CHAPTER THREE

Musicat's Domain

Introduction

Each program based on the FARG architecture — Copycat, Tabletop, Letter Spirit, Phaeaco, Seqsee, and so on — operates in its own microdomain, as was discussed in the previous chapter. The use of *micro*domains in these models is not only humble and realistic, it also helps programs focus on particular aspects of cognition without being overwhelmed by real-world domain complexity. This chapter describes Musicat's microdomain and gives examples of the types of music the program is intended to handle. In brief, the microdomain is made up of simple, singable Western melodies, including folk tunes such as "Twinkle, Twinkle, Little Star" and "Row, Row, Row Your Boat"; popular melodies such as "On the Street Where You Live" (from the musical *My Fair Lady*) are also part of the domain, albeit on the more sophisticated end of the spectrum of what Musicat is meant to "listen" to.

Although the spirit of FARG projects is to work in microdomains, for the sake of musical interest we have expanded Musicat's microdomain over time enough that the word "micro" seems less applicable than it did for previous FARG programs. Therefore, sometimes (as in this chapter title) I omit the "micro-" prefix and simply write "domain". Compared with music in general, however, this domain is certainly "micro" — just think of the

complexity of Bach's Inventions, not to mention his *Well-Tempered Clavier*, or Monteverdi's Madrigals, Chopin's Nocturnes, Shostakovich's Preludes and Fugues, Mahler's Symphonies, Rachmaninoff's Piano Concertos, Wagner's Operas, and so on and so forth. And if we expand this list further to include atonal music, contemporary popular music, jazz, musical theatre, minimalism, electronica, or even different musical traditions entirely such as Balinese *gamelan,* Bhutanese *zhungdra*, or Brazilian *choro*, we see that music is almost impossibly diverse even for a human listener — let alone a computer model — to understand.

A Note on the Universality of Musicat's Domain

Musicat is designed to simulate the experience of listening to relatively simple, monophonic, Western, tonal melodies; it does not attempt to handle the wide variety of musical styles listed above. In this sense, Musicat operates in a very limited musical space. Much of the complexity of listening to music, however, is encountered even in simple-sounding melodies, so Musicat's domain is still highly complex, the specific restrictions described in the next section notwithstanding.

What is it about modeling the listening process in Musicat's domain that may be applicable to modeling listening for other musical styles? Certainly, some aspects of Musicat's approach are style-specific. But in general, the approach should transfer well, because it is based on many features that all types of music-listening have in common, including music's temporal nature, human memory systems, gestalt principles, and the centrality of analogy-making. Anything we call music is based on structured sound events arranged in time, so Musicat's handling of issues of temporal perception should be applicable to all types of music. Human memory systems are, obviously, implicated in all music listening (and in general all of cognition). Gestalt principles that inform perception of musical groups and

musical expectations also typically operate at the lower levels of perception, and should cross over to most musical contexts. Finally, analogy is central to Musicat's listening strategy, and if analogy is the core of cognition, as Douglas Hofstadter argues, then it follows (trivially) that analogy is also the core of *music* cognition (see Chapter 4). Musicat is better equipped to listen to extremely simple Western melodies such as "Twinkle, Twinkle, Little Star" for the time being, and it incorporates domain knowledge that it specific to Western tonal music, but its foundations should apply to music-listening in general.

Microdomain Specifics

Melody Characteristics

Musicat was originally conceived as an extension of Steve Larson's Seek Well project. It has diverged significantly since the first version (as described in Chapter 9), but it will be helpful here to explain how Seek Well influenced Musicat's microdomain. Indeed, Larson used "Seek Well" not only as the name of his program, but also as the name of the microdomain itself, as in "the creative microdomain Seek Well" (Larson, 1997a). Melodies in the Seek Well microdomain followed the set of rules described in Chapter 2. Musicat's domain is less restrictive, in the following ways:

- Only one note sounds at once, just as in Seek Well. In early versions of Musicat a simple bass line was allowed to describe chord progressions, but this is not done in the most recent version.
- Notes can have the following durations: whole, half, quarter, eighth, sixteenth, and dotted versions of each of these except for sixteenth. Notes can be tied to immediately following notes. Note that notes shorter than

sixteenths, triplets, and other more complex types of durations are not supported.

- The only attributes that distinguish one note from another are *pitch* and *duration*. In earlier versions of Musicat, notes could be explicitly accented in the input.
- Rests are allowed, of any of the durations available for normal notes.

Extra Information Included with Melodies

One of the interesting features of the Seek Well microdomain was that key signature and time signature were not specified; programs had to infer these properties themselves (just as human listeners do). For simplicity, we provide key and time signatures to Musicat. The metric position of the first note, which may be an upbeat, is also provided (see Chapter 8 for more details). The key and time signature must stay the same throughout the melody.

In a similar vein, note that Musicat is given musical input in a symbolic form much like sheet music. The program knows from the outset that it is hearing a "quarter-note C" or "whole-note F"; it doesn't listen directly to sound waves, and no audio processing via Fourier transforms or their ilk is required. Audio analysis at these lowest levels is out of the scope of this work.

Two Examples

The next section gives detailed descriptions of what it is like for humans to listen to melodies of the sort expected by Musicat. As a preview, here are two of the melodies, which abide by the rules listed above.

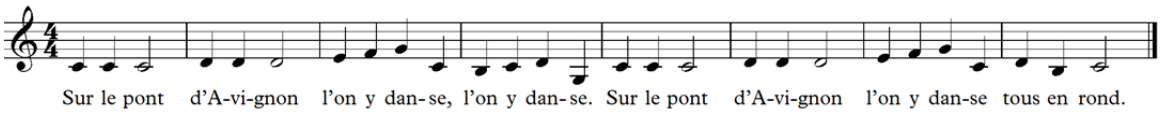

Figure 3.1: "Sur le pont d'Avigon", a typical melody in Musicat's domain.

The French folk tune in Figure 3.1 is a typical melody in Musicat's domain. The melody is relatively simple both rhythmically and melodically. The rhythm is made up of just quarter notes and half notes, and the melody is tonal and easy to sing. Note that the lyrics are certainly not part of the domain, and they are not given to Musicat. I include lyrics in the figures simply to aid the reader in remembering the melodies.

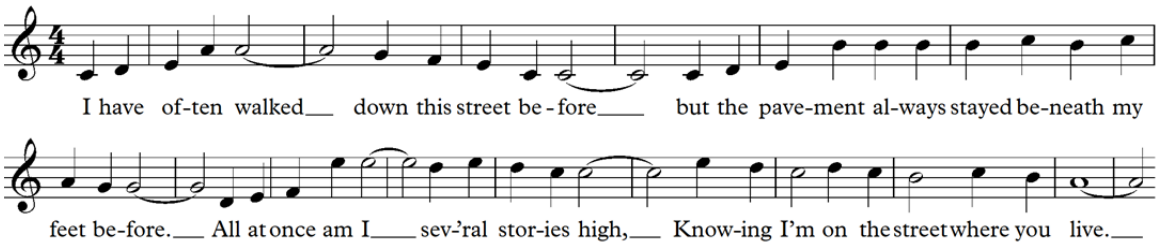

Figure 3.2: "On the Street Where You Live", a more complicated melody in the domain.

The melody from "On the Street Where You Live" (Figure 3.2), as was mentioned previously, is a more sophisticated melody in this microdomain. Only the first 16 measures of the melody are pictured here; the full melody is much longer. Note the use of ties across bar lines. There are just two beats in the first measure (a "pickup measure" containing two "pickup notes", or "upbeats"), but the rest of the measures, except for the final measure, have four beats each. Not pictured explicitly is the more complex tonal structure of the melody. Musicat "hears" melodies as if they were performed *a cappella* — no chords are given — but melodies tend to imply harmonic structures. The implied harmonies in this example are more complicated than those in Figure 1. The grouping structure of the melody is also more sophisticated. (An exercise for the reader: where do you hear groups starting and ending in this melody?)

Example Melodies and Human Listening Performances

In this section I discuss three melodies (including the two melodies above) from the point of view of human listening. Chapters 6, 7, and 9 will examine Musicat's performance on a variety of melodies including these three, but for now I focus on how I, as a human listener, hear these melodies. Some of the things I describe in this section can be modeled by Musicat, but some can't.

To be clear, these are introspective walkthroughs of my own subjective experience of listening to these melodies, and are not based on EEG, fMRI, or other laboratory methods that some readers might expect. If this sounds suspect for some reason, please realize that I am simply describing my own experience in order to introduce the challenge of music-listening, in preparation for a careful examination of Musicat's behavior to follow in later chapters. Moreover, these melodies are simple enough that I am sure there is not much variation in how listeners who are experienced in Western tonal music will hear them, at least in terms of low-level grouping structure.

Mystery Melody

For my first example, I will reproduce a melody that Steve Larson played for me in 2003 to demonstrate some aspects of listening. He played it one note at a time, and we stopped to discuss what I had heard and what I expected after each note. I will try to replicate the experience here, by adding a single note to each successive figure in the next several pages. I encourage the reader to try not to look ahead — we can't "look ahead" when listening, after all! — but instead, sing the melody to yourself and think about your own expectations. I will use the pronoun "we" in this section to bring you, the reader, along with me as we listen to this short melody. See the next figure for the first note.

Figure 3.3: Mystery Melody, Note 1.

We have heard a single note with a quite typical duration. A person with absolute pitch (also known as "perfect pitch") might have heard this as a middle C — for the rest of us, it just sounds like a note with a pitch from somewhere in the middle of a piano keyboard (in the rest of this section I will use the note names for simplicity, but we can think of "D" as simply meaning "the note above the first one"; when I use the note names I am not meaning to suggest that typical listeners know the names of the notes they hear). Do we have any expectations about what notes might come next? It certainly doesn't seem there is much information to go on. Let's listen to the next note:

Figure 3.4: Mystery Melody, Notes 1–2.

The melody is now a bit more interesting — two notes are better than one! We have heard two notes that are close together in pitch: the second note (D) is a whole step higher than the first. What do we expect to come next? In my case, I expect the next note to be another step higher still; I expect to hear an E. (See Appendix B for a pilot study investigating expectations that are formed after two notes such as these are heard.) Another very reasonable expectation is for the melody to go back down to C.

By this point, I also have heard the first note not simply as "C" but also as "the tonic note of the scale" (I wouldn't *consciously* think about the note C being the "tonic", but I believe my internal representation of the first two notes would be somehow associated with

this concept). I have implicitly heard this as a melody in the key of C — more specifically, C major. Also, I have heard the first note as a downbeat; the C sounds stressed to me and the D sounds unstressed. Even when I have a computer play the two notes for me with equal volume, I hear the C as stressed (this entire melody should be performed with equal volume and articulation for every note; any stresses should not be present, but rather should be implied and ascribed to notes in the "mind's ear" of the listener).

We have default assumptions that kick in rapidly as we listen, but the melody may turn out to prove us wrong. Nothing is forcing the melody to sound like C major; it might turn out to be in the key of A minor, for instance. There is also no guarantee that the first note is a downbeat; it might be a pickup note, with a downbeat D, and perhaps the melody would continue in D minor (imagining this hypothetical way of hearing the first two notes instantly reminded me of the song "Summer Love" by Neil Diamond, from the movie *The Jazz Singer*, whose melody starts this way: the first note in the orchestra is a lowered seventh scale degree, and the second note is the tonic, accompanied by a minor tonic triad).

Figure 3.5: Mystery Melody, Notes 1–3.

The third note makes me more confident in my default assumptions and expectations. Indeed, the E that I expected came next. This certainly sounds like the start of a C-major melody, and the first note sounds like it was indeed a downbeat.

Different people will have different expectations, but I expect the melody to continue up the scale at this point with F and then G. Maybe it will even continue up to the C an

octave higher. I could also imagine that the melody will turn around and come back down to C instead of going up, but my first expectation was for an F to come next.

What will come next? Will the melody go up to F or drop back down to D? Or will it drop immediately down to C as in the song (coming up in Chapter 6) "*Frère Jacques*"?

Figure 3.6: Mystery Melody, Notes 1–4.

Another E — what a surprise! I find this second E to be confusing. It is remarkable that this extremely simple-looking 4-note pattern causes surprise and confusion, but it does, at least for me. What do we expect to come next now? The second E seems to bring a halt to the upwards motion established by the first three notes (according to Larson's theory of musical forces, we might say that the melody's upwards-pointing inertia was overcome by the force of magnetism or gravity). I find it uncomfortable to stop for long here after the repeated note, so let's listen to the next note.

Figure 3.7: Mystery Melody, Notes 1–5.

The melody has continued its upwards climb after a brief pause on E. But this is too simple a description of what the melody sounds like. At this point, many interesting things are happening (or have already happened) to me as a listener, although I will delay describing them in detail because I don't want to ruin the experience for you, the reader. For me, the

melody is still a little confusing here, but after another note all will be resolved. What do we expect to come next?

Figure 3.8: Mystery Melody, Notes 1–6.

The previous two notes, E, and F, left me feeling confused for some reason, but by the time I heard the G, everything made sense. I hear these six notes as two copies of a 3-note pattern: C–D–E followed by E–F–G. The pattern is made up of three notes that are successive steps in the C-major scale. The first instance of the pattern started on C, while the second instance started on E (which was the same note that the first pattern ended on).

Perceptually, the elements of each pattern form a group: I hear the first three notes as being grouped together, and likewise the last three notes sound like a group. Furthermore, when I say that they sound like instances of a pattern, I am really making an analogy between the two groups: C–D–E is like E–F–G.

Why was I so confused by the repeated E? Besides the E denying my expectations for the motion to continue up or down, I was surprised that even though the two E's had the same pitch, and sort of sounded like they should be grouped together, they eventually turned out to be members of *different* groups. The first E turned out to be the end of a group, while the second E was the start of a different group.

The real surprise, though was that the groups I heard had three elements each. When I heard the initial notes C–D–E, I heard the C as a downbeat, the D as a weaker beat, and the E as a stronger beat. That is, I heard them in duple meter. Specifically, I expected the fourth note to the last beat of a measure in 4/4 time, and I expected another strong

downbeat to occur at the point in the melody where we heard the F, as in the following figure, in which implied downbeats are marked with accent marks, and a smaller implied stress, typical in 4/4 meter, is indicated with a stress mark over the first E:

Figure 3.9: Mystery Melody, Notes 1–5, heard in 4/4.

However, the F doesn't sound like a strong downbeat at all! Instead, the second E sounds like a downbeat, like so:

Figure 3.10: Mystery Melody, Notes 1–5, heard in 3/4.

In short, I was surprised because I started off hearing the melody in 4/4, but over the course of two or three notes (from the first E through the G) I had to revise my hearing to hear everything in triple meter (3/4) instead (this shift in metric interpretation is similar to the one I described in Chapter 1 for the opening of the song "Ants Marching").

The mental processes involved in listening to these six notes are representative of those that Musicat is intended to model. To hear these notes as two similar 3-note groups involves grouping, analogy-making, expectation-generation, and revision of previous perceptions.

After hearing the first six notes, what do we expect to come next? I recommend trying to sing the entire melody from the start and to improvise a conclusion. Will the melody fulfill our expectations?

Figure 3.11: Mystery Melody, Notes 1–7.

For me, this was just what I expected. Because I heard two 3-note groups, one starting on C and another starting on E, I expected the pattern to continue with a group starting on G and continuing upwards.

Figure 3.12: Mystery Melody, Notes 1–8.

The note A was just what I expected; nothing is surprising here.

Figure 3.13: Mystery Melody, Notes 1–9.

The B was also expected. At this point, I have heard three copies of the 3-note pattern, and furthermore, I have heard the music in triple meter, as in the following figure:

Figure 3.14: Mystery Melody, Notes 1–9, heard in 3/4.

Notice that at this point something subtle occurs in our expectations. It would be ridiculous to expect another copy of the pattern starting on B as follows:

Figure 3.15: An unrealistic expectation for the next three notes.

We don't expect the pattern to repeat, starting on B, despite the seemingly mathematical precision of the first nine notes. However, a computer program (such as Seqsee, for instance) might generate such an expectation if it "heard" the pattern in the following way, with pitches replaced by numbers with no musical meaning (expectation shown in parentheses):

1–2–3 3–4–5 5–6–7 (7–8–9)

Instead of this expectation, we hear the B as a leading tone that "wants" to resolve to the tonic (the C a half step above). Moreover, the downbeat notes C–E–G outline a C-major triad, and we expect motion through a C-major triad (*i.e.*, an arpeggiation of the C-major chord) to lead directly to the C on the next downbeat, even though it breaks the surface-level pattern established by the first nine notes.

Figure 3.16: Mystery Melody, Notes 1–10 (complete).

Indeed, this short melody continues (and concludes) as expected, on the note C. In my 3/4 hearing, the melody would be notated as follows (the accent marks, as before, simply indicate implied accents):

Figure 3.17: Opening theme of Bach's Organ Prelude in C major, BWV 547.

It turns out that our mystery melody is the start of a Bach prelude, from the Organ Prelude and Fugue in C major, BWV 547. Henceforth I will refer to this melody by a shorthand name, "Triplet Scale", to emphasize its 3/4 nature. This melody is of the sort that early versions of Musicat were designed to listen to. However, Musicat evolved to focus on larger-scale musical structures and expectations, so the note-by-note description of my experience of listening to this melody is not indicative of what the latest version of Musicat does (but see Chapter 9 to learn more about the early versions of Musicat and to see it running on this melody).

Although this melody is very short and simple, there are complex cognitive phenomena involved in forming default expectations, in making groups and analogies, and in switching from a duple- to a triple-meter way of hearing somewhere between notes 4 and 6. These same sorts of complications arise in most music, so I find this melody to be quite inspiring and even provocative: what a challenge it poses for music cognition research, despite its simplicity! A program that could be just as confused as I was after hearing the note F would be well on its way to modeling human music-listening.

Sur le pont d'Avignon

Whereas for the "Triplet Scale" I described my listening process one note at a time, for "Sur le pont d'Avignon" I will proceed by roughly a measure at a time to speed things along.

Figure 3.18: "Sur le pont d'Avignon", measure 1.

This melody seems even simpler, so far, than the previous melody. Instead of C–D–E, we have heard C–C–C. The repetition is almost too obvious to mention, but the presence of exact repetition of pitch is quite salient; we will automatically notice that these notes are linked together. There is one difference between the notes, however: the third note is a half note. Of course, when we first hear the third C, it doesn't sound any different from the others, but eventually we realize that it is sustained for a longer time than each of the first two notes. The longer duration of the third C makes it likely for us to hear a grouping boundary after it. Therefore, as a result of all the pitches being the same and the long duration after the third note, I hear these three notes as a group.

Figure 3.19: "Sur le pont d'Avignon", measures 1–2.

The second measure has the same rhythm as measure 1, and it also consists of three notes at a single pitch. Even after I have heard just the first or second D, I notice the similarity to the first measure and form an analogy between them, so I expect the third note to be the half note D that does indeed occur.

At this point, I have heard an *analogy* between the two measures, and also heard each measure as forming a *group* consisting of three notes. I'm not sure if I have strong expectations for measure three, but if anything I would expect the pattern to continue on an

E. The second measure sounds like the first measure transposed up by one step; it is easy to imagine this pattern continuing up another step.

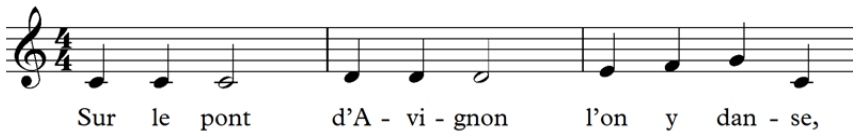

Figure 3.20: "Sur le pont d'Avignon", measures 1–3.

The third measure starts out just as we might expect, with a quarter-note E. However, it diverges on the second beat and we hear a succession of different pitches; this is quite different from the pitch repetition of each of the first two measures. The notes E–F–G go together and sound like they form a group because they are a scale fragment. The next note, C, is a bit of a surprise, because it is the first non-stepwise motion in the melody. After an ascent all the way from the initial C in measure 1 to the G in measure 3, this note sounds like a sudden fall back down to the tonic, where the melody started. Because this fall lands on the tonic, I hear it as a conclusion to the group started on the note E. Thus, I hear all four of these notes as constituting a single group, although perhaps I hear the first three notes as a subgroup inside this group.

In any case, the notes of this third measure form a group. The rhythm and the pitch contour are quite different from those of the first or second measure, however. Therefore, when I hear measure 3, I also realize that the first two measures together form a unit, but measure 3 is starting something new. The groups in measures 1–2, then, seem to form a meta-group, and perhaps a new meta-group will start at the start of measure 3.

Figure 3.21: "Sur le pont d'Avignon", measures 1–4.

Measure 4 sounds like measure 3; the similarity is obvious and jumps out, even without my consciously thinking about the precise intervals used in measure 4. Starting on the B, we hear another scale fragment going up, just as in measure 3. At some point in the measure, I notice the similarity, and perhaps I even expect the large leap down that happens after beat 3. (I can't be sure about the details, because the first time I heard this melody was many years ago and now I know it very well, but it seems that early on in measure 4, any listener will hear the similarity of that measure to the start of the measure 3.)

At the end of the measure, the leap down to G is analogous to the leap down to C in the previous measure; measure 4 is an exact transposition (down a perfect fourth) of measure 3. Additionally, even if I wasn't thinking consciously about harmonic implications in the first three measures, measure 4 clearly evokes dominant-chord harmony. The most salient notes in the measure are the B, D, and G — the notes of a G-major triad. The C sounds like a passing tone between B and D and it occurs on a weak beat (beat 2). The G is also on a weak beat, but it is salient because the melody lands on it after leaping down a perfect fifth. Applying the same reasoning to measure 3, we see that it is also easy to hear as emphasizing a C-major chord (the notes E, G, and C are salient). Hearing measure 3 with C-major harmony and measure 4 with G-major harmony makes measure 4 sound very strongly like a half-cadence, especially in conjunction with the grouping structure we have heard so far. An analogy was heard between measures 1 and 2 and they also formed a meta-group; similarly, an analogy was heard between measures 3 and 4, and they seem to constitute a second meta-group. The half-cadence in measure 4 makes this measure sound like the end of a group. Furthermore,

the cadence serves to close off other potential higher levels of groups: I hear the entire first four measures as forming a single group, which is brought to a close at this half-cadence on a G-major chord.

There is a slightly subtler higher-level analogy here. I can't be sure that a typical human listener hears it in exactly this way, but I'll mention it for completeness. Measures 1 and 2 sound like two instances of a musical idea, related by transposition. Likewise, measures 3 and 4 sound like two copies of the same thing, also related by transposition. In the first case, the transposition was a simple movement up by a whole step. In the second case, the transposition was larger: a movement down by a perfect fourth. But these two measure pairs sound analogous to me: the relationship between measures 1 and 2 is somewhat related to the relationship between measures 3 and 4, since both pairs involve transpositions.

I think that the analogy is even stronger, however, because of a stepwise connection linking measure 1 to measure 2 and a very similar connection linking measure 3 to measure 4. At the end of measure 1, the melody moves *up* by one step (this is trivial because we have been thinking of measure 2 as a transposition) as the C goes up to D, while at the end of measure 3, the melody moves *down* by one step: the C moves down to B. This is a much stronger connection than simply hearing measure 4 as a transposition: the entire measure is indeed a transposition, but at the note level there is a stepwise connection. I think of this is a meta-analogy: measures 1 and 2 are connected in a very similar way to how measures 3 and 4 are. Also, there is a nice symmetry in motion between these pairs of measures: the melody moves up to go from measure 1 to 2, but it moves down to go from measure 3 to 4. The rhythm also contributes to this symmetry: the initial ascent from C to E in measures 1–3 is quite slow, involving two half notes, but measures 3–4 consist exclusively of quarter notes that reach the high point of G but then rapidly come back down to C — and then fall past C

all the way to a lower G. Metaphorically, it seems almost as if it took a lot of work for the melody to "climb" up to the mountain peak of G, but this climb was followed by an effortless fall back down to the starting point, crashing through the C to bounce off B, only to fall again all the way down to the low G, which acts as a safety net that finally stops the fall. (Of course, it is hard to articulate or even to know exactly how we hear this pattern, but the distinctive shape of measures 3 and 4 surely evokes some unique feeling that contrasts with how we hear measures 1 and 2, and I suspect the rising and falling motion in the shape causes most listeners to hear something vaguely akin to what I described.)

Let's keep listening — the next figure gives the next three measures:

Figure 3.22: "Sur le pont d'Avignon", measures 5–7.

After just a few notes in measure 5, it sounds like the melody is repeating from the start. We hear the three C notes with the quarter–quarter–half rhythm and it's easy to remember that this is the same thing we heard at the very start of the melody. As the melody continues into measures 6 and 7, there is no surprise; we hear an exact repetition.

The same groups form in these measures as in the first ones, although it is perhaps a little easier to make sense of the melody now, the second time around (it was easy enough the first time, but now we really know what to expect). In a trivial way, we form analogies between these measures and the first measures: measure 5 is exactly like measure 1, measure 6 is the same as measure 2, and measure 7 is the same as measure 3.

There is one slight difference, though. We have already heard the first four measures as a group ending on a half-cadence. The current group of measures sounds like a repetition, but familiarity with tonal music suggests that we might reach a *full* cadence soon, even

though so far we have exact repetition. In cases such as this one, I expect a tonic chord to arrive soon, perhaps at the start of measure 8, so I am waiting to hear the note C. However, measure 7 already ends with a C. It falls on a weak beat, however, so I anticipate a stronger C (perhaps reached via stepwise motion, as happens quite often in strong cadences).

Figure 3.23: "Sur le pont d'Avignon", complete.

In the final measure, this expectation is satisfied nicely. The first two notes are still members of the dominant chord, D and B (over an implied chord root G), but the third note is a solid half-note C, which was indeed reached by a step and which sounds much more conclusive than the C in measure 7.

After hearing this final C, the group structure is quite clear: I hear the final two measures as forming a group (just as I did measures 5–6), at a higher level I hear the final four measures, 5–8, as a group, and finally I hear the entire eight measures of the melody as one big group. Because of the near-repetition of the first four measures (yielding the final four measures), the implied dominant harmony in measure 4 followed by the tonic harmony and melodic closure in measure 8, the first four measures sound like an antecedent phrase and the last four sound like its consequent phrase. A strong analogy links measures 1–4 with measures 5–8.

All in all, nothing was particularly surprising here, aside from the elegant meta-analogy that we might hear linking the pair of measures 1–2 with the pair of measures 3–4. (Notice that although measure 8 was not a transposition of measure 7, it was still linked to measure 7 by step, so it was not completely unlike measure 4 and its connection to measure 3.) This little melody seems rather easy to understand, and it induces some straightforward expectations along the way that mostly come true. I am quite fond of this little tune, especially because of the way measures 3 and 4 are linked together and because of the way the nonstop quarter notes in those measures, with their quick fall from a high G all the way to a lower G, contrast with the slow climb in measures 1–3.

On the Street Where You Live

The melody of "On the Street Where You Live" is much more sophisticated than the previous one. Listening to this melody is therefore quite a bit more complex, so I will focus on just a handful of the most interesting issues that arise.

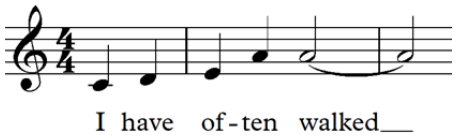

Figure 3.24: "On the Street Where You Live", first 5 notes.

This melody starts out with the same three notes of the Triplet Scale melody, except that the first two notes are pickup notes: the E is on a downbeat instead of the C. Once the E has been heard, the first three notes will likely be heard as a cohesive group — an ascending-scale fragment in C major — just as they did in the melody earlier. At this point I might might expect the scale to continue upwards to G or to descend back to C.

Therefore the next note, A, may come as a small surprise: instead of a continuation of the scale, there is a leap upwards by the interval of a fourth. This gap in pitch — the wide

interval — causes me to hear the A as rather separate from the first three notes. In a way, I hear the A as potentially forming the start of another group of notes — it naturally forms a group with the long A that immediately follows it, because they have the same pitch. However, I also hear the first A as a continuation of the initial motion, partly it feels unnatural to think of a group ending immediately after hearing the first downbeat of the melody. Therefore, for me the most natural hearing involves hearing the first five notes, C–D–E–A–A, as one a single group. The group certainly ends after the second A, however, because of its quite long relative duration.

Also worth noting is the tension that I hear in the melody at this point. The upwards leap, coming right on the heels of a small climb from C to E, makes it feel like the melody is suspended up in the air, waiting to come back down. I feel tension as a listener at this point, because the melody is at an unstable place. Not only is the pitch suddenly high, but also the note A, in C major, is tonally less stable than the strong C and E notes that came before.

Figure 3.25: "On the Street Where You Live", first phrase.

Sure enough, in the next bit of music, the A resolves downwards, relieving the tension. The next notes, G–F–E, form a descending group. Because they are moving with stepwise motion, I hear them as notes in a scale. I also notice (perhaps only semi-consciously) a similarity to the initial C–D–E scale fragment. This analogy, though simple-seeming, is subtle in several ways. The C–D–E group is ascending, whereas G–F–E is descending. The pattern of absolute distances in terms of half-steps is different: C–D–E is made up of the interval pattern "+2, +2" (ascending by two half-steps between each pair) while G–F–E

consists of the pattern "–2, –1" (the interval between F and E is smaller than the G–F interval). Nonetheless, we effortlessly hear these as similar patterns, because we can hear the notes in terms of their position within the major scale: C–D–E is heard as 1–2–3 and G–F–E as 5–4–3, making interval patterns of "+1, +1", and "–1, –1". Moreover, both of these segments involve three notes, which start "stably" (on C or G), then pass through a less stable note (D or F), and both end on the relatively stable note E.

At the end of measure 3, the note C is heard again; the melody leaps down from E to C, much as it leaped up from E to A earlier. Even though this leap involves a different interval and direction (down a third instead of up a fourth), I still hear it as a leap in the same direction as the previous notes were moving. The C is then repeated (just as the A was) and I hear this second set of five notes as another group; the C, along with all the downward motion, has relieved the tension of the high A earlier, and the melody has reached a point of repose and closure.

This 5-note group, G–F–E–C–C, is analogous to the first one, in that they have very similar contour, as long as we substitute downwards motion for upwards motion. In addition, their rhythmic structure is identical: four quarter-notes followed by a whole note. Table 1 summarizes the similarities.

C–D–E (scale degrees 1–2–3)	G–F–E (scale degrees 5–4–3)
Ascending scale	Descending scale
Leap up (E → A)	Leap down (E → C)
Stable–unstable–stable	Stable–unstable–stable
A–A repetition	C–C repetition
	

Table 1: Similarities between the two halves of the first phrase of "On the Street Where You Live".

A higher level of structure also emerges here. I hear a meta-group comprised of the first (ascending) 5-note group and the second (descending) 5 note group. This is quite similar

to the structure of the first four measures of "Sur le pont d'Avignon": two measures ascend, and then two measures descend forming a symmetric large structure. In the present example, though, the analogy between the two halves of the meta-group is of a quite different nature from the analogy described for the previous melody.

So far, we have heard four measures worth of the melody. I'll call this group a "phrase" for the purposes of this discussion. Let's continue listening. I will only show successive phrases in the following figures because this melody is longer than the previous ones; a figure later in this section shows the melody as a whole.

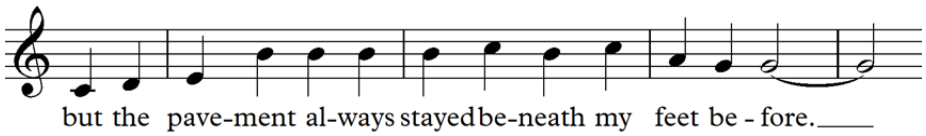

Figure 3.26: "On the Street Where You Live", second phrase.

Since measure 5 begins with the same three notes as measure 1 (C–D–E), this might trigger a naïve expectation that the melody will continue in the same way as measures 1–4 did (after all, in "Sur le pont d'Avignon", the first three measures did indeed repeat exactly). The note after the E, however defies this expectation. That note, B, is indeed the result of an upwards leap from E, as expected, but the leap is larger this time: E goes up to B, not A. This note, B, sounds even more unstable than the A did, because B is the leading tone in C major.

The B is repeated, so that the first five notes heard are C–D–E–B–B. At this point, I expect the B to be sustained for four beats, just as the A was at the start of the melody. I am expecting to hear these five notes, ending with the long-duration B, as a group analogous the the C–D–E–A–A group. However, another surprise is in store: the B is only a quarter note long, and it is followed by another B. This is not too hard to understand, though: I expected the B to last for four beats, so I can hear this "extra" B simply as part of the expected long-

duration B, with an added rearticulation. Several identical notes in a row aren't terribly different from one long note, aside from short breaks in the sound and rhythmic emphasis occurring at extra points.

As I continue to listen, I hear yet another B — so far, there have been four B's in a row (instead of the expected two). But again, it sounds something like what I expected. Instead of a whole-note's duration worth of the B (after the initial quarter note B) I've heard three quarter-note B's. Maybe there is just one more left to go to fill out the duration of the whole note.

The next note, however, is a surprise: instead of another B, the melody ascends a half step more to the C an octave higher than where we started! (In the remainder of this section I denote this higher C with a prime symbol: C′.) Perhaps it is unsurprising that the leading tone resolved upward to C′, but it increases the feeling of the melody having reached quite a high point, since after an initial ascent by step from C to E, followed by a large leap up from E to B, the melody is pushing even higher. Granted, the C′ is the tonic and has some stability, but there is still a lot of tension here due to the height of the melody. (And although lyrics are not part of Musicat's domain, notice the text-painting here: the melody has soared up to C′ just as the lyrics describe how "the pavement always stayed beneath my feet before". The "pavement" of the lower C is an entire octave below the current lofty height of the melody.)

A particularly interesting phenomenon here is the way in which this phrase is heard in relation to the first phrase. It is clear that I hear an analogy between the first phrase and this one, and the analogy informs my expectations about what is about to happen (such as my expectation for the B to be a long note lasting four beats). But even more interestingly, the current push up to C′ starts to strain the analogy: the disparities between the two phrases are growing as the melody continues, yet there is still an obvious relationship between them.

The feeling I have at this point is similar to what I felt in the Triplet Scale, when I had to reinterpret what I heard as being in triple meter instead of 4/4. It is a bit confusing and the melody is not proceeding in the way I expected.

It's not only the analogy that is being strained or stretched here: the group starting with C–D–E–B–B is also being expanded to the right. I had expected the long note B to be the end of the group, but now "extra" notes keep coming, and there is no obvious breakpoint, so I mentally keep adding these notes to the group. The group so far contains C–D–E–B–B–B–B–C′. When will the group end?

The next note is another B (we are now on the second syllable of the word "beneath"). By this point, the expected four-beat B should have ended, but the melody is pushing along, still in B territory. Maybe it's finally time for the tension to be resolved and for the melody to descend as it did in the first phrase.

But we have yet another surprise: it's right back up to C′ yet again! I feel a bit uneasy here as a listener: the succession of B's and C's is getting overwhelming as I yearn even more for the tension to be released. There is tension from the leading tone B's and from the ever-expanding group in my mind (now it has been extended to C–D–E–B–B–B–B–C′–B–C′), and I'm still waiting for the expected descent to happen.

Finally, the melody leaps down to A, bypassing the B. It's a welcome relief. The grouping structure I hear here is a bit unclear; perhaps I hear a new group starting on the A, or perhaps it continues to expand the ever-growing group that contains all the B's and C's.

After the A, the melody continues right on to G on the word "before", and the G repeats. Suddenly, it sounds like we've reached the end of the second phrase, because the G is a long note, just like the C in the first phrase was. The very end of this second phrase sounds like the end of the first phrase to me, even though the three-note descending scale was

skipped this time. Also, the phrase doesn't end on the "pavement" of C, but rather on the alternate, fairly stable, dominant base of G. Because G is the second most stable note in the C scale, it is easy to make the analogy: both phrases end with a descent to a stable note that is repeated.

Here I will take a moment to discuss the grouping structure of the melody so far. At this point, I have heard the music as one large 8-bar meta-group consisting of two 4-bar meta-groups, with numerous relationships linking the 4-bar groups to each other. The first 4-bar meta-group is heard as being made up of 2-bar groups that are related to each other, and each of these may be further divided into yet smaller pieces just two or three notes long. Note how the perceived structure is built up of ever-larger structures, where each larger structure is roughly twice the size of its highest-level components. The tendency of simple melodies to be constructed in this way, with groups often having lengths that are a power of two (2, 4, 8, 16, *etc*.) is known to the Musicat program and provides a helpful heuristic for how to hear groups.

Figure 3.27: "On the Street Where You Live", third phrase.

The doubling in size of structural components described above continues apace, starting with measure 9. Again there are three scale notes starting the measure, but they are D–E–F this time instead of C–D–E. Despite this small change in pitch, I effortlessly hear an analogy linking these three notes to the three at the start of the first phrase, or maybe to the notes at the start of the second phrase (or both). These notes are followed "as usual" by a leap, but this time it is an even larger leap (a major seventh) all the way to E′, which is much higher than the "pavement". (More text painting happens here: the large leap up to E′

perfectly echoes the idea that "all at once am I several stories high". It's a major 10th high, to be exact.) Next comes the analogy-based expected repetition of this note, and, as in the first phrase, it repeats just once, with a whole-note duration. I hear an analogy between this part of the melody and the group at the start of the first phrase. The melody is suspended at the dizzying height of E′, and it sounds like this is the conclusion of another group: D–E–F–E′–E′.

Now I expect the melody to descend again to the "pavement" C, but instead it moves down just a bit to C′, an octave above the C. The end of the phrase is still recognizable and reminds me of the end of the first phrase, despite many differences. Both phrases end on a repeated tonic note (C) and the last note of each phrase has a long duration. So, I hear a strong analogy between this phrase and the first one.

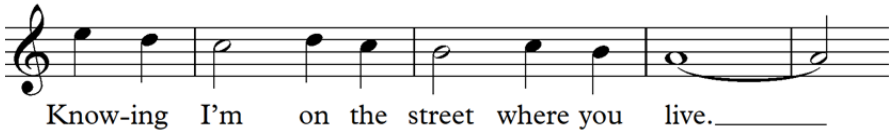

Figure 3.28: "On the Street Where You Live", fourth phrase.

In the final phrase, a new rhythmic pattern begins. The first three phrases were quite similar to each other, except that in each successive phrase the high note was higher, and the second phrase had an interesting bunch of "extra" B notes. So where did this fourth phrase come from? It is very different from the first three phrases, but it doesn't sound out of place.

Consider the first three notes of this phrase, E′–D′–C′. They seem to fit in well because they remind me of the notes that just came before. The same notes occur close to the end of the previous phrase, on the lyric "(sev)-ral stories high". The rhythm is different in the third note, however: the C′ is a half note. This change in duration causes me to hear a group boundary as occurring after the C′; the notes E′–D′–C′ ("Knowing I'm") form a small group.

At this point, several small groups and analogies form in succession. The next notes, D′–C′–B ("on the street") are a simple transposition of the previous three notes. The pattern continues with the next three notes, C′–B–A ("where you live"). This constitutes a musical *sequence*: The three-note seed, E′–D′–C′, appeared three times, each time shifted down by a scale degree, and then the final A had a longer duration to finish the phrase. Naturally, I hear this entire sequence as a group, and there are straightforward analogies among all three components of the sequence.

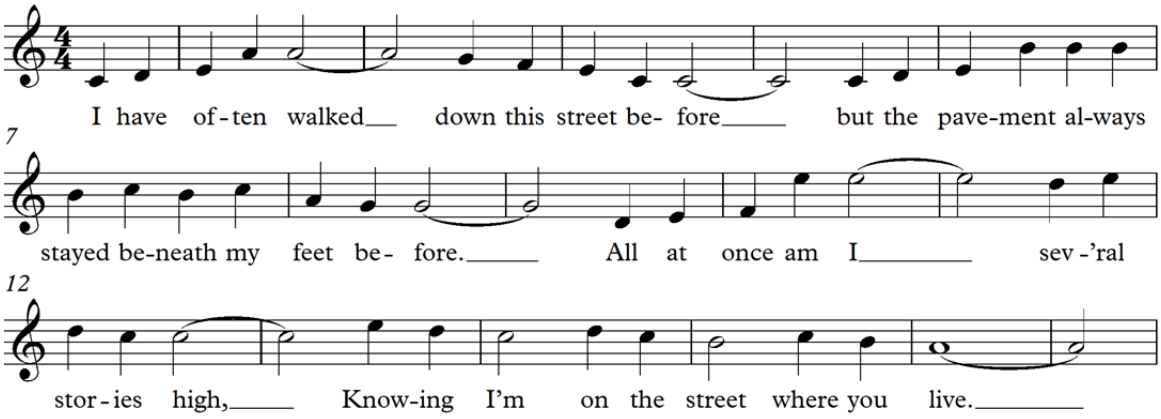

Figure 3.29: "On the Street Where You Live", phrases 1–4.

The melody isn't complete yet: this excerpt ends on A, and eventually, of course, it will return all the way to the tonic and make a much more final conclusion to the song. The sequence ending on A yields only a temporary feeling of closure, but it does relieve much of the tension that was created by the high E′ in the third phrase (as well as by the gradual climb up through A and B in the first two phrases).

There is much more to say about melody of the "On the Street Where You Live", but this description should illuminate some of the complexity of the listening process, and in turn, hint at aspects of listening that I intended to model with Musicat.

A Sneak Preview of Musicat

Chapters 5, 6, and 7 describe Musicat's listening in detail, for many songs, just as I have described my own listening in this chapter. But this seems an opportune moment to present a preview of what the program does. The final figure in this chapter shows the mental structures that Musicat formed after listening to "Sur le pont d'Avignon" on one particular run. The ellipses represent groups and meta-groups, while the thick parabolic arcs below the boxes represent analogies linking groups and linking meta-groups. (The thin arcs will be explained in Chapter 5.) Several (but not all) of the structures and connections I described above in my own listening experience were also "heard" by Musicat, and can be found in this diagram.

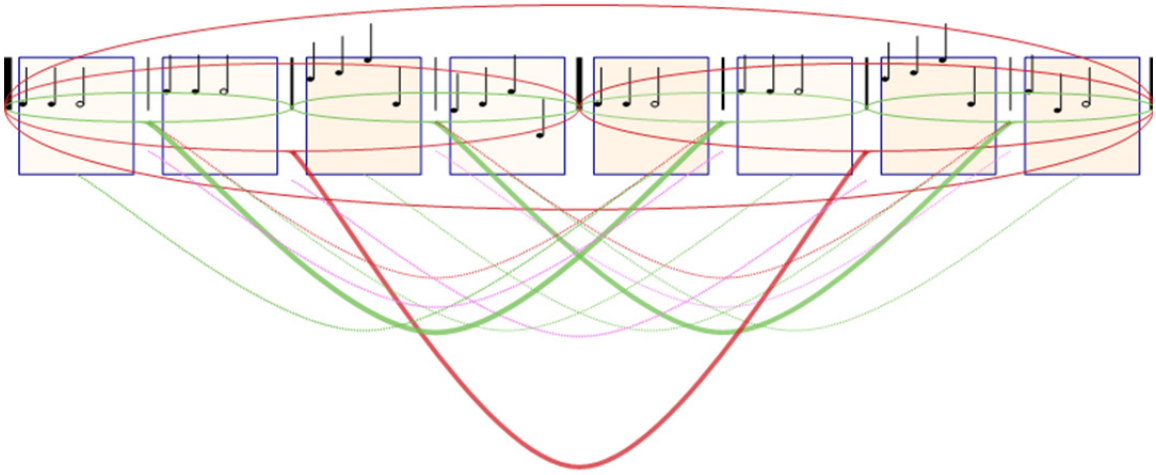

Figure 3.30: One of Musicat's listening performances for "Sur le pont d'Avignon".