University of Huddersfield Repository

Jan, Steven

Meme Hunting with the Humdrum Toolkit: Principles, Problems, and Prospects

Original Citation

Jan, Steven (2004) Meme Hunting with the Humdrum Toolkit: Principles, Problems, and Prospects. Computer Music Journal, 28 (4). pp. 68-84. ISSN 0148-9267

This version is available at http://eprints.hud.ac.uk/id/eprint/1485/

Steven Jan
University of Huddersfield
School of Music and Humanities
Queensgate, Huddersfield
HD1 3DH United Kingdom
s.b.jan@hud.ac.uk

Meme Hunting with the Humdrum Toolkit: Principles, Problems, and Prospects

Introduction: Theorizing a Memetics of Music

> [A *meme* is] a unit of cultural transmission, or a unit of *imitation*. . . . Examples of memes are tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches. Just as genes propagate themselves in the gene pool by leaping from body to body via sperms or eggs, so memes propagate themselves in the meme pool by leaping from brain to brain via a process . . . [that], in the broad sense, can be called imitation. (Dawkins 1989, 2nd ed., p. 192)

So wrote Dawkins over a quarter of a century ago, drawing together various strands of nature–culture analogizing from the previous two centuries and more and recasting them in the crucible of his powerful *selfish gene* hypothesis—the notion that, ultimately, the driving force of evolution is the single gene, whose *phenotypic* effects (i.e., those on the organism's morphology and behavior) influence the reproductive prospects of that gene in ways that justify Dawkins's metaphor of apparent selfish intentionality.

In its use of music as the first example of a meme substrate, Dawkins's definition is a provocative invitation to the development of a *memetics* of music—a subdiscipline of musicology that would attempt systematically to apply the insights of *Universal Darwinism* (Plotkin 1995) to the medium of music to trace pattern transmission and evolution over time. Such an application has strong intuitive attraction: after all, music appears to support discrete, "digital" patterns within the fluid, "analog" continuity of the sound stream—on a smooth continuum from pointed instances of intra- and inter-composer quotation to the myriad stan-

dardized clichés and gestures that are the "connective tissue" of a musical style. In what is essentially a memetic study (he terms it "referential analysis"), Cope offers a formalization of this continuum moving from "quotations" and "paraphrases" through "likenesses" to "frameworks" and "commonalities" (2003, p. 11), and he describes his software, Sorcerer, that can detect such inter-composer references.

Nearer to the former than to the latter end of this continuum, an example of cadence-pattern replication such as that shown in Figure 1 suggests that what was probably imitated from J. C. Bach (a common figure in his style) was a closed and cognitively salient unit for Mozart. It is, Dawkins would argue, a selfish meme that hijacked Mozart's neuronal mechanisms in the service of its own replication.

It appears that, in brief, a memetics of music would need to address the following three broad conceptual issues (for fuller treatments, see Jan 2000a, 2000b, 2002, and 2003):

1. The ontological basis of the musical meme, perhaps in terms of an analogy with the *genotype–phenotype* distinction in biology, although its application to memetics is still controversial (Blackmore 1999, p. 63).
2. The nature of the musical meme, or how the continuum of musical elements is segmented into discrete particles that can be related, by presumed replication, to equivalent particles in other contexts. This issue also encompasses the replication of musical memes at different hierarchic locations.
3. The evolutionary dynamic of musical memes, accounting for the continuous change of musical style over time as a consequence of the differential transmission and survival of *mutant* memes.

One problem with exploring the second and third points above is the sheer volume of music that

Figure 1. Replicated Patterns in Works of J. C. Bach and Mozart. (a) J. C. Bach, Keyboard Concerto in E-flat Major, Op. 7, No. 5 (C 59; 1770), I, mm. 52–55; (b) Mozart, Die Entführung aus dem Serail, *KV 384 (1782), No. 11, "Martern aller Arten," mm. 13–15.*

must be investigated for statistically significant conclusions to be drawn from it—although suggestive conclusions may often be drawn from non-empirical studies. Clearly, computer-aided study is a viable method of making the empirical investigation of such a corpus of music manageable. For assistance with the development of a theory of memetics and the investigation of memes in music, one software package in particular seems the most suitable currently available—the virtues of Cope's Sorcerer program notwithstanding—namely the Humdrum Toolkit, conceived and developed since 1989 by David Huron (1997, 2002).

This article investigates some uses of the Humdrum Toolkit in shaping and empirically testing aspects of a theory of musical memetics. It begins

with a discussion of the psychological processes that affect pattern perception—and therefore what patterns might constitute a musical meme. After a brief overview of the Humdrum Toolkit, two case studies, using contrasting but related methodologies, are investigated to determine to what extent certain patterns are propagated in the music of the late 18th and early 19th centuries and how Humdrum might be used to detect them. Finally, the prospects for computer-aided memetic analysis using Humdrum are assessed.

Pattern Replication in Music: A Gestalt/Implication-Realization Perspective

In formulating a memetics of music, our starting point must be to define what we understand by Dawkins's "unit of cultural transmission . . . of imitation" (point 2 in the list in the previous section), because memetics is fundamentally a discipline that studies particulate entities in a variety of substrates and their movement between brains via external media, such as scores, sound waves, and recordings. Such a definition is clearly fundamental for implementing investigations of the memetic paradigm in computerized searches. To this end, and in a further subdivision of point 2, one might propose three hypotheses:

(1) In a stream of musical information, memes acquire definition from their surrounding information as a result of partly learned but principally innate attributes of human perceptual and cognitive architecture—Meyer's level of *laws* (1989, pp. 13–14). If incoming information is subject to this filter, then only certain privileged configurations—those "which . . . *approach* the ideal of indivisible particulateness" (Dawkins 1989, p. 33; emphasis his)—will be able to pass through and be consciously attended to or stored in memory, both essential preconditions for memetic transmission and evolution.

(2) As multiparametric "molecules" made up of a number of "atoms" of cultural information—individual pitches, rhythms, and other uniparametric entities—musical memes appear to exist at several structural-hierarchic levels. The Schenkerian

model offers a useful perspective, terminology, and graphic symbology to describe them, but, in an inversion of the Schenkerian orthodoxy, it appears that memes at higher hierarchic levels are generated and expressed by patterns at lower levels. Other theoretical models have been formulated to explain this relationship between shallow-middleground-level structure and surface-level patterning, including Narmour's distinction between *style structures* and *style forms/shapes* (1977, pp. 173–174; 1990, p. 30ff), and the broadly comparable *schema–feature* dualism invoked by Gjerdingen (1988, pp. 45–46). Neither of these models appears incompatible with a memetic perspective.

(3) The *cultural salience* of a meme—the attribute that affects its propensity to imitation, its *fecundity*, as Dawkins calls it (1989, p. 17)—and its *differential fitness*—its fecundity as compared with that of its memetic *alleles* or rivals—appear to be partly a function of its relationship to its *antecedent* (i.e., the form from which it derives). A more chromatic, rhythmically syncopated *consequent* (mutation) of a given antecedent may, up to a certain degree of variation, be more likely to be imitated than the antecedent itself—although simpler versions may sometimes be more memorable. The element of variation is generally to be found in the surface elements of the meme, the underlying structure assuring the sameness of the two patterns. Thus, as a general principle, a consequent mutant meme is essentially the same meme as its antecedent if the structural level above that which is modified remains unaltered.

Although a complete computer-aided empirical testing of the above hypotheses—and others developed from the list in the introduction—is beyond the scope of this article, a preliminary exploration of them might use two distinct but related methodologies, illustrated later in the two case studies using the Humdrum Toolkit.

The first methodology identifies a natural candidate pattern and then evaluates its prevalence in a given *idiom* or *dialect* (Meyer 1989, pp. 23–24) according to the criterion of fit against the template pattern. If matches are found, then the pattern is, by definition, a meme, existing in the form of several copies. It may then be fruitful to modify the

original template pattern in the light of the original search to locate other patterns that, while structurally similar to the template, nevertheless manifest certain evolutionary deviations from it.

The second methodology synthesizes an artificial candidate pattern to predict what types of configurations are likely to attain memetic status. Then, attempts are made to locate real instances of the candidate in a given repertory. As with the first methodology, the hierarchic location and evolutionary profile of such figures may also be an element of the search strategy.

It is well known that issues pertinent to the first hypothesis above were effectively theorized by members of the Gestalt school of psychology in the 1920s and 1930s. In a reinvigoration of this tradition, their work informs two avenues of music-theoretical investigation developed during the last twenty years. Much of Lehrdahl and Jackendoff (1983) relies on the insights of Gestalt psychology for the formulation of metrical and grouping *well-formedness* and *preference rules*. Temperley's more recent model (2001) of grouping structure in music—a development of the preference-rule methodology of Lehrdahl and Jackendoff—also draws on Gestalt insights but, unlike Lehrdahl and Jackendoff, subjects them to empirical testing via computer models.

Perhaps the most comprehensive application of Gestalt principles to music analysis, however, is Narmour's implication-realization (hereafter "i-r") model (1977, 1984, 1989, 1990, 1992, 1999), which draws strongly on the groundwork of Meyer (1956, 1973, 1989). Reviewing Gestalt grouping principles and categorizing them into innate and acquired phenomena, Narmour notes that

> the separate registral and intervallic aspects of small intervals . . . are said to be implicatively governed from the bottom up by the Gestalt laws of similarity, proximity, and common direction. . . . [W]hat is important to notice about the invocation of such Gestalt laws is (1) that they have been shown to be highly resistant to learning and thus may be innate . . . (2) unlike the notoriously interpretive, holistically supersummative, top-down Gestalt laws

of "good" continuation, "good" figure, and "best" organization . . . the Gestalt laws of similarity, proximity, and common direction are measurable, formalizable, and thus open to empirical testing. (1989, pp. 46–47)

It seems sensible to begin a computerized investigation of the grouping and "closural" structure of music with the bottom-up Gestalt principles of similarity, proximity, and common direction, as implemented in the i-r model. The top-down laws of good continuation, good figure, and best organization are less consistent, being culturally conditioned and therefore variable aspects of perception and cognition. The benefit of these insights to memetic research is that we can use the i-r model to predict what forms memes are likely to take, and then we can use these criteria to help design computerized search tools to locate them in real musical contexts—the second pattern-searching methodology discussed above.

In brief, the i-r model offers a means of assessing the bottom-up, note-to-note implicative flux of a passage and identifying its points of procession and closure at various hierarchical levels. Narmour proposes a number of basic i-r functions, to which he applies a distinctive terminology and symbology. His fundamental premise is that, in melodic motion, an interval of a perfect fourth or smaller moving in a given registral direction implies a continuation of similar intervallic magnitude in the same direction (a phenomenon termed *process* and symbolized by "P"); whereas an interval of a perfect fifth or larger moving in a given registral direction implies a continuation of smaller intervallic magnitude in the opposite direction (termed *reversal*, "R"). Most importantly for this study, these implicative potentials may not be fulfilled, being terminated by *durational interference* ("d"), *harmonic interruption* ("h"), or *metric (beat) differentiation* ("b"), which serve to define the boundaries of a perceptual/cognitive grouping (Narmour 1989, pp. 45–51; 1990).

Although Narmour's theory is not universally accepted, in their evaluation of the predictive power of the i-r model (using the Humdrum Toolkit as

their analytic engine), Thompson and Stainton conclude that ''a combination of these principles can predict much of how Bach and Schubert composed (or 'realized,' in Narmour's terminology) notes following implicative intervals. . . . [C]omposers may be abiding by some basic rules, learned or innate, regarding melodic expectancy'' (1996, p. 33). If the results of this and other studies (such as Schellenberg 1996) are to be believed, then a Gestalt perspective, refined by the precepts of the i-r model, would appear to be a fundamental element in both pure and empirical memetic research.

Overview of the Humdrum Toolkit

The Humdrum Toolkit is a suite of software tools for UNIX systems. It encapsulates a central UNIX design philosophy in that, although each of the over 70 tools is fairly modest in its individual effects, great analytical sophistication can be achieved by connecting several tools in pipelines. The tools can also be incorporated into shell scripts to facilitate the processing of lengthy pipelines and expedite large processing tasks.

Beyond the tools themselves, Humdrum provides a syntax for representing music—that is, a series of formal conditions that stipulate how music may be represented in a manner comprehensible by the tools. The default representation scheme is ''**kern,'' which provides a means of encoding the fundamental elements of common-practice Western notation. Essentially, pitches are represented by their letter names (CC–BB, C–B, (middle) c–b, cc–bb, etc.), rhythm values are indicated by integers (''4'' = quarter-note, ''8.'' = dotted eighth-note, etc.), and parts/voices are organized in *spines* (tab-separated columns of data), with leftmost spines representing lower parts and rightmost representing upper parts. Each horizontal line of the **kern score, called a *data record*, represents a simultaneity in the music, being divided into as many *data tokens* as there are spines (Figure 3c later in this article shows two bars of music represented in **kern). More detailed overviews of Humdrum and the **kern encoding are given in Huron (1997, 2002).

It is worth noting that effective use of the toolkit in its ''raw'' state requires a degree of facility with UNIX that many musicologists without technical backgrounds are unable or unwilling to devote time to acquiring. Andreas Kornstädt's JRing program (2001) circumvents this difficulty by providing a graphical front-end to the Humdrum tools which should increase their accessibility, albeit with the loss of some flexibility. The investigations described below, however, use Humdrum in its original form.

Pattern Searching with Humdrum

In investigating musical memes, two of the Humdrum tools have particular significance. The broadly similar patt and pattern tools, as their names imply, are pattern-locating utilities. Both tools search the representation of the music under investigation using a template that defines the pattern sought. Templates use UNIX regular expression syntax to stipulate what the target file may or may not contain in order for a *positive* (i.e., a match or hit; see Temperley 2001, p. 74) to be registered.

The chief difference between patt and pattern is in the latter's implementation of *record-count metacharacters*, i.e., characters that specify how many records containing a sought pattern may occur in the target file for a positive to be registered—namely one or more (specified by '' + ''), zero or more (''*''), or zero or one (''?''). Patt does not support record-count metacharacters, interpreting them as literals.

A useful feature of patt is its *echo* option, which is not supported by pattern. This shows the results of the search by displaying only those segments of the **kern score containing the found pattern. An alternative option is to specify that patt *tag* the score under investigation by the insertion of an additional rightmost ''**patt'' spine, which places a user-specified keyword (such as ''meme'') at those points in the encoding where the sought pattern occurs. Pattern, by contrast, simply outputs a list of line numbers indicating where the pattern is found, and the researcher must then peruse the input to locate these occurrences.

The essence of a successful search using `patt` or `pattern` lies in designing a parsimonious and efficient template. Having provisionally formulated such a template, it is essential to test it out on passages—either synthetic examples or extracts taken from real pieces—in which the sought pattern is known to exist. This is to verify that the template will not register a *false negative*, i.e., will fail to find an instance of the pattern, despite its occurring in the music. Such a test may also indicate if the template would record a *false positive*, i.e., would register a pattern match for configurations that the musician's judgment would reject as unconvincing. Given these two potential problem classes, it is unwise to begin investigations of real repertories without such simulations conducted in controlled and circumscribed conditions.

Meme Hunting with the Humdrum Toolkit: Principles and Problems

I now describe two case studies that illustrate the application of Humdrum to some of the theoretical problems outlined above and exemplify the two methodologies for computer-aided pattern replication analysis outlined in the second section.

Case Study 1: The "Glass Harmonica" Pattern

A pattern occurs in mm. 1–2 of Mozart's Adagio in C Major for Glass Harmonica, KV 356 (617a) from 1791, that is essentially a chromatically altered interrupted cadence. We might hypothesize that figures such as this derive from evolutionarily simpler forms—those outlining a diatonic interrupted cadence and, ultimately, those outlining a perfect cadence (with or without a terminal $\hat{4}$–$\hat{3}$ appoggiatura)—such as are represented in Figure 2.

This distinctive figure seems a good starting point for a study following the first pattern-searching methodology, and it might well be a meme in the non-trivial ("strong") sense of having been propagated within the dialect of the European late 18th century. (The trivial—"weak"—sense ascribes memetic status to it on the grounds that the pattern is replicated in contemporary musical culture in the minds of present-day listeners to the Adagio.)

My segmentation of the music is based on the fact that, as well as exemplifying a clear melodic process, the figure clearly satisfies all three of Narmour's criteria for closure (i.e., termination of implicative potential). First, in Figure 3a, the left-hand dotted half-note in m. 2 creates durational interference; second, the unit is harmonically interrupted owing to the resolution of the vii^7/vi in the fourth beat of m. 1 to vi in the first beat of m. 2; and finally, the rest in the right hand in m. 2 effects metric (beat) differentiation. This segmentation also satisfies Temperley's "PSPR 1 (Gap Rule)" and "PSPR 3 (Metrical Parallelism Rule)" (2001, pp. 68–71), the latter on account of the recurrence of the unit in the same metrical position at mm. 5, 21, and 25.

Figure 3 shows three representations of these measures, the first a score with an overlaid i-r analysis, the second a voice-leading reduction showing foreground and shallow-middleground elements (to which I shall return later), and the third a `**kern` encoding.

When designing a template to search for the "Glass Harmonica" pattern, we begin with the premise that it is likely to be independent of key, and therefore the search process must be conducted in terms of scale degree and not absolute pitch. The repertoire under investigation must be converted to, and the template conceived in terms of, the Humdrum "`**deg`" representation. The encoding in Figure 3c appears as in Figure 4 when converted to `**deg`, in which all rhythmic information is removed (numbers being used instead to indicate scale degrees) and in which "+" and "−" indicate chromatically raised and lowered versions, respectively, of the diatonic pitch. Additionally, data are propagated downwards into *null* (empty) tokens to indicate the effect of the `ditto` command.

Having arrived at a `**deg` version of the pattern, a `patt` template suitable for locating the configuration shown in Figure 4, which might be stored as a file called `glass.patt`, may be extrapolated from it, as shown in Figure 5.

This template would only locate examples of the

(a)

(b)

Andante

pattern in the tonic key, unless they occurred in a passage whose departure from the tonic was so extensive or significant as to warrant a new key *tandem interpretation* (e.g., "*G:*") indicating the changed tonal center (and therefore the status of the new local tonic note as $\hat{1}$, as far as `deg` is concerned). To locate examples in the dominant key—the most probable other key in which the pattern might occur—then an analogous "transposed" template would have to be employed.

For ease of use, the commands that implement this search may be placed in a shell script, `memefinder1.ksh`. The script is available online at mitpress2.mit.edu/e-journals/Computer-Music-Journal/memefinder.tgz. If this archive is not expanded automatically, the Unix command `tar`

`xzvf memefinder.tgz` will extract its contents. `Memefinder1.ksh` repeats its operations on all `.krn` files in the current directory, rendering it suitable for searches of a large repertory. The findings are successively appended to a single `memeresults` file that may be inspected by the researcher after processing is completed. The operation of the script is summarized in Figure 6.

In conducting a search for the pattern it seems reasonable to restrict an initial investigation to Viennese music of the late 18th century. It was decided to use those string quartets of Haydn and Mozart that are publicly available from the CCARH Muse Data Web site (www.musedata.org), which contains the vast majority of Haydn's quartets (apart from some of the early works) and all of

Figure 3. Three representations of Mozart, Adagio in C Major for Glass Harmonica, KV 356 (617a; 1791), mm. 1–2.

(a) Implication-Realization Analysis; (b) Voice-Leading Reduction; (c) ``**Kern`` Encoding.

(a)

(b)

Mozart's quartets, together with a small selection of various other chamber works of Mozart.

A ``patt`` search using ``memefinder1.ksh`` and the template ``glass.patt`` produced the results shown in Table 1.

The examples from Mozart's KV 298 and KV 478 are the most salient, being clearly demarcated by durational interference, harmonic interruption, and metric (beat) differentiation. The passages from Mozart's KV 458 and Haydn's Op. 71, No. 1 are, by contrast, relatively open, "processive" tonicizations of vi in the context of larger descending I–vi–IV sequences (possibly itself memetic) and lack the demarcation engendered by Narmour's three closural criteria. Our musical judgment might be inclined to reject these latter examples for not representing clear examples of the pattern. As for its "strong" memetic status, we might reserve judgment on the "Glass Harmonica" pattern at this stage, insisting

(c)

**kern	**kern	**kern
*clefG2	*clefG2	*clefG2
*k[]	*k[]	*k[]
*C:	*C:	*C:
*M4/4	*M4/4	*M4/4
=1-	=1-	=1-
(4c/	2.ee\	(2.gg/
4e/	.	.
4g/	.	.
4g#/)	4dd\	4ff/)
=2	=2	=2
(2.a/	2dd\	(2ff/
.	4cc\	4ee/)
4gn/)	4r	4r
*-	*-	*-

`**deg`	`**deg`	`**deg`	`^1.*5$`
`*clefG2`	`*clefG2`	`*clefG2`	`^3`
`*k[]`	`*k[]`	`*k[]`	`^5`
`*C:`	`*C:`	`*C:`	`^(5+).*4$`
`*M4/4`	`*M4/4`	`*M4/4`	`^6.*4$`
`=1-`	`=1-`	`=1-`	`^6.*3$`
1	3	5	
3	3	5	
5	3	5	
5+	2	4	
=2	=2	=2	
6	2	4	
6	1	3	
`*-`	`*-`	`*-`	

that a copy be found in the work of another composer, thus implying connection in a nexus of imitation.

As noted earlier, after a preliminary search based on the configuration of a candidate pattern, it may be instructive to revisit the search template and consider ways of modifying it to account for variants that are derived from the antecedent form during the course of a diachronic-evolutionary process. As mentioned, it might be hypothesized—according to the rule given under point 3 in the second section of this article—that such variants will retain the shallow-middleground-level structure of the meme (Figure 3b) but will manifest a degree of variability in the foreground-level patterning that generates that middleground structure. Such variability of surface patterning might, for instance, affect the rhythmic synchronization of components of the pattern.

The template shown in Figure 7, `glass.pattern`, is designed to find versions of the pattern that, although broadly based on the middleground structure shown in Figure 3b, synchronize the two component linear strata in various ways. One might in particular speculate on a possible variable alignment of the upper-voice $\hat{4}$ against its supporting harmony, or the omission of the diatonic lower-voice $\hat{5}$; both seem likely evolutionary alterations. The use of record-count metacharacters in the template—essential to account for these two scenarios—necessitates the use of `pattern` as the search tool.

If such a `pattern` template is derived from a `patt` template to account for unpredictable surface elaborations of an underlying structure, then it will lose some of the precision of the `patt` template owing to its use of record-count metacharacters to account for the unknown extra elements—the diminutions that are symptomatic of memetic evolutionary change. Although Humdrum `pattern` searches are in principle suitable for locating stable middleground patterns beneath a changing foreground, such searches inevitably increase the chance of producing false positives. By this, we are referring to patterns that, although perhaps appearing to be descended from the antecedent at the foreground level, do not correspond at the middleground level, and are therefore—according to the rule given under point 3 in the second section—not consequents of the given antecedent. As may be seen by comparing Figure 5 with Figure 7, all but the first line of the `pattern` template contain variously the "+," "*," and "?" symbols or the Bool-
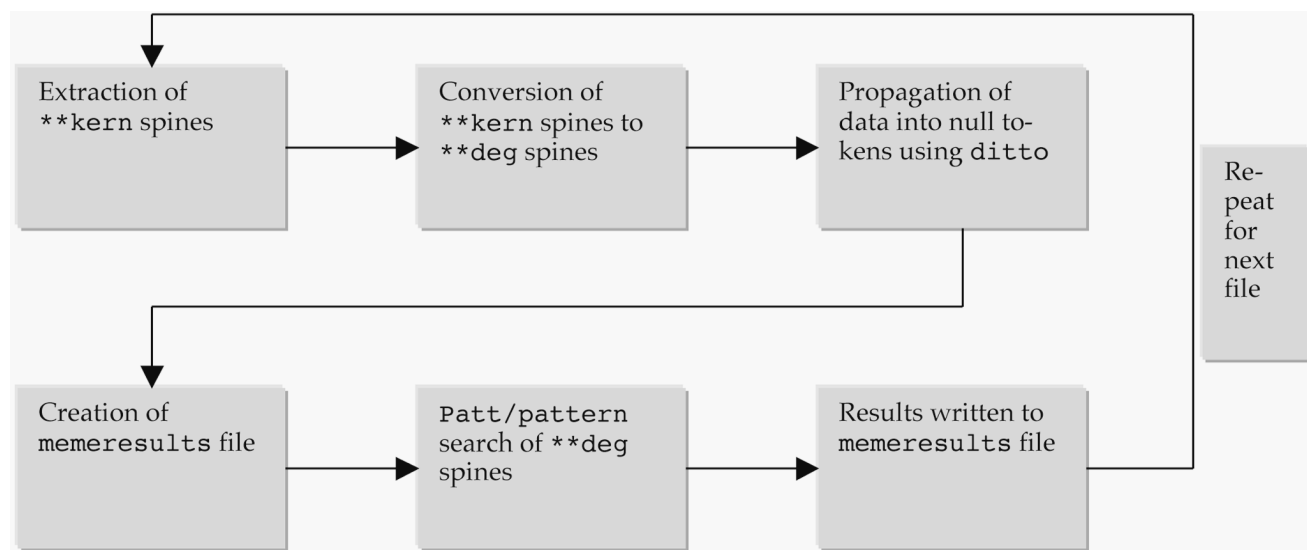
Table 1. Search Results (I) for the "Glass Harmonica" Pattern

Haydn	*Pattern found in*
String Quartet in B-Flat major, Op. 71, No. 1 (1793), III	mm. 37–38

Mozart	*Pattern found in*
Flute Quartet in A major, KV 298 (1778/1786?), II	mm. 5–6
String Quartet in B-Flat major KV 458 ("Hunt"; 1784), II	mm. 4–5; 24–25
Piano Quartet in G minor KV 478 (1785), II	mm. 1–2

ean OR operator ("|"). Clearly, any modifier that allows for a multiplicity of often very loosely defined elements, the absence of an element, or "alternativity" between elements, will tend to register as positives a great variety of middleground structures.

Humdrum ultimately cannot "know" the structural-hierarchic location of a given pitch (to my knowledge, no software yet available can perform Schenkerian analysis); `patt` and `pattern` do not have the higher-order algorithms necessary for this task. This type of complex determination, which often vexes seasoned Schenkerians, must be the task of the researcher. However, given that middleground elements often have greater metrical (accentual) and durational (agogic) weight than elements at the foreground level, then searches may be adjusted to account for these factors. The posi-

tion of the bar line—which often directly precedes a structurally important pitch—can easily be factored into a search template. For this reason, the template shown in Figure 7 adds barline signifiers ("^ =") before "^1.*5$" and before "^6.*4$" because these elements are markers of pitches interpreted, in Figure 3b, as having middleground status.

Nevertheless, the second of these components of the template is qualified by a metacharacter in such a way as to make its presence non-obligatory. `Glass.pattern` deliberately casts the net widely, yet this lack of precision has its advantages: the template registers a positive for the decorated and syncopated version of the pattern in mm. 25–26 of Mozart's Adagio that `glass.patt` cannot detect. Owing to this attribute, and because a middleground $\hat{3}$/vi may not necessarily directly follow a bar line, use of this template necessitates a subse-

Figure 7. The
`glass.pattern` *template*
for the "Glass Harmonica"
pattern.

```
^=

^1.*5$       +

^1    *

^3    +

^5    *

^(5\+).*(3|4)$ +

(^=)|(^6.*3$)

(^6.*4$)|(^6.*3$)    *

.*    ?

.*    ?

^6.*3$       *
```

quent manual interpretation of the results. Ultimately, in making such judgments, one is treading the fine line between false positives and false negatives: sifting and eliminating false positives arising from a less-discriminating template makes more effort for the researcher (undermining, to some degree, the benefits of computer-assisted investigations); however, false negatives are more serious, compromising the integrity of a set of results.

Using this revised template, `glass.pattern`, a `pattern` search using `memefinder1.ksh` (Figure 6) produced the results shown in Table 2.

The second search found an additional example

of the pattern, from the opening of the slow movement of Haydn's String Quartet in F major, Op. 74, No. 2. Given the distinct instances of the pattern in Mozart's KV 298, 478, and 356 (617a), and in this work of Haydn's, we are now justified in regarding it as a meme in the "strong" sense, for it is a particulate unit that would appear to have been transmitted (spontaneous generation aside) by imitation from Mozart to Haydn (or to Haydn by means of some other nexus of imitation).

The passage from Haydn's Op. 74, No. 2 demonstrates the attributes mentioned above in connection with likely evolutionary variants of the pattern (i.e., variable alignment of the upper-voice $\hat{4}$ against its supporting harmony and omission of the diatonic lower-voice $\hat{5}$). In finding a passage that `glass.patt` missed but which our musical judgment tells us is a clear example of the unit, `glass.pattern` appears to strike an appropriate balance between an overly narrow and an overly broad specification of the sought configuration.

Case Study 2: The "I-R Process" Pattern

To begin, one must devise a musical pattern that contains some combination of realized and/or denied implications and satisfies one or more of Narmour's criteria for closure. Figure 8 shows a suitable pattern.

This simple shape (which may exist, as indicated by the permutation slashes ["/"] in the example, in a variety of chromatic and rhythmic forms) is demarcated before its initial pitch and closed after its terminal pitch by rests, which enforce closure by

Table 2. Search Results (II) for the "Glass Harmonica" Pattern

Haydn	Pattern found in
String Quartet in B-Flat major, Op. 71, No. 1 (1793), III	mm. 37–38
String Quartet in F major, Op. 74, No. 2 (1793), II	mm. 1–2; 5–6; 75–76

Mozart	Pattern found in
Flute Quartet in A major, KV 298 (1778/1786?), II	mm. 5–6
String Quartet in B-Flat major, KV 458 ("Hunt") (1784), II	mm. 4–5; 24–25
Piano Quartet in G minor, KV 478 (1785), II	mm. 1–2

Figure 8. Candidate pattern.

metric (beat) differentiation, establishing it as a discrete perceptual/cognitive unit. As a rising scale pattern, it exemplifies Narmour's notion of process, the implication of the initial motion c–d(flat) being realized by continuation of the ascent to e(flat) and then f(sharp). Further continuation to g is suppressed owing to the durational interference of the dotted quarter-note (or longer) f(sharp) and/or the metric (beat) differentiation of the rest, which neutralizes the implicative charge and closes the grouping. Clearly, the pattern is not as distinctive or salient as that of Figure 3—it is a "commonality" on Cope's continuum—but its closural attributes nevertheless allow it to exist as a meme.

It is most efficient to conceive this "I-R Process" pattern in terms of its intervallic structure and (as indicated by the permutations in Figure 8) to allow for a variable sequence of major and minor seconds, so the pattern essentially represents a tetrachord starting on any degree of the diatonic or chromatic scale. One could specify more accurately its intervallic structure, but for present purposes, a wide variety of configurations will be accommodated; later, finer distinctions will be made. In conceiving a template, it is again necessary to convert the musical information to a representation independent of absolute pitch: most parsimonious would be to use the melodic interval or "**mint" representation, which indicates the direction of motion ("+" = up, "−" = down), interval category ("M" for major, "m" for minor, etc.), and interval size of a melodic progression. Figure 9 shows the pattern in both **kern and **mint representations.

An appropriate template file must contain regular expressions capable of identifying patterns in which the terminal pitch is a quarter note or longer followed by a rest. Without this capacity, patterns that terminate on note values shorter than a quarter note and/or with no terminal rest—and which

are therefore not, in Narmour's view, closed by durational interference and/or metric (beat) differentiation—will also be recorded as positives. On the basis of our working hypothesis, such patterns may not have sufficient end-closure for independent memetic existence.

To find patterns that match the sought intervallic criteria and closural properties, it is necessary to use the Humdrum `assemble` tool vertically to align the original **kern spine against the derived **mint spine (the result resembling Figure 9), and then to devise a template that can distinguish between the presence of integers in the **kern spine (which signify rhythmic values) and integers in the **mint spine (which signify intervallic steps).

A suitable template, `i-r_process.pattern`, is shown in Figure 10. In the penultimate line, it specifies that the first "4," "4.," "2," or "2." to be encountered must be followed directly by a pitch. That is, it must be only a **kern quarter note, dotted quarter-note, half-note, or dotted half-note and not, therefore, a **mint intervallic signifier, the desired values of which are specified at the end of this line by "((\+M2)|(\+m2))." For this search, `pattern` must be used, because record-count metacharacters are employed in several lines of the template, including "*" in lines 2, 4, 6, and 8 to account for the presence of any null data tokens interspersed between **mint signifiers. These, in this context, would indicate the presence in other spines of rhythmic values shorter than those of the sought pattern.

As with the first case study, the pertinent commands may be placed in a shell script, `memefinder2.ksh` (also available at mit-press2.mit.edu/e-journals/Computer-Music-Journal/memefinder.tgz), a variant of `memefinder1.ksh` (see Figure 6). Unlike `memefinder1.ksh`, however, this script extracts each spine of the input file in turn and subjects it individually to the actions of `pattern` to ensure that the sought pattern exists within a single spine and is not spread "diagonally" across two or more parts. The operation of the script is summarized in Figure 11.

Using the same repertoire as that employed for the first case study, a `pattern` search using `memefinder2.ksh` and the `i-r_process.pattern` template produced the results shown in Table 3. To

(a)

```
**kern

*clefG2

*k[]

*C:

*M4/4

=1

(4r

4c

4d

4e

=2

2f)

2r
*-
```

(b)

```
**mint

*clefG2

*k[]

*C:

*M4/4

=1

r

[c]

+M2

+M2

=2

+m2

r
*-
```

distinguish between the various forms of this figure (which the i-r_process.pattern template specifically did not), values in square brackets after bar numbers indicate the internal interval classes of the particular pattern, the careted superscript before these indicating its starting scale degree.

Ignoring rhythmic factors for present purposes, all intervallic variants of this pattern appear to be memes in the "strong" sense. Every type found in the Haydn movement also occurs in the three Mozart examples, suggesting their connection in a nexus of imitation—albeit a very broad one, given that this is such a basic figure, and one surely not restricted to Haydn and Mozart. Correlating this observation with initial scale degree, we see that $\hat{7}$1–2–2 and $\hat{5}$2–2–1 are common patterns in this sample. Nevertheless, there is clearly insufficient evidence to extrapolate wider conclusions about

the prevalence of this pattern and its variants in the dialect as a whole, given the small size of the results group and the distorting effect of intra-work repetition. Moreover, one should perhaps be careful not to make too much of such fine structural distinctions, owing to the limited number of configurations available within the diatonic tetrachord, a function of the intervallic characteristics of the major and minor modes.

Conclusion: Prospects

It seems clear that there is much to commend the use of the Humdrum Toolkit in developing a theory of memetics. The following conclusions can be drawn from the above investigations. First, a theory of musical memetics will be complex and multifac-

Figure 10. The
i-r_process.pattern
template.

Figure 11. Flow chart illus-
trating the operation of
Memefinder2.ksh.

```
r       +

^(\.)*

[a-gA-G]

^(\.)*

.*((\+M2)|(\+m2))$

^(\.)*

.*((\+M2)|(\+m2))$

^(\.)*

^=      ?

^(((4)|(4\.)|(2)|(2\.))[a-gA-G]).*((\+M2)|(\+m2))$

r       +
```

Figure 10

| Extraction of **kern spine 1 | Conversion of **kern spine 1 to **mint spine 1 | Alignment of **kern spine 1 with **mint spine 1 | Repeat for **kern spine 2, etc; then move to next file |

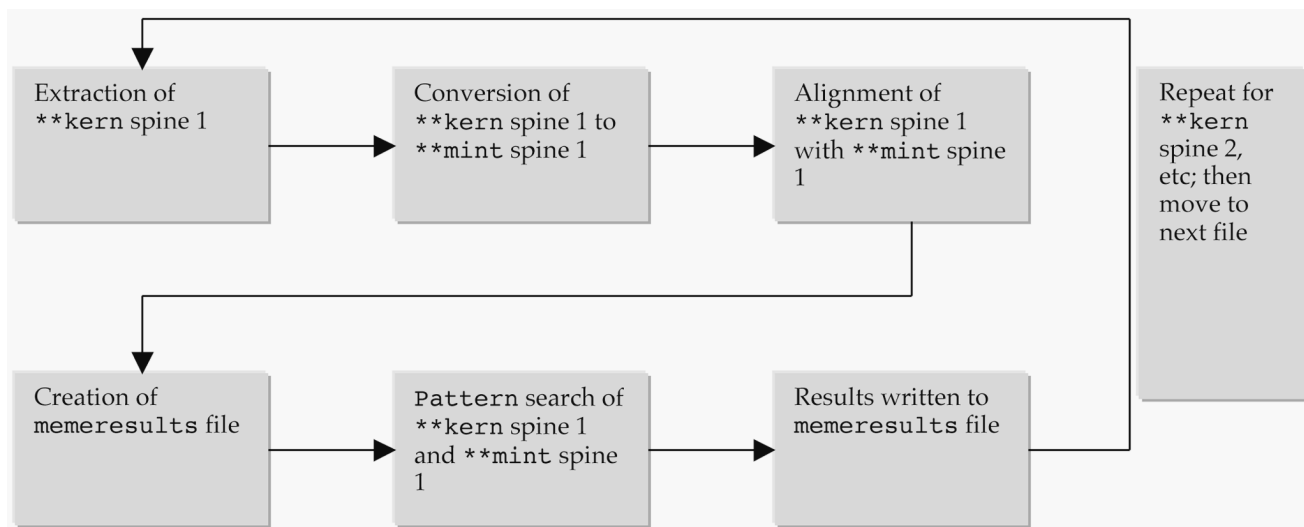| Creation of memeresults file | Pattern search of **kern spine 1 and **mint spine 1 | Results written to memeresults file |

Figure 11

Table 3. Search Results for the "I-R Process" Pattern

Haydn	Pattern found in
String Quartet in F-Sharp minor, Op. 50, No. 4 (1787), I	mm. 5 [$^{\flat}$2–2–1; $^{\sharp}$2–1–2], 15 [$^{\flat}$7–1–2–2], 120 (= m. 5) [$^{\flat}$2–2–1; $^{\sharp}$2–1–2]

Mozart	Pattern found in
String Quartet in F major, KV 158 (1773), III	mm. 28 [$^{\flat}$2–2–1], 29 [$^{\flat}$7–1–2–2]
String Quartet in A major, KV 464 (1785), IV	m. 4 [$^{\flat}$2–1–2]
String Quartet in D major KV 575 (1789), I	mm. 6 [$^{\flat}$7–1–2–2; $^{\flat}$2–2–1], 14 [$^{\flat}$7–1–2–2; $^{\flat}$2–2–1], 122 (= m. 6) [$^{\flat}$7–1–2–2; $^{\flat}$2–2–1], 130 (= m. 14) [$^{\flat}$7–1–2–2; $^{\flat}$2–2–1]

eted, requiring for its formulation both synchronic-structural and diachronic-evolutionary perspectives. Computers offer sophisticated means of facilitating research in both dimensions.

Second, one of the central problems in developing a theory of musical memetics is the necessity to formulate a robust, psychologically valid model of pattern perception. Narmour's i-r model, and the Gestalt principles upon which it is based, are arguably the best candidates at present.

Third, two useful strategies for verifying the memetic status of candidate patterns are 1) to identify a natural candidate pattern and then evaluate its prevalence in a given repertory; and 2) to synthesize an artificial candidate pattern using Gestalt/i-r principles and then attempt to locate real instances of it in a given repertory.

Fourth, the Humdrum Toolkit's `patt` and `pattern` utilities are powerful tools for searching representations of musical information and locating memes and are amenable to both the above strategies. The regular expression syntax they implement is well suited to memetic investigations on account of its flexibility and precision.

Finally, the results of such searches may be used both qualitatively and quantitatively. That is, they can foster analytical and critical discussion of specific musical works and the memes they contain, and they can facilitate statistical evaluation of the frequency and distribution of memes within idioms and dialects.

I hope to have shown that it is in the nature of computer-aided research that, perhaps more important than making a complex task more manageable, the highly structured approach required by the enterprise forces the researcher to interrogate assumptions, to rethink methodologies, and to refine the hypotheses upon which the investigation is based. In considering the prospects for this research, three issues in particular deserve mention.

The Geographical Distribution of Memes and the Generic Mapping Tools

The second part of the last point above suggests that memes might be tracked as they spread geographically. A mechanism for implementing this already exists, for, as described in Aarden and Huron (2001), Humdrum can be integrated with Wessel's and Smith's Generic Mapping Tools (available online at gmt.soest.hawaii.edu). This linkage might, in theory, permit the detailed memetic profiling of musical styles, allowing their transmission and evolution over time to be charted.

The Humdrum `Simil` Tool

As suggested earlier, the "Glass Harmonica" meme might be regarded as a mutation of a normative diatonic V–I or, less modified, a V–vi progression (Figure 2). The concept of mutation can usefully be described in terms of the notion of *edit-distance*, that is, the number of changes (additions, deletions, and substitutions) required to move from the antecedent to the consequent form. This scale, the *Damerau-Levenshtein edit distance* metric, has al-

ready been implemented by Keith Orpen in the Humdrum `simil` tool (Orpen and Huron 1992) and may prove useful in memetic research.

Investigating the Hierarchic Location of Memes

A third development—touched upon in the first case study—would be to search for patterns at deeper levels of structure. The patterns examined here have been short figures of foreground orientation based on simple shallow-middleground-level frameworks. The use of record-count metacharacters with `pattern` facilitates the identification of more distributed, ''virtual'' configurations that, if replicated, can still be regarded as memetic.

Acknowledgments

References

Aarden, B., and D. Huron. 2001. ''Mapping European Folksong: Geographical Localization of Musical Features.'' *Computing in Musicology* 12:169–183.

Blackmore, S. J. 1999. *The Meme Machine*. Oxford: Oxford University Press.

Cope, D. 2003. ''Computer Analysis of Musical Allusions.'' *Computer Music Journal* 27(1):11–28.

Dawkins, R. 1989. *The Selfish Gene*, 2nd ed. Oxford: Oxford University Press.

Gjerdingen, R. O. 1988. *A Classic Turn of Phrase: Music and the Psychology of Convention*. Philadelphia: University of Pennsylvania Press.

Huron, D. 1997. ''Humdrum and Kern: Selective Feature Encoding.'' In E. Selfridge-Field, ed. *Beyond MIDI: The Handbook of Musical Codes*. Cambridge, Massachusetts: MIT Press, pp. 375–401.

Huron, D. 2002. ''Music Information Processing Using the Humdrum Toolkit: Concepts, Examples, and Lessons.'' *Computer Music Journal* 26(2):11–26.

Jan, S. 2000a. ''Replicating Sonorities: Towards a Memetics of Music.'' *Journal of Memetics—Evolutionary Models of Information Transmission* 4(1). Available online at jom-emit.cfpm.org/2000/vol4/jan_s.html.

Jan, S. 2000b. ''The Memetics of Music and Its Implications for Psychology.'' In J. Sloboda et al., eds. *Proceedings of the Sixth International Conference on Music Perception and Cognition*. Keele, UK: University of Keele Department of Psychology.

Jan, S. 2002. ''The Selfish Meme: Particularity, Replication, and Evolution in Musical Style.'' *International Journal of Musicology* 8:9–76.

Jan, S. 2003. ''The Evolution of a 'Memeplex' in Late Mozart: Replicated Structures in Pamina's 'Ach ich fühl's.' '' *Journal of the Royal Musical Association* 128(2):330–370.

Kornstädt, A. 2001. ''The JRing System for Computer-Assisted Musicological Analysis.'' *Proceedings of the Second Annual International Symposium on Music Information Retrieval*. Bloomington: Indiana University Press, pp. 93–98.

Lerdahl, F., and R. Jackendoff. 1983. *A Generative Theory of Tonal Music*. Cambridge, Massachusetts: MIT Press.

Meyer, L. B. 1956. *Emotion and Meaning in Music*. Chicago: University of Chicago Press.

Meyer, L. B. 1973. *Explaining Music: Essays and Explorations*. Chicago: University of Chicago Press.

Meyer, L. B. 1989. *Style and Music: Theory, History, and Ideology*. Philadelphia: University of Pennsylvania Press.

Narmour, E. 1977. *Beyond Schenkerism: The Need for Alternatives in Music Analysis*. Chicago: University of Chicago Press.

Narmour, E. 1984. ''Toward an Analytical Symbology: The Melodic, Harmonic, and Durational Functions of Implication and Realization.'' In M. Baroni and L. Callegari, eds. *Musical Grammars and Computer Analysis: Atti del Convegno*. Florence: Olschki, pp. 83–114.

Narmour, E. 1989. ''The 'Genetic Code' of Melody: Cognitive Structures Generated by the Implication-Realization Model.'' In S. McAdams and I. Deliège, eds. *Music and The Cognitive Sciences*. London: Harwood, pp. 45–63.

Narmour, E. 1990. *The Analysis and Cognition of Basic Melodic Structures: The Implication-Realization Model*. Chicago: University of Chicago Press.

Narmour, E. 1992. *The Analysis and Cognition of Melodic Complexity: The Implication-Realization Model*. Chicago: University of Chicago Press.

Narmour, E. 1999. ''Hierarchical Expectation and Musical Style.'' In D. Deutsch, ed. *The Psychology of Music*, 2nd ed. San Diego, California: Academic Press, pp. 441–472.

Orpen, K. S., and D. Huron. 1992. ''Measurement of Similarity in Music: A Quantitative Approach for Non-Parametric Representations.'' *Computers in Music Research* 4:1–44.

Plotkin, H. 1995. *Darwin Machines and the Nature of Knowledge: Concerning Adaptations, Instinct and the Evolution of Intelligence*. London: Penguin.

Schellenberg, E. G. 1996. ''Expectancy in Melody: Tests of the Implication-Realization Model.'' *Cognition* 58:75–125.

Temperley, D. 2001. *The Cognition of Basic Musical Structures*. Cambridge, Massachusetts: MIT Press.

Thompson, W. F., and M. Stainton. 1996. ''Using Humdrum to Analyze Melodic Structure: An Assessment of Narmour's Implication-Realization Model.'' *Computing in Musicology* 10:24–33.