Toward a Unified Framework for Realizable Entities: Dispositions, Plans, and Beyond

Stefan Schulz[1,2,*], Michel Dumontier[3], Remzi Çelebi[3], Katerina Serafimova[4] and Catalina Martínez Costa[5]

[1]Institute for Medical Informatics, Statistics and Documentation, Medical University of Graz, Austria
[2]Averbis GmbH, Freiburg, Germany
[3]Institute of Data Science, Maastricht University, The Netherlands
[4]Graphwise, Sofia, Bulgaria
[5]University of Murcia, Spain

Abstract

This paper extends the concept of a *realizable entity* (RE). Traditionally, REs include dispositions, functions, and roles. We propose an enriched typology that also encompasses plans and information content entities (ICEs). We term these *structured* and *epistemic* REs, respectively, based on shared characteristics: they depend on a bearer and can exist regardless of whether the entity they refer to is actually instantiated. To capture this commonality, we introduce a unified relation, '**refers to**', which generalizes the notions of realization, concretization, and aboutness. The '**refers to**' relation enables reference to both individuals and repeatables (i.e., universals, types, concepts), including hypothetical, future, or even impossible entities. To address cases of *generic reference*—where the referent is a repeatable and not a particular instance—we introduce *collections* as extensions of repeatables. After presenting the motivation for this unified relation, we incorporate it into a series of OWL-based design patterns. Some of these patterns require the expressive power of the OWL-DL profile, while others are compatible with the more scalable OWL-EL. These patterns support, or approximate, the representation and reasoning over REs, with particular attention to references to repeatables, represented as collections. We mention several open issues, including the semantics of collections, the disambiguation of reference types, and complexities introduced by multi-referential information content entities. Our approach aims to support practical ontology engineering across diverse domains—especially biomedicine and clinical informatics—where modeling uncertainty, risk, hypothetical states, and plans is essential. Our proposal offers a lightweight, user-friendly framework for computable reasoning while remaining largely agnostic with respect to foundational metaphysical assumptions. Further validation is needed to assess its effectiveness and to address its current limitations.

Keywords

Formal ontology, Description logics, Realizable entity, Information entity

1. Introduction

Much scientific, technical, and everyday discourse relies not only on reference to actual entities, but also to potential or hypothetical ones. We speak of potential extraterrestrial life, disease risks, possible side effects of drugs, an expected thunderstorm, and future travel plans. This kind of discourse would not be possible without the reference to realizable entities (REs) and another key category, *viz.* information content entities (ICEs). Realizable entities (REs) are defined as entities that are ontologically dependent on a bearer, and whose essential nature is determined by their realization in processes of specific, correlated types in which the bearer participates [1]. Typical examples include *capabilities* (e.g., the ability to metabolize glucose), *dispositions* (e.g., fragility) [2], and *functions* (e.g., to drive in a nail or to pump blood). The status of *roles* (e.g., being a teacher) as realizable entities (REs) remains debated [1], as does the question of whether functions are siblings or subtypes of dispositions [3]. In the classical view, REs are manifested—i.e., *realized*—through processual entities under certain trigger conditions.

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XI: The Sicilian Summer under the Etna, co-located with the 15th International Conference on Formal Ontology in Information Systems (FOIS 2025), September 8–9, 2025, Catania, Italy
*Corresponding author: S. Schulz, Med Uni Graz, Neue Stiftingtalstraße 6, 8010 Graz, Austria, stefan.schulz@medunigraz.at
000-0001-7222-3287 (S. Schulz); 0000-0003-4727-9435 (M. Dumontier); 0000-0001-7769-4272 (R. Çelebi); 0009-0004-0797-7313 (K. Serafimova); 0000-0003-1857-1744 (C. M. Costa)

However, REs still exist, even if these trigger conditions never become manifest. Yet, the boundaries of what exactly falls in the category ‘realizable entity’ (RE) are under-explored. So is ‘realization’ sometimes associated with plans or plan specifications [4], i.e., kinds of entities referred to as semiotic, epistemic, representational, or informational. We will refer to them as information content entities (ICEs), according to the Information Artifact Ontology (IAO) [5].

This paper addresses the challenges encountered in an upper-level ontology engineering project. SULO (Simplified Upper-Level Ontology) is an emerging ontology that emphasizes maximal simplicity and robustness across divergent metaphysical foundations. It has been proposed as a parsimonious and user-friendly upper level for data integration, primarily in healthcare and biomedical research settings [6]. However, its specification has revealed differing views on the ontological status of realizable entities (REs) and their distinction from ICEs (information content entities). We pursue the following objectives:

- To expand the boundaries of what counts as a RE by suggesting a broadened typology including intentional and representational entities, justified by similar representational patterns and committed to the goal of developing a simplified foundational ontology;
- To propose a unified relation for relating REs in this extended sense with target entities of the most general type, encompassing universals (types, concepts, repeatables) and particulars (individuals);
- To find solutions to the phenomenon that classes of domain entities need to be defined concerning properties or states of affairs that are associated with them but do not necessarily exist while the definiendum exists;
- To develop and analyze representational patterns and their reasoning power in two dialects of description logics, viz. OWL-DL and OWL-EL [7], motivated by the need for scalable representational formalisms in an upper-level ontology;

2. Background

2.1. The Classical Framework of Realizables

Realizable entities (REs) have three important characteristics in common, viz. (i) they always inhere in some material or informational bearer, (ii) they may exist without being realized, and (iii) their realization typically occurs through a process. This classical perspective emphasizes a connection between the world’s static structure, i.e. continuants or endurants, with the dynamical structure of types of possible and actual causal processes [8]. Material entities have dispositional properties, such as fragility, many kinds of technical artefacts have built-in functions. Genetic predispositions are typically latent and probabilistic as they indicate an elevated likelihood of disease manifestation. The assessment of bodily and molecular functions forms a cornerstone of medical diagnostics. Risk factors, often expressed in statistical terms (e.g., a 40% increased risk of cardiovascular disease), represent conditional probabilities rather than actual disease states, but are relevant to underpin predictive and preventive healthcare models. Dispositions such as allergic sensitivities remain silent until triggered by exposure to specific allergens.

The salience of dispositions and functions in life science and health care is demonstrated in the large clinical ontology SNOMED CT [9], where capabilities of biomolecules are explicitly named ‘disposition’ (404 concepts), where the concept ‘Propensity to adverse reaction’ has 1413 descendants (among which are all allergic dispositions), ‘Function’ has 2710 descendants, and where 295 concepts are descendants of ‘Finding of increased risk level’. It is also noteworthy that SNOMED allergy concepts—as a typical example of dispositions of clinical relevance—are defined by existential clauses[1]:

[1]For the meaning of the object property ‘**role group**’ cf. [10].

'*Allergy to nickel*' equivalentTo
  '*Propensity to adverse reaction*' and '**role group**' some
    ('**has realization**' some '*Allergic process*' and '**causative agent**' some *Nickel*)

On face value, this pattern implies that for each instance of nickel allergy as a propensity (an RE), some instance of a causative agent (nickel) and some instance of allergy manifestation, i.e., an allergic process must exist. This assumption contradicts the ontological premise that REs can exist without being realized, exemplified by the fact that allergic dispositions can be diagnosed without any allergic process being active. You may not be able to point to any particular quantity of nickel atoms in case a (silent) allergy to nickel is asserted. Thus, unrealized REs are phenomena whose existence is determined with respect to an essential absence; according to T. Deacon [11], they can be characterized as bearing "absential features". We can describe this scenario as follows: The particular (individual) entity **b** (an instance of *B*) is the bearer of the particular realizable entity **r** (an instance of R_P). All instances of R_P are characterised by being realizable through processes of the type *P*. Nevertheless, instances of R_P may also exist in the absence of any instance of *P* connected to them. Reference not only to *particular entities* (tokens, individuals, instances) but also to *repeatable entities* (types, concepts, universals) is therefore a characteristic feature of RE.

2.2. Plans as Structured Realizable Entities

Plans are similar to classical REs in that they inhere in some object, e.g., a piece of paper or a computer drive, can exist without being executed/realized, and their execution manifests itself by processes of the type *Plan execution*. In addition, the result of plan execution/realization may be a tangible object (e.g., a bridge) or an intangible outcome, (e.g., such as a computer program or a scientific paper.

Existing ontologies diverge in considering plans as REs or ICEs, or they introduce a distinction between plan specifications and plans proper [12]. In this paper, we consider plans as structured REs, as they exhibit an internal temporal or logical structure. This characterization makes it possible to abstract away from a distinction between plans and their specifications, while still permitting such a distinction when needed. The analogy to classical REs is obvious:

The particular (individual) entity **b** (an instance of *B*) is the bearer of the plan **p** (an instance of P_E). All instances of P_E are characterised by being realizable by plan execution processes of the type *E*; they also exist in the absence of any instance of *E*. Optionally, we can also consider the outcome of *E*, e.g. a type of material or informational entity, a realization of **p**.

2.3. Epistemic Realizable Entities

Plan specifications are already ICEs, characterised as 'realizable directives' by IAO, an ontology under BFO[13], which defines ICEs as entities that stand in a relation of **aboutness** to some other entity, i.e., they bear intentional or referential content. For example, a radiological image is **about** a morphological structure in an organism, and a geographic map is **about** features of a region on Earth's surface. Although not being REs in the classical sense, ICEs also depend on bearers and, like classical REs and plans, can exist independently of the real-world existence of their targets. A painting of an elephant, for instance, need not refer to any particular elephant. Moreover, there are cases where the existence of the referent is uncertain (e.g., preons, i.e., hypothesized quark sub-components), was only confirmed much later than described (e.g., Higgs bosons, discovered in 2012), or is known to be fictional (e.g., a unicorn). Conversely, some ICEs—such as photographs or voice recordings—maintain a persistent aboutness to a specific individual or event throughout their existence, even if the referent ceases to exist. In contrast to classical REs, which always target processes, there are no such restrictions for ICEs, as our examples show. In our extended view on realization, we consider ICEs as *epistemic realizables*. Again, we can apply the same pattern for ICEs as with classical REs and plans:

The particular (individual) entity **b** (an instance of *B*) is the bearer of the ICE **i** (an instance of I_R).

All instances of I_R, are about referents of type *R*. They can exist in the absence of any instance of *R*, and they can even point to target referents that are mere concepts or ideas, including those that violate physical or metaphysical constraints such as perpetual motion machines or round triangles.

3. Methods

By broadening the scope of REs to include (i) classical, (ii) structured, and (iii) epistemic REs, we bring together phenomena that share the following characteristics:

- The source entity inheres in some non-processual, mostly material entity, (e.g., a gene, an organ, an artefact, a piece of paper, a brain, or a computer drive). We also allow ICEs as bearers, e.g., algorithms as bearers of certain computational functions;
- Each source realizable class can be fully defined by reference to its corresponding target class: '*Increased risk of Stroke*' can be stated as equivalent to some combination of *Increased risk* and *Stroke*; *Painting of Unicorn*, to some combination of *Horse*, *Horn* and *Painting*, and *Manned Mars mission* by combining *Mars*, *Space mission*, and *Human*, (cf. [14]);
- The target entity is not required to exist as an instance or even as a non-empty class for the source entity to exist or to be defined. Thus, repeatables are allowed as target entities.

This leads us to the following methodology: We first clarify the ontological nature of what we informally refer to as 'target entity', inspired by concrete examples, covering the three varieties of the phenomenon of (extended) realization as introduced above. To this end, we will collect more examples of the three types of REs. We will then choose an appropriate name for the relation that points to this target entity, inspired by relations found in the literature.

Based on this analysis, we then propose several ontology design patterns[15], using the description logic OWL[2], which has established itself as a standard for computable ontologies. We then discuss the different patterns from both foundational and logical points of view. We refrain from the use of punning, i.e., using the same IRI to refer to both a class and an individual, as permitted by OWL Full [7], because it is ontologically shallow and would lead to second-order logic or require workarounds that are unsupported in widely used modeling frameworks [17]. In particular, we will use the OWL dialects OWL-DL and OWL-EL. Only OWL-EL has polynomial time complexity, which makes it scalable, making it the preferred language for large-scale domain ontologies such as SNOMED CT.

4. Results

4.1. The overarching relation 'refers to'

Table 1 provides examples from different domains. They show a common pattern in the sense that a dependent source entity points to some target entity whose existence is not required for the source entity to exist. They also show contentious cases, including not only plans and specifications, but also risks. Notably, risks could be seen as dispositional by some and epistemic by others, in the sense that a specialist assigns a certain probability to participate in a future instance of the 'at-risk' situation. As for a unifying relation applicable to all these cases, we find, in related literature, several relations that link REs to their target entities (see Table 2). As we expect, there is a distinction between REs proper (dispositions) and information content entities (ICEs), including plans. This table also displays the subtle distinction between different flavors between the different top-level ontologies.

[2]We use OWL Manchester Syntax [16] for our examples.

Table 1: Examples of different types of realizable phenomena

Classical Realizable Entities (REs)

- All glasses have the disposition to break when thrown on the floor, but not all glasses will break;
- All smartphones have the capability to connect to the internet, but in some instances, users have never activated that function;
- The function of all instances of the class uterus is to support a developing fetus, but some organisms never get pregnant and never fulfill this function;
- All vehicle airbags have the function to cushion passengers during a collision, but some airbags never deploy because the vehicle is never involved in an accident;
- All Space Shuttle ejection seats have the function of ejecting an astronaut from the spacecraft, although there is no record of any ejection of a person from a spacecraft via an ejection seat;
- All sorting algorithms have the function of arranging data in a specific order. A new algorithm has this function even before executed.

Structured Realizable Entities (REs)

- All urban development plans are characterized by a structured description of zoning and construction projects, but some plans are never put into practice;
- All business plans include structured descriptions of actions, deployments, and investments, but some of them are never executed;
- All manned Mars mission plans include structured descriptions of all necessary processes and actions, including their participants, but no Mars mission plan has been accomplished up until now;
- All plans about human whole-brain transplantation include structured descriptions of processes and actions, including their participants. However, no single human whole-brain transplantation has ever been accomplished;
- All fitness plans are characterized by regular exercise and diet instructions, but some plans are never followed.

Epistemic Realizable Entities (REs)

- The word ‘stroke’ is mentioned in a patient’s health record, but the patient has never had a stroke. The stroke discourse is only hypothetical or the expression of a false diagnosis;
- The MoCA (Montreal Cognitive Assessment) questionnaire screens patients for cognitive impairment. All instances of MoCA are about the *topic of* cognitive impairment, but many of them are about patients without any cognitive impairment;
- A Chinese Medicine health record includes references to Qi, a vital energy that animates the body by flowing through so-called meridians. According to many Western scholars, however, Qi can only be seen as a metaphor and has no real existence;
- There is some literature about human whole brain transplantation, although this is far from being a realistic therapeutic option;
- The word ‘theft’ is used in a testimony, but only as part of a hypothetical scenario discussed by the witness, not an actual crime;
- The word ‘unicorn’ is used as the subject of a work of art, although unicorns do not exist;
- Multiverses are the subject of scientific speculation, and it is yet unknown whether they exist.

Table 2: Relations between dependent entities and their targets in different ontologies

Dependent entity	Relation(s)	Ontologies / Frameworks
Dispositions / Functions	'**realized in**, '**has realization**', '**is manifested in**'	BFO[3], RO[4], OBI[5], UFO[6], GFO[7]
Information Content Entities (ICEs)	'**is about**', '**refers to**', **denotes**, **represents**	IAO[8], DUL/DnS[9], RO
Plans / Plan specifications / Descriptions	**describes**, '**is concretized in**', '**has goal**'	DUL/DnS
Mental States / Intentional Acts	'**is concretized in**', '**has subject**',	DnS, UFO

We propose the relation

'**refers to**'
(inverse '**is referred to by**')

as a general relation applying to realizable, informational, and intentional entities as its domain and imposes no restrictions on its range. Thus, '**refers to**' is an abstraction of more specific relations such as '**has realization**', '**is about**', **represents** and others. What all of these relations have in common is not only that the domain entities are inherent in some object but particularly that their range encompasses not only particulars but also things that are only potentially or hypothetically existent, which would require referring to repeatables (types, universals, concepts).

After motivating the use of a single relation and providing evidence that the three categories of (extended) REs exhibit comparable ontological structures, we now turn to the implementation of these structures using the ontology languages OWL-DL and OWL-EL. We give preference to OWL-EL, due to its polynomial time complexity (PTIME), which makes it scalable even with huge T-boxes such as SNOMED CT. With the ELK reasoner, it classifies the whole SNOMED CT ontology in a few seconds [18]. In contradistinction, OWL-DL is NEXPTIME-complete and therefore of limited scalability, at least when complete reasoners are used. The new relation '**refers to**' will be used throughout the following examples. The examples vary, but the pattern is the same. According to the examples, the '**inheres in**' clause is optional.

4.2. Targets referred to by existential restriction

This pattern is common across all three types of REs, especially when performance considerations require restricting to OWL-EL.

Pattern SOME: R_T equivalentTo R and '**refers to**' some T and '**inheres in**' some B

with R_T being the (extended) RE to be defined, R its parent (specifying an RE subtype), R the target class and B its bearer, where needed.

[3] https://basic-formal-ontology.org/bfo-2020.html
[4] https://obofoundry.org/ontology/ro.html
[5] https://obi-ontology.org/
[6] https://nemo.inf.ufes.br/ufo
[7] https://www.onto-med.de/ontologies/gfo
[8] https://obofoundry.org/ontology/iao.html
[9] http://www.ontologydesignpatterns.org/ont/dul/DUL.owl

'*Increased risk of Stroke*' equivalentTo
'*Increased risk*' and '**refers to**' some *Stroke* and '**inheres in**' some *Brain*

'*Whole brain transplantation plan*' equivalentTo
Plan and '**refers to**' some '*Whole brain transplantation*'

However, the use of existential restrictions in this pattern leads to unintended representations: We derive from it that some instance of *Stroke* is required for each instance of '*Increased risk of Stroke*', as well as an instance of '*Whole brain transplantation*' for each instance of '*Whole brain transplantation plan*'.

4.3. Targets referred to by value restriction

The following pattern has been recommended by various authors [19, 17] with regard to all three types of REs. Other than the SOME pattern, it does not claim the existence of any instance of the target class, as it only restricts the range of potential target classes.

Pattern ONLY$_1$: R_T equivalentTo R and '**refers to**' only T and '**inheres in**' some B

'*Increased risk of Stroke*' equivalentTo
'*Increased risk*' and '**refers to**' only *Stroke* and '**inheres in**' some *Brain*

'*Whole brain transplantation plan*' equivalentTo
Plan and '**refers to**' only '*Whole brain transplantation*'

As both examples show, OWL's open world assumption needs to be taken into account, because these axioms do not prevent other realizations from being possible unless explicitly declared as disjoint. This is particularly the case with epistemic REs: a piece of information may have different kinds of target, e.g., in a medical diagnosis statement: such a statement is not only about a given condition, but also about a diagnoser, a patient, a diagnostic certainty value etc. This is the rationale for an extension of this pattern with an additional negated clause, making up an implication: if the stroke diagnosis is about a clinical condition, then it can only be about a condition of the type stroke.

Pattern ONLY$_2$[10]: R_T equivalentTo R and '**refers to**' only (T or (not Cat_T)) and '**inheres in**' some B

'*Stroke diagnosis*' equivalentTo
'*Medical diagnosis*' and '**refers to**' only (*Stroke* or (not '*Clinical condition*'))

However, the use of value restrictions leads to scalability challenges in large-scale ontologies. As mentioned above, OWL-DL reasoning tasks (e.g., classification, consistency checking, instance checking) are computationally expensive. Reasoning engines such as HermiT, Pellet, or FaCT++ [20] become exponentially slower or may fail to terminate at all. This is the reason why large ontologies such as SNOMED CT as well as some OBO ontologies (Gene Ontology, Foundational Model of Anatomy) use the OWL-EL profile, keeping reasoning in polynomial time, despite lower expressiveness, as disjunctions (or), negations (not), and value restrictions (only) are not supported.

[10]With Cat_T being a high-level category of which T is a subclass

4.4. Universals as targets

Most ontologies are explicitly or implicitly, ontologies of individuals (particulars). Universals or types are only granted a distinct status in GFO [21] with the class *Category* or in UFO [22] with the class *Type*. Instances of these classes are, therefore, technically OWL individuals. In BFO [13] and SIO [23], the top class *entity* leaves room for interpretation, which means that it is not the same as *owl:Thing*. Also in the BFO-dependent Information Artifact Ontology (IAO, cf. [5]), the class '*Information content entity*' (ICE) implies the clause '**is about**' some *owl:Thing*'. This would allow 'non-entity things', whatever this means. Accordingly, the reference to a target entity would be a reference to an instance of the class *Universal*, i.e., $\mathbf{t_{univ}}$ Type *Universal.* This means that within a description logics perspective, universals are represented as OWL particulars, which then allows, syntactically, the reference to an individual (an instance of the class *Universal*) via the OWL constructor 'value'.

Pattern UNIV$_1$: R_T equivalentTo R and '**refers to**' value **t** and '**inheres in**' some B

The stroke risk example would be expressed as:

> '*Increased risk of Stroke*' equivalentTo
> '*Increased risk*' and '**refers to**' value $\mathbf{Stroke_{univ}}$ and '**inheres in**' some *Brain*

This modeling pattern would require a parallel hierarchy of universals, as already mentioned in [17]. We propose a mechanism that can circumvent this:

Pattern UNIV$_2$: R_T equivalentTo R and
'**refers to**' some (*Universal* and '**extends to**' some T) and '**inheres in**' some B

We state that a universal (technically an OWL individual) is characterized by extending to a class of 'real' particulars:

> $\mathbf{T_{univ}}$ Type (*Universal* and '**extends to**' some T)

Thus, the reference to OWL individuals is no longer explicit:

> $\mathbf{Stroke_{univ}}$ Type (*Universal* and '**extends to**' some *Stroke*)

then we can bring the two statements together as follows:

> '*Increased risk of Stroke*' equivalentTo '*Increased risk*' and
> '**refers to**' some (*Universal* and '**extends to**' some *Stroke*) and '**inheres in**' some *Brain*

However, this approach has its limitations whenever the target is a mere concept or idea, which does not fulfil the requirements of being a universal, *viz.*, being capable of multiple instantiation, existing independently of particular minds, and having a consistent identity across all its instances [24]. This would preclude, e.g., the following:

> '*Whole brain transplantation plan*' equivalentTo
> '*Plan*' and '**refers to**' some
> (*Universal* and '**extends to**' some '*Whole brain transplantation*')

simply because there is no instance of a whole brain transplantation, and, consequently, no universal that extends to it. Similar problems would arise with objects of dubious or unknown existence, as discussed in [14]. Thus, the UNIV patterns depends on the acceptance of universals as parts of reality, a position that is not necessarily shared by all ontologists, but probably deemed as irrelevant by most ontology users. Realism (or Universalism) holds that universals exist independently of our minds. Conceptualism, by contrast, sees universals as mental constructs, which reflect how humans organize and understand the world. Nominalism denies the existence of universals altogether, treating general terms as mere names or linguistic conveniences. A more inclusive approach would therefore be to abstract away from universals to 'repeatables', a term coined by Armstrong [24], which stands for concepts, categories, ideas, types, kinds, and universals of several flavors. How such an ontology extension could look was demonstrated in KOSonto [25].

4.5. Collections as extensions of repeatables

However, what the three positions 'Conceptualism', 'Nominalism', and 'Universalism/Realism' have in common is that universals, concepts, or names can be formally described (intensions) and extend to collections of particular things in the domain of interest (extensions). The last pattern we propose no longer sticks to the notion of universals and focuses, instead, on the extensions of repeatable, here introduce as the class *Collection*, regardless of any underlying school of thinking.

Pattern COLL: R_T equivalentTo R and '**refers to**' some (*Collection* and '**has member**' some T)

In the case of the stroke risk example, we have

'*Increased risk of Stroke*' equivalentTo
'*Increased risk*' and '**refers to**' some
(*Collection* and '**has member**' some *Stroke*) and '**inheres in**' some *Brain*

Thus we would continue having OWL-EL expressiveness without requiring a concrete stroke referent. Instead, we have all strokes in the world as a collective referent, which supports us in characterising the nature of a possible stroke in the case of a given patient about whom a stroke risk was asserted. It means that such a potential stroke resembles the elements of my collection class. In contrast, in the case of our whole brain transplantation example, the collection of whole brain transplantations is currently empty, i.e., the OWL class '*Whole brain transplantation*' has no instances even if we define it as follows.

'*Whole brain transplantation*' equivalentTo
Transplantation and '**has participant**' some '*Whole brain*'

Although the open world assumption, which holds for the interpretation of description logics, remains agnostic regarding class extensions unless we explicitly state it, we want to avoid dealing with empty collections, so that we do not have to solve the puzzle whether all empty collections are equivalent to manned Mars missions, unicorns, or whole brain transplantations or not. On the other hand, such a class would no longer be empty in case a future whole brain transplantation is accomplished. Under this assumption, the intersection of the class of all transplantations with the class of all processes with whole brains participating would no longer be empty. We could nevertheless define '*Whole brain transplantation plan*') in our ontology as such:

'*Whole brain transplantation plan*' equivalentTo
'*Plan*' and '**refers to**' some
(*Collection* and '**has member**' some '*Whole brain transplantation*')

The acceptability of this approach in scientific ontologies is contentious among ontologists. Therefore, it is preferable to avoid dealing with classes known to be empty. Taking again the brain transplant example, we redefine it using the intersection of two target collections, each of which is known to be non-empty, *viz.*, the collection of all transplantations and the collection of all processes with participation of a whole brain:

‘*Whole brain transplantation plan 2*’ equivalentTo
Plan and
‘**refers to**’ some (*Collection* and ‘**has member**’ some *Transplantation*) and
‘**refers to**’ some (*Collection* and ‘**has member**’
some (‘**has participant**’ some *Whole Brain*))

The class ‘*Whole brain transplantation plan*’ is recognized by a DL reasoner as a subclass of ‘*Whole brain transplantation plan 2*’. However, if we use closure axioms, accepting the OWL-DL complexity, then these two classes are recognized as equivalent. To this end, we introduce the pattern COLL-C:

Pattern COLL-C: R_T equivalentTo R and ‘**refers to**’ only (*Collection* and
‘**has member**’ some T and ‘**has member**’ only T)

‘*Whole brain transplantation plan*’ equivalentTo
‘*Plan*’ and ‘**refers to**’ only (*Collection* and
‘**has member**’ some ‘*Whole brain transplantation*’ and
‘**has member**’ only ‘*Whole brain transplantation*’)

‘*Whole brain transplantation plan 2*’ equivalentTo
Plan and ‘**refers to**’ only (*Collection* and
‘**has member**’ some *Transplantation* and
‘**has member**’ only *Transplantation* and
‘**has member**’ some (‘**has participant**’ some ‘*Whole brain*’) and
‘**has member**’ only (‘**has participant**’ some ‘*Whole brain*’))

A description logics reasoner that supports OWL-DL expressiveness, e.g. HermiT, states:

‘*Whole brain transplantation plan 2*’ equivalentTo ‘*Whole brain transplantation plan*’

5. Discussion

The proposed introduction of the relation ‘**refers to**’ provides a neutral and versatile mechanism for linking a broad spectrum of entities—from classical realizable entities (REs) such as dispositions and functions to information content entities (ICEs)—to their referents. The conceptual range spanning from functions to information content entities is not new; for example, T. Deacon [11] coined the term “ententionality” (sic!) to describe all phenomena that are about something, have a purpose, or exhibit function. These entities share the feature that the ontological status of their targets is uncertain, deliberately underspecified, or simply incomplete. Such phenomena arise whenever reference is made generically to a *type* of entity relevant to describing a state of affairs—where that type may not be instantiated, may be unknown if instantiated, or may not currently apply to any existing particular entity. Examples include diagnostic hypotheses (“the patient may have the disease D”), risk assessments (“the patient has a high probability of developing D in the future”), scientific hypotheses (“we posit the existence of an elementary particle E with precisely these properties”), as well as plans and intentions (“a manned Mars mission”, “a whole brain transplantation”), and references to non-existent entities (“a statue representing a minotaur”). Thus, our proposal supports epistemic flexibility, allowing for the representation of the hypothetical, speculative, future, and even the impossible. At the same time, it ensures ontological neutrality by avoiding commitment to specific metaphysical interpretations, e.g., whether a referent is understood as a Platonic type, an Aristotelian universal, a mental concept, or a linguistic expression. Finally, it promotes representational parsimony by offering modeling patterns with a minimal set of logical constructors, ensuring high scalability, particularly within the OWL-EL profile, although this comes at a price because OWL-EL lacks closure axioms involving value restrictions.

5.1. Realizables (REs) and Informational Entities (ICEs): A Convergent View

The rationale for treating realizable entities (REs) and information content entities (ICEs) in a unified manner lies in the nature of their targets: both can refer to either particulars or repeatables. By employing a single relation, **refers to**, across all such cases, we also avoid the need to enforce a strict boundary between realization and denotation—a boundary that is often conceptually ambiguous. This ambiguity is especially pronounced when dispositions or functions are not regarded as entities with intrinsic ontological status, but rather as conceptual or linguistic constructs, as in DOLCE [26], a perspective aligned with nominalist or empiricist approaches, and typically opposed by ontological realism. Recognizing that creators of domain ontologies are generally less concerned with deep philosophical distinctions, our approach offers the advantage of minimal ontological commitment and, therefore, maximal modeling flexibility. At the same time, ontology developers who find the **refers to** relation too general are free to refine it by introducing subrelations such as **is realized by** or **is about**. Otherwise, they may opt not to decide whether a given RE is dispositional or referential—particularly in cases where the ascription of functions or dispositions is grounded in observed behavior and serves primarily to support scientific predictions. This may be the case in the following two examples: *Increased risk of stroke* can be understood as referring to the class *Stroke*, functioning as a shorthand for an empirical observation about a group of individuals who experience strokes at significantly higher rates than the general population. Similarly, attributing the disposition *Elasticity* to a class of objects that return to their original shape after deformation can be interpreted as rather referential—especially when contextual parameters (e.g., pressure or temperature) are unspecified, e.g., polyethylene is brittle at $-200°$C and plastic at $+100°$C).

5.2. Collections as Referents

A key novelty of the proposed modeling approach is the use of collections as targets in cases where no particular entity serves as a referent. But what does it mean to '*refer to*' a collection when the relevant instance may or may not be included in it? This issue becomes especially pertinent when the collection is empty. From an intensional modeling perspective, it is the defining properties of membership in a collection—rather than its current extension—that matter. This allows us to maintain meaningful reference even when the referent has never been instantiated or may only be instantiated in the future. Nevertheless, by adopting this view, the ontological status of collections remains open. In the spirit of "minimal ontological commitment", whether collections are interpreted as abstract classes or concrete mereological sums, or seen as extensions of universals, concepts, or names, does not affect modeling decisions, provided we restrict ourselves to the relation '**has member**' between a collection and each of its constituents. And even if we know that the referent is empty at a given time, we can refer to it by the conjunction of two or more non-empty referents. As long as its emptiness is not explicitly stated, different intensional meanings of currently empty referents can be distinguished and and employed in inference.

5.3. Limitations and Open Issues

While the relations '**refers to**' and '**has member**' offer a promising framework for simplifying and generalizing reference modeling, several open issues remain:

- **Collection Semantics:** The assumption that a formal clarification of the exact nature of collections is unnecessary is an untested assumption;
- **Ambiguity:** It remains an open question whether introducing '**refers to**' introduces ambiguity compared to more precise relations such as '**is about**', '**realized by**', etc., as well as with respect to conflations between plans and plan specifications—as strictly separated in OBI [12], based on IAO;
- **Reasoning:** It remains to be demonstrated in detail which patterns support which use cases and at what (computational) cost. The patterns we have introduced cover only the cases of reference, where the referent is not concrete. Cases of reference to concrete entities are not included within our patterns for abstract reference. Notably, reference to a particular entity **e** does not entail reference to the collection of entities of type *E*. This could be addressed by a disjunctive definition that combines the case of "reference to an individual of type *R*" with that of "reference to the type *R*" (as collections of individuals of type *R*). Whether this would allow for defining REs that are never realized should be subject to further investigation;
- **Roles and Intentions:** As particular kinds of REs they have been bypassed by our examples and discussions, although they would deserve further analysis;
- **Epistemic realizables:** It remains to be investigated whether the proposed patterns can accommodate diverging referents of an ICE. For instance, can the same ICE be simultaneously about a general topic and about a distinct truth statement at the instance level [27]? As an example, a physician discussing the possibility of cancer (as a topic) with a known healthy patient (known absence of cancer) should not lead to a contradiction within the ontological framework;
- **Lack of Empirical Use Cases:** No practical validation has yet been conducted through domain-specific modeling patterns or ontology design case studies guided by competency questions.

These concerns suggest the need for further theoretical refinement, practical modeling guidelines, and evaluation of the proposed approach in Applied Ontology settings.

6. Conclusion and Outlook

We proposed a minimal set of patterns based on the relations **refers to** and **has member**, together with the class *Collection*, to express generic reference, i.e., reference to repeatables (types, concepts). The relation **refers to** provides a unified and ontologically neutral mechanism for linking both realizable entities (REs) and information content entities (ICEs) to either particular instances or collections. This enables flexible modeling of reference in contexts of epistemic uncertainty, hypothetical scenarios, or absent instantiations, without committing to specific metaphysical positions on universals or concepts. By allowing reference to collections understood as extensions of repeatables, the model supports meaningful representation even when the existence of concrete instances is unknown or uncertain. We further showed that reference to non-existent entities can be modeled via intersections of collections, relieving ontology engineers from the need to introduce classes that are known or presumed to be empty. A practical advantage of this approach is that the notions of realization and reference—as well as intermediate cases such as plans—can be captured using a single relation within general design patterns. These patterns preserve neutrality with respect to competing philosophical views. Finally, we demonstrated that the approach can be supported—albeit with some compromises—by scalable and tractable reasoning, meeting a core requirement for ontology-based applications in domains such as biomedicine and clinical informatics, where large ontologies demand efficient and computable representations.

Funding

This work is supported by the Horizon Europe Framework Program under Grant Agreement Nos. 101057062(AIDAVA), 101112022 (iCare4CVD), 101095435 (REALM), 101057603 (RES-Q+), 101080875 (CMC), and the Spanish Research Agency Grant RYC2020-030190-I.

Declaration on Generative AI

During the preparation of this work, the main author used ChatGPT for grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] F. Toyoshima, A. Barton, L. Jansen, J.-F. Ethier, Towards a unified dispositional framework for realizable entities, in: Formal Ontology in Information Systems, IOS Press, 2021, pp. 64–78.

[2] S. Choi, M. Fara, Dispositions, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Spring 2021 ed., Metaphysics Research Lab, Stanford University, 2021.

[3] J. Röhl, L. Jansen, Why functions are not special dispositions: an improved classification of realizables for top-level ontologies, Journal of biomedical semantics 5 (2014) 27.

[4] W. D. Duncan, M. Diller, D. Dooley, W. R. Hogan, J. Beverley, Concretizing plan specifications as realizables within the OBO Foundry, Journal of Biomedical Semantics 15 (2024) 15.

[5] W. Ceusters, B. Smith, Aboutness: Towards foundations for the information artifact ontology, in: W. Ceusters, B. Smith (Eds.), Proceedings of the Sixth International Conference on Biomedical Ontology (ICBO), CEUR vol. 1515, 2015, pp. 1–5.

[6] M. Dumontier, R. Çelebi, K. Gilani, I. de Zegher, K. Serafimova, C. M. Costa, S. Schulz, SULO – a simplified upper-level ontology, in: Proceedings of FOUST'25: Workshop on Foundational Ontology, Catania, Italy, 2025. Sep 08–12, 2025.

[7] P. F. Patel-Schneider, OWL 2 web ontology language new features and rationale (2nd edition), 2012. URL: https://www.w3.org/TR/owl2-new-features/.

[8] J. Röhl, L. Jansen, Representing dispositions, Journal of biomedical semantics 2 Suppl 4 (2011) S4.

[9] SNOMED International, SNOMED CT starter guide, 2025. URL: https://confluence.ihtsdotools.org/display/DOCSTART.

[10] S. Schulz, J. T. Case, P. Hendler, D. Karlsson, M. Lawley, R. Cornet, R. Hausam, H. Solbrig, K. Nashar, C. Martínez-Costa, Y. Gao, SNOMED CT and Basic Formal Ontology – convergence or contradiction between standards? the case of "clinical finding", Applied ontology 18 (2023) 207–237.

[11] T. W. Deacon, Incomplete nature: How mind emerged from matter, WW Norton & Company, 2011.

[12] A. Bandrowski, R. Brinkman, M. Brochhausen, M. H. Brush, B. Bug, M. C. Chibucos, K. Clancy, M. Courtot, D. Derom, M. Dumontier, et al., The ontology for biomedical investigations, PloS one 11 (2016) e0154556.

[13] J. N. Otte, J. Beverley, A. Ruttenberg, BFO: Basic Formal Ontology, Applied Ontology 17 (2022) 17–43.

[14] S. Schulz, M. Brochhausen, R. Hoehndorf, Higgs bosons, mars missions, and unicorn delusions: How to deal with terms of dubious reference in scientific ontologies., in: ICBO 2011 – 2nd Int. Conference on Biomedical Ontologies, CEUR vol. 833, 2011, pp. 183–189.

[15] A. Gangemi, V. Presutti, Ontology design patterns, in: Handbook on ontologies, Springer, 2009, pp. 221–243.

[16] M. Horridge, P. F. Patel-Schneider, OWL 2 web ontology language Manchester syntax (second edition), 2021. URL: https://www.w3.org/TR/owl2-manchester-syntax/.

[17] S. Schulz, C. Martínez-Costa, D. Karlsson, R. Cornet, M. Brochhausen, A. Rector, An ontological analysis of reference in health record statements, IOS Press, 2014.

[18] Y. Kazakov, M. Krötzsch, F. Simancik, ELK reasoner: Architecture and evaluation, in: I. Horrocks, M. Yatskevich, E. Jimenez-Ruiz (Eds.), Proceedings of the OWL Reasoner Evaluation Workshop 2012 (ORE), CEUR vol. 858, 2012.
[19] M. Boeker, I. Tudose, J. Hastings, D. Schober, S. Schulz, Unintended consequences of existential quantifications in biomedical ontologies, BMC Bioinformatics 12 (2011) 456.
[20] A. Khamparia, B. Pandey, Comprehensive analysis of semantic web reasoners and tools: a survey, Education and information technologies 22 (2017) 3121–3145.
[21] F. Loebe, P. Burek, H. Herre, GFO: The General Formal Ontology, Applied Ontology 17 (2022) 71–106.
[22] G. Guizzardi, A. B. Benevides, C. M. Fonseca, D. Porello, J. P. A. Almeida, T. P. Sales, UFO: Unified Foundational Ontology, Applied Ontology 17 (2022) 167–210.
[23] M. Dumontier, C. J. Baker, J. Baran, A. Callahan, L. Chepelev, J. Cruz-Toledo, N. R. Del Rio, G. Duck, L. I. Furlong, N. Keath, et al., The Semanticscience Integrated Ontology (SIO) for biomedical research and knowledge discovery, Journal of biomedical semantics 5 (2014) 1–11.
[24] D. M. Armstrong, Universals: An opinionated introduction, Boulder, San Francisco, London: Westview Press, London, England, 1989.
[25] J. N. Nikiema, F. Mougin, V. Jouhet, S. Schulz, KOSonto: An ontology for knowledge organization systems, their constituents, and their referents, in: International Conference on Biomedical Ontologies 2023 (ICBO 2023), volume 3603, 2023, pp. 130–141.
[26] S. Borgo, R. Ferrario, A. Gangemi, N. Guarino, C. Masolo, D. Porello, E. M. Sanfilippo, L. Vieu, DOLCE: A descriptive ontology for linguistic and cognitive engineering1, Applied Ontology 17 (2022) 45–69.
[27] S. Yablo, Aboutness, Princeton University Press, Princeton, NJ, USA, 2017.