

Righetti G, Porello D, Troquard N, Kutz O, Hedblom M, Galliani P. To appear in: Neuhaus F, Brodaric B, eds. Proceedings of the 12th Int'l Conf. (FOIS 2021), *Frontiers in Artificial Intelligence and Applications* by IOS Press.

Asymmetric Hybrids: Dialogues for Computational Concept Combination

Guendalina RIGHETTI^a, Daniele PORELLO^b, Nicolas TROQUARD^a,
Oliver KUTZ^a, Maria M. HEDBLOM^c, Pietro GALLIANI^a

^a *Free University of Bozen-Bolzano, Italy*

^b *University of Genova, Italy*

^c *University of Bremen, Germany*

Abstract. When people combine concepts these are often characterised as “hybrid”, “impossible”, or “humorous”. However, when simply considering them in terms of extensional logic, the novel concepts understood as a conjunctive concept will often lack meaning having an empty extension (consider “a tooth that is a chair”, “a pet flower”, etc.). Still, people use different strategies to produce new non-empty concepts: additive or integrative combination of features, alignment of features, instantiation, etc. All these strategies involve the ability to deal with conflicting attributes and the creation of new (combinations of) properties. We here consider in particular the case where a Head concept has superior ‘asymmetric’ control over steering the resulting concept combination (or hybridisation) with a Modifier concept. Specifically, we propose a dialogical approach to concept combination and discuss an implementation based on axiom weakening, which models the cognitive and logical mechanics of this asymmetric form of hybridisation.

Keywords. Concept Combination, Hybridisation, Axiom Weakening, Dialogues, Compositionality

1. Introduction

Meredith: This is like a haunted coffeehouse thing?

Michael: No. Dwight is confusing you. That - it's, it's more of a disco.

Andy: It's like a haunted disco.

Michael: ... with coffee but without the haunted.

Phyllis: It's a combo dance house coffee bar.

Michael: It's a daytime disco on the ground floor of an industrial office building.

Erin: It's a cafe disco.

Michael: Exactly.

Kevin: So, like, a disco cafe?

Michael: Wha - No. No. Not even close.

The Office, Season 5, “Cafe Disco” [1]

The scene above demonstrates an interesting phenomenon. Namely, that concepts can be interpreted in different ways based on the weights of their attributes. The differentiation between a ‘disco cafe’ and a ‘cafe disco’ is determined by which of the two concepts has the more prominent role in the compound. While this difference might be quite intuitive for a native English speaker, it is a non-trivial problem to construct and explain in a formal setting. There exist different

views of what concepts are and how they should be represented. The logic-based view aims to represent concepts in term of definitions or, more precisely, as sets of individually necessary and jointly sufficient conditions [2]. In this setting, the combination of two or several concepts is commonly understood in terms of set theoretic operations. This view presents advantages for classic knowledge representation, mostly because it offers a compositional and well-understood semantics that is in line with mainstream reasoning systems. Unfortunately, empirical evidence in psychology and cognitive science has shown that many concepts lack precise definitions, being subject to various degrees of indeterminacy as well as to typicality effects [3]. Moreover, as we will further discuss, a number of cognitive phenomena linked to concept combination are difficult to reconcile with a straightforward modelling of concepts using Boolean extensional logic [4].

This paper will analyse the case of “incompatible” combinations, based on the empirical research on *impossible combinations* [5, 6, 7] and hybridisation [8], focusing on *asymmetric combinations*. Impossible conceptual combination (e.g. the combination of Fish and Vehicle, of Furniture and Fruit,...) has been studied in the context of experimental psychology to investigate the flexibility and adaptability of concept meaning. If we look at concepts simply from an extensional point of view, when combining concepts without obvious similarities or shared features, the intersection will often be empty. Still, people use different strategies to produce creative non-empty concepts: alignment of features, instantiation, features emergence etc. These strategies involve the ability to deal with conflicting attributes and the creation of new properties: simply put, a sort of game of meaning negotiation.

In order to elucidate and model the cognitive and logical mechanics in this kind of asymmetric concept combination, we here propose a computational framework based on three essential ingredients:

- (1) a computational model of concept combination taking into account cognitive aspects [9];
- (2) formal approaches based on axiom weakening [10];
- (3) an agent-based dialogical implementation¹ combining (1) and (2) to simulate meaning negotiation and construction in the asymmetric combination as it is found in the literature on hybrid concepts [7, 8].

Our approach to hybrid concept combination is related to conceptual blending (e.g. [11]) and, in particular, the distinction between Head and Modifier is reminiscent of asymmetric amalgams [12]. Please note, however, that our approach does not rely on the identification of a shared structure between the two input spaces via anti-unification — i.e. it does not require the identification of a *generic space* to steer the combination process. In contrast, our focus is on the asymmetric roles of Head and Modifier in the combination process [6], the integration process, and the resulting hybrid aspects of the combined concept [7, 8].

A similar distinction between Head and Modifier concepts is used in [13], where the authors propose an algorithm for concept combination based on a

¹The notion of dialogue here is quite abstract and the choice of the term “dialogue” is justified by the intention to point to the literature in multiagent systems and dialogical logics.

non-monotonic description logic of typicality. In contrast to our approach, which relies on axiom weakening, the proposed algorithm discards all the axioms of the Modifier which are inconsistent with the Head concept. The only constraint is that the combined concept should not be trivial, i.e., should not include *all* the features associated with the Head concept. It is not obvious how such a procedure could account for the kind of impossible combination proposed here.

The agent-based dialogical concept combination proposed in this paper is taking concept combination approaches further particularly regarding the so far rather monolithic debugging techniques for inconsistent blends, see [14].

More in line with our work is [15], which employs epistemic logic to negotiate the debugging of aligned ontologies. In terms of cognitive heuristics, the work proposed here goes beyond the plain distinction between Head and Modifier, and instead presents a model of steering the dynamics of cognitive concept combination as suggested by the results of Hampton [6].

2. Forms of Concept Combination: Hybridity and Impossibility

Knowledge Representation (KR) systems are usually characterised by their compositional behaviour. Compositionality is the principle according to in which any complex concept or expression is understood as a function of the parts it is composed by, plus a set of syntactic operations to combine them. This perspective became a cornerstone of classical logic and moved from there to be also a paradigm in description logic. In this setting, where concepts are essentially considered in terms of sets, the combination of two (or several) concepts is mostly understood in terms of set theoretic operations. The compound concept “Tool Weapon” would be understood as the intersection of the set of tools and the set of weapons (the component concepts). Compositionality is sometimes used to explain, at least in part, the ease and prolific ability by which humans create and understand new and meaningful phrases, arguably, part of its theoretical strength. In KR, in particular, it offers the advantage of having a clear and well understood semantics. Related to compositionality, one beneficial feature of many KR systems is attribute inheritance. Namely, for each class A in an ontology, the instances of sub-classes $B \sqsubseteq A$ would inherit all the attributes from the super-class. For combined concepts this would mean that what lies in the intersection of two concepts would inherit all the features normally associated to any conjunct (see [16] for a recent simulation study on inheritance illustrating the complexity of the issues involved).

The process of concept combination has been extensively studied in the field of cognitive science and experimental psychology. This led to several distinct accounts of concept combination, diverging widely from what is expressible in terms of intersections of sets [4, 8, 17, 18].

Hybridity. For instance, it is possible to distinguish between different kinds of combinations depending on whether we consider adjective-noun combinations or noun-noun combinations. Although, in simpler logical modellings, they are often treated in the same way, it is at the same time normally assumed that noun-noun combinations involve much more semantic change in the compound

concept [18]. Looking at noun-noun combinations in English, two parts can be distinguished, the Head and the Modifier, depending on the syntactic construction of the compound (this has been extensively studied in Linguistics, see e.g. [19]). Considering again “Tool Weapon”, the noun “weapon” would play here the role of the Head, whereas “tool” would be the Modifier. As the names suggest, the Head provides the base category of the combined concept, whilst the Modifier alters the attributes of the Head. This means that humans interpret “Weapon Tool” (e.g. a certain repair tool for the Avtomat Kalashnikova) significantly different from a “Tool Weapon” (e.g. James Bond’s typical screwdriver-shaped flame thrower). Clearly, any formal system employing compositional and commutative conjunction for such purposes would not be able to distinguish the two cases. Accounting for the difference in attribute inheritance is an important but logically challenging problem.

According to Wisniewski [8], there exist at least three ways to interpret noun-noun combinations: 1) The first is the *relation-linking* interpretation, where some kind of relation between the components is highlighted (using Wisniewski’s example [8, p. 168] a *robin snake* is a snake that eats robin). 2) The second is the *property* interpretation, where one or more properties of the Modifier noun apply to the Head concept (a *robin snake* is a snake with a red under-belly [8, p. 169]). 3) The third is called *hybridisation*, where the result of the combination corresponds essentially to a ‘mesh-up’ or ‘blend’ of both components.

We focus here on the third kind of combination interpretation, *Hybridisation*, and give a formal definition and computational account of it. In [8], the author refers to this last kind as a “combination of the two constituents [...] or a conjunction of the constituents” (p. 169). Conceptually, this corresponds to the combinations analysed in [4]. In [4]’s experiments, people were asked to interpret noun-noun combinations expressed with a that-clause (e.g. a Tool Weapon is expressed as a “Weapon that is also a Tool”). This was done to encourage people to think of the combination in conjunctive terms [6]. Hampton’s experiments [4] are of particular interest because he analysed the combination of ordinary concepts in terms of a logical interpretation. He found that, although it was possible to identify predictable patterns in the relation between compound and components, people are often not consistent with the rules of set theory.²

Impossibility. In a series of experiments [5, 6, 7], Hampton asked people to combine concepts that usually would not be combined, leading to impossible, or at least imaginary, objects. The aim of the investigation was to analyse the underlying principles for concept combination in a setting free from bias and prior knowledge of the concepts involved in order to study how adaptable and flexible concepts are. In [6], people were presented with a list of concept pairs (e.g. “Furniture” and “Fruit”, “Vehicle” and “Fish”, “Bird” and “Kitchen Utensil”, etc.). Then, they were asked to imagine the objects resulting from the combination (e.g. “a Vehicle which is also a Fish”), and to describe, or draw, the attributes they would expect such an object to have. If analysed in terms of set-theoretic operations, the intersection of the concepts involved would be the empty set, and the set of axioms associated to both component concepts would likely

²For a formal analysis of these Hampton phenomena in weighted logics, we refer to [20, 21].

be inconsistent. Still, subjects showed a great variability of strategies to solve incompatible combinations. In this context, the process of properties *alignment* is particularly interesting. In order to select the ‘right’ properties for the compounds, people try to align properties and functions of the two component concepts. This was also noticed by [8], which, particularly in relation to hybridisation, proposed a comparison and alignment model where the Head and Modifier concepts are first aligned, so that the properties of one concept are put in correspondence with the properties of the other, and then are compared to find connections. This process can then go in quite different directions. First, subjects can find some *alignable differences* [17]: once the two representation are aligned, it is possible to find differences wrt. some of the aligned properties (or dimensions). In [6, table 9], for instance, when people analysed the concept “a Fruit which is a kind of Furniture”, they tended to align the skin of the fruit with the fabric of a sofa. Or, when asked about “a Vehicle which is a kind of Fish”, they reasoned that both could move, or made the analogical mapping between fuel and food. This alignment process corresponds to either finding commonalities in the differences that helps to understand which properties of the Modifier are integrated into the Head concept, or to find strong incompatibilities between the concepts. For instance, subjects could notice that, whereas a piece of furniture is made to last, fruit is perishable. Likewise, while a vehicle is normally controlled, a fish is likely to be ‘self-motivated’. In these cases, people react to incompatibilities producing new, or *emergent* attributes [6]. Emergent attributes are defined as features that were not listed as true neither of the Head concept, nor of the Modifier, but that appear, or emerge, for the combined concept.

Another strategy observed by Hampton in his experiments is the process of *instantiation*: when asked to combine two super-ordinate categories (such as Vehicle and Fish or Fruit and Furniture), people would find it easier to come up with a solution “instantiating” one of them to a more basic and well-known category (combining e.g. Boat and Fish, or Banana and Furniture). Moreover, some of the categories were more likely to be instantiated than others (for instance, the category of Vehicle was almost always instantiated, whereas this very rarely happened to the one of Fish). The phenomenon of instantiation does not have an obvious explanation, since a more general concept would pose fewer constraints on the combination. Likely this is due to the fact that basic categories are easier to be imagined and more familiar to subjects, with more concrete properties to combine [7].

Aside from all these strategies, it is possible to observe some asymmetries between the Head and Modifier even in the case of impossible combinations. Hampton [7] shows that the solutions elaborated from the subjects usually bear more similarity to the Head noun than to the Modifier. Also in the case of impossible combinations, subjects keep the Head noun as a base to be modified by means of the Modifier.

3. Concept Refinements and Axiom Weakening

We here consider ontologies as sets of formulae in an appropriate logical language with the purpose of describing a particular domain of interest. The choice

of a specific logic used is not crucial to our general approach, but we here employ the well-known description logic \mathcal{ALC} both for illustrating examples and in our experimental prototype implementation. Syntactically, the \mathcal{ALC} concept language is based on two disjoint sets N_C and N_R of *concept names* and *role names*, respectively. The set of \mathcal{ALC} concepts is generated by the grammar $C ::= A \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \forall R.C \mid \exists R.C$, where $A \in N_C$ and $R \in N_R$. We denote by $\mathcal{L}(N_C, N_R)$ the set of \mathcal{ALC} concepts built over N_C and N_R . Axioms are *general concept inclusions* (GCIs) or *individual assertions*. GCIs are of the form $C \sqsubseteq D$, where C and D are concepts. Finite sets of GCIs are called *TBoxes* which constitute the general, terminological knowledge of the ontology. Finally, individual assertions are of the form $C(a)$ where C is a concept and a an individual name; such statements constitute *ABoxes*. We assume standard DL syntax and semantics [22]. By $\text{sub}(C)$ we denote the set of subconcepts of C , defined recursively as usual. The set of subconcepts in an axiom $C \sqsubseteq D$ is $\text{sub}(C \sqsubseteq D) = \text{sub}(C) \cup \text{sub}(D)$; also, $\text{sub}(C(a)) = \text{sub}(C)$, and $\text{sub}(O)$ denotes the set of all the subconcepts of all the axioms in O .

Refinement operators [10, 23] have been used for ontology aggregation in [24]. Here, they will be instrumental for concept combination. Given the quasi-ordered set $\langle \mathcal{L}(N_C, N_R), \sqsubseteq \rangle$, a generalisation refinement operator satisfies $\gamma_O(C) \subseteq \{C' \in \mathcal{L}(N_C, N_R) \mid C \sqsubseteq_O C'\}$. A specialisation refinement operator satisfies $\rho_O(C) \subseteq \{C' \in \mathcal{L}(N_C, N_R) \mid C' \sqsubseteq_O C\}$. Generalisation refinement operators take a concept C as input and return a set of descriptions that are more general than C by taking an ontology O into account. A specialisation operator, instead, returns a set of more specific descriptions. We can now define the notion of axiom weakening. The set of weakenings of an axiom wrt. a reference ontology O is defined as follows. Given an inclusion axiom $C \sqsubseteq D$ of O , the set of (least) *weakenings* of $C \sqsubseteq D$ wrt. O , denoted by $g_O(C \sqsubseteq D)$, is the set of all axioms $C' \sqsubseteq D'$ such that $C' \in \rho_O(C)$ and $D' = D$ or $C' = C$ and $D' \in \gamma_O(D)$. Given an individual assertion $C(a)$ of O , the set of (least) *weakenings* of $C(a)$, denoted $g_O(C(a))$, is the set of all axioms $C'(a)$ such that $C' \in \gamma_O(C)$.

The proposal laid out in this paper can make use of any refinement operator. When specific refinement operators are needed, as e.g. in examples and the implementation, we will be using the refinement operators from [10].

4. Dialogues for Concept Combination

We assume two agents, h and m , are interacting, trying to build a consistent compromise ontology R describing a concept. Each agent has an ontology associated, O_h and O_m , describing their initial version of R . Moreover, they each have a preference orderings \prec_h and \prec_m over the axioms of their respective ontology. The preferences represent the importance of the axiom for describing the concept.

In the dialogue, the agents are proposing in turn axioms coming from their ontology to be added to the ontology under construction R , weakening them when necessary. This is inspired by an approach from [24] for ontology aggregation. When the axioms proposed by the agents turn out to render the devised ontology inconsistent, the axiom weakening procedure is called to solve that inconsistency. A dialogue protocol is described informally in Algorithm 1.

Algorithm 1 Combination($O_{init}, O_h, \prec_h, O_m, \prec_m, prob_h$)

```
1:  $R \leftarrow O_{init}$ 
2: TreatedAxioms  $\leftarrow \emptyset$ 
3: FinishedAgents  $\leftarrow \emptyset$ 
4:  $ag \leftarrow \mathbf{RandomPickOneAgent}(prob_h, \{h,m\})$ 
5: while FinishedAgents  $\neq \{h,m\}$  do
6:   if  $\forall$  axiom  $\varphi \in O_{ag}$ :  $\varphi \in$  TreatedAxioms or  $R$  entails  $\varphi$  then
7:     FinishedAgents  $\leftarrow$  FinishedAgents  $\cup \{ag\}$ 
8:   else
9:      $\varphi \leftarrow \mathbf{FavoriteNextAxiom}(\prec_{ag}, O_{ag}, \text{TreatedAxioms}, R) \triangleright$  Favorite axiom of agent  $ag$ ,
not in TreatedAxioms, and not entailed by  $R$ 
10:    TreatedAxioms  $\leftarrow$  TreatedAxioms  $\cup \{\varphi\}$ 
11:    while  $R \cup \{\varphi\}$  is inconsistent do
12:       $\varphi \leftarrow \mathbf{RandomPickOneWeakening}(g_R(\varphi))$ 
13:       $R \leftarrow R \cup \{\varphi\}$ 
14:       $ag \leftarrow \mathbf{RandomPickOneAgent}(prob_h, \{h,m\})$ 
15: return  $R$ 
```

The algorithm takes a few parameters: an initial ontology O_{init} , an ontology O_i for each agent $i \in \{h, m\}$, a (strict) preference order \prec_i over the set of axioms O_i for each agent i , and a probability $prob_h$ of agent h to take a turn.

The algorithm iteratively builds an ontology R for the combined concept, initialised with O_{init} . The two agents take turns randomly following the probability distribution $(prob_h, 1 - prob_h)$. When it is her turn, agent i will choose her preferred axiom φ in O_i according to \prec_i , and not already entailed by the combination R . As long as φ can not be added to R without causing an inconsistency, it is replaced by one of its weakenings in $g_R(\varphi)$, i.e. a weakening of φ wrt. the current combination R . As soon as φ can be added to R without causing an inconsistency, the combination R is augmented with φ . When all the axioms of an agent have been considered or are already entailed by the current combination, this agent is finished. This iterative process continues until all agents have finished. At the end, the combination R is returned.

In the experiments, we also consider a **bounded variant** of this algorithm, where a maximum number max_turns of turns is added as a parameter, and where the **while**-loop exits after at most max_turns iterations. We now state a few formal properties of these two algorithms. It is easy to see that the returned ontology R is always consistent.

Proposition 1 (Consistency). *If O_{init} is consistent and Algorithm 1 (or its bounded variant) returns R , then R is consistent as well.*

Also, as a corollary of [25, Th. 2], we can show that the algorithm almost always terminates when using the refinement operators of [10].

Proposition 2 (Termination). *If $prob_h \notin \{0, 1\}$, then Algorithm 1 (and its bounded variant) terminates with probability 1.*

Moreover, we can formulate a sufficient condition for the combination R to be maximal in the following sense:

Proposition 3 (Maximality). *Let R be an ontology returned by Algorithm 1 (or by its bounded variant with $\text{max_turns} \geq |O_h \cup O_m|$) and let φ be an axiom in $O_h \cup O_m$. Then $R \cup \{\varphi\}$ is either inconsistent or equivalent to R .*

We can readily use the algorithm for asymmetric concept combination of a Head concept H described by an ontology O_h with a Modifier concept M described by an ontology O_m . The result is an ontology intended to describe the target concept MH, which is the asymmetric hybridisation of the Head concept H with the Modifier concept M .

Probability prob_h . The asymmetry of the hybridisation must be enforced by a suitable weight given to the Head and to the Modifier, and an appropriate probability to take turns in the dialogue. In the asymmetric case, the Head agent h will be given a greater weight than the Modifier agent, agent m ; it will have relatively more opportunities to insert its information into the hybridisation. One then needs to translate these weights into a probability for the Head and Modifier agents to take turns. In practice, one needs to consider the granularity of the information contained in both ontologies. At the time of the combination, an agent with a high granularity ontology, i.e. many detailed axioms, is likely to require more turns to add its information into the blend. To this end, we need a way to quantify the amount of information in an ontology. We take into account the logical axioms in a given ontology and also the inferred class hierarchy. Given an ontology O , we define the set of axioms in the *inferred class hierarchy* as follows: $\text{inf}(O) = \{C \sqsubseteq D \mid C, D \in N_C \cap \text{sub}(O), C \sqsubseteq_O D\}$, where $\text{sub}(O)$ is the set of subconcepts appearing in O . This provides us with a useful measure to evaluate the ‘amount’ of information contained in an ontology O , namely defined as the quantity $|O \cup \text{inf}(O)|$. Combining two ontologies O_h and O_m with weights w_h and w_m , agent h will play $w_h \cdot |O_h \cup \text{inf}(O_h)| = f_h$ turns when agent m will play $w_m \cdot |O_m \cup \text{inf}(O_m)| = f_m$ turns. The probability passed as a parameter to the algorithm is then $\text{prob}_h = \frac{f_h}{f_h + f_m}$.

Preferences \prec_h and \prec_m . The preferences of the agents represent the importance of their axioms in expressing certain features of the concept at issue, for the purpose of the specific combination. We take them here as given inputs, provided by the agents, and they partially determine the ‘direction’ of the combination.

Initial Ontology O_{init} . The choice of the initial ontology is motivated by the goal of combining two concepts. So, when combining H and M , the initial ontology O_{init} will contain the two axioms: $\text{MH} \sqsubseteq H$ and $\text{MH} \sqsubseteq M$, where MH is the target hybrid concept. This is enough to bootstrap the formalisation of the requirement that the hybrid concept is an H that is also an M . Moreover, to avoid the trivial result where the resulting ontology is consistent but the hybrid concept is unsatisfiable we must also add an axiom $\text{MH}(a)$ for a fresh individual name a , making sure that some MH’s do exist.

5. Computational Simulations of Impossible Combinations: The FishVehicle

We illustrate how the two versions of our algorithm work in the case of an *impossible combination*. Namely, we simulate the combination of the concepts

Fish and Vehicle as it is described in Hampton’s experiments [6] by means of our dialogue implementations.³ We start with a consistent initial ontology, which will guide our weakening procedure. To provide some of the high level ontological distinctions needed for representing the input concepts, we include in our initial ontology an excerpt of the taxonomy of DOLCE, formulated in \mathcal{ALC} . DOLCE (i.e. Descriptive Ontology for Linguistic and Cognitive Engineering) is a cognitively oriented Foundational Ontology, which enables a fine-grained analysis of concepts. As such, it provides a number of basic ontological and cognitive distinctions required to represent and confront the common sense concepts at issue. The most general categories of DOLCE include *Perdurants*, *Endurants*, *Abstracts* and *Qualities* (see [26]). We are mostly interested in *Endurants*, and in particular in the distinctions related to *Physical Endurants*, which branch into *Amount of Matter*, *Features* and *Physical Objects*. Fish and Vehicle are indeed *Physical Objects*. Food and Fuel, for instance, which appear in our input ontologies in axioms such as “fish eats food”, “vehicles need fuel”, are here included in *Amount of Matter*. To state that “fish has slippery surface”, we include *Slippery Surface* in *Feature*. We also use *Qualities* from DOLCE, to represent, for instance, the shape, the colour, the spatial locations of Fish and Vehicles. DOLCE *Abstract* and *Space Region*, i.e. values for spatial location qualities, are required to classify places such as *Air*, *Ground*, and *Water*, which we use in axioms such as “fish swims in water”. Moreover, *Abstracts* allow for introducing conceptual spaces, to model the values of the individual qualities (e.g. colour, shape, weights) of fish and vehicles.

DOLCE serves to provide, via alignment, the ontological distinctions needed to represent and reason about the possible incompatibilities between the two concepts to be combined. Aside from DOLCE, the initial ontology contains also two additional axioms, which directly relate to the concept we want to build (as described in the section above): $\text{FishVehicle} \sqsubseteq \text{Fish}$ and $\text{FishVehicle} \sqsubseteq \text{Vehicle}$. To ensure the concept *FishVehicle* is not empty, we also add an instance of the concept: $\text{FishVehicle}(\text{Wanda})$.

We need then two ontologies, which represent the concepts of Fish and Vehicle respectively, before the combination can be started. These can be seen as micro-theories, little domain ontologies, modelling the two concepts involved. In our setting, they are associated with two different agents, and each axiom corresponds to a possible move in the dialogue. The specific content of the ontology of Vehicle (resp. Fish) is partly reverse-engineered using the information contained in Hampton’s experiments described above (e.g. a fish *eats food*, is *autonomous* and can *swim*; a vehicle *needs fuel*, is *controlled*, and can *move*, etc.). Additional information (e.g. body parts, vehicle component, etc.) is inspired by the Leuven Database ([27]), which collects psychological, commonsense data on a feature generation task of 15 concepts (including fish and vehicle).

In order to make the two input ontologies of fish and vehicle interoperable, they are aligned to the common upper level provided by DOLCE (to achieve that, the classes of the domain ontologies are subsumed under the pertinent categories of DOLCE). DOLCE is particularly well-suited to account for some of the distinctions observed by Hampton in his experiments. E.g., it can

³See <https://bitbucket.org/troquard/ontologyutils/src/master/>.

capture the distinction between the agency of the Fish (modelled as a subclass of `AgentivePhysicalObject`) contrasting the non-agency of the Vehicle (i.e. a `NonAgentivePhysicalObject`), asserted indirectly through the possibility to control it.

The goal of the procedure is to build the concept of `FishVehicle`, which should share both features of the concept of `Fish` and features of the concept of `Vehicle`. When the algorithm starts, at each turn the agents will try to add their favourite axioms to the initial ontology. If the axiom cannot be added without causing inconsistency, it is weakened by the procedure.

We have two agents: `agent_h` represents the Head concept (in this case, `Vehicle`) and `agent_m` represents the Modifier (`Fish`). To implement the asymmetry of the combination, we do not distribute the turns equally between the two agents. At each round, the weight for `agent_h` to play is higher than the one for `agent_m`. Having the possibility to play its favourite axioms sooner, `agent_h` is more likely to add less weakened information to the initial ontology.

The last important aspect to consider is the preference order that we put on the axioms. As already mentioned, if an axiom is preferred and added sooner to the initial ontology, it will be less likely that it causes an inconsistency and is weakened. We consider three different preference orders. Firstly, we consider an order which enforces the strength of the ontological distinctions, i.e. the link between the ontologies of `Vehicle` (resp. `Fish`) with DOLCE. Secondly, we consider the opposite situation, i.e. where the specific axioms of `Vehicle` (resp. `Fish`) were preferred. Finally, for the domain ontologies we here follow a preference order aiming at replicating the process of instantiation as described by Hampton [7] and outlined in Section 2. In this case, `agent_h` preferred all the axioms containing information related to `Car`. In contrast, we left the `Fish` order as a random order.

The unbounded version of our algorithm ends, in any case, when both agents have done all their possible moves, and we obtain a maximally informative ontology R about `FishVehicle`. The bounded version ends after the selected number of moves, returning a consistent ontology R for `FishVehicle`.

6. Evaluating Asymmetric Concept Hybridisations

Distinguishing between good and bad hybridisations is neither a straightforward task nor an entirely new one. Research on it appears both in computational creativity when evaluating machine-generated combinations [28], and in cognitive psychology, where identifying human strategies and cognitive heuristics are the focus [5]. In his experiments, Hampton asked two independent judges to evaluate on a 1 to 10 scale the ‘success’ of the responses given by the subjects to an impossible combinations task [7]. However, this does not tell us much about how to effectively evaluate computationally the outcome of an impossible combination.

Lacking this kind of information, what we can measure is what kind of effects our strategies show on the output of the algorithm. We therefore consider next two parameters, namely the *asymmetry* of the combination and its *hybridity*.

Asymmetry. The asymmetry of the combination represents the relative effect of the Head concept (e.g., `Vehicle`) and the Modifier concept (e.g., `Fish`) in the result ontology R . To measure this asymmetry we exploit a ratio of preserved

information from [24]: a ratio with values between 0 and 1, measuring how much information from an ontology O_i is present in another ontology R . The *ratio of preserved information* from O_i in R is

$$\text{rpi}(O_i, R) = \frac{|\{\varphi \in O_i \cup \text{inf}(O_i) \mid \varphi \text{ is entailed by } R\}|}{|O_i \cup \text{inf}(O_i)|}.$$

To measure the asymmetry of the combination R , we then use the difference between the ratios of preserved information, specifically, the difference between the ratio of preserved information from O_h (the ontology of the Head) in R and the ratio of preserved information from O_m (the ontology of the Modifier) in R . The measure of *asymmetry* of R is thus the number: $\text{rpi}(O_h, R) - \text{rpi}(O_m, R)$.

The larger the absolute value of the asymmetry, the more one of the two concepts dominates the result ontology R . Further, a positive value indicates that the Head does dominate the Modifier in the result ontology R , whilst a negative weight instead means that the Modifier does.

Hybridity. Intuitively, a hybrid of Fish and Vehicle should share some of its features with Fish and some others with Vehicle. To formally capture this intuition, we introduce here the notion of *hybrid description*. A hybrid description of Fish and Vehicle is something like a ‘shark that needs fuel’ or ‘is made of metal and has fins’. More generally, we define the set of hybrid descriptions for a Head concept H and a Modifier M as the set of conjunctions $C_h \sqcap C_m$ such that C_h (resp. C_m)

- (1) is an ascendant or descendant of H within O_h (resp. of M within O_m), and
- (2) is **not** a (sub-)concept appearing in the ontology O_m describing M (resp. O_h describing H).

Formally, let O_h be the ontology defining the head concept H (e.g., Vehicle) and O_m the ontology defining the modifier concept M (e.g., Fish). Let MH be the target hybrid concept (e.g., FishVehicle), defined through the result ontology R . For a concept D within ontology O , define the set of ascendants and descendants (the ‘lineage’ of D in O) as $\text{lin}(O, D) = \{C \in \text{sub}(O) \mid C \sqsubseteq_O D \text{ or } D \sqsubseteq_O C\}$. Then, to measure the hybridity of the concept MH, we count the number of times in which $\text{MH} \sqsubseteq C_h \sqcap C_m$, for $C_h \in \text{lin}(O_h, H) \setminus \text{sub}(O_m)$ and $C_m \in \text{lin}(O_m, M) \setminus \text{sub}(O_h)$, over the total number of hybrid descriptions. Notice, crucially, that we exclude the ascendants and descendants which are shared by the two concepts. Formally, with $\Lambda_H = \text{lin}(O_h, H) \setminus \text{sub}(O_m)$ and $\Lambda_M = \text{lin}(O_m, M) \setminus \text{sub}(O_h)$, the measure of *hybridity* of R is the number

$$\frac{|\{(C_m, C_h) \mid C_h \in \Lambda_H, C_m \in \Lambda_M, \text{ and } \text{MH} \sqsubseteq C_h \sqcap C_m \text{ is entailed by } R\}|}{|\Lambda_H| \times |\Lambda_M|}.$$

We evaluated the output of our algorithms on the preference orders and parameters introduced above. We present our findings next.




Figure 1. Asymmetry and hybridity values for the two preference orders, with varying weight of the Head. Bounded $\sim 50\%$ refers to the bounded variant of Algorithm 1.

Controlling the Effects of DOLCE. In order to enforce constraints from DOLCE, agents prefer all the axioms that bridge the classes of DOLCE within our O_{init} and the classes pertaining more strictly to the ontology of Vehicle (resp. Fish). For instance, agents prefer axioms such as $\text{Artefact} \sqsubseteq \text{NonAgentivePhysicalObject}$. Agents also prefer the constraints imposed on the roles (Domain and Range). We expected that enforcing the link with DOLCE, and emphasising hard ontological distinctions, would have had a negative effect on the hybridity value.

The opposite strategy enforces the specific axioms for Vehicle (resp. Fish). The preference gave priority to all the axioms containing the concept Vehicle (resp. Fish) on the left or right side of the subsumption relation (e.g. $\text{Vehicle} \sqsubseteq \text{hasComponent.VehiclePart}$ or $\text{Car} \sqsubseteq \text{Vehicle}$). Enforcing first the specific information for the concepts to be combined, and only establishing the link with DOLCE at the end of the procedure, was expected to enhance the number of hybridisations, and thus increase the hybridity value.

At the same time, increasing the weight associated to a specific concept, i.e. increasing that agent’s probability to play, was expected to increase the asymmetry between the two concepts.

This was confirmed by our experiments, as can be seen in Figure 1. Increasing the weight of the first player tends to increase the asymmetry value. This is particularly evident in the bounded cases. The value of hybridity was affected by the preference order, and prioritising the link with DOLCE critically decreases the number of hybridisations. The weights of the two players do not affect the hybridity values. By contrast, by setting boundaries, hybridity decreases.

Simulating Hampton’s Findings. As described in Section 2, Hampton observed the tendency of people to instantiate (or specialise) general classes in order to find a solution to an impossible combination task. When instantiating the concept Vehicle into, e.g., Car, the combined concept should show some of the distinctive features of the instantiated concept. The effectiveness of an instantiation strategy should then be evaluated on the capability of the combined concept to satisfy the specific features of the instantiated concept. Similarly to the methodology of competency question employed in knowledge engineering [29], we selected to this

	Priority DOLCE	Priority Vehicle/Fish	Instantiation
FV goes on ground	4%	2%	100%
FV has brake	0	0	100%
FV has motor	0	0	100%
FV has steering wheel	0	0	100%
FV has wheel	0	0	100%
FV does not go on air or water	16%	2%	100%
FV has not wing	18%	4%	100%

Table 1. Percentage of combinations satisfying the instantiation questions, over 50 runs.

end the questions described in Table 1, which correspond to the description of the concept `Car` within the ontology of `Vehicle`.

To replicate this phenomenon within our set-up, we first include an additional axiom in our initial ontology, enforcing the `FishVehicle` to be also a sub-concept of `Car`. Then, the agent `_h` prefers all the axioms containing information related to `Car` (i.e. the concept `Car` occurs on the right or left hand side of a subsumption). The effectiveness of this strategy is shown in Table 1⁴. Namely, in all of 50 runs, the `FishVehicle` showed *all* the features associated to `Car` (i.e. 100%). In our tests, the weight for agent `_h` is set to be 3, whereas the one for agent `_m` is set to be 1; a probability of about 0.75 for agent `_h` to take a turn. In the case of the unbounded procedure, the hybridity value was high, with an average of 0.7. The 50% bounded version, at about 0.36, cuts that value of hybridity to about half. Although one might have expected such an effect given the reduced opportunity to impose ‘hybrid’ information, the effect is here surprisingly strong.

Another phenomenon observed by Hampton [6] is the use of alignments. According to his findings, people tend to align the fact that fish eat food with the fact that vehicles need fuel; or, the fact that both have the capacity to move. To replicate this phenomenon, the following set of axioms was then added to the initial ontology: `Food` \equiv `Fuel` ; \exists `needs.Fuel` \equiv \exists `eats.Food` ; \exists `swimsIn.Water` \equiv \exists `goesOn.(Water` \sqcup `Air` \sqcup `Ground)` ; `Water` \equiv `Air` ; `Water` \equiv `Ground` ; `Ground` \equiv `Air` . The alignments were not, as was to be expected, consistent with the ontologies of the two players. Therefore, it was part of the weakening procedure to integrate them consistently.

We expected that introducing the alignments within our procedure would have had a positive effect on the hybridity value. This was, however, not observed within our dataset. Looking at the effects of the alignments, the main benefit observed was in terms of *feature emergence*. Introducing the alignments between `Fuel` and `Food` and between `Air`, `Ground` and `Water` produced in fact some mixed axioms, which were present neither in the ontology of `Fish`, nor in the ontology of `Vehicle`. Table 2 shows an example of this effect.

	Alignment	No alignment
FV eats fuel	21	0
FV swims on air or ground	12	0

Table 2. Number of emergent features over 50 runs, on a random order, unbounded procedure.

⁴We report here the values for the unbounded procedure, but the result is analogous for the bounded one.

7. Conclusions and Future Directions

We developed a dialogue-based algorithm for the computational generation of hybrid, sometimes considered ‘impossible’, combinations. Our method is inspired by the empirical research in psychology identifying human heuristics for combining concepts that lack any obvious similarities. To explore the dynamics involved in our dialogue games and to experimentally evaluate the human heuristics mentioned, we defined and implemented a number of measures including ratio of preserved information (*rpi*), hybridity, asymmetry of the combination, and impact of alignment with an upper ontology (here, DOLCE). In general, the unbounded dialogue game allows for the construction of ‘almost perfect conjunctions’ in the sense that preservation of information remains high whilst hybridity increases. This is a feature of interest more generally to ontology engineering. Further, the lack of high asymmetry in the unbounded combination can be traced back to the fact that, as Prop. 3 showed, we construct maximal combinations in this case. In contrast, the use of a bounded procedure permits to build highly asymmetrical combinations, arguably more in line with the distinction between Head and Modifier as described in cognitive psychology. As may be expected, this is obtained at the cost of a decrease in hybridity.

We also showed the flexibility of our algorithm in reproducing some of the phenomena observed in the cognitive psychology of impossible combinations, namely the use of alignments and instantiation. To simulate human concept combination in a subtler way, a more fine-grained protocol regarding evaluation of preferences, prioritisation strategies and resource-bounding should be investigated further.

References

- [1] The Office, Season 5; 2009. Episode: Cafe Disco. Warren Lieberstein and Halsted Sullivan (writers) and Randall Einhorn (director).
- [2] Murphy GL. The Big Book of Concepts. MIT press; 2002.
- [3] Rosch E. Principles of Categorization. In: Margolis E, Laurence S, editors. Concepts: Core Readings. vol. 189. MIT press; 1999. p. 189-206.
- [4] Hampton JA. Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*. 1987;15(1):55-71.
- [5] Gibbert M, Hampton JA, Estes Z, Mazursky D. The Curious Case of the Refrigerator-TV: Similarity and Hybridization. *Cognitive Science*. 2012;36:992-1018.
- [6] Hampton JA. Emergent Attributes in Combined Concepts. In: Ward TB, Smith SM, Viad J, editors. Creative thought: An investigation of conceptual structures and processes. Washington DC: APA Press; 1997. p. 83-110.
- [7] Hampton JA. Compositionality and Concepts. In: Hampton JA, Winter Y, editors. Compositionality and Concepts in Linguistics and Psychology. Cham: Springer International Publishing; 2017. p. 95-121.
- [8] Wisniewski EJ. When concepts combine. *Psychonomic Bulletin & Review*. 1997;4:167-83.

- [9] Confalonieri R, Kutz O. Blending under Deconstruction: The roles of logic, ontology, and cognition in computational concept invention. *Annals of Mathematics and Artificial Intelligence*. 2020;(88):479-516.
- [10] Troquard N, Confalonieri R, Galliani P, Peñaloza R, Porello D, Kutz O. Repairing Ontologies via Axiom Weakening. In: *Proc. of AAAI 2018*; 2018. p. 1981-8.
- [11] Ontañón S, Plaza E. Amalgams: A Formal Approach for Combining Multiple Case Solutions. In: Bichindaritz I, Montani S, editors. *Case-Based Reasoning. Research and Development*. Berlin, Heidelberg: Springer Berlin Heidelberg; 2010. p. 257-71.
- [12] Besold TR, Kühnberger KU, Plaza E. Towards a computational- and algorithmic-level account of concept blending using analogies and amalgams. *Connection Science*. 2017;29(4):387-413.
- [13] Lieto A, Pozzato GL. A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics. *J Exp Theor Artif Intell*. 2020;32(5):769-804.
- [14] Eppe M, Maclean E, Confalonieri R, Kutz O, Schorlemmer M, Plaza E, et al. A computational framework for conceptual blending. *Art Int*. 2018;256:105-29.
- [15] van den Berg L, Atencia M, Euzenat J. Agent ontology alignment repair through dynamic epistemic logic. In: *Proc. of AAMAS 2020*. ACM; 2020. p. 1422-30.
- [16] Thorn PD, Schurz G. How Category Selection Impacts Inference Reliability: Inheritance Inference From an Ecological Perspective. *Cognitive Science*. 2021;45(4):e12971.
- [17] Markmar AB, Gentner D. Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*. 1993;32:517-35.
- [18] Costello FJ, Keane MT. Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*. 2000;24:299-349.
- [19] Jackendoff R. English Noun-Noun Compounds in Conceptual Semantics. In: *The Semantics of Compounding*. Cambridge University Press; 2016. p. 15–37.
- [20] Righetti G, Porello D, Kutz O, Troquard N, Masolo C. Pink Panthers and Toothless Tigers: Three Problems in Classification. Manchester, Sept. 10–11: CEUR-WS, Vol-2483; 2019. *Proc. of the 5th Int. Workshop on Artificial Intelligence and Cognition*.
- [21] Porello D, Kutz O, Righetti G, Troquard N, Galliani P, Masolo C. A Toothful of Concepts: Towards a theory of weighted concept combination. CEUR-WS, Vol-2373; 2019. *Proceedings of the 32nd International Workshop on Description Logics*.
- [22] Baader F, Horrocks I, Lutz C, Sattler U. *An Introduction to Description Logic*. Cambridge University Press; 2017.
- [23] van der Laag PRJ, Nienhuys-Cheng SH. Completeness and properness of refinement operators in inductive logic programming. *The Journal of Logic Programming*. 1998;34(3):201-25.
- [24] Porello D, Troquard N, Penaloza R, Confalonieri R, Galliani P, Kutz O. Two Approaches to Ontology Aggregation Based on Axiom Weakening. In: *Proc. of 27th Int. Joint Conf. on Artificial Intelligence and 23rd European Conf. on Artificial Intelligence*; 2018. p. 1942-8.
- [25] Confalonieri R, Galliani O, Pietro Kutz, Porello D, Righetti G, Troquard N. Towards Even More Irresistible Axiom Weakening. CEUR-WS, Vol-2663; 2020. *Proc. of 33rd International Workshop on Description Logics*.
- [26] Masolo C, Borgo S, Gangemi A, Guarino N, Oltramari A. *WonderWeb Deliverable D18*. CNR; 2003.
- [27] De Deyne S, Verheyen S, Ameel E, Vanpaemel W, Dry MJ, Voorspoels W, et al. Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behavior Research Methods*. 2008;40(4):1030-48.

- [28] Veale T. From Conceptual Mash-ups to Badass Blends: A Robust Computational Model of Conceptual Blending. In: Computational Creativity. Springer; 2019. p. 71-89.
- [29] Grüninger M, Fox MS. Methodology for the Design and Evaluation of Ontologies. IJCAI-95, Montreal; 1995. Workshop on Basic Ontological Issues in Knowledge Sharing.