A General Framework for Shape Similarity

Stefano Borgo[1] and Oliver Kutz[2]

[1] Laboratory of Applied Ontology (ISTC-CNR), Trento & KRDB (FUB), Bolzano, Italy
borgo@loa-cnr.it
[2] Research Center on Spatial Cognition (SFB/TR 8), University of Bremen, Germany
okutz@informatik.uni-bremen.de

Abstract. We propose a formal framework to clarify the import of and the relationship among different notions of similarity. The framework, borrowing from the general $\mathcal{E}$-connections approach, allows to use multiple similarity notions in the same system by relating the domain of discourse to specialized sets of entities even across different logics. In practice, we start form the assumption that (Leibnizian) relative identity is the only *local* form of similarity, and show that more sophisticated notions can be conceptually understood as the result of transformations across heterogeneous logics. These transformations, called *global* similarities, are usually motivated by arguments external to the starting system, and can be intertwined to generate a prolific family of similarity relations. We exemplify this framework in the context of shape similarity whose comprehension motivated our work. Finally, the approach developed in this paper sets the basis for a general study of similarity as a group of transformations generated by (primitive) global similarity relations.

Key words: Similarity, Shapes, Counterparts, $\mathcal{E}$- and $\mathcal{S}$-Connection

1 Introduction

The search for similarities across objects is pervasive in human thought and reasoning, and similarity analysis is a main cognitive tool for the classification of entities. As Quine puts it, the notions of kind, similarity and resemblance "seem to be variants or adaptations of a single notion" and "[...] a sense of similarity or of kinds is fundamental to learning in the widest sense—to language learning, to induction, to expectation" [18, p. 129]. Since there are different ways in which things can be said to be similar, similarity analysis makes room for reasoning according to different perspectives. It thus has a crucial role even in our scientific understanding of the world, as L. R. Goldstone and colleagues observe "An account of what make problems, memories, objects, and words similar to one another often provides the backbone for our theories of problem solving, attention, perception, and cognition." [6].

In this paper we look at similarity in the domain of objects with spatial properties, and in particular we focus on shape. Research on shape is carried out for a variety of reasons; from mathematical aspects on how to model shape [15] to qualitative issues on how to describe shape [2], from the neuro-psychological problem of how the mind processes shape [3] to the computational issues of shape classification and handling [23], from shape detection [16] to the formal relationship between shape and aesthetics [1].

These approaches rely on different forms of shape description, approximation and reduction, which we suggest can be compared if understood as different applications of shape similarity. Clearly, not all aspects of these approaches can be taken into account when analysed in terms of similarity but this is not our concern here. Indeed, our main goal is to find a way to organize and relate theories of shape via the notion of similarity they embrace. Whatever in these theories goes beyond similarity will need to be dealt with by specific extensions of the framework proposed here.

In a nutshell, what we propose is a logical framework based on [12,9] in which the modelling of each environment for shape similarity is done via codification in a possible world. We find this strategy suitable to model and compare different approaches to shape understanding, granularity, classification and recognition and in agreement with claims that similarity is a cluster of notions rather than a unitary concept [21]. Although this paper introduces just the basic ideas, without discussing important details, the rephrasing of well-known approaches to similarity within our framework in Section 3 and the presentation of a paradigmatic example in Section 4 should suffice to exemplify the flexibility and usefulness of the framework.

An interesting aspect of our work is the clear-cut distinction, nicely embedded in the framework, between what we call *ontological* and *epistemic* similarities (see Section 2 for technicalities). This distinction serves to explain why some similarities are proper of the space where one is working in, e.g., all circles are similar in Euclidean geometry but not in the domain of shape approximation where size is relevant. In our view, *ontological similarities* are the byproduct of the space definition, i.e., they arise from the structure of the possible world itself and thus 'belong' to the world because of the way it is (we call this also the natural or canonical similarity of a space). Other forms of similarity need cross-world relations to be modelled and these relations are motivated by external considerations, generally related to specific tasks. In this latter case we talk about *epistemic similarity* since these similarity forms, although added into the system, remain motivated by external knowledge.

Finally, it should be clear to the reader that our goal leads us to take a very general notion of similarity. It is well known that standard formalisations of similarity are based on some type of equivalence (sub)structure. But which one? It is undisputed that transitivity is too strong for similarity [7]. Still, most (formal) work on similarity retains the axiom of reflexivity and often even that of symmetry. However, in many areas such as e.g. in cognitive science these postulates are rejected, and therefore all three are given up in general in our framework. Indeed, a particular example for the failure of symmetry is given by similarity assessments of shapes (see, e.g., [4]).

Our approach allows us to be neutral with respect to these constraints when similarity is modelled across worlds. Indeed, what we call epistemic similarity is a general form of connection between entities (or classes) living in different worlds. Taking this approach we are not forced to introduce a reflexive link from a world to itself (which is needed to model reflexivity in epistemic similarity). Analogously, if we posit a similarity λ_1 from a world W_1 to a world W_2 (here understood as picking out the most similar objects in world W_2, i.e. the counterparts, see Section 2), there is no constraint to enforce a similarity relation λ_2 from W_2 to W_1. Furthermore, even if we have a λ_2 from W_2 to W_1, this does not need to take an entity $\lambda_1(a)$ back to a. That is, we can very well have $\lambda_2(\lambda_1(a)) \neq a$, so symmetry of similarity is not enforced either.

The other form of similarity, previously dubbed ontological similarity, is stronger. Entities are similar in this sense because of the structure of the world in which they live, and thus reflexivity is enforced by the logical setting. Symmetry is not, however. This point is relevant especially when similarity is used for filtering: when searching for, say, convex shapes, an arbitrary shape A can be classified as similar to a convex shape B but the opposite does not hold unless A is convex as well.

2 Counterparts, Cross-World Similarity, and $\mathcal{S}$-Connection

David Lewis provided the first formal theory of counterparts [14], a two-sorted first-order theory, whose sorts are objects and worlds, and which has four predicates: $W(x)$ says that x is a world, $I(x, y)$ that x is in the world y, $A(x)$ that x is an actual object, and $C(x, y)$ that x is a counterpart of y. The postulates of Lewis' counterpart theory codify that every object is in one and only one world, that counterparts of objects are objects, that no two different objects of the same world can be counterparts of each other, any object is a counterpart of itself, and that there is a world that contains all and only the actual objects and which is non-empty.

He described the basic intuition underlying the idea of counterparthood as follows:

> I prefer to say that you are in the actual world and no other, but you have counterparts in several other worlds. Your counterparts resemble you closely in content and context in important ways. They resemble you more closely than do the other things in their worlds. But they are not really you. For each of them is in his own world, and only you are here in the actual world. Indeed we might say, speaking casually, that your counterparts are you in other worlds, that they and you are the same; but this sameness is no more literal identity than the sameness between you today and you tomorrow. It would be better to say that your counterpart are men you *would have been*, had the world been otherwise. [14], p. 27–28

The connection to the standard modal language is obtained by translating sentences of modal predicate logic to the counterpart theory via the so-called **Lewis translation**; for details, see [14].

The general idea of counterpart relations being based on a notion of *similarity across worlds* also lies at the heart of similarity-based knowledge representation as proposed in this paper, and was a major inspiration for the design of 'modular languages' such as $\mathcal{E}$-connections [12], which we discuss next. An overview and discussion of counterpart-theoretic semantics can be found in [11].

$\mathcal{E}$-Connections as Counterpart Theory. In $\mathcal{E}$-connections, a finite number of formalisms talking about distinct domains are 'connected' by relations between entities in the different domains, capturing different aspects or representations of the 'same object'. For instance, an *abstract object* o of a description logic $\mathcal{L}_1$ can be related via a relation R to its *life-span* (a set of time points) in a temporal logic $\mathcal{L}_2$ as well as to its *spatial extension* (a set of points in a topological space, for instance) in a spatial logic $\mathcal{L}_3$. Essentially, the language of an $\mathcal{E}$-connection is the (disjoint) union of the original languages enriched with operators capable of talking about the link relations. The possibility of having multiple relations between domains is essential for the versatility of this framework, the expressiveness of which can be varied by allowing different language constructs to be applied to the connecting relations. A main feature of $\mathcal{E}$-connections is that, similar to description logics, they offer an appealing compromise between expressive power and computational complexity. The coupling between the combined logics is sufficiently loose for proving general results about the transfer of decidability: if the connected logics are decidable, then their connection will also be decidable. More importantly in our present context, they allow the combination of heterogeneous logical formalisms without the need to adapt the semantics of the respective components.

Note that, differently from the disjointness of the *formal languages* of the component logics, the requirement of disjoint domains is not essential. What this boils down to is the following simple fact: while more expressive $\mathcal{E}$-connection languages allow to express various degrees of *qualitative* identity, for instance by using number restrictions on links to establish partial bijections, they lack means to express 'proper' *numerical* trans-module identity. This issue, clearly, is closely related to the problem of trans-world identity well known from counterpart theory; we will expand on this below when introducing $\mathcal{S}$-connections.

For lack of space, we can here only roughly sketch the formal definitions, but compare [12]: we assume that the **languages** L_1 and L_2 of two logics $\mathcal{L}_1$ and $\mathcal{L}_2$ are pairwise disjoint. To form a connection $\mathcal{C}^{\mathcal{E}}(\mathcal{L}_1, \mathcal{L}_2)$, fix a non-empty set $\mathcal{E} = \{E_j \mid j \in J\}$ of binary relation symbols. The **basic $\mathcal{E}$-connection language** is then defined by enriching the respective languages with operators for talking about the link relations. A structure

$$\mathfrak{M} = \langle \mathfrak{W}_1, \mathfrak{W}_2, \mathcal{E}^{\mathfrak{M}} = (E_j^{\mathfrak{M}})_{j \in J} \rangle,$$

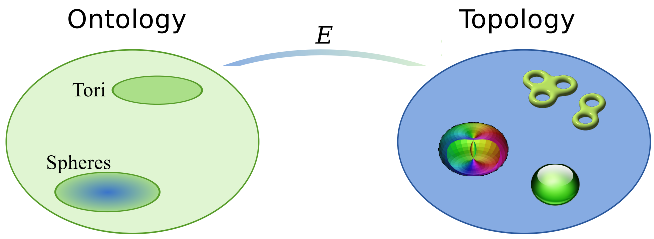

Fig. 1. A two-dimensional connection.

where $\mathfrak{W}_i = (W_i, .^{\mathfrak{W}_i})$ is an interpretation of $\mathcal{L}_i$ for $i \in \{1, 2\}$ and $E_j^{\mathfrak{M}} \subseteq W_1 \times W_2$ for each $j \in J$, is called an **interpretation** for $\mathcal{C}^{\mathcal{E}}(\mathcal{L}_1, \mathcal{L}_2)$. Given terms C_i of logics $\mathcal{L}_i$, $i = 1, 2$, denoting subsets of W_i, the semantics of the basic $\mathcal{E}$-connection operators $\langle\cdot\rangle^1$ and $\langle\cdot\rangle^2$ is

$$(\langle E_j\rangle^1 C_2)^{\mathfrak{M}} = \{x \in W_1 \mid \exists y \in C_2^{\mathfrak{M}} \ : (x, y) \in E_j^{\mathfrak{M}}\}$$
$$(\langle E_j\rangle^2 C_1)^{\mathfrak{M}} = \{y \in W_2 \mid \exists x \in C_1^{\mathfrak{M}} \ : (x, y) \in E_j^{\mathfrak{M}}\}$$

Fig. 1 displays the connection of an ontology (classifying shapes according to various qualities) with a spatial logic that distinguishes shapes by topological properties, with a single link relation E interpreted as the relation 'is the spatial extension of', i.e. relating abstract objects of the ontology with spatial objects.

As follows from the complexity results of [12], $\mathcal{E}$-connections add substantial expressivity and interaction to the component formalism.

$\mathcal{S}$-Connections: Similarity-based $\mathcal{E}$-Connections. Research on similarity is of a rather broad nature, including work in areas such as philosophy and general cognitive science, (description) logics, bio-informatics, and information retrieval, among others. Technically, the notion of similarity is closely related to fuzziness, as [8] discusses. By attaching fuzzy-values to link-relations, we can say that y is in the spatial extension $E(x)$ of point x with degree $p \in [0, 1]$, etc.[3]

Here, we concentrate on modelling a notion of *heterogeneous similarity*, i.e. similarity of objects drawn from conceptually different domains, specified by means of (heterogeneous) similarity measures which are closely modelled on the notions of distance functions and metrics. The notion of similarity-based $\mathcal{E}$-connections defined below thus combines the ideas of $\mathcal{E}$-connections [12], distance logics [13], and similarity logics [19].

We next define the notions of similarity and heterogeneous similarity spaces: let $\mathbb{R}^+_{0,\infty}$ denote the positive real numbers including zero and the symbol ∞ (denoting 'infinity'), and for $i = 1, 2$, we set $\bar{i} = 1$ if $i = 2$ and $\bar{i} = 2$ if $i = 1$.

[3] This natural idea has been studied for instance in the work of Suzuki on graded accessibility relations [22]. Also, Williamson [24] pursued similar semantic ideas when developing his propositional logics of clarity.

The following is a slight generalisation of the definition given in [9] to account for all similarity measures that appear in the literature, including partiality and non-reflexive ones.

Definition 1. *A **similarity space** $\mathbb{S} = \langle \mathcal{S}, \int \rangle$ (*sim*-space for short) consists of a set $\mathcal{S}$ together with a **similarity measure** $\int$ (read 's'), i.e. a partial function $\int : \mathcal{S} \times \mathcal{S} \to \mathbb{R}^{+}_{0,\infty}$ with $\mathsf{dom}(\int) \subseteq \mathcal{S} \times \mathcal{S}$. If $\int$ satisfies $\int(x,x) = 0$ for all $\langle x, x\rangle \in \mathsf{dom}(\int)$, it is called **reflexive**. In case $\forall \langle x,y\rangle \in \mathsf{dom}(\int) : \int(x,y) = 0 \iff x = y$ holds, we call $\int$ **discrete**. If $\int$ satisfies $\forall\langle x,y\rangle \in \mathsf{dom}(\int) : \int(x,y) = \int(y,x)$, we call $\int$ **symmetric**, and if it satisfies $\forall\langle x,y\rangle, \langle y,z\rangle, \langle x,z\rangle \in \mathsf{dom}(\int) : \int(x,y)+\int(y,z) \geq \int(x,z)$, we call it **triangular**. If $\int$ is a total function, discrete, symmetric and triangular, and $\infty \notin \mathsf{range}(\int)$, it is also called a **metric**, and $\langle \mathcal{S}, \int\rangle$ is called a **metric space**.*

Here, $\int(x,y) = 0$ means that x is *perfectly similar* to y. Note that, contrary to other formal approaches to similarity, closeness in the similarity space (i.e. a low value of the similarity measure) corresponds to high similarity: this intuition derives from the spatial interpretation of metric spaces.

However, note that perfect similarity implies identity only in the case of discrete spaces, i.e. that discreteness implies reflexivity but not *vice versa.* $\int(x,y) < \int(x,z)$ means that x is *more similar* to y than to z, and $\int(x,y) = \int(x,z)$ means that x is *equally similar* to y and z. Moreover, we say that x is *discernibly similar* to y if $\int(x,y) < \infty$ and *indiscernibly similar* otherwise, i.e. if $\int(x,y) = \infty$. For example, suppose as spatial objects you consider convex polygons in the plane together with circles of different sizes. Now a very simple similarity measure would, in the case of polygons, compute the difference in the number of vertices (i.e., the smaller this difference the more similar the polygons are) and, for circles, compute the difference in diameter (i.e., the smaller that difference the more similar the circles are). Comparing polygons with circles, however, remains undefined (think of circles as having an infinite number of vertices). Then, clearly, all pairs of polygons (and all pairs of circles) are discernibly similar, whilst all pairs of a polygon and a circle are indiscernibly similar.

For $X, Y \subseteq \mathcal{S}$ sets (rather than just elements), similarity is defined by extending $\int$ as follows:

$$\int(X,Y) := \begin{cases} \inf\{\int(x,y) \mid x \in X, y \in Y, \langle x,y\rangle \in \mathsf{dom}(\int)\}, \text{if } \mathsf{dom}(\int) \cap X \times Y \neq \emptyset \\ \infty, \text{ otherwise} \end{cases}$$

If in fact the minimum exists for all non-empty sets X and Y, $\mathcal{S}$ is also called a min-space, compare [19]. Clearly, whenever a space is finite, it is a min-space; this obviously is the case in many practical applications.

In general, a space can be equipped with several similarity measures, some of which can be epistemic in the sense of Section 1. But in particular, we assume that the basic underlying ontology and logic of a given space, i.e. what the logic 'quantifies over' (used here in the sense of [17]), defines a corresponding **natural** or **canonical similarity space**. Here, the **canonical similarity measure** is given by relative identity, i.e. by measuring the number of shared properties available in a given logical language (in the sense of Leibniz; see, e.g., [11]), and therefore in particular always satisfies reflexivity. Canonical similarity is ontological in the sense discussed in Section 1.

When relating *different* sets of objects, e.g. shapes as topological objects, as objects in Euclidean or projective space etc. (see Section 4), the above definitions need to be adapted. For simplicity, we here restrict our attention to the case of only two such sets.

Definition 2. *A (2-dim)* ***heterogeneous similarity space*** *(*hsim*-space for short) is a quadruple* $\mathbb{H} = \langle \mathbb{S}_1, \mathbb{S}_2, \int_1^2, \int_2^1 \rangle$ *consisting of, for* $i = 1, 2$*,* sim*-spaces* $\mathbb{S}_i = \langle \mathcal{S}_i, \int_i \rangle$*, and* ***heterogeneous similarity measures****, which are partial functions* $\int_i^{\bar{i}} : \mathcal{S}_i \times \mathcal{S}_{\bar{i}} \mapsto \mathbb{R}^+_{0,\infty}$*.* $\mathbb{H}$ *is* ***het-symmetric*** *if whenever* $\langle x, y \rangle \in \mathsf{dom}(\int_1^2)$ *and* $\langle y, x \rangle \in \mathsf{dom}(\int_2^1)$ *we have* $\int_1^2(x, y) = \int_2^1(y, x)$*. It is* ***het-triangular*** *if for all* $x, z \in \mathcal{S}_i$ *and* $y \in \mathcal{S}_{\bar{i}}$ *we have* $\int_i^{\bar{i}}(x, y) + \int_{\bar{i}}^i(y, z) \geq \int_i(x, z)$ *(for* $i = 1, 2$*), whenever all measures are defined.*

In the heterogeneous case, *perfect similarity* now means that $x \in \mathcal{S}_1$ and $y \in \mathcal{S}_2$ are indistinguishable from the perspectives of both heterogeneous similarity measures, $\int_1^2$ and $\int_2^1$, i.e., that $\int_1^2(x, y) = \int_2^1(y, x) = 0$. Note that the notion of *discrete* similarity measure makes no immediate sense in the heterogeneous case as (numerical) identity is typically not available. However, the notion can be ‘simulated’ by replacing identity with an independently defined notion of trans-module identity, ‘equalising’ cross-domain elements whilst respecting the similarity measures.

Counterparts in $\mathcal{S}$-Connections. Note that, in this setting, the problems of transworld identity and counterparthood can be neatly separated: transworld identity may be taken to be synonymous with perfect similarity as defined above. Counterparthood understood as *maximal similarity* is a looser notion, and may be explicated by the following principle (see [5]).

> For $x \in \mathcal{S}_i$ and $y \in \mathcal{S}_{\bar{i}}$, y is a counterpart of x only if nothing in $\mathcal{S}_{\bar{i}}$ is more similar to x as it is in $\mathcal{S}_i$ than is y as it is in $\mathcal{S}_{\bar{i}}$.

We take this principle as the defining criterion for counterparthood in similarity spaces:

Definition 3 (Counterparts). *Let* $\mathbb{H} = \langle \mathbb{S}_1, \mathbb{S}_2, \int_1^2, \int_2^1 \rangle$ *be a* hsim*-space. We call* $b_{\bar{i}} \in \mathcal{S}_{\bar{i}}$ *an* $\bar{i}$***-counterpart of*** $a_i \in \mathcal{S}_i$ *if* $\int_i^{\bar{i}}(a_i, b_{\bar{i}}) = \inf\{\int_i^{\bar{i}}(a_i, b) \mid b \in \mathcal{S}_{\bar{i}}\} < \infty$*, which we also write as* $\mathsf{Cp}_i^{\bar{i}}(a_i, b_{\bar{i}})$*. This gives us two relations:* $\mathsf{Cp}_i^{\bar{i}} \subseteq \mathcal{S}_i \times \mathcal{S}_{\bar{i}}, i = 1, 2$*. Moreover, for* $X \subseteq \mathcal{S}_i$*, we denote by* $\mathsf{Cp}_i^{\bar{i}}(X)$ *the set* $\{y \in \mathcal{S}_{\bar{i}} \mid \exists x \in X. \mathsf{Cp}_i^{\bar{i}}(x, y)\}$*.*

Note that counterparts thus defined may or may not exist, and if they exist may or may not be unique. Moreover, $b_{\bar{i}}$ may be an $\bar{i}$-counterpart of a_i without a_i being an i-counterpart of $b_{\bar{i}}$; counterparthood is *directional.* Although counterparts need not be unique, in applications it is often desirable to select amongst the elements with maximal similarity a unique element, according to certain *external* criteria. We here solve this problem by incorporating into the structures an explicit choice function selecting a counterpart.

Definition 4 (Counterpart choice). *A* hsim*-space* ***with choice*** *is a triple* $\langle \mathbb{H}, \lambda_1, \lambda_2 \rangle$*, where* $\mathbb{H} = \langle \mathbb{S}_1, \mathbb{S}_2, \int_1^2, \int_2^1 \rangle$ *is a* hsim*-space, and, for* $i = 1, 2$*,* $\lambda_i : \mathcal{S}_i \longrightarrow \mathsf{Cp}_i^{\bar{i}}(\mathcal{S}_i)$ *are (partial)* ***choice functions*** *such that, for all* $x \in \mathcal{S}_i$*,* $\lambda_i(x) \subseteq \mathsf{Cp}_i^{\bar{i}}(x)$ *is a singleton if* $\mathsf{Cp}_i^{\bar{i}}(x) \neq \emptyset$*, and undefined otherwise.*

Of course, often the λ_i are uniquely determined by the similarity measures $\int_i^{\bar{i}}$, in which case we call λ_i a **deterministic choice function**.

Apart from the elements with maximal similarity, i.e. the counterparts, it is also of interest to be able to refer to elements of a foreign domain that are similar *to some degree* (i.e. discernibly similar). This can be achieved by simulating the notion of *link relation* from $\mathcal{E}$-connections as follows:

Definition 5 (Link-relation). *Given a* hsim*-space* $\mathbb{H} = \langle \mathbb{S}_1, \mathbb{S}_2, \int_1^2, \int_2^1 \rangle$*, we define the* ***induced link relations*** $E^1_{\mathbb{H}}, E^2_{\mathbb{H}}, E_{\mathbb{H}} \subseteq \mathcal{S}_1 \times \mathcal{S}_2$ *by setting, for all* $x \in \mathcal{S}_1$ *and* $y \in \mathcal{S}_2$*:*

$$\begin{aligned} E^1_{\mathbb{H}}(x,y) &\iff \textstyle\int_1^2(x,y) < \infty; \\ E^2_{\mathbb{H}}(x,y) &\iff \textstyle\int_2^1(y,x) < \infty; \\ E_{\mathbb{H}}(x,y) &\iff \min\big(\textstyle\int_1^2(x,y), \int_2^1(y,x)\big) < \infty (= E^1_{\mathbb{H}} \cup E^2_{\mathbb{H}}). \end{aligned}$$

Intuitively, the relation $E_{\mathbb{H}}(x,y)$ holds if x and y are discernibly similar from at least one 'viewpoint', and $E^i_{\mathbb{H}}(x,y)$ holds if x and y are discernibly similar from the point of view of $\int_i^{\bar{i}}$. This way we can recover standard $\mathcal{E}$-connections as shown in [9]. A basic logic for these semantic structures was also introduced in [9], but we omit the details here and instead focus on informally studying examples that can be modelled within this framework.

3 Approaches to Similarity and $\mathcal{S}$-Connections

In this part, we classify in $\mathcal{S}$-Connections some traditional approaches to similarity.

Goldstone and colleagues list four classes of models that have been used to formalise similarity. These are dubbed (see [6] for further information and references): *geometric*, *feature-based*, *alignment-based*, and *transformational*.

The geometric models take as input similarity judgments like "A is 90% similar to B" or "A is more similar to B than to C". The model represents similarity among entities by selecting a solution of the set of constraints given by the similarity judgments. The solution associates each entity (an individual or a class) to a point (an individual) in the world, say W_G, which is typically a vector space (applications to non-linear manifolds exist also).

In our framework these judgments are modelled as constraints on primitive similarity relations. Given a relation Sim, a judgment "A is 90% similar to B" is formalised in W_G as, say, $\int(A,B) = 0.1$ (recall that 0 stands for perfect similarity). Comparative judgments like "A is more similar to B than to C", are modelled by direct comparison, i.e. $\int(A,B) < \int(A,C)$. In general, the goal of geometrical models is to systematise similarity information in order to allow further analysis of this kind of data. For this reason, in our framework one would typically work within a single possible world whose language is driven by the provided types of judgments.

Note that geometrical models are here classified as generating ontological similarities since the judgments are taken as primitive elements of the world W_G. This conclusion is justified by the fact that one looks at what is assumed within that world and the outcome of these assumptions on similarity are ontological consequences of what that world takes to exist. In order to see a geometrical model as a case of epistemic similarity, one should introduce a world W_i for each similarity judgment (or group of them), a new world W'_G, and introduce cross-world relations $\int_1^2$ with $W_i = \mathbb{S}_1$ and $W'_G = \mathbb{S}_2$. These relations embed all the given constraints into W'_G which is thus equivalent to the single world W_G we discussed earlier. However, since W'_G is now collecting those constraints indirectly via cross-world relations, the similarity relation $\int'_G$ of W'_G is now an example of epistemic similarity.

The feature-based models characterise similarity in terms of a feature-matching process. Entities are represented as a collection of features and similarity is established by weighting common and distinctive features. In our framework, this is modelled in a single possible world W_F where the language has all the needed properties to model the entities' features (a feature can correspond to a property or to a combination of these). Differently from the geometrical case, here the similarity relation $\int$ is not primitive but defined by the formulas used in feature-based approaches (generally a linear combination of the common and distinctive properties). For instance, in a world including property *Ang* (having an angle), *Str* (having a straight border), *Cir* (being a circle) and *Pol* (being a polygon) we can define different similarity relations $\int_{F_1}, \int_{F_3}, \ldots$ to the result that a sector of a circle is similar to a polygon and not to a circle according to $\int_{F_1}$, *vice versa* if we use $\int_{F_2}$, and is similar to both if we use $\int_{F_3}$. Of course, if other relations are available in the language, like in a metric space, one can further refine the similarity based, say, on the proportion between length of the radius and the sector angle. Feature-based models are typically systems of ontological similarity: the similarities arising within these models are built out of information purely contained in each world structure W_F.

The third traditional approach to similarity is based on alignment models. In these models, comparison requires to determine how elements and properties correspond to, or align with, one another. For instance, we can associate red in the world of colours to danger in the commonsense world. Then, a further association of red to fire would lead to state a similarity in the commonsense world between fire and danger. The alignment model approach allows also to contextualise some judgments or features and is particularly suited when epistemic connections, i.e. cross-world relations like in the red-danger example, are used.

Basically, positively aligned features are crucial to increase similarity, negatively aligned features are used to decrease similarity, and non-aligned features can be ignored (the presence of features of one entity that cannot be aligned at all with features of another can be considered also). Clearly, alignment models generate what we called epistemic similarities.

Finally, transformational models in our framework enhance the similarity relationships across different worlds by adding a new parameter to measure complexity. Cross-world relations of form $\int_i^{\bar{i}}$ are indeed now weighted with 'transformation complexity' values. A cross-world relation that is complex has a minor impact on the definition of $\int$ in the word W_T than a simpler cross-world relation. Usually complexity of transformations can be established in different terms, e.g., one can look at the combinatorics of basic similarity transformations like rotation, scaling, reflexion, convex-hull generation, topological transformation etc. Being a generalisation of the alignment approach, transformational models are also classified as modelling epistemic similarities.

4 $\mathcal{S}$-Connection of Euclidean and Projective Geometry

We have seen in Section 3 how general approaches and uses of shape similarity can be reformalised in our framework. This is not always trivial though: whilst approaches like [3] are naturally reconstructed, more involved uses of similarity like [23] and [10] are more demanding.

In this last section we describe a different example to show how we can use the framework to compare and integrate different similarity relations. In accordance with the goals of this paper, we do not dwell upon formalisation details.

Assume we have a world W_E capturing Euclidean geometry, and a world W_P for Euclidean projective geometry (the world of 3D polygons). The ontological similarity relations $\int_E$ of $\mathbb{S}_E$ and $\int_P$ of $\mathbb{S}_P$ agree on the common subset of polygons, and there is a natural relation $E_{\mathbb{H}}$ of perfect similarity between polygons in W_E and polygons in W_P.

The problem of evaluating shape-based matching and retrieval algorithms for (generally 3D) polygonal models has led to the Princeton Shape Benchmark (PSB), a framework to compare 3D shape matching algorithms [20]. The PSB provides a mechanism to specify partitions of the 3D models in classes that can be as generic as "animals", "flying creatures", "birds" and "birds in a flying pose". Multiple classifications are possible and thus can be averaging, e.g. "roughly spherical". Let us call W_{PSB} the world of PSB with its predicates for the classes.

In our framework, the different shape matching algorithms are seen as 'implicit' definitions of cross-world similarity relations: a cross-world relation $E^1_{\mathbb{H}}$ goes from the world of polygons W_P (partially instantiated by the PSB database of 3D models) to the partitions in the W_{PSB} world in agreement with the classification provided by the associated algorithm.

From the theoretical viewpoint, the PSB system implements a similarity graph with initial node the world W_E of actual shapes, connected to W_P via

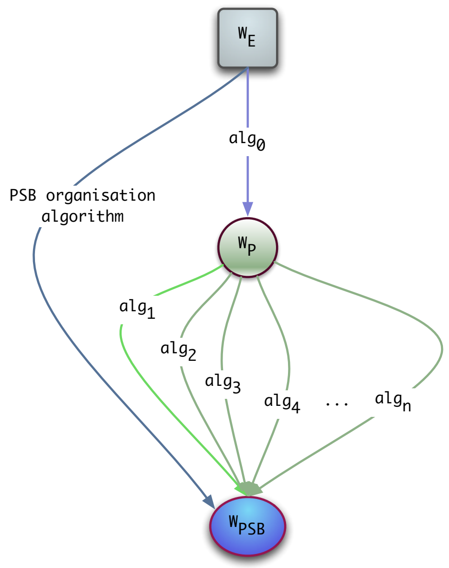

l

Fig. 2. The Princeton Shape Benchmark as $\mathcal{S}$-Connection.

some approximation algorithm. W_P is then the source of a series of similarity links ending into W_{PSB}. The explication of the algorithms' steps would reveal a network of worlds that are positioned midway between W_P and W_{PSB}. A further link connects directly W_E to W_{PSB} to model the shape classification developed by the PSB organisers and that is at the base of the algorithms' evaluation.

A typical evaluation by the PSB system would end up with a judgment of type "method X is better for this type of object and method Y is better for that type of object, etc" [20]. With our framework, we can now motivate and systematise the PSB classification by comparing the classifications in the midway worlds. One can thus understand where the actual grouping happens and the corresponding computational cost, so to have a better view of the tradeoff between quality of shape classification and efficiency of the procedure. But more importantly, one can see the role of each similarity transformation used by each algorithm, how these transformations impact on the result if applied at an earlier or later stage, and finally it becomes possible to establish the best strategy to combine them in order to optimise the system.

Acknowledgements.

We gratefully acknowledge the financial support of the Deutsche Forschungsgemeinschaft through the Collaborative Research Center on Spatial Cognition (SFB/TR 8). The authors would also like to thank the anonymous reviewers for their feedback, and especially Dirk Walther for detailed comments on a draft of this paper.

References

1. Birkhoff, G.D., *Aesthetic Measure*, Cambridge Massachusetts' University Press, 1933.
2. Clementini, E, and P Di Felice, 'A global framework for qualitative shape description', *GeoInformatica*, 1 (1997), 1, 11–27.
3. Edelman, S, F Cutzu, and S Duvdevani-Bar, 'Similarity to reference shapes as a basis for shape representation', *Proceedings of 18th Annual Conf. of the Cognitive Science Society*, (1996), 260–265.
4. Edelman, Shimon, 'Representation is representation of similarities', *Behavioral and Brain Sciences*, 21 (1998), 449–498.
5. Forbes, G., *The Metaphysics of Modality*, Oxford Clarendon Press, Oxford, 1985.
6. Goldstone, R.L., S. Day, and J.Y. Son, 'Comparison', *Towards a Theory of Thinking*, (2010), 103–121.
7. Goodman, Nelson, 'Seven strictures on similarity', in *Problems and projects*, York : Bobbs-Merril, 1975.
8. Hájek, P., 'Why Fuzzy Logic?', in Dale Jacquette, (ed.), *A Companion to Philosophical Logic*, Blackwell, 2002, pp. 595–607.
9. Hois, J., and O. Kutz, 'Counterparts in Language and Space—Similarity and $\mathcal{S}$-Connection', in C. Eschenbach, and M. Grüninger, (eds.), *Formal Ontology in Information Systems (FOIS 2008)*, IOS Press, 2008, pp. 266–279.
10. Jowers, Iestyn, David Hogg, Alison McKay, Hau Chau, and Alan de Pennington, 'Shape detection with vision: implementing shape grammars in conceptual design', *Research in Engineering Design*, 21 (2010), 235–247.
11. Kracht, M., and O. Kutz, 'Logically Possible Worlds and Counterpart Semantics for Modal Logic', in D. Jacquette, (ed.), *Philosophy of Logic, Handbook of the Philosophy of Science*, vol. 5, Elsevier, Amsterdam, 2007, pp. 943–996.
12. Kutz, O., C. Lutz, F. Wolter, and M. Zakharyaschev, '$\mathcal{E}$-Connections of Abstract Description Systems', *Artificial Intelligence*, 156 (2004), 1, 1–73.
13. Kutz, O., H. Sturm, N.-Y. Suzuki, F. Wolter, and M. Zakharyaschev, 'Logics of metric spaces', *ACM Transactions on Computational Logic (TOCL)*, 4 (2003), 2, 260–294.
14. Lewis, D. K., 'Counterpart Theory and Quantified Modal Logic', *Journal of Philosophy*, 65 (1968), 113–126.
15. Lysenko, Mikola, Vadim Shapiro, and Saigopal Nelaturi, 'Non-commutative morphology: Shapes, filters, and convolutions', *Computer Aided Geometric Design*, (2011), 1–34.
16. Mimaroglu, S., and E. Erdil, 'Asod: Arbitrary shape object detection', *Engineering Applications of Artificial Intelligence*, (2011).
17. Quine, Willard van Orman, 'Variables explained away', in *Selected Logic Papers*, Random House, New York, 1960, pp. 227–35. Reprinted.
18. Quine, W.V., *Ontological relativity and other essays*, vol. 1, Columbia Univ Pr, 1969.
19. Sheremet, M., D. Tishkovsky, F. Wolter, and M. Zakharyaschev, 'A Logic for Concepts and Similarity', *J. of Logic and Computation*, 17 (2007), 3, 415–452.
20. Shilane, P., P. Min, M. Kazhdan, and T. Funkhouser, 'The Princeton Shape Benchmark', in *Proceedings of Shape Modeling International*, IEEE Computer Society, Washington, DC, USA, 2004, pp. 167–178.
21. Sovran, Tamar, 'Between similarity and sameness', *Journal of Pragmatics*, 18:4 (1992), 329–344.

22. Suzuki, Nobu-Yuki, 'Kripke frames with graded accessibility and fuzzy possible world semantics', *Studia Logica*, 59 (1997), 249–269.
23. Tierny, Julien, Jean-Philippe Vandeborre, and Mohamed Daoudi, 'Geometry flavored topological skeletons: applications to shape handling, segmentation and retrieval', in *Second DELOS Conference on Digital Libraries*, 2007.
24. Williamson, T., *Vagueness*, The Problems of Philosophy, Routledge, London, 1994.