Fuzzy Logic and Elemental Saturation in Musical Pitch-Class Set Theory

Michael Buchler
University of Iowa
Iowa City, Iowa USA

Abstract

Over the past twenty-five years, the rhetoric of pitch-class set theory has been pervasive in analyses of atonal—and particularly modernist—musical compositions. The tools associated with this form of analysis are primarily generalized labeling systems and operations that map one realization of a set type onto another (see Forte, 1973; Lewin, 1977 and 1987; and Morris, 1987). Relatively few associational strategies for comparing different set types (e.g., similarity measures) move beyond the bivalent correspondences that are most commonly drawn (see Morris, 1979-80; Lewin, 1979-80; Rahn, 1979-80 and 1989; Isaacson, 1990; Castrén, 1994; Hermann, 1994; and Buchler, 1997).

This paper describes a generalized fuzzy version of Lewin's embedding function that examines the degree to which any set X is saturated with some subset Y. This function forms the theoretical basis for a number of new multivalent tools for comparing two musical set types.

Keywords

Atonal music; music theory; fuzzy logic; pitch-class set; saturation; similarity relations.

In music theory, we implicitly (and perhaps unwittingly) invoke fuzzy categories quite frequently. Consider statements such as "those two passages sound very similar" and "this movement is in sonata form." The terms "very similar" and "sonata form" refer to fuzzy sets that many people have vainly attempted to unfuzzify (or make "crisp"). Admitting that such vagueness exists can help us identify the types of common features that lead people to categorize things in a particular way. Which features does a movement really need (and which features could it lack) in order to fall into the fuzzy set of "sonata forms?" (Incidentally, I think that the fuzzy approach has great pedagogical advantages, and I frequently utilize it when addressing tonal forms with my students.)

Traditionally, however, the tools of atonal analysis have been crisp. Pitches or pitch classes are either members or non-members of a particular set; two sets either enjoy a given relationship (e.g., transformational or inclusional) or they do not. One of the benefits of a crisp system is that it facilitates the application of specific labels, although such labels might not be particularly meaningful without further interpretation. Indeed, it could be argued that the early atonal theorists were essentially nominalists and, to a lesser extent, logical positivists. The problems with such bivalent systems of similarity resemble the problems found in other crisp systems: they can bind the analyst to a rigid notion of what is and is not related. "Nearly related" is simply not a viable option.

The application of such a crisp system can distort the gray areas that exist in nearly all music-analytical decisions. Much of what will be advocated here is essentially a fuzzier version of some well-known pre-existing concept. To that end, let us begin with the very basic idea of inclusion. Given two pitch-class sets {C, C♯, D} and {C, C♯, D, E♭, E}, we can say that the former is a subset of the latter—or, to use Allen Forte's term, there is a K relation (Forte, 1973).

Let us now consider the trichord {C, D, E}, which is also a subset of {C, C♯, D, E♭, E}. Both {C, C♯, D} and {C, D, E} are subsets of {C, C♯, D, E♭, E}, but most people would agree that the *chromatic* trichord is more similar to the chromatic pentachord than is the *whole-tone* trichord. Nonetheless, both pairs share exactly the same Fortean K relation. A K relation is a bivalent thing: either two sets exhibit the relationship or they don't. There is no such thing as a strong K relation or a weak K relation.

David Lewin's EMB(*/X/*, *Y*) function (short for "embedding") provides a bit more information. Here, */X/* denotes set class *X* and *Y* denotes a set of pitch classes. The embedding function indicates how many members of the set class are subsets of the larger set (Lewin, 1977). In the case of EMB([0,1,2], {C, C♯, D, E♭, E}), there are three distinct forms of set class [0,1,2] that are subsets of {C, C♯, D, E♭, E}: {C, C♯, D}, {C♯, D, E♭}, and {D, E♭, E}. The value returned is therefore 3. On the other hand, there is only one form of [0,2,4] (the whole-tone trichord) that is embedded in {C, C♯, D, E♭, E}, so EMB in this case returns a value of 1. If we assume that larger embedding values indicate greater similarity between two sets, then the embedding function reflects our intuition that the chromatic trichord more closely resembles the chromatic pentachord than does the whole-tone trichord.

Indeed, many similarity measures are based upon this model, and this is appropriate if one believes that the more embedded forms of *X* there are in *Y*, the more similar *X* is to *Y*. However, such numbers can be deceiving. For example, there are three distinct forms of [0,1,2] in {C, C♯, D, E♭, E}, as shown above, but there are twice that number in the ten-note set {C, C♯, D, E♭, E, F♯, G, G♯, A, B♭}. Using the logic employed by many subset-based similarity measures, the chromatic trichord apparently has twice the affinity to this decachord as it does to the chromatic pentachord, despite the fact that the trichord could be seen as the progenitor of the pentachord but not of the decachord. (In other words, overlapping transpositions of [0,1,2] will eventually form the [0,1,2,3,4] pentachord, but not the [0,1,2,3,4,6,7,8,9,10] decachord.)

The remedy that I propose for this type of problem is what I call a measure of saturation. To figure out the degree to which {C, C♯, D, E♭, E} is saturated with [0,1,2], we take the embedding number (3) and divide it by the maximum number of [0,1,2]s that could be found embedded in *any* five-note set. In other words, we figure out which five-note set is most packed with [0,1,2]s and see exactly how many [0,1,2]s it embeds. In this case, the answer happens to be [0,1,2,3,4], which (as we've seen) contains three [0,1,2]s. Dividing the actual embedding number by the maximal embedding number, we obtain the degree of saturation. So, a saturation version of Lewin's embedding function goes beyond telling us that there are three [0,1,2]s in {C, C♯, D, E♭, E}, instead informing us that {C, C♯, D, E♭, E} contains as many [0,1,2]s as possible, given its cardinality. By contrast, a ten-note set could embed as many as eight forms of [0,1,2]. The ten-note set given earlier contains only six forms, and is therefore comparatively less saturated with [0,1,2]s than is the chromatic pentachord.

Actually, the degree of saturation is almost that simple, but not quite. In order for this to be a valid measure, we should also consider the minimum possible number of [0,1,2]s contained in any pentachord and any decachord. In the case of the pentachord, the minimum number is zero (consider, for example, the whole-tone pentachord). It is impossible, however, to construct a ten-note set that doesn't contain at least one three-note chromatic segment—in fact, it is impossible to include fewer than six distinct [0,1,2]s in any decachord. So, if the minimum number of [0,1,2]s in a ten-note set is six and the maximum number is eight, then rather than focusing on how our decachord embeds twice as many [0,1,2]s as does our pentachord, we can instead say that the decachord embeds the minimum possible [0,1,2]s and the pentachord embeds the maximum possible [0,1,2]s, given their respective cardinalities.

Like the fuzzy degree of membership, the degree of saturation is expressed as a number between 0 and 1, and is essentially a percentage. 0 represents the minimum and 1 represents the maximum. A .5 degree of saturation falls halfway between the two extremes, and is equivalent to the point of fuzzy entropy. My saturation version of Lewin's embedding function is defined formally as follows:

$/X/$ denotes set class X; Y denotes a set of pitch classes; $\#Y$ denotes the cardinality of set Y.
min$(/X/, \#Y)$ = the minimum number of $/X/$ that can be embedded in *any* set of cardinality Y.
max$(/X/, \#Y)$ = the maximum number of $/X/$ that can be embedded in *any* set of cardinality Y.

$$\mathrm{SATEMB}(/X/, Y) = \frac{\mathrm{EMB}(/X/, Y) - \min(/X/, Y)}{\max(/X/, Y) - \min(/X/, Y)}$$

Saturation is fuzzy in spirit because it produces a scaled, rather than categorical, response. Using these new means, we can "fuzzify" most (if not all) analytical tools. To illustrate this, let us set aside Lewin's EMB function and take up Forte's interval-class vector.

The hexachord class [0,1,2,3,4,6] is nearly chromatic: five elements are adjacent in a chromatic scale and one is set apart by a whole tone. The interval-class (ic) vector of [0,1,2,3,4,6] is <4,4,3,2,1,1> (that is, there are four instances of ic1, four instances of ic2, three instances of ic3, and so on). Obviously, [0,1,2,3,4,6] does not embed the maximum number of ic1s possible for a hexachord; after all, this set class was described only as *nearly* chromatic. The entirely chromatic hexachord class ([0,1,2,3,4,5]) has five ic1s, not four. The minimal number of ic1 that is possible in a hexachord is zero, as found in the whole-tone collection ([0,2,4,6,8,10]). So, rather than simply showing that there are four ic1s, our fuzzified vector indicates that [0,1,2,3,4,6] contains 80% of the maximal ic1 content for any hexachord.

This kind of a relative value enables us to draw connections between sets of different sizes more readily because it always describes the interval content in terms of what is possible. The entire saturation vector of [0,1,2,3,4,6] is <0.80,0.67,0.60,0.00,0.20,0.33>. Notice that there are just as many ic1s as ic2s, yet the fuzzy values in the first and second columns are different. This is because a hexachord can contain as many as six ic2s (again, think of the whole-tone scale), whereas it can contain no more than five ic1s. Notice also that the degree of saturation for ic4 is 0—despite the fact that there are two ic4s in the set. This is no mistake: it is impossible to construct a hexachord with fewer than two ic4s, and the entry in the new vector clarifies this minimal saturation.

Set Class	IC Vector	Interval-Class Saturation Vector
[0,1,2,4,5,7,9]	< 3,4,4,4,5,1 >	< 0.25, 0.50, 0.50, 0.33, 0.75, 0.00 >
[0,2,3,5,7,9]	< 1,4,3,2,4,1 >	< 0.20, 0.67, 0.60, 0.00, 0.80, 0.33 >

Possible pitch realization of [0,1,2,4,5,7,9]:

Possible pitch realization of [0,2,3,5,7,9]:

Figure 1. Two "almost diatonic" set classes

Figure 1 shows two "almost diatonic" set classes of different cardinalities. The first set class, [0,1,2,4,5,7,9], *contains* a six-note subset of the diatonic collection; the second set class, [0,2,3,5,7,9], *is* a subset of the diatonic collection, but not a subset of the first set class. Examining their ic vectors alone, one might notice very little similarity between the two except that they have the same amount of ic2 and ic6. This fact is deceptive, however, since the set classes differ in cardinality. Use of the saturation vector illustrates that they have nearly the same amounts of four of the six interval classes, and the values in the ic4 and ic6 places are exchanged. This is a fuzzy adaptation of Forte's R_1 relation (Forte, 1973), and one which allows set classes of different cardinality to be compared meaningfully. (Forte's R relations can only be used for set classes of identical cardinality when using the standard ic vector as data.) Moving beyond the R relations, we could create a fuzzy measure of exactly *how* similar these two saturation vectors are simply by averaging the differences in each of their respective arguments. For example, they are 5% different in ic1 content, 17% different in ic2 content, 10% different in ic3 content, and so forth. The average difference is 17%. By this particular measure, then, these two set classes are only 17% different.

One of the greatest problems with some analytical methodologies is an inadequate contextualization of their terms. Consider, for example, the assertion "Sets *X* and *Y* are related by inclusion." We might ask how many other sets of the same sizes as *X* and *Y* are related by inclusion? How many are *not* related by inclusion? How *many* *X*s are embedded in *Y*? How many *X*s *can be* embedded in any set the same size as *Y*? How many *X*s *must* be embedded in a set the same size as *Y*?

Furthermore, no matter how precise a number may be, its practical value depends on our ability to interpret it meaningfully. When I once mentioned to a Ugandan friend that a particular baseball player batted .333 last season, he was initially unimpressed: why should a professional baseball player hit the ball successfully only a third of the time? Once I explained that the American League batting champion hit only .347 and that the worst batter in the league hit around .175, he understood that, relatively speaking, .333 is a pretty good average. Similarly, my earlier assertion that the two set classes in figure 1 are 17% different becomes more meaningful with an understanding of the larger context. When comparing six- and seven-note sets using the simple similarity measure that I outlined, the average difference is 25.9% and it is possible to find two sets that are as much as 62.5% different. A 17% difference is therefore relatively small.

Considering elemental saturation focuses our attention, helps us to convey observations more clearly and efficiently, and can diminish our reliance on jargon. For instance, we can say that set *Y* is 90% saturated with set *X*, or even that set *Y* is *almost maximally saturated* with set *X*. Given a list of saturation values, we can use fuzzy logic to decide the point at which "almost maximally saturated" becomes "fairly saturated," then "moderately saturated," then "not particularly saturated," and so forth. Our ideas of where these points lie can be as fuzzy as more commonly discussed fuzzy categories such as "hot," "warm," "mild," and "cold" (Kosko, 1993). These dividing points can even shift as our intuitions change or a musical context is altered.

In addition to allowing theorists to express essentially subjective ideas more clearly and precisely, fuzzy relations can also open doors that previously seemed locked. Much (if not most) music theory exists in the domain of what Robert Morris calls pitch-class space—a conceptual dimension where register does not matter (Morris, 1987). Up to this point, I have only discussed relationships among pitch classes and pitch-class sets, avoiding comments about register, contour, and orchestration. Pitch-class space, with only 12 circularly-related elements, is a conveniently finite realm, and formally considering parameters other than pitch class introduces a multitude of problems. For example, once we accept register along with pitch, we have a virtually infinite number of sets and very few ways to relate them.

Robert Morris (1995) has developed several pitch-space tools that consider the number of common tones between two sets as well as the number of common intervals. Like most of Forte's initial constructs for pitch-class set analysis, Morris's tools are only useful for sets of the same cardinality—an obvious drawback when dealing with music of any complexity. Using saturation, however, it is also possible to relate pitch sets of different sizes. Like the saturation-based pitch-class relations, these measures simply take into account the range of possibility for each set of any given cardinality. With pitch sets, though, it is first necessary to consider the distance between the outer voices in order to factor in the potential sizes and quantities of intervals available. We can then analyze pitch sets on their own terms and compare them to other pitch sets that have been similarity analyzed. While space prohibits a detailed explanation of the pitch-space algorithm, it is sufficient to say that the saturation principles remain precisely the same as those described earlier.

These excursions into both pitch and pitch-class space illustrate just a few of the ways that we can apply fuzzy principles to fundamental analytical problems. The bivalent (or "crisp") tools created by Forte and others in the sixties and seventies provided a solid foundation for the young field of atonal analysis. Fuzzy modifications to these tools enable us to maintain the rigorous approach that their creators brought to the field while allowing us to express a much wider range of musical relationships.

Buchler, M. (1997); Relative Saturation of Subsets and Interval Cycles as a Means for Determining Set-Class Similarity; Ph.D. diss., University of Rochester.

Castrén, M. (1994); RECREL: A Similarity Measure for Set-Classes; Studia Musica 4, Sibelius Academy

Forte, A. (1973); The Structure of Atonal Music; Yale

Hermann, R. (1993); A General Measurement for Similarity Relations: a Heuristic for Constructing or Evaluating Aspects of Possible Musical Grammars; Ph.D. diss., University of Rochester.

Isaacson, E. (1990); Similarity of Interval-Class Content between Pitch-Class Sets: The IcVSIM Relation; Journal of Music Theory Vol. 34, No. 1 (pp. 1-28)

Klir, G. J. and B. Yuan (1995); Fuzzy Sets and Fuzzy Logic: Theory and Applications; Prentice Hall

Kosko, B. (1993); Fuzzy Thinking: The New Science of Fuzzy Logic; Hyperion

Lewin, D. (1977); Forte's Interval Vector, My Interval Function, and Regener's Common-Note Function; Journal of Music Theory Vol. 21, No. 2 (pp. 194-237)

Lewin, D. (1987); Generalized Musical Intervals and Transformations; Yale

Morris, R. D. (1979-80); A Similarity Index for Pitch-Class Sets; Perspectives of New Music, Vol. 18 (pp. 445-460)

Morris, R. D. (1987); Composition with Pitch Classes: A Theory of Compositional Design; Yale

Morris, R. D. (1995); Equivalence and Similarity in Pitch and their Interaction with Pcset Theory; Journal of Music Theory, Vol. 39., No. 2 (pp. 207-243)

Rahn, J. (1979-80); Relating Sets; Perspectives of New Music Vol. 18 (pp. 483-498)

Rahn, J. (1989); Toward a Theory for Chord Progression; In Theory Only Vol. 11, No. 1 (pp. 1-10)