Asymmetric Hybrids: Dialogues for Computational Concept Combination (Extended Abstract)*

Guendalina Righetti[1] , **Daniele Porello**[2] , **Nicolas Troquard**[1] ,
Oliver Kutz[1] , **Maria Hedblom**[3] and **Pietro Galliani**[1]
[1]Free University of Bozen-Bolzano, Italy
[2]University of Genova, Italy
[3]Jönköping University, Sweden.
guendalina.righetti@unibz.it, daniele.porello@unige.it, nicolas.troquard@unibz.it, oliver.kutz@unibz.it,
mariamagdalena.hedblom@ju.se, pietro.galliani@unibz.it

Abstract

When considering two concepts in terms of extensional logic, their combination will often be trivial, returning an empty extension. Consider e.g. "a Fish Vehicle", i.e. "a Vehicle which is also a Fish". Still, people use sophisticated strategies to produce new, non-empty concepts. All these strategies involve the human ability to mend the conflicting attributes of the input concepts and to create new properties of the combination. We focus in particular on the case where a Head concept has superior 'asymmetric' control over steering the resulting combination (or hybridisation) with a Modifier concept. Specifically, we propose a dialogical model of the cognitive and logical mechanics of this asymmetric form of hybridisation. Its implementation is then evaluated using a combination of example ontologies.

1 Introduction

There exist different views of what concepts are and how they should be represented. The logic-based view aims to represent concepts as sets of individually necessary and jointly sufficient conditions [Murphy, 2002]. In this setting, the combination of concepts is commonly understood in terms of set-theoretic operations. This view presents advantages for classic Knowledge Representation (KR), mostly because it offers a compositional and well-understood semantics that is in line with mainstream reasoning systems. Unfortunately, a number of cognitive phenomena linked to concept combination are difficult to reconcile with a plain modelling of concepts using Boolean extensional logic [Hampton, 1987].

This paper summarises [Righetti *et al.*, 2021b], which analyses the case of "incompatible" combinations, based on empirical research on *impossible combinations* and hybridisation, with a focus on *asymmetric combinations*, accounting for the distinction between Head and Modifier concepts as studied in cognitive psychology [Gibbert *et al.*, 2012; Hampton, 1997; Hampton, 2017; Wisniewski, 1997].

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, feature emergence, etc. These strategies involve the ability to deal with conflicting attributes and the creation of new properties: simply put, a certain kind of meaning negotiation game. 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 [Confalonieri and Kutz, 2020]; (2) a formal approach based on axiom weakening to deal with conflicting attributes [Troquard *et al.*, 2018]; (3) an agent-based dialogical implementation combining (1) and (2) to simulate the meaning negotiation and construction in asymmetric combinations, as it is approached in the literature on hybrid concepts [Wisniewski, 1997; Hampton, 2017].

Our approach is related to conceptual blending [Eppe *et al.*, 2018; Ontañón and Plaza, 2010; Besold *et al.*, 2017] as well as to the system described in [Lieto and Pozzato, 2020]. We refer to [Righetti *et al.*, 2021b] for a wider discussion of related work.

2 Forms of Concept Combination

KR systems are usually characterised by their compositional behaviour. Compositionality is the principle according to which the meaning of any complex concept or expression is understood as a function of the meanings of the parts it is composed of. 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. 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

---

[*] This is an extended abstract of a paper that won the FOIS Best Paper Award at the 12th International Conference on Formal Ontology in Information Systems (FOIS 2021).

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 means that what lies in the intersection of two concepts would inherit all the features normally associated to any conjunct.

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 simply in terms of intersections of sets [Hampton, 1987; Markmar and Gentner, 1993; Wisniewski, 1997; Costello and Keane, 2000].

Asymmetry and Hybridity. Looking at noun-noun combinations in English, two parts can be distinguished, the Head and the Modifier, depending on the syntactic location of the noun [Jackendoff, 2016]. Considering the concept "Tool Weapon", the noun "weapon" plays the role of the Head, whereas "tool" is 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 would not be able to distinguish the two cases.

According to Wisniewski [1997], 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 Modifier and the Head components is highlighted (a *robin snake* is a snake that eats robin [Wisniewski, 1997, p. 168]).[1] (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 underbelly [Wisniewski, 1997, 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 *hybridisation*, and give a formal definition and computational account of it. Wisniewski [1997] 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 by Hampton [1987]. Hampton's experiments [1987] 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 aiming at investigating human's concept-forming abilities [Hampton, 1997; Gibbert *et al.*, 2012; Hampton, 2017], Hampton asked people to combine concepts that usually would not be combined, leading to impossible, or at least imaginary, objects. In [Hampton, 1997], people were presented with a list of concept pairs (e.g. "Vehicle" and "Fish", etc.), and were then asked to imagine and describe the objects resulting from the combination (e.g. "a Vehicle which is also a Fish"). If analysed just 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 be inconsistent. Still, subjects showed a great variability of strategies to solve incompatible combinations. Firstly, in order to select the 'right' properties for the compounds, people try to align properties and functions of the two component concepts. This **alignment** process corresponds to finding commonalities in the differences, thus supporting the integration of Head and Modifier concepts: when asked about "a Vehicle which is a kind of Fish", subjects could notice that while Vehicle needs Fuel, Fish needs Food (i.e. both need some kind of Energy to move). Alternatively, it may lead to identify strong incompatibilities between the concepts, that need then some strategy of conflict resolution (e.g. they may notice that 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** [Hampton, 1997].

Another strategy observed by Hampton in his experiments is the process of **instantiation**: when asked to combine two super-ordinate categories (e.g. Vehicle and Fish), people would find it easier to come up with a solution "instantiating" one of them to a more basic and well-known category (combining instead e.g. Boat and Fish). The phenomenon of instantiation does not have an obvious explanation, but it is likely due to the fact that basic categories are easier to be imagined and more familiar to subjects.

Aside from these heuristics, asymmetries between the Head and Modifier have been observed even in the case of impossible combinations: subjects keep the Head noun as a base to be modified by means of the Modifier.

3 Dialogues for Concept Combination

We here consider ontologies as sets of axioms in an appropriate logical language with the purpose of describing a particular domain of interest. We employ the well-known description logic $\mathcal{ALC}$ and assume standard DL syntax and semantics [Baader *et al.*, 2017]. Full details and definitions of what follows can be found in [Righetti *et al.*, 2021b], Sec. 3 and 4.

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. They each have a preference ordering $\prec_h$ and $\prec_m$ over the axioms of their own ontology. In the dialogue, the agents are proposing in turn axioms coming from their ontology to be added to the ontology under construction, R. When the axioms proposed by the agents turn out to render the devised ontology inconsistent, an *axiom weakening* procedure is called to solve that inconsistency.

Axiom weakening is a procedure which allows one to weaken inconsistent axioms instead of removing them. Axioms are *general concept inclusions* (GCIs) or *individual assertions*. GCIs are of the form $C \sqsubseteq D$, where C and D are

---

[1] See [Hedblom *et al.*, 2021] for a formal modelling and discussion of this kind of conceptual combination.

[2] For a formal analysis of these Hampton phenomena in weighted logics, we refer to [Righetti *et al.*, 2019; Porello *et al.*, 2019; Righetti *et al.*, 2021a].

concepts. Intuitively, a GCI can be weakened by either specialising the concept C to a smaller class, or by generalising the concept D to a larger class (both wrt a given reference ontology), exploiting appropriate refinement operators [Confalonieri *et al.*, 2020]. An individual assertion of the form $C(a)$ may be weakened into $C'(a)$, where C' is a generalisation of C.

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 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: $\mathsf{MH} \sqsubseteq H$ and $\mathsf{MH} \sqsubseteq M$, where MH is the target hybrid concept. Moreover, to avoid the trivial result we must also add an axiom $\mathsf{MH}(a)$ for a fresh individual name a, making sure that some MH's do exist.

The two agents take turns randomly following the probability distribution $(prob_h, 1 - prob_h)$. The asymmetry of the hybridisation can be enforced by a suitable weight given to the Head and to the Modifier, which induces an appropriate probability to take turns in the dialogue. In the asymmetric case, the Head agent h will be given a greater probability to play than the Modifier agent, agent m; it will have relatively more chances to insert his information into the hybridisation.

When it is its turn, agent i will choose its preferred axiom φ in O_i according to $\prec_i$, and not already entailed by the combination R. 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, and they partially determine the 'direction' of the combination. As long as φ can not be added to R without causing an inconsistency, it is replaced by 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 are 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, imposing a maximum number of moves to the agents.

We now state a few formal properties of these two algorithms. The returned ontology R is always consistent. Also, as a corollary of [Confalonieri *et al.*, 2020, Th. 2], we can show that the algorithm almost surely (i.e. with probability 1) terminates when using the refinement operators of [Troquard *et al.*, 2018]. Moreover, let R be an ontology returned by the algorithm described above (or by its bounded variant with $max\_turns \geq |O_h \cup O_m|$) 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.

4 The Case of the Vehicle and the Fish

We illustrate how the two versions of our algorithm work in the case of an *impossible combination* by simulating the combination of the concepts Fish and Vehicle as it is described by Hampton [1997] by means of our dialogue implementation.[3] We start with a consistent initial ontology, which will guide our weakening procedure. We include in our initial ontology an excerpt of the taxonomy of DOLCE [Borgo *et al.*, 2022], a cognitively oriented Foundational Ontology, formulated in $\mathcal{ALC}$. DOLCE was used to provide some of the high level ontological distinctions needed for reasoning about the possible incompatibilities between the input concepts.

Aside from DOLCE, the initial ontology contains two additional axioms, which directly relate to the concept we want to build: $\mathsf{FishVehicle} \sqsubseteq \mathsf{Fish}$ and $\mathsf{FishVehicle} \sqsubseteq \mathsf{Vehicle}$. To ensure the concept FishVehicle is not empty, we also add an instance of the concept: FishVehicle(Wanda).

Next, we need two ontologies which represent the concepts of Fish and Vehicle respectively, before the combination can be started. These can be seen as small 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. To make the two input ontologies of Fish and Vehicle interoperable, they are aligned to the common upper level provided by DOLCE.

The goal 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 her favourite axioms sooner, agent_h is more likely to add less weakened information to the ontology being built.

The last important aspect to consider is the preference order that we put on the axioms. 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) are preferred. Finally, we follow a preference order aiming at replicating the process of instantiation as described by Hampton [2017] and outlined in Section 2. In this case, agent_h prefers all the axioms containing information related to Car. In contrast, we leave agent_m's preferences as a random order.

The unbounded version of our algorithm ends when both agents have done all their possible moves, and we obtain a maximally informative ontology R about FishVehicle. The

---

[3] See https://bitbucket.org/troquard/ontologyutils/.

bounded version ends after the selected number of moves, returning a consistent ontology R for FishVehicle.

5 Evaluating Asymmetric Hybridisations

In order to evaluate the result of an asymmetric concept hybridisation we consider two parameters, namely the *asymmetry* of the combination and its *hybridity*.

The *asymmetry* aims at capturing the uneven influence of the Head and of the Modifier concepts. To measure the asymmetry of the combination R, we first measure the *ratio of preserved information* from O_i in R, i.e. we measure how much information from an ontology O_i is present in another ontology R [Porello *et al.*, 2018]. Then, we use 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 to evaluate the asymmetry of the combination.

The *hybridity* aims at capturing how much information in the combined ontology actually comes from both input ontologies. To measure the hybridity, we count the number of *hybrid descriptions* satisfied by the resulting ontology R. Intuitively, a hybrid description of Fish and Vehicle is something like 'is made of metal and has fins', i.e. a conjunction of features coming from different ontologies/concepts. Notice, crucially, that we exclude the features which are shared by the two concepts in the count.

The influence of preferences and weights. We evaluated: (1) the influence of the preference orders on the hybridity of the resulting combinations; and (2) the influence of the assignment of weights on the combinations' asymmetry.

The first preference order we consider prioritises the constraints of the ontological distinctions coming from DOLCE. In this case, agents prefer all the axioms that bridge the classes of DOLCE and the classes pertaining more strictly to the ontology of Vehicle (resp. Fish). We expected that this would emphasise 'hard' ontological distinctions, and would have had a negative effect on the hybridity value. The second preference order prioritises 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. Enforcing first the specific information for the concepts to be combined was expected to enhance the hybridity value.

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.

Our hypotheses were largely confirmed by our experiments, see Figure 1. The asymmetry remains lower when using the unbounded version of the algorithm.

Simulating Hampton's findings. 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 is then evaluated on the capability of the combined concept to satisfy the specific features of the instantiated concept, through appropriate competency questions [Grüninger and Fox, 1995].

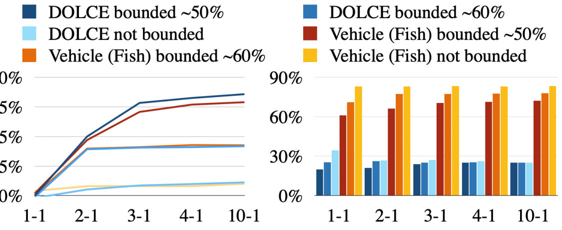

(a) Asymmetry (b) Hybridity

Figure 1: Asymmetry and hybridity values for the two preference orders, with varying weight of the Head (x-axis). Bounded $\sim$50% (or $\sim$60%) refers to the bounded variant of the algorithm.

To replicate the phenomenon of **instantiation** 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, agent_h prefers all the axioms containing information related to Car. Our experiments showed the effectiveness of this strategy: in all our runs, the FishVehicle showed *all* the features associated to Car.

Another phenomenon observed in [Hampton, 1997] is the use of **alignments**. To replicate this phenomenon, we added to the initial ontology a set of axioms inspired by the experiments described in Section 2 (e.g. Food $\equiv$ Fuel). 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 **attribute emergence**. Introducing the alignments produced in fact some mixed axioms, which were present neither in the ontology of Fish, nor in the ontology of Vehicle (e.g. a FishVehicle eats Fuel).

6 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.

The unbounded version of the dialogue game allows for the construction of 'almost perfect conjunctions' for which the two ratios of preserved information and the hybridity are high. This is particularly interesting for ontology engineering. However, since the ratios of preserved information are both very high, the asymmetry remains low. In contrast, the bounded version of the dialogue game permits to build highly asymmetrical combinations. This is in compliance with the distinctive role of the Head and of the Modifier that is observed 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. Finally, a more quantitative analysis will be necessary to more decisively evaluate our results.

References

[Baader *et al.*, 2017] Franz Baader, Ian Horrocks, Carsten Lutz, and Uli Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017.

[Besold *et al.*, 2017] T. R. Besold, K.-U. Kühnberger, and E. Plaza. Towards a computational- and algorithmic-level account of concept blending using analogies and amalgams. *Connection Science*, 29(4):387–413, 2017.

[Borgo *et al.*, 2022] S. Borgo, R. Ferrario, A. Gangemi, N. Guarino, C. Masolo, D. Porello, E. M. Sanfilippo, and L. Vieu. DOLCE: A descriptive ontology for linguistic and cognitive engineering. *Applied Ontology*, 17(1), 2022.

[Confalonieri and Kutz, 2020] R. Confalonieri and O. Kutz. Blending under Deconstruction: The roles of logic, ontology, and cognition in computational concept invention. *Annals of Mathematics and Artificial Intelligence*, 88:479–516, 2020.

[Confalonieri *et al.*, 2020] R. Confalonieri, P. Galliani, O. Kutz, D. Porello, G. Righetti, and N. Troquard. Towards even more irresistible axiom weakening. In *Proc. of DL*. CEUR-WS, Vol-2663, 2020.

[Costello and Keane, 2000] F.J. Costello and M.T Keane. Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, 24:299–349, 2000.

[Eppe *et al.*, 2018] Manfred Eppe, Ewen Maclean, Roberto Confalonieri, Oliver Kutz, Marco Schorlemmer, Enric Plaza, and Kai-Uwe Kühnberger. A computational framework for conceptual blending. *Art. Int.*, 256:105–129, 2018.

[Gibbert *et al.*, 2012] M. Gibbert, J.A. Hampton, Z. Estes, and D. Mazursky. The Curious Case of the Refrigerator–TV: Similarity and Hybridization. *Cognitive Science*, 36:992–1018, 2012.

[Grüninger and Fox, 1995] Michael Grüninger and Mark S. Fox. Methodology for the Design and Evaluation of Ontologies, 1995. Workshop on Basic Ontological Issues in Knowledge Sharing.

[Hampton, 1987] J. A. Hampton. Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, 15(1):55–71, 1987.

[Hampton, 1997] J.A. Hampton. Emergent attributes in combined concepts. In T.B. Ward, S.M. Smith, and J.Viad, editors, *Creative thought: An investigation of conceptual structures and processes*, pages 83–110. APA Press, Washington DC, 1997.

[Hampton, 2017] J. A. Hampton. Compositionality and concepts. In James A. Hampton and Yoad Winter, editors, *Compositionality and Concepts in Linguistics and Psychology*, pages 95–121. Springer International Publishing, Cham, 2017.

[Hedblom *et al.*, 2021] Maria M. Hedblom, Guendalina Righetti, and Oliver Kutz. Deciphering The Cookie Monster: A Case Study in Impossible Combinations. In *Proc. of the Twelfth International Conference on Computational Creativity (ICCC-21)*, pages 222–225, 2021.

[Jackendoff, 2016] Ray Jackendoff. *The Semantics of Compounding*, chapter English Noun-Noun Compounds in Conceptual Semantics, page 15–37. Cambridge University Press, 2016.

[Lieto and Pozzato, 2020] A. Lieto and G. L. Pozzato. A description logic framework for commonsense conceptual combination integrating typicality, probabilities and cognitive heuristics. *J. Exp. Theor. Artif. Intell.*, 32(5):769–804, 2020.

[Markmar and Gentner, 1993] A. B. Markmar and D. Gentner. Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, 32:517–535, 1993.

[Murphy, 2002] Gregory L. Murphy. *The Big Book of Concepts*. MIT press, 2002.

[Ontañón and Plaza, 2010] Santiago Ontañón and Enric Plaza. Amalgams: A formal approach for combining multiple case solutions. In Isabelle Bichindaritz and Stefania Montani, editors, *Case-Based Reasoning. Research and Development*, pages 257–271. Springer, 2010.

[Porello *et al.*, 2018] Daniele Porello, Nicolas Troquard, Rafael Penaloza, Roberto Confalonieri, Pietro Galliani, and Oliver Kutz. 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*, pages 1942–1948, 2018.

[Porello *et al.*, 2019] Daniele Porello, Oliver Kutz, Guendalina Righetti, Nicolas Troquard, Pietro Galliani, and Claudio Masolo. A Toothful of Concepts: Towards a theory of weighted concept combination. In *Proc. of DL-19*. CEUR-WS, Vol-2373, 2019.

[Righetti *et al.*, 2019] Guendalina Righetti, Daniele Porello, Oliver Kutz, Nicolas Troquard, and Claudio Masolo. Pink Panthers and Toothless Tigers: Three Problems in Classification, 2019. Proc. of the 5th Int. Workshop on Artificial Intelligence and Cognition.

[Righetti *et al.*, 2021a] Guendalina Righetti, Claudio Masolo, Nicolas Troquard, Oliver Kutz, and Daniele Porello. Concept Combination in Weighted Logic. In *Proc. of the Joint Ontology Workshops (JOWO-21)*. CEUR-WS, Vol-2969, 2021.

[Righetti *et al.*, 2021b] Guendalina Righetti, Daniele Porello, Nicolas Troquard, Oliver Kutz, Maria Hedblom, and Pietro Galliani. Asymmetric Hybrids: Dialogues for Computational Concept Combination. In *12th International Conference on Formal Ontology in Information Systems (FOIS 2021)*, pages 81–96. IOS Press, 2021.

[Troquard *et al.*, 2018] Nicolas Troquard, Roberto Confalonieri, Pietro Galliani, Rafael Peñaloza, Daniele Porello, and Oliver Kutz. Repairing ontologies via axiom weakening. In *Proc. of AAAI 2018*, pages 1981–1988, 2018.

[Wisniewski, 1997] E.J. Wisniewski. When concepts combine. *Psychonomic Bulletin & Review*, 4:167–183, 1997.