property *inheritance* (3.7) that allows no exceptions

Note 1 to entry: Some systems may declare that elephants are grey and regard a white elephant as an exception. Under strict inheritance it would hold that: if all elephants were grey, then a white elephant could not be an elephant. Obviously, not all elephants are grey; to be grey is not part of the intension of the concept elephant but an optional property. This document applies strict inheritance as a normalization principle.

3.27
subclass
specialization of another *class* (3.1), i.e. the superclass

Note 1 to entry: Specialization or the IsA relationship means that all instances of the subclass are also instances of its superclass, that the intension of the subclass extends the intension of its superclass, i.e. its traits are more restrictive than that of its superclass, and that the subclass inherits the definition of all of the properties declared for its superclass without exceptions (strict inheritance), in addition to having none, one or more properties of its own. A subclass can have more than one immediate superclass and consequently inherits the properties of all of its superclasses (multiple inheritance). The IsA relationship or specialization between two or more classes gives rise to a structure known as a class hierarchy. The IsA relationship is transitive and may not be cyclic. In some contexts (e.g. the programming language C++) the term derived class is used synonymously with subclass. For example, every Person IsA Biological Object, or Person is a subclass of Biological Object. Also, every Person IsA Actor. A Person may die. However, other kinds of Actors, such as companies, don’t die (c.f. 2). Every Biological Object IsA Physical Object. A Physical Object can be moved. Hence, a Person can be moved also (c.f. 3).

3.28
subproperty
specialization of another *property* (3.18), i.e. the superproperty

Note 1 to entry: Specialization or IsA relationship means that all instances of the subproperty are also instances of its superproperty, that the intension of the subproperty extends the intension of the superproperty, i.e. its traits are more restrictive than that of its superproperty, that the domain of the subproperty is the same as the domain of its superproperty or a subclass of that domain, that the range of the subproperty is the same as the range of its superproperty or a subclass of that range, and that the subproperty inherits the definition of all of the properties declared for its superproperty without exceptions (strict inheritance), in addition to having none, one or more properties of its own. A subproperty can have more than one immediate superproperty and consequently inherits the properties of all of its superproperties (multiple inheritance). The IsA relationship or specialization between two or more properties gives rise to the structure called a property hierarchy. The IsA relationship is transitive and may not be cyclic. Some object-oriented programming languages, such as C++, do not contain constructs that allow for the expression of the specialization of properties as sub-properties.

Note 2 to entry: Alternatively, a property may be subproperty of the inverse of another property, i.e. reading the property from range to domain. In that case all instances of the subproperty are also instances of the inverse of the other property, the intension of the subproperty extends the intension of the inverse of the other property, i.e. its traits are more restrictive than that of the inverse of the other property, the domain of the subproperty is the same as the range of the other property or a subclass of that range, and the range of the subproperty is the same as the domain of the other property or a subclass of that domain. The subproperty inherits the definition of all of the properties declared for the other property without exceptions (strict inheritance), in addition to having none, one or more properties of its own. The definitions of inherited properties have to be interpreted in the inverse sense of direction of the subproperty, i.e. from range to domain.

3.29
superclass
generalization of one or more other classes, i.e. the subclasses

Note 1 to entry: A superclass subsumes all instances of its subclasses, and that it can also have additional instances that do not belong to any of its subclasses. The intension of the superclass is less restrictive than any of its subclasses. This subsumption relationship or generalization is the inverse of the IsA relationship or specialization. For example, “Biological Object subsumes Person” is synonymous with “Biological Object is a superclass of Person”. It needs fewer traits to identify an item as a Biological Object than to identify it as a Person.

3.30
superproperty
generalization of one or more other properties, i.e. the subproperties

Note 1 to entry: A superproperty subsumes all instances of its subproperties, and that it can also have additional instances that do not belong to any of its subproperties. The intension of the superproperty is less restrictive than any of its subproperties. The subsumption relationship or generalization is the inverse of the IsA relationship or specialization. A superproperty may be a generalization of the inverse of another property.

3.31
symmetric property
binary relation R on a set X that, for all elements a and b in X, whenever R relates a to b then R also relates b to a

Note 1 to entry: Symmetry is defined in the standard way found in mathematics or logic: A property P is symmetric if the domain and range are the same class and for all instances x, y of this class the following is the case: If x is related by P to y, then y is related by P to x. The intention of a property as described in the scope note will decide whether a property is symmetric or not. An example of a symmetric property is “E53 Place. P122 borders with: E53 Place”. The names of symmetric properties have no parenthetical form, because reading in the range-to-domain direction is the same as the domain-to-range reading.

 PROOF/ÉPREUVE

3.32
transitivity
binary relation R on set X that, for all elements a, b, and c in X, whenever R relates a to b and R relates b to c, then R relates a to c

Note 1 to entry: Transitivity is defined in the standard way found in mathematics or logic: A property P is transitive if the domain and range is the same class and for all instances x, y, z of this class the following is the case: If x is related by P to y and y is related by P to z, then x is related by P to z. The intention of a property as described in the scope note will decide whether a property is transitive or not. For example, the property *P121 overlaps with* between instances of E53 Place is not transitive, while the property *P89 falls within (contains)* between instances of E53 Place and the property *P46 is composed of (forms part of)* between instances of E18 Physical Thing are both transitive. Transitivity is especially useful when the ISO 21127 ontology is implemented in a system with deduction.

3.33
universal
entities that can have *instances* (3.8) in a possible world

Note 1 to entry: The fundamental ontological distinction between universals and particulars can be informally understood by considering their relationship with instantiation: particulars are entities that have no instances in any possible world; universals are entities that do have instances. Classes and properties (corresponding to predicates in a logical language) are usually considered to be universals. (after Reference [87], pp. 166-181).

4 Objectives

The primary role of this document is to enable the exchange and integration of information from heterogeneous sources for the reconstruction and interpretation of the past at a human scale, based on all kinds of material evidence, including texts, audio-visual material and oral tradition. It starts from, but is not limited to, the needs of museum documentation and research based on museum holdings. It aims at providing the semantic definitions and clarifications needed to transform disparate, localised information sources into a coherent global resource, be it within a larger institution, in intranets or on the Internet, and to make it available for scholarly interpretation and scientific evaluation. These goals determine the constructs and level of detail of this document.

More specifically, it defines, in terms of a formal ontology, the underlying semantics of database schemata and structured documents used in the documentation of cultural heritage and scientific activities. In particular, it defines the semantics related to the study of the past and current state of our world, as it is characteristic for museums, but also or other cultural heritage institutions and disciplines. It does not define any of the terminology appearing typically as data in the respective data structures; it foresees, however, the characteristic relationships for its use. It does not aim at proposing what cultural heritage institutions should document. Rather, it explains the logic of what they actually currently document, and thereby enables semantic interoperability.

This document intends, moreover, to provide a model of the intellectual structure of the respective kinds of mentioned documentation in logical terms. As such, it has not been optimised for implementation specific storage and processing factors. Actual system implementations may lead to solutions where elements and links between relevant elements of our conceptualizations are no longer explicit in a database or other structured storage system. For instance, the birth event that connects elements such as father, mother, birth date, birth place may not appear in the database, in order to save storage space or response time of the system. This document provides a conceptual and technical means to explain how such apparently disparate entities are semantically and logically interconnected, and how the ability of the database to answer certain intellectual questions is affected by the omission of such elements and links.

This document aims to support the following specific functionalities:

— inform developers of information systems as a guide to good practice in conceptual modelling, in order to effectively structure and relate information assets of cultural documentation;