CHAPTER SIX

Musicat Listens to Simple Melodies

This chapter demonstrates Musicat listening to the most basic type of melodies we had in mind when we first envisioned the program. These are rather short melodies with structures that are easy to understand for a human listener; many of these melodies are children's songs. Of course, music listening is hard, as was discussed in Chapter 1, and even simple melodies like these can exhibit complex and elusive structures. In short, even the very simple melodies discussed here constitute enormous challenges.

A key feature of Musicat is that, like other programs using the FARG architecture, it follows different pathways and comes up with different results on different runs. Thus, it is important to see how the program behaves on multiple runs. I showed multiple runs on some of the Bad Melodies in the previous chapter, but in this chapter it will be even more important to show how Musicat behaves on different runs. It would not be sufficient for me to simply show the results from a single run for each melody, because you, the reader, would not be assured that I had not simply picked the "best" run from a dozen or a hundred different runs. Moreover, there is not, in general, a single "best" way of listening, just as in the Copycat domain there is no single "correct" answer to a given analogy problem. Musicat notices different musical features and forms different groups and analogies on different runs. We hope, however, that obvious and strong relationships will stand out and will be formed by the program on nearly every run. Therefore, for each melody in this chapter, I have presented results from at least two different runs, and I took care to show *typical* results, which give a realistic picture of how the program typically behaves, rather than outlier performances that might give an exaggerated sense of Musicat's "talent" at listening to and understanding music.

TWINKLE, TWINKLE, LITTLE STAR (RUN 2)

This section shows the results from another run of the program on "Twinkle, Twinkle", our sample melody from the start of the previous chapter. However, in this run I will skip the measure-by-measure displays, and will just show screenshots of the program at a few selected points during its listening.



Figure 6.1: Twinkle Twinkle (run 2), measure 4.

After measure 4 has been heard, the expected groups have formed, as has an analogy between them. In the previous run, the program was confused at this point; in this run, on the other hand, everything is proceeding as expected.



Figure 6.2: Twinkle, Twinkle (run 2), measure 9.

After measure 9, many analogies have been formed (notice how this strongly contrasts with the results for the "bad melodies" in the previous section). In the first run of "Twinkle, Twinkle", at this point in the listening performance, the analogies discovered were $(1-2)\leftrightarrow(7-8)$ and $(1-2)\leftrightarrow(5-6)$. Both involved the initial group (1-2). In this run, by contrast, there is only one analogy involving the initial group: $(1-2)\leftrightarrow(3-4)$. But notice how similar *all* of the two-measure groups are after the first two measures: (3-4), (5-6), and (7-8) look virtually identical. Indeed, (5-6) and (7-8) are exactly the same, each of them being a simple transposition of (3-4) up one step. Thus it is not so surprising that in the earlier run, Musicat made an analogy from (1-2) to (7-8), while on this run (1-2) was

mapped to (3-4) instead. — one could imagine swapping most of these measure-pairs around in the melody, and regardless of the order, there are obvious relationships between all pairs. Measure (1-2) is the most unique in this group, so it requires a bit more flexibility in analogy-making to map it onto any of the other groups.

In this run, Musicat finally saw the strongest and most obvious analogy amongst this set of 2-measure groups, which was missing in the previous run: it made the analogy between the two identical measure-pairs (5–6) and (7–8). It also noticed the analogy based on the transposition mentioned above: $(3-4)\leftrightarrow(7-8)$. But why, one might then wonder, didn't it also make the analogy $(3-4)\leftrightarrow(5-6)$? Indeed, it might find this analogy some of the time. However, the program slightly prefers making a correspondence between the two groups ending on measures 4 and 8, respectively, because they both end just before a thick bar line. In other words, it notices how these two groups are situated similarly in terms of the metric hierarchy of the piece; each of them forms the end of a larger structure. Group (5–6), in contrast, is at the start of a large structure, and so there is some justification for Musicat's making the $(3-4)\leftrightarrow(7-8)$ analogy more often than the $(3-4)\leftrightarrow(7-8)$ analogy.



Figure 6.3: Twinkle, Twinkle (run 2), measure 12.

In the figure above, all 12 measures have been presented to the program. In addition to the analogies in place after measure 9 was heard, two new analogies have formed, linking the new groups in the final four measures to earlier groups. The grouping structure has also changed slightly. Group (5-8) has disappeared (or is extremely weak and essentially invisible). The large eight-measure group (1-8) is still in place (despite the missing or weak inner group (5-8)), but the expected final group (9-12) is missing (indeed, since the final four measures are identical to the first four, we might well expect Musicat to create the same grouping structure). However, this listening performance is not yet complete: the program always goes on to "think" about the music it has just heard for the time of several additional

measures. And in this particular run, it turns out that a lot will happen during this final processing time.



Figure 6.4: Twinkle, Twinkle (run 2), measure 12 (+3 measures of extra time).

After three more measures of time have passed, group (9-12) has formed, most gratifyingly, and even more gratifyingly, a new long-distance analogy (colored red) links this new group to the group comprised of the first four measures: $(1-4) \leftrightarrow (9-12)$. This picture is quite similar to the one at the end of the previous run, although in this run there are more

small analogies (colored green). Also, the lack of group (5–8) is conspicuous; that central group was involved in two analogies in the first run.



Figure 6.5: Twinkle, Twinkle (run 2), measure 12(+4).

After one more measure of time has passed, something new has happened: the entire melody has been enclosed in one large group, (1-12). Recall that this didn't happen on the previous run. Moreover, Musicat now expects the entire melody to repeat as a whole: the large purple structure at the right indicates that the program expects this same structure to occur in the next 12 bars. Because the melody has ended, however, the program will soon destroy all expectations automatically.

Many green analogies are visible, although the strongest is the relatively new analogy $(1-4)\leftrightarrow(9-12)$, linking the first four measures to their exact repetition at the end of the melody. This analogy was visible in Figure 4 but it has gained considerably in strength since then, thanks to the extra processing time after the melody's final note "sounded".



Figure 6.6: Twinkle, Twinkle (run 2), end of run.

At last, the listening is complete; the figure above shows the discovered structures at the end of the run. Because there are so many analogies, it is a bit hard to make sense of the picture. Therefore, I will show how this picture varies when I adjust the program's detail level slider, just as I did in the Bad Melodies section. However, in this section the goal is not to identify *weak* structures as much as it is to help make the *strongest* structures more visible.



Figure 6.7: Twinkle, Twinkle (run 2), outer group (very low level of detail).

When I turn down the detail level to a very low setting, the only remaining structure (aside from the measures, bar lines, and notes, which can never disappear) is the group encompassing the entire melody.



Figure 6.8: Twinkle, Twinkle (run 2), big analogy (low level of detail).

Turning up the detail level slightly results in the strong long-distance analogy appearing. This picture helps us understand what Musicat has heard: fundamentally, it has heard the melody as a large 12-measure structure, with two strong internal groups, (1-4) and (9-12), which are connected by a strong analogy.



Figure 6.9: Twinkle, Twinkle (run 2), medium detail level.

Turning up the detail level a bit more results in this figure. The strongest analogies can be seen here; some of these are components of the large analogy from the previous figure. Interestingly, one of these sub-analogies involves the two-measure group (3-4), which has been mapped onto *both* (9-10) and (11-12), while the other important sub-analogy, $(1-2)\leftrightarrow(9-10)$, is too weak to be visible in this view.



Figure 6.10: Twinkle, Twinkle (run 2), medium detail level. Analogy in blue singled-out.

In this figure, I left the detail level at "medium" but have hidden all the green analogies, in order to highlight a curious analogy, $(5-6)\leftrightarrow(9-12)$, which was made between two different-sized structures. Even though the groups involved are of different lengths, the program recognized a strong similarity. This is not too surprising; as I mentioned earlier, (5-6) is a transposition of the melody in (3-4), or, equivalently, its copy at the end of the melody (11-12).) The figure shows (with a light blue line) that the fundamental reason for this analogy is indeed the mapping $(5-6)\leftrightarrow(11-12)$.





Figure 6.11: Row, Row, Row Your Boat.

The "Row, Row, Row Your Boat" melody presented a surprisingly significant challenge to Musicat: when we first gave this melody to an earlier version of the program, the results seemed meaningless. In this section I will present the results of five runs (using the current version of the program) to illustrate various things that the program notices, and fails to notice, when listening to the melody. I encourage the reader to stop for a moment, to hum or sing this melody, and to reflect on what groupings and analogies come to mind, before proceeding.

A comment on the figures in the rest of this chapter: from this point forward, all figures will use the highest detail setting unless otherwise noted. Also, because of space constraints in the user interface (recall that these pictures are simply screen captures of the program as it runs), the top of the display is often cut off for long melodies. For example, we cannot see the very top of the blue measure-link arcs in many of the figures below, but since these are the least significant structures generated by the program, that lack should not affect our understanding at all.



Figure 6.12: Row, Row, Row Your Boat (run 1).

This figure shows the state of the program after listening was complete for the first run. In this picture, we see that a great deal of structure has been perceived. Every successive pair of measures has been chunked into a group. The beginning and the end of the melody also exhibit 4-measure meta-groups. All the groups seem reasonable, although I expected to see more higher-order groups (every successive group of four measures sounds like a group to me).

The first and strongest analogy in the figure, $(3-4) \leftrightarrow (5-6)$, is interesting and insightful. The program has noticed a contour relationship between the two groups

(indicated by the small pink arc with endpoints that are aligned horizontally with the centers of the groups involved). It is easy to see this relationship: the melody in (3-4) is moving up by steps (C–D–E), which is similar to the upwards motion of (5-6) if we ignore the first note in (E–D–E–G). The analogy is also based on the similar rhythmic pattern of the two groups: (5-6) starts with the same rhythm as (3-4) and then is extended by the addition of one more note (the quarter note at the end of measure 6). To put it in terms of the lyrics, the rhythm for the words "row your boat" is essentially reused with the following words, "gently down". But whereas the word "boat" was set to a dotted half note, the word "down" is set to a regular half note, making room for the word "the" on a quarter note. The reuse (with slight modification) of this rhythm gives structure to the melody and propels it forward. I will have more to say about this analogy in run 4 below.

One of the most salient features of this melody to me is that I hear the notes of "gently down the stream" (5–8) as rising-upwards to the climax on "stream", echoed by the later descent of "life is but a dream" (13–16). That is, I hear the analogy (5–8) \leftrightarrow (13–16). In this run, the individual two-measure groups are members of analogies, just as I would expect: (5–6) \leftrightarrow (13–14) and (7–8) \leftrightarrow (15–16). However, there is no four-measure group (5–8), just two smaller groups, and thus no single analogy (5–8) \leftrightarrow (13–16). In other words, Musicat *almost* formed the main structure that I myself hear and that I had hoped it would find, but it fell slightly short.

There was one more analogy made in this run: $(11-12) \leftrightarrow (13-14)$, but this one is harder to explain. I had expected measures 9–12 to be involved in a different way, because their rhythm (all quarter notes) is so distinct from other parts of this melody. The very low happiness values for measures 9–10 indicate that the program was aware that it had missed something in that section, but that problem was not fixed in this run. Let's examine a second run to see what the program does differently.



Figure 6.13: Row, Row, Row Your Boat (run 2).

In this run, the same groups were formed as in the previous run, but, curiously, the analogies are completely different. Ironically, the strongest analogy, $(9-10)\leftrightarrow(11-12)$, involves the measures that were ignored in the first run. This is the part of the melody with the lyrics "merrily, merrily, merrily, merrily", so not only are all four measures identical in

terms of rhythm and pitch contour (each measure has a flat contour), but also, if we know the words, we hear the exact repetition of the word "merrily". Now, I think I hear this section as four measures with three notes each, where the melody is descending through a C-major arpeggio, and so my hearing differs from Musicat's hearing of two measures that are mapped onto the following two measures. However, the analogy found by the program is a justifiable way of hearing the passage, especially if one thinks hypermetrically, with measures 9 and 11 being more heavily stressed, and measures 10 and 12 being softer.

The program found two more analogies. Analogy $(1-2)\leftrightarrow(3-4)$ is based on a rhythmic similarity between the two groups: each measure starts with a long note on the downbeat, and in measures 1, 2, and 4, this is the only note in the measure (a dotted half note). Only in measure 3 is the rhythm slightly different, with the first note of the measure shortened and a quarter note added on beat 3. Also, measures 1 and 3 both start with the note C, strengthening the analogy.

The final analogy maps two groups at quite a long distance: $(3-4) \leftrightarrow (13-14)$. This is slightly different from the analogy found in the previous run, $(5-6) \leftrightarrow (13-14)$. But (3-4)and (5-6) are quite similar, so it is not surprising that the program can discover both of these analogies. Indeed, both analogies make intuitive sense. (Analogies are not mutually exclusive in Musicat, but for some reason it didn't find both of these in the same run.) Ignoring the program's results for a moment and considering my own listening, I can easily imagine hearing an analogy between the first four and the last four measures of the melody. In this analogy, the first two measures would map onto the long, tied note at the very end of the melody, while the inner measures would be related to each other: the ascending line ("...row your boat") in measures 3-4 is echoed at the end of the melody by the descending line ("life is but a...") in measures 13-14. Incidentally, this imagined analogy is reminiscent of simple letter-string analogy we might make in the Copycat domain: in the mapping $abc\leftrightarrow cba$, the letter **a** at the *start* of the sequence is mapped onto the **a** at the *end* of the sequence, and the *ascending* pattern $a\rightarrow b\rightarrow c$ is mapped onto the *descending* sequence $c\rightarrow b\rightarrow a$. In the current run of "Row, Row, Row Your Boat", Musicat did not create the entire analogy $(1-4)\leftrightarrow(13-16)$ that I expected, but it did make the analogy $(3-4)\leftrightarrow(13-14)$, which is half of my desired analogy — it maps the initial ascending line to the final descending line.



Figure 6.14: Row, Row, Row Your Boat (run 3).

In a third run, many of the analogies created in the first two runs were created again. In addition, all of the groups present before were created in this run, along with two additional groups: the very strange meta-group (7-10), which awkwardly stretches across the thick bar line at the very center of the melody, and the large, strangely-sized group (1-12).

A new analogy in this run was $(1-2) \leftrightarrow (7-8)$, mapping the first notes of the melody onto the long note at the end of the first half of the melody. This analogy is followed by the analogy $(7-8) \leftrightarrow (15-16)$, which was created in the first run as well. In this run, it makes a nice, cohesive picture to see both of these analogies together, because the analogies form a sort of sequence: $A \leftrightarrow B$ followed by $B \leftrightarrow C$ (the program, however, doesn't have the ability to hear a relationship such as this one between two analogies). This pair of analogies helps to emphasize the large-scale structure of the melody: the opening notes are developed in such a way that they lead up to the cadence in measures 7–8, these first 8 measures form an antecedent phrase, and then the next 8 measures are a consequent phrase that ends on a cadence that is related to that which formed the end of the antecedent phrase. (The melody also has a stereotypical arch form: it rises in pitch to the middle of the melody and then falls back down to a stable base at the end of the melody. This idea about the arch shape is not noticed by Musicat, although the relationship $A \leftrightarrow B \leftrightarrow C$ contains the germ of this idea, because A and C consist of notes on the tonic (the note C), whereas the B material is on the dominant (the note G).



Figure 6.15: Row, Row, Row Your Boat (run 4).

This fourth run is also similar to previous runs, with a slightly different collection of analogies in place at the end of the run. This particular run gives a good picture of the types of relationships typically heard by the program, and, I believe, sheds some light on how people may hear this melody. I have already discussed most of the analogies present, but for a moment, reconsider the initial analogy $(1-2)\leftrightarrow(3-4)$, as well as the three strongest analogies in the figure above: $(3-4)\leftrightarrow(5-6)$, $(9-10)\leftrightarrow(11-12)$, and $(3-4)\leftrightarrow(12-13)$. Taken together, these four analogies tell a nice story and are evidence of a good listening performance

The initial analogy indicates that we hear a connection between the first and the second pairs of measures. (The grouping structure also indicates another type of connection; each of these measure pairs, 1–2 and 3–4, forms a connected unit, but in this discussion I focus on the analogies rather than the groups created.) Already, this contrasts greatly with the way the program listened to the random Bad Melody #1, in which no strong analogies or groups were heard.

The second analogy, $(3-4) \leftrightarrow (5-6)$, indicates that the third small group (the third pair of measures) was heard as similar to the second, and the strength of the analogy shows that this relationship was quite salient. Now, it is important to notice what has happened: group (1-2) was mapped onto (3-4), which was itself mapped forward onto (5-6). This chain of two analogies linking three groups is striking in how it indicates the cohesiveness of the components in this melody, even though as the melody proceeds from measure 1 through 6, the rhythmic speed is increasing. Notice how different group (1-2) appears from (5-6). No *direct* analogy has formed between these two groups in this run (a weak analogy between them does form in run 5, below), but they are nevertheless linked indirectly through the analogy-chain, with the intermediate group (3-4) acting as a stepping-stone.

But what is the significance of this little chain of analogies? I claim that, taken together, these two analogies indicate one of the critical features of this melody that make it sound more *melodic* than the Bad Melodies earlier in this chapter: it has been constructed based on the principle of motivic development. Arnold Schönberg is best known as the father of *atonal* music, but he also taught *tonal* music theory and composition. His book *Fundamentals of Musical Composition* (1967) instructs students to build melodies by beginning with a short motive and developing it through a series of transformations as the melody moves forward. His book catalogs many types of such transformations and provides

many examples of how famous classical melodies can be heard in terms of these transformations. The little chain of analogies $(1-2)\leftrightarrow(3-4)\leftrightarrow(5-6)$ made by Musicat on this run shows that it sometimes hears these six measures in terms of this sort of motivic development, and this analysis is consistent with how Schoenberg suggests that we understand simple melodies.

Continuing to an analogy later in the melody, we find $(9-10) \leftrightarrow (11-12)$, which I discussed earlier. Notice how, in this hearing, the first six measures are linked together, followed by two measures which (unfortunately) were not heard as part of a larger group. Then, starting in measure 9, we encounter new material, and this analogy indicates how all four of these measures, 9–12, are heard in a linked-together manner.

The final strong analogy, $(3-4) \leftrightarrow (12-13)$, also was discussed earlier. But in this hearing it is particularly instructive to see how, after encountering the contrasting melodic material in measures 9–12, suddenly the program has heard a return to an earlier theme. Measures 12–13 remind the program of measures 3–4, and then the melody ends on a long note (although unfortunately, in this hearing that final note was not linked to something earlier, as it was in run 1).

This run was lacking some elements, but overall it is indicative of a quite reasonable overall way of hearing the melody. The perceived analogies have done a relatively good job of indicating how different melodic ideas are arranged and linked together within the melody. This is also an appropriate moment to remind the reader to consider what the fundamental purpose of the program is. I just stated that this run was "lacking some elements", but our goal is not to model a "perfect" listening performance, as such a thing does not make sense. Human attention and perception are resource-limited and we can only attend to a certain number of things at once and they will not always be the same ones. Musicat's explicit perception of an analogy in one run but not in another reflects this variable attention and the uniqueness of each listening experience. (Of course, extremely salient features of a melody *should* be heard by the program on every run, and Musicat admittedly misses many very salient analogies and group boundaries that a good human listener would never miss.)



Figure 6.16: Row, Row, Row Your Boat (run 5).

In a final run on this melody, the grouping structure is different from that in earlier runs, and a larger analogy has formed than we have yet seen for this melody. A few of the familiar analogies from earlier runs are present, although they seem weaker in this run. Thus I lowered the detail level to show more clearly which structures were strong in this run; see the next figure.



Figure 6.17: Row, Row, Row Your Boat (run 5), low detail level.

When the detail level for this run was reduced to "low", many of the analogies and groups disappeared. The strongest analogy was the large one linking the first two 4-measure groups: $(1-4) \leftrightarrow (5-8)$. It is nice to see that the program has heard all eight measures of the melody's first half as forming a coherent unit (indicated by the green ellipse surrounding all of them), and it is consistent with my hearing that there is an analogy between the two 4-measure groups that make up this opening half of the full melody.

The red group (1-10), however, is a mystery, and it does not seem justifiable as a way of hearing the melody. In addition to this problem, all the smaller structures in the second half of the melody have disappeared at this detail level; they were all weak, indicating that the program did not create any strong structures in the second half.



Figure 6.18: Row, Row, Row Your Boat (run 5), medium detail level.

I raised the detail level to "medium" in order to bring back into the picture some of the structures from the second half of the melody. The analogy $(9-10)\leftrightarrow(11-12)$ is visible again, although it is quite odd that the right-hand edge of the large red group cuts this analogy in half. The final note of the melody is now linked by analogy to the final note of the first half (as in run 1), which seems very reasonable. However, the descending line in measures 13–14 is not connected by any analogy to any earlier material.

This particular run was disappointing, but overall the program tended to find many of the expected analogies in this melody. Run 4 was particularly "good" in that it corresponded well with my own experience of listening to the melody. While the foregoing snapshots of these five runs are no substitute for seeing the program in action, I think that, taken together, they give a balanced picture of what Musicat perceives when it listens to "Row, Row, Row Your Boat".

SUR LE PONT D'AVIGNON



Figure 6.19: Sur le pont d'Avignon.

This short French children's song has a simple structure that is visible just by looking at the music in the figure above. The first four measures are aligned above the second four measures, making it easy to see that the two halves of the melody are nearly identical. The only difference is in measures 4 and 8. Measure 4 ends with a half-cadence on an implied G major chord (the dominant chord in C major), whereas measure 8 ends on C with an implied tonic chord. The program listens to melodies with no chords provided, but behindthe-scenes, it imagines chords implied by the melody, much as people do. These chords are not shown here, but they may help the program understand the melody, especially in obvious cases such as this one. A very simple implied harmony that we might imagine (with no prior knowledge of which chords are typically used to accompany this melody) is given below, with one chord per measure (except for the final measure, which has two chords):

$$C G C G C G C (G - C)$$

This melody is simpler than "Row, Row, Row Your Boat", so I consider just two different listening performances below.



Figure 6.20: Sur le pont d'Avignon (run 1), just after hearing the final measure.

Musicat has heard the melody as having a completely regular binary grouping structure, and this grouping makes intuitive sense. Moreover, before the program realizes that the melody is over, it expects the grouping structure for the whole melody to be repeated exactly once again (shown in purple to the right).

To make the structure even clearer, I modified the detail level to produce the next figure (I also waited for the run to complete — Musicat continues "listening" for a little while after the last note arrives — but the picture did not change significantly).



Figure 6.21: Sur le pont d'Avignon (run 1), low detail level.

This figure shows the strongest structures heard in this run. The melody has been divided into two nearly-identical halves — the program has noticed this obvious property of the melody, which I mentioned above — and a strong analogy (shown in red) has been made between the two halves. The analogy is supported by a variety of contour relationships shown in thinner lines in the figure.



Figure 6.22: Sur le pont d'Avignon (run 1), medium detail level.

Moving the detail-level slider up to a medium setting makes some green analogies and groups visible. Each half of the melody is itself divided in half again (each red metagroup contains two smaller green two-measure groups).

Turning to the green analogies, we see two analogies that cover nearly the same material: $(1-2)\leftrightarrow(3-4)$, and the very similar $(5-6)\leftrightarrow(7-8)$. I say that these two analogies are very similar to each other because the two halves of the melody are almost identical. Why did the program form these analogies? It is easy to see that measure 1 is similar to measure 2 or that measure 3 is similar to measure 4, but the relationship between measures 1 and 3 or 2 and 4 is less obvious. The light green lines in the figure linking measure 1 to measure 3 and measure 2 to measure 4 indicate that Musicat found *rhythmic* relationships between these pairs of measures. Measure 1 differs rhythmically from measure 3 only by one quarter note, and the same holds for measures 2 and 4. Because rhythmic relationships are very salient for Musicat, it makes these analogies, even though the pitch contour of measures 1 and 2 is very different from that of measures 3 and 4.

The third analogy in the picture, $(3-4) \leftrightarrow (5-6)$, actually follows the same logic. Indeed, because measures 5 and 6 are identical to measures 1 and 2, this is the same analogy as the two just discussed, except that the temporal order of the two groups involved has been reversed.

This series of three green analogies in a row is reminiscent of the analogies at the start of run 4 on "Row, Row, Row Your Boat". However, they are less interesting to me in this case because instead of indicating a true motivic development in the melody, here they merely indicate that Musicat has noticed that the first two groups (1-2) and (3-4) are rhythmically similar, and that the whole first half was repeated. The larger red analogy in this picture seems more significant to me, and is more in tune with how I hear this melody.



Figure 6.23: Sur le pont d'Avignon (run 1), high detail level.

Returning the detail slider to the "high" level, we notice that one more green analogy is visible: $(3-4) \leftrightarrow (7-8)$. This analogy has formed even though the two groups involved end in slightly different ways. In this case, Musicat probably noticed that the left-hand group ended on the dominant note (and dominant chord), whereas the right-hand group ended on the tonic note and chord. The dominant—tonic relationship present in this green analogy most likely strengthened the large red analogy: Musicat is aware of this sort of simple dominant-to-tonic progression.



Figure 6.24: Sur le pont d'Avignon (run 2), just after hearing the final measure.

A second run on this melody resulted in a very similar picture. Two big differences here are that the middle analogy $(3-4)\leftrightarrow(5-6)$ is not present. Instead, the analogy $(1-2)\leftrightarrow(5-6)$ has formed (although it is quite weak at this point).

Recall that the program keeps listening for several more measures of time after the last note is heard. In this case the extra time made a difference in structure strength. The next figure shows the program's state at the conclusion of this extra listening time.



Figure 6.25: Sur le pont d'Avignon (run 2), final state.

Notice that the analogy $(1-2)\leftrightarrow(5-6)$ is stronger now, compared with the previous view. Indeed, the only weak structures are the blue measure links and the "next-door-neighbor" analogies $(1-2)\leftrightarrow(3-4)$ and $(5-6)\leftrightarrow(7-8)$. These disappear when the detail level is turned down, as is shown in the next figure.



Figure 6.26: Sur le pont d'Avignon (run 2), low detail level.

At a low detail level, only the two strong long-distance green analogies and the large red analogy remain. Again, notice how much this situation contrasts with the one earlier involving the Bad Melodies, in which nearly all structures disappeared at low detail.



Figure 6.27: Sur le pont d'Avignon (run 2), medium detail level.

This medium-detail view shows the important structures formed in this run. I think that this listening performance was slightly better than the previous run; the long-distance analogies $(1-2)\leftrightarrow(5-6)$ and $(3-4)\leftrightarrow(7-8)$ are salient and important to a human listener.

Even though this run seems like a good performance, Musicat still missed some important features of this melody that human listeners would likely notice. Most significantly, there are stepwise relationships between measures that are not explicitly noticed by the program, but which contribute significantly to the smoothness and singability of the melody line. For instance, the first measure consists of three C's while the second measure has three D's. The program can recognize these measures as having the same rhythm and pitch contour, but it would not affect Musicat's hearing much if each D in the second measure were replaced with E or even F. The relationship of a single scale step between measures 1 and 2, however, helps a human listener to group the two measures together. More importantly, the interesting melodic pattern in measure 3 is repeated, with transposition, in measure 4. Musicat may notice this transposition and repetition (although it does not currently display these things with any special notation, such as the blue zigzaggy lines used to indicate sequences of three or more items related by transposition), and when they are noticed, Musicat is more likely to form a group consisting of two measures involved. However, the program would not notice the smooth stepwise connection between the final note of measure 3 (C) and the first note of measure 4 (B). A stepwise connection is also present at the end of the melody between measures 7 and 8, although in this case the C is connected to the D one step above. These types of melodic connections are very important in human listening, but unfortunately the present version of Musicat does not notice them.

FRÈRE JACQUES





Our next example melody is another French folk tune (but this one also is well known to English speakers: "Are You Sleeping, Brother John?"). A glance at the music above reveals that it also has a simple structure, but this structure is different from that of "Sur le pont d'Avignon". In that melody, the second half was a nearly exact repetition of the first half. In "Frère Jacques", however, every single measure is followed by an exact repetition. Instead of having a block of four measures repeated one time, we have four single measures, each immediately repeated. To make this point clear, I re-composed "Frère Jacques" to follow the formal pattern of "Sur le pont d'Avignon":



Figure 6.29: Frère Jacques d'Avignon.

These two melodies look and sound quite different. How will Musicat's listening performance on Frère Jacques differ from its perception of "Sur le pont d'Avignon" in the previous section? (Incidentally, I have not run Musicat on the re-composed melody, although I would guess it would hear it in a very similar way to the way it heard "Sur le pont d'Avignon".)



Figure 6.30: Frère Jacques (run 1).

This run produced the same grouping structure as we saw in "Sur le pont d'Avignon", but no analogies were found. Musicat does indeed hear these melodies in a different way (at least in this run).

Although no analogies were found in this run, it is worth reminding the reader that the blue lines above the figure also indicate relationships between measures, and we can certainly call these "analogies" as well. For technical reasons described in Chapter 9, these simple relationships between measures, based solely on rhythmic similarity, are handled separately from the more complex analogies drawn below the measures. However, in this melody there is, of course, plenty of analogy-making at work, as always (recall the discussion of the pervasiveness of analogy in Chapter 4).

Because this melody consists of exact repetitions of each measure, the most obvious analogies are the relationships between adjacent measures. Analogies drawn below the measures in Musicat usually span more than two measures, but the blue measure links above show that Musicat has noticed the rhythmic repetition. For instance, the blue arc between measures 5 and 6, just like the one between measures 7 and 8, indicates exact rhythmic repetition. Musicat almost certainly created arcs such as these between measures 1 and 2 and between measures 3 and 4 during the run, but by the end of the run they have faded away. Those arcs, however, likely helped to create all the groups consisting of measure pairs in the melody: (1–2), (3–4), (5–6), and (7–8). In summary, these repetitions were not unnoticed by the program, even though that fact is not obvious in the figure above showing the program's final state.



Figure 6.31: Frère Jacques (run 2).

A second run resulted in the same *grouping* structure, but in addition to groups we now have three green *analogies*. This picture is slightly reminiscent of the pictures of "Sur le pont d'Avignon", but there is no large red analogy linking the first half of the melody to the second half. Also, two of the green analogies are weak. The next figure uses the detail level to highlight the strongest structures:



Figure 6.32: Frère Jacques (run 2), strongest structures.

At this low detail level, only one green analogy remains, and some of the groups have also disappeared. Notice that this is a surprising analogy, because it is between differentlysized structures: $(3-4)\leftrightarrow(5-8)$. The thin green arc indicates that the analogy is partly motivated by a rhythmic relationship: the two measures 3–4 have the same rhythm as 7–8. The pink arc indicates some sort of non-rhythmic relationship, which most likely is the dominant→tonic relationship implied by the left-hand group (3–4) ending on the dominant note G and the right-hand group (7–8) ending on the tonic, C.

This run, like the previous one, resulted in different behavior from runs on "Sur le pont d'Avignon". However, because some of the analogies were similar in both performances, I ran it one more time to verify that Musicat hears these two melodies in a significantly different way.



Figure 6.33: Frère Jacques (run 3).

On the third run, a large red analogy formed, and again we have the same group structure as before. Did the program hear this melody as if it were equivalent to "Sur le pont d'Avignon" after all? Or are the analogies formed here still reasonable ways of hearing musical similarity, even though the melody is not constructed as a large-scale repetition of four initial measures? The analogy $(3-4) \leftrightarrow (7-8)$, at least, looks very reasonable. I lowered the detail level to remove the green analogies from the picture so I could investigate the red analogy more closely.



Figure 6.34: Frère Jacques (run 3), low detail.

This analogy connects the two halves of the melody, just as in "Sur le pont d'Avignon", run 1. However, there are far fewer supporting components of this analogy in this run on "Frère Jacques" than there were in that run on "Sur le pont d'Avignon": there are just two pink arcs (left and center), indicating non-rhythmic relationships, and one light red arc (right) indicating a rhythmic relationship. This rhythmic relationship is the same one discussed in run 2; that is, both halves of this melody end with the rhythm QQH QQH. The possible domaint \rightarrow tonic relationship between the two groups was also mentioned above, and it may be part of the analogy here (unfortunately, although these details are available to the user of the program while it is running, they are not visible in the pictures). Another possible source of this relationship is a contour relationship between the first two measures of each half of the melody — that is, between measures 1–2 and measures 5–6. Specifically, each of measures 1 and 2 has a rising and then falling contour, in which the upwards motion is by *step* and the downwards motion is by *leap*. Measures 5–6 have a different number of notes but the contour is related to that of measures 1–2: there is a one-note stepwise ascent followed by a stepwise descent and finally a downwards leap. Musicat considers these

contours to be just slightly similar; human listeners may even hear them as more similar than the program does.



Figure 6.35: Frère Jacques (run 3), medium detail.

This figure shows the final state with the detail level turned back up to "medium". Turning our attention to the first green analogy, $(1-2) \leftrightarrow (3-4)$, we see that this analogy looks similar to the first analogy in "Sur le pont d'Avignon". Perhaps surprisingly, the first four measures of "Frère Jacques" are quite similar to the first four of "Sur le pont d'Avignon". In both melodies, measure 1 has the same rhythm as measure 2. Likewise, measures 3 and 4 share a rhythm. In "Frère Jacques" the first rhythm is four quarter notes (QQQQ), followed by the exact repetition, and then QQH (and another QQH). In "Sur le pont d'Avignon", these same two rhythms are also involved, just in a different order: QQH QQH QQQQ QQQQ. "Frère Jacques" features exact pitch repetition, while "Sur le pont d'Avignon" uses transposition of the pitches in the melody, but this is a minor difference. It is not surprising, then, that this large-scale analogy is found by Musicat in both melodies.

In sum, this listening performance by Musicat indeed makes sense; perhaps these two melodies are more similar than it first appears. The nature of the large red analogy, however, is quite different for the two melodies. In "Sur le pont d'Avignon", the analogy was due to a near-exact repetition of the two melody halves. In this run of "Frère Jacques", on the other hand, it was due to a few different factors, including rhythmic repetition and contour similarity.

SICILIENNE (FAURÉ)



Figure 6.36: Fauré's Sicilienne for cello and piano, op. 78.

The first four measures of this short excerpt from the Sicilienne for cello and piano by Gabriel Fauré look very similar to the last four measures, just as was the case with "Sur le pont d'Avignon". Thus we would expect to hear the strong analogy $(1-4)\leftrightarrow(5-8)$ in this melody. However, the analogy is slightly subtler here: in "Sur le pont d'Avignon", the first three measures of the first half were exactly the same as the corresponding measures in the second half. In this Sicilienne, however, the two halves have the same rhythm but different pitches. The melody *contour*, however, is the same in the two halves. Will Musicat recognize the similarities?



Figure 6.37: Sicilienne (run 1).

In this first run, the program *almost* found the analogy I expected. In the final output (Figure 37), there is no large-scale analogy linking the two halves of the melody, but a reasonable grouping structure has appeared, and all the components of the two halves have formed the expected analogies. That is, (1-2) was mapped onto (5-6), and (3-4) was mapped onto (7-8). Additionally, the weaker analogy $(3-4)\leftrightarrow(5-6)$ has formed because of some perceived rhythmic similarities.

In the previous paragraph, I stretched my notation in order to accommodate the presence of upbeats (or pickup notes) in the melody: measure "1" does not indicate the first full measure after the thick bar line, as one might expect, but rather it stands for the first *measure-length segment* of the melody line (*i.e.*, in this case, the first 6-eighth-notes time span in the melody, starting on the low eighth note E and ending just after the quarter note E an octave higher). Likewise, groups always start and end just before final eighth-note upbeats

(groups are composed of measures and must also accommodate upbeats.) In the figure above, the ellipses representing groups include the pickup notes and hence are shifted to the left.

That Musicat accommodates the upbeats in this melody should not be seen as a sign of sophistication; to the contrary, it forces groups to include upbeats in a completely mechanical way. For melodies having an upbeat at their start, such as this one, Musicat simply assumes that all melodic phrases it hears will have a regular structure in which every group formed will start or end just before the metric position indicated by the first upbeat. Also recall that groups in Musicat always have a duration that is an integral multiple of the measure length. A more flexible version of the program would relax these constraints and would allow group boundaries to occur anywhere within measures (this would also be necessary for the program to make sense of melodies having internal time-signature changes). Chapter 8 has a discussion of a previous version of Musicat that did feature this flexibility in grouping.



Figure 6.38: Sicilienne (run 2), final state.

At the very end of a second run, the picture is similar (Figure 38). Analogy $(1-2)\leftrightarrow(5-6)$ did not form this time, although a different analogy, $(5-6)\leftrightarrow(7-8)$, formed. The desired big-picture analogy, $(1-4)\leftrightarrow(5-8)$, however, is not visible. But let's take a look at the final stages of the run before it completely stopped listening:



Figure 6.39: Sicilienne (run 2), measure 8(+1).

After one measure of "post-melody" listening time had passed, only one (very weak) analogy was visible...



Figure 6.40: Sicilienne (run 2), measure 8(+2).

...but after one more measure, our desired large-scale analogy was formed: $(1-4)\leftrightarrow(5-8)$, colored red in the figure.



Figure 6.41: Sicilienne (run 2), measure 8(+5).

After a few more post-melody measures, more green analogies had formed, and the program had generated a large-scale purple expectation. However, in the program's final state, just moments later, the big red analogy suddenly disappeared:



Figure 6.42: Sicilienne (run 2), final state, medium detail.

Why did the big analogy disappear? Viewing the final state at a medium detail level (as well as taking a closer look at the figure that included the big analogy) yields an

explanation: the root cause was that group (1-2) was too weak. Even though the group was formed, it was still weak (notice its light color in the high-detail figures earlier). Because (1-2) was weak, analogy $(1-2)\leftrightarrow(5-6)$ did not form during this run. In Musicat, an analogy is strongest when all of its *components* (all the groups on its left- and right-hand sides) are involved in the mapping. In the figure above, where the large analogy did form momentarily, we can see that the group (1-2) was not involved in the mapping, and hence the big analogy was too weak and was destroyed.

One of the behind-the-scenes difficulties in this case is that group (1-2) must be formed and must achieve a high-enough strength value fairly early in the run. Musicat, remember, focuses its perceptual energy on the most recent several measures of music that it has heard. If, as in this example, a structure formed at the start of a melody is not strong enough, it is possible that the structure will remain weak when it eventually moves out of the simulated working-memory area ("the mind's ear"). After that, the program cannot go back and re-hear the structure. In this example, group (1-2) simply never was perceived as a strong group, and this weakness prevented the formation of the key small-scale analogy $(1-2)\leftrightarrow(5-6)$, and finally, this missing analogy in turn resulted in a lack of evidence to support the large analogy $(1-4)\leftrightarrow(5-8)$.



Figure 6.43: Sicilienne (run 3).

On a final run, Musicat *did* discover the big analogy between the two halves of the melody, and the analogy remained stable and strong all the way through the very end of the run. But even here, notice that the sub-analogy $(1-2)\leftrightarrow(5-6)$ is still much weaker than $(3-4)\leftrightarrow(7-8)$.

These three runs give a sense for the difficulty of Musicat's domain. To a human listener, the large analogy between the two halves of this melody is so obvious that it seems trivial. Even a non-musician looking at the music notation for this melody would be likely to recognize the similar shape formed by notes in the two halves of the melody. It is frustrating that Musicat still has trouble discovering and maintaining this analogy, but we can see that the analogy's formation depends on many substructures being perceived in a certain way. It may be tempting to pessimistically imagine that Musicat's analogy-formation is akin to building a house of cards, vulnerable and at risk of collapsing if subjected to the slightest breeze. Musicat's analogies are *not* this fragile, fortunately: analogies, sub-analogies, and groups mutually reinforce each other, so that the whole structure can be stronger than the sum of its parts. However, the temporal nature of the domain complicates the picture, as we saw most clearly in run 2 above: if certain structures do not form quickly enough, their nonexistence (or weak existence) may prevent the formation of larger structures that are only perceivable later on. Even though this causes problems, as in run 2, where it was difficult for the program to create the large analogy, it is a natural consequence of the program's modeling of temporal perception. If a human listener, after all, were distracted during listening to the first few measures and failed to form a strong mental representation of the first half of the melody, we would not expect that person to notice an analogy between the two halves. Even the listening performances of the first two runs for this melody, then, may illustrate behavior that is "correct" in that it exhibits the expected sorts of missing large-scale structures after the end of the second half of the melody, given listening "mistakes" it made in early measures.

Even the simple-sounding melodies in this chapter posed formidable challenges to Musicat. The next chapter gives several examples of even more complex melodies, which will pose even more problems for Musicat. These will illustrate the kinds of groups and analogies it can make consistently, as well as indicate more areas for future improvement.