Minds and Machines manuscript No.
(will be inserted by the editor)

Unifying conceptual spaces

Concept formation in musical creative systems

Jamie Forth · **Geraint A. Wiggins** · **Alex McLean**

Received: date / Accepted: date

Abstract We examine Gärdenfors' theory of conceptual spaces, a geometrical form of knowledge representation (Gärdenfors 2000), in the context of the general Creative Systems Framework introduced by Wiggins (2006a, b). Gärdenfors' theory offers a way of bridging the traditional divide between symbolic and sub-symbolic representations, as well as the gap between representational formalism and meaning as perceived by human minds. We discuss how both these qualities may be advantageous from the point of view of artificial creative systems. We take music as our example domain, and discuss how a range of musical qualities may be instantiated as conceptual spaces, and present a detailed conceptual space formalisation of musical metre.

Keywords Conceptual Spaces · Creativity · Search · Geometry · Musical Rhythm · Similarity

Introduction

In this paper, we examine the relationship between the Creative Systems Framework (CSF), formalised by Wiggins (2006a, b) from the more abstract ideas of Boden (2004, 1998), and the theory of Conceptual Spaces proposed by Gärdenfors (2000). Such a comparison is interesting for two reasons: first, Gärdenfors' theory (which addresses general cognition and meaning, rather than creativity) is a good candidate to instantiate some of the abstractions in Boden's theory, and, second, it allows us to extend that theory with a relatively well-defined (though not uncontroversial) notion of *semantics* and thence to apply it in real creative contexts. This may allow us to address some of the difficult questions surrounding the creation of meaning in computational creativity; further, an important aspect of Gärdenfors' theory is that it yields empirically testable predictions—a valuable step forward in rigour for the field. We will also involve recent work by Thornton (2007) in our study, which will admit what Gärdenfors calls *relational concepts* (Gärdenfors 2000, p. 92) into the CSF, allowing the reasoning of which the CSF is capable to extend beyond abstractions of objective concepts into the world of relations. We will base a detailed example on the work of Justin London (2004), on musical rhythm.

We begin by introducing the theories involved in the study, and then discuss how they fit together. As an initial example, we outline the conceptualisation of musical melody using existing methods of description to show that the framework gives an intuitively reasonable account. We then give a more detailed, novel application of the conceptualisation of musical metre. This yields a formal theory of metrical similarity which is amenable to empirical study.

We conclude that the unification of these theories of conceptual space has considerable potential for future development of creative systems and for scientific understanding of human cognition.

Jamie Forth, Geraint A. Wiggins & Alex McLean
Department of Computing &
Centre for Cognition, Computation and Culture
Goldsmiths, University of London
Tel.: +44-20-7919 7861
Fax: +44-20-7919 7853
E-mail: {g.wiggins,j.forth,a.mclean}@gold.ac.uk

Sources and Background

Boden's *Creative Mind*

The Creative Mind (Boden 2004) has attained an almost biblical status in the field of computational creativity, despite having its critics (Haase 1995; Lustig 1995; Perkins 1995; Ram et al. 1995; Schank and Foster 1995; Turner 1995). It has been summarised elsewhere (Wiggins 2006a), and therefore we focus here on precisely the points that are relevant to the current study, referring the reader to the original for a fuller story. We focus on Boden's theory because it takes a modern and more formal approach to cognitive science, as compared with the earlier theories of Wallas (1926), Koestler (1964) and Guilford (1967). As such, it is more directly amenable to implementation, and in any case may well encompass the more cognitive-mechanistic aspects of the older theories[1]; it is these aspects that interest us here.

Boden's key notion is of a *conceptual space*, inhabited by concepts related to a particular topic; since there are many possible topics, there are many conceptual spaces. Boden presents her concepts as essentially *generated* (Thornton 2007), but does not specify a generation mechanism, other than naming it *exploration* of the space; in consequence, the idea has been compared with traditional AI state space search, though it is not the same (Wiggins 2006b). It does, however, stand in a clear AI tradition of knowledge representation, and can be instantiated with any one of a range of standard AI methods, such as frames. Boden does not present her space as having intrinsic properties: it is merely a collection of related ideas, grouped together so as to allow reasoning about them. However, reasoning about concepts and their relation with the conceptual space is not meaningful in Boden's formulation, because there is nowhere for concepts to exist, but within their conceptual space. Notwithstanding this, Boden introduces the notion of *transformational* creativity, where the space itself is changed, corresponding with a redefinition of the topic with which the conceptual space was associated.

Important in Boden's theory is a notion of *evaluation* (Boden 1998). This is the process whereby the quality of a creation is measured; membership of the conceptual space certifies the nature of the creation, while quality measures whether it is any *good*—which, of course, begs an appropriate criterion of quality. This distinction allows us to identify something as, for example, a joke, while separately deciding whether or not it is funny.

To simplify our exposition, we note the equivalence between a quasi-Platonic universe of all possible concepts (including partial ones), *traversed* by an appropriate procedure, and meaningful structures, *constructed* from building blocks by an appropriate procedure; the same duality exists in state space search. In what follows, we will take the former view, which is more readily amenable to mathematical formulation.

[1] For example, Boden's *combinatorial creativity* seems to coincide precisely with Koestler's *bisociation of matrices*. On the other hand, while Guilford's descriptions of divergent and convergent thinking may evidently be applicable *within* a framework such as that which Boden provides, they are necessarily vague descriptions of the high-level behaviour of a very complex system—sufficiently high-level not to be helpful in actually defining it.

Finally, it is worth precluding a misunderstanding, which we have encountered in several contexts, of Boden's proposal of exploratory creativity. A common argument against that proposal is that regular systematic exploration is not creativity at all, but merely enumeration and selection. The argument is that such enumeration and selection does not match with the notion of "inspiration" which is so heavily intertwined with the Romantic (and present-day) conception of creativity. This argument misses a significant point, however: the sensation of "inspiration" is caused when a human creator *becomes conscious* of his or her new idea. Now, not all brain/mind activity is accessible to consciousness at all times, and linguistic creativity (of the small kind involved in generating sentences of everyday speech) has been shown to begin ahead of conscious awareness of the resulting utterance (Carota et al. 2009). So one possible mapping from Boden's stated theory onto the mind would be such that exploratory mechanisms operate at a level *inaccessible* to consciousness, with an evaluation mechanism acting as the trigger to *attract conscious attention*, thus producing the sensation of "inspiration". Therefore, it is a reasonable proposition that even a mechanism as naïve as uninformed search might be operating on a non-conscious level, with its outputs becoming available to consciousness only on completion, and therefore *appearing* to be original insight or inspiration to the person experiencing it. It follows that such naïve mechanisms cannot be casually ruled out as "obviously incorrect". A corollary of this is that we cannot make the appropriate distinction in a creative system without a functional model of consciousness, such as Global Workspace Theory (Baars 1988)—and completion of such a model is probably some way off.

In summary, Boden's theory still provides the most general—but also most strongly grounded—framework for discussing the issues, especially when formalised as discussed in the next section.

Wiggins' Creative Systems Framework

Wiggins (2006a, b) defines the *Creative Systems Framework* (CSF) for describing and reasoning about creative systems, based upon the work of Boden (2004). The CSF defines a creative system as "a collection of processes, natural or automatic, which are capable of achieving or simulating behaviour which in humans would be deemed creative" (Wiggins 2006a, p. 451). The CSF specifies seven key symbols and functions, shown in Table 1. Given these, we can define Boden's conceptual space:

$\mathscr{C}$ A conceptual space, a set of concepts selected by $[\![\mathscr{R}]\!](\mathscr{U})$

Using the CSF, it can be shown that Boden's transformational creativity (Boden 2004) is in fact the same process

Table 1: The symbols and concepts of the Creative Systems Framework (Wiggins 2006a)

$\mathscr{U}$	The universe of all possible concepts
$\mathscr{L}$	A language in which to express the elements of $\mathscr{U}$ and functions and predications over them
$\mathscr{R}$	Rules defining valid concepts, expressed in $\mathscr{L}$
$[\![.]\!]$	An interpreter for $\mathscr{L}$; given a rule set, returns a function which evaluates the degree to which the rules are true of a concept
$\mathscr{T}$	A traversal strategy for stepping between concepts within $\mathscr{U}$
$\mathscr{E}$	Rules which evaluate the quality or desirability of a concept
$\langle\langle .,.,. \rangle\rangle$	An interpreter for $\mathscr{L}$, returning a function which maps the three rule sets, $\mathscr{R}$, $\mathscr{T}$, $\mathscr{E}$. to a function which operates upon an ordered subset of $\mathscr{U}$ (to which it has random access) and outputs another ordered subset of $\mathscr{U}$.

as exploratory creativity, but at a meta-level with respect to the creative system; that is to say, it is exploratory creativity in the conceptual space of conceptual spaces (Wiggins 2006a). This yields a more parsimonious theory, and raises interesting questions about other possible meta-levels, not yet considered.

A key difference between the CSF and Boden's theory is the inclusion of a *traversal* mechanism, expressed as a set of traversal rules, $\mathscr{T}$, which may refer to the definition of the conceptual space, $\mathscr{R}$, and the evaluation rules, $\mathscr{E}$, with an interpreter $\langle\langle .,.,. \rangle\rangle$, to apply them. Creative systems search for valued concepts by iteratively applying $\langle\langle \mathscr{R},\mathscr{T},\mathscr{E} \rangle\rangle$ over an ordered set of (possibly partial) concepts. Given a $\mathscr{C}$ with $\mathscr{E}$-valued but as yet undiscovered concepts, the success of the system depends on the ability of $\mathscr{T}$ to navigate the space. However, a crucial feature of the CSF, which broadens Boden's conceptualisation, is that creative search is defined as an operation over $\mathscr{U}$, and not limited only to $\mathscr{C}$. This allows the search to lead *outside* $\mathscr{C}$, so that application of $\mathscr{T}$ can result in concepts not conforming with $\mathscr{R}$. Such effects are termed *aberrations* and they can invoke transformational creativity. If an aberration contains only $\mathscr{E}$-valued concepts, it is *perfect aberration* and $\mathscr{R}$ should be transformed to include the concepts in $\mathscr{C}$. If no aberrant concepts are valued, it is *pointless aberration* and $\mathscr{T}$ should be transformed to avoid them. If some are valued and others not, it is *productive aberration* and both $\mathscr{R}$ and $\mathscr{T}$ should be transformed. In this way, a creative system is able dynamically to manipulate its conceptual space and its manner of searching in response to the concepts it finds.

These behaviours allow description of various different kinds of creative behaviour, in terms of concepts. However, they say little about the concepts themselves, and make no attempt to structure the conceptual space beyond the implicit structuring imposed by $\mathscr{T}$. Therefore, there is no means of distinguishing different kinds of concept—for example, "thing" *vs.* "action". Next, we summarise an approach to addressing this issue, introducing different types of concept (in particular, concepts which are essentially relational), including derived types.

Thornton's Taxonomy of Concept Types

Thornton (2007) extends Boden's conceptual space with some notions from symbolic AI admitting a small but significant taxonomy of concept types. Thornton takes the constructive view of his conceptual space, and therefore describes the construction of his concepts as either *categorical*, where one concept is an *instance* of another or *compositional*, where one concept is created from a combination of others. This approach requires some primitives: there must be basic, atomic concepts, so that the semi-lattice produced by the constructions eventually bottoms out, and there must be relations, which instantiate the construction method for the compositional concepts. Thornton gives an example based on clothing: given a SHIRT, a JACKET and some TROUSERS, each item is a GARMENT; but also, and differently, one can *compose* them into a UNIFORM. The former is a categorical construction; the latter is compositional.

This example, however, is not straightforward. The concept of UNIFORM is not entirely captured by the constituents of an archetypal suit of clothes: there is an aspect of UNIFORM which connotes identification of the wearer as a member of a social group; some uniforms connote authority, some connote belief, and so on. Therefore, it seems likely that many, if not most, useful compositional concepts will themselves be instances of, or identical with, categorical ones—as was so with the *is-a* and *part-of* hierarchies of symbolic AI.

The taxonomy is further complicated by the fact that one must allow arbitrary stacking of both kinds of relation, just as in the older AI hierarchies, and, worse, that the relations used in compositional concepts are themselves concepts. The space explodes quickly, but Thornton gives a heuristic which may help in navigating it.

The examples given in Thornton's presentation are fundamentally symbolic, and he does not attempt to explain how his basic concepts are given their meaning, nor how their symbols might be defined; equally, consideration is not given to concepts which consist of continuous dimensions—the work is essentially symbolic, and the definition of the symbols themselves is presupposed. This is not a failing of the theory, because its contribution is on a different level from the current proposal, but it does leave a gap to be filled, to which we return below.

Thornton's ideas form an important extension of Boden's theory, which can straightforwardly be formalised within the CSF. They also match neatly (Thornton 2009) with aspects of the Theory of Conceptual Spaces proposed by Gärdenfors (2000), which we introduce next.

Gärdenfors' Conceptual Spaces

Before introducing Gärdenfors' theory, we must defuse a potential confusion in terminology: in the CSF, a conceptual space is an abstract representation of a component of a creative system, corresponding broadly with Boden's and Thorton's structure of the same name. However, such a correspondence with Gärdenfors' conceptual spaces should not be *assumed*; indeed, our aim here is to explore the *apparent* relationship between the two. For this section, then, the term "conceptual space" is exclusively reserved for Gärdenfors' ideas.

Gärdenfors (2000) argues that conceptual structures should be represented using geometry on what he terms the *conceptual level*. This level of representation is situated between the *symbolic level*, which includes, for example, formal grammar, and the *sub-conceptual level* of high-dimensional representations such as neural networks. An important aspect of the theory is that these three levels of representation should be understood as complementary: Gärdenfors sees his conceptual level as *uniting* the other two, which are often presented in the literature as an irreconcilable dichotomy. One interpretation of the difference between the geometrical conceptual level and the non-symbolic sub-conceptual one may be that the basic dimensions of the former (described below) are perceptually immediate (in the sense that they correspond directly with experienced perception), though this is not made explicit in Gärdenfors' exposition.

Gärdenfors' theory of conceptual spaces begins with an atomic but general notion of *betweenness*, in terms of which he defines *similarity*, represented as (not necessarily Euclidean) distance. This allows models of cognitive behaviours (such as creative ones) to apply geometrical reasoning to represent, manipulate and reason about concepts. Similarity is measured along *quality dimensions*, which 'correspond to the different ways stimuli are judged to be similar or different' (Gärdenfors 2000, p. 6). An archetypal example is a colour space with the dimensions hue, saturation (or chromaticism), and brightness. Each quality dimension has a particular geometrical, topological or ordinal structure. For example, hue is circular, whereas brightness and saturation correspond with measured points along finite linear scales. Identifying the characteristics of a dimension allow meaningful relationships between points to be derived, and it is important to note that the values on a dimension need not be numbers—though how an appropriate algebra is then defined is not discussed.

Quality dimensions may be grouped into *domains*. A domain is a set of *integral* (as opposed to *separable*) dimensions, meaning that no perceptually-meaningful value can be chosen in one dimension without every other dimension in the domain being defined. Therefore, hue, saturation, and brightness in the colour model, above, form a single domain. It follows that Gärdenfors' definition of a conceptual space is simply 'a collection of one or more domains' (Gärdenfors 2000, p. 26). Another example, more relevant to the current work, is musical PITCH, which seems to form a two dimensional domain, one dimension, *pitch height*, being linear, and the other, *chroma* (broadly, note name), being circular (Shepard 1982).

Since the quality dimensions originate in betweenness, similarity is directly related to proximity, though not necessarily Euclidean proximity. Such spatial representations naturally afford reasoning in terms of spatial regions. For example, in the domain of COLOUR, one can identify a region that corresponds with the RED. Boundaries between regions are fluid, an aspect of the representation that may be usefully exploited by creative systems searching for new interpretations of familiar concepts. Similarly, there are various ways to tune musical scales: the pitch of a note can be moved a surprising distance from its perceptual centroid before it becomes a different note (and so we can sing "out of tune").

Gärdenfors identifies various types of regions with differing topological characteristics. *Convex* regions allow us to define *natural* properties:

CRITERION P A *natural property* is a convex region of a domain in a conceptual space. (Gärdenfors 2000, p. 71)

Again taking the example of RED in the domain of colour: given any two shades of RED, any shade between would also be RED. Therefore, the region corresponding to RED must be convex. These convex regions in conceptual domains can be closely related to basic human perceptual experience. The same is true of PITCH: there is a region (and therefore concept) for each note of the Western chromatic scale, and each region is convex, though allowing variation at the boundaries, depending on context. A C$\natural$ is very close to a C$\sharp$ in the context of purely scalic comparison: the two are directly next to each other on the spiral produced by the integral constraints of the domain's two dimensions.

For straightforward domains such as COLOUR and PITCH, we can think of concepts as natural properties. However, more complex concepts exist, over multiple domains. In music, the TONALITY domain, for example, is more problematic: its various regions are composed of different regions of PITCH linked together around a distinguished PITCH region corresponding with tonal centre; it is very hard to see how such a structure could be convex. In this context, a D$\natural$ (two chromatic notes away from C$\natural$) is sometimes—but not always, depending on tonal centre—much closer to a C$\natural$ than is C$\sharp$. To admit these more complex structures, Gärdenfors defines a *natural concept* as follows:

CRITERION C A *natural concept* is represented as a set of regions in a number of domains together with an assignment of salience weights to the domains and information about how the regions in different domains are correlated. (Gärdenfors 2000, p. 105)

Our interpretation of Criterion C is that a natural concept is a set of one or more natural properties and salience weights for the dimensions.

Gärdenfors' conceptual spaces, then, are hierarchical: they are constructed from primitives, using geometrical notions, the constructs that inhabit them are typed, and they can be viewed either as constructed or as Platonic, in the same sense as Thornton's concepts. Included in the hierarchy are higher-order, structured concepts (for example, shapes made from more than one other shape), and even truth-functional predicates such as "longer than", and it is shown how all these can be represented geometrically (Gärdenfors 2007). It is not clear how non-perceptual concepts (such as truth itself) would be conceptualised; perhaps such concepts exist at the symbolic level only. At its other end, the hierarchy bottoms out via properties into quality dimensions.

The theory also has a semantic component. Taking a strictly cognitivist stance, Gärdenfors argues that the generation (or the experience?) of semantics is strictly a property of mind/brains, and resists the notion that words have their own semantics. This position is perhaps reconcilable with the more common usage of "semantics" in linguistics if we allow that "the semantics of word W is S" is shorthand for "word W is a symbol which stimulates the experience of semantics S in the mind of the listener"; but in any case, semantics is in the mind, not in the world. This leads us to an understanding of the relationship between language, the symbolic level of Gärdenfors' theory and the semantics of the symbols in each in terms of conceptual spaces: words denote symbolic-level structures which are in correspondence with conceptual-level regions. It is not clear if this is Gärdenfors's intention: the issue is not addressed directly in the primary statement of the theory (Gärdenfors 2000), where the symbolic level seems almost to serve as a direct proxy for language. However, semantic grounding in the theory clearly arises from the quality dimensions, which correspond directly with perceptual mechanisms; thus, semantics is fundamentally grounded in perception of the world, as one might expect, with higher-level meaning being constructed from primitives. Below this and beyond the scope of the current discussion, the higher-dimensional subconceptual level models wetware implementation.

Two approaches to the mathematical formalisation of Gärdenfors' theory of conceptual spaces appear in the literature, both building on an initial formalisation by Aisbett and Gibbon (2001). One strand of research, based on fuzzy set theory, is presented in detail by Rickard et al. (2007a), drawing on previous work by Rickard (2006) and Rickard et al. (2007b). Another strand, employing vector spaces, is presented by Raubal (2004), with subsequent related work by Schwering and Raubal (2005a, b) and Raubal (2008a, b). Here, we follow a similar approach to Raubal's vector spaces formalisation. However, we have aimed to simply the mathematics as far as possible, by encoding only the minimal structures needed for the operation of the theory, and leaving pure nomenclature unstated.

Semantic distance between concepts is determined by calculating the distance between points in the space. As a rule of thumb, Gärdenfors suggests Euclidean distance is appropriate for integral dimensions, while the city-block metric is appropriate for separable dimensions (Gärdenfors 2000, pp. 24–26).

A prerequisite for the calculation of meaningful distances between concepts is the *normalisation* of the quality dimensions; the operation defines an appropriate mapping from a dimension's range, to the range $[0,1]$. So called *min-max* normalisation (Jain et al. 2005) can be appropriate where values are uniformly distributed over the dimension. However, otherwise, for example, when the bulk of the distribution is concentrated within a narrow range, min-max normalisation can result in poor discrimination between the most typically occurring values. Therefore, any decision on the propriety of a normalisation method should be informed by knowledge about the distribution of values in that dimension. An overview of normalisation techniques is provided by Jain et al. (2005). We use the technique of distributional distance, investigated in detail for computational music analysis by Müllensiefen (2009), whereby distance is defined in terms of a cumulative distribution function describing the distribution of values across the dimension. Distance is simply defined as the absolute difference between the cumulative distribution values corresponding to the individual values of the dimension. For two values, x_i and x_j, in quality dimension c, the distributional distance is thus:

$$d(x_i, x_j) = |F_c(x_i) - F_c(x_j)|, \tag{1.1}$$

where $F_c(x)$ is the appropriate continuous, or discrete, cumulative distribution function.

Distributional distance has the property of maximising the discrimination between values that occur most frequently. The psychological rationale for adopting this approach is that, through exposure, we become most attuned to detecting differences between stimuli that occur most frequently.

Towards a unified theory of conceptual space

Relationships between the theories

We now explore the relationship between the theories outlined in Section . We contend that Thornton's theory is located on Gärdenfors' top, symbolic level, while Boden's encompasses all levels but does not distinguish between them. Wiggins' CSF was originally motivated primarily at the symbolic level, but lends itself naturally to operation at the conceptual level as well. Thus, the CSF gives us a theoretical structure within which we can unify Thornton's theory of explicit creative construction with Gärdenfors' theory of conceptual cognition and semantics, perhaps fulfilling Boden's original conception more completely than before.

To distinguish the different uses of "conceptual space", we introduce the following superscripts, to differentiate between Gärdenfors' three levels of representation:

- ^{s} Symbolic level
- ^{c} Conceptual level
- ^{sc} Sub-conceptual level

Each suffix may be applied to a CSF set symbol in order to make explicit reference to elements of the CSF at a particular level of representation. For example, a subset of symbolically represented concepts in a conceptual space $\mathscr{C}$ would be denoted by $\mathscr{C}^s$, and the rules specifying this subset being denoted by $\mathscr{R}^s$. The ability to refer explicitly to layers of representation with different representational properties helps clarify the situation where multiple levels of representation are simultaneously available, and used, in $\mathscr{C}$. The metre space defined in our application, below, is one such.

To begin with the self-evident: Thornton's symbolic, constructive approach can be modelled in the CSF, at the symbolic level, by inclusion of appropriate symbols to express the various relations involved in the language $\mathscr{L}^s$, in which the creative system is expressed. Those relations can be made explicit in whatever $\mathscr{T}^s$ works over $\mathscr{C}^s$, and used in its reasoning. Whether we view the production of concepts as construction of new things, or traversal of a Platonically extant space of possibilities is immaterial, as before.

The next component is newly explicit: a mapping, $\mathscr{M}$, between concept symbols in $\mathscr{C}^s$ and concept regions in $\mathscr{C}^c$. It is worth noting here that we would expect this mapping, though partial, to be one-to-one, since (we suggest) $\mathscr{C}^s$ contains concepts and not words, so ambiguity would be representationally pathological. $\mathscr{M}$ gives the symbols in $\mathscr{C}^s$ semantics in terms of the implicit perceptual semantics of $\mathscr{C}^c$—and not *vice versa*, as might be expected in a symbolic AI formulation, where symbols are primary, serving as proxies for their own interpretation by the user.

An example of this relationship can be given using the colour space introduced earlier. $\mathscr{C}^s$ contains the concept RED, which is a symbol. It has several instances (in Thornton's terminology), including SCARLET, CRIMSON and VERMILION, each linked in $\mathscr{C}^s$ by the INSTANCE relation. Each one of these symbolic instances is mapped by $\mathscr{M}$ to a different convex region of $\mathscr{C}^c$, though the boundary between them may be fuzzy (Rickard 2006); and $\mathscr{M}$ maps RED, the more general categorically defined concept, to a larger region which contains all the instance regions. The relation in $\mathscr{C}^c$ corresponding with INSTANCE in $\mathscr{C}^s$ is geometrical inclusion, which comes as no surprise to the reader familiar with description logics (Brachman and Levesque 1985). At the more complicated level of natural concepts, the concept of CHORD can be captured as a combination of regions in PITCH, and then related to TONALITY by the addition of a tonal centre, as before. Here, however, mere inclusion is not the generalisation method, because simply taking the union of all the regions so defined would yield the entire domain. More research is required here: models such as that of Chew (2000) may help.

It is important to note that $\mathscr{M}$ is not absolute, and may vary across individuals. For example, language influences the development of perceptual categories (Roberson et al. 2006), and at this fundamental level, there are no concepts other than perceptual categories. Therefore, people from different cultures, at least, may have different $\mathscr{C}^s$, with correspondingly different $\mathscr{C}^c$ and $\mathscr{M}$, with correspondingly different linguistic distinctions, although some commonality will presumably co-exist at some level.

It seems, then, that Gärdenfors' conceptual level comfortably fits directly underneath the AI-style symbolic conceptual space, with $\mathscr{M}$ anchoring the meaning of concepts in $\mathscr{C}^s$ to the concepts in $\mathscr{C}^c$ which, according to Gärdenfors, have their own intrinsic perceptual semantics.

Since we now have a two-layer conceptual space, we must ask what happens to the other CSF components at the new $\mathscr{C}^c$ level. First, we have a new language, $\mathscr{L}^c$, in which the space of concepts is represented and manipulated. Gärdenfors allows for a unified, multidimensional space encompassing all quality dimensions, by the device of allowing each dimension an *undefined* value, and we take this approach too, maximising uniformity of the formalism and straightforwardly admitting the set $\mathscr{U}$ as the conceptual space generated by all possible quality dimensions, which seems to correspond with Gärdenfors' most general usage of "conceptual space".

The nature of $\mathscr{R}^c$, the rule set defining each conceptual space, is of particular interest. The introduction of basic perceptual concepts, such as colour, seems to require, simply, that $\mathscr{R}^c$ selects a region of the relevant domain, such as RED, which is more or less fixed, and amenable only to exploration, and not transformation. On levels involving more complex structure, however, transformation becomes possible, at least in principle. We return to this issue below.

Traversal rules, $\mathscr{T}^c$, need to be formulated, and, because the conceptual level is different in kind from the symbolic, its traversal rules are likely to be different in kind also. For example, in the conceptual space of colour, the idea of lightening a colour is hard to express in $\mathscr{C}^s$, because it is a continuous notion and $\mathscr{C}^s$ is by definition discrete, so transition in $\mathscr{C}^s$ is necessarily by step. But in $\mathscr{C}^c$ of colour, lightening (which is itself a relational concept) is simply vector addition, and so can be trivially applied. This raises the interesting question of *whether* and *when* a new symbol should be added to $\mathscr{C}^s$ ("I have created a new colour which I will think of as..."[2]), since that new colour was always implicitly extant in $\mathscr{C}^c$. Thus, the multi-layer $\mathscr{C}$ yields new detail of transformational creativity: transformations may perhaps be in one or other space, or in both; in turn, new kinds of aberration (Wiggins 2006a) are potentially available. In general, then, the geometrical nature of $\mathscr{C}^c$ may afford transitions which were not available in symbolic $\mathscr{C}^s$, and vice versa. Here, the theory models reality neatly: the symbolic level seems to model categorical perception appropriately, while maintaining an explanatory account of the phenomenon at the conceptual level.

A very interesting possibility, given the geometrical nature of $\mathscr{C}^c$, is to view traversal of the space as the production of an optimised trajectory; this is commensurate with the AI search view. However, the geometrical nature of the space admits the possibility of optimising this trajectory with methods more powerful than mere search. An example of the kind of traversal strategy one might use is given by Wiggins et al. (2009), where $\mathscr{C}^s$ is a set of melodic note sequences, and $\mathscr{C}^c$ is defined in terms of pitch and time spaces (as we illustrate below). Because $\mathscr{R}$ is supplied by a statistical model, in this work, Metropolis sampling can be used (MacKay 1998).

Evaluation rules, $\mathscr{E}^c$, perhaps become more accessible, since $\mathscr{C}^c$ affords a similarity metric. Maybe this can become a component of a theory of aesthetics, since it is at least known that familiarity (with what is already known, or, in our terms, conceptualised) drives aspects of preference (Zajonc 1968). Furthermore, a similarity metric over $\mathscr{C}^s$ may, in future, be derived from the combination of $\mathscr{C}^c$ and $\mathscr{M}$.

Transformational creativity in the geometrical $\mathscr{C}^c$ is interesting. We have already noted the duality between, on one hand, the notion of a symbolic level space which is constructed during traversal, and, on the other, one which is notionally Platonically extant, from which components are merely selected during traversal. Since Gärdenfors' geometrical spaces are in some sense more clearly Platonically extant than Boden's conceptual ones, on account of their strongly and regularly mathematical nature, we must ask whether it makes sense to think of constructing one in a way analogous to Thornton's symbolic construction. We must also ask whether there can in fact be a notion equivalent to Boden's transformation of such a space[3], the point being that the dimensions of the space are universal (in the same sense as Wiggins' universe, $\mathscr{U}$) so there is nowhere for them to be transformed to. The definition of a new concept in $\mathscr{C}^c$ consists in the identification of a new, bounded convex region; and since that region (like all the others) is continuous, and defined in terms of *betweenness*, that act of definition has much more of the flavour of labelling something already existing than of creating something new. However, this is perhaps unsurprising. Notwithstanding graphic designers' rhetoric, it is hard to see how one can *create* a new colour—rather, the creativity involved in such an act is in the *selection* of a new region of colour space, and perhaps in the choice of a name for it; or, more positively, the creativity may be personal to the individual (Boden's P-creativity) in that they have *imagined* a colour they have never seen before, and then perhaps mixed it, using paint. So, if we would not expect transformational creativity at this perceptual level, perhaps transformation of $\mathscr{C}^c$ would consist in transformation of the vector space defined by the quality dimensions themselves: perhaps the transformation from HSL colour space to RGB or CMYK (Raubal 2004) is a transformational creativity step, since it allows new things to be done with colour (displaying and printing, respectively, as opposed to seeing). To find stronger, more transformational creativity, it seems likely that we will need to search in the higher-level, relational concepts introduced in the more advanced parts of Gärdenfors' theory, and since those parts are still developing, a good intermediate strategy may be to locate transformational creativity as primarily a symbolic-level operation.

Nevertheless, Gärdenfors seems to have supplied all the components needed to instantiate the conceptual space in the CSF, and also to enrich the possibilities gained therefrom. On balance, it seems likely that we still need the symbolic level for combinations of perceptual concepts which are themselves more abstract than perceptual. To extend Thornton's uniform example: even if it is possible to represent the concepts of the physical objects SHIRT, JACKET and TROUSERS and their physical function, such as keeping warm, geometrically, it is still very hard to see how one can naturally represent the social function of these objects (precluding nudity, enhancing appearance, etc.) in the same way. When we begin to consider the highly abstract concept of UNIFORM, which is defined compositionally (in Thornton's sense) but also in terms of a social function which is not a function of the physical objects independently, and which is primarily a predication of the wearer, we begin to see that an abstract symbolic layer, albeit partly grounded in the cognitive semantics of $\mathscr{C}^c$, is still strictly necessary. Therefore our mapping, $\mathscr{M}$, will not cover the whole of both levels, $\mathscr{C}^s$ and

[2] Photoshop users evidently do this all the time: there is a tool for sampling the colour of a region, which allows the user to refer to "the colour of that region" without having a name for it. "The colour of that region" is, however, as much a symbolic concept as *taupe*, *turquoise* or *turnip*.

[3] It is clear that geometrical transformations can change the nature of the space itself, perhaps rendering it non-Euclidean or changing the magnitude of its dimensions, but this, in our understanding, is not the same kind of transformation as Boden's transformational creativity, which involves a change in the *content* of $\mathscr{C}^s$.

$\mathscr{C}^c$, with the consequence that there must also be a direct relationship, possibly via other layers, between $\mathscr{C}^s$ and $\mathscr{C}^{sc}$, the sub-conceptual level, so that the symbolic layer can be directly implemented in wetware where appropriate.

While Gärdenfors strives to represent more external-world semantics in his geometrical system (Gärdenfors 2007), it is perhaps worth considering a compromise, in which a creative system works at *both* the conceptual *and* the symbolic levels simultaneously, but synchronised and constrained by $\mathscr{M}$. Having the semantic information embedded under the symbolic structure of the knowledge representation is likely to enhance the process of creative and inductive inference, and may contribute to the discovery of more highly valued artifacts. In more advanced creative reasoning, metaphor is likely to figure strongly: many mathematicians report diagrammatic visualisation when they reason, and perhaps this can be characterised as a metaphorical mapping from $\mathscr{C}^s$ to $\mathscr{C}^c$, based on Gärdenfors' own theoretical approach to metaphor (Gärdenfors 2000, §5.4); however, we leave this can of worms on the shelf, for the moment.

Illustration: musical pitch and melody

By way of example, we outline how a $\mathscr{C}^c$ representation of musical tonal melody might be imagined in Gärdenfors' way of thinking, while maintaining symbolic conceptual representation in $\mathscr{C}^s$. Common terminology from Western music theory gives clues to the key symbolic concepts; and it has been shown elsewhere how a machine learning system can be used to simulate a process of creative composition in these terms (Wiggins et al. 2009). In our example, we ignore the quality dimension(s) associated with timbre (Caclin et al. 2006), which is currently not tractable, though we do not claim that this omission is musically adequate. It is, however, a useful feature of the Gärdenfors' conceptual spaces that one can neatly project out certain quality dimensions on which one wishes to focus.

We need just pitch and time to represent melodies in a cognitively realistic way (Wiggins et al. 1989); here, we will use metrical time, and thus filter out expressive aspects of performance. First, the sensation of pitch, as mentioned earlier, which, for the simplest representation of melody, can be viewed as a purely linear structure, modelled as a linear Abelian group. Each note of the standard Western scale corresponds with a segment of the one-dimensional conceptual space thus formed, with a centroid of each region at the accurately-tuned value for each note. The conceptual space of metrical time is substantially more complicated, and we consider it in detail below, but for the current example, we need only its basic group properties, which allow us to represent perceived sequence and *inter-onset interval* (IOI: the time between successive note beginnings) in in terms of (possibly low-integer-fractional) multiples of a consistent pulse (the *tactus*). These four quality dimensions fuse into one integral natural concept of NOTE. Music theory gives a name, and therefore a symbolic concept, to each and every one of these components, and for each valid value on each quality dimension.

Notes, though, are not enough to represent the experience of melody. For one thing, melodies can be shifted in pitch without changing their basic identity. For another, the relationships between consecutive notes in a melody define their harmonic function in the context of a scale. Both of these properties are enjoyed by a representation based not on pitch, but on pitch-difference or, in musical terms, *interval*, coupled with the notion of tonal centre, to which we return below. The conceptual space of MELODIC INTERVAL is, in Gärdenfors' terms, a relational one, in two senses: first, there must be a first note and second note (so a sequence relation is involved); and, second, there is a quality dimension, *pitch interval*, which corresponds exactly with distances between pairs of points on the pitch dimension, and so is also relational. A natural representation of a melody, given these quality dimensions, is as a sequence of intervals, paired with a sequence of IOIs. Again, music theory names each of these atomic concepts, and each value on each implicated quality dimension.

Given such a representation, we can abstract further, to notions such as *melodic contour*, where we replace the pitch intervals with magnitudeless indicators of direction (up, same, or down); contour has been shown empirically to be significantly implicated in melodic memory (Dowling 1978). Interestingly, at this point, traditional music theory ceases to supply terminology[4], presumably because the concepts have become too imprecise to be useful in practical discussion between musicians, notwithstanding their perceptual import: it is left to compilers of music reference books to construct a specialist notation, which seems to be used mostly in the context of database indexing (Parsons 1975; Lemström and Wiggins 2009).

At a level of representation which is in some sense further from the absolute notes, there is the level of *harmonic function*. Each tonal melody has a tonal centre (usually its last note, but in any case normally deducible from at most the first three or four notes), and each note in the scale, and each interval experienced in context of the tonal centre, has a different harmonic function—each harmonic function *sounds different* from the others in a way which is utterly inexplicable in other perceptual terms[5]. Arguably, it is these harmonic functions in sequence, and not the actual pitches,

[4] While "up" and "down" are standard terminology, *sequences* of them are not.

[5] For a tutorial on the difference between major and minor tonality, listen to "Ev'ry Time We Say Goodbye" by Cole Porter: the perceptual experience of tonal function is available to every listener, even if he or she has not studied music theory, nor become consciously aware of the differences in sound so produced.

that give the melody its character. Here we see a strongly relational concept (relating PITCH and/or MELODIC INTERVAL with TONAL CENTRE, which is itself strongly related to the circular dimension of PITCH mentioned above), that is also strongly perceptual, and that some theorists would argue contributes directly to any semantics music may convey. Again, music theory has terms and concepts for each quality dimension of harmonic function, and for each standard value on it. An open, perhaps unanswerable, question is that of whether an untrained listener, capable of perceiving harmonic function, but not of identifying it or its components, has those concepts at the symbolic level. Our position is negative: such symbolic-level concepts would be created by the act of identifying the functions in question (either as perceptual experiences or as theoretical constructs). Thus, concept-formation *per se* is potentially a weak kind of creativity, as one might expect.

In the context of musical creativity, only one of the above representations can be said to admit creation in itself. For example, the use of a non-standard pitch would be seen by a musician as exactly that, and not as the creation of a new pitch: it might be a creative *selection* in context of a melody, but the pitch itself would be seen as already extant (even if it were of a precise frequency never before used in any musical culture); equally, new time-divisions would be considered creative only in context. The clearly distinguished level at which one begins to exhibit music-compositional creativity is that of the sequences of notes or intervals mentioned above. It is at this level only, where one is not merely choosing a point in conceptual space bounded by given quality dimensions, but choosing a point that conforms to subtle and complex external constraints, which are evidently learned and culturally constructed, that we can make a clear distinction between, on one hand, an $\mathscr{R}^c$ which selects exactly a natural concept (e.g., RED[6]) and, on the other, an $\mathscr{R}^c$ which selects a *domain* (in Gärdenfors' terms), from which a more specific concept may be chosen (by its own $\mathscr{R}^c$): as the concept of MELODY is chosen from within the domain of NOTE SEQUENCES. There are points in the domain constructed by these quality dimensions which do not conform to the constraints required of a tonal melody: for example, certain notes and certain intervals are precluded by the tonal centre; and melodies are expected to end on one of three particular scale notes—indeed, usually, the tonal centre itself. This region is certainly not convex, for there are holes in it. And here, finally, is a justification for $\mathscr{R}^c$, a CSF rule set which defines a Boden-style conceptual space of melody as a region of the larger Gärdenfors-style quality-dimensional domain of note sequences. Here, perhaps, is the point at which friction between the edges of the two theories produces the most interesting sparks: maybe the boundary between these two kinds of conceptual space corresponds with the elusive boundary between perceptual and cognitive phenomena.

[6] In Gärdenfors' terms, this is a natural property, too, but since we are considering it, here, in its own right, and not as part of a more complicated concept, of something which might have RED as a property, we think of it as a concept.

Summary

In this section, we have attempted to show, at a discursive level, how Gärdenfors' theory of conceptual spaces may instantiate Wiggins' Creative Systems Framework, and extend it in useful ways. Drawing the two theories together has led us to examine each in new light. Gärdenfors' theory, in particular, yields many interesting questions when thus dissected, which we have attempted to highlight along the way.

A fully unified theory of conceptual space offers the field of computational creativity several advantages: a coherent, explanatory basis for modelling of Boden's abstract specifications of creative cognition; a framework which can be implemented directly, at least in part; a means by which the creation of semantics may be explicable; and the beginnings of a route to explication of the mysterious evaluation rule set, $\mathscr{E}$.

In the next section, we show how a leading theory of musical rhythm can be brought to bear in this framework to yield a conceptual space suitable for exploration by the methods outlined above.

Application: Musical Metre

Preamble

We now present a formalisation of Gärdenfors' theory of conceptual spaces (Gärdenfors 2000) in the domain of musical metre. The primary motivation underlying this work is the assumption that, in order for creative systems to begin to emulate human creative behaviour, they require rich knowledge structures comparable to those of their human counterparts. Inevitably, this is no trivial problem, even within very confined domains. We conjecture that perceptual groundedness is one of the most important aspects for any representation over which a creative agent operates. The importance of similarity in mental processing is long established (Shepard 1987), and is the guiding principle underlying the computational theory proposed here. In short, representations that afford efficient and flexible manipulation of concept similarities may prove useful in the pursuit of machine simulated creativity.

Music, and art in general, are interesting application domains in which to investigate Gärdenfors' theory of conceptual spaces, not least because of the primacy of subjective experience within them. Music is notoriously difficult to describe with language, although humans have very little difficulty distinguishing between music and non-music when heard. Music exists over time, and there is a delicate interplay between what has gone before, and what might come next (Pearce and Wiggins 2006). Furthermore, all musical experiences are shaped by past musical experiences. In principle, the theory of conceptual spaces seems to offer a viable approach for representing such fluid, yet richly structured, phenomena. Below, we address questions concerning relational and temporal concepts, and draw from work carried out in the field of music psychology in order to define a perceptually grounded representation of some very basic musical concepts.

In principle, an approach to the representation of music based on conceptual spaces should not need to be confined to any one specific conceptualisation of music. In fact, the conceptual spaces theory itself supports an elegant model of learning, which accords with evolutionary views of musical development (Bown and Wiggins 2009), in which the process of developing understanding of unfamiliar concepts is modelled by extending a conceptual space with additional quality dimensions, affording greater discrimination between novel stimuli. Furthermore, the notion of dimensional salience, modelled by weightings associated with each quality dimension, allows for the possibility of adapting conceptual spaces to take into account individual musical backgrounds and experience. Therefore, it is assumed that differing conceptualisations of music, such as those evident across Western classical or pop music, Balkan folk music, or Ghanian drumming (Patel 2008, pp. 97–99), can be represented consistently within the conceptual spaces theory.

Our primary reference is Justin London's authoritative account of musical metre (London 2004). London provides a detailed theory of musical time, drawing together a range of insights from music theory, musicology, and psychology. Importantly, the theory offers considerable generality as a result of its foundation upon basic human perceptual and physiological constraints, and provides many examples from both Western and non-Western musical traditions. Of course, the experience of music as a whole is primarily dependent on cultural context, and as such can radically differ between cultures, and between individuals within cultures. However, at a very basic level of music conceptualisation, such as the experience of periodicity, London argues that commonality across many musical practices can be found. The present computational theory similarly concentrates on low-level musical concepts, addressing some of what might be considered as primitives of musical conceptualisation.

Grounding in musical representation

Gärdenfors' initial theory of conceptual spaces concentrated primarily on tangible properties and concepts, where quality dimensions typically relate to attributes directly available to our sensory system, for example, the colour of apples. Although musical phenomena are closely linked with physical events in the world, which are experienced via the senses, the *experience* of musical stimuli cannot be equated with the physical stimuli itself. One consequence of this for any representation of musical experience is the necessity for relatively abstract quality dimensions—relative at least to the dimensions required to represent tangible physical properties or concepts. Despite a long tradition within musicology to concentrate on notationally objective musical structure, typically derived from Western notation, Gabrielsson (1993) points out that there is very often broad agreement between music theoretical and music psychological studies. Therefore, much insight into music conceptualisation can be gained from music theory, and usefully for the theory of conceptual spaces, music theory offers a vocabulary for distinguishing between what may be relatively abstract concepts, and provides clues as to possible structures of quality dimensions within which they may be represented—which may then be tested empirically.

In the case of music—that is, music as *perceived*—the issue of representational grounding is mediated by subjective experience, since the phenomenon itself, or any 'objects' which one might consider meaningful, are psychological or cultural in nature. To be explicit, the purpose of the representational theory pursued here should not be confused with the representation of musical scores or of physical musical sound. In both these cases, which are sometimes confusingly referred to as 'music', there are clear concrete referents, whose correspondence with the cognitive constructs may not be straightforward. The aim of the following formalisation is precisely to capture the cognitive constructs associated with musical experience.

Towards a conceptual space of musical time

The conceptual space described in this section captures musical metre, rather than all aspects of musical timing. A common distinction made in the literature is between musical metre and rhythm (more generally: serial and periodic concepts, as discussed below), although there is debate over the extent to which they can be treated independently (Benjamin 1984; Cooper and Meyer 1960; Hasty 1997). London (2004, p. 4) defines rhythm as involving 'patterns of duration that are phenomenally present in the music'. Duration here refers not to note lengths, but to the *inter-onset interval* (IOI) between successive notes. Rhythm is therefore a theoretical construct describing the arrangement of events in time. However, this objective description does not necessarily accord with perceived musical structure.

The perceptual counterpart to rhythm is metre:

> [M]etre involves our initial perception as well as subsequent anticipation of a series of beats that we abstract from the rhythmic surface of the music as it unfolds in time. In psychological terms, rhythm involves the structure of the temporal stimulus, while metre involves our perception and cognition of such stimulus. (London 2004, p. 4)

The perception of metre can, therefore, be considered a form of categorical perception, where the surface details of the temporal stimuli, such as the particular structure of the rhythmic pattern, or any expressive performance timing, are conceptualised with reference to a hierarchical organisation of regular beats, itself induced from the stimuli. Rhythmic patterns also tend to be characterised by a bimodal distribution of 1:1 and 2:1 duration ratios (Clarke 1999), which typically coincide with metrical locations (thus making syncopation possible (Steedman 1977)) and further emphasising the interrelatedness of the concepts of rhythm and metre.

It is also necessary to draw a distinction between the different timescales that operate within music. Relationships across time may be comprehended on all levels of musical organisation, from IOIs lasting a few hundred milliseconds, to relationships between patterns of notes spanning entire works. A boundary between rhythm and form is usually defined as being the duration of the perceptual present: up to about 10 seconds (Fraisse 1978; Clarke 1999). The comprehension of form is considered to require deliberate cognitive effort involving long-term memory. Below we consider only concepts that are bounded by the temporal extent of the perceptual present, reserving larger-scale musical concepts for future research.

Formalising a conceptual space of musical rhythm

A minimal conceptual space representation of rhythmic structure would seem to require at least the ability to represent rhythmic stimuli at the level of the perception of events in time, as well as more abstract and stable concepts, which emerge from the event level, and persist over the extent of the perceptual present. The lower level might be thought of as representing cognitive primitives, which are strongly grounded concepts, closely associated with direct sensory experience, such as the basic perception of time intervals, or the sensation of periodicity. As the level of abstraction increases, conceptualisation becomes less rooted in perception, and concepts such as tempo, or grouping, may emerge, and so on to still higher culturally specific concepts. As the level of abstraction changes, it is assumed that the quality dimensions of the space may also need to change, thus requiring the ability to map points between alternative sets of dimensions, as discussed by Raubal (2004). It may become impossible to represent some concepts adequately within quality dimensions, with symbolic representations proving more appropriate, as discussed above. In any case, both points and symbols representing higher level concepts can be understood as mapping to regions of lower level concepts. This paper focuses on lower level concepts, since the definition of the structure of this level is a prerequisite for building representations of more abstract concepts.

Following Parncutt (1994), primitive level rhythmic concepts are here divided into two classes: *periodic* and *serial*. Periodic and serial concepts represent qualitatively distinct sensations arising from the same physical stimulus pattern of IOIs. For the present, we consider only periodic concepts, which is possible by assuming an independence between the two, as is common in the literature (Lerdahl and Jackendoff 1983; Povel 1984; Parncutt 1994). However, perception is known to be more stable when the boundaries of periodic and serial groundings coincide, and, indeed, much of the interest in rhythmic patterns is the result of the interaction between the two (Lerdahl and Jackendoff 1983). Future work is necessary to address the geometrical representation of serial concepts, which present particular challenges to the conceptual space theory due to their strictly sequential nature.

Periodic concepts

Periodic concepts are essentially the building blocks of *metre*, or ‘regular temporal structure’ (Steedman 1977, p. 555). Metre can be defined as the grouping of perceived beats or pulses into equivalence classes, which is typically expressed as the ‘regular alternation of strong and weak beats’ (Lerdahl and Jackendoff 1983, p. 12). The basic distinction between periodic and serial concepts is that periodic concepts depend on the ‘relative timing and perceptual properties of *nonadjacent* events’ (Parncutt 1994, p. 412), rather than consecutive events. Metre can also be thought of as concerning durationless points in time, whereas serial concepts inherently concern the relationships between events of specific duration (Clarke 1999, p. 478).

The periodic nature of metre means that it can be represented graphically in the form of a circle. Following the convention developed by London (2004, p. 64–69), time flows clockwise, and the dots on the circumference mark peaks of *attentional energy*. The 12:00 position marks the downbeat. In Fig. 1a, two levels of metric organisation are represented—the total time-span of the cycle, or *measure period*, and the intervals between the individual beats. Three levels of metre are represented in Fig. 1b—the measure period, the beat level, and the duple subdivision of the beat.

The total number of dots around the circumference of a circle defines the *cardinality* of the metre. This cyclic component, referred to as the *N-cycle*, is the lowest level (fastest

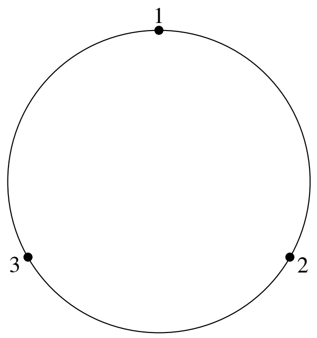

a: $\frac{3}{4}$ with no beat subdivision.

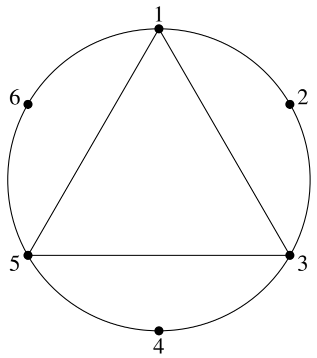

b: $\frac{3}{4}$ with quaver beat subdivision.

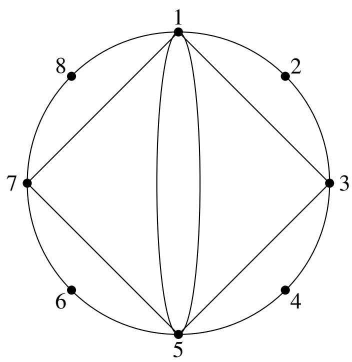

c: $\frac{4}{4}$ with quaver beat subdivision.

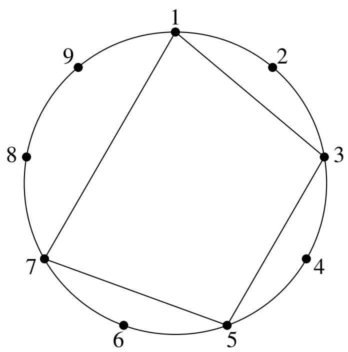

d: Non-isochronous 2-2-2-3 beat pattern

Fig. 1: Cyclical representation of metre, after (London 2004, pp. 64–69).

moving) cycle in any metrical hierarchy. The concept of the N-cycle can be used as a basis for distinguishing individual *metrical types* (London 2004, pp. 73–75). Figure 1a is therefore a 3-cycle metre, which in this case is an N-cycle component that also corresponds to the *beat cycle* or tactus. Figure 1b is a 6-cycle, which in this case corresponds to a subdivision of the beat-cycle. It is more common for a beat-cycle to be a *subcycle* of an N-cycle, because most metres include at least one level of subdivision (London 2004, p. 35). We can therefore refer to the metre represented in Fig. 1b as a 6-cycle metre, with a 1-3-5 component. Metres may have further levels of organisation, as shown in Fig. 1c, which contains four levels of periodic motion, and may also include non-isochronous cycles, as in Fig. 1d, where the beat cycle consists of three short beats followed by one long beat.

One further piece of information is needed to specify a metre fully within this framework: the measure period. Specifying the total time interval of a metrical structure determines the IOIs between the timepoints of the component cycles, resulting in a *tempo-metrical type* (London 2004, pp. 76–77). Drawing on the psychological literature, London defines the maximum period of the N-cycle as between 5 and 6 seconds, and the IOI between timepoints on the N-cycle as at least 100 milliseconds. This range defines the '*temporal envelope* for metre' (London 2004, p. 27). The range of tactus IOI is from 200 ms to 3000 ms (corresponding to 20–300 beats-per-minute), with a preference around 600 ms (100 bpm). Note that these limits are not arbitrary constraints, defined in order to simplify the representation of metre: they are derived from empirical perceptual and cognitive limitations. Such quantitative understanding of perceptual phenomena is central to our present purpose of constructing perceptually valid conceptual space representations.

In addition to the above constraints on the N-cycle, London provides further constraints that apply to the internal subcycle structure of metre. The following summarises the set of metric well-formedness constraints (London 2004, p. 72).

WFC 1: IOIs between time points on the N-cycle must be nominally isochronous, and at least $\approx$ 100 ms.

WFC 2: All cycles must form a closed loop.

WFC 3: All cycles must have the same phase.

WFC 4: All cycles must span the same amount of time (the measure period), which should not be greater than $\approx$ 5 seconds. In the formalisation below, we define the upper limit of the measure period as 6 seconds.

WFC 5: Each subcycle must connect nonadjacent time points on the next lowest cycle.

WFC 6: All subcycles must be maximally even (or as maximally even as possible in some non-isochronous metres). The principle of maximal evenness requires that all time points of a cycle be as evenly distributed as possible, which avoids ambiguous or pathological metrical structures (London 2004, p. 103).

The above constraints on metric well-formedness define a large space of possible metres, which London is confident encompasses the vast majority of metres present across all musical cultures (London 2004, p. 114). Importantly, the constraints allow us to exclude from consideration the much larger space of all possible hierarchical cyclic structure that do *not* correspond with a subjective experience of metre. In the following, we define a conceptual space representation of well-formed metres.

The domain of isochronous metre

We now construct a domain of isochronous metre, following London's specification, working from the bottom up. Individual cycles may be either isochronous, or non-isochronous. Here we will focus on isochronous cycles, leaving non-isochronous metres for future work.

First, we define a PULSE_IOI quality dimension, to represent the IOI between the isochronous time points of a cycle, measured in milliseconds, and defined over 100–6000 ms:

$$\text{PULSE\_IOI} = \{x \in \mathbb{R} \mid 100 \leq x \leq 6000\}. \tag{1.2}$$

We use the usual $<$ from $\mathbb{R}$ to provide an order over PULSE_IOI.

The range of PULSE_IOI accords with London's definition of the temporal envelope for metre (London 2004, p. 27), and there should be greater discrimination between lower values of PULSE_IOI than between higher values. To achieve this, we normalise the IOI values using a distributional distance function, S, based on Parncutt's pulse-period salience (PPS) function, shown in (1.3) and Fig. 2a (Parncutt 1994, p. 438).

$$PPS = \exp\left\{-\frac{1}{2}\left[\frac{1}{\sigma}\log_{10}\left(\frac{\text{PULSE\_IOI}}{\mu}\right)\right]^2\right\}. \tag{1.3}$$

We set the parameter $\mu = 600$, defining 600 ms (100 bpm) as the most salient pulse rate. We also use the typical value $\sigma = 0.2$, which concentrates the bulk of the distribution in the 200–2000 ms range. $S_{\text{PULSE\_IOI}}$ is defined as the corresponding cumulative distribution function, shown in Fig. 2b. The distance between points x and y is calculated thus[7]:

$$x -_{\text{PULSE\_IOI}} y = |S_{\text{PULSE\_IOI}}(x) - S_{\text{PULSE\_IOI}}(y)| \tag{1.4}$$

where vertical bars denote absolute value.

[7] We adorn functions specialised to a dimension and/or domain with a subscript identifying that dimension. Since these dimensions are subsets of infinite sets, these operations may not be everywhere defined.

Next, we define CARDINALITY, a quality dimension representing the cardinality of a metrical cycle, defined over the range $[1..60]$ timepoints, where only integers are allowed:

$$\text{CARDINALITY} = \{x \in \mathbb{Z} \mid 1 \leq x \leq 60\}. \tag{1.5}$$

CARDINALITY is defined to accommodate the theoretical minimum and maximum number of timepoints that may be present on a metrical cycle, according to the metrical well-formedness constraints (WFCs; see above). London (2004, pp. 68–69) employs this term principally with reference to N-cycles; however, in order to improve the uniformity of the resulting formalism, here the term can be applied to metrical cycles at any level in the hierarchy. We include 1 as a degenerate case, corresponding to the measure period. The CARDINALITY value 2 corresponds with the tactus cycle of the most basic metre of alternating strong and weak beats, with no subdivision. This metre only has two hierarchical levels of organisation—the measure period, and a 2-cycle beat cycle, which also corresponds with the N-cycle. However, since it is much more common for the beat-cycle also to contain at least one level of subdivision (London 2004, p. 68), a value CARDINALITY < 4 is unlikely because 4 is the minimum number of time points necessary for three levels of metrical organisation: 4-cycle (N-cycle); 2-cycle (beat-cycle); and the 1-cycle (measure period). The maximum CARDINALITY, 60, is the theoretical maximum number of periodic timepoints perceptible within the temporal envelope of metre, derived by assuming a minimum PULSE_IOI value of 100 ms, within a maximum measure period of 6000 ms; this tight relation between CARDINALITY and PULSE_IOI illustrates their integral nature in this context. The ordering, $<$, on CARDINALITY is the usual one over $\mathbb{Z}$.

The simplest normalisation approach for CARDINALITY would be min-max normalisation (Jain et al. 2005). However, a linear mapping from the range $[1..60]$ to $[0,1]$ presupposes that the perceptual notion of distance remains constant over the range. Intuitively this is not the case, since we are more likely to discriminate more finely between lower values of CARDINALITY, which will also typically span shorter time periods. Furthermore, since the largest value of CARDINALITY for any cycle in a metre is always the N-cycle, and all other cycles are subsets of the N-cycle according to the WFCs, there will necessarily be a higher density of lower CARDINALITY values across any corpus of well formed metres. Therefore, in the absence of a psychologically tested normalisation function, a distribution-based distance measure may serve as an appropriate approximation. Ideally, such a distribution should be derived from a corpus containing detailed and perceptually verified metrical annotation, which unfortunately is not available to us. As an alternative, we simple scale the range logarithmically, using (1.6), shown in Fig. 3. This simple normalisation approach per-

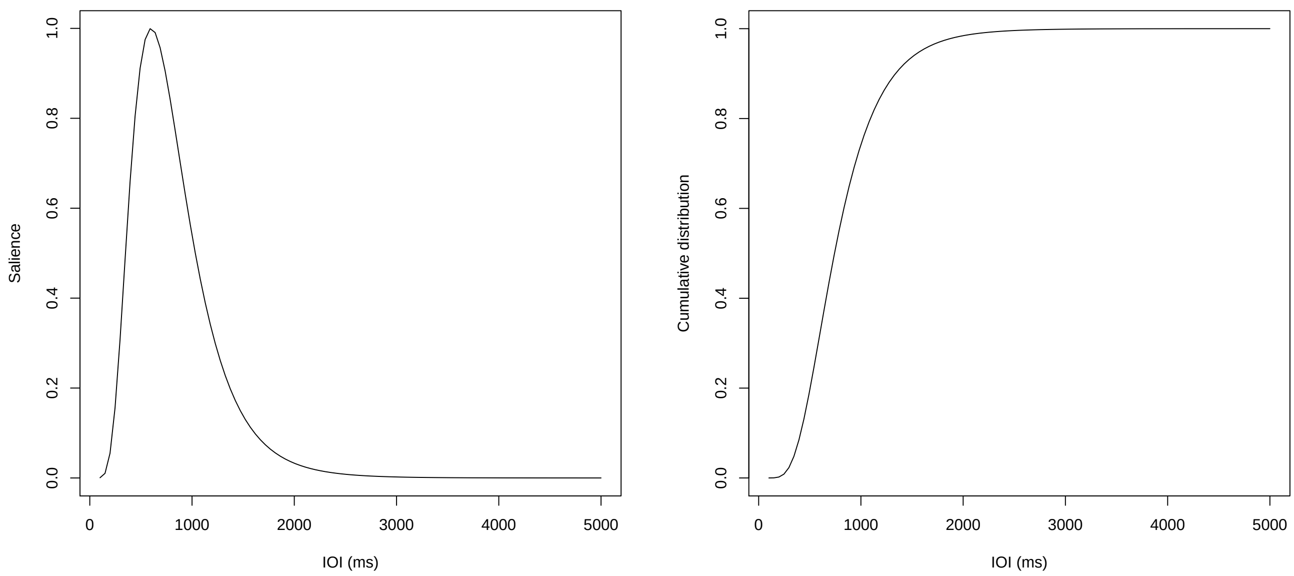

a: Pulse period salience, $\mu = 600$, $\sigma = 0.2$

b: Cumulative distribution for pulse period salience.

Fig. 2: Distribution-based normalisation for PULSE_IOI, based on Parncutt's pulse-period salience (Parncutt 1994, p. 438).

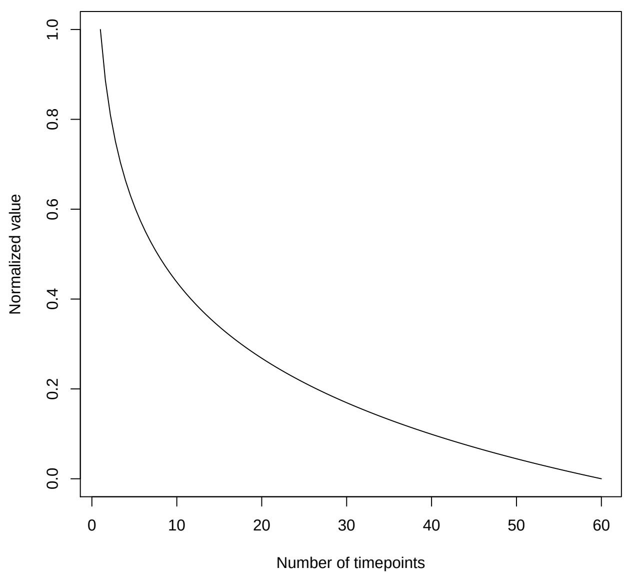

Fig. 3: Normalisation function for CARDINALITY.

formed adequately during subsequent empirical testing (see Discussion section below).

$$S_{\text{CARDINALITY}}(x) = 1 - \frac{\log(x)}{\log(60)} \tag{1.6}$$

Calculating the distance between dimension values is defined analogously to (1.4):

$$x -_{\text{CARDINALITY}} y = |S_{\text{CARDINALITY}}(x) - S_{\text{CARDINALITY}}(y)| \tag{1.7}$$

To reconstruct London's definitions, we first require a domain of PERIODICITY, as follows:

$$\text{PERIODICITY} = (\text{PULSE\_IOI} \times \text{CARDINALITY}) \cup \{\top\} \tag{1.8}$$

The value $\top$ denotes "undefined". The ordering of PERIODICITY is defined in terms of PULSE_IOI:

$$\langle p_i, c_i \rangle <_{\text{PERIODICITY}} \langle p_j, c_j \rangle \text{ iff } p_i <_{\text{PULSE\_IOI}} p_j \tag{1.9}$$

$$\top <_{\text{PERIODICITY}} \langle p_i, c_i \rangle \tag{1.10}$$

where $\langle p_i, c_i \rangle \in$ PERIODICITY.

Given PERIODICITY, we construct a domain comprising 9 PERIODICITY domains, which we call L-METRE, to distinguish it from our preferred METRE representation, below, as follows:

$$\text{L-METRE} = \text{PERIODICITY} \times \text{PERIODICITY}^4 \times \text{PERIODICITY}^4 \tag{1.11}$$

We divide the 9 PERIODICITY domains into 3 groups of 1, 4 and 4 respectively, for notational clarity. The first, a single domain, denotes the tactus (main beat), which must always be defined, the first vector of 4, whose first value must always be defined, denotes metrical levels larger than the tactus, higher metrical levels to the right, and the second vector, all of whose values may be undefined, denotes metrical levels lower than the tactus, lower metrical levels to the right. For both tactus groupings and subdivisions, 4 is the maximum number of levels possible given the metrical WFCs. There are further constraints: in either group, moving from left to right, undefined can only be followed by undefined; and, again moving from left to right, each element of a group must be a non-unity integer multiple of the previous one. Finally, there are salience weights on each dimension, but we simplify the notation slightly by omitting these here. Thus, L-METRE is a natural concept, because it conforms to Criterion C.

Given this definition, for example, the Western metre $\frac{3}{4}$ at 100 beats per minute (bpm), subdividing into quavers (giving a repeating group of 3 groups of 2 pulses), would be denoted by

$$\langle\langle 600,3\rangle,\langle\langle 1800,1\rangle,\top,\top,\top\rangle,\langle\langle 300,6\rangle,\top,\top,\top\rangle\rangle$$

while $\frac{6}{8}$ (which has a repeating group of 2 groups of 3 pulses) would be

$$\langle\langle 600,2\rangle,\langle\langle 1200,1\rangle,\top,\top,\top\rangle,\langle\langle 200,6\rangle,\top,\top,\top\rangle\rangle$$

However, this representation leaves something to be desired, because of the absolute values used. A natural way to represent these levels is relative to each other (as in the text, above). It turns out, also, that for a natural geometrical representation to model perception accurately, one also needs to represent the relationships between cycles. To see this, compare the representation of $\frac{3}{4}$, above, with that of $\frac{2}{4}$ and $\frac{4}{4}$, at the same tempo:

$$\begin{array}{ll} \frac{2}{4} & \langle\langle 600,2\rangle,\langle\langle 1200,1\rangle,\top,\top,\top\rangle,\langle\langle 300,4\rangle,\top,\top,\top\rangle\rangle \\ \frac{4}{4} & \langle\langle 600,4\rangle,\langle\langle 1200,2\rangle,\langle 2400,1\rangle,\top,\top\rangle,\langle\langle 300,8\rangle,\top,\top,\top\rangle\rangle \end{array}$$

We would like a Euclidean (or near-Euclidean) space for our representation. But $\frac{2}{4}$ is substantially closer, perceptually, to $\frac{4}{4}$ than it is to $\frac{3}{4}$, and any Euclidean notion of subtraction will not work here. To solve this problem, we construct a new dimension CARDINALITY_RATIO, to represent the relationship between two CARDINALITY dimensions:

$$\text{CARDINALITY\_RATIO} = \{2,3\} \cup \{\top\} \tag{1.12}$$

The CARDINALITY_RATIO value for a cycle is calculated in terms of its containing cycle, and thus ultimately the N-cycle, for which CARDINALITY_RATIO is undefined. Therefore, there will be one fewer defined CARDINALITY_RATIO dimensions than the number of CARDINALITY dimensions, as illustrated by the following values for $\frac{12}{8}$ with semi-quaver beat subdivision:

$$\begin{array}{ll} \text{CARDINALITY} & \langle 24,12,4,2,1\rangle \\ \text{CARDINALITY\_RATIO} & \langle \top,2,3,2,2\rangle \end{array}$$

We constrain the values of CARDINALITY_RATIO to equivalence classes 2 and 3, thus constraining the regions in the resulting domain of METRE, defined below; these restricted values are chosen following Lerdahl and Jackendoff (1983), and so model metre in Western tonal music. The formalism can be extended by adding further small prime numbers, as supported by empirical data from other cultures, as necessary.

Given this new dimension, we define

$$\begin{aligned} \text{PERIODICITY\_RATIO} = (&\text{PERIODICITY} \times \\ &\text{CARDINALITY\_RATIO}) \\ &\cup\{\top\} \end{aligned} \tag{1.13}$$

Now, a minimal conceptual space of METRE, capturing London's notion of periodic flow of attentional energy, can be constructed thus, by analogy with (1.11):

$$\begin{aligned} \text{METRE} = {} & \text{PERIODICITY\_RATIO} \times \text{PERIODICITY\_RATIO}^4 \times \\ & \text{PERIODICITY\_RATIO}^4 \end{aligned} \tag{1.14}$$

Now, we are in a position to compare values. To do so, we must define the result of subtracting, or subtracting from, $\top$: this is a positive value near 0, which we write ε; however, $\top - \top = 0$. In the new conceptual representation, the three example duple metres discussed above are denoted as follows:

$$\begin{array}{ll} \frac{2}{4} & \langle\langle\langle 600,2\rangle,2\rangle,\langle\langle\langle 1200,1\rangle,2\rangle,\top,\top,\top\rangle,\langle\langle\langle 300,4\rangle,\top\rangle,\top,\top,\top\rangle\rangle \\ \frac{3}{4} & \langle\langle\langle 600,3\rangle,2\rangle,\langle\langle\langle 1800,1\rangle,3\rangle,\top,\top,\top\rangle,\langle\langle\langle 300,6\rangle,\top\rangle,\top,\top,\top\rangle\rangle \\ \frac{4}{4} & \langle\langle\langle 600,4\rangle,2\rangle,\langle\langle\langle 1200,2\rangle,2\rangle,\langle\langle 2400,1\rangle,2\rangle,\top,\top\rangle, \\ & \qquad \langle\langle\langle 300,8\rangle,\top\rangle,\top,\top,\top\rangle\rangle \end{array}$$

The Euclidean distance between $\frac{2}{4}$ and $\frac{3}{4}$ is 1.02; that between $\frac{2}{4}$ and $\frac{4}{4}$ is $\sqrt{0.086+\varepsilon^2}$ (which is less than 1); and that between $\frac{3}{4}$ and $\frac{4}{4}$ is $\sqrt{1.07+\varepsilon^2}$ (which is greater than 1). These distances rank in the same order as the corresponding perceptual and musicological distances.

Discussion: Isochronous metre

Equipped with the basic framework outlined above, it is now possible to represent well-formed isochronous metres within perceptually grounded quality dimensions. Here we will provide examples of how various metres, common to Western music, appear within the conceptual space, and will also show how it is possible to calculate the distance between sets of points in the space of METRE, which corresponds to a plausible notion of similarity.

Figure 4 shows two results from a series of investigations designed to visualise the semantic distances between points in the multidimensional conceptual space of METRE. The space of METRE comprises 27 component quality dimensions, organised into sub-domains; however, in practice, many quality dimensions remain undefined. Various datasets were created, consisting of a variety of metres designed to illustrate different aspects of higher-level conceptual similarity. Pairwise distance matrices were calculated for each dataset, and then projected into 2- and 3-dimensional spaces using classic (parametric) *multidimensional scaling*.[8] Here we present two such examples, which demonstrate that a plausible notion of conceptual similarity is maintained when comparing various metres at both the same and across a range of tempi. A full empirical study involving human subjects is the only way to evaluate this computational model, but this method provides insightful and instant feedback, which is particularly appropriate given the exploratory nature of this research.

Tactus was held constant in Fig. 4a, with the aim of visualising the distances between a range of equal tempo metres common to Western music. It is difficult to give a simple interpretation in perceptual terms of the dimensions of this projection; however, meaningful clusters of metres can be observed. The space reveals two distinct regions divided by the diagonal between the x and y axes. The region $x < -y$ contains compound metres (triplet subdivisions of the tactus pulse), and $x > -y$ contains all simple metres (duple subdivision of the tactus). Within each region, further clusters of metres can be observed corresponding to duple and triple groupings of the tactus beat. Within the simple metre region, $\frac{2}{4}$, $\frac{4}{4}$ and $\frac{3}{2}$ (one, two and three groups of two tactus beats, respectively) form the duple cluster. The second cluster corresponds to metres with triple tactus groupings: $\frac{3}{4}$ and $\frac{6}{4}$ (one and two groups of three tactus beats, respectively). The z dimension provides further discrimination within clusters, which can be interpreted as reflecting the higher level groupings of tactus groupings. For example, $\frac{2}{4}$ is approximately equidistant from $\frac{4}{4}$ and $\frac{3}{2}$ on the z axis (although is slightly closer to $\frac{4}{4}$). $\frac{2}{4}$ has no grouping above the tactus grouping level, except the implicit measure period. $\frac{4}{4}$ has a duple grouping of tactus groups, while $\frac{3}{2}$ has a triple grouping of tactus groups. Therefore, the distance between $\frac{4}{4}$ and $\frac{3}{2}$ is greater than the distance between $\frac{2}{4}$ and either $\frac{4}{4}$ or $\frac{3}{2}$, which accords with musicological intuition.

Figure 4b visualises how the distance between three simple metres changes over a range of different tempi. Regions corresponding to the three metres are clearly evident in the projection, with $\frac{2}{4}$ being closer to $\frac{4}{4}$ than $\frac{3}{4}$, as would be expected. The distance between individual tempi is nonlinear, which is a consequence of normalising the range of PULSE_IOI according to pulse period salience. The distance between 80 bpm and 120 bpm here appears larger than the distance between 120 bpm and 160 bpm, which is consistent with the intuition that perceived dissimilarity should be greater between more highly salient periodicities.

[8] MDS was carried out using the `cmdscale` function from the R statistical package (R Development Core Team 2010)

Conclusion

Creativity is manifest across a wide range of human endeavour, and real progress is being made within the computational creativity community to understand this seemingly fundamental, yet elusive, capacity. The work of Boden (2004, 1998) has become a cornerstone on which much of this work resides, and subsequent formalisation (Wiggins 2006a, b) has revealed further convincing insight into the nature of creative processes. However, many issues still remain. In this paper we have attempted to pursue one approach to the problem of meaning within creative systems. To this end, we have shown how a geometry-based knowledge representation, as proposed by Gärdenfors (2000), can be usefully employed within Wiggins' (Wiggins 2006a, b) creative systems framework, as one possibility for establishing a perceptually grounded semantics.

We also discussed the relationship between Gärdenfors' geometric representation and traditional symbolic representations from the perspective of creative systems. Gärdenfors' theory, which is situated at the conceptual level, provides a link between sub-symbolic and symbolic representations. We formally defined this relationship within the CSF, and explored some of the implications for reasoning over both conceptual and symbolic representations. It appears that the two levels are not only compatible, but also greatly complementary. We involved Thornton's (2007) representational theory, which provides a means of describing relations between concepts, as a further example of how mappings between representational levels can support efficient reasoning.

Building on Gärdenfors' theory of conceptual spaces, we have attempted to construct the beginnings of a geometrical formalisation of basic musical concepts. We have provided descriptive examples of a range of musical concepts, including pitch, duration, melody and harmony, and how they might be instantiated in a musical conceptual space. We have also provided a detailed conceptual spaces formalisation of musical metre. Even with a high degree of domain knowledge, this problem proves far from trivial. Metre is only one aspect of music's temporal structuring, and the specification of anything approaching a complete conceptual spaces representation of 'musical time' is some way off. However, progress has been made, which is an essential first step on the road to richer levels of conceptualisation. Given a metrical representation of the kind developed here, we can consider how the geometry of the space can inform the process of creative traversal, and thus perhaps build systems which are more musically creative in a more human-like way.

Most importantly, the present work fulfils and adds to Gärdenfors' laudable aim of generating testable hypotheses,

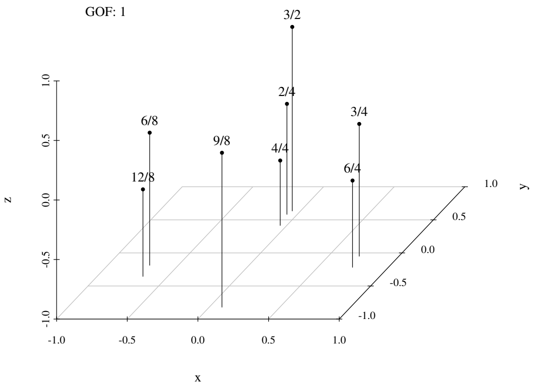

a: MDS of the distances between metres corresponding to common simple and compound time signatures. All metres are at tactus = 600 (100 bpm), and include one level of tactus subdivision.

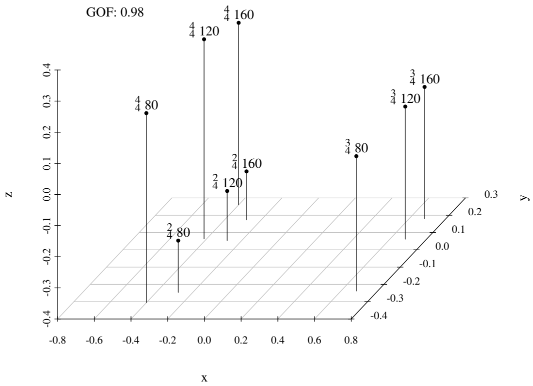

b: MDS of the distances between three simple metres, each with one level of subdivision, across the tempo range 80–160 bpm.

Fig. 4: Multidimensional scaling of distance matrices produced by calculating the pairwise Euclidean distance between metres represented within the conceptual space of METRE. The value of ε was fixed at 0.2. Goodness of fit is a measure of how well the distances between objects in the lower dimensional projection reflect the original data. Values of GOF above 0.8 are generally considered to be acceptable projections.

ultimately contributing to the body of evidence for or against cognitive semantics, in the creative context.

Acknowledgements

We are grateful to our colleagues in the Intelligent Sound and Music Systems group at Goldsmiths, and the computational creativity community at large, for many years of richly fruitful discussion. We also thank several anonymous reviewers for constructive comments on earlier drafts. The work reported here was supported by an Arts and Humanities Research Council Doctoral Studentship awarded to the first author and an Engineering and Physical Sciences Research Council DTA Doctoral Studentship awarded to the third author.

References

Aisbett, J. and Gibbon, G. (2001). A general formulation of conceptual spaces as a meso level representation. *Artificial Intelligence*, 133(1–2):189–232.

Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.

Benjamin, W. E. (1984). A theory of musical meter. *Music Perception*, 1(4):355–413.

Boden, M. A. (1998). Creativity and artificial intelligence. *Artificial Intelligence Journal*, 103:347–356.

Boden, M. A. (2004). *The Creative Mind: Myths and Mechanisms*. Routledge, London, UK, 2nd edition.

Bown, O. and Wiggins, G. A. (2009). From maladaptation to competition to cooperation in the evolution of musical behaviour. *Musicæ Scientiæ*. Special Issue on Evolution of Music.

Brachman, R. J. and Levesque, H. J., editors (1985). *Readings in Knowledge Representation*. Morgan Kaufmann.

Caclin, A., Brattico, E., Tervaniemi, M., Näätänen, R., Morlet, D., Giard, M.-H., and McAdams, S. (2006). Separate neural processing of timbre dimensions in auditory sensory memory. *J. Cognitive Neuroscience*, 18(12):1959–1972.

Carota, F., Posada, A., Harquel, S., Delpuech, C., Bertrand, O., and Sirigu, A. (2009). Neural Dynamics of the Intention to Speak. *Cereb. Cortex*, page bhp255.

Chew, E. (2000). *Towards a Mathematical Model of Tonality*. PhD thesis, Department of Operations Research, MIT, Cambridge, Massachusetts.

Clarke, E. F. (1999). Rhythm and timing in music. In Deutsch, D., editor, *Psychology of Music*, pages 473–500.

University of California, San Diego, CA, 2 edition.

Cooper, G. W. and Meyer, L. B. (1960). *The Rhythmic Structure of Music*. University of Chicago Press, Chicago, IL.

Dowling, W. J. (1978). Scale and contour: Two components of a theory of memory for melodies. *Psychological Review*, 85(4):341–354.

Fraisse, P. (1978). Time and rhythm perception. In Carterette, E. and Friedmans, M., editors, *Handbook of Perception*, volume 2, pages 203–254. Academic Press, New York, NY.

Gabrielsson, A. (1993). The complexities of rhythm. In Tighe, T. J. and Dowling, W. J., editors, *Psychology and music: The understanding of melody and rhythm*, pages 93–120. Lawrence Erlbaum Associates, Inc., Hillside, NJ.

Gärdenfors, P. (2000). *Conceptual Spaces: the geometry of thought*. MIT Press, Cambridge, MA.

Gärdenfors, P. (2007). Representing actions and functional properties in conceptual spaces. In Ziemke, T., Zlatev, J., and Frank, R. M., editors, *Body, Language and Mind*, volume 1, pages 167–195. Mouton de Gruyter, Berlin, Germany.

Guilford, J. (1967). *The Nature of Human Intelligence*. McGraw-Hill, New York.

Haase, K. (1995). Too many ideas, just one word: a review of Margaret Boden's *The Creative Mind: Myths and Mechanisms*. *Artificial Intelligence Journal*, 79:69–82.

Hasty, C. F. (1997). *Metre as Rhythm*. Oxford University Press, Oxford.

Jain, A., Nandakumar, K., and Ross, A. (2005). Score normalization in multimodal biometric systems. *Pattern Recognition*, 38(12):2270–2285.

Koestler, A. (1964). *The Act of Creation*. Hutchinson & Co., London.

Lemström, K. and Wiggins, G. A. (2009). Formalizing invariances for content-based music retrieval. In Tzanetakis, G. and Hirata, K., editors, *Proceedings of ISMIR 2009*.

Lerdahl, F. and Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. MIT Press, Cambridge, MA.

London, J. (2004). *Hearing in time: psychological aspects of music metre*. Oxford University Press, Oxford, UK.

Lustig, R. (1995). Margaret Boden, *The Creative Mind: Myths and Mechanisms*. *Artificial Intelligence Journal*, 79:83–96.

MacKay, D. J. C. (1998). Introduction to Monte Carlo methods. In Jordan, M. I., editor, *Learning in Graphical Models*, NATO Science Series, pages 175–204. Kluwer Academic Press, Dordrecht, The Netherlands.

Müllensiefen, D. (2009). Fantastic: Feature analysis technology accessing statistics (in a corpus). Technical report, Goldsmiths, University of London, London, UK.

Parncutt, R. (1994). A perceptual model of pulse salience and metrical accent in musical rhythms. *Music Perception*, 11:409–464.

Parsons, D. (1975). *The Directory of Tunes and Musical Themes*. S. Brown, Cambridge, UK.

Patel, A. D. (2008). *Music, Language, and the Brain*. Oxford University Press, Oxford.

Pearce, M. T. and Wiggins, G. A. (2006). Expectation in melody: The influence of context and learning. *Music Perception*, 23(5):377–405.

Perkins, D. (1995). An unfair review of Margaret Boden's *The Creative Mind* from the perspective of creative systems. *Artificial Intelligence Journal*, 79:97–109.

Povel, D.-J. (1984). A theoretical framework for rhythm perception. *Psychological Research*, 45(4):315–337.

R Development Core Team (2010). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Ram, A., Wills, L., Domeshek, E., Nersessian, N., and Kolodner, J. (1995). Understanding the creative mind: a review of Margaret Boden's *Creative Mind*. *Artificial Intelligence Journal*, 79:111–128.

Raubal, M. (2004). Formalizing conceptual spaces. In Varzi, A. and Vieu, L., editors, *Formal Ontology in Information Systems, Proceedings of the Third International Conference (FOIS 2004)*, volume 114 of *Frontiers in Artificial Intelligence and Applications*, pages 153–164, Amsterdam, NL. IOS Press.

Raubal, M. (2008a). Cognitive modeling with conceptual spaces. In Ragni, M., Schultheis, H., and Barkowsky, T., editors, *Workshop on Cognitive Models of Human Spatial Reasoning, Freiburg, Germany*, pages 7–11.

Raubal, M. (2008b). Representing concepts in time. *Spatial Cognition VI. Learning, Reasoning, and Talking about Space*, pages 328–343.

Rickard, J. T. (2006). A concept geometry for conceptual spaces. *Fuzzy Optimization and Decision Making*, 5(4):311–329.

Rickard, J. T., Aisbett, J., and Gibbon, G. (2007a). Knowledge representation and reasoning in conceptual spaces. In *Foundations of Computational Intelligence, 2007. FOCI 2007. IEEE Symposium on*, pages 583–590.

Rickard, J. T., Aisbett, J., and Gibbon, G. (2007b). Reformulation of the theory of conceptual spaces. *Information Sciences*, 177(21):4539–4565.

Roberson, D., Davidoff, J., Davies, I. R. L., and Shapiro, L. R. (2006). Colour categories and category acquisition in Himba and English. In Pitchford, N. and Bingham, C., editors, *Progress in Colour Studies*, pages 159–172. John Benjamins, Amsterdam.

Schank, R. and Foster, D. (1995). The engineering of creativity: A review of Boden's *The Creative Mind*. *Artificial Intelligence Journal*, 79:129–143.

Schwering, A. and Raubal, M. (2005a). Measuring semantic similarity between geospatial conceptual regions. In *GeoSpatial Semantics - First International Conference, GeoS 2005*, volume 3799 of *LNCS*, pages 90–106. Springer-Verlag, Berlin, Germany.

Schwering, A. and Raubal, M. (2005b). Spatial relations for semantic similarity measurement. In *Perspectives in Conceptual Modeling*, volume 3770 of *LNCS*, pages 259–269. Springer-Verlag, Berlin, Germany.

Shepard, R. N. (1982). Structural representations of musical pitch. In Deutsch, D., editor, *Psychology of Music*, pages 343–390. Academic Press, New York.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237(4820):1317–1323.

Steedman, M. J. (1977). The perception of musical rhythm and metre. *Perception*, 6(5):555–569.

Thornton, C. (2007). How thinking inside the box can become thinking outside the box. In Cardoso, A. and Wiggins, G., editors, *Proceedings of the 4th International Joint Workshop on Computational Creativity*.

Thornton, C. (2009). Personal communication.

Turner, S. R. (1995). Margaret Boden, *The Creative Mind*. *Artificial Intelligence Journal*, 79:145–159.

Wallas, G. (1926). *The Art of Thought*. Harcourt Brace, New York.

Wiggins, G. A. (2006a). A preliminary framework for description, analysis and comparison of creative systems. *Journal of Knowledge Based Systems*, 19(7):449–458.

Wiggins, G. A. (2006b). Searching for computational creativity. *New Generation Computing*, 24(3):209–222.

Wiggins, G. A., Harris, M., and Smaill, A. (1989). Representing music for analysis and composition. In Balaban, M., Ebcioğlu, K., Laske, O., Lischka, C., and Soriso, L., editors, *Proceedings of the Second Workshop on AI and Music*, pages 63–71, Menlo Park, CA. AAAI.

Wiggins, G. A., Pearce, M. T., and Müllensiefenllensiefen, D. (2009). Computational modelling of music cognition and musical creativity. In Dean, R., editor, *Oxford Handbook of Computer Music and Digital Sound Culture*. Oxford University Press.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology, Monograph Supplement*, 9:1–27.