

CHAPTER FOUR

Music and Analogy

Analogies All the Way Down

Douglas Hofstadter has argued that analogy is the core of cognition (Hofstadter, 2001; Hofstadter & Sanders, 2013, forthcoming). His view is that analogy-making is not a fancy cannon that we wheel out only on rare occasions to solve a tricky problem, but rather, that analogy-making happens all the time, involving everything from unconscious fleeting thoughts to conscious deliberation. Not only does making a statement such as “Iraq is another Vietnam” involve an analogy, but so does seeing a chair and mentally classifying it as a chair, or seeing the shapes **Z**, **Ż**, **Ẑ**, and possibly **Z̃**, and recognizing them as instances of the letter Z, or even my nearly unconscious deciding that this very clause should start with the word “or” while I was typing twenty words or so ago. According to Hofstadter, every word choice — even deciding between simple alternatives such as “and” or “or”, “the” or “a”, and so on — involves analogy-making. Analogies occur everywhere in thought and perception.

Given this viewpoint about analogy, it is natural to try to apply it to the particular case of analogy-making in music perception. We can imagine every bit of music perception as analogy-making, from recognizing a simple, primitive sound (a single note or a simple

sound in nature), to hearing a chord as either major or minor, to recognizing the repeating four-note “Fate” motif in Beethoven’s Fifth Symphony, to recognizing an entire performance as a reinterpretation of an older, well-known song, to hearing the style of a Chopin nocturne as similar to the style of a John Fields nocturne. In the following section, I will point out how analogy-making pervades the act of listening to Chopin’s “Raindrop” prelude, Op. 28 No. 15. The examples ascend through ever greater levels of complexity, from small-scale details to large-scale musical forms, and they thereby show how analogy is a fundamental component of music-listening at each hierarchical level. In most cases I will show examples of how two musical items are heard as “the same thing” even when there are real differences between the two items — the listener creates a musical analogy between the two chunks of sound and thus hears them as “the same”.

Ascending Through the Levels

To a music theorist, the note might seem to be the smallest unit of music and thus the logical starting point for an examination of analogy-making on many hierarchical levels. See, for example, Leonard Bernstein’s lectures comparing music to language, which start by likening notes to phonemes, motifs to words, and so on (Bernstein, 1976). However, an “atomic” note is not really the smallest unit after all. Instead, I begin by breaking the note into its constituent parts: the parameters that make up a note, such as its frequency, its timbre, and its perceived pitch, before ascending to higher levels of musical organization such as rhythmic structure, motifs, phrases, and sections.




Figure 4.1: A pure tone at 440 Hz.

SOUND WAVES

The simplest possible musical tone is arguably a basic sine wave, as shown in Figure 4.1. The amplitude corresponds to air pressure in a sound wave, shown here with a frequency of 440 Hz (the note A in modern tuning). If the sound wave starts and then stops a second or two later, we can effortlessly hear this as a note, and even if we don't have perfect pitch part of our auditory system resonates with the particular frequency of the note. For such low-level perception, can we really claim that analogy-making is taking place? The physical structure of the ear causes fluid in the cochlea to vibrate in sympathy with the 440 Hz sound wave, and then hair cells in the inner ear corresponding to this frequency are stimulated, and the frequency and intensity of the sound are transmitted up towards the auditory cortex. One might think "Surely this isn't analogy-making — it's just a biological process." This is a reasonable position, but I claim that even this process of decoding a sound wave into a frequency (or set of simultaneous frequencies, in the case of a more complex timbre) is, in a tiny way, an example of analogy-making.

Analogy-making is simply the process of seeing something as similar to another, previously known thing. In the case of auditory processing, we hear a pitch (as opposed to a noise) when an extremely short-duration waveform is repeating periodically. In Figure 4.1, the sine wave sounds like a pure pitch at 440 Hz because the 1 complete period of the sine

wave shape (a peak followed by a valley) is repeating 440 times per second. Mechanisms in the ear, somewhat miraculously, can detect these 440 repetitions of the same “shape” each second. By detecting this repetition, a mechanical process in the ear is implicitly perceiving a sort of sameness: each of these patterns is like the others. Even though we don’t perceive repetitions or 440-ness in a normal cognitive sense, something inside us is responding to the sameness. This is a trivial and mechanical process, but we might still consider it an example of super low-level analogy-making (we might make the extreme statement that there are 440 analogies per second here, although compared with the much more sophisticated analogies people are capable of, perhaps we should call these micro-analogies or nano-analogies).




Figure 4.2: A composition of two sine waves at 440 and 880 Hz.

Lest the reader think that detecting sine-wave frequency is too trivial to be called even a nano-analogy, consider some more difficult tasks that our auditory system also does automatically. For instance, if we superimpose two sine waves of different frequencies we get a more complex picture (Figure 4.2), yet the ear is still able to detect both frequencies present in this more complicated shape.




Figure 4.3: A small portion of the waveform of the note A played on a piano.

Finally, consider a more complex shape such as that produced by a piano playing a single note (Figure 4.3). The signal doesn't repeat exactly, but approximately, and the roughly-repeated shape is complicated, yet the brain is still able to detect the frequency of repetition and to make sense out of the complexity. The skeptical reader might point out that the ear is simply performing an operation analogous to the Fourier transform instead of really constructing an analogy. However, while the Fourier transform (a computational process for decomposing a signal into a sum of several simpler component signals) is a rather straightforward mathematical process, without the flexibility we associate with analogy-making in general, I like to think of this basic perceptual process as the lowest-level form of analogy-making (or nano-analogy-making) that we do when listening to music. However, even a single note can involve high-level analogy-making: we can recognize piano-ness and clarinetness and even Frank-Sinatranness in a single note. This is *high-level* analogy-making on a very low-level temporal scale. As we move upwards to consider larger levels of structure, higher-level analogies will be even more apparent and the analogy processes involved will become more complex and fluid. But these higher-level analogies rely on the foundation provided by our excellent pitch-detecting nano-analogy makers in the auditory system.

TIMBRE

Several years ago I had the sheet music to Chopin's "Raindrop" prelude sitting open on my desk, when my friend Alia saw it and asked, "Which piece is this?" I pointed to the

opening measure (Figure 4.4) and sang “Ya, da daaaaa, da daaaa... it’s Chopin’s Raindrop prelude.” She nodded in recognition.



Figure 4.4: The beginning of Chopin’s “Raindrop” prelude.

Naturally, even though I was singing the notes instead of playing them on the piano, my friend still recognized the melody. My singing created an implicit analogy between the sound of my voice and of a piano. Neither of us thought about it at the time, since such analogy-making is so natural, but it occurred nonetheless.

PITCH

One of the most important musical attributes of a note is its pitch. Still, it is easy to think of one pitch as being “the same as” another even when the frequency of the two sounds are different (excepting individuals with perfect pitch, of course). We are more sensitive to the position of a note in the tonal context of a scale than to its absolute pitch. The fundamental frequency of a note itself is not even so well-defined when we consider, for example, an opera singer with a large-amplitude vibrato.

The most basic example of analogy involving pitches is that we hear notes separated by octaves as very similar to each other. Octave equivalence is so directly based on acoustics that it almost seems like a property of the physics of sound instead of kind of analogy-

making. When a choir of men and women sings a passage in unison, for instance, it sounds like both groups are singing the same pitches, even though the men typically sing an octave lower.

When we notice that a pitch is repeated, we might consider this act of recognition to be analogy-making. One of the most characteristic features of the Raindrop Prelude is the nearly constant repetition of the $A\flat$ throughout. Each time one hears the $A\flat$, it is in fact slightly different from the others: the attack and loudness will inevitably differ from note to note, the previously sounding vibrations in the string will give the sound a different character, and the other notes present at the same time will affect the sound. Even more important is the harmonic context in which each $A\flat$ occurs. Indeed, although the frequency of the note does not alter one bit at the end of the first section of the piece, in measure 29 of the score it is replaced with its enharmonic equivalent, $G\sharp$, indicating that the tonal underpinnings have changed. The music modulates from $D\flat$ Major to $C\sharp$ Minor (Figure 4.5).



Figure 4.5: The key change to $C\sharp$ Minor between measures 28 and 29.

As soon as the bass chords begin at this key change, the $A\flat/G\sharp$ simultaneously sounds “the same” as before — it is, after all, a sound with the same frequency as before — but it also sounds darker in this new context. Indeed, as I listen to the measure before the key change, I anticipate the coming measure and notice that the $A\flat$ gradually takes on a more

sinister character. Paradoxically, the repeated note sounds like it is *remaining* the same (Ab) at the same time as it is *changing* into something else (G#). Here, analogy is a unifying force, keeping the perception of repetition intact to some degree even as the context changes.

RHYTHM

As soon as we contemplate musical structures larger than single notes, we encounter the concept of rhythm. Often, a characteristic rhythm becomes a unifying musical feature that suggests similarity between two otherwise disparate structures. A simple example is the incessant repetition of the eighth-note G# octaves from measure 36 through measure 40. The rhythm of repeated eighth notes is almost the simplest rhythm possible, yet in the next measure (measure 41) it becomes crucial to continuity as the most stable element of the piece suddenly changes. The G# octaves become B octaves (Figure 4.6).



Figure 4.6: G# becomes B in measure 41.

Were it not for the repetition of the eighth-note rhythm here, I might have written that the G# octaves are *replaced* by B octaves. However, the rhythmic pulse keeps going forward and the eighth notes keep repeating, making the analogy between the previous measures and measure 41 extremely clear. Even though this measure seems distinct from earlier material, in the sense that it feels like a climactic moment in the piece, it is nonetheless composed of the same sort of musical material. The minor chords of this dark

C#-minor section suddenly transform into a loud E-major and the unexpected B octaves appear, but we can hear a strong analogy to what came earlier, thanks in large part to the constant rhythmic pattern. In that sense, the G# octaves *become* (as opposed to *being replaced by*) B octaves,

The opening motif of the piece (Figure 4.4) also has a characteristic rhythm: dotted-eighth – sixteenth – half. The melody in measure 3 sounds very similar to that in measure 1 even though the pitches are quite different (except, of course, for the bass’s repeating Ab). Rhythmic similarity also helps link the final measure here (measure 4) with measure 23. Both the initial phrase (measures 1–4) and a slightly modified version of this phrase (measures 19–23) end with a seven-note tuplet in the right hand. The second time it occurs, though, Chopin modifies the final seven pitches, and does not include the grace note at the end of the third beat that was present in measure 4 (Figure 4.7).



Figure 4.7: Measure 23 (analogous to measure 4).

“This phrase” appears again after the return to Ab-major near the end of the piece, and this time the tuplet has ten notes instead of seven, but it certainly sounds like more of the same thing. The only difference is that the flourish at the end of the phrase has become more elaborate, thanks to the extra notes (Figure 4.8).



Figure 4.8: Measure 85 (analogous to measure 23).

MOTIF

Continuing upwards to larger-scale structures, we encounter motifs — characteristic recurring short patterns of notes and/or rhythms. Familiar examples are found in Bach’s inventions, which are in large part based on the contrapuntal treatment of short motifs that return in all sorts of transpositions, inversions, and other modifications. Beethoven’s Fifth Symphony starts with the famous “Fate” motif and develops it throughout the rest of the first movement. Chopin may not make such obvious and pervasive use of short motifs as Bach and Beethoven in general, but the Raindrop prelude certainly does have one very significant and obvious motif: the repetition of $A\flat/G\sharp$. It is also possible to identify several other short motifs that gain significance during the course of this piece. I will examine some of these motifs before moving on to larger scale forms. During this analysis it is useful to keep in mind that motifs are highly relevant to the present discussion because when we notice a non-literal repetition of a motif, it indicates that we have formed an analogy.

$A\flat/G\sharp$ Revisited

I already mentioned the *fortissimo* moment where the repeated $G\sharp$ changes to a repeated B. Here, the rhythm and repetition of the motif make it clear that the B is still part of the same $G\sharp$ repetition motif, even though the pitch is different. This happens in a few

other places in the piece as well. Later in the minor middle section, G \sharp moves to F \sharp and then to A. The G \sharp switches between octaves a few times as well; sometimes it is lower, sometimes higher, and sometimes it is doubled. It even becomes C \sharp for a moment in measure 71 before the end of the minor section, as it moves from the upper to the lower octave. The A \flat in the first section of the piece also slips down to G \flat in measure 11, anticipating a descent through G \flat in measure 14 to become an F for five measures (Figure 4.9).

The image shows a musical score for piano, consisting of two systems of staves. The first system starts at measure 10 and the second system starts at measure 15. The key signature is B-flat minor (three flats) for the first system and changes to F major (no sharps or flats) for the second system. The time signature is 4/4. The score features a piano accompaniment with a steady eighth-note bass line and a melody in the right hand. The motif of a repeated eighth note on the dominant is maintained across the modulation.

Figure 4.9: A \flat becomes F.

In all these cases, we still hear the repetition motif regardless of the difference in pitch. Even when it changes to F, it's still the “same thing” as the A \flat . The analogy is particularly strong here because the change to F occurs as the piece modulates to the key of B \flat -minor. F is the dominant pitch in B \flat , just as A \flat was the dominant of D \flat (and G \sharp is the dominant of C \sharp). The motif thus seems to be of the form “repeated eighth note on the dominant”; when perceived in this way, the analogy is very convincing.

Later in the piece, once the motif has been well established, even a single repetition of an eighth note in the proper context is enough to evoke the motif (Figure 4.10).



Figure 4.10: Another variant of the repetition motif.

Dotted-Eighth – Sixteenth – Half

The very first three notes of the piece in the right hand present an important motif with a characteristic rhythm (Figure 4.11).



Figure 4.11: Rhythmic motif.

When the same rhythm appears in the third measure, the similarity is immediately noticeable, even though this version has more notes playing simultaneously and moves *upwards* by *scale steps* instead of moving *downwards* by *leaps* (Figure 4.12).



Figure 4.12: Rhythmic motif appearing in another form.

This rhythmic motif also appears at the end of other measures such as measure 11, although here it is preceded by three grace notes and is in a different metric position — the last beat instead of the first beat — and does not evoke the original motif as strongly (Figure 4.13).



Figure 4.13: Another appearance of the motif.

Perhaps the original version of the motif (that is, the motif including the original pitches, not just the rhythm) appears relatively rarely in order to maintain its distinctiveness whenever it makes an entrance. Each time the F–D–A \flat version is heard at the beginning of a measure, it is a clear signal that the opening melody has returned. This is particularly important at the end of the minor section, when this motif signals the end of the middle section of the piece and the return home to the major key.

High Descending Notes

The next most salient motif to me (aside from the ornamental tuplet flourish discussed earlier, in the Rhythm section) starts with a series of eighth notes descending from the high points of the phrases in the B-section of the first page of the piece (measures 9–19), for example, the G \flat –F–E \flat –D \flat eight notes in measure 9. The motif concludes each time with a return to a longer note that is a bit higher than the final eighth note. The first instance is in measure 10, starting on the high G \flat (Figure 4.14).



Figure 4.14: Measures 8-11: High-descending-notes motif.

The next measure (measure 11) is the second instance of the motif. It is several steps lower in pitch, moves through a minor scale, and starts with a leap of a third, but it still sounds like the same thing. This motif comes back in a different form at the climax of the short coda of this prelude. The first note in the motif becomes much longer and higher (two quarter notes tied together on a high B \flat in measure 87), while the others become quarter notes and involve a large leap to and from the note C in measure 88 (Figure 4.15).



Figure 4.15: Return of the high descent motif: a ray of sunlight or “a star in the jaws of the clouds”.

This is a quite significant moment in the piece, as it the only time that there is a lull in the incessant repetition of the eighth-note motif. If I think of the “Raindrop” subtitle of the prelude, this moment makes me think of a ray of sunlight. And that image, in turn, makes me think of a vivid image described by Victor Hugo in *Les Misérables*, “a star in the jaws of the clouds” (“une étoile dans les gueules des nuages”) (Hugo, 1862, 1987). The analogy-making doesn’t stop.

PHRASES

A note to the reader: this section, on musical phrases, and the next, on large-scale form, will be easier to understand if you have the sheet music for the “Raindrop” prelude handy, because I will not reproduce the score here. I recommend finding a copy to use in following along, although it’s not strictly necessary.

The opening part of the prelude divides nicely into phrases. The first four measures make up a phrase (call it **a**). The next four measures are quite similar, but the phrase cadences and comes to a pause instead of moving forward with the seven-note tuplet of **a**. Otherwise the phrases are exactly the same. The analogy between them is almost trivial, especially since the second is a truncation of the first. I'll label this phrase **a'**. Next, a contrasting phrase (**b**₁) begins, although it only has three measures, followed by (**b**₂) for four measures and (**b**₃) for another four measures. The relationships between these phrases are more complicated and hard to pin down, even though it is easy to hear them as closely related. Each of these phrases involves a combination of the descending high-note motif and another vaguer motif involving quarter notes (measures 12, 16, and 18). After these **b** phrases, the **a** phrase returns two more times. The first return is the same phrase as before except for a modification of the pitches in the seven-note tuplet, while the final **a'** is much like **a**, only shorter and with a more conclusive coda followed by a transition into the minor section.

My grouping of measures into phrases and pointing out that **b**₁, **b**₂, and **b**₃ are related is based on my noticing of similarities between motifs, rhythms, textures, etc. in the phrases. That is, making analogies between phrases is based on the analogy-making occurring at smaller scales. It is more complicated to describe precisely the relationships between the **b** phrases than between **a** and **a'** because of the large number of elements that make up their similarity.

LARGE-SCALE FORM

The phrases of the first section group naturally into larger structures ("periods") so we can describe the first section of the piece as **ABA'**. The minor middle section has a more complicated larger form, although **CCD** is a good approximation. Finally, the coda might be labeled **A''E** (a variant of the original two phrases followed by an ending). Just as the **b**

phrases involved more sophisticated analogies between many elements, the minor section of the piece (CCD) is quite complicated and would take a good deal of prose to describe well. However, in listening to the middle section, it is easy to hear relationships between the phrases.

Overall, the piece itself follows an **ABA'** structure. This structure is easy to spot, thanks to the key changes that divide the piece into three sections, and the final **A'** is quite clearly related to **A** due to near-literal repetition of the initial phrase. Perhaps the piece is easiest to describe at this top level, rather than at middle or low levels, because so many details have been abstracted away by lower-level analogies. The **ABA'** structure indicates that we have an initial section, a contrasting middle, and a final section related to the beginning. Noticing this contrast involves analogy-making while we group all the pieces of **A** together and notice how they are distinct from **B**, but it seems that the most interesting analogy-making in this example happens at lower levels.

RAINDROPS KEEP FALLING...

Although Chopin himself did not give this prelude the “Raindrop” subtitle, it is difficult for me to dissociate the piece from thoughts of rain since the subtitle is virtually always attached to the piece. Naturally, this evokes a large-scale analogy wherein the music is compared with falling rain. When I hear the piece, the dotted-eighth opening note evokes thoughts of a raindrop collecting on a tree branch or overhang of a roof. The following sixteenth sounds like the point at which the mass of collected water becomes too great and the droplet breaks free of surface tension and falls down to the half note $A\flat$. It continues to the octave below, splashing into a puddle of many $A\flat$ eighth notes: “plunk, plunk, plunk”.

This interpretation of the opening motif is of course my own idiosyncratic view based on my personal experience, while the incessant “plunk, plunk, plunk” of the repetition motif seems likely to be a more universal interpretation. I find it likely that because of the name that has been associated with the piece, most people will automatically hear this “plunk, plunk, plunk” image of raindrops falling. Additionally, the minor section in the middle evokes predictably gloomy, stormy images. It is remarkable that music can evoke such concrete images about the real world. Listening to this prelude leads me to form detailed analogies between notes in the piece, between various rhythmic elements, musical shapes, motifs, phrases, sections, and so forth. These elements of analogy work at various levels in the piece to make sense of the structure and unify it into a whole piece that makes sense even at a representational level. I can listen to the prelude and, without difficulty, hear it as a light rain shower followed by a more serious windy storm, followed by a lightening of the sky, a burst of sunlight, and a final few plunks of raindrops before the end of the storm.

However, although raindrop analogies are easy to make to go along with this prelude, it is also enjoyable to listen to the piece without rain in mind. Indeed, in the *Unanswered Question* lectures, Bernstein (1976) challenges his audience to listen to a performance of Beethoven’s “Pastorale” Symphony without thinking of the cartoon images from the Disney movie *Fantasia*, in which animated stories are invented to go with the music of several classical pieces including the “Pastorale”. It’s a difficult challenge! But if you can strip away the animations from *Fantasia* or the subtitle “Raindrop” from the Chopin prelude, you can *still* hear *emotional* content. The act of perceiving various structures, motifs, analogies, expectations, tensions, resolutions, and so forth evokes emotional responses on its own.

Analogy as the Core of Music Cognition

We have seen many examples of musical elements that can sound like other elements due to analogies between the elements. How, though, do these examples support the thesis that analogy is the core of music cognition? What is meant by this statement? Specifically, I claim that

- 1) we *make sense of music* by making analogies between musical elements to aid in grouping and understanding; and
- 2) we *appreciate music* by enjoying the similarities and differences that go into each analogy.

An entire musical piece is a large, complex structure, which is digestible only in terms of smaller chunks. These smaller chunks are formed by grouping sections, phrases, measures, or notes together by virtue of similarities making up analogies. The act of perceiving analogies within a piece (or between pieces, or even between musical events and real-world events) helps give meaning to music. Musical semantics is different from the semantics of natural language because in general it is not representational, but when we listen, simply perceiving that a new musical element is similar to something heard before makes the new element sound meaningful. For example, the return of the opening motif at the return of the major section at the end of the “Raindrop” prelude sounds significant and means something to the listener just because of the happy recognition of the dotted-eighth motif. Much of our enjoyment of music comes from the analogies we spontaneously make when listening. For instance, the recognition that the $A\flat$ has turned into a sinister $G\sharp$ is a wonderful moment where the induced affect is crucially based on the listener perceiving an analogy between the notes.

Analogy is central in music cognition because it occurs spontaneously at every level of musical listening, from individual sounds to the form of an entire piece and even further “upwards”, connecting music with the external world. Additionally, parallel analogy-making at various different scales facilitates the building-up of larger and larger compound structures of musical understanding.

A subtle difficulty with analogy in music-cognition research is that analogy-making is so automatic for human listeners that it is hard to recognize that analogies are being made at all. My purpose here has been to try to illustrate the prevalence of analogy-making in music-listening in hopes that future work in the field may be motivated by recognition of linkthe centrality of analogy.

Analogy as the Core of Musicat

THE IMPORTANCE OF ANALOGY IN MODELS OF MUSIC COGNITION

Musicat without Analogy

Initial attempts at building Musicat did not include analogy as a central notion; likewise, most computational models of music cognition have little or no notion of analogy. Without analogy, however, Musicat was a collection of mostly low-level perceptual rules, which failed to make sense of musical structure. Even for structures as small as a musical phrase of four or eight measures, it was hard to make the program understand the grouping structure without reference to analogy. Once analogy was added, however, the program worked much better. Chapter 9 discusses this in more detail.

GTTM and Parallelism

In their seminal book *A Generative Theory of Tonal Music*, Lerdahl and Jackendoff do not use the term “analogy” *per se* but they do point out the importance of parallelism in musical grouping, which is a closely related notion. The relevant rule is “GPR 6”:

Grouping Preference Rule 6 (Parallelism) Where two or more segments of the music can be construed as parallel, they preferably form parallel parts of groups. (Lerdahl & Jackendoff, 1983, p. 51)

To “construe as parallel” is the heart of the matter here, and it is not obvious how to formalize this intuitive idea. Lerdahl and Jackendoff give a three-page digression (pp. 52–55) on the importance of parallelism, beginning with this passage:

The importance of parallelism in musical structure cannot be overestimated. The more parallelism one can detect, the more internally coherent an analysis becomes, and the less independent information must be processed and retained in hearing or remembering a piece. However, our formulation of GPR 6 still leaves a great deal to intuition in its use of the locution “parallel.”

Of particular interest is the authors’ humble admission that the details of parallelism constitute an important lack in their theory:

It appears that a set of preference rules for parallelism must be developed, the most highly reinforced case of which is identity. But we are not prepared to go beyond this, and we feel that our failure to flesh out the notion of parallelism is a serious gap in our attempt to formulate a fully explicit theory of musical understanding.

I claim that Musicat begins to fill this gap. Musicat presents an explicit model of analogy-making between musical structures, and parallelism can be considered as a particular application of analogy to the understanding of musical structure. And not only is formalizing analogy-making important to music cognition, but also it is important to cognition in

general, if we agree with Hofstadter's claim that analogy is the core of cognition. Optimistically, I hope that my work (like all the models developed by FARG) will have applicability beyond its particular (micro)domain. Lerdahl and Jackendoff have similar sentiments:

The problem of parallelism, however, is not at all specific to music theory; it seems to be a special case of the much more general problem of how people recognize similarities of any sort — for example similarities among faces... the hope of developing a solution to the musical problem in terms of the preference-rule formalism suggests that such a formalism may be more widely applicable in psychological theory. (Lerdahl & Jackendoff, 1983, p. 53)

Of course, my approach uses the FARG architecture instead of preference rules, but I share Lerdahl and Jackendoff's desire to apply lessons from work in music to other fields. Modeling analogy-making in music is essential to music (and Musicat), but it is also relevant to cognition in general.

This chapter has presented some evidence for the importance of analogy-making in music-listening, and therefore in the modeling of music cognition. The following three chapters demonstrate how these ideas have been put to use: in them we will see many concrete examples of analogies (and other structures) generated by the program Musicat.

