



Taking Up Thagard's Challenge: A Formal Model of Conceptual Revision

Sena Bozdag^{1,2} · Matteo De Benedetto^{1,2}

Received: 22 February 2021 / Accepted: 13 December 2021 / Published online: 15 January 2022
© The Author(s) 2021

Abstract

Thagard (1992) presented a framework for conceptual change in science based on conceptual systems. Thagard challenged belief revision theorists, claiming that traditional belief-revision systems are able to model only the two most conservative types of changes in his framework, but not the more radical ones. The main aim of this work is to take up Thagard's challenge, presenting a belief-revision-like system able to mirror radical types of conceptual change. We will do that with a conceptual revision system, i.e. a belief-revision-like system that takes conceptual structures as units of revisions. We will show how our conceptual revision and contraction operations satisfy analogous of the AGM postulates at the conceptual level and are able to mimic Thagard's radical types of conceptual change.

Keywords Conceptual revision · Thagard · Conceptual change · Belief revision · AGM postulates · Conceptual structures

1 Introduction

Thagard [36] developed a fine-grained cognitivist model of scientific theory change centered around transformations in conceptual systems. Conceptual systems are complex structures similar to frames [9, 22]. They are made of concepts and objects nodes

Sena Bozdag and Matteo De Benedetto contributed equally to this work.

✉ Sena Bozdag
a.senabozdag@gmail.com

Matteo De Benedetto
matteo.debenedetto@lrz.uni-muenchen.de

¹ Munich Center for Mathematical Philosophy, Ludwig-Maximilians-Universität München, München, Germany

² Fakultät für Philosophie, Wissenschaftstheorie und Religionswissenschaft, Munich Center for Mathematical Philosophy (MCMP), Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1, D-80539, München, Germany

connected via different kinds of links such as kind-links, instance-links, rule-links, and part-links. Changes in science then correspond to different modifications of these links. Specifically, scientific revolutions involve major transformations in part-links and in kind-links inside a conceptual system.

Thagard defended his concept-based model and the autonomy of conceptual change arguing that these revolutionary changes cannot be modeled by belief-revision theories. This supposed impossibility of modeling radical conceptual change within a belief-revision framework has been dubbed *Thagard's challenge* [24]. Specifically, Thagard's challenge claims that strong kinds of conceptual change are irreducible to belief-revision types of changes, because the formers involve holistic recombinations of links and nodes in a given conceptual system that cannot be modeled by any piece-meal belief-revision operation. This irreducibility shows for Thagard how frame-based representation of knowledge, despite being expressively equivalent to first-order logic, is procedurally different [33, 34].

Despite the enormous expansion of the belief-revision literature in the last thirty years [12] and recent work connecting it with philosophy of science [23], Thagard's challenge has not received so much attention¹. The main aim of this work is to suggest a way of taking up Thagard's challenge by developing a belief-revision-like framework capable of modeling the radical types of conceptual change described by Thagard. Specifically, we will present a conceptual revision framework in which we can revise and contract conceptual structures, i.e. set-theoretic representations of Thagard's conceptual systems. Our change operations will be reminiscent of the ones used in base-generated belief change theories [12, 29], but working on conceptual structures instead of belief bases.

Our choice of units of revision, i.e. conceptual structures, makes our system differ from other applications of belief-revision to the problem of scientific change. Traditionally, belief-revision theories deal with piece-meal changes in a belief set similar to the kind of changes happening in normal science (cf. [10]). In applying these theories to the problem of scientific change, logicians have focused on mirroring changes in scientific theories as changes in (usually structured) belief sets [3, 6, 13, 20, 32]. This belief-centered take on scientific change is exactly the reason why Thagard claims that belief revision theories are not adequate for representing conceptual change [34, 35]. We instead chose to model conceptual change at its native level of abstraction, without any reference to the belief level. We achieve this by lifting the methodology of belief revision theories to the conceptual level. As a result, the aim of our change operations will then be the preservation of the consistency of conceptual structures. This consistency is understood as the satisfaction of some structural constraints on the components of a conceptual structure that ensure the overall consistency of the knowledge represented by it. The knowledge represented via our conceptual structures is similar to the content represented by description logics [39], since they also represent knowledge about concept hierarchies².

¹For a couple of exceptions that suggest ways of expanding belief revision to treat certain aspects of scientific conceptual change, and thus could be considered implicit partial replies to Thagard's challenge, see [7, 31] and [11].

²AGM-style and base-generated revision theories in description logics are also proposed in [26] and in [27].

We will show how eight out of Thagard's nine degrees of conceptual change can be adequately represented in our conceptual revision framework. Specifically, we will mirror all changes that reflect transformations in the structure of conceptual systems, leaving out what Thagard calls *tree switching*, i.e. a more radical kind of change involving a gestalt-like switch in the external interpretation of a given conceptual system. We will demonstrate how each of these eight degrees of conceptual change is mirrored by a specific case of our conceptual revision and contraction operations.

By taking up Thagard's challenge, we intend to show the limitations of belief-centered approaches to scientific change. Our framework shows that, in order to have a satisfactory belief-revision-like account of conceptual change, we have to work at the conceptual level. The extension of traditional belief revision systems into our conceptual revision framework involves rethinking some of the core notions of belief revision such as consistency and completeness. We move from a set-theoretic notion of consistency (and completeness) to a structural one, and we show how these conditions of consistency organize conceptual knowledge. By shifting to the conceptual level we therefore create new opportunities for more general systems of belief revision which can work on different levels of reasoning. Furthermore, using belief revision as a background framework allows us to bridge the logic-oriented approaches to conceptual change with the more cognitive-oriented ones, therefore achieving a logical taxonomy of types of conceptual change.

In Section 2, we will present Thagard's account of scientific conceptual change. In Section 3, we will present our belief-revision-like model of conceptual revision. More specifically, we will present a revision and a contraction operation that work on conceptual structures. In Section 4, we will show how our conceptual revision model satisfies several rationality postulates analogous to the AGM ones for belief revision theories [1]. In Section 5, we will demonstrate how our revision and contraction operations are able to mirror several kinds of conceptual changes depicted in Thagard's framework. Finally we will draw some general conclusion on the results and limitations of the present article and we will sketch some directions for future work.

2 Thagard's Model of Scientific Conceptual Change

Thagard's model of conceptual change in science is built upon the notion of a conceptual system [36]. A conceptual system is a set of nodes interconnected via various kinds of links, a structure that closely resembles frames [9, 22]. More specifically there are two kinds of nodes and four kinds of links that can figure in a conceptual system. Nodes can be concept nodes or object nodes, mirroring respectively concepts and objects. Concept nodes can be connected with other concept nodes via three kinds of links (kind-links, part-links, rule-links) and with other object nodes via another kind of links, i.e. instance-links³:

³Note that Thagard in presenting his framework mentions also a fifth kind of link, property-links [36, p. 31]. This kind of links is supposed to mirror the information of a given object possessing a given property, but it does not seem to play any role into Thagard's model of conceptual change. It is in fact not mentioned in his abstract presentation of the model [36, pp. 34-39] nor in any of the case studies [36, pp. 131-224]. We chose therefore to omit this kind of link from our discussion.

- Kind-links (from concepts to concepts): intuitive reading ‘is a kind of’, example ‘the canary is a kind of bird’.
- Part-links (from concepts to concepts): intuitive reading ‘a whole has a given part’, examples ‘the beak is a part of birds’, ‘fins are part of fishes’.
- Rule-links (from concepts to concepts): intuitive reading as expressing generic relations between concepts, example ‘canaries are yellow’.
- Instance-links (from objects to concepts): intuitive reading ‘is an instance of’, example ‘Tweety is a canary’.

The most important kinds of links are the ones between conceptual nodes. Kind-links and part-links specify what the constituents of (a part of) the world are according to a given conceptual system. Concepts within conceptual systems are organized in kind-hierarchies and part-hierarchies, i.e. sets of kind-links and part-links that are constrained in a tree-like form in order to give a consistent picture of (a part of) the world. Rule-links instead represent factual information and default reasoning mechanisms codified within the conceptual system. They are not organized in a hierarchy, but they can be divided between weak-rules and strong-rules depending on the strength of the information they represent.

Conceptual changes on a given conceptual system are then ordered by Thagard [36, p. 35] in terms of how radical they are, from the least to the most radical:

1. Instance-addition: adding an instance relation saying that a given individual is an instance of a given concept, e.g. ‘that blob in the distance is a whale’.
2. Rule-addition: adding a rule relation, e.g. ‘whales can be found in the Arctic ocean’ or ‘whales eat sardines’⁴.
3. Part-addition: adding a new part-relation, e.g. ‘whales have spleens’.
4. Kind-addition: adding a new kind-relation, e.g. ‘a dolphin is a kind of whale’.
5. Concept-addition: adding new concept, e.g. ‘sound-wave’ or ‘narwhal’.
6. Kind-collapse: collapsing part of a kind-hierarchy, abandoning a previous distinction, e.g. when Darwin collapsed species and varieties within a species distinction.
7. Hierarchy-reorganization: shifting concepts or parts of the kind and part-hierarchies to another part of the hierarchies, i.e. *branch-jumping* such as Darwin’s shift of humans to the animal-mammal part of the kind-hierarchy. It may also involve transformation of part-relations onto kind-relations and *vice versa*.
8. Tree-switching: changing the organizational principle of the kind-hierarchy, e.g. Darwin’s switch from a morphological kind-hierarchy to an evolutionary one.

The aforementioned Thagard’s challenge consists of the claim that belief revision systems can model just the first two degrees of conceptual change, i.e. instance-addition and rule-addition, but not the other six [36, p. 36]. Both instance-addition

⁴Note that Thagard actually divides the rule-addition kind of conceptual change in two distinctive sub-types: weak-rule and strong-rule addition. Since Thagard’s distinction between weak and strong rules is entirely pragmatic [36, p. 35], being it based on the problem-solving power of a rule, we collapse in our framework these two types of changes in one.

and rule-addition represent in fact piecemeal additions that do not involve any recombination in the part- and kind-hierarchies of a given conceptual system. These two kind of changes can then be adequately mirrored as changes at the belief-level, revising for instance the extension of a predicate and its prototypical instances [32]. The other six, more radical kinds of conceptual changes are more holistic types of changes, since they involve the adjustment of the part- and kind-hierarchies (as well as rule and instance-relations) of the whole conceptual system. These changes represent in fact how scientists in revolutionary times add new concept, delete old concepts, drastically reorganize kind and part-hierarchies, and sometimes they even change the organizational principle of the hierarchical tree. Due to their holistic character, these changes cannot be easily mirrored as changes at the belief level like the first two. These revolutionary changes, then, are for Thagard [36, p. 28] evidence that conceptual change is irreducible to belief-revision.

3 A Conceptual Revision Model

In the previous section we described Thagard's model of scientific conceptual change. In this section, we present a formal model of conceptual revision that is able to model the kind of changes described by Thagard, including the more radical ones. Our system is equipped with a change mechanism similar to the one of base-generated belief revision frameworks, but the units of change are structure mirroring Thagard's conceptual systems rather than belief bases. In this way, we can mirror Thagard's changes at their native level of abstraction, namely the conceptual level.

Our framework takes as its units of changes set-theoretic entities which we call *conceptual structures*. We define two different domains, one for concepts and one for individual objects, as the primary elements of a conceptual structure. Our conceptual structures enrich these two basic domains with different relations between elements of these domains. Mirroring Thagard's system, we define three two-place relations between elements of the concept domain (kind-relation, part-relation, rule-relation) and one two-place relation between elements of the object domain and elements of the concept domain (instance-relation).

Conceptual Structures and Conceptual Hierarchies Formally, a *conceptual structure* is defined as follows:

Definition 1 $CS = \langle C, O, K, P, R, I \rangle$ is a conceptual structure iff,

- C and O are (possibly empty) finite domains of (respectively) atomic concepts and individual objects.
- $K = \{\langle x, y \rangle, \dots\}$ and $P = \{\langle x, y \rangle, \dots\}$ are two-place irreflexive relations between elements of the concept domain, such that $x, y \in C$ and $\langle x, y \rangle$ is an ordered pair. They represent respectively Thagard's kind and part links between concept nodes. If $\langle x, y \rangle \in K$, we write $x \sqsubset_K y$ (same for $x \sqsubset_P y$, if $\langle x, y \rangle \in P$).
- $R = \{\langle x, y \rangle, \dots\}$, with $x, y \in C$ and $\langle x, y \rangle$ is an ordered pair, is a two-place anti-symmetric relation between elements of the concept domain. It represents Thagard's rule links between concept nodes.

- $I = \{\langle a, x \rangle, \dots\}$ with $a \in O$ and $x \in C$ and $\langle a, x \rangle$ is an ordered pair, is a two-place anti-symmetric relation between elements of the object domain and elements of the concept domain. It represents Thagard's instance links between object and concept nodes.

We can then single-out specific kind-relations and part-relations through a tree-like structural requirement. Relations satisfying this requirement are then called respectively *kind-hierarchies* and *part-hierarchies*. This requirement is our way of rationally reconstructing Thagard's implicit structural requirements on conceptual systems. Similarly, we introduce criteria to single out certain rule and instance relations as *consistent* rule and instance relations. With these further criteria we mirror common constraints on how knowledge is represented in a consistent way by frames (cf. [2, 9]). Then, a conceptual structure is a *conceptual hierarchy* iff its kind relation is a kind-hierarchy, its part relation is a part-hierarchy, its rule relation is a consistent rule relation, its instance relation is a consistent instance relation, and all the concepts and objects occurring in its relations are members respectively of the concept domain or the object domain.

Definition 2 $CH = \langle C, O, K_h, P_h, R_{cons}, I_{cons} \rangle$ is a conceptual hierarchy iff,

- C and O are (possibly empty) finite domains of respectively concepts and objects, which include all the concepts and objects that appear in the relations.
- K_h is a kind-hierarchy, i.e. a transitive kind-relation $K = \{\langle x, y \rangle, \dots\}$ that, if non-empty, has a top element and from any other element of the ordering there exists a unique path to this top element modulo transitivity.
- P_h is a part-hierarchy, i.e. a transitive part-relation $P = \{\langle x, y \rangle, \dots\}$ that, if non-empty, has a top element and from any other element of the ordering there exists a unique path to this top element modulo transitivity.
- R_{cons} is a consistent rule-relation, i.e. a rule-relation $R = \{\langle x, y \rangle, \dots\}$ such that $\forall x, y, z \in C$ if $\langle x, y \rangle \in R_{cons}$ and $z \sqsubset_K x$, then $\langle z, y \rangle \in R_{cons}$.
- I_{cons} is a consistent instance-relation, i.e. an instance-relation $I = \{\langle a, x \rangle, \dots\}$ such that $\forall x, y \in C$ and $\forall a \in O$ if $\langle a, x \rangle \in I_{cons}$ and $x \sqsubset_K y$, then $\langle a, y \rangle \in I_{cons}$; $\forall x, y \in C$ and $\forall a \in O$ it holds that $\langle a, x \rangle \in I_{cons}$ and $\langle a, y \rangle \in I_{cons}$ only if $x \sqsubset_K y$ or $y \sqsubset_K x$.

The transitivity requirement says about a kind-hierarchy and a part-hierarchy that $\langle x, y \rangle, \langle y, z \rangle \in K_h(P_h)$ implies that $\langle x, z \rangle \in K_h(P_h)$. We define the *top element in a kind-relation (part-relation)* as follows: given a conceptual structure H , the concept domain C_H and the kind-relation K_H (part-relation P_H) of H , a concept $a \in C_H$ is a top element in K_H (P_H) iff for all concepts $t \neq a$ in C_H which occur in a pair in K_H (P_H), it holds that $\langle t, a \rangle \in K_H$ ($\in P_H$). Hence, an ordering on concepts determined by a kind-hierarchy (part-hierarchy) is upward closed. By a *unique path to the top element from any other element modulo transitivity* we mean that, given a is a top element in a kind-relation (or in a part-relation), for all $t \neq a$ which occur in a pair in the kind-relation (part-relation), if $\langle t, y \rangle$ and $\langle t, z \rangle$ are pairs in the kind-relation (part-relation), and $y \neq z$, then either $\langle z, y \rangle$ or $\langle y, z \rangle$ is also a pair in


the same relation. In other words, kind-hierarchies and part-hierarchies do not allow upward branchings (Fig. 1). The requirements on the consistent rule-relation and the consistent instance-relation say that these relations are transitively closed modulo the relevant kind-hierarchies. We have also required that a consistent instance-relation does not allow upward branching.

3.1 Revision on Conceptual Structures

In this section we will describe how a conceptual structure should be revised in our framework. Revising a conceptual structure means adding new elements (of a conceptual structure) to an existing conceptual structure, while preserving (or restoring) the consistency of the revised conceptual structure. The information we want to add (or delete) can be a concept, an object, a kind-relation, a part-relation, a rule-relation, or an instance-relation. Consistency is characterized by the idea of conceptual hierarchies, i.e. by structural restrictions on the different kinds of relations connecting concepts and objects within a conceptual structure. Therefore, the goal of our conceptual revision framework is to define change operations which preserve these structural limitations.

We start with identifying the form of the eligible arguments for a revision. Suppose we want to revise an existing conceptual structure with an instance (i.e. an instance-addition in Thagard's framework). Suppose we want to add that Bob is an orka ($I\langle Bob, orka \rangle$). If the existing structure does not already contain the concept of an orka and the object Bob, simply adding the instance link would not make sense. Hence we require that the arguments of conceptual revision are formulated as proper conceptual structures. That is, as an argument of a revision, we express the above instance link as a conceptual structure (let us call it C) which consists of the following: $\mathcal{C}_C = \{orka\}$, $\mathcal{O}_C = \{Bob\}$, $I_C = \{\langle Bob, orka \rangle\}$, $K_C = P_C = R_C = \emptyset$.

Fig. 1 A conceptual hierarchy made of four concepts and four kind-links between them



Next, we have to choose what kind of revision operation we want in our framework. The consistency of the revised (conceptual) structure could for instance be preserved while making the additions, or it could instead be restored after the addition process [13]. The former approach it is typical of the AGM belief revision paradigm [1], while the latter is common amongst base-generated revision theories [12, 29]. Our revision operations on conceptual structures will resemble the approach of base-generated revisions. Given a conceptual structure and an argument of revision, we will first add them on top of each other, while restoring the transitivity of the relation (by adding new links) in the resulting conceptual structure. Next, we will retrieve the consistent parts of this structure to build the revised (consistent) conceptual structure.

In base-generated belief revision, one adds the new information to the existing body of beliefs via a set-theoretical union operation.⁵ The potential inconsistency in the expanded belief set is caused by too much information. To eliminate this inconsistency, the (less entrenched) beliefs responsible for it are deleted from the new belief set.

In our framework, the inconsistency of a conceptual structure may be caused by too much information or by too little information. In fact, the pivotal requirement of transitivity for the relations of a conceptual hierarchy could be lost during revision. In order to restore the consistency of a conceptual structure, we need to eliminate the inconsistent parts and to repair the transitivity of its relations. We will deal with the transitivity issue in the first step of our revision, i.e. the addition of new information to a conceptual structure (which constitutes the operation of conceptual expansion). In the second step of our revision, we will instead deal with the problem of inconsistent parts, proposing a mechanism that retrieves consistent parts of the expanded conceptual structure⁶.

Conceptual expansion We will define conceptual expansion as the process of adding new information to a conceptual structure, while restoring the transitivity of the relation in the resulting conceptual structure. More specifically, conceptual expansion will be performed via the fusion models.

⁵We significantly simplify the mechanism of base generated revisions. In fact the new information is added on top of a structure called a *belief base* which is the foundation of an agent's beliefs. Moreover, the addition of the new information is set-theoretical only if we are dealing with *flat* belief bases which are not ordered by a preference or entrenchment relation. Full theories of belief revision usually include such orderings to account for the rationality of changing beliefs. Adding beliefs in the AGM paradigm also goes further than the union operation as it involves taking the deductive closure of the new belief set.

⁶While we could keep the expansion process simple and deal with the (possibly lost) transitive closure in the process of retrieving information, we believe it is more natural to restore the transitivity required for consistency as part of the expansion operation. One reason is that restoring transitive closure will be done by adding new links, and keeping all additions as part of the expansion and limiting the process of retrieval of information to elimination of some (less entrenched) parts of the expanded structure which contribute to the inconsistency allows simpler definitions for the two processes. Another reason is that this allows us to characterise a conceptual expansion operation which results in structures which resemble conceptual hierarchies to an extent that they are somehow useful in practice.

Definition 3 A tuple $CS^\oplus = \langle \mathcal{CS}, \oplus \rangle$ is fusion model on conceptual structures iff \mathcal{CS} is a set of conceptual structures that is closed under the total conceptual fusion function \oplus from $CS \times CS$ to CS , uniquely determined by the following:

- $\mathcal{C}_{A\oplus B} = \mathcal{C}_A \cup \mathcal{C}_B$
- $\mathcal{O}_{A\oplus B} = \mathcal{O}_A \cup \mathcal{O}_B$
- $K_{A\oplus B} = TC(K_A \cup K_B)$
- $P_{A\oplus B} = TC(P_A \cup P_B)$
- $R_{A\oplus B} = TC|K_{A\oplus B}(R_A \cup R_B)$
- $I_{A\oplus B} = TC|K_{A\oplus B}(I_A \cup I_B)$

TC stands for the transitive closure operation on our sets of pairs. For instance, the transitive closure of a kind-relation K is the smallest transitive set of pairs that contains K such that if $\langle a, b \rangle$ and $\langle b, c \rangle$ is in $TC(K)$ then $\langle a, c \rangle \in TC\{K\}$. Transitive closure on rule-relations and instance-relation are via transitivity modulo the kind-relation. Thus, an instance-relation I is transitively closed modulo the relevant kind-relation K ($TC|K$) iff given $\langle b, c \rangle \in K$ and $\langle a, b \rangle \in I$, then also $\langle a, c \rangle$ is in I .

The above model specifies how to add a full conceptual structure on top of another one.

We show with an example how conceptual expansion via conceptual fusion models works. Let H be the conceptual hierarchy depicted in Fig. 1 such that

$$C_H = \{[mammal], [human], [whale], [orka]\}, O_H = P_H = R_H = I_H = \emptyset.$$

$$K_H = \{\langle human, mammal \rangle, \langle whale, mammal \rangle, \langle orka, whale \rangle, \langle orka, mammal \rangle\}.$$

Let A be a conceptual structure consisting of: $C_A = \{[fish], [orka]\}$, $O_A = \{(Bob)\}$, $K_A = \{\langle orka, fish \rangle\}$, $I_A = \{\langle Bob, orka \rangle\}$, $P_H = R_H = \emptyset$. Then, by the fusion model, we can obtain the conceptual structure $H \oplus A$ determined by the following elements:

$$C_{H\oplus A} = \{[mammal], [human], [whale], [orka], [fish]\}, O_{H\oplus A} = \{(Bob)\}$$

$$K_{H\oplus A} = \{\langle human, mammal \rangle, \langle whale, mammal \rangle, \langle orka, whale \rangle,$$

$$\langle orka, mammal \rangle, \langle orka, fish \rangle\}, P_{H\oplus A} = R_{H\oplus A} = \emptyset,$$

$$I_{H\oplus A} = \{\langle Bob, orka \rangle, \langle Bob, fish \rangle, \langle Bob, whale \rangle, \langle Bob, mammal \rangle\}.$$

The instance pairs $\langle Bob, mammal \rangle$, $\langle Bob, whale \rangle$ and $\langle Bob, fish \rangle$ in $I_{H\oplus A}$ are additions to the simple union of I_H and I_A via the transitive closure operation (Fig. 2).

Conceptual revision Conceptual expansion may not always produce a conceptual hierarchy. Since our aim is to obtain conceptual hierarchies as the result of revisions, we propose a consistency check mechanism for restoring the consistency of conceptual structures. This mechanism is a modified version of the consolidation operation described in relation to the base generated revisions [29, p. 40]. In what follows,

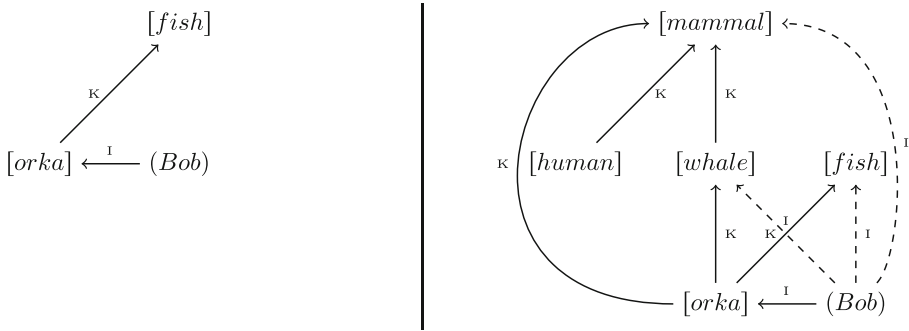


Fig. 2 The conceptual structure A (on the left) and the conceptual structure $H \oplus A$ (on the right)

we use the notions of a *substructure* and a *maximal hierarchy within a conceptual structure*:

Definition 4 Given two conceptual structures H, H' , we say that H is a substructure of a conceptual structure H' (denoted by $H \subseteq H'$) iff the domains and relations defined in H are subsets of the respective domains and relations defined in H' .

Definition 5 Given two conceptual structures H, H' , we say that H is a maximal hierarchy within H' iff H is a substructure of a conceptual structure H' and any non-trivial expansion of H within H' is not a conceptual hierarchy.

A substructure of a conceptual structure is then a conceptual structure, the components of which (such as objects, concepts, and relations) are subsets of the respective components of the other conceptual structure. For instance, consider the conceptual structure $H \oplus A$ in Fig. 2 above. The conceptual structure below is a substructure of $H \oplus A$:

$$\begin{aligned} \mathcal{C} &= \{[mammal], [whale], [orka]\}, \mathcal{O} = \{(Bob)\}, \\ K &= \{\langle whale, mammal \rangle, \langle orka, whale \rangle, \langle orka, mammal \rangle\}, \\ I &= \{\langle Bob, orka \rangle, \langle Bob, whale \rangle, \langle Bob, mammal \rangle\}, P = R = \emptyset. \end{aligned}$$

A substructure that cannot be (non-trivially) expanded to a conceptual hierarchy, within the given conceptual structure, is a maximal hierarchy. For instance, the above example of a substructure is not a maximal hierarchy, within $H \oplus A$. This is because it can be expanded to a larger hierarchy that includes the concept [human] and the kind link $\langle [human], [mammal] \rangle$. This expanded structure is instead a maximal hierarchy within $H \oplus A$.

While revising a conceptual structure, we first expand it with the argument of the revision. Since it might be the case that the expansion operation fails to produce a conceptual hierarchy, we need a mechanism which marks off the best candidates for the revised conceptual structure. In order to meet some rationality criteria, such selection mechanisms usually rely on an ordering of the alternatives based on the

preferences of the selecting agents. Hence we include in our conceptual revision model a preorder among conceptual structures:

Definition 6 A conceptual revision model is a tuple $CS^{\oplus \leq} = \langle CS, \oplus, \leq \rangle$, such that

- $\langle CS, \oplus \rangle$ is a fusion model on conceptual structures, and
- \leq is a connected preorder on a set of conceptual structures.

The preorder between conceptual structures is a preference ordering. If X, Y are conceptual structures in a set, and $X \leq_{CS^{\oplus \leq}} Y$, we say the conceptual structure X is at least as preferred as the conceptual structure Y given the model $CS^{\oplus \leq}$. The most preferred conceptual structures in a set are the ones that are minimal under \leq , i.e., X is most preferred with respect to \leq in a set S of conceptual structures iff for all Y in S , it holds that $X \leq_{CS^{\oplus \leq}} Y$.

We propose that during the revision operation, the preference ordering on a set of conceptual structures changes as follows: let CS be a non-empty set of conceptual structures, and $H \in CS$ the argument of a revision, and let \leq be the pre-revision preference ordering on CS and \leq' be the revised preference ordering, then,

- for all $A, B \in CS$, if $H \subseteq A$ and $H \not\subseteq B$ then $A \leq' B$ and $B \not\leq' A$, and if $H \subseteq B$ and $H \not\subseteq A$ then $B \leq' A$ and $A \not\leq' B$,
- otherwise, $A \leq' B$ iff $A \leq B$.

In our model, then, we give priority to new data, i.e. the argument of the revision. More specifically, given the revised preference ordering, the parts of the expanded conceptual structure which include the new data are strictly more preferred to the parts which exclude the new data. Apart from this change, the preference ordering remains the same. The revised conceptual structure will then be given by the intersection of the most preferred maximal hierarchies within the expanded conceptual structure, according to the revised preference ordering. The maximality principle concerning these hierarchies is assumed in order to preserve as much information as possible while revising.

A preference ordering may, in fact, rate multiple conceptual hierarchies as the best ones. In the belief revision literature, these cases are commonly solved by taking the intersection of the selected alternatives, following the *partial meet contraction and revision* operations introduced within the AGM paradigm. However, as we will show with an example, intersecting multiple conceptual hierarchies may generate inconsistent conceptual structures. As a solution to this problem, we propose a repetitive revision operation, where the intersection mechanism is repeated until a conceptual hierarchy is obtained.

Since our revision operation involves changing the preference order in a revision model, it is essentially a model-changing operation. Therefore, even if the expansion of the initial conceptual structure with the argument of the revision is a conceptual hierarchy, the revision operation does not reduce to expansion, since changes on the preference ordering are significant for iterating any change operation on the new conceptual structures.

Before moving on to a simple example of a conceptual revision, let us say more about the preference ordering for conceptual structures that we assume in our model. In the scientific context, there are many different factors that might be considered for determining such a preference ordering. One could, for instance, establish a preference ordering based on corroboration, where more empirically confirmed parts of a conceptual structure are preferred (cf. Hansson's scientific corpus model, [13]). Alternatively, another option could be to prioritize parts of a conceptual structure involving empirical concepts in comparison to parts involving theoretical terms, obtaining a sort of empiricist conceptual revision. It could also be possible to hold a specific part of a conceptual structure as the most preferred, identifying it with the constitutive part of a scientific theory à la Friedman [8] or with what Lakatos called the hard-core of a scientific research program (cf. [17]). Moreover, a suitable preference ordering could be determined virtually by any specific bundles of theoretical values discussed in debates over theory choice in science, such as simplicity, fruitfulness, empirical adequacy, and the like. In the related literature in philosophy of science, one finds contrasting bundles of epistemic values defended (e.g. [16, 18]).

Given this plethora of possible criteria for determining a preference ordering for our conceptual structures, we decided here to stay neutral on which specific criteria we prefer. Consistently, we will assume an arbitrary preference ordering for each set of conceptual structure prior to revision and focus on how such an ordering change when the conceptual structure is revised.

Since we are dealing with revolutionary changes in particular, an important assumption of our model is, as we have already mentioned, to give priority to new data. In revising a given conceptual structure with some new information, our framework will re-order the preference ordering by giving priority to the new information, i.e. to the argument of the revision. This is, of course, not the only viable option. In the related belief revision literature, one could find several preference re-ordering operations, ranging from extremely conservative options to truly radical ones (cf. [30]).

Now, in the scientific context, just like we found many different possible factors for determining a preference ordering, we can find several, often contrasting, factors for choosing a specific way of re-ordering the preference ordering of our conceptual structures. Our choice of giving priority to new information could be, for instance, generally supported by a Popperian confutationalist stance (cf. [25]) or by the general idea that conservativity is an enemy of scientific progress. Several scholars, including Thagard [38], have in fact argued that the principle of minimal change should be abandoned in the scientific context in favor of a more radical revision that prioritizes new information. However, a different picture of scientific activity and rationality could favor a radically different preference re-ordering operation. A conservative operation, prioritizing the old parts of a conceptual structure, could perhaps be preferred by philosophers of science influenced by Kuhn [15] and Lakatos [17], who believed that mature science always includes some degree of conservativity. Various finer-grained distinctions on this re-ordering could also be made, such as restricting this conservativity to a specific part (e.g. the hard-core of a theory) of the conceptual structure or to specific concepts (e.g. the empirical ones).

In general, the choice of the preference re-ordering operation, just like the choice of the criteria determining such a preference ordering, seems extremely dependent on the specific view of scientific dynamics that one favors. We chose here to prioritize novel information, because this seems to us the choice closer to the spirit of Thagard's framework⁷. That said, it should be clear to the reader that this is by no means the only interesting way in which the preference ordering of a conceptual structure can be re-ordered.

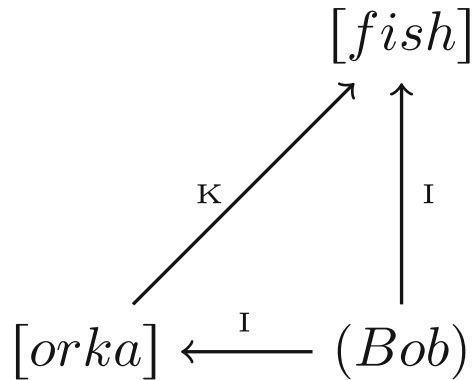
To see how the revision operation is applied to conceptual structures, consider the conceptual structures H and A in the above example. Suppose we want to revise H with A instead of simply expanding the former with the latter. The first step of the conceptual revision process consists of the aforementioned expansion $H \oplus A$ (depicted in Fig. 2). Note that $H \oplus A$ is not a conceptual hierarchy. This is because, its kind-relation is not a kind-hierarchy and its instance-relation is not a consistent instance-relation. The kind-relation of $H \oplus A$ is not a kind-hierarchy because: i) although it is non-empty, it does not have a top element, since the concept $[mammal]$ and the concept $[fish]$ are the top elements of two distinct conceptual substructures; ii) the kind-links $\langle orka, whale \rangle$ and $\langle orka, fish \rangle$, both of which occur in the kind relation of the structure, generate an upward branching. The instance-relation of $H \oplus A$ is not a consistent instance-relation because the pairs of instance-links $\langle Bob, whale \rangle$, $\langle Bob, fish \rangle$ and $\langle Bob, mammal \rangle$, $\langle Bob, fish \rangle$ generate upward branchings.

Since $H \oplus A$ is not a conceptual hierarchy, we continue the revision operation by identifying and intersecting the best conceptual hierarchies within $H \oplus A$, determined by the revised preference ordering. As the new data in this revision is the conceptual structure A , after the revision, the conceptual hierarchies of which A is a substructure are strictly more preferred over the ones which exclude some part of A . Then, an easy way to identify the best maximal hierarchies within $H \oplus A$ is to start with A as the base structure and expand it within $H \oplus A$ until reaching a maximal conceptual hierarchy. However, it might be the case that there are no maximal hierarchies within the expanded structure that include the new data. Then, one identifies all maximal hierarchies within the expanded structure; their preference ordering is based on the 'otherwise' clause in our definition above.

Once we accept A within $H \oplus A$, we can only expand it with the instance link $\langle Bob, fish \rangle$. This is in particular due to accepting that orkas are fishes. This information is inconsistent with the information that orkas are whales, and that orkas are mammals, since fishes are neither whales nor mammals given the expanded structure. The part of $H \oplus A$ concerning the kind-links $\langle human, mammal \rangle$ and $\langle whale, mammal \rangle$ does not directly contradict with the information given in A .

⁷We should note that technically Thagard's framework does not involve any preference ordering. Although, it has an evaluation mechanism between different conceptual systems based on the notion of explanatory coherence [37]. This mechanism is however used only to compare fully finished conceptual structures after the changes have taken place. We leave the study of the relations between Thagard's evaluation mechanism and our preference ordering for future work.

Fig. 3 The conceptual hierarchy M



However, accepting also this part would mean that the resulting structure does not have a top element since fishes are not mammals, and mammals are not fishes. Consequently, the following is the unique conceptual hierarchy M (Fig. 3) within $H \oplus A$ which fits the description: $\mathcal{C}_M = \{[fish], [orka]\}$, $\mathcal{O}_M = \{(Bob)\}$, $K_M = \{\langle orka, fish \rangle\}$, $I_M = \{\langle Bob, orka \rangle, \langle Bob, fish \rangle\}$, $P_M = R_M = \emptyset$. The revision of H with A finalizes here⁸.

In the case that there are more than one maximal hierarchies in the expanded structure which includes the argument of the revision, we pick up the most preferred among these, and intersect them. It might be the case that the first iteration of this last step do not result in a conceptual hierarchy. In that case, we repeat by intersecting the most preferred maximal hierarchies within the conceptual structure we have obtained as the result of the latest iteration. This is done until one reaches a conceptual hierarchy. Our next example shows how the revision operation is applied repetitively.

Suppose we want to revise a conceptual structure X such that $\mathcal{C}_X = \{a, b, A, B\}$, $K_X = \{\langle a, A \rangle, \langle b, A \rangle, \langle a, B \rangle, \langle b, B \rangle, \langle B, A \rangle, \langle A, B \rangle, \}$ and $\mathcal{O}_X = P_X = R_X = I_X = \emptyset$ with the empty conceptual structure $\{\emptyset\}$. There are two maximal hierarchies within the expanded conceptual structure $X \oplus \{\emptyset\}$ whose kind relations are the following, $K_{X1} = \{\langle a, A \rangle, \langle b, A \rangle, \langle a, B \rangle, \langle b, B \rangle, \langle B, A \rangle\}$ and $K_{X2} = \{\langle a, A \rangle, \langle b, A \rangle, \langle a, B \rangle, \langle b, B \rangle, \langle A, B \rangle\}$. Since the empty conceptual structure is a substructure of both, the preference ordering remains untouched and therefore the two conceptual hierarchies are equally preferred. Then, their intersection

⁸This example of conceptual revision reveals a significant amount of information loss as a result. This is connected to the revolutionary aspect of the scientific changes we want to represent. As it was famously stressed by Kuhn [15], scientific revolutions involve often the loss of information in the transition from one scientific theory to its successor, a phenomenon commonly known in philosophy of science as Kuhnian loss. A more conservative conceptual revision may also be defined in our framework. We could, for instance, weaken our requirements for conceptual hierarchies, eliminating the necessity of having a top element while maintaining the ban of upward-branching. This change would allow a conceptual hierarchy to consist of several disconnected conceptual hierarchies. This alternative definition would make the revised conceptual structure in the above example to consist of the original result M and the following conceptual hierarchy: $\mathcal{C} = \{[human], [mammal], [whale]\}$, $K = \{\langle human, mammal \rangle, \langle whale, mammal \rangle\}$, $\mathcal{O} = P = R = I = \emptyset$.

includes a kind-relation which does not have a top element, i.e. $K_{X1 \cap X2} = \{\langle a, A \rangle, \langle b, A \rangle, \langle a, B \rangle, \langle b, B \rangle\}$, that cannot be the kind-relation of a conceptual hierarchy. We therefore repeat the revision operation, first determining the best maximal hierarchies within the conceptual structure $X1 \cap X2$. These are the hierarchies with the kind-relations $K_Y = \{\langle a, A \rangle, \langle b, A \rangle\}$ and $K_Z = \{\langle a, B \rangle, \langle b, B \rangle\}$. If they are preferred equally, then the revised conceptual hierarchy has in its concept domain only the concepts a and b , together with an empty kind-relation

We can then define our conceptual revision operation as follows:

Definition 7 Given a conceptual revision model $CS^{\oplus \leq} = \langle CS, \oplus, \leq \rangle$, and H, A conceptual structures, H revised with A (let us denote it with $H * A$, and the preference ordering after revision with \leq') is determined by the following cases:

1. $H * A \equiv H' = \cap \{B : B \text{ is a maximal hierarchy within } H \oplus A \text{ and for all maximal hierarchies } I \subseteq H \oplus A \text{ it holds that } B \leq' I\}$, if H' constitutes a conceptual hierarchy;
2. $H * A \equiv H'' = \cap \{C : C \text{ is a maximal hierarchy within } H' \text{ and for all maximal hierarchies } I \subseteq H' \text{ it holds that } C \leq' I\}$, if H' does not constitute a conceptual hierarchy and H'' constitutes a conceptual hierarchy;
3. repeat case 2 with the maximal hierarchies within the resulting conceptual structure (e.g. starting with H'' in the first repetition) until reaching a conceptual hierarchy as the result of the intersection, if otherwise.

3.2 Contraction on Conceptual Structures

Contracting a conceptual structure means eliminating a part of it. While our contraction operation is defined based on conceptual revision, it differs from revision significantly in terms of how the argument of a contraction should be formulated or expressed. Suppose we want to contract an instance-pair $\langle x, y \rangle$ from a conceptual structure. In regards to the arguments of revision, we required that they are formulated as conceptual structures. An analogous way of formulating the argument of contraction would be the following: $C = \{y\}, O = \{x\}, I = \{\langle x, y \rangle\}, K = P = R = \emptyset$. However, it is not (always) necessary to eliminate the concept and the object in order to eliminate the instance-link. Hence, a well-formed argument for our contraction operation does not have the limitations we proposed for revision. An argument of contraction can be a part of a conceptual structure as well.

In order to formalize conceptual contraction, we introduce in our revision models a set-theoretical elimination operation \ominus :

Definition 8 Given that A and B are conceptual structures, $A \ominus B = A/B$, such that

- $\mathcal{C}_{A \ominus B} = \mathcal{C}_A / \mathcal{C}_B$
- $\mathcal{O}_{A \ominus B} = \mathcal{O}_A / \mathcal{O}_B$
- $K_{A \ominus B} = K_A / K_B$
- $P_{A \ominus B} = P_A / P_B$
- $R_{A \ominus B} = R_A / R_B$

- $I_{A \ominus B} = I_A / I_B$

Our elimination operation simply eliminates B from A . For simplicity, we define our elimination and contraction operations on conceptual structures. When the argument of the contraction is not a (complete) conceptual structure, the excluded elements of a conceptual structure will be regarded as empty. For instance, if an argument of a conceptual contraction A consists solely of an instance relation, then K_A , P_A , R_A , C_A and O_A are taken to be empty. Note that the elimination operation does not include taking the transitive closures of the resulting relations (as opposed to the expansion operation). Our contraction operation is defined as follows:

Definition 9 A conceptual revision and contraction model is a tuple $CS^{\oplus \ominus \leq} = \langle \mathcal{CS}, \oplus, \ominus, \leq \rangle$, such that

- $\langle \mathcal{CS}, \oplus \rangle$ is a fusion model on conceptual structures,
- \ominus is the set-theoretical elimination operation on conceptual structures, and
- \leq is a connected preorder on a set of conceptual structures.

We propose that during the contraction operation, the preference ordering on a set of conceptual structures changes as follows: let \mathcal{CS} be a non-empty set of conceptual structures, and $H \in \mathcal{CS}$ the argument of a contraction, and let \leq be the pre-contraction preference ordering on \mathcal{CS} and \leq' be the new preference ordering, then,

- for all $A, B \in \mathcal{CS}$, if $H \not\subseteq A$ and $H \subseteq B$ then $A \leq' B$ and $B \not\leq' A$, and if $H \subseteq B$ and $H \subseteq A$ then $B \leq' A$ and $A \not\leq' B$,
- otherwise, $A \leq' B$ iff $A \leq B$.

Note that, after the elimination operation, the remaining part of the conceptual structure does not have any substructures that include the argument of the contraction. Thus, when we focus on the relevant portion of the new preference ordering, that is the portion which orders the conceptual structures within the eliminated conceptual structure, the new preference ordering looks identical to the initial preference ordering, due to the second clause of our description.

Our contraction operation is quite radical, since we not only decrease the preference ordering of the conceptual structures that include the argument of the contraction, but we eliminate the argument from the new conceptual structures irrevocably. Just like in the revision case, this specific contraction operation is of course not the only possible one. In the belief revision literature, we can find more conservative ways of contracting information, such as operations that abstain from irrevocable elimination and give priority to parts of a structure that do not include or imply the argument of the contraction over the parts that do. One could also imagine more radical ways of reordering the preferences by a contraction operation, such as requiring that the parts of a structure that are consistent with the contracted information

are at most as preferred as the the parts which are not consistent with it⁹. The choice between these different contraction operations appears, analogously to the one between different revisions operations, strongly dependent on the specific views on scientific rationality that one favors. It is easy, in fact, to envisage philosophical criteria that would justify any of the alternatives above. Thus, we decided to stay as neutral as possible here, choosing the contraction operation most similar to our revision operation.

As in the case of revision, in our framework the outcome of a contraction operation on a conceptual structure ought to be a conceptual hierarchy. It should also be the case that contraction operation do not expand the contracted structures with novel relations, concepts or objects¹⁰. As we will see, even if nothing is added to a conceptual structure through contraction, the hierarchical structure may be lost. For instance, contracting a structure with respect to a kind-link may affect the transitivity of the kind-relation hence breaking the hierarchical structure. We restore the consistency of contracted conceptual structures as we did for revised structures. That is, we pick up the most preferred maximal conceptual hierarchies within the eliminated conceptual structure, according to the new preference ordering, and apply the intersection mechanism, just as described for revisions, until we obtain a conceptual hierarchy.

As an example of a contraction, consider the conceptual hierarchy H' , such that $C_{H'} = \{[mammal], [whale], [orka], [narwhale]\}$, $\mathcal{O}_{H'} = I_{H'} = P_{H'} = R_{H'} = \emptyset$
 $K_{H'} = \{\langle whale, mammal \rangle, \langle orka, mammal \rangle, \langle orka, whale \rangle,$
 $\langle narwhale, whale \rangle, \langle narwhale, mammal \rangle\}$.

Consider also the part of a conceptual structure A' : $C_{A'} = \{[orka], [narwhale]\}$, $K_{A'} = \{\langle narwhale, whale \rangle, \langle orka, whale \rangle, \langle orka, mammal \rangle, \langle narwhale, mammal \rangle\}$. Suppose we want to contract H' with respect to A' . We start with the simple elimination of A' from H' , obtaining $H' \ominus A'$, such that

$$C_{H' \ominus A'} = \{[mammal], [whale]\}, K_{H' \ominus A'} = \{\langle whale, mammal \rangle\},$$

$$\mathcal{O}_{H' \ominus A'} = R_{H' \ominus A'} = P_{H' \ominus A'} = I_{H' \ominus A'} = \emptyset.$$

The output of this particular contraction $H' - A'$ (Fig. 4) is equal to $H' \ominus A'$, aside from the changes in the preference ordering. This is because, the latter is a conceptual hierarchy, and the relevant portion of the new preference ordering is determined completely by the *otherwise* clause in our description of preference reordering by the contraction operation. If it were the case that $H' - A'$ is not a conceptual hierarchy, the consistency of the resulting conceptual structure would be recovered by iteratively intersecting the most preferred maximal conceptual hierarchies within $H' \ominus A'$, according to the new preference ordering, until reaching a conceptual hierarchy.

⁹To see the counterparts of these contraction operations, compare for instance, moderate contraction, conservative contraction and severe withdrawal in [30].

¹⁰This is another reason in favour of keeping the operation of adding relations or links to recover transitivity as part of the conceptual expansion. This way, we can use the exact process defined for revisions in order to retain consistency after conceptual contraction without making any additions to the conceptual structure.

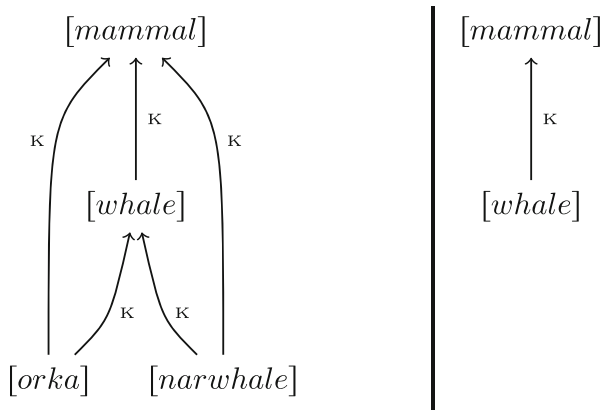


Fig. 4 The conceptual hierarchy H' (on the left) and the conceptual hierarchy $H' - A'$ (on the right)

Definition 10 Given a conceptual revision and contraction model $CS^{\oplus\ominus\leq}$, and given H is a conceptual structure and A is a (part of a) conceptual structure, H contracted with A (let us denote it with $H - A$, and the preference ordering after contraction with \leq') is determined by the following cases:

1. $H - A \equiv H' = \cap\{B : B \text{ is a maximal hierarchy within } H \ominus A \text{ and for all maximal hierarchies } I \subseteq H \ominus A \text{ it holds that } B \leq' I\}$, if H' constitutes a conceptual hierarchy;
2. $H - A \equiv H'' = \cap\{C : C \text{ is a maximal hierarchy within } H' \text{ and for all maximal hierarchies } I \subseteq H' \text{ it holds that } C \leq' I\}$, if H' does not constitute a conceptual hierarchy and H'' constitutes a conceptual hierarchy;
3. repeat case 2 with the maximal hierarchies within the resulting conceptual structure (e.g. starting with H'' in the first repetition) until reaching a conceptual hierarchy as the result of the intersection, if otherwise.

4 Rationality Postulates for Conceptual Change

In this section we will show how our conceptual revision models satisfy several rationality postulates analogous to the AGM ones for belief revision [1]. Since our system works at the conceptual level of abstraction, we cannot straightforwardly apply the AGM postulates to it. Thus, for each AGM revision postulate we will try to develop an analogous postulate at the conceptual level. We will also discuss rationality postulates for conceptual contraction, trying to comprehend the counterparts of the conceptual revision ones.

First, we show that a conceptual counterpart of the AGM closure and consistency postulates for revision is satisfied in our framework¹¹. We will call this first conceptual revision postulate the *hierarchy postulate*. This postulate amounts to the claim that a conceptual revision operation always results in a conceptual hierarchy. Recall in fact that for a conceptual structure to be a conceptual hierarchy, the information represented by the relations of the structure should not be contradictory. A conceptual hierarchy is furthermore closed, in the sense that none of the links needed for the transitive closures of the relations is missing. Hence, in our framework the consistency of a conceptual structure is intertwined with its completeness. Our framework satisfies this postulate thanks to the conjunction of the following properties: all conceptual structures have at least one maximal conceptual hierarchy as their substructure (due to their finiteness), the preference ordering always yields some minimal (most preferred) conceptual hierarchy (due to its connectedness), and a conceptual hierarchy can always be reached in finitely many iterations of our revision operation.

Theorem 11 *For all conceptual structures H, A , the product of revising H with A ($H * A$) is a conceptual hierarchy.*

Proof Let $CS^{\oplus\ominus\leq} = \{\mathcal{C}\mathcal{S}, \oplus, \ominus, \leq\}$ be a conceptual revision contraction model. We need to show that, for all conceptual structures H, A , after expanding H with A ($H \oplus A$) and reordering the preferences, we can obtain a conceptual hierarchy in finitely many steps, based on the definition of conceptual revision operation. To do this, we show that given an arbitrary conceptual structure, the operation of intersecting the most preferred maximal conceptual hierarchies within H (call this operation \cap^* and the result of it on H with $f(H)$) yields a conceptual hierarchy in finitely many iterations on H (that is, performing \cap^* on the resulting conceptual structure $f(H)$, and on $f(f(H))$, etc.).

(1) We start with showing that for all conceptual structures H , there is an $A \subseteq H$ s.t A is a maximal conceptual hierarchy within H . To show this, the following suffices:

- i) The empty conceptual structure (\emptyset) is a conceptual hierarchy. This directly follows from the definition of a conceptual hierarchy.
- ii) For all conceptual structures H , the empty conceptual structure is a substructure of H ($\emptyset \subseteq H$). This directly follows from the definition of a substructure.
- iii) For all conceptual structures H , it holds that H is finite. To show this, we define the size of a conceptual structure as follows: $\#H = \#C_H + \#O_H + \#K_H + \#P_H + \#R_H + \#I_H$ where $\#C_H$ and $\#O_H$ are to the number of elements in C_H and in O_H respectively. $\#K_H, \#P_H, \#R_H, \#I_H$ are the number of pairs in respective relations of H (i.e, we

¹¹Our consistency claim is stronger than what is required by the AGM consistency postulate. The AGM postulate assumes that the new information is not a contradiction

equate the size of a conceptual structure with the sum of the cardinalities of its elements). Since C_H and O_H are assumed to be finite by the definition of a conceptual structure, and since all relations in a conceptual structure are products of these domains and hence finite themselves, it follows that all conceptual structures are finite.

- (2) Next we show that for all conceptual structures H , there is an $A \subseteq H$, such that A is a maximal conceptual hierarchy within H and A is at least as preferred as all $B \subseteq H$, where B is a maximal conceptual hierarchy within H . It follows from (1) above that for all conceptual structures H , there are finitely many maximal conceptual hierarchies within H , and there is at least one maximal conceptual hierarchy within H . By the model assumptions for the preference ordering, namely by its connectedness, it holds that for all $A, B \subseteq H$ with A, B maximal conceptual hierarchies within H , it holds that either $A \leq B$ or $B \leq A$. These two facts establish our claim.
- (3) Lastly, we show that for all conceptual structures H , finitely many iteration of the operation \cap^* can be performed on H before reaching a conceptual hierarchy as the product ($f(H)$). To achieve this, we refer to the description of size of a conceptual structure from point (1) above, and prove the following:
 - iv) If S is the empty conceptual structure (\emptyset), $f(S)=S=\emptyset$. Since \emptyset is a conceptual hierarchy, \cap^* can be performed exactly $n = 1$ times before reaching a conceptual structure.
 - v) If S is a non-empty conceptual hierarchy, $f(S) = S$. This is because S is the unique maximal hierarchy within S (for all $A \subseteq S$ such that A is a hierarchy and $A \neq S$, it is possible to expand A into S , hence A is not maximal). Hence, \cap^* can be performed exactly $n = 1$ times before reaching a conceptual structure.
 - vi) If S is a conceptual structure, that is not a hierarchy, it holds that $f(S) \subset S$, hence $\#f(S) < \#S$. Since $\#S$ is finite, we can iterate the operation \cap^* on S at most $n = \#S$ times before reaching a conceptual hierarchy within S . To see this, let A, B the best preferred maximal conceptual hierarchies within S . Given that S is not a conceptual hierarchy, it holds that $A, B \subset S$. Then, $A \cap B \subset S$ by the definition of substructure. It follows that $\#(A \cap B) < \#S$. Since S is arbitrary, $\#f(S) < \#S$ holds for all such S .
 - vii) Since $\#S$ is finite, in $n \leq \#S$ iterations of \cap^* on S , we obtain a conceptual hierarchy, in the limit case we obtain the empty conceptual structure, with $\#\emptyset = 0$, which is a conceptual hierarchy.

Combining what is established above with the definition of our conceptual revision operation, we conclude that the product of a conceptual revision is always a conceptual hierarchy. □

Next, we show that our framework satisfies a *success postulate*, i.e. the claim that if the argument of a conceptual revision is a conceptual hierarchy, the argument

becomes a substructure of the revised conceptual structure. This postulate corresponds to a weakened version of the AGM success postulate for revisions¹². For the satisfaction of this postulate, it suffices that the argument of the revision is among the minimal conceptual structures in the (revised) preference ordering. This is achieved since our revision mechanism involves exactly this reordering of the preferences when revising a conceptual structure.

Theorem 12 *For all conceptual structures H and for all conceptual hierarchies A , $A \subseteq H * A$.*

Proof Let $CS^{\oplus\ominus\leq} = \{CS, \oplus, \ominus, \leq\}$ be a conceptual revision contraction model. Let H be conceptual structure and let A be a conceptual hierarchy. Suppose we revise H with A ($H * A$). By the definition of expansion, $A \subseteq H \oplus A$. Let H' be the conceptual structure $H \oplus A$ with the revised preference ordering in accordance with revision of H with A . It follows via (1) above, that there is a maximal conceptual hierarchy A' within H' , and since A is a conceptual hierarchy, it holds that $A \subseteq A'$. It follows from the description of reordering the preferences by the conceptual revision operation, that for all $B \subseteq H'$ such that B is a maximal hierarchy within H' , it holds that $B \leq' A'$ (after the revision) only if $A \subseteq B$. Thus, for all $B \subseteq H'$ that are picked up for intersecting, it holds that $A \subseteq B$. It follows that, $A \subseteq \bigcap \{B : B \text{ is a maximal hierarchy within } H' \text{ and for all maximal hierarchies } I \subseteq H' \text{ it holds that } B \leq' I\}$.

Similarly, for all $n \geq 1$ and for all B' such that $B' = f^n(H')$ (i.e., B' is the result of n -times iterating \cap^* on H'), it holds that $A \subseteq f^n(H')$. Since A itself is a conceptual hierarchy, it is not possible that $f^n(H') \subset A$. Therefore, $A \subseteq H * A$. \square

The third rationality postulate we consider is the *vacuity postulate*, i.e. the requirement that if the expansion of a conceptual structure is a conceptual hierarchy, this expansion is equal to the output of the revision process, aside from the reordering of the preference relation. This requirement corresponds to the vacuity postulate in the AGM theory and it is satisfied by our framework¹³.

Theorem 13 *For all conceptual structures H , A , it holds that $H * A \subseteq H \oplus A$ and $H \oplus A \subseteq H * A$ given that $H \oplus A$ is a conceptual hierarchy.*

¹²The AGM success postulate requires inclusion of the new belief without an antecedent that says it is a consistent belief. On the other hand, the success postulate required for base-generated beliefs by Rott [29] and Hansson [12] has that antecedent. We consider the weaker version of this postulate due to the strong consistency claim we established. Otherwise we have a contradiction saying the result of a conceptual structure is always consistent and if we revise a conceptual structure with a contradiction, the contradiction is part of the revised structure.

¹³In the AGM theory, one initially starts with a belief set and a preference ordering, yet as the result of revision or contraction, obtains a new belief set. The result of the change operations do not include a preference ordering. For this reason, while constructing our version, we state that the equality holds, aside from the changes in the preference ordering. The same reasoning applies to the vacuity postulate for contractions and to the recovery postulate, both states below.

Proof Let $CS^{\oplus\ominus\leq} = \{\mathcal{CS}, \oplus, \ominus, \leq\}$ be a conceptual revision contraction model. Let H, A be conceptual structures, and suppose $H \oplus A$ is a conceptual hierarchy. Let $(H \oplus A)'$ be $H \oplus A$ with the revised preference ordering. It holds that $H * A = f^n((H \oplus A)'),$ i.e. $H * A$ is the result of n -times iterating \cap^* on $(H \oplus A)'$. By (3) in theorem 15, it holds that $f(S) \subseteq S$ for all conceptual structures S . Therefore, $H * A \subseteq (H \oplus A)'$. Since $H \oplus A$ and $(H \oplus A)'$ are equivalent besides the preference ordering, it holds that $H * A \subseteq H \oplus A$. Note that this holds even when $H \oplus A$ is not a conceptual hierarchy.

The other direction holds since $H \oplus A$ is the unique maximal hierarchy within itself, and $f((H \oplus A)') = (H \oplus A)' = H * A$ as shown in (v) above. Since $H \oplus A$ and $(H \oplus A)'$ are equivalent besides the preference ordering, it holds that $H \oplus A \subseteq H * A$ \square

Next, we consider the *inclusion postulate*, i.e. the requirement that the outcome of a conceptual revision is a substructure of the expansion of the original conceptual structure with the argument of the revision. This postulate corresponds to the AGM inclusion postulate for revisions. This requirement makes sure that a conceptual structure is not expanded further than what is needed to consistently include the argument of the revision. This postulate is satisfied in our framework since all the steps of our revision operation involve only substructures of the expanded conceptual structure¹⁴.

Theorem 14 *For all conceptual structures H, A , it holds that $H * A \subseteq H \oplus A$.*

Proof See the proof for Vacuity above. Note however, the result of the operation $H \oplus A$ is different from the basic union of H and A ($H \cup A$) since conceptual expansion fixes transitivity of the relations. It does not hold that $H * A \subseteq H \cup A$. \square

After we mapped and analyzed some rationality postulates for conceptual revision framework, let us briefly discuss the corresponding contraction postulates. The first conceptual contraction postulate requires the result of a conceptual contraction to be a conceptual hierarchy. Since conceptual contraction involves the same consistency-recovery mechanism of conceptual revision, this principle is satisfied for reasons analogous to the revision case.

Theorem 15 *For all conceptual structures H and A (A might also be a part of a conceptual structure), the product of contracting H with A ($H \ominus A$) is a conceptual hierarchy.*

¹⁴It should be noted that there are three other basic AGM rationality postulates we did not discuss here. One is the extensionality postulate which states that revision of a belief set with classically logically equivalent arguments lead to logically equivalent revised belief sets. Since we did not comment on identity principles concerning the conceptual structures, we cannot map this requirement to our framework for now. The other two postulates are about revisions with conjunctions. We do not consider these as relevant for our current conceptual revision framework, since we did not discuss relations between structures which would correspond to logical connectives.

Proof We need to show that, for all conceptual structures H , after eliminating a conceptual structure (or a part of a conceptual structure) A from H ($H - A$), we can obtain a conceptual hierarchy in finitely many steps, based on the definition of conceptual contraction operation. The proof is same as theorem 15 above. \square

The success postulate for conceptual contraction requires the argument of the contraction (a conceptual structure or a part of one) to not be a substructure of the the result of the contraction. A weaker version of this principle, which limits the argument of the contraction to non-empty conceptual structures or their parts, is satisfied in our framework. This is because, once the argument of the contraction is deleted from the initial conceptual structure, nothing is added to the resulting structure while rebuilding consistency.

Theorem 16 *For all conceptual structures H and for all non-empty conceptual structures A (A might also be a part of a conceptual structure), it holds that $A \not\subseteq H \ominus A$.*

Proof Let $CS^{\oplus\ominus\leq} = \{CS, \oplus, \ominus, \leq\}$ be a conceptual revision contraction model. Let H be a conceptual structure and let A be a non-empty (part of a) conceptual structure. By definition of conceptual elimination we know that $A \not\subseteq H - A$. By the definition of contraction, we have $H \ominus A = f^n(H - A)$, according to the new preference ordering. By (3) in theorem 15, it holds for all conceptual structures S , $f(S) \subseteq S$, hence $H \ominus A \subseteq H - A$. It follows that $A \not\subseteq H \ominus A$. \square

We can also show that the operation of contracting a conceptual structure does not expand the initial conceptual structure in any way. This is a counterpart of the AGM inclusion postulate for contractions.

Theorem 17 *For all conceptual structures H and A (A might also be a part of a conceptual structure), it holds that $H \ominus A \subseteq H$.*

Proof Let $CS^{\oplus\ominus\leq} = \{CS, \oplus, \ominus, \leq\}$ be a conceptual revision contraction model. Let H be a conceptual structure and let A be a (part of a) conceptual structure. By the conceptual elimination operation we defined in definition 10, it holds that $H - A \subseteq A$. Since our contraction operation functions via intersections (\cap^*), and with (3) in theorem 15 (for all conceptual structures S , it holds that $f(S) \subseteq S$), it holds that $H \ominus A \subseteq A$. \square

The vacuity postulate for conceptual contraction states that, if the argument of the contraction does not occur in the initial conceptual structure, then no changes are made to this structure. In our framework, this requirement is not satisfied, since it is possible that the initial conceptual structure changes in the process of consistency-recovery. A weaker version of this requirement, assuming that the initial conceptual structure is a conceptual hierarchy, is however satisfied, aside from the reordering of the preferences.

Theorem 18 *For all conceptual hierarchies H and for all conceptual structures A (A might also be a part of a conceptual structure), if $A \not\subseteq H$ then $H \ominus A \subseteq H$ and $H \subseteq H \ominus A$.*

Proof Let $CS^{\oplus\ominus\leq} = \{CS, \oplus, \ominus, \leq\}$ be a conceptual revision contraction model. Let H be a conceptual hierarchy and let A be a (part of a) conceptual structure. Suppose $A \not\subseteq H$. It follows that $H - A = H$. Since H is the unique maximal conceptual hierarchy in H , it holds that $f(H - A) = H$ see point (v) in theorem 15 above. Thus, $H \ominus A$ differs from H only in terms of the preference ordering between conceptual structures in the model. Therefore, our claim holds.

Let us also show that this does not hold unless the initial structure is a conceptual hierarchy. Let H be a conceptual structure and let A be a (part of a) conceptual structure. Suppose $A \not\subseteq H$. It follows that $H - A = H$. Given that H is not a conceptual hierarchy, intersecting the most preferred maximal hierarchies within H yields a proper substructure of H . Hence, it holds that $f(H \ominus A) \subset H$, see point (vi) in theorem 15 above. \square

Next, we consider a counterpart of the AGM recovery postulate for contractions. Our version of the postulate requires that the result of the contraction operation is such that, if it is expanded with the argument of the contraction, the initial conceptual structure is recovered. As we will see this requirement does not hold in our framework. This should not come as a surprise. Counterparts of this principle have in fact been widely rejected in alternative theories of belief revision, including base-generated revisions [12] and belief withdrawals [19], and in application of such belief revision strategies to scientific change, for example in abductive belief revision in science [31]. In order for the recovery principle to be satisfied, a rational change theory should first and foremost adhere to the principle of minimal change. In particular, the principle states that so much information is retained in a contracted theory, up to the point that the initial theory can be recovered by a simple expansion with the argument of the contraction. It is often argued that there are other important rationality postulates, concerning the preferences of the agents, which should not be overwhelmed by the minimal change principle (cf. [28]). The failure of this postulate is then again in line with our stance concerning the revolutionary nature of scientific change modeled in this paper. In our framework, given that a conceptual structure H is contracted by a conceptual structure (or a part of a conceptual structure) A , it might be the case that $H \ominus A \subset H - A$ (see theorem 15), and expanding $H \ominus A$ with A does not necessarily recover the information that is lost in the transition from $H - A$ to $H \ominus A$.

Theorem 19 *Given that H is a conceptual structure and A is a (part of a) conceptual structure, it might not hold that $(H \ominus A) \oplus A \subseteq H$ and $H \subseteq (H \ominus A) \oplus A$.*

Proof We show with a counterexample. Let $CS^{\oplus\ominus\leq} = \{CS, \oplus, \ominus, \leq\}$ be a conceptual revision and contraction model, with \leq such that for all $A, B \in CS$, $A \leq B$ iff A includes the instance pair $\langle Bob, mammal \rangle$. Let $H = \langle C_H = \{\{orka\}, [mammal]\}, O_H = \{(Bob)\}, K_H = \{\{orka, mammal\}\}, I_H = \{\{Bob, orka\},$

$\langle Bob, mammal \rangle \rangle, P_H = R_H = \emptyset \rangle$ be a conceptual structure in \mathcal{CS} . Let $A = \{\langle orka, mammal \rangle\}$ be a part of a conceptual structure in \mathcal{CS} . Thus, $H \ominus A = \langle C_{H-A} = C_H, O_{H-A} = O_H, K_{H-A} = \emptyset, I_{H-A} = \{\langle Bob, mammal \rangle\}, P_{H-A} = P_H, R_{H-A} = R_H \rangle$.

On the other hand, $H' = (H \ominus A) \oplus A = \langle C_{H'} = C_H, O_{H'} = O_H, K_{H'} = K_H, I_{H'} = \{\langle Bob, mammal \rangle\}, P_{H'} = R_{H'} = \emptyset \rangle$. Note that $H \not\subseteq H'$ since the instance pair $\langle Bob, orka \rangle$ is included in H but not in H' . Therefore, our claim holds. \square

Lastly, we will mention the *Levi-identity*, which reduces AGM belief revision to an operation of AGM belief contraction followed by an expansion. In particular, Levi-identity says that revision of a theory H with a piece of information A can be performed by first contracting H with the negation of A , hence making space for consistent incorporation of A , and then expanding this contracted theory with A . In our framework, an analogous identification of conceptual revision with a sequence of conceptual contraction and expansion is not possible. This can be shown with the help of a simple informal example. Suppose we want to revise a conceptual structure that includes the kind-relation representing the information that orkas are fishes, with the new, contrasting, kind-relation expressing the information that orkas are mammals. Assuming the Levi-identity, we would need to first remove from the initial theory whatever contradicts with this latter kind-relation. This might include not only the kind-relation between orkas and fishes, but also (possibly) other related parts of a conceptual structure, such as rules relations encoding important information about fishes. This additionally contracted parts are not necessarily recovered by an expansion operation¹⁵.

This alleged failure of Levi-identity in our framework should not be surprising. Similar to the case for the recovery postulate above, while contracting a conceptual structure, we may end up eliminating more information than what is required for the consistent incorporation of the new information. Moreover, in the scientific context, it can be argued that Levi-identity should not hold. For instance, Schurz [31] states that, in the context of his abductive belief revision framework, combining ordinary belief expansion and abduction generation based on a contracted theory does not describe abductive revision in science. This is because the information provided by the initial theory that gets lost in the contraction is not necessarily recovered in this way. The same rationale explains the failure of (an alleged translation of) the Levi-identity in our model of conceptual revision.

5 Taking up Thagard's Challenge

We have now presented our conceptual revision model and we have shown how our revision and contraction operations satisfy several rationality postulates for conceptual change. In this section, we will show how we can mirror the dynamics of

¹⁵We conclude our discussion of the failure of the Levi-identity here, since a formal counterexample requires formal tools we have not introduced, such as negating a conceptual structure or a part of it.

Thagard's conceptual systems in our system. Specifically, we will demonstrate how almost every kind of change described by Thagard can be adequately represented in our framework via a suitable (combination of) change operation(s) on conceptual structures.

5.1 Mirroring Thagard's kinds of Changes in our Conceptual Revision Model

As we saw in Section 2, Thagard described a hierarchy of nine degrees of changes applicable to conceptual systems, ordered by their increasing strength: instance-addition, rule-addition, part-addition, kind-addition, concept-addition, kind-collapse, hierarchy-reorganization, and tree-switching.

In what follows, we will discuss each of these degrees of change one by one, from the weakest to the most radical one. With the exception of tree-switching, whose case will be completely different from all the others, the structure of our discussion will take the following form. We will first present how a given kind of change operates on one of Thagard's conceptual systems. Then, we will explain informally how this kind of change can be represented in our framework. After that, we will give a formal definition of the degree of change under focus, showing how it can be seen as a special case of (a series of applications of) our revision and/or our contraction operations. Finally, we will present a toy-example of this kind of change in our framework in order to make clearer our proposed formalization.

Instance-addition The addition of an instance-link is the least radical kind of change described by Thagard. It consists in the addition of a single instance link between one object node and one conceptual node of a given conceptual system, representing the information that a given individual is an instance of a given concept.

In our framework, we can mirror instance-addition via our conceptual revision operation, revising a given conceptual structure with another conceptual structure that includes a non-empty instance-relation. In particular, we can define three different forms of instance-addition as three different constraints on the argument of revision. The most general form, what we will call general instance-addition, consists of requiring the argument of the revision to include a non-empty instance relation. A more specific form of instance-addition, i.e. pure instance-addition, requires the argument of the revision to have instance-relation as its only non-empty relation (concept and object domains can be non-empty as well). Finally, we have an atomic instance-addition when the argument of the revision of a pure instance-addition has a single instance-pair as its instance-relation. This last form corresponds to (our interpretation of) Thagard's understanding of instance-addition.

More formally, a conceptual revision operation $H * A$ represents a *general instance-addition* iff $I_A \neq \emptyset$. A conceptual revision operation $H * A$ represents a *pure instance-addition* iff $I_A \neq \emptyset$ and $K_A = P_A = R_A = \emptyset$. A conceptual revision operation $H * A$ represents an *atomic instance-addition* iff $|I_A| = 1$ and $K_A = P_A = R_A = \emptyset$. For an example of a general instance-addition, see the conceptual revision example presented in Section 3.1.

Rule-addition The second kind of change described by Thagard consists in adding a rule-link between two concepts nodes of a given conceptual system. This change represents adding the information that a generic holds between two concepts.

In our framework, rule-addition is represented similarly as we treated instance-addition, i.e. by requiring the argument of our revision operation to include a non-empty rule-relation. As in the previous case, three different forms of rule-addition can be defined, differing in terms of generality: general rule-addition, pure rule-addition, and atomic rule-addition.

More formally, a conceptual revision operation $H * A$ represents a *general rule-addition* iff $R_A \neq \emptyset$. A conceptual revision operation $H * A$ represents a *pure rule-addition* iff $R_A \neq \emptyset$ and $K_A = P_A = I_A = \emptyset$. A conceptual revision operation $H * A$ represents an *atomic rule-addition* iff $|R_A| = 1$ and $K_A = P_A = I_A = \emptyset$.

As a simple example of rule-addition, let H be composed by:

$$\begin{aligned} C_H &= \{[mammal], [whale], [orka]\}, O_H = I_H = P_H = R_H = \emptyset \\ K_H &= \{\langle whale, mammal \rangle, \langle orka, mammal \rangle, \langle orka, whale \rangle\}. \end{aligned}$$

and let A be composed by $C_A = \{[mammal], [air]\}$, $O_A = K_A = P_A = I_A = \emptyset$, $R_A = \{\langle mammal, air \rangle\}$ (intuitive interpretation: mammals breath air)¹⁶. The output of this revision operation expands the rule-relation of H with R_A and the pairs $\langle whale, air \rangle, \langle orka, air \rangle$. We then have $H \oplus A = M$ where:

$$\begin{aligned} C_M &= \{[mammal], [whale], [orka], [air]\}, O_M = P_M = I_M = \emptyset \\ K_M &= \{\langle whale, mammal \rangle, \langle orka, whale \rangle, \langle orka, mammal \rangle\} \\ R_M &= \{\langle mammal, air \rangle, \langle whale, air \rangle, \langle orka, air \rangle\}. \end{aligned}$$

Since M is a conceptual hierarchy, we have $H * A = M$, modulo the revised preference ordering (Fig. 5).

Part-addition The third kind of change described by Thagard is called part-addition or decomposition. It consists in adding a part-link between two concept nodes of a given conceptual system, representing the information that a relation of part-hood holds between the concepts denoted by these nodes.

In our framework, part-addition is represented similarly as we treated instance-addition and rule-addition, i.e. by requiring the argument of our revision operation to include a non-empty part-relation. As in the previous cases, three different forms of part-addition can be defined, differing in terms of generality: general part-addition, pure part-addition, and atomic part-addition.

More formally, a conceptual revision operation $H * A$ represents a *general part-addition* iff $P_A \neq \emptyset$. A conceptual revision operation $H * A$ represents a *pure part-addition* iff $P_A \neq \emptyset$ and $K_A = R_A = I_A = \emptyset$. A conceptual revision operation $H * A$ represents an *atomic part-addition* iff $|P_A| = 1$ and $K_A = R_A = I_A = \emptyset$.

¹⁶Note that it would be possible in our framework to differentiate rules in terms of their intended interpretation, so that for instance the rule *breath* is represented differently from other rules (e.g. *swim*) that may be added to a given conceptual structures. We decided to follow Thagard in leaving the interpretation of the rules outside our framework, considering all rules as uninterpreted rule-pairs.




Fig. 5 The output of the rule-addition example (on the left) and the output of the part-addition example (on the right)

As a simple example of part-addition, take H to be such that:

$$C_H = \{[mammal], [whale], [orka]\}, O_H = P_H = R_H = I_H = \emptyset$$

$$K_H = \{\langle whale, mammal \rangle, \langle orka, mammal \rangle, \langle orka, whale \rangle\}.$$

Let A be composed by $C_A = \{[mammal], [lungs]\}$, $P_A = \{\langle mammal, lungs \rangle\}$ (intuitive interpretation: mammals have lungs), and $K_A = R_A = I_A = O_A = \emptyset$. The output of this revision operation expands the part-relation of H with P_A and the pairs $\langle whale, lungs \rangle, \langle orka, lungs \rangle$, in order to recover transitivity. We then have $H \oplus A = M$ where:

$$C_M = \{[mammal], [whale], [orka], [lungs]\}, O_M = R_M = I_M = \emptyset$$

$$K_H = \{\langle whale, mammal \rangle, \langle orka, mammal \rangle, \langle orka, whale \rangle\}$$

$$P_M = \{\langle mammal, lungs \rangle, \langle whale, lungs \rangle, \langle orka, lungs \rangle\}.$$

Since M is a conceptual hierarchy, we have $H * A = M$ (Fig. 5).

Kind-addition The fourth kind of change described by Thagard consists in adding a kind-link between two concept nodes of a given conceptual system, representing the information that a relation of kind-hood holds between the concepts denoted by these nodes. Furthermore, Thagard, following Carey’s terminology for conceptual change in child psychology [5], distinguishes two special cases of (series of) kind-addition(s): *coalescence* and *differentiation*. The former type of kind-addition happens when we add a superordinate conceptual node linked via a series of kind-links with some concept nodes that had no superordinate kinds before. The latter denotes instead the addition of some subordinate conceptual nodes connected via a series of kind-links with a conceptual node that before had no subordinate kinds.

In our framework, kind-addition is represented by requiring the argument of our revision operation to include a non-empty kind-relation. As in the previous cases, three different forms of kind-addition can be defined, differing in terms of generality: general kind-addition, pure kind-addition, and atomic kind-addition. Coalescence and differentiation can then be represented as specific cases of general or pure kind-addition.

Formally, a conceptual revision operation $H * A$ represents a *general kind-addition* iff $K_A \neq \emptyset$. A conceptual revision operation $H * A$ represents a *pure kind-addition* iff $K_A \neq \emptyset$ and $P_A = R_A = I_A = \emptyset$. A conceptual revision operation $H * A$ represents an *atomic kind-addition* iff $|K_A| = 1$ and $P_A = R_A = I_A = \emptyset$. Furthermore, a general or pure kind-addition $H * A$ is a *coalescence* iff there exists a $x \in C_A$ such that $\langle y, x \rangle \in K_A$ and there is no w such that $\langle y, w \rangle \in K_H$. A general or pure kind-addition $H * A$ is instead a *differentiation* iff there is an $x \in C_A$ such that $\langle x, y \rangle \in K_A$ and there is no w such that $\langle w, y \rangle \in K_H$. For an example of a general kind-addition, see the conceptual revision example presented in Section 3.1.

Concept-addition The fifth kind of change described by Thagard consists in adding a new concept node to a given conceptual system. This type of change represents the addition of a new concept to a given scientific theory¹⁷.

In our framework, concept-addition is represented by requiring the argument of our revision operation to include a new concept. Several further restrictions can be imposed. For instance, we present here two more specific forms of concept-addition: unique concept-addition and connected concept-addition. We have a unique concept-addition when there is only one new concept in the argument of the revision (it may also include non-empty relations). We have a connected-concept addition when each new concept in the argument figures in at least one relation.

Formally, a conceptual revision operation $H * A$ is a *concept-addition* iff there is an $x \in C_A$ such that $x \notin C_H$. A concept-addition $H * A$ is then a unique concept-addition iff there is only one $x \in C_A$ such that $x \notin C_H$. A concept-addition $H * A$ is then a connected concept-addition iff for all $x \in C_A$ such that $x \notin C_H$ there exists a $y \in C_A \cup O_A$ such that $\langle x, y \rangle$ or $\langle y, x \rangle$ is in $K_A \cup P_A \cup R_A \cup I_A$. For an example of a unique and connected concept-addition see the conceptual revision example in Section 3.1, the rule-addition example, or the part-addition example above.

Kind-collapse The sixth change described by Thagard is kind-collapse, i.e. the removal of a (series of) kind-link(s) from a given conceptual system. More specifically, Thagard says that kind-collapse is the inverse change of differentiation, so that kind-collapse denotes removing all subordinate kinds of a given conceptual node.

In our framework, kind-collapse is a specific case of our contraction operation, namely, the contraction of a given conceptual structure with respect to a set of kind-pairs all of which have the same element as their second element and such that in the contracted structure this element has no subordinate kinds.

Formally, a conceptual contraction operation $H - A$ is a *kind-collapse* iff $\exists x \in C_H$ such that $K_A = \{\langle j_1, x \rangle, \dots, \langle j_n, x \rangle\}$ and $\neg \exists y \in C_H \cup C_A$ such that $\langle y, x \rangle \in K_{H-A}$. This definition of a kind-collapse makes it the inverse process of a differentiation, just like in Thagard's system. For an example of a kind-collapse, see the contraction example in Section 3.2.

¹⁷Thagard also stresses how concept-addition sometimes involves combining two simple concepts into a complex one [36, pp. 35-36]. This combination aspect of concept-addition is outside the scope of the present version of our framework, since we assumed for simplicity that the concept universe is constant.

Hierarchy-reorganization The seventh kind of change in Thagard's theory is the general process of hierarchy-reorganization or *branch-jumping*, i.e. moving a set of concept and object nodes from one part of a conceptual system to another one, thus changing (some of) their relations. This change is typical of many scientific revolutions, such as the Copernican revolution in which the earth branch-jumped from being a unique entity to a kind of planet.

In our framework branch-jumping is a specific series of our contraction and revision operations that does not involve changes to the concept-domains of the conceptual structures involved. The output of such combination is the transportation of certain parts of a given conceptual structure to a different part of it, involving some change in its relations.

Formally, we say that the sequence of contraction and revision operations $(H - A_1) * A_2$ represents a *hierarchy-reorganization* iff $\mathcal{C}_H = \mathcal{C}_{(H-A_1)*A_2}$, $\mathcal{O}_H = \mathcal{O}_{(H-A_1)*A_2}$ and either $K_H \neq K_{(H-A_1)*A_2}$ or $P_H \neq P_{(H-A_1)*A_2}$ or $R_H \neq R_{(H-A_1)*A_2}$ or $I_H \neq I_{(H-A_1)*A_2}$. Note that we leave completely open how the relations between the objects and concepts involved in this type of change are transformed. Specific kinds of hierarchy-reorganization, such as part-kind transformation, can then be defined by imposing further constraints on the relations in the contraction and in the revision operation.

As an example of a hierarchy-reorganization, take H to be such that:

$$\begin{aligned} \mathcal{C}_H &= \{[animal], [fish], [mammal], [whale], [orka]\}, \\ \mathcal{O}_H &= I_H = P_H = R_H = \emptyset \\ K_H &= \{\langle whale, fish \rangle, \langle orka, fish \rangle, \langle orka, whale \rangle, \langle orka, animal \rangle, \\ &\quad \langle whale, animal \rangle, \langle mammal, animal \rangle, \langle fish, animal \rangle\}. \end{aligned}$$

Let A_1 be the part of a conceptual structure $\{\langle whale, fish \rangle, \langle orka, whale \rangle, \langle orka, fish \rangle\}$ and A_2 be composed by $\mathcal{K}_{A_2} = \{\langle whale, mammal \rangle, \langle orka, whale \rangle\}$, $\mathcal{C}_{A_2} = \{[whale], [orka], [mammal]\}$, $\mathcal{O}_{A_2} = \{\emptyset\}$, and $P_{A_2} = R_{A_2} = I_{A_2} = \emptyset$.

The output of the hierarchy-reorganization $H - A_1 * A_2$ is then equal to the structure H' (Fig. 6) where:

$$\begin{aligned} \mathcal{C}_{H'} &= \{[mammal], [whale], [orka], [fish], [animal]\}, \\ \mathcal{O}'_H &= R'_H = P'_H = I'_H = \emptyset \\ K'_H &= \{\langle whale, mammal \rangle, \langle orka, mammal \rangle, \langle orka, whale \rangle, \langle orka, animal \rangle, \\ &\quad \langle whale, animal \rangle, \langle mammal, animal \rangle, \langle fish, animal \rangle\}. \end{aligned}$$

Tree-switching The last change described by Thagard is tree-switching, i.e. the change of the organizing principle of the whole hierarchy. This change implies thus re-interpreting any kind-relation and part-relation. An example of this kind of change is the Darwinian revolution, a revolution that involved the re-interpretation of kind-relations of biological entities as historical kinship and not as they were before as morphological similarities. This is the most radical change that can happen in science for Thagard, up to the point that it is sufficient but not necessary for having a

conceptual revolution. Only certain scientific revolutions that are particularly radical exemplify tree-switching.

Since tree-switching is not really about changing the structure of a conceptual system, focusing instead on the external interpretation of the conceptual system, it would be at least unclear how to frame this kind of change in our framework. Using an epistemological metaphor, modeling tree-switching in our framework would be like implementing a gestalt-operation in traditional belief revision that changes the meaning of the logical consequence between beliefs. We therefore do not treat this kind of change in the present work, focusing only on the first eight changes that affect the internal-structure of conceptual system, confident that we do not lose too much in generality, since as Thagard himself acknowledges many scientific revolutions do not even exemplify tree-switching.

6 Conclusion

Let us recap the main steps of the present work. Starting from Thagard's model of scientific conceptual change, we saw his taxonomy of nine degrees of conceptual change and his claim that belief revision theories can only account for the first two of them. We then presented our system of conceptual revision, i.e. a belief-revision-like system for conceptual structures. We showed how our conceptual revision and contraction operations satisfy several rationality postulates analogous to the AGM ones. We then demonstrated how our system, working at the conceptual level of abstraction, is able to mirror eight out of nine kinds of conceptual changes described by Thagard.

More generally, our framework shows how belief revision theories can be mapped to the conceptual level in order to obtain a logical interpretation of radical conceptual change. The present work is only a first step towards a better understanding of the relationships between belief change and conceptual change. Several directions of future work naturally present themselves. Interesting ways of extending our framework include investigating specific preference orderings and alternative ways of changing them, reconstructing case studies from the history of science as series of conceptual revision and contractions, working with expanding domains to




Fig. 6 The input (on the left) and the output (on the right) of the hierarchy-reorganization example

model conceptual combination, adding the possibility of revising conceptual structure with complex information (such as negative one, for instance) to further model logical relationships between elements of a conceptual structure, having a way of comparing differing conceptual structures in order to model Thagard's explanatory coherence notion, and also augmenting our conceptual structures in order to mimic more elaborate approaches to theory-change (e.g. [2, 4, 14, 15, 21]). These extensions would allow to model even Thagard's most radical type of conceptual change, i.e. tree-switching. It would also be interesting to merge conceptual structures with (structured) belief sets, in order to have a revision system capable of revising beliefs and concepts at the same time. Such a conceptual-plus-belief-revision system would be able to model (some of) the interesting connections between conceptual change and belief change, thereby offering a more fine-grained logical reconstruction of scientific change.

Acknowledgments We are indebted to Hannes Leitgeb, Diego Tajer, and an anonymous reviewer for helpful comments on previous drafts of this paper.

Funding Open Access funding enabled and organized by Projekt DEAL. This research was funded by the LMU Munich Institutional Strategy LMU excellent within the framework of the German Excellence Initiative.

Declarations

Conflict of Interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2), 510–530.
2. Andersen, H., Barker, P., & Chen, X. (2006). *The cognitive structure of scientific revolutions*. Cambridge: Cambridge University Press.
3. Andreas, H. (2011). A structuralist theory of belief revision. *Journal of Logic Language and Information*, 20(2), 205–232.
4. Balzer, W., Moulines, C. U., & Sneed, J.D. (1987). *An Architectonic for Science: the structuralist program*. Dordrecht: D. Reidel Publishing Company.
5. Carey, S. (1985). *Conceptual change in childhood*. Cambridge: MIT Press.
6. Cresto, E. (2008). A model for structural changes of belief. *Studia Logica*, 88, 431–451.
7. Enqvist, S. (2011). A Structuralist Framework for the Logic of Theory Change. In E. J. Olsson, & S. Enqvist (Eds.) *Belief Revision meets Philosophy of Science*, (Vol. 21 pp. 105–136). Heidelberg: Springer.
8. Friedman, M. (2001). *Dynamics Of reason: The 1999 Kant Lectures at Stanford University*. Stanford: CSLI Publications.

9. Gärden­schlag, T., et al. (2013). *Frames and Concept Types: applications in language and philosophy. Studies in Linguistic and Philosophy* (94). Berlin: Springer Science & Business Media.
10. Gärden­fors, P., & Rott, H. (1995). Belief Revision. In D. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.) *Handbook of logic in artificial intelligence and logic programming, (Vol. 4): epistemic and temporal reasoning* (pp. 35–132). Oxford University Press.
11. Hansson, B. (2011). A Note on Theory Change and Belief Revision. In E. J. Olsson, & S. Enqvist (Eds.) *Belief Revision meets Philosophy of Science, (Vol. 21 pp. 155–162)*. Heidelberg: Springer.
12. Hansson, S. O. (1999). *A Textbook of Belief Dynamics: theory change and database updating, Applied Logic Series (11)*, Kluwer Academic Publishers.
13. Hansson, S. O. (2011). Changing the Scientific Corpus. In E. J. Olsson, & S. Enqvist (Eds.) *Belief Revision meets Philosophy of Science, (Vol. 21 pp. 43–59)*. Heidelberg: Springer.
14. Kormmesser, S., & Schurz, G. (2018). *Analyzing Theories in the Frame Model*. Erkenntnis. <https://doi.org/10.1007/s10670-018-0078-5>.
15. Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago: Chicago University Press.
16. Kuhn, T. S. (1977). Objectivity, value judgment, and theory choice. In *Thomas kuhn, the essential tension: Selected studies in scientific tradition and change* (pp. 320–339). Chicago University Press.
17. Lakatos, I. (1970). Falsification and the Methodology of Scientific Research Programmes. In I. Lakatos, & A. Musgrave (Eds.) *Criticism and the Growth of Knowledge* (pp. 91–195). Cambridge: Cambridge University Press.
18. Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton: Princeton University Press.
19. Makinson, D. (1987). On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 383–394.
20. Martin, E., & Osherson, D. (1998). Belief revision in the service of scientific discovery. *Mathematical Social Sciences*, 36(1), 57–68.
21. Masterton, G., Zenker, F., & Gärden­fors, P. (2017). Using conceptual spaces to exhibit conceptual continuity through scientific theory change. *European Journal for Philosophy of Science*, 7(1), 127–150.
22. Minsky, M. (1975). A framework for representing knowledge. In P. H. Winston (Ed.) *The Psychology of Computer Vision* (pp. 211–277). New York.
23. Olsson, E. J., & Enqvist, S. (2011). *Belief Revision meets Philosophy of Science, Vol. 21*. Heidelberg: Springer.
24. Park, W. (2010). Belief revision vs. conceptual change in mathematics. In *Model-Based Reasoning in Science and Technology* (pp. 121–134). Berlin: Springer.
25. Popper, K. R. (1963). *Conjectures and refutations: The growth of scientific knowledge*. London: Routledge.
26. Ribeiro, M. M., & Wassermann, R. (2007). Base revision in description logics-preliminary results. In *Proceedings of the International Workshop on Ontology Dynamics (IWOD-07)* (vol. 6982).
27. Ribeiro, M. M., & Wassermann, R. (2009). AGM revision in description logics. *Proceedings of ARCOE*.
28. Rott, H., & Pagnucco, M. (1999). Severe withdrawal (and recovery). *Journal of Philosophical Logic*, 28(5), 501–547.
29. Rott, H. (2001). *Change, choice and inference: A study of belief revision and nonmonotonic reasoning (No. 42)*. Clarendon Press.
30. Rott, H. (2006). Shifting Priorities: Simple Representations for Twenty-seven Iterated Theory Change Operators. In D. Makinson, J. Malinowski, & H. Wansing (Eds.) *Towards Mathematical Philosophy* (pp. 269–296). Dordrecht: Springer.
31. Schurz, G. (2011). Abductive belief revision in science. In E. J. Olsson, & S. Enqvist (Eds.) *Belief Revision meets Philosophy of Science, (Vol. 21 pp. 77–104)*. Heidelberg: Springer.
32. Strößner, C. (2020). Predicate change: a study on the conservativity of conceptual change. *Journal of Philosophical Logic*, 49, 1159–1183.
33. Thagard, P. (1984). Frames, knowledge, and inference. *Synthese*, 61, 233–259.
34. Thagard, P. (1988). *Computational philosophy of science*. Cambridge: MIT Press.
35. Thagard, P. (1990). The conceptual structure of the chemical revolution. *Philosophy of Science*, 57(2), 183–209.
36. Thagard, P. (1992). *Conceptual revolutions*. Princeton: Princeton University Press.

37. Thagard, P. (2000). *Coherence in Thought and Action*. MIT Press.
38. Thagard, P., & Findlay, S. (2011). Changing Minds About Climate Change: Belief Revision, Coherence, and Emotion. In E. J. Olsson, & S. Enqvist (Eds.) *Belief Revision meets Philosophy of Science*, (Vol. 21 pp. 329–345). Heidelberg: Springer.
39. Wolter, F., & Zakharyashev, M. (1999). Multi-dimensional description logics. *IJCAI*, 99, 104–109.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.