



Aesthetics and Philosophy of the Arts

Style and Supervenience

*Igor Douven
Utrecht University*

ABSTRACT: Cope's *Computers and Musical Style* (1991) describes a computer program that allegedly can represent and replicate musical styles solely on the basis of compositions that have been entered into it. If this claim is correct, then it must be that an oeuvre's stylistic characteristics locally supervene on its textual features, which roughly means that its stylistic properties are entirely determined by its textual properties. In my paper I argue that stylistic properties do not locally supervene on textual properties, and thus that neither Cope's program nor any other that essentially works like it can represent or replicate styles.

Cope (1991) describes a computer program that allegedly can represent and replicate musical styles solely on the basis of compositions that have been entered into it (ix, xi—all page-references are to Cope 1991). If this claim is correct, then it must be that an oeuvre's stylistic characteristics locally supervene on its textual features, which roughly means that its stylistic properties are entirely determined by its textual properties. This paper argues that stylistic properties do not locally supervene on textual properties, and thus that neither Cope's program nor any other that essentially works like it can represent or replicate styles.

1. Cope's Composing Computer

David Cope is a composer and music theorist who got interested in the applications of computer science to music. The direct cause of his interest was a composer's block; this made him turn to computer programming in the hope to find a 'composing partner' (18). Eventually his search resulted in a program he termed *Experiments in Musical Intelligence* (EMI). Let me briefly sketch what EMI does and how it works.

What EMI does is easily explained: musical data, like for instance a number of Mozart piano-sonatas, are (in coded form) fed into the computer, which then outputs new musical material. This new material is then hoped to be and, according to Cope, also often *is*, in the style of the music that was entered. How the program works is not so easily explained, but the following simplified account will do for the purposes of this paper (cf. 152ff for details). The two most important components of EMI are a pattern-matcher and a so-called Augmented Transition Network. The first searches for common patterns in the works that have gone into the computer, and stores these in a 'style dictionary' (together with a weight, indicating how 'common' they are). The second generates musical phrases by

transforming, in systematic ways, following constraints from the style dictionary, basic ‘context free’ musical material (stemming from a kind of automated part-writing device, which is also part of EMI).

Cope’s book contains a number of musical examples that have been generated by EMI. For instance, there is a first movement of an imitation Mozart piano-sonata, a slow movement of an imitation Prokofiev piano-sonata, a Bach invention imitation, a Joplin rag imitation, and much more. Although none of these computer compositions is musically really satisfactory, anyone familiar with previous attempts to do more or less the same as Cope will find the overall achievement very impressive; what Cope offers here is by far the best that has been brought forward in the field. But exactly because the results are so admirable, there is a danger that we interpret them too easily as constituting a kind of proof of the view on the metaphysical status of style (or stylistic properties) as it emerges from Cope’s text, which, I believe, they do not.

The view on style we find in Cope’s book is not an uncommon one, as is also witnessed by some of the quotes given in chapter 2. It basically says that style is something ‘located in’ a composer’s (or author’s or painter’s) oeuvre (or in the art of some period), something that ‘inheres’ in the texts of these works. Let us call any view on style that agrees to this a *localistic* view. It need not be part of such a view that stylistic features are *reducible* to textual features, i.e., it is not implied that the former can be 1-1 correlated with the latter. Localists need merely hold that an oeuvre’s stylistic features *locally supervene* (to be explained below) on its textual properties. (Strictly speaking, localists need not even hold this much; they may believe that stylistic features inhere ‘magically’ in works of art, and are not in any comprehensible way related to the textual properties of the works they are stylistic features of—for obvious reasons I will ignore this option.) However, Cope does hold the stronger reducibility thesis: for him stylistic features are to be equated with statistically significant features, with recurring patterns (see e.g. 30).

As I just said, Cope’s view on style is not uncommon. What makes his book nevertheless interesting even from a philosophical perspective, is that the computer project described in it seems to *empirically establish* that view. (Another part of philosophy naturalized?!) For if EMI indeed does what Cope claims it does, *viz.*, representing and replicating musical styles on the basis of nothing but some coded musical material, then localistic views on style must be right: it must be that stylistic properties are entirely determined by textual properties, for what else had EMI to go on?

In the sequel I will argue that localistic views on style are misguided and that Cope’s program does not show the contrary. More particularly it will be argued that already the weaker localistic views are false: stylistic properties do not locally supervene on textual properties, let alone that the former are reducible to the latter. To facilitate discussion of these matters, however, I now first briefly explain the concept of supervenience.

2. Locally vs. Non-Locally Supervenient Properties

Supervenience can either be described as a relation between higher- and lower-level properties (for instance psychological and physical properties) or, equivalently, as a relationship between higher- and lower-level conceptual systems or theories, i.e., as a relation between the theories that attempt to describe these properties. I here opt for the first alternative, but it is straightforward to rephrase the following in terms of relations between theories.

Imagine a device—a huge device—for transforming our world (readers familiar with the Calvin and Hobbes comic strips may think here of a gigantic and somewhat more sophisticated version of Calvin’s transmogrifier). The device can change any part of the world in all respects, as well as add things to the world and remove things from it. However, we can also determine that certain parts of the world will in the process of transformation remain unchanged in certain respects. Now suppose we give our magical device the command to transform the world with the exception of part *p*; that should at least in all *R*-respects remain unchanged. Then a higher-level, non-*R*, property *P* is said to supervene locally on *p*’s *R*-properties if, no matter how much the rest of the world is transformed, *p* is guaranteed to possess *P* also after the transformation-process.

Other notions of supervenience can now be defined in terms of what apart from *p* has to remain unchanged in order to secure that some higher-level property is retained in the process. For instance, we may have to determine that not just *p* remains unchanged in all *R*-respects, but that also a certain segment *s* of the environment in which *p* is embedded is kept fixed in those respects so as to ensure that *p* still has the higher-level property after the process; we might in that case say that the particular higher-level property supervenes regionally: it supervenes on the *R*-properties of *p* together with those of the region *s*. It may even be the case that *p* is guaranteed to have *P* after the transformation only if the transformation leaves everything as it is in all *R*-respects; in that case *P* might be called a globally supervenient property: it supervenes upon the *R*-properties of the world as a whole. Of course one can go on to define ever more fine-grained notions of supervenience in this way. However, for our purposes it suffices to distinguish between local and non-local supervenience.

3. Do Stylistic Properties Locally Supervene on Textual Properties?

According to localistic views on style the answer to this question is evidently *yes*. For the basic idea underlying these views can, bearing the foregoing in mind, be reformulated as: by keeping the textual properties fixed we automatically keep the stylistic properties fixed, no matter how much everything else is transmogrified. In this section I will argue that this idea is incorrect and that the answer to the above question should be negative.

In his (1975: 269) Laske complains that ‘the notion of ‘style’ ¼ is one of the most ill-defined’. It seems to me that more than twenty years later this is still true. To have at least something to go on in the ensuing discussion, but without any pretension of giving a satisfactory definition nor of being original, I suggest that it is part of the meaning of ‘style’ that stylistic features are those features of a work or oeuvre that somehow help us to place it, as being made by a certain artist, or as being created in a certain period; a work’s stylistic features are those features that are typical or characteristic for its author and that thus help us identify that author, or that in any case give us some indication as to when the work was created. It’s these properties that we want to state when we attempt to describe the style of an artist (or of a school or period). Notice that this doesn’t pre-empt any questions concerning the metaphysical status of stylistic properties; the features characteristic of an oeuvre may, for all that is said so far, still locally supervene on, or even simply be, textual properties of the works constituting that oeuvre.

The first doubts about localism arise when we reflect upon what it means to characterize or typify something. What is typical of your face, does not just depend on your own but also on other people’s faces. Likewise, what makes Mozart’s piano-sonatas ‘typically Mozart’ does not only depend on these sonatas but also on works by other composers. It is not hard

to come up with some thought-experiments which show how this affects the question we seek to answer in this section. The recipe for such experiments is simple: Either take some feature *F* of a composer's *C* work that you consider to be a stylistic characteristic of that work, transmogrify the world in such a way that in the resulting world there are many different oeuvres that all exhibit *F*, and ask yourself whether you would still call *F* a stylistic property of *C*'s work in the resulting world. Or take a feature *G* of *C*'s work that you do not at all regard as a stylistic characteristic of that work, transmogrify the world so that after the process none of the other oeuvres that also exhibit *G* is left, and ask yourself whether you still would not regard *G* as a stylistic feature of *C*'s work in that world.

Here is an example of a thought-experiment of the second kind (for an example of the first kind see towards the end of the paper): We program our transmogrifier so that it keeps the textual properties of Mozart's works and of those of the 20th century German composer Boris Blacher fixed, but 'eliminates' all other music. Blacher wrote music in what he called a *variabile Metrik*; in his works the time changes in every new bar according to some system that need not concern us here. Note, now, that in the resulting world the fact that Mozart does not change time every new bar—a feature of his music that we, in our world, wouldn't dream of calling typical of him—has become a stylistic feature of his work. Someone who is told nothing about Mozart's music but that it has *that* feature will already be able to unfailingly identify Mozart's works (i.e., will be able to unfailingly distinguish Mozart's from Blacher's works).

The conclusion I want you to draw from thought-experiments like this is, of course, that stylistic properties do not locally supervene on textual properties. But you may object that these at most show that in different worlds different features may be *highlighted* as being characteristic of an oeuvre, and that I'm over-emphasizing the pragmatic element in the notion of style. Surely, you may say, a predicate may express *a* characteristic of an oeuvre—in the sense that the predicate is true of all or most of the works that oeuvre consists of—without expressing something characteristic of that oeuvre (something which, given the way the world is, distinguishes it from other oeuvres). If, then, we identify style with the totality of characteristics not in the second but in the first sense (of which one sub-collection may be highlighted in one world, another sub-collection in another world, etc.), it seems Cope can still maintain that, by capturing the statistically significant features, his program succeeds to capture those characteristics, and thus succeeds to capture musical styles. I strongly doubt the more liberal conception of style hinted at here. However, even if it should be correct, the thesis of this paper still stands. For, as the following considerations show, there are cases where features are undeniably part of an artist's style (i.e., are characteristic in both the first and second sense), but that are neither in any relevant sense statistically significant textual features, nor locally supervenient on such features, nor even locally supervenient on textual features in general.

Consider EMI's Mozart sonata. Cope himself is clearly particularly proud of it. And it must be admitted that this piece exhibits remarkably many of the features that are quite commonly regarded as 'typically Mozart': the scale-like figures, the *Seufzer*, the Alberti-basses, the broken chords in the right-hand part; in addition, the harmonic progressions are, overall, quite 'Mozartian'. There are some all too shrewd modulations perhaps, but of course it might be just a matter of more research and more programming to overcome such shortcomings. But although these features indeed often suffice to pick out Mozart's music, this is in some sense just a coincidence, largely due to the fact that virtually no piano-composition of, say, Hummel or Kozeluch is played anymore nowadays. I'm sure that, had these works not gone into oblivion, we wouldn't find Cope's result so 'typically Mozart',

since all the features that can be found in Cope's Mozart sonata can also be found in the works of these composers. (One might say that, for all but the expert in 18th century piano-music, our world is epistemically indistinguishable from a world transmogrified so that in it Clementi, Hummel, Dussek, Kozeluch, and many other composers of the same era have never existed.) On the other hand, what is typical for Mozart's music in contrast to the works of these composers is something that is not found in the EMI sonata and, in fact, is something a program like EMI cannot get a grip on.

For what's typical is that, where these other composers almost exceptionally 'scholastically' follow the rules, stay within the schematic, Mozart often surprises us, moves us, makes us laugh, by breaking the rules, by breaking out of the schemata. It is these capricious, unpredictable elements in Mozart's music that distinguish it from the works of many of his contemporaries. Now the problem these surprising events pose for programs that like EMI attempt to analyze style via some pattern-matching procedure—and for the reductionist view on style that underlies EMI—is this: the only pattern that can be discovered in these events is 'that the rules are broken'. There is nothing underlying them that could be expressed in terms of some advanced music theory, at least nothing informative. That is to say, there is no harmonic transition, or melodic turn, or anything else of that sort common to them. So these events—which together *do* constitute a stylistic feature of Mozart's music—escape analysis by a pattern-matcher. For although the pattern 'that the rules are broken' may be detectable, it is so vague and general that it cannot be said to adequately represent the aspect of Mozart's style at issue (it's *how* and *when* these rules are broken that is so typical for Mozart), unless it can be explicated in further steps. But this is what cannot be done—because of the very nature of these events ('the pattern behind the unpredictable' surely sounds like a contradiction in terms). Thus stylistic properties cannot simply be equated with statistically significant properties.

It is also easily seen that stylistic features of the kind just mentioned do not locally supervene on statistically significant textual properties. This time program the transmogrifier so that it keeps the statistically significant features, i.e., the patterns, of Mozart's music fixed, including the pattern that the rules are broken every so often, but transforms all else. If Cope's conception of style were correct, the resulting music would again have to be in Mozart's style. However, although we made sure that in it there will also be non-schematic elements, nothing guarantees that not some of these now are ridiculous or even sheer musical nonsense instead of witty or exciting or moving. To put the point concisely: Thanks to Cope we now know how to get the patterns into the computer, but we still have no idea as to how to get the genius in. And hence there is no guarantee that the result will be 'typically Mozart'.

Finally, these non-schematic elements in Mozart's music can also serve to underpin the stronger thesis we want to establish in this section, *viz.*, that stylistic features do not even locally supervene on textual properties in general (i.e., on statistically significant and insignificant ones taken together). Take for instance the final part of Mozart's symphony no. 31, composed and first performed in Paris, about which Mozart writes in a letter to his father: 'Since I heard that here [i.e. in Paris] all final Allegros $\frac{1}{4}$ begin with all instruments together and mostly *unisono*, I began with only the second violins *piano*' (letter of July 3, 1778; my translation). As one can also read in that letter, this 'inversion' had an enormous effect on the Parisian public. Now we can readily program our transmogrifier in such a way that this surprise, as well as any of the other surprising events in Mozart's music, is no longer striking, exciting, witty, ingenious, but becomes dull; simply adding many symphonies to the world with final parts starting not with an orchestral *tutti* but in the same

fashion as the *Pariser* will take (part of) the wit from this work, will make it (more) conventional; we can do similarly for any of the other surprises of course.

In sum: I have argued that stylistic properties do not locally supervene on textual properties. If my arguments were correct, then there is no hope that we will ever have a computer program that in some form or other is able to give a stylistic analysis of an oeuvre or a part thereof just on the basis of the texts of these works. So then in particular EMI cannot be such a device.

References

Cope, D. (1991) *Computers and Musical Style*, Oxford: Oxford University Press.

Laske, O. (1975) 'On psychomusicology', *International Review of the Aesthetics and Sociology of Music*, 269-281.