The Mereology of Concepts: Preliminary Explorations

Guendalina Righetti[1]

[1]Department of Philosophy, Classics, History of Art and Ideas, University of Oslo, Blindernveien 31 Georg Morgenstiernes hus 0313 Oslo

Abstract

Do concepts have parts? In this paper, we take the simplifying assumption that concepts can be understood as (more or less structured) bundles of features and discuss the distinction between features and concepts in terms of the distinction between parts and wholes. In doing so, we will leverage different ideas related to Fine's theory of parthood.

Keywords

Concepts, Features, Parts, Constructional Approach

1. Introduction

The notion of concepts has been deeply studied and developed across different fields, such as philosophy, logic, and cognitive science. Although the questions investigated in the diverse fields are somehow orthogonal (what are concepts, are they just in our mind or do they exist in the reality? How should we better represent them formally? How do humans conceptualise, categorise and combine concepts?), this has led to distinct accounts, and explanations for the role concepts have in our reasoning over, and understanding of, the world.

In the literature, the notion of concept often relies on a subsidiary one, that of *feature*. The features, sometimes called attributes or characteristics, of a concept are usually assumed to play a role in defining or describing it. In the context of psychological theories of concepts [1], in particular, various theories suggest different ways in which humans structure and represent concepts in our semantic memory. While these theories may differ in many aspects, many of them converge on the use of the notion of features.

According to the Prototype Theory, concepts are represented in our semantic memory as prototypes. Prototypes consist of sets, or lists, of *features*, each assigned a certain degree of importance for describing the concept. This list of features encapsulates and describes the most typical exemplars of the concept [2, 3]. Examples of features are

> discrete attributes such as *has legs*, *you drive it*, or *the letter B is a member*. These are the kinds of features of natural semantic categories which can be most readily reported and the features normally used in definitions of categories by means of lists of formal criteria. [2, p. 576]

Proceedings of the Joint Ontology Workshops (JOWO) - Episode X: The Tukker Zomer of Ontology, and satellite events co-located with the 14th International Conference on Formal Ontology in Information Systems (FOIS 2024), July 15-19, 2024, Enschede, The Netherlands.

guendalina.righetti@ifikk.uio.no (G. Righetti)
0000-0002-4027-5434 (G. Righetti)

The knowledge view offers a different representation of concepts, but still makes use of the notion of feature: concepts consist of micro-theories registering the complex network of relations existing between the *features* describing a concept and our prior or background knowledge.

> Most people think of birds as being feathered, two-legged creatures with wings, which fly, lay eggs in nests, and live in trees. [...] With simple, mundane knowledge, one can explain many of these features. Let's start with flying. In order to fly, the bird needs to support its weight on wings. The feathers are important as a very lightweight body covering that also helps to create an aerodynamic form. Thus, wings and feathers enable flying. By virtue of flying, the bird can live in nests that are in trees, because it can easily fly into and out of the trees. [1, p.143]

The distinction between features and concepts stems, arguably, from philosophy. A similar pattern is at least endorsed by supporters of the so-called Bundle theory, proposed initially by Hume in his Treatise [4]. Accordingly, it would be possible to distinguish between *simple perceptions*, which are basic, indivisible and unanalysable, and *complex* ones, which can be analysed in their constituent parts[1] :

> Tho' a particular colour, taste, and smell are qualities all united together in this apple, [...] they are not the same, but are at least distinguishable from each other" [4, pp.7-8]

This tradition has been picked up by Kant. According to Kant, a concept is a collection of marks, or characteristics [5, 6], something quite close to what elsewhere is called a feature. This view emerges in the discussion of synthetical judgments such as "all bodies are heavy":

> though at first I do not at all include the predicate of weight in my conception of body in general, that conception still indicates an object of experience, a part of the totality of experience, to which I can still add other parts; and this I do when I recognize by observation that bodies are heavy. I can cognize beforehand by analysis the conception of body through the characteristics of extension, impenetrability, shape, etc., all which are cogitated in this conception. [7]

A concept is thus built, through experience, as a collection of marks, and, significantly, can in turn be used as a mark itself [6]. While there appears to be a distinction between features and concepts, ultimately, they can both be regarded as "marks" to be combined with other marks.

Not all concepts can be understood as bundles of features, nor do all agree that concepts can be represented as some kind of feature aggregations. In philosophy, an example is Frege discussion of the concept of *relation* [8]. In cognitive science, an example is the exemplar theory [9], which endorses a purely extensional representation of concepts. Notwithstanding this, the bundle theory has had quite a success in concept representation and turns out particularly useful in the context of intensionally-oriented representation of concepts.

The precise definition of what qualifies as a *feature*, and the distinction between features and concepts, is typically taken for granted. Notwithstanding this, a common thread among

[1] Although here Hume had in mind particular objects, the distinction applies also to simple and complex *ideas*, and thus translate also to the case of concepts, at least to what he calls imaginary concepts.

the perspectives discussed above appears to be the idea that concepts are comprised of various features, which consequently serve as *parts* of the concepts.

That mental representation can be seen as constituted by parts has already been noted by Leech [10]. Perhaps unsurprisingly, this discussion has been undertaken in the context of Kant's analysis of concepts, to distinguish between intentions (particular perceptions) and concepts.

Similarly, although not as explicitly, cognitive theories of concepts too, often, look at the features as *parts* of the description of the concepts:

> For example, dogs have four legs, bark, have fur, eat meat, sleep, and so on. Some subset of these features might be part of the definition, rather than only one characteristic.[1, p.16]

And, analogously, when building up concept descriptions in description logics and formal ontology, we often speak of the different *parts* composing the definition of the concept. As a matter of fact, logical operators are called concept constructors in DL, reflecting the idea that concept definition can be built from different concept parts.

In this paper, we take the simplifying assumption that concepts can be understood as (more or less structured) bundles of features and discuss the distinction between features and concepts in terms of the distinction between parts and wholes. In doing so, we will leverage different ideas related to Fine's theory of parthood [11, 12, 13].

2. Kinds of Parts

The possibility of studying the mereology of abstract objects has been advocated by several authors, among which a place of honour is held by Fine. In Fine's words [13, p.561],

> Philosophers have often supposed the notion of part only has proper application to material things or the like and that its application to abstract objects such as sets or properties is somehow improper and not sanctioned by ordinary use. But I suspect that this is something of a philosopher's myth.

More specifically, the possibility of describing a mereology of mental representations has already been advocated by Leech [10]. Declining her discussion to the case of concepts, if we assume concepts are mental representations of reality, and if we believe that reality is organised in a part-whole fashion, then it is not so implausible to suppose that the features describing a concept reflect the part-whole structure of the reality it represents. Nevertheless, when considering concepts as (some kind of) sums of their features, this involves considering as parts also things that would not be looked at as parts in the physical realm. An example is the case of qualities. If I describe *fruit* as sweet, and consider this one of the features of the concept fruit, then the quality of being sweet becomes a *part* of the concept fruit.

The idea of looking at concepts as representational devices made out of parts, possibly contributing in different ways to the description of the concept as a whole, has been discussed also in the context of Knowledge Representation, see [14].[2]

We are assuming here a not too precisely specified notion of concept. The quite intuitive (and rather vague) description would correspond to a mental representation consisting of a bundle of

[2]For a discussion of the representation of mereological relations in the same context see in particular [15].

features, without specifying too much how the latter are arranged, or structured. How features are structured will depend on the specific theory or perspective one assumes on concepts. If we adopt a cognitive stance and embrace prototype theory, concepts will be organised into weighted lists of (more or less) independent features. If we assume the knowledge view, concepts and their features are organised in some kind of network structure. In this scenario, concepts and their features are clustered according to their corresponding theories, integrated into a larger network interconnecting various parts of our knowledge.

A similar argument can be applied to formal ontologies. Concepts in the ontologies are built through definitions. Although we do not directly refer to them as composed of features, definitions can also be seen as construed by different parts. In this case, parts are organised and structured through the use of logical connectives.

If the relation between features and concepts is one of part-whole, what kind of part are the features? Following Fine, parts can be classified in at least two dimensions [12, 13]. First, a way of being part can be either *basic* or *derivative*. A basic part-whole relation is not definable in terms of other, more basic, part-whole relations. Examples are how a letter is part of a word, or a set-member is part of a set. A derivative parthood relation depends instead on more basic parthood relations. For instance, saying that the weekend is a *small part* of the week, involves a relation of parthood that only holds if the weekend is part of the week [13]. Second, parthood can be *absolute*, or timeless, when it is not subject to time. As an example, Fine suggests the way in which a pint is part of a quart. Otherwise, it can be *relative*, when it changes over time and is thus subject to circumstances. For instance, a tyre that was once part of the car [12].

If we accept that features can be parts of concepts, then it seems safe to assume that they are parts in a basic way. Fine delineates two ways by which a parthood relationship may be derived: through subsumption, as observed in the case of small parts, or via chaining, wherein parthood is established through a concatenation of parthood relations [13]. Neither of them seems to reflect the intuitive way in which a feature is part of a concept.

The difference between absolute and relative parthood is not as clear. If we consider concepts as bundles of features, then necessarily concepts will be constituted by some features at any time. In this broad sense, then, features are part of the concept in an absolute way, as much as pints are parts of gallons or members are parts of sets.

Nevertheless, when considering everyday concepts, it is also quite evident that they can change over time. Arguably, their features-parts change as well. As concepts are representations of our reality, and reality changes over time, this doesn't sound very surprising. A case in point is conceptual evolution, namely the evolution of conceptual meaning over time, due to societal or cultural changes. Consider the concept of *marriage* as an example. Rather obviously, the features used to describe the concept two centuries ago were quite different from the ones one would use to describe it today (at least in some parts of the world). This argument has been used by Richards [16] to defend his position on concept evolution, according to which concepts are comparable to biological species evolving over time. Another example may be the one of *game*. Although some features are probably the same as they were in the '50s, many of them have changed (think of video games).

This phenomenon may initially appear primarily associated with the use of natural language. Nevertheless, formal definitions of concepts within formal ontologies (or their *parts*) change as well, for several reasons. For instance, one might find that the initial definition was overly rigid and failed to encompass the intended domain of application adequately. Additionally, revisions may be prompted by scientific advancements that necessitate a reassessment of our understanding and, consequently, our definitions. For this reason, ontologists have developed strategies to keep track of successive versions of their ontologies [17].

At least two aspects are worth noticing in the present discussion. First, conceptual evolution opens the problem of finding appropriate identity criteria for the identification of concepts through time. If concepts are subject to changes as are physical objects, it seems reasonable to wonder how much a concept can change while remaining the same. How many features-parts is it possible for a concept to lose, add or change without losing its identity? Or, in a similar fashion but with a different terminology, for how long is it possible for a concept in an ontology to maintain its identifier through ontology revisions? Different views exist in philosophy on similar issues, from positions endorsing the possibility of concepts to change through maintaining their identity to people denying this possibility (see [18], specifically Chapter 10). Despite the similarity with the case of physical objects, we are looking at concepts in terms of mental representation, and thus as (some kind of) abstract objects. Given the nature of concepts and some previous considerations, assuming purely extensional criteria of identity may easily bring us in the wrong direction.

A second (arguably related) aspect to note is that in many cases, although some features may change over time, others are persistent to the point of appearing as timeless parts of the concept. Let us consider the concept of game again. Arguably, although many of its features have changed over the years, it can quite stably be described as a kind of activity. Being an activity is a stable (one may say necessary) feature of game across the evolution of the concept. These kinds of features may be seen as the semantic core of concepts and could be considered to play a role in selecting the correct identity criteria for concepts.

3. Rigid, Variable and Semi-Variable Embodiments

When discussing the distinction between timeless and temporary parthood, Kit Fine introduces two new kinds of objects/wholes, namely rigid and variable embodiments [12].

Rigid embodiments are introduced to exemplify objects such as sandwiches or water molecules. In such cases, Fine claims, a certain form or structure (e.g. the form of a ham sandwich, being structured as two pieces of bread with the ham *between* them) is embodied in the fixed, rigid, matter constituting the object (in the case of the ham sandwich, the two slices of bread and the ham). In such kinds of wholes, both the structure (or *principle of rigid embodiment*, in Fine's terminology) and the fixed matter constitute timeless parts of the rigid embodiment.

Variable embodiments are instead used to describe wholes that are extended in time and have hence (possibly) temporary parts. Fine's example is about the variable quantity of water in a river. Accordingly, there exists a principle, or function, that picks out the amount of water in the river at any time in which the river exists. There exists then a kind of whole, the variable water, that corresponds to that principle. The variable embodiment corresponds in this case to the quantity of water picked out by the principle at any given time t. The principle picking out the water is called the *principle of variable embodiment*, and the various objects picked out by the principle are the *manifestations* of the variable embodiment.

The two theories can be combined. According to Fine's example, a particular car is, at a given time t, a rigid embodiment, constituted by several parts arranged according to a specific structure (or principle of embodiment), the one characteristic of a car. However, each of these parts can be subject to change, which then affects, and changes, the rigid embodiment as a whole as well. One can then consider the car as the variable embodiment whose manifestations are all the rigid embodiments corresponding to these changes. Moreover, each part of the particular car is constituted, in turn, by several parts (the motor is constituted by a cylinder head, pistons, etc), which are subject to change as well. The shift between rigid and variable embodiments applies to them as well.

The theory of rigid and variable embodiment has had tremendous success (although it did not escape some criticism, e.g. [19]), both in philosophy and in applied ontology, where it has been used to model many different objects and phenomena, such as organisations [20], processes and events [21, 22], collective intentionality [23], powertypes [24] groups [25], and even noun phrases [26].

Can the theory be applied to concepts as well? As already discussed, concepts vary in time: they can lose (or gain) some of their parts due to societal or cultural changes or even knowledge revision. Similarly to the variable water of a river, they seem to be describable in terms of variable embodiments. As some of the parts of a concept may vary, the manifestations of the concept-variable embodiment will correspond to different rigid embodiments at different times.

The question is then what would be the principle of variable embodiment in the case of concepts. This discussion would require a dedicated paper, but one can try to give some intuitions. Similarly to what was observed by Ferrario et al. in the context of organisations [20], it is not implausible to imagine that the principle of variable embodiment needs to be constrained by the history of the concepts. A possibility is thus to consider the concept classification, and the classification coherence (in the sense of [27]), to play a role in defining the principle of variable embodiment applying to concepts. What we mean by classification, is not, in this case, necessarily a matter of classification of instances (concepts may vary their extensions due to the conceptual changes), but the classification of concepts in higher categories. For example, consider the concept of *marriage*, and suppose it is categorised at a higher level as a *contract*. Then, although the concept may vary, it will still do it in a way that is coherent with this classification: otherwise, it would be a different concept. Another possibility is to constrain the principle through the notion of *topic*: as argued by Cappelen [18], although concepts may vary, their topic remains the same. Further work is needed to settle this question.

Even when one does not want to take a specific position over the precise structure of concepts, it may be uncontroversial to state that concepts do have some structure, that being in our semantic memory or our ontology. As mentioned above, if one assumes prototype theory, the (weighted) features of concepts will be structured in the form of a list. In the case of the knowledge view, the structure will be that of a network. In the case of concept definitions in an ontology, some suggest employing Aristotelian definition [17], thus imposing a structure of genus-differentia on concept representation. Considering the different manifestations of a concept in terms of rigid embodiments can account for the structure of the representation one assumes. One can see all of the different representational choices as taking a stance on the principle of rigid embodiment giving form to concepts.

We concluded the previous section by mentioning a peculiar aspect of concepts, namely the fact that some of their parts do not change over time. In other words, taking all of the possible manifestations (rigid embodiment) of concepts, some of the parts will be the same across time[3].

(At least some) concepts have parts that are somehow essential to their description, or definition - features that once changed, completely change the meaning of the concept.

Cognitive theories of concepts often deny the possibility of providing clear-cut, stable definitions for everyday, common-sense concepts. In specific domain, specifically when we are *not* dealing with ordinary concepts, though, it is indispensable to define concepts through definitions. A case in point are scientific domains, e.g. when defining biological (or mathematical, etc) concepts. In such cases, concept definitions will have necessary parts. Definitions are, after all, often assumed to consist of set of necessary and sufficient conditions for instances classification - a view that can be traced back to Aristotle. Arguably, in the case of concept definitions, changing the set of necessary and sufficient conditions would lead to a completely different concept. One may thus imagine that, in the case of some scientific domains, concepts are more properly represented as rigid embodiments.

Nevertheless, also cognitive theories of concept which reject the possibility of providing sharp definitions for concepts, often assume some parts of a concept description to be necessary [3]. Sloman et al. [28, p.191] seem to assume something similar in the following:

> Note that the mutability of a feature is concept-relative. For example, roundness is a mutable feature of oranges. Even if all oranges were round, our notion of orange would be substantially unaltered if we imagined one that was not. But roundness is an immutable feature of wheels. If a wheel is not round, then it has to be completely reconceived to retain its (mental) status as a wheel.

There seem to be different variations in the case of concepts. In some domains or contexts (e.g. mathematics, etc), one may want to treat concepts as rigid embodiments, always consisting of the same parts[4]. Once the parts change, the concept changes as well: it is a different concept. And other domains (as in the case of everyday concepts) where this would be a too-strong assumption. In such cases, concepts have both parts subject to change and timeless parts. We might call objects of this kind *semi-variable embodiments*: embodiments that (can) vary in almost all respects except for some essential parts.

Are concepts the only objects of this kind? There seem to be other cases.

Let us take the case of organisations. The first example involves (certain) rock bands. In many cases, all band members are replaceable except for the frontman. The departure of the frontman could lead to the dissolution of the band itself (consider, for a concrete example, the case of Nick Cave and the Bad Seeds [5]: many members of the band have changed through the years, but what would happen if Nick Cave left the band?).

---

[3]One may argue that all the features describing a concept may vary as well. For instance, we argued that being an activity is an essential part of games. Nonetheless, *activity* is also a complex concept, whose features can change across time. What we are arguing, though, is that it is possible, stepping up in this hierarchy of concepts, to find a semantic core that is stable.

[4]There might be exceptions in these cases as well. Consider, for instance, how the conception of *set* has changed from Frege to Gödel.

[5]See https://en.wikipedia.org/wiki/Nick_Cave_and_the_Bad_Seeds

Another example is some dictatorships, where all major figures involved in the governments (advisers, generals, high-ranking officials) can change, but the dictator, whose elimination would cause the cease of (at least) that specific dictatorship.

Organisations are not the only examples, as one can think of events that have essential parts as well[6]. Consider a birthday party. Almost every part can be variable (the guests may come and go, the location may change due to bad weather, etc), but the birthday girl is a stable, essential part of the event (if she gets sick, the birthday party is cancelled). Concerts follow a similar pattern: anyone in the audience can decide not to show up or to leave the concert, but if the singer ditches the concert, or has a problem during the concert, the event is cancelled (or ends).

As mentioned above, Fine's discussion of variable embodiment assumes the possibility of nesting variable (and rigid) embodiment (see again the car example).

A similar nesting can (and often do) occur with concepts. Take the concept of a *car*. One could define it as the class of vehicles that normally have four wheels and a motor, run on roads, can seat up to eight people and are used mostly to transport people (and not cargo)[7]. By extending this description with more contextual features, one could see the concept of *car* as a (semi-)variable embodiment, that, at different times, can have different features as part (think of the evolution of cars in the last 100 years, and how the corresponding concept has changed). Nevertheless, by looking at the features of the concept *car* presented above, *being a vehicle* is also a complex concept composed of several features, and as such can be seen as a (semi-)variable embodiment, nested into the concept car.

This is a critical aspect of the representation of concepts. In the previous example, we have a concept playing the role of a feature for another concept. What is the relation of parthood in this case? And, generalising, how does the relation of parthood behave in the case of the combination of two or more concepts?

4. Concepts, Concept Combination and Part Transitivity

The smallest characterisation of the parthood relation is Core Mereology, which assumes the relation to be a partial ordering (a reflexive, antisymmetric and transitive relation) [29]. Where P is the parthood relation, this is expressed through the following axioms.

$$\text{Reflexivity: } \forall x(Pxx) \tag{1}$$

$$\text{Antisymmetry: } \forall x \forall y(Pxy \wedge Pyx \implies x = y) \tag{2}$$

$$\text{Transitivity: } \forall x \forall y \forall z(Pxy \wedge Pyz \rightarrow Pxz) \tag{3}$$

Applied to the case of features (and concepts) this means that: 1) each feature is part of itself; 2) if a feature f_1 is part of another feature f_2, and the second feature f_2 is part of feature f_1 as well, then the two features are the same; and that 3) if some feature f_1 is part of another feature f_2, and feature f_2 is part of feature f_3, then also feature f_1 is part of feature f_3. Let us focus on

[6]The discussion may be more tricky in the case of events, where different perspectives assume different positions on what is *part* of the event.

[7]https://en.wikipedia.org/wiki/Car

axiom 3[8], and let us consider the example of the concept *car* again. We said the concept *car* can be described as being a vehicle with a bunch of other features. Let us take f_v to correspond to the feature *being a vehicle*, and let us f_c be the concept of *car*. Being a vehicle is used here as a feature of *car*, but it is a concept that can be described as a bunch of other features $f_{v_1}, ..., f_{v_m}$. We have thus $f_{v_i} P f_v$ $(1 < i < m)$ and $f_v P f_c$, and by transitivity $f_{v_i} P f_c$. If I say that a vehicle is a machine designed for self-propulsion[9], then this description become part of the concept *car*.

This may sound relatively unproblematic, as when we are describing the concept *car* as a vehicle, we want exactly to use the meaning of *vehicle* to describe cars. Peculiar to concepts is the possibility of combining them into more complex ones, and it is, after all, the idea behind compositionality that we can understand complex expressions based on their constituents' meaning. Nevertheless, when considering compound concepts, the transitivity of features may sound like a too strong requirement, and in some cases blatantly false. As an example consider the concept of *car bomb*. Although *is used mostly to transport people* is one of the features that is part of the concept *car*, we sure wouldn't say that it is part of the concept *car bomb*. The reason is that compounds often cannot be understood conjunctively as the union of the features of the two concepts, but involves more subtle semantics.

One may thus believe the case of conjunctive combination, putatively taking the union of the features of two concepts as describing the compound, to be easier to understand in terms of parthood relation. Even in the case of conjunctive combinations, however, assuming features transitivity does not work: it is easy to think of examples where transitivity fails, for instance, due to non-monotonic phenomena.

Consider a conjunctive concept such as *pet fish*. While, ordinarily, a *pet* may be described as an animal that lives in the house and is kept for a person's company or entertainment, a *fish* does not really fit the description of an entertaining animal[10]. In the combination of two concepts, some of the features may get lost, and some new features might be gained. Then we get to the point where we may want to consider the features of *pet* to be part of the concept *pet*; we may want to consider the concept *pet* to be part of the concept *pet fish*; but then we cannot guarantee the transitivity of the features from the first to the second.

Do we have to give up on the idea of treating features as parts? Another possibility is to argue that the features are part of a concept in a different sense in which a concept is part of a compound [11], and this because they are the effect of different kinds of *operations*, or *constructions*.

When discussing his theory of parthood [13], Kit Fine has advocated the priority of the composition operation (standing behind the whole) with respect to the relation of part-whole. More specifically, he has argued the importance of taking the "operation of composition as primitive rather than the more familiar relation of part-whole" [13, p.565]. This lies, more generally, in his constructive approach to ontology [11]. Accordingly, one's ontology can be generated through a process of construction. In this process, a set of constructors can be

[8]Depending on the choice of the primitive relation, axioms 1 and 2 may be dropped. Proper parthood, for instance, is irreflexive and asymmetric.

[9]https://en.wikipedia.org/wiki/Vehicle

[10]The case of *petfish* has been used in the literature as a counter-example for the representation of concepts as prototypes [30, 31].

[11]This is also done, sometimes, in the context of material objects: a finger is part of my hand, I am part of the department of Oslo, but my finger is not part of the department of Oslo.

iteratively applied to some basic elements, the *givens* in Fine's terminology, to generate new elements in the ontology.

We might then try to apply similar ideas to the discussion conducted so far. In this context, we consider the *features* as ontological givens. We have then (at least) two kinds of operations: a "feature aggregation" operation, which aggregates features into meaningful concepts description; and (at least) a "concept combination" operation, operating on concepts (hence more complex elements, available only at a later stage).

5. Towards Determining CLAP Profiles for Concepts

In the following, we will distinguish between a concept/feature aggregation operation and a concept composition operator, and discuss the CLAP principles on them [11]. CLAP principles (aka Collapse, Leveling, Absorption and Permutation) have been introduced by Fine as a way to determine identity conditions for wholes emerging from different composition operations. The principles are: 1) Collapse: the whole composed of a single component is identical to that very component; 2) Leveling: the embedding of components is irrelevant to the identity of the whole; 3) Absorption: the repetition of components is irrelevant to the identity of the whole; 4) Permutation: the order of the components is irrelevant to the identity of the whole [13, p.573].

In [13], Fine studies how such principles behave in the context of different composition, or construction, operations, namely sums, sets, strings and sequences. He shows, in particular, how the four operations differ, and how the resulting wholes can be distinguished by analysing the different CLAP principles they obey. The intuition behind the subsequent discussion is that concept aggregation and concept composition obey different CLAP profiles, and are thus different kinds of operations. Rather obviously, the operations analysed here are not as precisely defined as the cases discussed by Fine. In the following, we will be mainly sketching some intuitions starting from some simplifying assumptions. To discuss the CLAP principles, we will be putting aside the possibility of concepts' evolution through time. After all, Fine developed the CLAP principles to discuss the absolute (timeless) notion of part. It is not implausible to believe that, to capture the dynamics of conceptual evolution, further principles would be required.

5.1. Concept Aggregation

Assumptions Let us simplify by assuming that concepts consist of bundles of features, namely, they can be represented as an unordered list of features (cf. Section 1). Concept aggregation amounts to the operation of *putting together* the features needed to describe the concept. In other words, given the set of features $F = \{x_1, ..., x_n\}$, concept aggregation is the operation of defining the concept y as consisting of a set of features $x_1, ..., x_k$. We will be using the following notation $y = [x_1, ..., x_k]$. The result of applying the operation of concept aggregation will be called *concept description.*

Collapse $[y] = y$
Namely, the concept formed out of a single feature (or concept) is identical to that feature (or concept). This seems to be a harmless condition when applied to concept aggregation.

Leveling $[[x_1, ..., x_k], x_m, ..., x_n] = [x_1, ..., x_k, x_m, ..., x_n]$
Namely, the embedding of (other) concepts is irrelevant to the identity of the concept description. As discussed in Section 3, the nesting of concepts into concept descriptions does often occur (consider again the concept *car* being described as a *vehicle*). Arguably, the goal of this operation is to use the features of the nested concept in the concept description they are embedded into. The substitution of the nested concept with the features describing it wouldn't harm the concept description. Leveling seems thus to hold to the case of concept aggregation.

Absorption $[x_1, ..., x_k, x_k, x_m, x_m, ..., x_n] = [x_1, ..., x_k, x_m, ..., x_n]$
Namely, the repetition of features/concepts is irrelevant to the identity of the concept description. This condition may seem too strong when considering the use of concepts in natural language. Consider someone describing *wrens* as small small birds, emphasising just how tiny they are. And thus one may believe that absorption may affect the identity of the whole after all. This seems however more an effect related to the use of natural language than to the representation of concepts. Consider the case where the repetition is due to the nesting of concepts among the features of the concept description, e.g. $y = [[x_1, ..., x_m], x_m, ..., x_n]$. For leveling, we would get $y = [x_1, ..., x_m, x_m, ..., x_n]$, and by absorption $y = [x_1, ..., x_m, ..., x_n]$.

Permutation $[x_1, ..., x_k, x_m, ..., x_n] = [x_1, ..., x_m, x_k, ..., x_n]$
Namely, the order of the features is irrelevant to the identity of the concept description. As we are assuming concepts to consist of unordered lists of features, permutation holds in the case of concept aggregation.

Following this discussion and starting from our assumptions, concept aggregation satisfies all of the $CLAP$ principles. It is thus a sum-like operation [13].

5.2. Concept Combination

While in the case of concept aggregation, it is possible to make a simplifying assumption and consider concepts as relatively unstructured entities, the case of combined ones is different. Compounds exhibit a higher level of organization, and their meaning is intricately connected with their specific structure. Composition can occur at different levels of complexity. It can involve the combination of adjectives and nouns (e.g. *black cats*) or noun-noun combinations. Noun-noun combinations are generally assumed to have more subtle semantics[12], involving complex patterns of relationship between the features of the concepts to be combined. Some of these issues are described by Wisniewski [32]. Accordingly, there exists at least three types of combination, namely: relation-linking ones, emphasising a relation between the two concepts (e.g. *robin snakes* are snakes that eat robins); property-mappings ones, where one of the salient properties of one of the concept is mapped onto the other (e.g. *tiger spiders* are spiders with stripes); and conjunctive combinations, combining the features of both concepts (e.g. *poet painters* have the features of both poets painters). As the contribution of the two concepts to the compounds varies in the three cases, it is not implausible to believe that the three combinations

[12]Although also some adjective-noun combinations are trickier than others: consider for instance *fake diamonds*. Such cases escape a plain conjunctive interpretation as well.

correspond to somehow different kinds of operations, giving rise to different parthood relations. We will simplify the discussion in what follows by analysing the case of conjunctive concepts, and leave the discussion of the other two kinds of combination for future work.

Assumptions We assume that, as a starting point, conjunctive combinations can be described as taking the union of the features of the constituent concepts. Leveraging on the literature on compounds in cognitive psychology and linguistics [3, 33, 32], we also make the following assumptions. First, some features may be lost during the combination due to *inheritance failure* [3], which occurs when features from one concept are incompatible with some of the other. A classic example is the one of *pet fish*, mentioned in Section 4. To use a less trite example, the concept of *tool weapon* might not retain features like *is used in war* from *weapon* or *it is used in construction* from *tool*, as they do not fit the combined definition. Second, the order of concepts matters. In English noun-noun combinations, the Head and Modifier roles determine the compound's structure and meaning [33]. The Head defines the main category, while the Modifier adjusts it. For instance, a *tool weapon* will be primarily a weapon, with some features of the tool, or some functions fitting the concept of *tool* -vice versa a *weapon tool*. In some cases, this inversion might also induce different interpretations of the concept: a *weapon tool* might be interpreted as a tool used on weapons (maybe to fix them), whereas the same would not be possible for a *tool weapon*.

We will be using the following notation to indicate the combination of concepts $z = C(x_1, .., x_m)$.[13] The result of applying the operation of concept combination will be called *compound*.

Collapse $C(y) = y$
Namely, the compound composed of a single concept is the same as the single concept. Given the assumptions, the principle holds. The discussion of this case is analogous to what said in the context of concept aggregation.

Leveling $C(C(y, z), w) = C(y, z, w)$
Namely, the embedding of compounds is irrelevant to the identity of the compound.

Given our assumption, this principle is against our intuition. One of the assumptions we have made concerns the potential loss of features during the combination referred to as inheritance failure. Let us make the slightly stronger assumption that, in the case of incompatibility of two features, both of them are lost.

Now let us assume $y = [x_1, x_2]$, $z = [x_3, x_4]$ and $w = [x_2, x_4]$, and x_1 is incompatible with x_4. In the case of $C(C(y, z), w)$ we have $C(C([x_1, x_2], [x_3, x_4]), w)$, and as x_1 is incompatible with x_4 we drop both of them. Now we combine the remaining concept(s) with w, namely $C([x_2], [x_3], [x_2, x_4])$, and the feature x_4 is thus preserved through the combination. In the case of $C(y, z, w)$, on the other hand, all of the occurrences of x_4 would be dropped, and we would get $C([x_2], [x_3], [x_2])$.

[13] The assumptions made here are reminiscent of [34], which presents a formal treatment of conceptual combination in weighted DL.

Absorption $C(y, y, z, z) = C(y, z)$
Namely, the repetition of concepts is irrelevant to the identity of the compound. Analogously to the case of Collapse, the principle holds given our assumptions.

Permutation $C(y, z) = C(z, y)$
Namely, the order of the concepts is irrelevant to the identity of the compound. The principle is in contrast with the second of our assumptions, accounting for the distinction between the head and modifier concepts.

Following this discussion and starting from our assumptions, concept combination corresponds to the following pattern of principles: $C\not{L}A\not{P}$.

6. Conclusion and Future Work

In this paper we have started discussing the mereology of concepts. By assuming the features describing a concept to be *part* of the concept, we have discussed what kind of parts are features, and what kind of whole are concepts.

In Section 3 we have suggested applying the theory of variable embodiment to concepts, and we have suggested the existence of *semi-variable* embodiments. A crucial issue, in this context, is to define the principle of (semi-)variable embodiment applying here. We have suggested the role of coherence in the classification as a possible way of constraining the principle. Further work is needed to make this idea precise, and implementable in a formal context.

The second half of the paper aims at defining an operational approach to concept construction. It differentiate between two kinds of operations, concept aggregation and concept combination, and analyse the different properties of the two leveraging on Fine's CLAP principles.

The discussion is conducted at an informal level. To develop a constructional approach to concept aggregation and combination in a formal setting is a matter of future work.

References

[1] G. Murphy, The big book of concepts, MIT press, 2004.
[2] E. Rosch, C. B. Mervis, Family resemblances: Studies in the internal structure of categories, Cognitive psychology 7 (1975) 573–605.
[3] J. A. Hampton, Inheritance of attributes in natural concept conjunctions, Memory & Cognition 15 (1987) 55–71.
[4] D. Hume, A treatise of human nature, Oxford University Press, 2000.
[5] R. Hanna, Kant's theory of empirical judgment and modern semantics, History of Philosophy Quarterly 7 (1990) 335–351.
[6] M. Lewin, Kant on philosophy as conceptual analysis, Con-textos Kantianos: International Journal of Philosophy (2023) 11–20.
[7] I. Kant, J. M. D. Meiklejohn, T. K. Abbott, J. C. Meredith, Critique of pure reason, JM Dent London, 1934.
[8] G. Frege, Boole's logical calculus and the concept-script, Posthumous writings (1979) 9–46.

[9] R. M. Nosofsky, The generalized context model: An exemplar model of classification, Formal approaches in categorization (2011) 18–39.
[10] J. Leech, X—the mereology of representation, in: Proceedings of the Aristotelian Society, volume 116, Oxford University Press, 2016, pp. 205–228.
[11] K. Fine, The study of ontology, Noûs 25 (1991) 263–294.
[12] K. Fine, Things and their parts, Midwest Studies in Philosophy 23 (1999) 61–74.
[13] K. Fine, Towards a theory of part, The Journal of Philosophy 107 (2010) 559–589.
[14] D. Porello, O. Kutz, G. Righetti, N. Troquard, P. Galliani, C. Masolo, A toothful of concepts: Towards a theory of weighted concept combination, in: M. Simkus, G. E. Weddell (Eds.), Proc. of the 32nd International Workshop on Description Logics, volume 2373 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2019.
[15] G. Sacco, L. Bozzato, O. Kutz, Representing mereological relations in weighted description logics, in: T. P. Sales, M. M. Hedblom, al. (Eds.), Proc. of the Joint Ontology Workshops 2022 Episode VIII: The Svear Sommar of Ontology, volume 3249 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022.
[16] M. Richard, Meanings as species, Oxford University Press, 2019.
[17] J. N. Otte, J. Beverley, A. Ruttenberg, Bfo: Basic formal ontology, Applied ontology 17 (2022) 17–43.
[18] H. Cappelen, Fixing language: An essay on conceptual engineering, Oxford University Press, 2018.
[19] C. Masolo, L. Vieu, R. Ferrario, S. Borgo, D. Porello, Pluralities, collectives, and composites, in: Formal Ontology in Information Systems, IOS Press, 2020, pp. 186–200.
[20] R. Ferrario, C. Masolo, D. Porello, et al., Organisations and variable embodiments., in: FOIS, 2018, pp. 127–140.
[21] R. Baratella, Processes and events as rigid embodiments, Synthese 202 (2023) 181.
[22] N. Guarino, G. Guizzardi, Processes as variable embodiments, Synthese 203 (2024) 104.
[23] T. Brouwer, R. Ferrario, D. Porello, Hybrid collective intentionality, Synthese 199 (2021) 3367–3403.
[24] G. Guizzardi, J. P. A. Almeida, N. Guarino, V. A. de Carvalho, Towards an ontological analysis of powertypes, in: O. Papini, S. Benferhat, L. Garcia, M. Mugnier, E. L. Fermé, T. Meyer, R. Wassermann, T. Hahmann, K. Baclawski, A. Krisnadhi, P. Klinov, S. Borgo, O. Kutz, D. Porello (Eds.), Proceedings of the Joint Ontology Workshops 2015 Episode 1: The Argentine Winter of Ontology co-located with the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015), Buenos Aires, Argentina, July 25-27, 2015, volume 1517 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2015. URL: https://ceur-ws.org/Vol-1517/JOWO-15_FOfAI_paper_7.pdf.
[25] G. Uzquiano, Groups: Toward a theory of plural embodiment, Journal of Philosophy 115 (2018) 423–452. doi:10.5840/jphil2018115825.
[26] F. Moltmann, Variable objects and truthmaking, in: M. Dumitru (Ed.), Metaphysics, Meaning, and Modality. Themes from Kit Fine, Oxford University Press, 2020.
[27] R. Confalonieri, O. Kutz, P. Galliani, R. P. Nyssen, D. Porello, M. Schorlemmer, N. Troquard, Coherence, similarity, and concept generalisation, in: Proceedings of the 30th International Workshop on Description Logics, volume 1879, CEUR-WS. org, 2017.
[28] S. A. Sloman, B. C. Love, W.-K. Ahn, Feature centrality and conceptual coherence, Cognitive

Science 22 (1998) 189–228.
[29] A. Varzi, Mereology, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Spring 2019 ed., Metaphysics Research Lab, Stanford University, 2019.
[30] D. N. Osherson, E. E. Smith, On the adequacy of prototype theory as a theory of concepts, Cognition 9 (1981) 35–58.
[31] J. Fodor, E. Lepore, The red herring and the pet fish: Why concepts still can't be prototypes, Cognition 58 (1996) 253–270.
[32] E. J. Wisniewski, When concepts combine, Psychonomic bulletin & review 4 (1997) 167–183.
[33] R. Jackendoff, English noun-noun compounds in conceptual semantics, The semantics of compounding (2016) 15–37.
[34] G. Righetti, P. Galliani, C. Masolo, Concept Combination in Weighted DL, in: European Conference on Logics in Artificial Intelligence, Springer, 2023, pp. 385–401.