

Dynamic Melodic Expectation in the Creative Microdomain *Seek Well*

CRCC Technical Report 138

Eric Nichols

Indiana University, Bloomington, IN 47408 USA

May 30, 2005

Purpose

Listening to music is a complex mental activity that invokes multiple high-level cognitive processes. Listeners quickly form mental representations of metric structure, key, relation of pitches to a tonic, and melodic form. These representations naturally induce expectations about how the music will continue. Steve Larson has developed several computer models of these melodic expectations based on his theory of musical forces. This paper reviews the elements of the theory and the existing models before suggesting some specific ways in which methods developed by the Fluid Analogies Research Group (FARG) can more effectively implement Larson's ideas about melodic expectation. A collection of musical examples motivates the suggested extensions by highlighting some of the interesting dynamics of the listener's experience not captured by the existing models.

The Creative Microdomain *Seek Well*

Projects developed by FARG in the past decades have utilized *microdomains* to study cognitive mechanisms (Hofstadter 1995). Microdomains are restricted domains typically involving a much smaller problem space than typical problem domains. Larson (1993b) uses the term *creative microdomain* to further require that the process of devising solutions must arguably involve creativity. A particular creative microdomain developed by FARG called *Copycat* (Hofstadter, Mitchell 1994) involves letter-string analogy problems such as:

“If *abc* goes to *abd*, what does *xyz* go to?”

This problem is deliberately restricted such that the solution involving “wrapping around” from ‘z’ to ‘a’ is disallowed. Creativity is implicit in interesting responses to this question. Indeed, many answers are possible besides the forbidden *xya*, such as *xyd*, *xyy*, and the surprising *wyz*. Creative microdomains generally do not admit a single correct response, as the process of coming up with an answer is the primary object of study. Other FARG microdomains include *Seek Whence*, *Tabletop*, *Jumbo*, and *Letter Spirit*, with subjects ranging from number sequences to font design (Hofstadter 1995).

FARG has devised the *Active Symbols Architecture* to model creative solutions to problems in these microdomains. Central to this architecture is the notion of simulating parallel subcognitive processes that mimic the conflicting internal pressures of human cognition. Mental representations are formed, rejected, modified, and eventually settled upon in a stochastic, flexible, working environment (Hofstadter 1995). Creative perception, cognition, and expectation in these models is an emergent result of a great deal of seething subcognitive behavior.

Steve Larson and FARG developed the creative microdomain *Seek Well* to study melodic expectation (Larson 1993b). To reduce the potentially overwhelming complexity of the domain while retaining key features of interest, *Seek Well* involves melodies in the classical tradition (i.e. Western tonal music) conforming to the following list of restrictions:

- Only one note sounds at once; the sound represented is monophonic
- All notes have the same duration
- All notes sound the same except for pitch (dynamics, articulation, etc. are not involved)
- Rests are not allowed (notes come immediately one after another until the melody ends)

Figure 1a is a typical melodic beginning in the *Seek Well* domain. Larson studied typical responses to this cue by asking listeners to sing a completion of any length (Larson 1997). Two common continuations are shown in Figures 1b and 1c. Note how the response 1b implies the listener heard the cue in C Major, while response 1c implies F Major.



Fig. 1

Larson's Theory of Musical Forces

THREE MUSICAL FORCES

The theory of musical forces states that “we tend to hear music as purposeful action within a dynamical field of musical forces” (Larson 1993a), making an analogy between physical motion through space and the perceived “motion” of a melodic line. The three forces involved are musical gravity, magnetism, and inertia.

Gravity refers to a tendency of notes heard above a stable platform to descend to that platform. For example, given the rising line C-D-E in C Major (Figure 2a), gravity would suggest a continuation back down to the stable tonic: C-D-E-D-C (Figure 2b). After the alternate beginning C-D-C (Figure 2c), however, gravity does not continue to pull the line down past C, because C has been established as a stable base.



Fig. 2

Magnetism refers to the perceived attraction between notes of unequal stability. For example, in the ascending octave that stops on the leading tone, C-D-E-F-G-A-B (Figure 3a), we feel the strong “magnetic” force of the stable upper octave C “pulling” on the unstable B, leading to the expected completion in Figure 3b. In Larson’s model, the strength of the force is inversely proportional to the distance between two pitches (Lerdahl [2001] proposes a related model with a more direct relationship to physics, where the force is inversely proportional to the *squared* distance). The B would also be attracted downward to the stable G, but the magnetic pull of the upper C is stronger because it is closer as measured in half steps.




Fig. 3

Inertia is the tendency of a musical pattern to continue “in the same way”. A simple example is the tendency of a moving line to continue moving in the same direction. Given the same beginning as in the gravity example, C-D-E (Figure 2a), inertia suggests that instead of falling back down to C, the melody may continue rising through F, G, A, etc. as in Figure 3.

Musical inertia corresponds to the physical statement “bodies in motion tend to stay in motion”. Another example of inertia is an Alberti bass, which tends to continue its characteristic pattern once set in motion.

These three forces act continuously on musical lines in a dynamically shifting musical context. A significant part of this context is provided by pairs of *reference alphabets* and *goal alphabets* (Larson 2004). We say that musical lines move through a reference alphabet between notes of a goal alphabet. Reference alphabets consist of collections of notes that make up musical lines, such as the C Major scale in the examples above. A goal alphabet is a subset of an associated reference alphabet that marks stable points of arrival, such as members of the tonic triad in these examples.

It is illustrative to reconsider the ascending octave line (Figure 3b) in the context of all three forces. One may hear the ascending octave in many different ways, but one possible interpretation follows. The melody starts out at rest on C, the stable base. Suppose that we hear the first two notes as if an external force started the motion by “pushing” the melody up to D. The external force subsides and the melody begins to be tugged in various directions. Gravity exerts a constant force pulling downwards, back towards C. Magnetic forces are pulling with equal strength up to the E and back down to C. The inertia of the initial push keeps the melody going upwards to E, overcoming the force of gravity. At this point the magnetic pull of G becomes prominent as it is closer than the lower C. Magnetism accelerates the melody up through F towards G, and once reaching G it sails past due to inertia. However, the force of gravity is still noticeable, as is the strong magnetic pull of G once the melody ascends to A. Only inertia enables the melody to rise a bit further to B. Once B has been reached, the magnetic force of the upper C is irresistible and overcomes both gravity and the magnetism of G, pulling the melody all the way to C.

The dynamic push and pull of the forces is significant in that it acknowledges some of the complexity of how we listen to music. The flexible interactions of the forces can also provide explanations for alternate melodic continuations. In the ascending line above, the melody might have changed direction before reaching C; perhaps the initiating “shove” was not strong enough to build up inertia to carry the melody past F, A, or even the initial D (Figure 4a,b,c below). In each of these cases, the inertial ascent “ran out of steam”, overcome by gravity, magnetism, or both. Musical forces can also explain more extended motion such as that of Figure 4d: the initial ascent to A is followed by a descent via gravity and magnetism back down through G, a brief lift back up due to the magnetic pull of G, and a final giving-in to gravity to return to the stable platform of C (after briefly dipping below the base due to inertia).



Fig. 4

Computer Models Incorporating Musical Forces

Larson's computer models of melodic expectation illustrate the theory by generating explicit predictions of melodic continuations. The *single-level* model is based on the three forces and built-in pairs of reference and goal alphabets. The *multi-level* model additionally uses information provided by an embellishment hierarchy (a simplified Schenkerian analysis).

SINGLE-LEVEL MODEL

The single-level model takes a melodic beginning as a cue and produces predicted completions of the cue based on gravity, magnetism, and inertia. For each of the three forces, the model analyzes the pitches in the cue to choose pairs of reference and goal alphabets from a predefined list. Several rules help ensure that the alphabets chosen make musical sense in the context of the force under consideration. For instance, when making a gravity or magnetism prediction, the cue must end on a pitch that is present in the reference alphabet but not in the goal alphabet (that is, the final note of the cue must be heard as unstable so that there is an impetus for continuation). Next, the model examines the final notes of the cue and produces predictions based on these notes and the pairs of alphabets under consideration. These predictions involve motion through the reference alphabet to a stable member of the goal alphabet, providing a sense of completion. The predictions are assigned probabilities according to the strength of the forces as defined by the theory. Thus, the output from the model is a set of different predictions and associated probabilities. Finally, if desired, the resulting predictions can be fed back into the system as new input so that it can predict what will follow each of those possible continuations, generating more extended predictions.

MULTI-LEVEL MODEL

The theory of musical forces applies not only to the surface-level notes in a melody, but also to deeper levels of an embellishment hierarchy (a Schenkerian analysis). In Figure 5 there are two levels in the hierarchy. Staff A represents the deepest (background) level, while staff B depicts the foreground level, made up of the notes actually heard by the listener. While there are only two levels of hierarchy in this example, there could be several middleground levels in more extended examples. Any discussion in this paper of relations between the background level and the surface applies equally well between any pair of levels.




Fig. 5

The current implementation of the multi-level model requires the user to provide not only the surface notes of the melodic beginning and the key of the melody, but also a specific type of Schenkerian analysis. In this case, both staves A and B would be supplied to the computer, along with a description of the embellishment function of those notes not present in the background level. For instance, the user might describe the initial note B as a “prefix chromatic neighbor” to the following C. This type of analysis differs from a typical Schenkerian analysis in its specificity: the precise embellishment function is explicitly described in the analysis. At the same time, the theory also claims that the analysis must be flexible in describing the embellishment function. For instance, in some musical contexts a description such as “suffix diatonic lower neighbor” might be appropriate, while in other cases the description might be generalized to “suffix diatonic neighbor”.

The multi-level model considers each level as an individual melody line and generates a completion of each line. In this example, the *multi*-level model would begin a prediction by applying the *single*-level model to staff A. The single-level model would generate middle C as a likely final note, due to the influence of both inertia and gravity. (The other possibility, discussed below, would be a return to G based on the magnetism between E and G.) Here the reference alphabet in use is the C Major triad, with a goal alphabet made up of only the tonic and dominant. Figure 6 demonstrates the partially completed prediction after this step.




Fig. 6

Having generated a completion for the background level, the multi-level model proceeds to complete the surface level. This step depends on the particular analysis supplied to the system. In the current example, if all the embellishment notes are described as “chromatic lower-neighbor prefixes”, the force of inertia indicates that to continue “in the same way” the final C should also have this type of prefix embellishment. The word “chromatic” indicates that the alphabet to be used for the prefix note is the chromatic scale (as opposed to the triadic alphabet used in the background level). Hence, the predicted embellishment note is the note B as shown in the analysis in Figure 7. In this diagram, the slurs connect embellishment notes (shown without stems) with notes present at a deeper level in the analysis.



Fig. 7

Improvements to the Multi-Level Model

The remainder of this paper focuses on additional features that could be added to the multi-level model. I claim that that these features may be implemented naturally using the FARG Active Symbols Architecture. Larson alludes to the possibility for such an implementation:

Although I have numbered these instructions, we should consider them as taking place in parallel, and influencing one another, until a potential completion is chosen. (Larson 2004)

The instructions mentioned describe in some detail how the model should generate predictions of melodic completions by automatically generating a hierarchical embellishment analysis including information about the key, mode, and meter, and then using this analysis to determine reference and goal alphabets and to make predictions based on the three musical forces. A complete implementation that incorporates all the instructions in parallel “will most likely be quite informative about music cognition.” (Larson 2004)

The model currently works in a serial fashion, generating predictions at the deepest structural levels first. However, there are cases where the surface levels can have a more direct impact on predictions. In general, musical forces working at different structural levels can conflict with one another, indicating the need for the model to consider the impact of such conflicts when making predictions. Additionally, other important musical factors such as implied metric structure and implied harmony affect how a listener hears an example and creates completions. Finally, the model is currently dependent on a provided embellishment analysis, but:

A complete specification of the theory might describe how hierarchical representations of embellishment structure may be built up and internally represented. (Larson 2004)

The next section presents a brief outline of a new model with an architecture to address these concerns. Future work will be required to describe the architecture in sufficient detail to allow computer implementation.

SYSTEM ARCHITECTURE

The desired extended version of the multi-level model is conceived as a recasting of the model in terms of a dynamic simulation. In particular, Larson's list of instructions is transformed into a set of behaviors of simulated subcognitive agents. The proposed system draws inspiration from other FARG models that function via a parallel terraced scan. The simulation involves a collection of codelets working together to 1) build up a shared representation of a short melodic fragment input to the system, called the **cue**, and 2) induce a natural expected continuation, treating the fragment as the beginning of a longer melody. The codelets act in parallel to explore various representations and continuations until the system judges the result to be satisfactory.

Although some names and ideas of the main components below are borrowed from other projects (Hofstadter 1995), each of those projects uses the components in a domain-specific manner. Likewise, in this case certain details are unique to the *Seek Well* microdomain. The main components of the architecture are the **Controller**, **Workspace**, **Coderack**, and the **Conceptual Memory**.

Controller

The Controller component coordinates the various parts of the system to provide overall control. It maintains a record of the system *temperature*, which indicates how satisfied the system is with its current predicted completion. When the simulation starts running, the Controller directs the other parts of the system to build up a representation of the given cue. This representation is built and stored in the Workspace. Then the Controller directs the system to generate possible completions of the cue. It stops the process when the temperature reaches a low enough value.

Workspace

The Workspace is the cognitive space where an internal representation of a musical cue is generated, stored, and modified. When the system starts running, the representation includes only the given cue, but as the system runs it builds up an embellishment analysis of the cue and eventually the expected continuation. The Workspace can be pictured as a bustling construction area: workers, called **codelets**, are busy working in the Workspace to construct an embellishment hierarchy in the same manner that a group of construction workers might work together to build a multi-story house. However, the analogy must be extended to allow for flexibility: instead of working from blueprints, imagine the construction workers building the house quickly in an improvisatory manner. These workers have no qualms about tearing down a wall if they decide it was badly planned. The musical workers build up structures

such as groups of notes, notes at deeper hierarchy levels, or analytical links between notes. They may also tear apart these structures when necessary to try different ideas (i.e., major changes would be made if the system temperature remained too high).

Coderack

The Coderack stores a large collection of codelets that each have extremely simple individual tasks to perform. Depending on the demands of the musical context, the Controller will select certain codelets from the collection to join the construction process in the Workspace. Some codelets are intended to focus on bottom-up tasks such as analyzing individual notes, while others deal with top-down ideas such as considering high levels of the embellishment hierarchy. Certain codelets will analyze groups of notes to determine the key or meter of a cue. Others will determine whether certain notes are more structural or more embellishing in function. Finally, some codelets will be involved with constructing actual predicted continuations: at each level of the hierarchy, there may be different codelets building continuations based on the forces of gravity, inertia, or magnetism.

Conceptual Memory

The Conceptual Memory is a fixed representation of the musical knowledge of the model and acts as a source of long-term memory. Examples of the knowledge in the memory might include the concept of interval size, stepwise vs. disjunct motion, the tendency for melodies to build to a climax on a high pitch, understanding of sequence, known pitch alphabets such as the major and minor scales, and simple harmonic schemes such as $V \rightarrow I$. Codelets will use knowledge in the Conceptual Memory to perform their tasks. Thus, changes in the Conceptual Memory will influence codelet behavior. Conversely, codelets may modify knowledge stored in the Conceptual Memory. For instance, codelets noticing interesting features of the music in the workspace may cause concepts in memory related to those features to become more active. In general, the Conceptual Memory provides an important top-down influence on the processes in the workspace: active concepts are persistent, in contrast to the ephemeral nature of the individual codelets.

ARCHITECTURE APPLICATIONS

This section illustrates how the proposed architecture could support the following features:

- Providing flexibility for embellishment descriptions
- Resolving conflicts between different hierarchy levels
- Inferring metric structure
- Generating a hierarchical embellishment structure

Flexibility of Embellishment Description

The completion given by the multi-level model for the previous example (Figure 7) was straightforward but simplistic. The theory of musical forces is applicable at all levels in the analysis, but in the example the only force evident at the surface level was inertia. Gravity and magnetism played a role in the background structure, but the constraints of inertia in the form of the embellishment analysis (diatonic lower-neighbor prefixes) dominated the surface level. Two observations play a role in liberating the surface level:

- The embellishment analysis could be less restrictive
- Musical forces at the surface level may counteract background level forces

In this example the analysis can be loosened by modifying or removing words in the description. For instance, changing the phrase “prefix chromatic lower-neighbor” into “prefix chromatic neighbor” yields two alternatives for the note preceding the final C. In the proposed architecture there would be codelets designed to modify the embellishment description as necessary to try out these alternatives. If a codelet chose D \flat , instead of B it would lead to the continuation in Figure 8.



Fig. 8

After adding the D \flat to the continuation, a codelet designed to detect linear motion might come along and find the notes E-D \flat -C in the completion. This codelet could draw on the conceptual memory to notice the augmented second between E and D \flat . If the codelet decided that interval was problematic, it could summon other codelets to fix the problem. For instance, it might request modification of the embellishment description to allow a different note than D \flat . Considering diatonic instead of chromatic neighbors, for instance, would produce the continuation of Figure 9.



Fig. 9

Codelets assigned to evaluate this continuation would notice that following the motion from E to D, directional inertia predicts a continuation that moves down from the unstable note D to the stable C. This provides additional confirmation of the earlier prediction that the motion C-G-E at the background level will continue down to C. Moreover, the final D-

C motion gives in to the force of gravity. Separate codelets would notice the stepwise linear motion E-D-C (indicated with a dotted line above). This additional structure would raise the system's satisfaction (i.e., lower the system temperature.) Thus, codelets making predictions based on different musical forces at both the background and foreground levels would agree that this is a likely completion.

Conflicts Between Levels

In Figure 9 the final note C was predicted at both levels by multiple musical forces. It is clear that in some cases individual forces may conflict with each other: inertia may indicate that a line will continue to rise while gravity pulls the line in the opposite direction. Returning to the initial cue in Figure 5, consider the effect of musical forces at the surface level. Inertia predicts that the line will continue to move upwards through the chromatic scale after the final two notes, D \sharp -E. The line might continue up to the more-stable note G, generating the motion from D \sharp through the chromatic reference alphabet to G as in Figure 10.



Fig. 10

This is a reasonable completion, but certain codelets in the Workspace might not be completely satisfied. In particular, a codelet concerned with metric structure could notice that the arrival on the final G has been delayed a beat longer than expected due to the five-note long chromatic line between D \sharp and G. This codelet could suggest dropping one of the notes to make the G into the 8th note of the cue, as expected. This leads to Figures 11 and 12, depending on the choice of notes to remove.



Fig. 11



Fig. 12

Figure 11 might be preferred if the active concepts in the Conceptual Memory include a high degree of chromaticism. Figure 12, on the other hand, would be favored by codelets looking for diatonic linear motion.

Note that in all three completions (Figures 10, 11, and 12), directional inertia at the surface level overcomes the force of inertia at the background level. It also overcomes the force of gravity at both surface and background levels. Although the continuation up to G in these figures contradicts the background-level prediction of moving down to C based on gravity and inertia (as in Figure 9), it does agree with the magnetism prediction at the background level (E is attracted to G more strongly than C because G is only three half-steps away). Due to the stochastic nature of the model, any of these completions might be generated, although some might be more likely than others.

We have seen several particular completions of the given melodic beginning. The choices made in the completions depend on the system's response to the conflicting pressures based on different musical forces operating at different levels, as well as pressure from other musical factors. The parallel terraced scan model of multiple codelets considering the possibilities in parallel provides an effective way for the system to explore many alternate continuations. When many separate codelets begin to agree on a particular continuation, confirmed at various hierarchical levels, this consensus translates into a lower Workspace temperature and a higher probability generated by the model for that continuation.

These examples have focused on how the architecture uses musical forces to generate and justify certain choices of pitches in the expected completion. Next, we turn to the question of simulating the perception of the metric structure of a cue.

Implied Metric Structure

Consider the following beginning of a melody:



Fig. 13

After hearing this beginning, a listener might expect the simple gravity or inertia predictions given earlier: C-D-E-D-C or C-D-E-F-G. Inertia might even predict a continuation all the way up through the octave: C-D-E-F-G-A-B-C.

Now, consider the same melody with one additional note:



Fig. 14

This might lead to gravity or inertia predictions as follows: C-D-E-E-D-C (Figure 15a) or C-D-E-E-F-G (Figure 15b). But these predictions seem to be in a different meter than those implied by the previous example. Indeed, continuing via inertia through the entire octave naturally yields C-D-E-E-F-G-G-A-B-C (Figure 15c).




Fig. 15

This pair of examples (Figures 13 and 14) gives a hint as to the subtlety of the metric problems that arise even in this restricted microdomain. To understand how the change in meter arises in Figure 14, we first reconsider Figure 13. Its predicted continuations imply that the melody is heard in duple meter (based on two or four beats per measure); some possibilities are shown in Figure 16.



Fig. 16

Even in this simple example, several factors combine to suggest duple meter. In the new model, there would be codelets and ideas in the Conceptual Memory corresponding to each of the following:

- **Primacy** – In the absence of other factors, there is a natural tendency to hear the first note as a downbeat (Toiviainen 2003).
- **Key** – This example starts on C and ends on E, suggesting the key of C Major. Because both of these notes are in the tonic triad, the listener may focus on these notes, making these likely candidates for metric emphasis.
- **Familiar pattern** – Similarly, the note D sounds like a passing tone between C and E triad tones, reinforcing the strength of the E note while deemphasizing the D.
- **Prevalence of duple meter** – Listeners of Western tonal music tend to assume that music is in duple meter unless factors suggest otherwise (Abecasis et al., 2005).

Returning to the second example (Figure 14), an additional factor is present that may suggest triple meter to a listener: the note E has been repeated. After the first three notes have been heard, predictions based on musical forces expect the line to continue rising or to reverse direction and descend. Especially if repeated notes are a significant idea in the Conceptual Memory, the model can determine that halting the rising motion and repeating the E indicates that something interesting is happening. This focuses attention on the second E and suggests that it may be in a stronger metric position such as a downbeat.

A codelet noticing the repeated E could summon other codelets to analyze its significance. For example, a codelet analyzing hierarchical structure could suggest that the second E is heard as a note from a deeper structural level, yielding the beginning of an analysis shown in Figure 17.




Fig. 17

At this point the model is beginning to form an analysis including meter, key, and hierarchical structure. The Workspace would now contain a representation of the given cue and the beginning of a hierarchical analysis, so the Controller would direct codelets to start building a continuation. We continue our discussion of the new model by examining three possible continuations based on this initial analysis of the cue.

CONTINUATION A

Considering the background level first, the C and E notes suggest a reference alphabet of the tonic triad, with the tonic and octave comprising the goal alphabet. Inertia and magnetism both point to continuing upwards to G instead of giving in to gravity and falling back to C. Thus, a likely continuation completes the arpeggiation C-E-G-C at the background level (Figure 18).




Fig. 18

At the surface level, the rules used by the single-level model suggest that the reference alphabet is the major scale, while the goal alphabet only includes the notes of the tonic triad. Although choosing alphabets was driven by rules in that model, in the new model there would be codelets assigned to pick out appropriate alphabets and communicate the alphabet choice to the rest of the Workspace. Similarly, the multi-level model would generate an inertia-based prediction for the next surface-level notes by using the analysis of the embellishment notes D and E and continuing to embellish the background notes in the same way. In the new model, analysis codelets would suggest different possibilities for how to analyze these notes. One likely analysis would be to describe the D and E as “moving up through two notes of the reference alphabet”. Then codelets would generate an inertial continuation by adding F-G after the background E in measure 2 followed by A-B in measure 3. Finally, after reaching the final C in the upper octave, the model would end the completion immediately instead of appending the suffix D-E to the end. Stopping after the C is indicated because both the surface-level and background-level lines have reached a pitch in the goal alphabet. Moreover, the melody has moved upwards for a long time without reversing direction, so there is pressure to stop or reverse direction. The final completion is shown in Figure 19.



Fig. 19

CONTINUATION B

A different continuation might result if a codelet using gravity predictions suggests that the line at the background level should give in to gravity at some point. For instance, after C-E-G the melody might turn around and fall back down to C. Filling in the surface level as before results in the problematic continuation of figure 20.



Fig. 20

A codelet looking for melodic continuity would notice that although the first three measures consist entirely of stepwise motion, there is a surprising leap from B to E at the end of measure 3. Inertia and magnetism predictions at the surface level indicate that the G-A-B line should continue up to the goal tone of C, just as in Figure 18. The lack of continuity due to the leap indicates a possible source of the problem: the analysis of the two suffix tones in each measure is too restrictive. The codelet could summon another codelet to generalize the analysis. Instead of requiring movement upwards, the analysis might be changed to “move through the reference alphabet for two steps.” Then the surface level notes would be free to move in the direction of the background notes to increase continuity, leading to the completion in Figure 21.



Fig. 21

This potential continuation lacks the strength of the first (Figure 19) due to its mundane symmetry (with a repeated high note in the very middle of the phrase) and the implied tonic harmony throughout. The shift in direction in measure 3 weakens the background structure, which was able to climb to the upper C in the first version. The first version also implied a stronger harmonic motion by a shift to dominant harmony in measure 3 followed by a cadence. This continuation would thus result in a higher system temperature than the first.

CONTINUATION C

Codelets analyzing the results of the previous continuation might be satisfied with the results, leading to a system temperature low enough to end the simulation. Alternately, they might notice possibilities for improvements and increase the system temperature to encourage more activity to fix the problems. Two issues in the continuation do seem promising as places for departure: the lack of harmonic changes and the lack of an implied cadential pattern. After an increase of system temperature, codelets could modify previous parts of the analysis, resulting in new possibilities for continuation that could address these issues. For example, a different reference alphabet could be selected for use at the background level. If the diatonic scale were used as the alphabet, the first two measures would suggest a Schenkerian initial ascent from C to E, which could be followed by stepwise motion through the scale.

In particular, consider motion to a neighbor tone of E after the second measure. If the second measure is the same as in the previous continuations (E-F-G) then it is important to preserve the continuity of motion in the third measure. Choosing D at the start of the third measure causes a surprising leap at the surface level. Starting that measure with F, however, provides a stepwise connection. This also implies non-tonic harmony. The background level can be completed by returning to E due to the forces of gravity and magnetism: C-E-F-E (see the stemmed notes of Figure 22). This E is a reasonable place to end the completion because it sounds relatively stable, although a longer completion that returns to the tonic is also likely. Because E is in the tonic triad, the last two measures will sound like an imperfect cadence. Of course, each of these steps would be modeled by codelets charged with noticing and fixing the problems.

Completing the surface level proceeds much as before: to increase continuity the particular embellishment description used might be “move through the diatonic alphabet by step towards the next background pitch”. Then the F in measure 3 would be followed by E-D. This leads nicely into the final E via the “circling” motion F-G-F-E-D-E starting in measure 2. This continuation causes the final E to sound as if it is exerting a magnetic pull on the melody as it circles in. The melody seems unable to stop on E in measure 3 due to inertia, but circles back around to come to rest in measure 4 (Figure 22).



Fig. 22

In this completion, note that musical forces have given rise to the musical gesture of circling around E, even though the theory itself does not contain any explicit notion of this

type of gesture. Codelets could be available to notice when a completion has resulted in such motion, making use of typical musical gestures stored in the long-term Conceptual Memory. Recognizing a typical gesture would lead to a reduction in system temperature and greater satisfaction with the results. Naturally, the simulation might stop when reaching this completion, but everything depends on the active concepts in memory. If the state of the simulation favors longer completions all four measures might be considered to be an antecedent phrase, and codelets would be set in motion again to build a consequent phrase. In this simple context, however, it is expected that the continuation in Figure 22 would result in a high degree of system satisfaction.

Lessons from the Examples

The example continuations explored in this paper illustrate how Larson's theory of musical forces functions in the context of the proposed extensions to the multi-level model. Examining how the model would generate melodic continuations highlighted several areas for further study and implicitly set up goals for an implementation of the new model. These areas include:

- *Flexibility of embellishment analysis.* When generating a continuation based on inertia, the particular analysis selected is essential. If surface-level problems occur in a prediction, examining the analysis and generalizing appropriately may resolve the issues while retaining the feeling of inertia.
- *Conflicts between forces and between analysis levels.* Musical motion at the surface level can put pressure on deeper structural levels, or vice versa. This supports the essential idea that musical forces generate a “dynamic field” in which musical lines move. The interactions between forces and levels can be complex, motivating the incorporation of the parallel terraced scan architecture into the model.
- *Implied metric structures.* Simple changes in a melodic beginning can imply different meters. It is essential to model how listeners hear the implied meter in order to generate musical continuations.

Future Work

The theory of musical forces suggests many avenues for further research. Of particular interest is the possibility of using musical forces to help determine the meter, determine the key, and generate an embellishment hierarchy for the beginning of a melody. Ideally, a model of melodic expectation would require as input simply a list of pitches, without the extra indication of the key, meter, or analysis. A future implementation of the suggested architecture will provide a means to test these ideas empirically.

References

- Abecasis, D., Brochard, R., Granot, R., & Drake, C. (2005). Differential Brain Response to Metrical Accents in Isochronous Auditory Sequences. *Music Perception* 22(3): 549-562.
- Hofstadter, D. R. & Fluid Analogies Research Group. (1995). *Fluid concepts & creative analogies: computer models of the fundamental mechanisms of thought*. New York, Basic Books.
- Hofstadter, D. R. & Mitchell, M. (1994). The Copycat Project: A Model of Mental Fluidity and Analogy-Making. In Holyoak, K. and Barnden, J. (eds.) 1994. *Advances in Connectionist and Neural Computation Theory Volume 2: Analogical Connections* pp 31-112. Norwood NJ: Ablex Publishing Corporation.
- Larson, S. (1993a). Computer Models of Melodic Implication and Key Determination in Tonal Music. Presented at the meeting of the Society for Music Perception and Cognition, Philadelphia, Available (from Center for Research on Concepts and Cognition; 510 North Fess; Bloomington, IN 47408) as CRCC Technical Report #77.
- Larson, S. (1993b). *Seek Well: A Creative Microdomain for Studying Expressive Meaning in Music*, Available (from Center for Research on Concepts and Cognition; 510 North Fess; Bloomington, IN 47408).
- Larson, S. (1997). *Continuations as Completions: Studying Melodic Expectation in the Creative Microdomain Seek Well*, Available (from Center for Research on Concepts and Cognition; 510 North Fess; Bloomington, IN 47408) as CRCC Technical Report #117.
- Larson, S. (2004). Musical Forces and Melodic Expectations: Comparing Computer Models and Experimental Results. *Music Perception* 21(4): 457-498.
- Lerhadl, F. (2001). *Tonal Pitch Space*. New York: Oxford University Press.
- Toiviainen, P. & Snyder, J. (2003). Tapping to Bach: Resonance-Based Modeling of Pulse. *Music Perception* 21(1): 43-80.