Chordify

Advanced Functional Programming for Fun and Profit

José Pedro Magalhães

http://dreixel.net

September 27, 2014
Berlin, Germany

Introduction

- Modelling musical harmony using Haskell
- Applications of a model of harmony:
  - Musical analysis
  - Finding cover songs
  - Generating chords for melodies
  - Generating chords and melodies
  - Correcting errors in chord extraction from audio sources
  - Chordify—a web-based music player with chord recognition

Demo: Chordify

Demo:

chordify®

http://chordify.net

What is harmony?

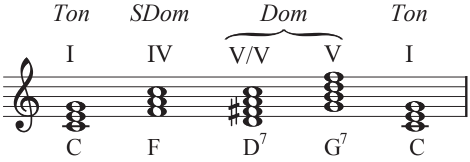

- Harmony arises when at least two notes sound at the same time
- Harmony induces tension and release patterns, that can be described by music theory and music cognition
- The internal structure of the chord has a large influence on the consonance or dissonance of a chord
- The surrounding context also has a large influence

What is harmony?

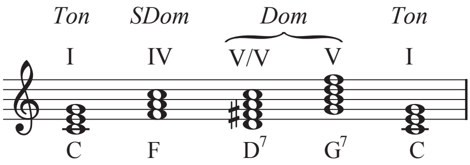

- Harmony arises when at least two notes sound at the same time
- Harmony induces tension and release patterns, that can be described by music theory and music cognition
- The internal structure of the chord has a large influence on the consonance or dissonance of a chord
- The surrounding context also has a large influence

Demo: how harmony affects melody

Simplified harmony theory I

- A *chord* is a group of tones separated by intervals of roughly the same size.
- All music is made out of chords (whether explicitly or not).
- There are 12 different notes. Instead of naming them, we number them relative to the first and most important one, the tonic. So we get I, II♭, II . . . VI♯, VII.
- A chord is built on a root note. So I also stands for the chord built on the first degree, V for the chord built on the fifth degree, etc.
- So the following is a chord sequence: I IV II^7 V^7 I.

Simplified harmony theory II

Models for musical harmony explain the harmonic progression in music:

- Everything works around the *tonic* (I).
- The *dominant* (V) leads to the tonic.
- The *subdominant* (IV) tends to lead to the dominant.
- Therefore, the I IV V I progression is very common.
- There are also *secondary dominants*, which lead to a relative tonic. For instance, II^7 is the secondary dominant of V, and I^7 is the secondary dominant of IV.
- So you can start with I, add one note to get I^7, fall into IV, change two notes to get to II^7, fall into V, and then finally back to I.

An example harmonic analysis

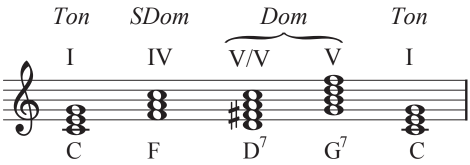

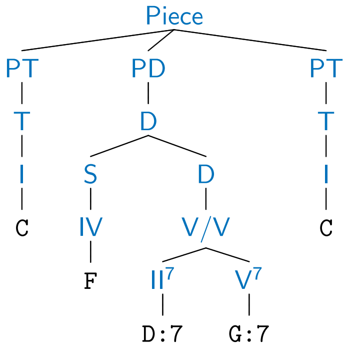

Why are harmony models useful?

Having a model for musical harmony allows us to automatically determine the functional meaning of chords in the tonal context.
The model determines which chords "fit" on a particular moment in a song.

Why are harmony models useful?

Having a model for musical harmony allows us to automatically determine the functional meaning of chords in the tonal context. The model determines which chords "fit" on a particular moment in a song. This is useful for:

- Musical information retrieval (find songs similar to a given song)
- Audio and score recognition (improving recognition by knowing which chords are more likely to appear)
- Music generation (create sequences of chords that conform to the model)

Application: harmony analysis

Parsing the sequence G_{min} C^7 G_{min} C^7 F_{Maj} D^7 G^7 C_{Maj}:

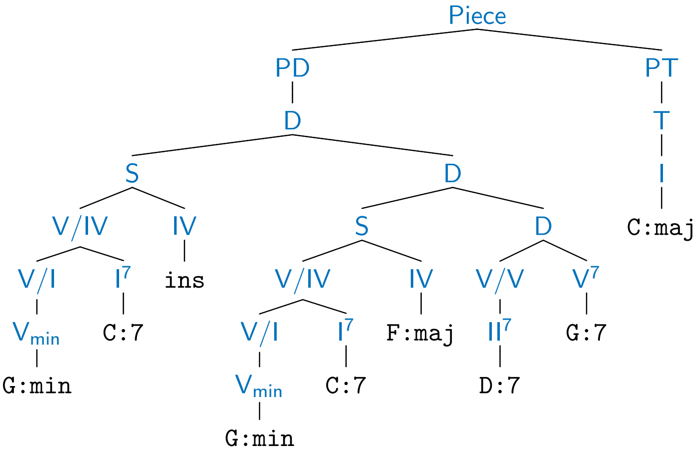

Application: harmonic similarity

- A practical application of a harmony model is to estimate harmonic similarity between songs
- The more similar the trees, the more similar the harmony
- We don't want to write a diff algorithm for our complicated model; we get it automatically by using a *generic diff*
- The generic diff is a type-safe tree-diff algorithm, part of a student's MSc work at Utrecht University
- Generic, thus working for any model, and independent of changes to the model

Application: automatic harmonisation of melodies

Another practical application of a harmony model is to help selecting good harmonisations (chord sequences) for a given melody:

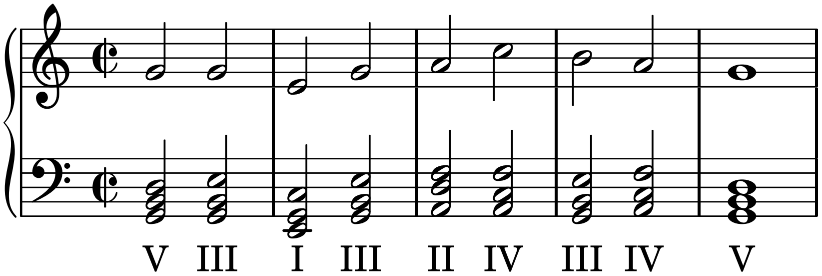

We generate candidate chord sequences, parse them with the harmony model, and select the one with the least errors.

Visualising harmonic structure

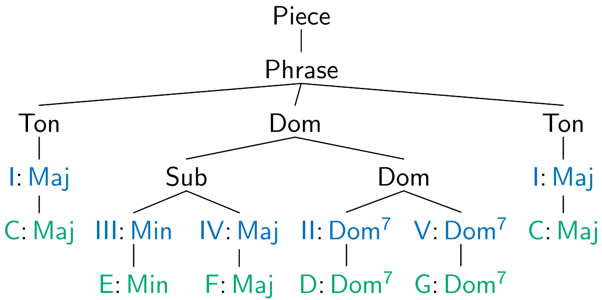

You can see this tree as having been produced by taking the chords in green as input. . .

Generating harmonic structure

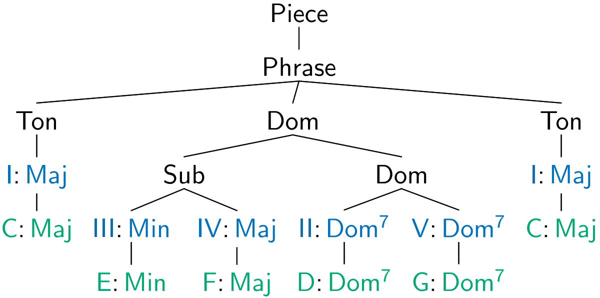

You can see this tree as having been produced by taking the chords in green as input. . . or the chords might have been dictated by the structure!

A functional model of harmony

$$\text{Piece}_{\mathfrak{M}} \rightarrow [\text{Phrase}_{\mathfrak{M}}] \qquad (\mathfrak{M} \in \{\text{Maj}, \text{Min}\})$$

A functional model of harmony

$$\text{Piece}_{\mathfrak{M}} \rightarrow [\text{Phrase}_{\mathfrak{M}}] \qquad (\mathfrak{M} \in \{\text{Maj}, \text{Min}\})$$

$$\begin{aligned}\text{Phrase}_{\mathfrak{M}} \rightarrow\ & \text{Ton}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}} \\ |\ & \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}}\end{aligned}$$

A functional model of harmony

$$\text{Piece}_{\mathfrak{M}} \rightarrow [\text{Phrase}_{\mathfrak{M}}] \qquad (\mathfrak{M} \in \{\text{Maj}, \text{Min}\})$$

$$\begin{aligned}\text{Phrase}_{\mathfrak{M}} \rightarrow\ & \text{Ton}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}} \\ |\ & \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}}\end{aligned}$$

$$\text{Ton}_{\text{Maj}} \rightarrow \text{I}_{\text{Maj}}$$
$$\text{Ton}_{\text{Min}} \rightarrow \text{I}^{m}_{\text{Min}}$$

A functional model of harmony

$\text{Piece}_{\mathfrak{M}} \rightarrow [\text{Phrase}_{\mathfrak{M}}] \qquad (\mathfrak{M} \in \{\text{Maj}, \text{Min}\})$

$$\begin{aligned}\text{Phrase}_{\mathfrak{M}} &\rightarrow \text{Ton}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}} \\ &\mid \quad \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}}\end{aligned}$$

$$\text{Ton}_{\text{Maj}} \rightarrow \text{I}_{\text{Maj}}$$
$$\text{Ton}_{\text{Min}} \rightarrow \text{I}^{m}_{\text{Min}}$$

$$\begin{aligned}\text{Dom}_{\mathfrak{M}} &\rightarrow \text{V}^{7}_{\mathfrak{M}} \\ &\mid \text{V}_{\mathfrak{M}} \\ &\mid \text{VII}^{0}_{\mathfrak{M}} \\ &\mid \text{Sub}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}} \\ &\mid \text{II}^{7}_{\mathfrak{M}}\ \text{V}^{7}_{\mathfrak{M}}\end{aligned}$$

A functional model of harmony

$$\text{Piece}_{\mathfrak{M}} \rightarrow [\text{Phrase}_{\mathfrak{M}}] \qquad (\mathfrak{M} \in \{\text{Maj}, \text{Min}\})$$

$$\begin{aligned} \text{Phrase}_{\mathfrak{M}} &\rightarrow \text{Ton}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}} \\ &\mid \quad \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}} \end{aligned}$$

$$\text{Ton}_{\text{Maj}} \rightarrow \text{I}_{\text{Maj}}$$

$$\text{Ton}_{\text{Min}} \rightarrow \text{I}^{m}_{\text{Min}}$$

$$\begin{aligned} \text{Dom}_{\mathfrak{M}} &\rightarrow \text{V}^{7}_{\mathfrak{M}} \\ &\mid \text{V}_{\mathfrak{M}} \\ &\mid \text{VII}^{0}_{\mathfrak{M}} \\ &\mid \text{Sub}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}} \\ &\mid \text{II}^{7}_{\mathfrak{M}}\ \text{V}^{7}_{\mathfrak{M}} \end{aligned}$$

$$\begin{aligned} \text{Sub}_{\text{Maj}} &\rightarrow \text{II}^{m}_{\text{Maj}} \\ &\mid \text{IV}_{\text{Maj}} \\ &\mid \text{III}^{m}_{\text{Maj}}\ \text{IV}_{\text{Maj}} \end{aligned}$$

$$\text{Sub}_{\text{Min}} \rightarrow \text{IV}^{m}_{\text{Min}}$$

A functional model of harmony

$$\text{Piece}_{\mathfrak{M}} \rightarrow [\text{Phrase}_{\mathfrak{M}}] \qquad (\mathfrak{M} \in \{\text{Maj}, \text{Min}\})$$

$$\begin{aligned}\text{Phrase}_{\mathfrak{M}} &\rightarrow \text{Ton}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}} \\ &\mid \quad \text{Dom}_{\mathfrak{M}}\ \text{Ton}_{\mathfrak{M}}\end{aligned}$$

$$\text{Ton}_{\text{Maj}} \rightarrow \text{I}_{\text{Maj}}$$
$$\text{Ton}_{\text{Min}} \rightarrow \text{I}^{m}_{\text{Min}}$$

$$\begin{aligned}\text{Dom}_{\mathfrak{M}} &\rightarrow \text{V}^{7}_{\mathfrak{M}} \\ &\mid \text{V}_{\mathfrak{M}} \\ &\mid \text{VII}^{0}_{\mathfrak{M}} \\ &\mid \text{Sub}_{\mathfrak{M}}\ \text{Dom}_{\mathfrak{M}} \\ &\mid \text{II}^{7}_{\mathfrak{M}}\ \text{V}^{7}_{\mathfrak{M}}\end{aligned}$$

$$\begin{aligned}\text{Sub}_{\text{Maj}} &\rightarrow \text{II}^{m}_{\text{Maj}} \\ &\mid \text{IV}_{\text{Maj}} \\ &\mid \text{III}^{m}_{\text{Maj}}\ \text{IV}_{\text{Maj}}\end{aligned}$$
$$\text{Sub}_{\text{Min}} \rightarrow \text{IV}^{m}_{\text{Min}}$$

Simple, but enough for now, *and easy to extend*.

Now in Haskell—I

A GADT encoding musical harmony:

$\mathbf{data}\ \mathsf{Mode} = \mathsf{Maj}_{\mathsf{Mode}} \mid \mathsf{Min}_{\mathsf{Mode}}$

$\mathbf{data}\ \mathsf{Piece} = \forall \mu :: \mathsf{Mode}.\mathsf{Piece}\ [\mathsf{Phrase}\ \mu]$

Now in Haskell—I

A GADT encoding musical harmony:

$\textbf{data}\ \mathsf{Mode} = \mathsf{Maj}_{\mathsf{Mode}} \mid \mathsf{Min}_{\mathsf{Mode}}$

$\textbf{data}\ \mathsf{Piece} = \forall \mu :: \mathsf{Mode}.\mathsf{Piece}\ [\mathsf{Phrase}\ \mu]$

$\textbf{data}\ \mathsf{Phrase}\ (\mu :: \mathsf{Mode})\ \textbf{where}$

$\quad \mathsf{Phrase}_{\mathsf{IVI}} :: \mathsf{Ton}\ \mu \to \mathsf{Dom}\ \mu \to \mathsf{Ton}\ \mu \to \mathsf{Phrase}\ \mu$

$\quad \mathsf{Phrase}_{\mathsf{VI}}\ :: \mathsf{Dom}\ \mu \to \mathsf{Ton}\ \mu \to \mathsf{Phrase}\ \mu$

Now in Haskell—I

A GADT encoding musical harmony:

$$\textbf{data}\ \mathsf{Mode} = \mathsf{Maj}_{\mathsf{Mode}} \mid \mathsf{Min}_{\mathsf{Mode}}$$

$$\textbf{data}\ \mathsf{Piece} = \forall \mu :: \mathsf{Mode}.\mathsf{Piece}\ [\,\mathsf{Phrase}\ \mu\,]$$

$$\begin{array}{l}
\textbf{data}\ \mathsf{Phrase}\ (\mu :: \mathsf{Mode})\ \textbf{where} \\
\quad \mathsf{Phrase}_{\mathsf{IVI}} :: \mathsf{Ton}\ \mu \rightarrow \mathsf{Dom}\ \mu \rightarrow \mathsf{Ton}\ \mu \rightarrow \mathsf{Phrase}\ \mu \\
\quad \mathsf{Phrase}_{\mathsf{VI}}\ :: \mathsf{Dom}\ \mu \rightarrow \mathsf{Ton}\ \mu \rightarrow \mathsf{Phrase}\ \mu
\end{array}$$

$$\begin{array}{l}
\textbf{data}\ \mathsf{Ton}\ (\mu :: \mathsf{Mode})\ \textbf{where} \\
\quad \mathsf{Ton}_{\mathsf{Maj}} :: \mathsf{SD}\ \mathsf{I}\ \mathsf{Maj} \rightarrow \mathsf{Ton}\ \mathsf{Maj}_{\mathsf{Mode}} \\
\quad \mathsf{Ton}_{\mathsf{Min}} :: \mathsf{SD}\ \mathsf{I}\ \mathsf{Min} \rightarrow \mathsf{Ton}\ \mathsf{Min}_{\mathsf{Mode}}
\end{array}$$

Now in Haskell—I

A GADT encoding musical harmony:

$$\textbf{data}\ \text{Mode} = \text{Maj}_{\text{Mode}} \mid \text{Min}_{\text{Mode}}$$

$$\textbf{data}\ \text{Piece} = \forall \mu :: \text{Mode}.\text{Piece}\ [\,\text{Phrase}\ \mu\,]$$

$$\begin{aligned}
&\textbf{data}\ \text{Phrase}\ (\mu :: \text{Mode})\ \textbf{where}\\
&\quad \text{Phrase}_{\text{IVI}} :: \text{Ton}\ \mu \to \text{Dom}\ \mu \to \text{Ton}\ \mu \to \text{Phrase}\ \mu\\
&\quad \text{Phrase}_{\text{VI}}\ \ :: \qquad\qquad\ \ \text{Dom}\ \mu \to \text{Ton}\ \mu \to \text{Phrase}\ \mu
\end{aligned}$$

$$\begin{aligned}
&\textbf{data}\ \text{Ton}\ (\mu :: \text{Mode})\ \textbf{where}\\
&\quad \text{Ton}_{\text{Maj}} :: \text{SD I Maj} \to \text{Ton}\ \text{Maj}_{\text{Mode}}\\
&\quad \text{Ton}_{\text{Min}} :: \text{SD I Min} \to \text{Ton}\ \text{Min}_{\text{Mode}}
\end{aligned}$$

$$\begin{aligned}
&\textbf{data}\ \text{Dom}\ (\mu :: \text{Mode})\ \textbf{where}\\
&\quad \text{Dom}_1 :: \text{SD V}\quad \text{Dom}^7 \to \text{Dom}\ \mu\\
&\quad \text{Dom}_2 :: \text{SD V}\quad \text{Maj}\quad \to \text{Dom}\ \mu\\
&\quad \text{Dom}_3 :: \text{SD VII Dim}\quad \to \text{Dom}\ \mu\\
&\quad \text{Dom}_4 :: \text{SDom}\ \mu \to \text{Dom}\ \mu \to \text{Dom}\ \mu\\
&\quad \text{Dom}_5 :: \text{SD II Dom}^7 \to \text{SD V Dom}^7 \to \text{Dom}\ \mu
\end{aligned}$$

Now in Haskell—II

Scale degrees are the leaves of our hierarchical structure:

```
data DiatonicDegree = I | II | III | IV | V | VI | VII
data Quality        = Maj | Min | Dom⁷ | Dim

data SD (δ :: DiatonicDegree) (γ :: Quality) where
  SurfaceChord :: ChordDegree → SD δ γ
```

Generating harmony

Now that we have a datatype representing harmony sequences, how do we generate a sequence of chords?

Generating harmony

Now that we have a datatype representing harmony sequences, how do we generate a sequence of chords?

QuickCheck! We give Arbitrary instances for each of the datatypes in our model.

Generating harmony

Now that we have a datatype representing harmony sequences, how do we generate a sequence of chords?

QuickCheck! We give Arbitrary instances for each of the datatypes in our model.

. . . but we don't want to do this by hand, for every datatype, and to have to adapt the instances every time we change the model. . . so we use *generic programming*:

Generating harmony

Now that we have a datatype representing harmony sequences, how do we generate a sequence of chords?

QuickCheck! We give Arbitrary instances for each of the datatypes in our model.

...but we don't want to do this by hand, for every datatype, and to have to adapt the instances every time we change the model...so we use *generic programming*:

$$\text{gen} :: (\text{Representable}\ \alpha, \text{Generate}\ (\text{Rep}\ \alpha)) \Rightarrow \text{Gen}\ \alpha$$

Generating harmony

Now that we have a datatype representing harmony sequences, how do we generate a sequence of chords?

QuickCheck! We give Arbitrary instances for each of the datatypes in our model.

. . . but we don't want to do this by hand, for every datatype, and to have to adapt the instances every time we change the model. . . so we use *generic programming*:

$$\text{gen} :: (\text{Representable}\ \alpha, \text{Generate}\ (\text{Rep}\ \alpha)) \Rightarrow [\,(\text{String},\text{Int})\,] \rightarrow \text{Gen}\ \alpha$$

Examples of harmony generation—I

$$
\begin{aligned}
&\mathsf{testGen} :: \mathsf{Gen}\ (\mathsf{Phrase}\ \mathsf{Maj}_{\mathsf{Mode}}) \\
&\mathsf{testGen} = \mathsf{gen}\ [(\texttt{"Dom4"}, 3), (\texttt{"Dom5"}, 4)] \\
&\mathsf{example} :: \mathsf{IO}\ () \\
&\mathsf{example} = \mathbf{let}\ \mathsf{k} = \mathsf{Key}\ (\mathsf{Note}\ \natural\ \mathsf{C})\ \mathsf{Maj}_{\mathsf{Mode}} \\
&\qquad\quad \mathbf{in}\ \mathsf{sample}'\ \mathsf{testGen} \ggg \mathsf{mapM\_}\ (\mathsf{printOnKey}\ \mathsf{k}) \\
&\mathsf{printOnKey} :: \mathsf{Key} \rightarrow \mathsf{Phrase}\ \mathsf{Maj}_{\mathsf{Mode}} \rightarrow \mathsf{IO}\ \mathsf{String}
\end{aligned}
$$

Examples of harmony generation—I

```
testGen  :: Gen (Phrase MajMode)
testGen  = gen [("Dom4", 3), ("Dom5", 4)]

example :: IO ()
example = let k = Key (Note ♮ C) MajMode
          in sample′ testGen >>= mapM_ (printOnKey k)

printOnKey :: Key → Phrase MajMode → IO String
```

```
> example
[C: Maj, D: Dom7, G: Dom7, C: Maj]
[C: Maj, G: Dom7, C: Maj]
[C: Maj, E: Min, F: Maj, G: Maj, C: Maj]
[C: Maj, E: Min, F: Maj, D: Dom7, G: Dom7, C: Maj]
[C: Maj, D: Min, E: Min, F: Maj, D: Dom7, G: Dom7, C: Maj]
```

Examples of harmony generation—II

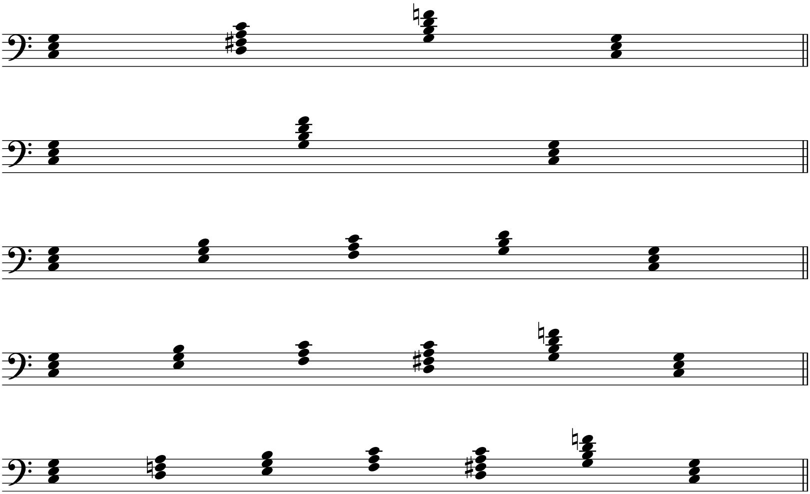

Back to Chordify: chord recognition

Yet another practical application of a harmony model is to improve chord recognition from audio sources.

Chord candidates		0.92 C	0.96 Em
		0.94 Gm	0.97 C
	1.00 C	1.00 G	1.00 Em
Beat number	1	2	3

How to pick the right chord from the chord candidate list? Ask the harmony model which one fits best.

Chordify: architecture

- Frontend
  - Reads user input, such as YouTube/Soundcloud/Deezer links, or files
  - Extracts audio
  - Calls the backend to obtain the chords for the audio
  - Displays the result to the user
  - Implements a queueing system, and library functionality
  - Uses PHP, JavaScript, MongoDB

Chordify: architecture

- Frontend
  - Reads user input, such as YouTube/Soundcloud/Deezer links, or files
  - Extracts audio
  - Calls the backend to obtain the chords for the audio
  - Displays the result to the user
  - Implements a queueing system, and library functionality
  - Uses PHP, JavaScript, MongoDB
- Backend
  - Takes an audio file as input, analyses it, extracts the chords
  - The chord extraction code uses GADTs, type families, generic programming (see the HarmTrace package on Hackage)
  - Performs PDF and MIDI export (using LilyPond)
  - Uses Haskell, SoX, sonic annotator, and is mostly open source

Chordify: numbers

- Online since January 2013
- Top countries: US, UK, Thailand, Philippines, Indonesia, Germany
- Visitors: 3M+ (monthly)
- Chordified songs: 1.5M+
- Registered users: 180K+

Summary

Musical modelling with Haskell:

- A model for musical harmony as a Haskell datatype
- Makes use of several advanced functional programming techniques, such as generic programming, GADTs, and type families
- When chords do not fit the model: error correction
- Harmonising melodies
- Generating harmonies
- Recognising harmony from audio sources

Play with it!

chordify®

http://chordify.net
http://hackage.haskell.org/package/HarmTrace
http://hackage.haskell.org/package/FComp