CHAPTER SEVEN

Musicat Listens to Complex Melodies

YOUNGER THAN SPRINGTIME (RODGERS AND HAMMERSTEIN)



Figure 7.1: Younger than Springtime (from the musical South Pacific), 32-measure excerpt.

"Younger than Springtime", with music by Richard Rodgers and lyrics by Oscar Hammerstein II, is our first example of a "Complex Melody". In contrast to the short melodies of the previous two chapters, this is a typical song-length melody: 32 measures of music are included here (I omitted an introduction section that is part of the complete song). The length poses some challenges to Musicat (see the Chapter 10 for more discussion of this point), so in this section I have the program analyze a few individual shorter segments: measures 1–8, then 1–16, and finally measures 17–24. Measures 25–32 are extremely similar to 1–8 so I left them out of these runs.

It's important to remember that the Musicat does not "hear" the lyrics — it is given only the notes. These lyrics provide many hints as to the grouping structure of the piece and suggest numerous analogies, but the program is working without these extra clues. For instance, the words "are you" at the end of each of the two phrase in "Younger than springtime are you / softer than starlight are you" help increase the feeling of parallelism between measures 1–2 and 3–4 (although the rhythmic and melodic similarity by itself is plenty enough for anyone to hear the analogy between the measures, without the lyrics). In the final 8 bars, moreover, the clever switch from "are you" to "am I" helps establish that this is the *final* section of the melody, wrapping up this 32-measure structure. In addition, the reprise of the words "younger than springtime" in measure 25 instantly indicates that this is the start of another section, related to the first one. It is thus almost trivial for a person who has access to the lyrics to hear that measure 25 is the start of a group. This final section is a minor variant, **A'**, of the first part of the melody, **A** (the large-scale structure can be heard as **AABA'**). Even simple grammatical features such as the end of the sentence at the start of measure 16 help people in understanding the grouping structure of the music.

Other textual features of the music might not play a large role in the grouping or analogy structures formed by Musicat, but still add greatly to the human listening experience. For example, the alliterations present in "softer / starlight", "warmer / winds", and "June / gentle" not only add to the poetry of the text, but also increase the cohesiveness of the measures and groups that these word-pairs are part of. Similarly, the text-painting in the melody, such as with the lyrics "heaven and earth" ("heaven" is symbolically set to a high note in the melody, while "earth" is on a low note) adds much to the listening experience that is not expressible in terms of the structures Musicat creates.

In any case, bear in mind that Musicat has none of the information conveyed by the text.

Younger than Springtime, 8-Measure Excerpt



Figure 7.2: Younger than Springtime, 8 measures.

I started by running Musicat once on the very first eight measures of the melody:



Figure 7.3: Younger than Springtime, 8 measures.

The first four measures have formed a nice, strong pair of groups enclosed by the meta-group (1-4). The next four measures also have the same structure, but group (7-8) and the meta-group (5-8) are very weak and hard to see in the figure. A strong analogy, $(1-2)\leftrightarrow(3-4)$, has formed. This was expected: (3-4) is an exact transposition of (1-2), one step down. A second analogy, $(3-4)\leftrightarrow(5-6)$, is weaker, and this also makes sense: measure 3 starts just like measure 5, but measure 5 continues up where measure 3 (like measure 1) leaped down to a tied note. Measures 5–6 constitute a development of the previous 2-measure pattern, and they drive forward, pushing with a quicker rhythm to reach higher notes. Thus, the weaker analogy makes sense: the previously detected pattern was altered significantly in measure 6, so this analogy should be weaker. Measures 7–8 are not associated with an analogy, which makes sense: they are quite different from anything that came before.

The analogy structure discovered by the program seems to represent a quite cogent hearing of this passage. The grouping in the final four measures, however, is much weaker than expected. Measures 5–6 seem to set up a momentum (think of Larson's theory of musical forces) that drives towards the cadence at the end of measure 8. This suggests that we hear (5–8) as a very strong group. Similarly, the $F\#^7$ at the end of measure 7 is so unstable that it seems to drive the music forward and should help to solidify group (7–8).

In the next section we see how Musicat does with a longer excerpt made up of the first 16 measures of the melody (the present 8 measures form the first half of the longer excerpt).



Younger than Springtime, 16-measure Excerpt

Figure 7.4: Younger than Springtime (16-measure excerpt).

The notes in measures 9–15 are identical to those in measures 1–7. Measure 16, however, is different from measure 8: the phrase reaches a conclusion on beat 1 in measure 16, and then the final 3 beats of the phrase are really pickup notes to the next section

⁷ Accidentals such as this "#" are not displayed in Musicat's screenshots, but the program does indeed "hear" them; this is just a simplification in the display.

(although for the runs that follow, the program is given these notes as the final notes of these 16 measures, and it will unfortunately try to include these notes in groups and analogies even though they logically belong to measure 17 as pickup notes). But aside from this small difference, the two halves are very similar. Will Musicat hear this excerpt that way?



Figure 7.5: Younger than Springtime (16 measures, run 1), after measure 9.

This figure shows the program mid-run, allowing us to see what has happened during the first eight measures before proceeding. Turning to groups first, we see that their overall structure is similar to that of the previous run, although in this case group (5-8), which was weak in the previous run, is so weak here that it cannot be seen in the figure. Group (3-4) is also very weak here, as is (7-8), just as in the previous run. However, in this run an 8-measure meta-group, encompassing most of the melody heard so far, did form (although it is also quite weak). The entire group structure has induced an expectation for these 8 measures to repeat.

Moving our attention to the analogies, we see that only one, $(1-2)\leftrightarrow(5-6)$, has formed so far. This stands in contrast to the previous run in which we saw the nice chain of analogies $(1-2)\leftrightarrow(3-4)\leftrightarrow(5-6)$. At this point in this run, the middle group, (3-4), has been left out of any analogies, probably because its strength was low.



Figure 7.6: Younger than Springtime (16 measures, run 1) after measure 15.

Moving forward several measures into the second half of the melody, we see that some long-distance analogies have formed between groups in the first half and groups in the second half. Recall that this second half (measures 9-16) is nearly identical to the first half (measures 1-8), so we might expect a mapping of every structure in the first half to its copy in the second half. At this point in the run, groups (1-2) and (3-4) in the first half have indeed been mapped onto their corresponding groups, (9-10) and (11-12), in the second half (green analogies in the figure). In addition, a most promising larger-scale (red) analogy has formed, involving the parent groups of those groups: $(1-4)\leftrightarrow(9-12)$. But notice that since group (3-4) was perceived as a weak group, the analogy involving it, $(3-4)\leftrightarrow(11-12)$, is also weak. This is reminiscent of the situation in run 2 of the Sicilienne, earlier, in which an early weak group resulted in a weak larger-scale analogy.

At this point in the run, some groups that we might expect are missing: the large group (1-8) from the previous run has not formed, nor have the smaller groups (5-8) or (13-14). With this latter group missing, the expected analogy involving it, $(5-6) \leftrightarrow (13-14)$, cannot form (even though the two groups involved are identical). So the program is still missing what seems quite obvious to a human listener: in the second half it is in the process of listening to what is, so far (up to measure 15), an exact copy of what it heard in the first 8 measures. The program did, however, make that big red analogy, thus showing its recognition of similarity of the first 4 measures of the two sections. Unfortunately, it has not made the leap of extending this idea forward and hearing measures (13-16) as a continuation of this repetition of the first half.

There is some hope for a group and then an analogy involving measures 13–14 to form soon, however. Strong blue measure links have formed involving each of these measures. Measure 14 has a strong link back to measure 6, just as expected. Measure 13, however, has a strong link back to measure 1, not the expected measure 5. The link makes sense in that the measures are indeed similar, but it is also, unfortunately, another clue that Musicat has not grasped the concept that this half of the melody is just like the first half.



Figure 7.7: Younger than Springtime (16 measures, run 1), end of processing.

At the end of the run, in a disappointing development, the promising large analogy $(1-4)\leftrightarrow(9-12)$ has disappeared: because of the weakness of the sub-analogy $(3-4)\leftrightarrow(11-12)$, the larger analogy did not have enough support. Indeed, that smaller analogy was also destroyed. Curiously, the program created another analogy, $(3-4)\leftrightarrow(9-10)$ instead. Notice how now there are two different red analogies connected to group (9-10). This looks silly to a human observer, since *two* different groups, (1-2) and (3-4), are now mapped onto (9-10), while group (11-12) has been left out in the cold. It seems obvious from the picture

that the analogy mapping (3-4) onto (9-10) is extraneous: it looks as if the analogy coming from (3-4) missed its mark on its way to (11-12), landing two measures early on (9-10).

One bit of positive news here is that Musicat did eventually make the analogy $(5-6)\leftrightarrow(13-14)$ that was missing earlier in the run; the link between measures 5 and 13 was discovered, even though measure 13 had previously been linked to measure 1 instead. This analogy, along with $(1-2)\leftrightarrow(9-10)$, is quite strong. If only the analogy $(3-4)\leftrightarrow(11-12)$ and the larger analogy $(1-4)\leftrightarrow(9-12)$ had persisted through the run! In that case we would probably be justified in claiming that the program had noticed the correspondence between the two halves of the melody. But this failure is important because it points out that Musicat can make good medium-size analogies but still remain oblivious to larger-scale patterns. If Musicat were capable of noticing something like "measures 9–11 are the same as measures 1–3... maybe we're in the middle of a repetition", then building the rest of the correspondences in the second half of this melody would have been trivial.



Figure 7.8: Younger than Springtime (16 measures, run 2).

I ran Musicat again on the 16-measure excerpt. In this second run, as in the previous run, it misses the essential first-half \leftrightarrow second-half correspondence, and makes another surprising pair of analogies, where two early groups are both mapped onto (11–12). And alas, a very bad group was present at the end of this run as well: not only does (14–15) straddle a very thick bar line, but it also defies the straightforward pattern of grouping every two measures that was established much earlier in the melody. The program has identified some problems in the last 4 measures (see the very low red happiness rectangles associated

with measures 13 and 16), but unfortunately it was unable to resolve its confusion about these measures by the end of the run.



Figure 7.9: Younger than Springtime (16 measures, run 3).

A final run on this 16-measure excerpt exhibits some similar behavior. In this run, the grouping structure was mostly good, but group (3–4) was so weak as to be invisible in the diagram, leading to missing analogies. Some good analogies were formed — indeed, all the analogies present are very reasonable and link identical or obviously-similar groups. The

long-distance analogy $(1-2)\leftrightarrow(11-12)$ is a bit puzzling, however: why didn't the program simply make the mapping $(1-2)\leftrightarrow(9-10)$? Or $(1-2)\leftrightarrow(3-4)$? This slightly-off analogy, $(1-2)\leftrightarrow(11-12)$, might serve as a metaphor for the essential problem Musicat is having with this melody: though it is making analogies, it is failing to see the larger-scale structure. The fundamental problem here is the lack of a top-down pressure to simply map measures 1-8 onto 9-16. If this directive were guiding the lower-level processing, it would be simple for the program to make all the smaller analogies such as $(1-2)\leftrightarrow(11-12)$, and it would also help it find the good grouping structure in the second half (providing the first half was heard in a reasonable way).





Figure 7.10: Younger than Springtime, 8-measure bridge.

I ran Musicat separately on the 8-measure "bridge" section that follows the first 16 measures. I have renumbered these measures here (from the original 17–24 to the new 1–8) for readers' convenience. The program was given measures 1–8 from the figure above, but not the three pickup notes shown in parentheses, since they belong to the earlier 16 measures. Notice that this melody is a case where Musicat would benefit from a more flexible way of grouping that does not limit groups to starting and ending at the same metric position in every measure. The natural grouping structure here would include the 3 pickup

notes before measure 1, and likewise a later group would include the three pickup notes before measure 5. But the final group in this excerpt would need to end at the very end of measure 8 (measure 24 in the context of the entire melody); there are no pickup notes in sight in this measure, and the very next group would start on the downbeat of the original measure 25. The excerpt above, then, should not be exactly 8 measures long, but I was forced to extract exactly 8 measures to give to the program, since it is limited to making groups that are comprised of an integral number of measures. The excerpt given to the program *should* be the 8³/₄ measures shown in the figure, but that is not possible in the current version of the program.



Figure 7.11: Younger than Springtime (middle 8 measures).

Despite the awkwardness of leaving off the initial 3 pickup notes, thereby making the quarter notes in measure 4 seem out of place, the program has created an analogy between the two halves of this melody. Just as the 16-measure excerpt was made of two nearly identical halves, so is this 8-measure excerpt composed of two nearly identical halves. Notice that, unlike in previous examples, the program made a large-scale analogy, $(1-4)\leftrightarrow(5-8)$, without having made any explicit sub-analogies, such as $(1-2)\leftrightarrow(5-6)$, which we might have expected because it would have linked two identical measure-pairs. However, even though "official" sub-analogies were not discovered, the program did indeed find similarities between groups, and, to be sure, finding similarity is indeed analogy-making — I simply use the word "analogy" in a very specific sense when discussing Musicat's diagrams. One similarity that Musicat saw, for example, was the contour relationship between (1-2) and (5-6) (they have exactly the same notes, and thus exactly the same contour). It also discovered a rhythmic relationship between (3-4) and (7-8), as well as a non-rhythmic relationship of some sort (unfortunately, the type is not indicated in the screenshot), which is surprising because it is between the two-measure group (3-4) and the four-measure group (5–8). Probably this latter relationship is a tonal relationship indicating the progression from the note D in measure 4 to the dominant, G, in measure 8. In any case, this collection of three relationships, rather than any "official" sub-analogies, was used to support the large analogy $(1-4) \leftrightarrow (5-8)$.

I ran the program again on this excerpt to see if the sub-analogy that I expected, $(1-2)\leftrightarrow(5-6)$, would be created, instead of just the contour relationship that was discovered in this run.



Figure 7.12: Younger than Springtime (bridge, run 2).

Indeed, on a second run the expected analogy, $(1-2)\leftrightarrow(5-6)$, was formed, as was another analogy, $(1-2)\leftrightarrow(5-6)$. Also, the program made the large red analogy $(1-4)\leftrightarrow(5-8)$ again, although this time it has slightly different supporting structures for the mapping. (The *contour relationship* between (1-2) and (5-6) is still present, but it has been augmented with the explicit *analogy* structure $(1-2)\leftrightarrow(5-6)$.) The rhythmic relationship between (3-4) and (7-8) is again present, although the tonal relationship linking (3-4) to (5-8) was not perceived this time.

This listening performance by Musicat seems quite reasonable at first glance, but it did miss one salient feature of the melody: measures 1–3 form a very easy-to-hear *sequence*. Not only is each measure a transposition of the previous measure, one step higher, but the last note of measure 1 is the same as the first note of measure 2, and the same thing happens between measures 2 and 3. Thus, when I hear this passage, I hear measures 1–3 as a

sequence; I certainly don't hear the break between measures 2 and 3 that is implied by the program's grouping. (Incidentally, Hammerstein's lyrics are also incompatible with a group boundary between these measures: the word "invade" straddles the bar line between measures 2 and 3.) In the program's defense, however, to hear this passage as a tonal sequence requires hearing it in the key of G (those F# notes are a big clue, of course, that we have modulated temporarily). Musicat has rudimentary knowledge of tonal functions, but no concept of local modulation, so it would not be able to hear this passage as a tonal sequence unless, perhaps, the G's in measures 2–3 were replaced by G#'s.





Figure 7.13: On the Street Where You Live (from the musical My Fair Lady).

For our next Complex Melody we consider a short excerpt from the song "On the Street Where You Live", with music by Frederick Loewe and lyrics by Alan Jay Lerner. This melody greatly influenced the development of Musicat. Not only is it a favorite melody of ours, but also it exhibits many of the features of melody in general that we hoped that the program would eventually be able to "hear". In its present state, Musicat certainly misses out on many of the interesting details in this melody, but let's nonetheless see how it fares, even though we know in advance that it will ignore some this melody's key features. Several different listening performances follow, to illustrate the various aspects of this melody that Musicat noticed at different times. But before proceeding, I again encourage the reader (as I did in Chapter 3) to stop and listen to or sing through this melody, and to think about what groups and analogies you might be forming during your *own* listening performance.

An important comment about measure numbering is in order: this melody starts on beat 3, so each group formed by Musicat will start on beat 3 of some measure and end just after beat 2 of some later measure. Measure numbering, then, requires extra care: one might think that "measure 1" refers to the two pickup notes at the start of the piece. However, as in earlier sections, measure numbers are indexed with respect to Musicat's potential grouping points. That is, when I write "measure 1" it will refer to the first 4 beats of the piece, "measure 2" will refer to the next 4 beats, and so forth. I made a metrically-shifted version of the melody (Figure 14) in which the measures exhibit the convention I have just described and the bar lines in this notation correspond to points where Musicat's group boundaries may occur. Refer to this figure for clarity whenever measures numbers are used in the text.



Figure 7.14: Metrically-shifted version of "On the Street Where You Live", illustrating the measurenumbering scheme used in the text.

In this melody, there is a small display problem that is apparent in the figures: the ellipses representing groups sometimes appear shifted a bit too far to the left. This problem is a result of an inaccuracy in the graphics-drawing code that shows up for longer melodies with pickup notes⁸, but remember that, despite appearances in some places in these figures, all groups start on the third beat of a measure, continue across a bar line into the next measure rectangle, and continue until the end of beat 2 at a later point in the melody.

⁸ The graphics-drawing code in general was not written to scale well to longer melodies. I had implemented zooming and scrolling capability in an earlier version of Musicat to avoid these problems (see Chapter 8), but the present version is lacking zooming and scrolling, and instead squishes everything horizontally into a fixed-size window.



Figure 7.15: On the Street Where You Live (run 1).

The first run had plenty of problems, but it was off to a great start. Measures 1–2 have been grouped together, as have been measures 3–4. A meta-group, (1-4), has formed around these two groups. These two small groups are also involved in the analogy $(1-2)\leftrightarrow(3-4)$. After reading the previous 100 pages or so of this thesis, filled with many examples of analogies made by Musicat, the reader might not be very surprised to see yet another mundane, run-of-the-mill 4-measure analogical structure created by Musicat. However, as simple as this analogy might look, it was surprisingly difficult to get previous

versions of Musicat to make it. The next chapter discusses this in more detail. For now, consider what this analogy means: Musicat has seen the *ascending* melodic line in measures 1-2 ("I have often walked...") as analogous to the *descending* melodic line in measures 3-4 ("...down this street before"). Specifically, Musicat has noticed that the contour of the first two measures is similar to the contour of the second two measures: they share a pattern of small steps, larger leaps, and note repetitions (step-step-leap-repeat), with the only difference being the direction of motion. Thus measures 1-2 have the contour:

step up, step up, leap up, repeat note

(notes: C–D–E–A–A)

while measures 3–4 have the opposite-direction contour:

step down, step down, leap down, repeat note

(notes: G-F-E-C-C)

Additionally, Musicat has noticed that measure 1 has the same rhythm as measure 3, and likewise that measure 2 has the same rhythm as measure 4. Furthermore, it has noticed that whereas the first group, (1-2), ends on a relatively unstable note (A), the second group, (3-4), ends on a stable note, the tonic (C).

Musicat's listening performance for measures 1-4 is just what we had hoped for. Measures 5–8, however, pose some problems in this run. First, they have been divided up into two groups, (5–6) and (7–8). This grouping in analogous to that in measures 1–4, and in that sense it is justifiable. However, group (1–2) ends with a long note, which makes the group boundary between measures 2 and 3 quite obvious. Group (5–6), on the other hand, consists entirely of quarter notes, and the notes B–C at the end of measure 6 are repeated at the start of the measure 7, causing a feeling of continuity that links measures 6 and 7 together. During the run, after measure 4, Musicat had an expectation for groups (5-6) and (7-8) to form. I agree that I have this expectation in listening, but between measures 6 and 7 I experience surprise as the tense leading-tone, B incessantly repeats without resolving to a more-stable tone such as C, and then when C finally appears, offering the possibility of tonal closure and a group ending, the melody *immediately* returns to B. All of these details move the melody forward and avoid establishing the expected group boundary between measures 6 and 7. Thus I don't hear a boundary until after measure 8; at that point, I have heard the group (5-8), with no subgroup boundary between measures 6 and 7.

The groups formed after measure 9 in this run are hard to understand. Unfortunately, measure 9 itself didn't end up as part of a group (the red happiness rectangle underneath the measures makes this problem obvious), and then after that point, several groups in a row are shifted over from what we would expect. That is, the groups associated with measures 10–15 would make more sense if they were shifted to the left to span measures 9–14 instead. Just in case the measure numbers are too confusing, I'll restate the problem in terms of the lyrics: the natural grouping I expected to hear in these measures separates the lyrics into these three groups (with groups indicated by parentheses, and meta-groups by larger parentheses):

((All at once am I) – (several stories high)) – (knowing I'm on the street where you live)

The grouping in this run, however, looks like this (including the non-grouped measure 9 at the start):

All at once am ((I several stories) (high, knowing I'm)) (on the street where you live)

Even though the program doesn't have access to the lyrics, this grouping still seems hard to justify when we think of the notes of the melody. Instead of trying to understand the bizarre structures the program made after measure 9, let's move on and look at another run.



Figure 7.16: On the Street Where You Live (run 2).

This run can best be described as "analogies galore!" It looks very different from the previous run, with a large number of groups and a tangle of analogies linking them. The fundamental difference is that here the program has made a group of every successive 2-measure pair: it has formed groups (1-2), (3-4), (5-6), and so forth, all the way through (15-16). (As a frame of reference in these figures, keep in mind that group (1-2) is the first group on the left, and it crosses the thick bar line and continues for another 6 beats after the bar line. Group (9-10) similarly crosses the thick bar line in the center of the melody. And group (15-16), naturally, is the last one in the melody, and hence does not cross the final bar

line.) Two meta-groups have also formed: the expected (1-4) as well as an unexpected group, (11-14). The latter group encompasses the lyrics "several stories high, knowing I'm on the street", which is a bit strange at the end, although the melody of "several stories high" does indeed sound like it flows nicely into the "knowing I'm" part of the melody, so it is not a completely indefensible group, even though other groupings seem much stronger to me.

Because Musicat heard so many groups in this run, it also had opportunities to make many analogies between the groups. There are so many that the figure above is confusing. I used the detail slider to simplify it:



Figure 7.17: On the Street Where You Live (run 2), low detail.

This figure shows only the strongest groups and analogies. The group (11-14) is not visible, which is consistent with the idea that this was not a very good group. Additionally, the final group (15-16), has disappeared. This is also not too surprising for me, because a more reasonable group would have been the 4-measure group (13-16) with no 2-measure subgroups. For a human listener, measures 13-15 form a sequence ("Knowing I'm" — "on the street" — "where you live"), and measure 16 involves no note attacks; it has single long

note that has been tied over from measure 15. Therefore, even though it makes sense to hear measures 15–16 as a small group, it makes much more sense for all four measures 13–16 to be grouped together, and ideally heard as a sequence.

In this low-detail figure, there are four analogies that link neighboring groups:

$$(1-2) \leftrightarrow (3-4)$$
 $(5-6) \leftrightarrow (7-8)$ $(9-10) \leftrightarrow (11-12)$ $(11-12) \leftrightarrow (13-14)$

There are also two analogies that span long distances:

$$(1-2) \leftrightarrow (9-10)$$
 $(1-2) \leftrightarrow (11-12)$

Notice how the first three neighbor-group analogies involve the first three instances of the main melodic theme. I already discussed the first neighboring-groups analogy, $(1-2) \leftrightarrow (3-4)$, because it appeared in the previous run. Musicat found it again this time. The next two analogies Musicat found involving neighboring groups are very similar to this one, because the melody does "the same thing" three times in a row: the first 4-measure group consists of a specific rising and then falling theme, with a characteristic rhythmic pattern. Each of the next two 4-measure segments of the melody is a variant on this initial theme. The analogy $(5-6) \leftrightarrow (7-8)$ is a bit weaker than $(1-2) \leftrightarrow (3-4)$, which makes sense because there is no long note in measure 6 that can be paired with the long note in measure 8. (Indeed, I think that I hear measures 4-8 as a group, and hear an analogy between (1-4) and (5-8) more than I hear the smaller analogy $(5-6) \leftrightarrow (7-8)$.) The next analogy, however, $(9-10) \leftrightarrow (11-12)$, is a strong analogy for Musicat, and it is similar to the strong $(1-2) \leftrightarrow (3-4)$ analogy: there is a long note in measure 10 that pairs with the long note in measure 12, in addition to the similar contour at the start and at the very end of the melody. All in all, then, Musicat's differing strengths for these three analogies are easy to understand: the first and third analogies are strong, and the second one is weaker. These three analogies, together, are part of a cogent way of hearing the first 12 measures, and the fact that they remain present even at the low detail level shows that Musicat has "heard" their significance, to some extent. The final neighbor-group analogy, $(11-12)\leftrightarrow(13-14)$, is different from the previous three: whereas the first three analogies each involved the internal structure of a 4-measure segment that was similar to the main theme of the melody, this analogy links the *end* of the third 4-measure segment to the *start* of the final descending sequence in the melody (although Musicat unfortunately did not perceive this final descending sequence). This analogy helps to explain the genesis of what one might think is brand-new musical material in the final sequence (measures 13–16): Musicat's analogy points out that the start of that sequence sounds somewhat like measures 11–12. This analogy also helped support the 4-measure group (11–14); even though I don't hear the grouping that way (because it is not as salient to me as the rival group (13–16)), Musicat's analogy and its grouping of measures 11–14 points out how the last four measures are connected to what came before.

The two strong long-distance analogies in this run both involve the first two measures of the melody, which is unsurprising because those first two measures contain the motif from which the rest of the melody derives. Indeed, the earlier high-detail figure shows many analogies from (1-2) to other groups. The long-distance analogy $(1-2)\leftrightarrow(9-10)$ sounds particularly strong to me as well as to Musicat: groups (1-2) and (9-10) both start just before a thick bar line. Both groups involve an initial stepwise ascent followed by a leap up, and their rhythms are identical. The analogy $(1-2)\leftrightarrow(5-6)$ is also quite important to me, and it was also created by Musicat, but it is not as strong and only shows up in the high-detail picture, because the rhythm of measure 2 is quite different from that of measure 6. The other strong long-distance analogy, $(1-2)\leftrightarrow(11-12)$ is a bit of a surprise, but it is easy enough to understand by thinking of transitivity: there is a strong analogy between (1-2) and (9-10), as well as between (9-10) and (11-12), and this chain of analogies,

 $(1-2)\leftrightarrow(9-10)\leftrightarrow(11-12)$, naturally suggests the $(1-2)\leftrightarrow(11-12)$ analogy. However, for a human listener, I think that the $(1-2)\leftrightarrow(5-6)$ analogy is heard much more strongly. Even though people may theoretically understand nearly every pair of measures of this melody in relation to the first two measures, this analogy $(1-2)\leftrightarrow(11-12)$ still strikes me as somewhat unnatural.

In this run, there was no indication, unfortunately, of any relationship *between* the first three neighbor-analogies. (Similarly, the current version of Musicat is not able to notice that the highest note of the first three 4-measure segments is getting progressively higher.) Even though there were many strong 2-measure analogies, the program didn't form anything larger than a 4-measure group, and the large-scale analogies I hoped to see between, say, (1-4) and (9-12) are not present. Will Musicat be able to "hear" these structures on other runs?



Figure 7.18: On the Street Where You Live (run 3).

In the third run, Musicat did indeed find some larger-scale structures (shown in red). Many of the small analogies found in the previous run were also found here (notice that these small analogies were colored red in the earlier figures, but in this figure we see that Musicat has automatically colored them green to allow the large analogy to be more visible). In this run, however, the grouping structure is very regular: not only are all successive measure pairs grouped together, but also every successive pair of these 2-measure groups is surrounded by a 4-measure meta-group. Furthermore, the first half of the melody has yet another level of grouping: an 8-measure group has formed. It is disappointing that the second half of the melody is not grouped analogously, and that no 16-measure group has formed, but otherwise the structure makes quite a lot of sense (although as I mentioned previously, I hear group (13-16) without any 2-measure subgroups, and likewise I don't hear a group boundary inside group (5-8).)



Figure 7.19: On the Street Where You Live (run 3), low detail.

I reduced the detail level to highlight the strongest structures in this run. The large analogy $(5-8)\leftrightarrow(9-12)$ and the large groups were evidently the structures perceived by Musicat as being the strongest (except for group (13-16), which doesn't appear in the low-detail figure). The importance and strength of the structures in this picture seem quite reasonable, although it is disappointing that group (1-4) is not involved in any analogies, since it is the first statement of the theme. The strongest analogy that I myself hear is $(1-4)\leftrightarrow(9-12)$, which is missing. Similarly, $(1-4)\leftrightarrow(5-8)$ is quite important but missing. That is, all three of these strong 4-measure groups are closely related, and Musicat has noticed only one of the three possible large relationships. The other important missing structure here is the descending sequence in the final four measures, which Musicat again fails to perceive (in general, the program needs improvement in the sequence domain).



Figure 7.20: On the Street Where You Live (run 3), medium detail.

I raised the detail level to "medium" to see why the large-scale analogies that I expected, involving group (1-4), did not occur. The figure shows that the analogy $(1-2)\leftrightarrow(9-10)$ formed again, as in the last run, but it was only of medium strength, and the program did not find the strong parallel analogy $(3-4)\leftrightarrow(11-12)$; if it had been found, it would have been easy for the program to make the larger analogy $(1-4)\leftrightarrow(9-12)$. Similarly, there are no analogies even of medium strength that link components of (1-4) to (5-8), so no larger-scale analogy had a chance of forming between these groups.

I was happy that the grouping structure of this run was fairly good, and happy about the large red analogy, but I was still hoping that Musicat would find more of these largerscale analogies, so I ran it again.



Figure 7.21: On the Street Where You Live (run 4).

At first glance, this run looks similar to the previous one, with the large red analogy in the center. However, in the center of the melody there is a large 8-measure *group* as well, straddling the thick bar line in the center; this is a strange grouping in comparison to the rival groups, (1-8) and (9-16), that might have been formed instead.

Moreover, the red happiness rectangles indicate a serious problem: measures 5 and 16 are not members of any groups. This is a clue to a deeper problem: a closer look at the large red group reveals that it is shifted one measure to the right of the expected grouping — all of the subgroups are out of phase with the regular structure that the program has perceived in the other runs of this melody. For example, the 4-measure red group above is (10-13), not the expected (9-12). This run was disappointing; I ran the program again.



Figure 7.22: On the Street Where You Live (run 5).

This run reminds me of run 3 because of the presence of the large analogy $(5-8) \leftrightarrow (9-12)$, the four 4-measure groups, the large 8-measure group (1-8), and many analogies. In this figure we see even *more* groups and analogies — what a jungle! The missing group (9-16) has finally appeared: Musicat has correctly perceived the melody as composed of two halves, each itself divided in two; there are four 4-measure groups. The entire melody has been grouped together into the very large red group (1-16). The only disagreement I have with this grouping structure is the same thing I mentioned on earlier

runs: I think that (5-8) and (13-16) should not be further subdivided, and instead I hear them as 4-measure groups (and in the case of (13-16), it should be heard as a sequence).

A plethora of analogies was formed in this run, and they seem reasonable, since so much of the melody derives from the opening 2-measure motif. I used the detail slider to focus on the strongest analogies in this tangle of green arcs (here, all the analogies are shown in green because the color red was reserved by the program for the largest structure: the 16-measure group).



Figure 7.23: On the Street Where You Live (run 5), low detail.

It is much easier to see the large analogies in this low-detail view. The analogies $(5-8)\leftrightarrow(9-12)$ and $(5-8)\leftrightarrow(13-16)$ have formed. The first of these is the strongest, as it should be: the analogy involving (13-16) is weaker because (13-16) is rather different from the first three-fourths of the melody. Quite disappointingly, however, the group (1-4) is still not involved in a large analogy. As I mentioned before, it is the first statement of the theme, and the rest of the melody should be heard in relation to it. Indeed, Musicat makes 2-

measure analogies involving (1-2) and (3-4), but misses the larger analogies. In this figure, however, we see that group (1-4) was heard as a weak group in this run, effectively preventing it from being involved in large analogies.



Figure 7.24: On the Street Where You Live (run 5), medium detail.

I raised the detail level to "medium" and saw an additional analogy involving 4measure groups: $(9-12)\leftrightarrow(13-16)$. Thus, all of the possible mappings involving the strong 4-measure groups have been made. That is, Musicat has heard each of these groups in terms of the others. Only the initial theme, (1-4), is missing, but as was explained previously, it is simply because that group was not heard as a strong group, unfortunately.

I decided to run Musicat one more time to see if that first 4-measure group might be involved in a large analogy next time.



Figure 7.25: On the Street Where You Live (run 6), with a very large analogy!

In this run, group (1-4) is again missing the expected analogy. However, in this run that lack seems a bit less important, because a very large analogy has formed! The red analogy here, $(1-8)\leftrightarrow(9-16)$, demonstrates that Musicat has the ability (at least given several listening opportunities) to make rather large-scale analogies. This is the largest structure I have ever seen Musicat create. Back in 2010, it was quite difficult to make Musicat see a small analogy involving just four measures. Therefore, I was extremely pleased in 2012 to see

Musicat create this figure, the likes of which seemed unreachable two years earlier. Just in case it seems trivial for a computer program to make this sort of connection (one might think "Eight measures on the left mapped to eight measures on the right — big deal!"), bear in mind that the program is restricted to a simulation of real-time listening and is attempting to model, albeit in a very coarse manner, some of the constraints of human short-term memory. For instance, on some hypothetical run, Musicat might hear measure 16 and then realize that if only group (1–4) were stronger then the whole red analogy could be strengthened (perhaps by forming the sub-analogy (1–4) \leftrightarrow (9–12)); however, the program is not allowed to go back and retroactively modify groups from the distant past. It is nontrivial for the program to make a 16-measure structure, because the correct substructures have to get put in place more or less in real time, when the relevant measures are first heard or are still fresh in memory. That it succeeded in generating this large analogy was quite gratifying to me.

I do, nonetheless, have two criticisms of this large analogy. First, it should have been stronger: the lack of analogy $(1-4) \leftrightarrow (9-12)$, as in previous runs, is disappointing, and indicates an area for future improvement in the program. Second, although the very large analogy makes some sense, I think a more informed listening performance would avoid mapping the whole first half of the melody onto the second half, and would instead involve hearing the large-scale structure of the melody as three sequence-like instances of the 4-measure theme, followed by a final 4-measure winding-down sequence. That is, I would like the program to hear the melody as having the following formal structure:

(A A' A'' B)

This form implies that large-scale analogies would be heard between all the A segments, and the entire melody would be heard as a group, with no extraneous grouping
that would break up the progression from A' to A''. Incidentally, the idea of hearing the entire melody in this way, with three related A parts followed by a final B part, is very similar to the idea of hearing the sequence in measures 13–16 as composed of three "copies" of a 1-measure group, followed by an extension in measure 16 to complete the phrase. This structural idea of "3 + 1" applies, obviously, to musical structures of various sizes; it is a pleasing coincidence that it appears here in the same melody in two different ways, with structures of two different sizes.

In any case, Musicat's performance on this melody is encouraging, in that it noticed quite large-sized groups and analogies, especially in comparison with what it saw in the previous chapter's Simple Melodies. Musicat, however, misses important aspects of this melody, and on some runs fails to make many good structures at all, so it still needs much improvement before it will be able, consistently, to hear the fundamental and salient structures of "On the Street Where You Live".

The successive runs in this section (starting from run 1 and continuing to this one, run 6) have resulted in pictures of generally increasing size of the groups and analogies formed. This kind of behavior would be expected in a hypothetical version of Musicat that remembered a melody over multiple runs and that was able to build up ever-larger structures with ever greater ease thanks to earlier acts of chunking and memorization of structures. However, the real Musicat forgets everything between runs; in this version of the program, any systematic progression perceived in Musicat's multiple performances is just an illusion. I also did not include in this section another run of Musicat that had few large groups and that was quite similar a run that had come at the start of the section. If the program had a persistent long-term memory, this kind of behavior — forming larger structures after repeated listenings — might occur systematically, and Musicat would make richer associations not only across runs on a *single* melody but between *different* melodies.⁹



TENNESSEE WALTZ (STEWART AND KING)

Figure 7.26: Tennessee Waltz.

Whereas the previous melody, "On the Street Where You Live", was used throughout the process of developing and testing Musicat, the remainder of the melodies in this section were not considered until the program had reached (or very nearly reached) its present state of development.

In an earlier section, I presented the 32-measure melody "Younger than Springtime" to Musicat in piecemeal fashion, only 8 or 16 measures at a time. Tennessee Waltz also has

⁹ Harry Foundalis's program Phaeaco, working in the domain of Bongard problems, demonstrates a long-term memory mechanism such as the one suggested here. I plan to implement the idea in a future version of Musicat.

32 measures. I tried letting Musicat listen to the entire 32 measures at once. The results illustrate some of Musicat's problems. The figure was far too tall for the program window, so a large portion of the image is cut off at the top, but this was the least of the program's problems. See the figure below.



Figure 7.27: Tennessee Waltz.

This strange picture looks more like the mysterious traces of subatomic particles seen in supercollider experiments, or a diagram drawn in some alien language, than it looks like an illustration of musical structure. In part this is due to the lack of support for large numbers of measures in the current drawing code, but it is also strange-looking because Musicat has created very strange structures. The first 8 measures actually do have a normallooking structure, and groups such as (1-2), (1-4), and (1-8), are visible, but after this we find the strange-length and quite strong group (1-14) (which is itself enclosed in the much weaker group (1-16), making a very unlikely-looking structure), an orange-colored analogy $(17-24)\leftrightarrow(25-28)$ between groups of quite different lengths (an 8-measure group linked to a 4-measure group), and some long-distance analogies that seem quite arbitrary. Musicat has trouble with melodies of this length, especially if they are complicated. I don't expect that the problems will forever remain insurmountable, but for the time being I decided to give Musicat shorter chunks than this. Therefore, I split "Tennessee Waltz" into two halves. It turns out, though, that this was still a tricky challenge for the program.





Figure 7.28: Tennessee Waltz, first half.

The numbering scheme I use for measures here, as in previous melodies, considers "measure 1" to start on the first note of the melody and to continue for the length of one measure — here, 3 beats. Measure 1 thus contains the notes C–D–E–G ("I was dancin"), measure 2 contains the next 4 notes, again C–D–E–G ("with my darlin"), and so on.

Before looking at the program's output, let's consider the structure of the melody for a moment. These 16 measures are made of two nearly-identical parts: measures 9–12 are an exact copy of measures 1–4. Measures 13–16 are slightly different from measures 5–8, but the rhythm pattern is nearly the same, as are most of the pitches. So the big picture expected here, at least by a human listener, is the analogy $(1-8) \leftrightarrow (9-16)$.



Figure 7.29: Tennessee Waltz, first half.

The large analogy I hoped for is missing in this run. Only two analogies have formed at all — this in striking contrast to all the analogies discovered in "On the Street Where You Live". I found this puzzling until I remembered all the blue measure links at the top of the picture, which indeed link many of the measures together. That is, Musicat perceived rhythmic similarities between many pairs of measures, and such noticing of similarities between measures is, as I have pointed out earlier, a limited form of analogy-making. However, for some reason, in this run, the program has not made very many links between structures longer than a single measure. Part of the problem may be the irregular grouping structure it created: the first and last 4-measure groups, (1-4) and (13-16), are the ones I expected. However, a central group, (7-10), has been formed. As happened in some previous runs on other melodies, this group, straddling the center bar line, leads to other strange groups that are hard for the program to match up with other groups in the melody. The two analogies found here, $(1-2)\leftrightarrow(3-4)$ and $(5-6)\leftrightarrow(13-14)$, do indeed link similar melodic structures together, but the obvious analogy — mapping the 8 measures constituting the first half onto the 8 measures constituting the second half — isn't possible when the grouping structure has turned out so strangely, with the single group (7-10) spanning both halves.

I tried a second run on this melody, and got a pleasant surprise (figure on next page).



Figure 7.30: Tennessee Waltz, first half (run 2).

In this second run, Musicat discovered a sequence. Not only was this the first time that I saw Musicat find a sequence that was strong enough to persist through an entire run, but I hadn't anticipated or even perceived any sequence myself! A very similar thing had happened during the run of Bad Melody #4 in chapter 5: Musicat discovered a "sequence" that does not involve successive measure transposition (despite the code in the program requiring each successive segment of a sequence to be a transposition of the previous one, by a number of scale steps that is constant for the duration of the sequence.) In Bad Melody #4, each measure in the sequence was identical to the one preceding it, so Musicat perceived a sequence with a 0-step "transposition" between successive measures. In this case, it's very similar. Measures 1 and 2 are identical, and measure 3 is exactly one octave higher, which Musicat also considered to be identical. Successive chunks in Musicat's sequences must be related by transposition, but Musicat looks only at pitch *classes*, not absolute pitch heights; in other words, octave information is discarded when the program is looking for a sequential pattern. This strategy allows Musicat to recognize as a sequence a passage that jumps between octaves in a circle-of-fifths progression, for instance. In measures 1-4 of "Tennessee Waltz", then, Musicat saw each measure (or group — the group (1-4) constitutes the third chunk of Musicat's perceived sequence) as a 0-step transposition of the previous, just as in Bad Melody #3; the octave leap between measures 2 and 3 was equivalent to a 0-step "transposition" for the program. Altogether, then, these 0-step "transpositions" allowed Musicat to "hear" a sequence here. It would be easy to fix this behavior if I considered it a bug, but in this case it serendipitously helped the program discover something interesting about the melody. I was quite surprised to see the blue sequence lines in the figure, because I hadn't consciously thought of measure 3 as being the same as measure 1 or measure 2. It sounded, to me, like a much-higher version of the same material, but as a listener I hadn't realized it was simply an exact octave transposition. This surprise was a gratifying instance of Musicat's listening performance informing my *own* understanding of a melody. Thanks, Musicat!

In this run, the program created a few interesting structures, but it missed many obvious things and created some strange structures as well, so the rest of the run was a disappointment. After discovering the surprising "sequence" in measures 1–4, the program *should* have heard measures 9–13 in the exact same way, as a sequence. Instead, it formed two groups, (9-10) and (11-12), instead of linking all four measures together. Worse, these belong to different meta-groups, (7-10) and (11-14). There is another group spanning the thick center bar line of the melody, (7-10), and another one just after it, (11-14), leading to the preposterous group (7-14). This group is truly bizarre, because it includes the cadence of

the first half of the melody, (7-8). Furthermore, Musicat did make the analogy $(7-8) \leftrightarrow (15-16)$, linking the cadences of each half, but it didn't realize that (7-8) and (15-16) were the ends of groups.



Figure 7.31: Tennessee Waltz, first half (run 3).

On a third run, Musicat again found a sequence (in measures 1–4), but the sequence disappeared early in the run, replaced with two 2-measure groups. But overall, the results were better: the program finally made the hoped-for large-scale grouping structure. The melody has been divided into an 8-measure first half and an 8-measure second half. It also

made two small analogies between the halves: $(3-4) \leftrightarrow (11-12)$ and $(5-6) \leftrightarrow (13-14)$. Unfortunately, though, it missed the most obvious repetition, $(3-4) \leftrightarrow (11-12)$, and the larger-scale analogies $(1-4) \leftrightarrow (9-12)$, and so on. Thus it didn't ever make the *large-scale analogy* $(1-8) \leftrightarrow (9-16)$.

Tennessee Waltz, Second Half



Figure 7.32: Tennessee Waltz, second half.

I run Musicat on the second half of the melody (the measures are renumbered here to 1–16 for convenience). Chopping the melody in half this way was frustrating, because one of the important features of this melody, and of many others, is that the final 8 measures here (measures 25–32 in the original non-chopped-up melody) are an exact copy of measures 9–16 in the first half. The program won't be able to see the large-scale **A A' B A'** form of the melody when it hears the two halves separately. Still, however, I was hoping that this run would prove interesting because there are connections to be found between the **B** and **A'** sections.



Figure 7.33: Tennessee Waltz, second half (run 1).

The group structure in this run looks reasonable, except that larger groups did not form. The 2-measure and 4-measure groups that did form seem reasonable (although, just as in "On The Street Where You Live", I myself don't hear any group boundary between measures 6 and 7, but here the program has formed groups (5–6) and (7–8)). The main theme in measures 1–4 formed a group, as did the winding-down end of the melody, (13– 16). The program found quite a few analogies, although because there were few large groups, all the analogies were between 2-measure structures (measures 5–8 are very similar to 13–16, so that would have been a very likely analogy candidate if the group (5–8) had formed). The first strong analogy in the figure, $(1-2)\leftrightarrow(9-10)$, identifies an important thematic connection between the **B** and **A** material in the melody. Even though measure 2 is quite different from measure 10, measures 1 and 9 are the same melody (with pitches one octave higher in measure 1 than in measure 9), and the program makes an analogy between the two groups because the groups share this starting material. Furthermore, this connection is essentially the same one that led to the sequence in the previous run, measures 1–4. I hoped that Musicat would make the similar analogy between (1-2) and (11-12) in this half because, although these two groups occur at different places in the metric hierarchy, measures 1 and 11 are both in the top octave of the melody and have very similar (but not identical) rhythms and pitch contour. Unfortunately, it didn't find this analogy in this run or the next, but doing so might provide a goal for a future version of the program.

The program found two other easy-to-understand analogies: $(5-6)\leftrightarrow(13-14)$ and $(7-8)\leftrightarrow(15-16)$. (Unfortunately, though, it did not make the obvious meta-analogy $(5-8)\leftrightarrow(13-16)$ here, although the two analogies it did make suggest that it came close.) These connections exemplify something I noticed about "Tennessee Waltz" that made it seems like a really good test for Musicat: the entire melody is made up of just a few basic motifs (exemplified by the rhythms and pitches of the chorus section, measures 1-4 in this half) which are developed slightly to form the rest of the melody (recall the earlier discussion of Arnold Schöenberg's book on melody and motif-development). The melody of "Tennessee Waltz", furthermore, has a casual and imprecise feel to it in spots: different singers (and different versions of the sheet music) use slightly different rhythms and pitches in some sections of the melody (except for critical places like the very start of the song and the chorus

section, which seem more standardized). Listening to two separate and melodically-varying performances of this melody and understanding them as instances of the same song requires — obviously! — a large amount of flexibility in one's *listening performance*. Hearing an analogies between measures 5–8 and 13–16 (or hearing the two sub-analogies found in this run) requires exactly this sort of flexibility. The following figure compares these two segments of music:



Figure 7.34: Measures 5-8 (top staff) and 13-16 (bottom staff).

These two segments look similar to each other when they are aligned vertically because they both have the same rhythm. This is a nice example of the utility of Musicat's bias that considers rhythmic somewhat more important than pitch in detecting similarity: even if no pitch information were available at all, these melodies would be heard as "the same". There are, though, some important pitch similarities: two of the downbeats have the same note in both segments: the C for lyrics "know" and "beau(tiful)" and the E for "much" and "Ten(nessee)". Most of the less important notes — those on the non-downbeats — are different in the two segments. Notice, for example, the curious swapping of notes A and G in the second full bar of the figure: notes A and G ("just how") in the top segment become G and A ("-ti-ful") in the bottom segment. The final downbeat in each segment (with lyrics "lost" and "Waltz") has the note D, suggesting a dominant chord, in the first segment, and the tonic C at the end, helping human listeners, as well as Musicat (potentially), to hear these two segments as parts of an antecedent \rightarrow consequent relationship.

Even though Musicat didn't see the entire analogy $(5-8) \leftrightarrow (13-16)$ in this run, the two smaller analogies it created for these measures show that it is flexible enough to hear similarities despite these melodic differences.



Figure 7.35: Tennessee Waltz, second half (run 2), after measure 12.

I ran Musicat on the second half of "Tennessee Waltz" one more time. Although in this chapter I have been focusing on the *final* results of each run, I think it will give some insight to the reader to see a few measure-by-measure screenshots of this run in progress. With simpler melodies such as "Twinkle, Twinkle" we have already seen how groups and analogies rapidly come in and out of existence as the program listens, but it is interesting to see this happening for larger structures. During the second run, after measure 12 the groups and analogies of the first 8 measures look good — the two red analogies here were not created in the first run, and the group (5-8), which was missing in that run, preventing the desired analogy, has been formed. The only strange thing is the large purple expectation, which is essentially a copy of the structure in the first 10 measures, which is predicted to occur starting on measure 11. This is unreasonable, because the important group here is the 8-measure group (1-8), not (1-10). Indeed, the strangely sized group (1-10) does not exist in the figure. It must have come into existence briefly — just long enough for the expectation to be generated — only to be destroyed soon after.



Figure 7.36: Tennessee Waltz, second half (run 2), after measure 13.

Fortunately, a measure later, the large purple expectation has vanished, and has been replaced by a more reasonable expectation for the next 4 measures (which will turn out to be the final measures of the melody). The 4-measure group (5–8) has, unfortunately, been destroyed, although another one, (9–3), has appeared, but this one is very weak at this point.



Figure 7.37: Tennessee Waltz, second half (run 2), after measure 14.

After another measure has passed, the group (9-13) has also vanished as has the 4measure purple expectation. Although these are frustrating developments, they are somewhat understandable as consequences of the dynamic nature of Musicat's perception. A more serious problem is that we have three new large structures, of which two make little sense. First, the red group (7-13) starts on measure 7 for no good reason, and is 7 measures long! Fortunately, the group is weak, as is indicated by the orange happiness rectangles underneath it. This red group has induced a ridiculous purple expectation for another 7-measure structure starting on measure 14; the previous 4-measure expectation was much more reasonable (and would have eventually helped to form the strong group (13-16), had the expectation not been destroyed and replaced by this one). Finally, a large group (1-14) has been formed. This group is not so strange, as long as in the future it is expanded to include the next 2 measures of the melody, as they are heard.



Figure 7.38: Tennessee Waltz, second half (run 2), after measure 15.

After one more measure the weird 7-measure group (7-13) has fortunately been destroyed, and a rival red group (1-8) has formed. I call it a "rival" group because measures 1–8 overlap with measures 7–13, and it is likely that while the program listened to measure 15, these two groups actually had a competition in which group (1-8) emerged victorious.

The strange 7-measure purple expectation is still in place, but notice that Musicat is, so far, sticking with a newly-perceived group (11-14), which overlaps with the expectation. I would have preferred measures 9–12 to have been heard as a *sequence*, as in run 2 of the first half of the whole "Tennessee Waltz" melody, but this group is not unreasonable.



Figure 7.39: Tennessee Waltz, second half (run 2), after measure 16(+3).

After the final measure was heard, Musicat kept running as usual for a few extra measures' worth of time. The figure above shows its state after the third extra measure. The

strange 7-measure expectation has finally been replaced with a 4-measure expectation, which is simply an expectation for the group structure of (13-16) to occur again. Notice that (13-16) is a new group, as is the weak group (9-12) — thus, a regular, repeated group structure has emerged for the entire melody. Unfortunately, the meta-group (9-16), which I expected, didn't form around these two new groups.



Figure 7.40: Tennessee Waltz, second half (run 2), end of processing.

After a bit more time, the program's listening performance is over. Sadly, the two 4measure groups in measures 9–16 have disappeared. Musicat is aware (see the red happiness rectangles) that measures 11–12 are missing any group structure. As often happens, it found a good grouping structure but threw it away in its constant search for something better. This strategy paid off well a few measures earlier in the run, when the strange 7-measure group (7-13) was replaced with the red group (1-8) that persisted all the way to the end of the run, but it did not work so well for the groups (9-12) and (13-16). Several analogies between the first part of the melody and the last part had been formed earlier, but, like these groups, they were also destroyed. One new analogy that was discovered near the end of the run links two phrase endings: $(7-8)\leftrightarrow(15-16)$, just as in the previous run, but there were many more potential relationships that were missing in this run.

What do these runs on "Tennessee Waltz" tell us about the program? In most of the runs (on each of the individual halves of the melody), Musicat eventually settled on an acceptable group structure at the end of the run, but one that was lacking in larger structures. The program did not find enough of the 4- and 8-measure groups that seem quite clear in the structure of this melody. It found several reasonable analogies in each run, but it seems that had it found better grouping structures, it would have found more analogies. When Musicat finds analogies, that fact influences the grouping structure by suggesting possible new groups and by strengthening groups involved in good analogies, but this mutuallyreinforcing effect was perhaps not strong enough, in this melody, to help the program quickly settle on a strong perceived grouping structure. If this mechanism worked better and if Musicat more consistently found a very strong set of 4-, 8-, and even 16-measure groups, I suspect it would have a much easier time listening to the entire 32-bar melody and would no longer generate "alien diagrams". A strong understanding of the first set of 16 measures would make it much easier to understand the second set of 16 measures. Indeed, if the program had settled on its grouping structure for measures 1-16, it could understand measures 17-32 in terms of the first 16 measures, making analogies galore between segments of the two halves. By the time a human listener has heard measures 1-24 (the **A A' B** parts of the melody), hearing the start of measure 25 makes the final measures, 26–32 almost trivially predictable, because the listener can simply expect to hear another instance of **A** or **A'**.

GOOD PEOPLE ALL (HOFSTADTER)



Figure 7.41: Opening measures of Good People All, from a cantata by Douglas Hofstadter.

This melody is from the piano introduction of the aria Good People All, by Douglas Hofstadter. I have included it in this Complex Melodies chapter because, although it is shorter and perhaps simpler-looking than other melodies in the chapter, it has several features that are especially complex with respect to Musicat's abilities. In my original work on Musicat I had intended the program to listen to melodies more like this one, whose musical tradition is more of the baroque or classical variety than it is folk or popular. I had also originally planned to use melodies in Steve Larson's "Seek Well" microdomain, in which all the notes are equal duration. This melody mostly fits in that microdomain: all the notes are eighth notes, except for a few exceptions in the final two measures.

Before looking at Musicat's listening performances, let's consider the structure of this melody. I anticipate several significant problems for Musicat that derive from the melody's relentless stream of eighth notes (a "motoric" rhythm reminiscent of those in Bach's Inventions.) Musicat's primary mechanisms for grouping and analogy-making rely heavily on the music having distinct rhythmic patterns, but nearly every measure here is rhythmically identical. Within this stream of eighth notes, however, there are obvious patterns that evoke more complex rhythms. Most obvious are the descending scale segments in measures 2 and 4. The first note of the start of each descending scale is extremely salient for a human listener, so it sounds like there is an accent on beat 2. In measure 1 (and measure 3), beats 1 and 3 also sound accented because they occur on strong beats and form the start of small groups. Finally, a repeating pattern of notes causes the notes at the end of measure 1 (or measure 3) to sound like they form a group that continues into the next measure, so beat 1 or measure 2 (or measure 4) is not accented. The next figure hopefully makes this implied rhythm (or equivalently, in this case, the implied grouping structure), clear:



Figure 7.42: Good People All (measures 1-4): starting points of groups indicated by upper voice.

In this figure, the upper-voice notes (stems up) indicate the notes I hear as the beginnings of groups. Their durations have been extended to correspond with the lengths of the implied groups. For instance, the high E in measure 2 is a dotted half note, indicating that I hear a group starting on that E that is a dotted half note in length (the 6 eighth notes descending in a scale). To make this even clearer, the following figure shows only the groups' starting notes:



Figure 7.43: Good People All (measures 1-4): groups represented by their starting notes.

The four measures in the figure have a distinctive rhythm when one hears them in this way — they are no longer simply a stream of eighth notes.

Now, remember that the preceding discussion has focused on a human hearing of this melody. Musicat is missing some of the essential abilities that would allow it to hear the melody in this way. Most importantly, recall that Musicat's groups must start on measure boundaries (unless the melody has pickup notes, in which case starting points are systematically shifted). The groups implied by the preceding figure, however, start in a variety of places: beat 1 and beat 3 in the first measure, beat 2 in the second measure, and similarly for measures 3 and 4. Musicat is not able to make these groups; for example, it can make a one-measure group containing measure 2, but not a partial-measure group containing only the 6 eighth notes of the descending scale in measure 2.

Not only does this melody suggest a *small*-scale grouping structure (lengths less than a measure) that Musicat cannot represent, but also the 10-measure length of the melody suggests a *large*-scale group that might pose problems (since 10 measures is longer than the more normal power-of-2 length of 8 measures) might pose problems for the program. (An amusing coincidence involving the number 10 in two different contexts: we might expect a grouping boundary on the bar line after the first 8 eighth-notes, between measures 1 and 2, but a strong group boundary actually occurs 2 notes later, after the first 10 eighth notes of the melody; similarly, we might expect a strong group boundary after measure 8, but then we see that the melody doesn't end until 2 measures later, after measure 10.) To illustrate how the group lengths in this melody — at both small and large scale — differ from the standard group lengths that Musicat expects, I *recomposed* the melody, creating a version with only 8 measures total, and likely group boundaries falling at measure boundaries:



Figure 7.44: Good People All, recomposed.

I claim that much of the interest in the original melody derives from the listener hearing the melody as a surprising variant of a more-expected melody such as my recomposed version. For example, in the recomposed version, the end of measure 1 corresponds to the end of a group, and this ending is expected, because the first two beats have been exactly repeated. In the original version, a group extends from measure 1 through the end of the first beat of measure 2, and this extension is a surprise. Similarly, the leap up to the E on beat 2 in the original melody is more surprising than the leap up to E on the downbeat in the recomposed version. Similar arguments apply at the end of the melody: in the recomposed version, it ends after 8 measures, but in the original version, phrase extension results in the melody continuing all the way to measure 10 before reaching a cadence.

A sophisticated, human-like listening performance for this melody would require the interplay between the expectation for something like my recomposed melody and the actually-heard melody. Tension would be created by groups extending longer than expected or starting on surprising beats. But these are issues for a future version of the program. Let's see what Musicat can "hear" at present in "Good People All".



Figure 7.45: Good People All (run 1).

Musicat performed remarkably well, considering its inability to make the smaller groups suggested above, such as the group of 6 eighth notes in the descending-scale in measure 2. Those smaller groups were not (and could not be) formed, but Musicat's smallest group here, (1-2), is a perfectly reasonable-sounding group. The only problem is that it doesn't have any subgroups that would highlight the interesting substructure. I will return to this point in a moment when looking at the analogies, but first, what about the *super*structure? That is, do the largest groups make sense? The first 4-measure group, (1-4), agrees with my hearing; measure 5 starts a distinct second part of the melody with new melodic material. The second 4-measure group, (5-8), however, doesn't sound like a group to me. Perhaps it would be fine if the group could be extended 1 beat to the right so that it could include the first note in measure 9. However, Musicat's groups can't extend in this way. Additionally, I am unclear about how I hear these measures. I may simply hear a large group that extends all the way to the end of the melody: (5-10). In this run, the final two measures

are separate from the rest of the melody, in a very weak group of their own (the group ellipse is present but difficult to see because the group is so weak). Musicat hears these two measures as a separate group, however, because they have very different rhythms from the other measures in the melody, and also because it (incorrectly) perceives a thick bar line after measure 8. Measure 9 ends with an ascending scale that connects smoothly to measure 9, but the group boundary indicates that Musicat doesn't hear this connection. If the program were able to make neighboring groups that overlap by a single note (group elision), then this listening performance could be made much more plausible by simply extending group (5–8) as mentioned above, and allowing this group to overlap slightly with group (9–10). Regardless, in this run the program unfortunately hasn't seen that there is phrase that extends all the way through measure 10. Even if measures 9 and 10 belong to a separate small group, they should be connected to the larger phrase. All in all, the large-scale grouping structure is not terrible, but it misses the subtleties of the connection between measures 8 and 9 and the extension of the phrase through the end of measure 10.

Returning to the start of the melody and the smaller groups, we see that the strongest analogy created was $(1-2) \leftrightarrow (3-4)$. This analogy is strong because Musicat noticed a contour relationship between the two groups: (1-2) has exactly the same contour as (3-4). The second group is a version of the first that has been transposed down by two steps (with an accidental, G[‡], in measure 3 to emphasize the A-minor chord). This contour, importantly, includes the highly-salient descending scale in measures 2 and 4. But Musicat doesn't notice the significance of the scale (in contrast, a previous version of Musicat would almost certainly have made groups containing each of these scales alone — see the next chapter).

Several other analogies were created. The second-strongest was $(3-4) \leftrightarrow (5-6)$, which indicates that Musicat heard (5-6) in terms of (3-4) — another instance, as in some earlier melodies, of the program hearing how initial musical material is developed to yield later material. In this case, the rhythm of (5-6) was exactly the same as (3-4), but in addition the program heard the contour of these groups as somewhat similar. Two more analogies were created, both involving (7-8), but these were weaker; the program heard these two this group as analogous to earlier groups, but the connection was much weaker than for other groups; I agree that the material in this group seems less related to what came before. And, finally, the last two measures (9-10), are not connected by analogy to any preceding groups, which makes sense considering how different the rhythm is: these are the only measures that are not made of a string of eighth notes. The relative strengths of all these analogies (and the lack of analogies at the end) seem accurate from my own listening perspective.



Figure 7.46: Good People All (run 2).

A second run was very similar to the first. One analogy is missing in this picture compared to the previous run, but the most important analogies are present: each 2-measure group is linked to the following 2-measure group, showing that Muscat hears a forwarddevelopment of the melodic material. Again, the analogies are slightly weaker as we proceed from left to right: the exact transposition linking the first two groups is much stronger than the rhythmic repetition and partial contour similarities involved in the later analogies.

This listening performance resulted in all the same groups as the first, with the addition of a 10-measure group enclosing the entire melody. I was happy to see this: it shows that the program has included the final two measures as part of the large structure. I wish that it had heard one additional layer of grouping and added the 6-measure group (5-10), but otherwise the grouping structure is fairly close to how I hear it.



Figure 7.47: Good People All (run 3).

In this final run, the same grouping structure was created as in the previous run. Two of the analogies from the first run are present: the strong analogy $(1-2)\leftrightarrow(3-4)$ as well as the

much weaker and longer-distance analogy $(3-4)\leftrightarrow(7-8)$, which was present in run 1 but missing in run 2. In this run, a larger analogy was formed, linking the two 4-measure groups together: $(1-4)\leftrightarrow(5-8)$. I don't hear this analogy, and it is a little hard to understand. The sub-analogy $(3-4)\leftrightarrow(7-8)$ contributed to this larger analogy, and in addition, Musicat made two rather strange non-rhythmic links (probably contour links, but it is not indicated in the notation): one between the 2-measure group (1-2) and the 4-measure group (5-8), and the other between (1-2) and (5-6). These links involving (1-2) seem strange to me, and helped the program form this larger analogy. All in all, this third run is harder for me to interpret, but it does show quite a bit of consistency with respect the previous two runs.

In summary, these three runs gave remarkably similar results. The grouping structure always had two 4-measure groups followed by a smaller 2-measure group at the very end, and in runs 2 and 3, the entire melody was included in a 10-measure group. The analogies found by the program were similar in the first two runs, whereas in the last run it found a slightly odd larger-scale analogy.

The biggest failings of the program were the ones predicted at the start of this section: the 10-measure length of the melody wasn't handled as well as possible: the 6-measure group (5-10) never formed on any of the runs. This lack of a large 6-measure group structure had the small-scale counterpart of missing 6-eighth note groups in measure 2 and 4 — no surprise, since Musicat's architecture doesn't allow this sort of group! But the lack of these small-scale groups was somewhat covered up by the strong 2-measure grouping structures that enclosed each of the areas where we would like to see sub-groups. Even though the smallest groups could not be formed, analogies involving the 2-measure groups were formed, and these analogies matcher my own hearing of the melody quite well. Overall, the program did better than I expected, given its inability to hear the smallest groups and the homogenous nature of the eighth-note rhythm.



SUN AND MOON (BOUBLIL AND SCHÖNBERG)

Figure 7.48: Sun and Moon (from the musical Miss Saigon).

The song "Sun and Moon" is the final melody in this chapter, and also was the last melody I decided to add; I didn't even think of including it until long after I had stopped modifying the program code. Along with the Bad Melodies (thanks to their random nature), and a late addition to the Simple Melodies section, "Frère Jacques", it is therefore one of the most "pure" tests of Musicat's listening abilities; nothing in the program's design was influenced by the knowledge that it would be listening to this melody. This will likely be obvious in its listening performances: I know *a priori* that there are many things about this melody that Musicat will be "deaf" to, and thus this melody should illustrate directions for future improvement.

"Sun and Moon" is one of my favorite songs from my favorite piece of musical theatre, *Miss Saigon*, composed by Claude-Michel Schönberg (a distant relative of Arnold

Schönberg), with lyrics by Alain Boublil. Since I have heard this piece — and the entire musical it is a part of — countless times, my own experience of listening to the song is deeply influenced by my memory. As I mentioned when describing Bad Melody 5, Musicat has no long-term memory of this sort, so it misses this part of the listening experience. For me, the most salient feature of the melody is the rhythmic figure in the first two measures. Measure 2 is a development of the rhythm in measure 1 — it simply is a version of measure 1 where the second note has been shortened and an additional quarter note has been added at the end of the measure (this same type of thing happened at the start of "Row, Row, Row Your Boat"). But more importantly, measure 2 is a restatement of the most important rhythmic figure in the entire two-hour score of *Miss Saigon*. This rhythm appears directly in the melody of "Now that I've Seen Her", and in the orchestral accompaniment to the songs "Last Night of the World", "Telephone Sequence", "I Still Believe", "The Fall of Saigon", "Ellen and Chris", and the song "Too Much for One Heart" (which was cut from the production). It also appears prominently in diminution (*i.e.*, the notes' durations have all been cut in half) in the very first notes of the whole score ("Opening Act One"):



Figure 7.49: Miss Saigon, opening notes.

The rhythm in question is evident when we listen to the accented notes. If we extract the accented notes and rewrite them as dotted quarters and quarters with a doubled tempo, the relationship to measure 2 in "Sun and Moon" is obvious:

Figure 7.50: Miss Saigon, opening accented notes, rewritten at double tempo.

Notice that the accent marks above were not strictly necessary — each of these notes occurs at a local high point in the melody, and it would thus sound somewhat accented without any deliberate articulation by a musician performing the piece. But in any case, Musicat currently notices only rhythms that *appear* in a score as notes with various durations, not those *implied* by accented notes.

The context provided by all these other examples of this rhythmic motif is, to me, *essential* in how I hear "Sun and Moon". Musicat will likely hear measure 2 as a rhythmic variant of measure 1 (a truly A. Schönbergian way of hearing this C.M. Schönberg melody!), and also will notice instances of these rhythms that are later repeated, but unfortunately cannot understand the significance and omnipresence of the measure 2 rhythm. By contrast, a human being listening to a production of *Miss Saigon* will practically be beaten over the head with this "dotted-quarter, dotted-quarter, quarter" figure for two hours, and would certainly notice its significance and the coherence it provides to the work.

Another aspect to my personal listening performance of this song is my understanding of the lyrics. These are completely unavailable to Musicat, by design, but it is helpful to remember that Musicat is missing out, through no fault of its own! Lyrics have the potential to influence our hearing of the grouping structure of a piece. If we consider the lyrics alone (paying attention to punctuation and rhyme), we may infer the following grouping structure (groups are indicated with phrase markings above the staff):



Figure 7.51: Sun and Moon, with phrasing implied by the text.

Another possible way of hearing the grouping structure, however, is suggested by considering the harmonic structure of the piece. While the chords are not shown in the following figure, I claim that this is a reasonable grouping structure:



Figure 7.52: Sun and Moon, with an alternate possible grouping suggested by the chords.

I will not argue for the "correctness" of either of these possible ways of hearing the grouping in this melody — both have some merit and might be reasonable listening

performances (likely influenced by the particular stage performance heard by the listener). Regardless of the grouping structure that is heard, there are some phrases consisting of an odd number of measures. In the first half of the melody (measures 1-12), both the lyric-based and chord-based grouping structures suggest the 3-measure group (10–12). In the lyric-based analysis we also have the group (1–3), while in the chord-based analysis we have the group (7–9). And at the very end of the melody, measures 21-25 form a 5-measure group. For this melody, then, finding one or more groups having an odd number of measures seems crucial for a cogent listening performance. We already saw in Bad Melody 4 that Musicat has trouble finding odd-length groups, so how will it hear this melody? (Keep in mind that the program has no knowledge of the lyrics or the chords.) In the runs that follow, I restricted the input melody to half the length of the melody at a time; first we consider measures 1-12.

Sun and Moon, First Half



Figure 7.53: Sun and Moon, first half.



Figure 7.54: Sun and Moon, first half (run 1).

This run seems to have started well: the first six measures are grouped into three red groups of two measures each. This grouping into pairs is consistent with the chord-based grouping suggested above, and also consistent with the observation that measure 2 sounds like a slight modification of measure 1. Musicat has also made groups of the measure pairs 3–4, 5–6, and 7–8. This grouping makes sense because the second measure of each pair can be heard as a slight variant of the pair's first measure, and Musicat tends to group together adjacent structures when they are similar to each other (grouping and analogy-making¹⁰ are two very different activities, but when analogous structures are adjacent, they are also good candidates for larger-scale grouping). This observation about measure-pairs applies not only

¹⁰ By "analogy-making" I intend to include Musicat's low-level ability to notice measure similarity. Even though measure-to-measure connections based on rhythmic or pitch similarity are not official "analogy" structures for Musicat, I include them in the grouping heuristic discussed here.

the rhythms but also to the pitches of these measures: the first measure of each group has exactly the same first two notes as the measure that follows. The red analogies between successive groups also make intuitive sense: each 2-measure group consists of a first measure that has only two notes, while the second measure of each group is an elaboration with additional notes at the end; Musicat notices the similar structure of these groups.

The 4-measure group (1-4) is less understandable to me: the end of measure 4 features a fast rhythm that moves the melody forward, connecting it with measure 5 and making it hard for me to hear a higher-level grouping boundary between these measures. We can see by the thin ellipse in the previous figure that the group is, fortunately, quite weak. More disappointing is the presence of group (7-8), where I had heard the longer group (7-9) instead. During the run, the presence of the short group (7-8) would be expected and even encouraged at first; it would simply continue the pattern established by the previous 6 measures. Indeed, after measure 6, the (7-8) group was present as a purple expectation:



Figure 7.55: Sun and Moon, first half (run 1), after measure 6.

But after measure 8 arrived and the expected group (7-8) was formed, a new expectation for group (9-10) formed. Skipping forward several measures until after measures
9 and 10 arrived, we see that this expectation was still present, although 9–10 remained ungrouped:



Figure 7.56: Sun and Moon, first half (run 1), after measure 10.

Here we can see that Musicat is unhappy with the situation in the final two measures. We can see three developments that seem to be preventing the formation of the desired group (7-9), however. First, group (7-8) is doing fine on its own, as far as the program is concerned. Musicat has found a slightly weak analogy, but an analogy nonetheless, between it and the previous 2-measure group. Second, (7-8) is also contained within a much larger group, (1-8), which is weak, but is nonetheless serving to insulate the group (7-8) somewhat from the following notes (in order to extend group (7-8) to the right, Musicat would first have to break group (1-8), to "free up" its child groups.) Third, there is a very thick bar line standing between measures 8 and 9, serving as a roadblock, or worse, a chasm

discouraging the making of groups that straddle the gap. Of course it is possible for groups to straddle such boundaries (see the red group (1-6) for an example), but boundaries make this much more difficult. Finally, Musicat doesn't see any relation between measures 7–8 and measure 9; no reason has emerged to join these into a 3-measure group. Indeed, it's a good question: why *should* they be grouped together?

Two good reasons jump to mind. The first is one that Musicat is aware of: the whole note in measure 9 is practically screaming out to be heard as the end of a group, simply because of its long duration. The importance of long time intervals between successive note onsets in establishing group structure has been well established in other models of music cognition, and plays a role in Musicat as well. Unfortunately, this pressure for the whole note to be the *end* of a group was not strong enough to overcome the forces exerted by the other pressures mentioned above that serve to maintain the status quo (*i.e.*, the happy status of the shorter group (7-8)) and the thick bar line that indicates that measure 9 should be the *start* of a group.

The second good reason for hearing a group that extends through measure 9 is one that Musicat is not aware of in its current incarnation: the note A in measure 9 is a continuation of a pattern of downward melodic motion present at a higher level of analysis (in a Schenkerian sense). The melody starts on D in measure 1 and continues down all the way to the D an octave lower in measure 12:



Figure 7.57: Sun and Moon, measures 1-12. Descending scale shown with large noteheads.

This pattern looks much like a Schenkerian *Urlinie*, although it doesn't qualify as such according to Schenker's strict definition requiring motion starting on scale degree 3, 5,

or 8, since this passage starts and ends on scale degree 2 and skips the note E entirely. In any case, the downward motion is very salient, and, importantly, measures 7 and 8 both start with the leading tone, B, which sounds unstable. The A in measure 9 provides a much more stable point of repose. (The program does not have access to the chord structure in the musical score, but that, too, would corroborate this analysis: the first beat in each of measures 7 and 8 is harmonized with the minor iii chord, E minor, while measure 9 has a more stable-sounding IV chord, F major.) Because of the long duration of the note A, measure 9 sounds very obviously (to human ears) like the end of a group. Musicat is aware of note stability, but in this example the important issue, with respect to grouping, is the stability of only those notes that are part of the high-level linear-descent structure. We need to focus on the notes B and A in these measures, temporarily ignoring the other notes, in order to hear how the linear structure and note stability strengthen this grouping choice. (Larson's multi-level Seek Well model, discussed in Chapter 2, deals with these sorts of issues, whereas Musicat, in its present state, does not.)



Figure 7.58: Sun and Moon, first half (run 1), medium detail.

At the end of the run, the program remains unhappy with measures 9–10, but unfortunately it never resolved the problem; the program seems confused by the final 4 measures. Although it seemed unlikely that things would get better, I tried several more runs just in case.



Figure 7.59: Sun and Moon, first half (run 2) (final processing).

In a second run, the grouping structure was still a problem. Larger meta-groups were formed, combining the shorter 2-measure and even 4-measure groups in not-toounreasonable ways, although these larger groups are hard to make sense of since the lowerlevel groups inside them seem wrong. In this run, many more analogies were discovered. The four successive green analogies from run 1 are present, and in this run many more green analogies were also formed, linking distant groups instead of just neighbor groups. A large red analogy was also formed during the final listening stage, although it disappeared at the very end of the run. The next figure shows the final state:



Figure 7.60: Sun and Moon, first half (run 2), medium detail.

There are many green analogies, and this is unsurprising because each of the two measure groups in measures 1–8 is so similar in terms of rhythm and melodic contour. The final 4 measures have a few analogies to earlier structures, but overall they seem not as strong; the presence of the whole notes in those final measures seem to confuse the program. (If the whole notes were heard as the final notes of 3-measure groups, the analogies involved would be much clearer, I believe.)

A third run turned out similar to the previous ones, so it is not shown here. The 3measure groups never formed. Now let's turn to the second half of this melody. Sun and Moon, Second Half



Figure 7.61: Sun and Moon, second half (measures 13-25, renumbered as 1-13 for simplicity).

Notice that I will refer to the measures in this second half as measures 1-13, for ease in counting. This half of the melody is best grouped, in my own listening, using a perfectly regular binary grouping structure for the first 8 measures, followed by the 5-measure group (9–13), or else a 3-measure group, (9–11), followed by a 2-measure group, (12–13), and then a meta-group, (9–13), containing both of these groups. Since Musicat didn't recognize any 3-measure groups in the first half of the melody, I don't expect odd-length groups in this half either, but I wanted to give it a chance.



Figure 7.62: Sun and Moon, second half (run 1).

As I had suspected, no 3- or 5-measure groups formed here, but there was a slight surprise: one odd-length group did form — a 1-measure group, (13)! Why did this happen? It turns out that before the last measure was heard, Musicat had generated an expectation for another 2-measure group, (13–14).When measure 13 arrived, it became part of the expected group, but the rest of the expected group never came. However, this nascent 1-measure group ended up getting linked to group (11–12) by an analogy, $(11–12)\leftrightarrow(13)$, and the strength of this analogy helped support the single-measure group (otherwise, it would likely have been destroyed). The analogy was strong-enough to form because Musicat saw measures 12 and 13 as rhythmically similar. Additionally, measure 13 consists of the stable tonic note C, and Musicat gives a strength bonus to analogies in which the right-hand side has a more stable ending than the left, because this is sometimes a clue to an antecedent \rightarrow consequent relationship between the two sides.

While I was glad that the final measure participated in some sort of group, the grouping structure of the final 5 measures was still disappointing. Measures 9-11, in particular, form something like a sequence, although it is not a strict sequence of the type Musicat can detect. If we consider only the *first* note in each of these measures, we see a descending sequence (A–G–F). The rest of the notes in each measure, however, behave differently, and do not move down by transposition along with the first notes of the sequence. This passage, then, is again one that suggests the utility of a Schenker-type extraction of a higher-level structure, just as I referred to in discussing the first half of the melody.

On a positive note, however, Musicat did form the expected grouping structure for the first 8 measures of this melody. There are small green groups of 2 and 4 measures, all enclosed in a red 8-measure meta-group. Two green analogies have formed that seem reasonable, although it would have been nice to see a larger-scale analogy such as $(1-4)\leftrightarrow(5-8)$; after all, the initial melodic ascent for 4 measures, followed by a 4-measure descent, is reminiscent of the ascent–descent pair of groups at the start of "On the Street Where You Live", so I expected this analogy to form. And in fact, it *did* form, momentarily, during the run...



Figure 7.63: Sun and Moon, second half (run 1), after measure 11.

... but unfortunately, it disappeared shortly after it was created. I tried several more runs, hoping for a possible 3- or 5-measure group to be found at the end of the melody or for the larger analogy (the green one in the figure above) to persist to the end of the run.



Figure 7.64: Sun and Moon, second half (run 2).

On a second run, the grouping structure of the first 8 measures was the same, but an analogy linking the *second* ascending group to the *first* descending group appeared this time: $(3-4) \leftrightarrow (5-6)$. Unfortunately, the large analogy did not form at all in this run.

In the final measures, however, we see a different picture than last time. The metagroup (9–12) has formed, which is a reasonable group (although I wish Musicat were able to hear it as something like a sequence, even though I know it's not possible with the current code). The final measure 13 has again been perceived as a 1-measure group, but this time it is part of the analogy $(7-8)\leftrightarrow(13)$, which is more interesting than the analogy $(11-12)\leftrightarrow(13)$ in the previous run. Measures 7–8 sound like the end of a phrase, and measure 13 does as well. Of course, if measure 13 were part of the hypothetical group (12-13), this would be an even stronger connection. Still, it's encouraging that the program heard an analogy between these two groups, despite their different lengths. Ideally, this analogy would have been a strong enough force to make Musicat shift the position of the thick bar line to fall *after* measure 13, instead of before it. This, in turn, might have helped the program to make sense of the grouping structure of those final measures.

One final positive feature of this run is the structure of the blue measure links. Recall that these links fade over time as the program focuses more on current and very recent measures. Thus, in a picture of the program's final state, links between two early measures usually will have faded out, although links involving at least one recent measure can remain strong. In this run, this has occurred as always, but we see a plethora of links from the recent measures extending back in time to previous measures, forming a distinctive shape when we look at links emanating from measure 10 and linking back to earlier measures. What does this shape indicate? It shows simply that Musicat has heard measure 10 as rhythmically-similar to almost every preceding measure! (Why measure 11 doesn't have just as many links is simply a feature of the stochastic nature of the program; for some reason, Musicat paid more attention to measure 10 during this run.) This large collection of links is gratifying because, as I mentioned earlier, the rhythm here (the one in measure 10, and also in measure 9 and 11 and many others) is one of the characteristic rhythms in the whole score to *Miss Saigon*; it's encouraging that the program has noticed all these similarities in this single 13-measure segment.



Figure 7.65: Sun and Moon, second half (run 3).

This is a final run on the same melody. Although the program never found the 3- or 5-measure groups I hoped for, this single run contains many of the best features from the previous two runs, and it finally did create and stick with the large red analogy $(1-4)\leftrightarrow(5-8)$. The 1-measure group formed again, along with the link to the other phrase ending, (7-8). Finally, in this run we see even more of the blue measure links connecting early measures to the measures in the range 9–11, all of which exhibit the characteristic rhythm discussed earlier. In short, Musicat has heard many of the things I hear in this melody, but it has totally missed the higher-level linear pattern that makes measures 9–11 sound something like a sequence, and it has not heard the final 5 measures as part of a single group (or meta-group).