Chapter 4

Basic Pitch Translations

Many musical processes entail some sort of data translation in which one form of representation is transformed into another form of representation. There are innumerable examples of such musical "translations." For example, we might rewrite guitar tablatures as notated pitches. Pitches might be translated to tonic-solfa syllables. Scale-degrees might be analyzed as Roman-numeral harmonies. Figured bass notations might be "realized" as pitches. Successive pitches might be characterized as melodic intervals. Intervals might be rewritten as semitone distances. Pitch-class sets might be transformed to interval-class vectors.

In this chapter we introduce some basic musical translations. We will limit this initial discussion to translating between various representations related to pitch. These will include French, German, and international pitch designations, frequency, semitones, cents, solfège, scale-degree and MIDI pitch representations. In later chapters we will describe additional types of translations.

ISO Pitch Representation

The best-known system for representing equally-tempered pitches is the International Standards Organization (ISO) format consisting of a letter (A-G) followed by an optional sharp or flat, followed by an octave number. According to the ISO representation middle C is designated C4. Octave numbers change between B and C so that the B a semitone below C4 is B3. Humdrum provides a predefined ISO-like representation called `**pitch` illustrated below. Here we see an ascending chromatic scale in the left spine, with a concurrent descending chromatic scale in the right spine:

```
**pitch   **pitch
C4        C5
C#4       B4
D4        Bb4
D#4       A4
E4        Ab4
F4        G4
F#4       Gb4
G4        F4
G#4       E4
```

```
A4      Eb4
A#4     D4
B4      Db4
C5      C4
*-      *-
```

Notice that only upper-case letters are used for pitch-name and that the flat is represented by the lower-case letter `b`. The small letter `x` can be used to indicate double-sharps and the double-flat is represented by two small successive letters `bb`. If a given pitch deviates from equal temperament, cents deviation can be indicated by trailing integers preceded by either a plus sign (tuned sharp) or a minus sign (tuned flat). The unit of the *cent* is one one-hundredth of a semitone, so the distance between C4 and C#4 is 100 cents. For example, the pitch corresponding to a tuning of A-435 Hz is 19 cents flat compared with the standard A-440, hence it is represented in `**pitch` as:

```
A4-19
```

Other pitch representations (such as `**kern`) can be translated to the ISO-inspired `**pitch` representation by invoking the **pitch** command. For example, consider the following `**kern` input:

```
**kern
=1
4g
4g#
4a
4cc
=2
*-
```

It can be translated to this `**pitch` output:

```
**pitch
=1
G4
G#4
A4
C5
=2
*-
```

using the following command:

```
pitch filename
```

Notice that the `**pitch` representation uses the same system for representing barlines as `**kern`. In fact, all of the pitch-related representations described in this chapter make use of the so-called 'common system' for representing barlines.

German Tonhöhe

The German system of pitch designations (*Tonhöhe*) is available in the Humdrum `**Tonh` representation. The German system is similar to the ISO system with the following exceptions. The pitch letter names include A,B,C,D,E,F,G,H and S. Sharps and flats are indicated by suffixes: '`is`' for sharps and '`es`' for flats, hence '`Cis`' for C-sharp and '`Ges`' for G-flat. Suffixes are repeated for double and triple sharps and flats. Special exceptions include the following: '`B`' for B-flat, '`H`' for B-natural, '`Heses`' for B double-flat (rather than '`Bes`'), and '`As`' and '`Es`' rather than '`Aes`' or '`Ees`'. '`S`' is a short-hand for '`Es`' (E-flat). As in the ISO pitch system, octaves are indicated by integer numbers with middle C represented as C4.

Although modern German practice has gravitated toward the ISO system, the traditional German system for representing pitches remains important in historically related studies, such as searching for 'B-A-C-H' and the pitch signature used by Dmitri Shostakovich ('D-S-C-H').

Data in the `**pitch` or `**kern` representations can be translated to `**Tonh` via the **tonh** command:

```
tonh filename
```

French Solfège

The common French system for pitch naming uses a so-called "fixed-do" method of diatonic pitch designations: *do, ré, mi, fa, sol, la* and *si* (rather than *ti*), where *do* corresponds to the English/German 'C'. In the Humdrum `**solfg` representation, solfège pitch names are used. Flats (*bémol*) and sharps (*dièse*) are abbreviated *b* and *d* respectively. When accidentals are encoded, the tilde character (~) is encoded following the solfège syllable and before the accidental. Double and triple sharps and flats are encoded via repetition. Hence, 'do dièse' (`do~d`) for C-sharp, 'la bémol' (`la~b`) for A-flat, 'sol double-dièse' (`sol~dd`) for G double-sharp, 'si double-bémol' (`si~bb`) for B double-flat, and so on. As with the German and ISO pitch representations, octave is designated by integers with `do4` representing middle C.

Frequency

For acoustic-related applications it may be helpful to translate to frequency. The Humdrum `**freq` representation can be used to represent frequencies for either pure or complex tones. Frequencies are encoded in *hertz* (abbreviated *Hz*) where 440 Hz means 440 cycles per second. In the `**freq` representation frequencies may be specified as integer or real values (with a decimal point).

Cents

The `**cents` representation provides a means for representing pitches in absolute units with respect to middle C (= 0 cents). In the `**cents` representation, all pitches are represented with respect to this reference. Thus C#4 is represented by the number 100, A4 is represented by 900, and A3 is represented by −300. As in the case of `**freq`, cents may be specified as integer numbers or as real values (with a decimal point).

Semitones

A related pitch representation is `**semits`. In this case, all pitches are represented in numerical semitones with respect to middle C (= 0 semits). An ascending chromatic scale beginning on C4 would be represented by the ascending integers from 0 to 12. Pitches below middle C are represented by negative values. Fractional values can be represented using decimal points.

MIDI

Another way of representing pitch is provided by the Humdrum `**MIDI` representation. This representation closely mimics the commercial MIDI specification. The `**MIDI` representation allows MIDI inputs and outputs to be exported or imported by various Humdrum tools. A complete description of `**MIDI` will be given in Chapter 7.

Scale Degree — ***solfa* and ***deg*

Two different Humdrum representations are provided to describe scale-degree related information: `**deg` and `**solfa`. Both of these representations emphasize slightly different aspects of scale-degree information. Both representations assume some established or pre-defined tonal center or tonic pitch.

The `**solfa` representation represents pitch according to tonic solfa syllables. Pitches are designated by the syllables *do, re, mi, fa, so, la* and *ti* or their chromatic alterations as indicated in the following table:

basic	raised	lowered
do (*doe*)	di (*dee*)	de (*day*)
re (*ray*)	ri (*ree*)	ra (*raw*)
mi (*me*)	my (*my*)	me (*may*)
fa (*fah*)	fi (*fee*)	fe (*fay*)
so (*so*)	si (*see*)	se (*say*)
la (*la*)	li (*lee*)	le (*lay*)
ti (*tee*)	ty (*tie*)	te (*tay*)

Summary of **solfa** *Signifiers*

The `**deg` representation identifies scale-degrees by the numbers 1 (tonic) to 7 (leading-tone). These values may be chromatically altered by raising (+) or lowering (–). The *amount* of chromatic alteration is not indicated; for example, both a raised supertonic and a doubly-raised supertonic are represented as `2+`. A lowered dominant is represented as `5-`.

The `**solfa` representation differs from `**deg` in that pitches are represented without regard to major or minor *mode.* For example, in the key of C major, `**deg` will characterize A-flat as a lowered sixth scale degree (`6–`), whereas the same pitch will be a normal (unaltered) sixth scale degree in the key of C minor (`6`). In the case of `**solfa`, the A-flat will be represented as `le` —

whether or not the key is C major or C minor. Like `**deg`, the amount of chromatic alteration is not represented in `**solfa`. Once a pitch is raised, raising it further will not change the representation. For example, if the tonic is B-flat, then both B-natural and B-sharp will be represented by `di` in the `**solfa` representation.

In the case of the minor mode, `**deg` characterizes scale degrees with respect to the *harmonic minor* scale only.

Another difference between `**solfa` and `**deg` is that the `**deg` representation provides a way for encoding *melodic approach.* The caret (^) denotes an ascending melodic approach to the current note, whereas the lower-case letter `v` denotes a descending melodic approach. Repeated pitches carry no melodic approach signifier.

Some of the differences between the `**solfa` and `**deg` representations are illustrated in Example 4.1. (The corresponding `**kern` representation is given in the first spine.) Notice that `**solfa` does not encode any octave information. The `**deg` representation does not encode the octave of the starting pitch, but it does indicate contour information using the caret (^) for ascending and the lower-case `v` for descending pitches. Notice also the different ways of characterizing accidentals.

Example 4.1

```
!! Comparison of pitch-related representations.
**kern          **solfa         **deg
*M2/4           *M2/4           *M2/4
*c:             *c:             *c:
8.cc            do              1
16dd            re              ^2
=1              =1              =1
8.ee-           me              ^3
16dd            re              v2
4een            mi              ^3+
=2              =2              =2
8r              r               r
8b-             te              v7-
8an             la              v6+
8cc             do              ^1
=3              =3              =3
2bn             ti              v7
==              ==              ==
*-              *-              *-
```

Pitch Translations

Humdrum provides a number of commands for translating between the various pitch-related representations described above. Typically, the command name is the same as the name of the output representation. For example, translating to the `**solfg` representation can be accomplished with:

```
solfg inputfile > outputfile
```

Translating to the German `**Tonh` representation:

```
tonh inputfile > outputfile
```

Translating to ISO `**pitch`:

```
pitch inputfile > outputfile
```

Similarly, the **freq** command translates pitch-related inputs to the `**freq` representation, the **cents** command translates appropriate inputs to the `**cents` representation, and so on.

In a few cases, the command names are slightly modified. All Humdrum command names employ lower-case letters only, so `**MIDI` output is generated by the **midi** command (rather than the **MIDI** command), and `**Tonh` output is generated by the **tonh** command.

Examples 4.2 and 4.3 compare several parallel representations of the same pitch-related information. In both examples, the pitch information has been derived from the `**kern` data shown in the left-most spine. The duration information in the `**kern` data is not available in the other representations. However, the 'common system' for barlines is used throughout.

Example 4.2 shows four pitch naming systems: ISO pitch, German Tonhöhe, French solfège, as well as `**kern`. Notice the different ways of treating accidentals such as the D-sharp and B-flat. Also note the German use of H for B-natural.

Example 4.2

```
!! Comparison of pitch-related representations.
**kern      **pitch     **Tonh      **solfg
*M2/4       *M2/4       *M2/4       *M2/4
*C:         *C:         *C:         *C:
8.cc        C5          C5          do5
16dd        D5          D5          re5
=1          =1          =1          =1
8.ee        E5          E5          mi5
16dd#       D#5         Dis5        re~d5
4ee         E5          E5          mi5
=2          =2          =2          =2
8r          r           r           r
8b-         Bb4         B4          si~b4
8a          A4          A4          la4
8c          C4          C4          do4
=3          =3          =3          =3
2bn         B4          H4          si4
==          ==          ==          ==
*-          *-          *-          *-
```

In Example 4.3 four of the more technical representations are illustrated, including frequency and cents. Notice that the `**MIDI` representation uses key-numbers to represent pitch: key-on events are indicated by positive integers (between two slashes) and key-off events are indicated by negative integers. More detail concerning `**MIDI` is given in Chapter 7.

Example 4.3

```
!! Comparison of pitch-related representations (continued).
**kern      **semits    **cents     **MIDI      **freq
*M2/4       *M2/4       *M2/4       *Ch1        *M2/4
*C:         *C:         *C:         *M2/4       *C:
*           *           *           *C:         *
8.cc        12          1200        /72/        523.25
16dd        14          1400        /-72/ /74/  587.33
=1          =1          =1          =1          =1
8.ee        16          1600        /-74/ /76/  659.26
16dd#       15          1500        /-76/ /75/  622.25
4ee         16          1600        /-75/ /76/  659.26
=2          =2          =2          =2          =2
8r          r           r           /-76/       r
8b-         10          1000        /70/        466.16
8a          9           900         /-70/ /69/  440.00
8c          0           0           /-69/ /60/  261.63
=3          =3          =3          =3          =3
2bn         11          1100        /-60/ /71/  493.88
==          ==          ==          ==          ==
.           .           .           /-71/       .
*-          *-          *-          *-          *-
```

Not all of the above pitch-related representations can be translated directly from one to another. Table 4.1 shows the possible translations supported by Humdrum Release 2.0 commands. The input representations are listed from right to left. Under each column, those commands that will translate *from* the given format are identified. For example, the `**cents` representation can be translated to `**freq`, `**kern`, `**pitch`, `**semits`, `**solfg`, and `**tonh`. Notice that `**deg` data cannot be translated to any other format since `**deg` representations do not encode absolute pitch height. Note also that when translating to the `**kern` representation, only pitch-related information is translated: duration, articulation marks, and other `**kern` signifiers are not magically generated.

Table 4.1

Input Representation

	**cents	**deg	**freq	**kern	**MIDI	**pitch	**semits	**solfa	**solfg	**Tonh
cents			•	•		•	•		•	•
cocho			•							
deg				•		•			•	•
freq	•			•	•	•	•		•	•
kern	•		•		•	•	•	•	•	•
midi				•						
pitch	•		•	•	•		•	•	•	•
semits	•		•	•	•	•			•	•
solfa				•		•			•	•
solfg	•		•	•	•	•	•		•	
tonh	•		•	•	•	•	•	•	•	

Transposition Using the *trans* Command

A common pitch-related manipulation is transposition. The **trans** command has the user specify a *diatonic offset* and a *chromatic offset*. The diatonic offset affects the pitch-letter name used to spell a note. The chromatic offset affects the number of semitones shifted from the original pitch height. The two types of offset are completely independent of each other. For common transpositions, both the diatonic and chromatic offsets will need to be specified. For example, in transposing up a minor third (e.g. C to E-flat), the diatonic offset is 'up two pitch-letter names,' and the chromatic offset is 'up three semitones.' The appropriate command invocation is:

```
trans -d +2 -c +3 input > output
```

The diatonic offset can be a little confusing because traditional terminology labels perfect unisons by the number 1 (e.g. P1) rather than zero. So transposing up a perfect fifth involves a diatonic offset of +4 letter names, and a chromatic offset of +7 semitones:

```
trans -d +4 -c +7 input > output
```

We can transpose without changing the diatonic pitch names. For example, the following command will transpose down an augmented unison (e.g. C# to C):

```
trans -d 0 -c -1 input > output
```

Conversely, we can respell the diatonic pitches without changing the overall pitch height. For example, the following transposition will transpose "up" a diminished second (e.g. from F-sharp to G-flat):

```
trans -d +1 -c 0 input > output
```

Modal transpositions are also possible by omitting the chromatic offset option. Consider, for example, the following C major scale:

```
**kern
c
d
e
f
g
a
b
cc
*-
```

We can transform this using the following diatonic transposition:

```
trans -d +1
```

The resulting output is the Dorian mode:

```
**kern
*Trd1
d
e
f
g
a
b
cc
dd
*-
```

When using the **-d** option alone, **trans** eliminates all accidentals in the input. This can be potentially confusing, but it is often useful. Suppose you have a passage in the key of E major which you would like to translate to E Dorian. First transpose so the tonic is D using only the **-d** option; then transpose exactly so the tonic is E again:

```
trans -d -1 Emajor | trans -d +1 -c +2 > Edorian
```

For some changes of mode (such as melodic to harmonic minor), you may need to use the **humsed** command described in Chapter 14 to modify accidentals for specific scale degrees.

Notice the addition of a "tandem interpretation" to the above example (`*Trd1`). Whenever **trans** is invoked, it adds a record indicating that the encoding is no longer at the original pitch. *Transposition tandem interpretations* are similar in syntax to the **trans** command itself. In the above example, `*Trd1` indicates a diatonic shift up one letter name. The tandem interpretation `*Trd-1c-2` would indicate that a score has been transposed down a major second. The **trans** command also provides a **-k** option that allows the user to specify a replacement key signature for the output.

The **trans** command can be used in conjunction with any of the appropriate pitch-related representations, such as `**pitch`, `**kern`, `**Tonh`, and `**solfg`.

Key Interpretations

In order for the **solfa** or **deg** commands to translate from other pitch representations, the encoded music must contain an explicit key indication. Keys are explicitly represented by a single asterisk, followed by an upper- or lower-case letter, followed by an optional accidental, followed by a colon. The octothorpe (#) indicates a sharp and the hyphen (–) indicates a flat.

Upper-case letters indicate major keys; lower-case letters indicate minor keys. By way of illustration, the following key interpretations indicate the keys of C major, C minor, B-flat major, and F-sharp minor:

```
*C:
*c:
*B-:
*f#:
```

Key interpretations usually appear near the beginning of a representation, and key interpretations can be redefined at any place in a score.

Pitch Processing

Apart from transposition, translating from one representation to another provides opportunities for different sorts of processing. Suppose, for example, we wanted to know whether the subdominant pitch occurs more frequently in one vocal repertory than in another repertory. We can use **solfa** in conjunction with **grep**'s **-c** option to count the number of occurrences. (For the following examples, we will assume that the inputs consist of only a single spine, that barlines are absent, and that appropriate interpretations are provided indicating the key of each work.) First we need to count the total number of notes in each repertory.

```
census -k repertory1.krn
census -k repertory2.krn
```

Next we translate the scores to the **solfa** representation and use **grep -c** to count the number of occurrences of the number 'fa':

```
solfa repertory1.krn | grep -c fa
solfa repertory2.krn | grep -c fa
```

The proportion of subdominant pitches can be calculated by simply comparing the resulting pattern count with the number of notes identified by **census**.

Recall that one of the differences between the `**solfa` and `**deg` representations is that the `**deg` output contains an indication of the direction of melodic approach. The caret (^) indicates approach from below, whereas the lower-case **v** indicates approach from above. Suppose we wanted to determine whether the dominant pitch is more commonly approached from above or from below. Assuming a monophonic input, we can once again use **grep** to answer this question. First let's count how many dominant pitches ('5') are approached from above ('v'):

```
deg repertory.krn | grep -c v5
```

The caret has a special meaning for **grep** which will be discussed in Chapter 9. We can escape the special meaning by preceding the caret by a backslash. In order to count the number of dominant pitches approached from below we can use the following:

```
deg repertory.krn | grep -c \^5
```

Recall that some scale tones are spelled differently depending on whether the mode is major or minor. For example, in A major the mediant pitch is C sharp; but in A minor the mediant pitch is C natural. The **deg** and **solfa** commands produce subtly contrasting outputs that make one or the other command better suited depending on the user's goal. The **deg** command would represent C sharp in A major, and C natural in A minor by the same scale degree — `3`. In the key of A major, C natural would be characterized as a lowered mediant (`3-`) and in A minor, C sharp would be characterized as a raised mediant (`3+`). By contrast, the **solfa** command characterizes pitches with respect to the tonic alone and ignores the mode. Hence, **solfa** would designate C sharp as '`mi`' whether the key was A major or A minor. Similarly, C natural would be designated '`me`' in both A major and A minor. The differences between **deg** and **solfa** allow users to distinguish chromatically altered scale tones in a manner appropriate to the task.

Uses for Pitch Translations

Occasionally it is useful to process a given representation to the *same* representation. The **kern** command translates various pitch-related representations to the `**kern` format. The **-x** option eliminates any input data that do not pertain to pitch. When applied to a `**kern` input, this option allows us to filter out durations, articulation marks, phrasing, and other non-pitch data. Suppose, for example, that we wanted to determine the proportion of successively repeated notes in a vocal melody: how often is a pitch followed immediately by the same pitch? We might begin by first determining the total number of notes in the melody using **census** with the **-k** option.

```
census -k melody.krn
```

We can use the **uniq** command to eliminate successive repeated pitches — but only if the note tokens are identical. First we can use **kern -x** to translate "from `**kern` to `**kern`" while eliminating non-pitch-related data. Then we need to remove barlines so they don't interfere with pitches that are repeated across the measure. Using **uniq** will then eliminate all of the successively duplicated records, so a sequence of six G's will be reduced to a single G. Finally, we pipe the output to **census -k** to count the total number of notes.

```
kern -x melody.krn | uniq | census -k
```

A variation on this approach would entail translating to a representation that does not distinguish enharmonic pitches. For example, translating our melody to `**semits` and then back to `**kern` will standarize all of the enharmonic spellings. If our melody contains a G-sharp that undergoes an enharmonic shift to A-flat, then the pitches will be deemed identical. The following command carries out the same task as above, but ignores possible enharmonic spellings:

```
semits melody.krn | kern | uniq | census -k
```

Incidentally, given `**semits` input, the **kern** command will spell pitches according to any key or key signatures it encounters. For example, if the key signature contains sharps, then G-sharp will be output; if the key or key signature contains flats, then A-flat will be output.

Reprise

In this chapter we have introduced a number of pre-defined pitch-related representations. Simple commands can be used to translate from one representation to another. Which representation is most appropriate depends on the user's goal.

There is a wealth of other representation formats related to pitch distances, tablatures, timing, and other types of musical information. These representations will be explored in later chapters. In addition, we'll describe how to design your own representations — representations that may be better tailored to a specific application. However, before we continue discussing further representations, this is an appropriate point to present a more formal description of the general Humdrum representation syntax.