DECOMPOSING MUSIC

MAKING THE HAPPY VALLEY BAND

BY DAVID KANT, COMPOSER AND BANDLEADER

“a shitty MP3 to MIDI converter”

“the Shaggs meets Guitar Hero”

“the best executed worst idea”

“James Brown backed by Sun Ra”

“substantive acousmatic research”

“pop music heard by a computer algorithm”

(You Make Me Feel Like) A Natural Woman, full score.

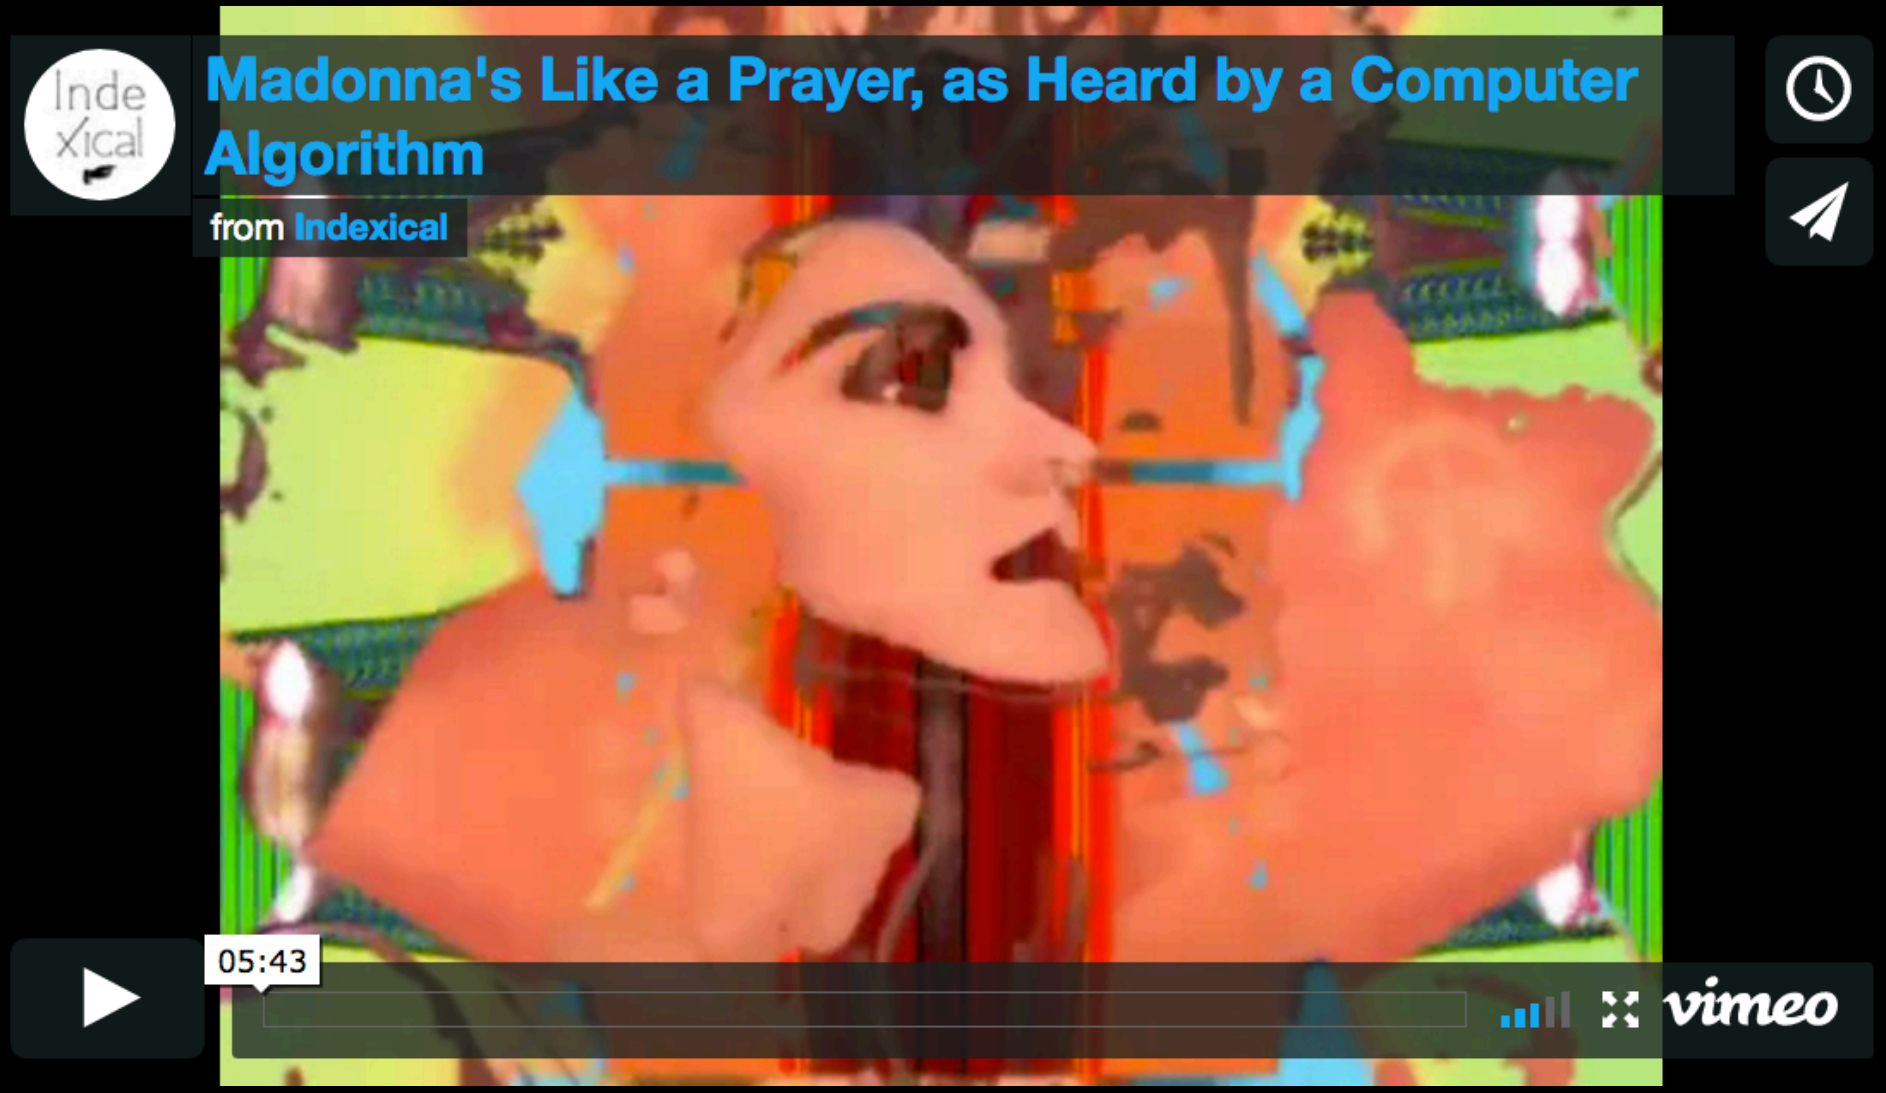
Madonna's Like a Prayer, as Heard by a Computer Algorithm
from Indexical
05:43
vimeo

[1] SEPARATION

[2] TRANSCRIPTION

[3] MUSICAL NOTATION

[4] PERFORMANCE

[1] SEPARATION

1) Xtrk

Spatial filtering based on stereo location cues:

ILD (inter-aural level difference)

$$|X(\omega)| = |X_l(\omega)| + |X_r(\omega)|$$

$$\Theta(\omega) = \frac{|X_l(\omega)|}{|X_l(\omega)| + |X_r(\omega)|}$$

ITD (inter-aural time difference)

$$\tau(\omega) = (\angle X_r(\omega) - \angle X_l(\omega) - \omega\tau)e$$

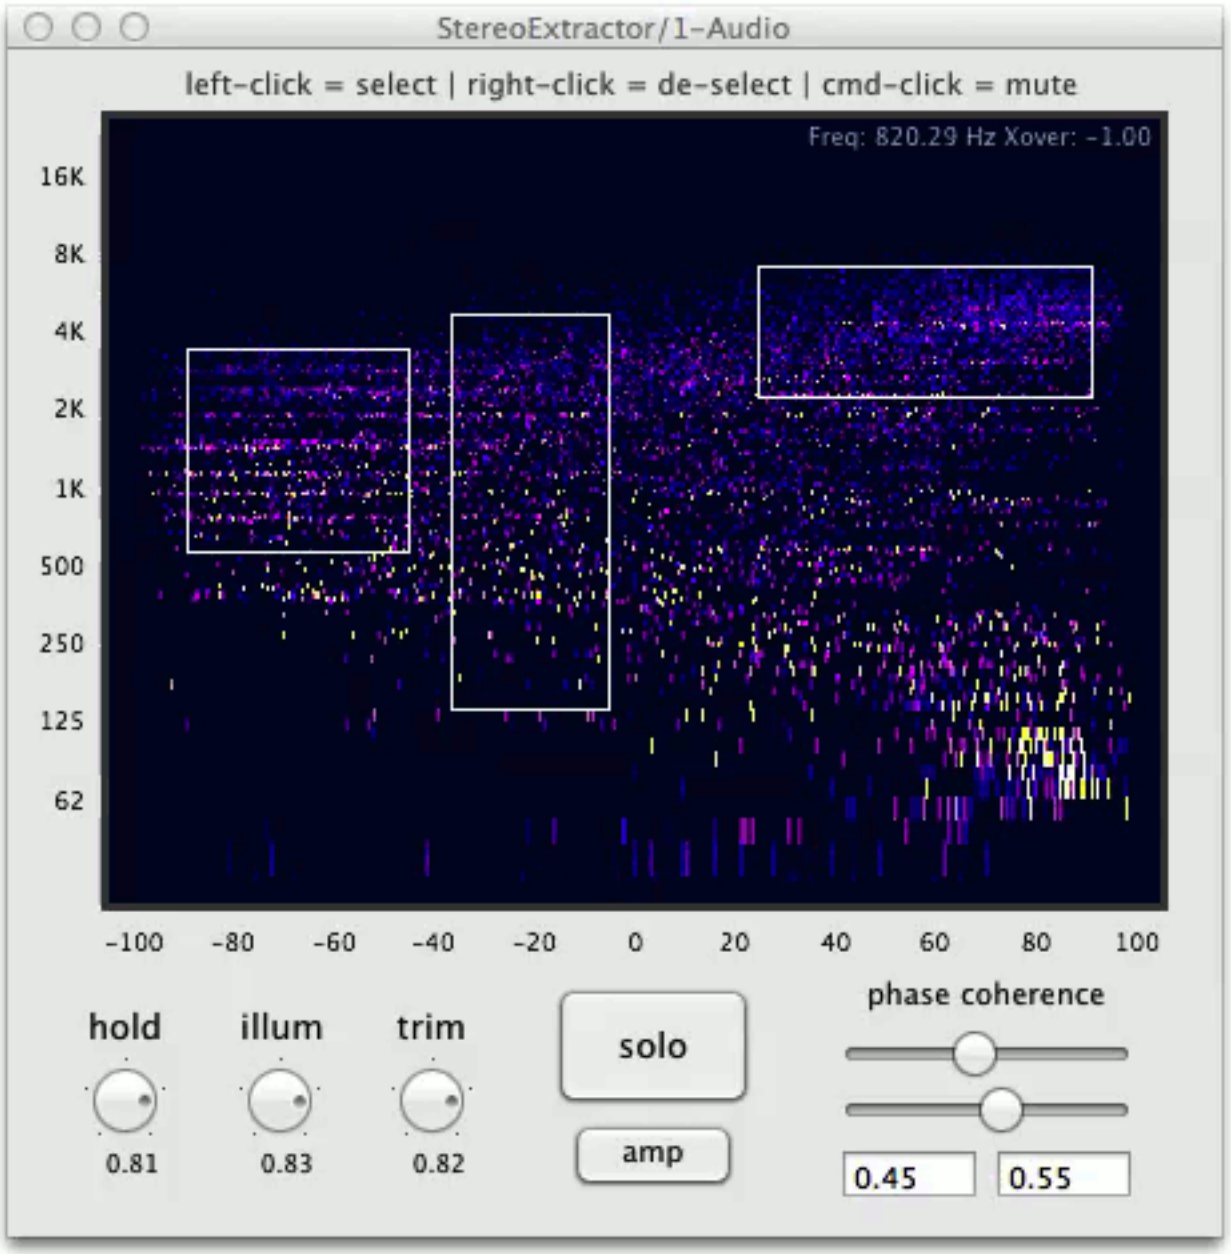

xtrk is an audiounit plug-in that allows real-time visualization and control of spatial filtering

2) harmonic / percussive separation

Full Mix

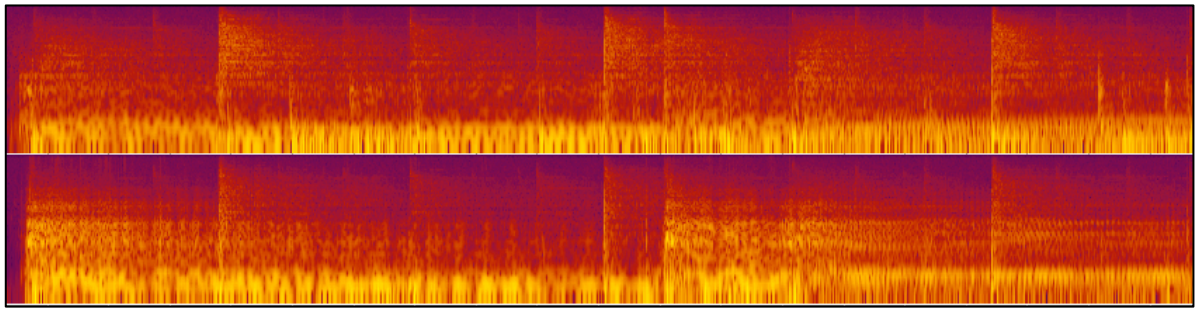

Harmonic Component

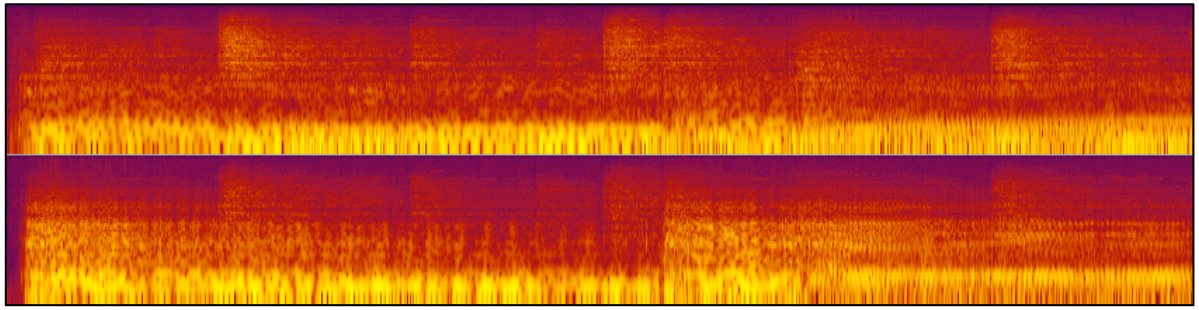

Percussive Component

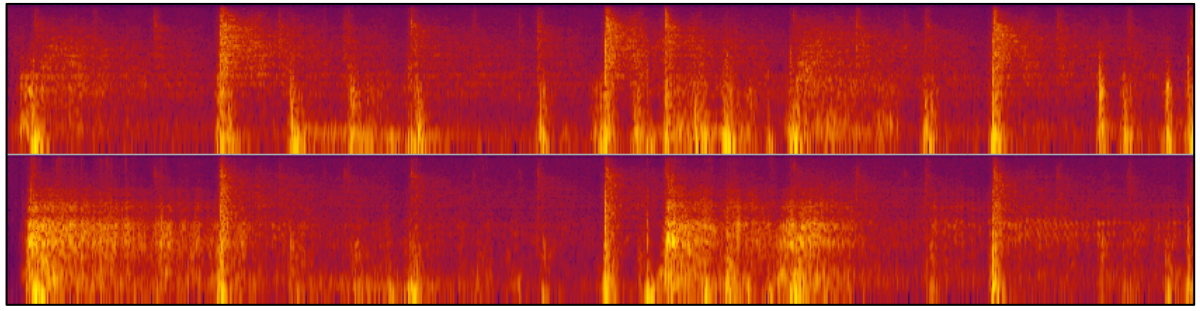

- median filter applied to spectrogram

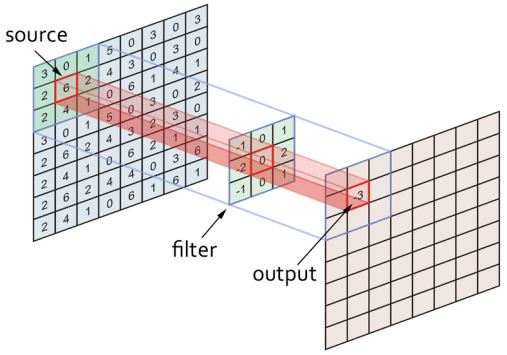

Harmonic: filter across successive frames (horizontal)

Percussive: filter across frequency bins (vertical)

3) PLCA

Probabilistic Latent Component Analysis

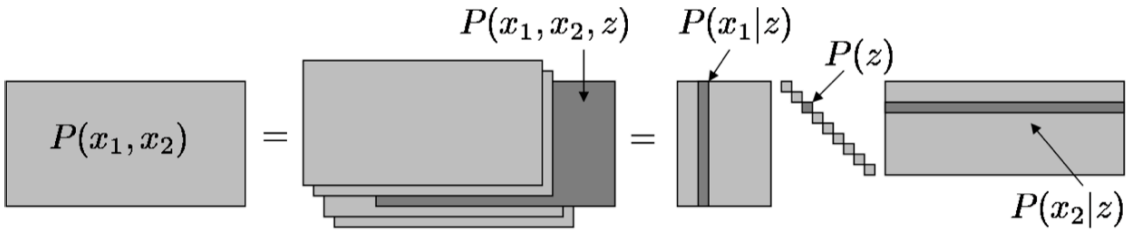

- ***semantically meaningful*** acoustic sources using a latent variable probabilistic model
- statistical formulation allows extensions to machine learning frameworks
  - priors, transformational invariances, and sparsity constraints

Probabilistic Interpretation

- mixture spectrogram is interpreted as a histogram describing the results of repeated draws from a number of different probability distributions
- a set of hidden (or latent) variables governs the chances of which probability distribution is drawn from each time
- each probability distribution correspond a source spectrogram, and the latent variable correspond to the presence of each source in the mixture

The Model

$$P(\mathbf{x}) = \sum_z P(z) \prod_{j=1}^{N} P(x_j|z)$$

- P(x) is represented as the sum of multiple marginal products
- proposed in 2006 by Smaragdis, Raj, and Shashanka

github.com/davidkant/sndsplt

PLCA voice separation on “After the Gold Rush”

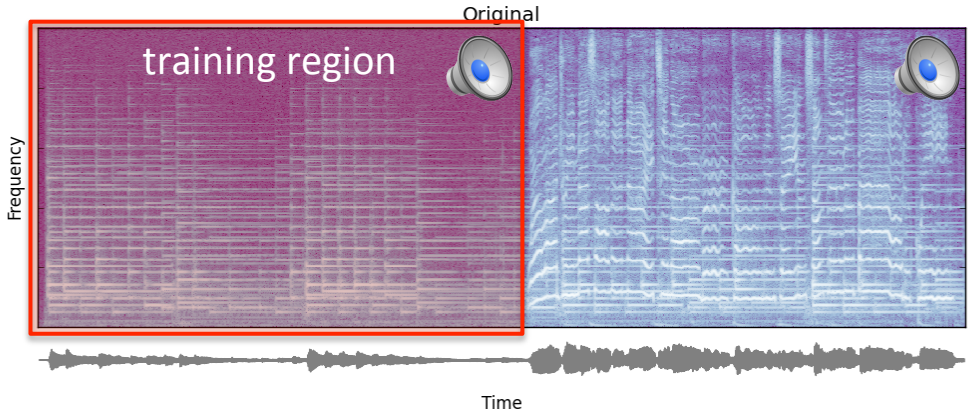

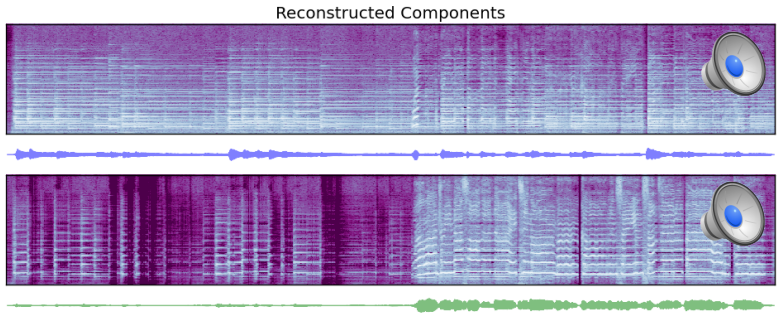

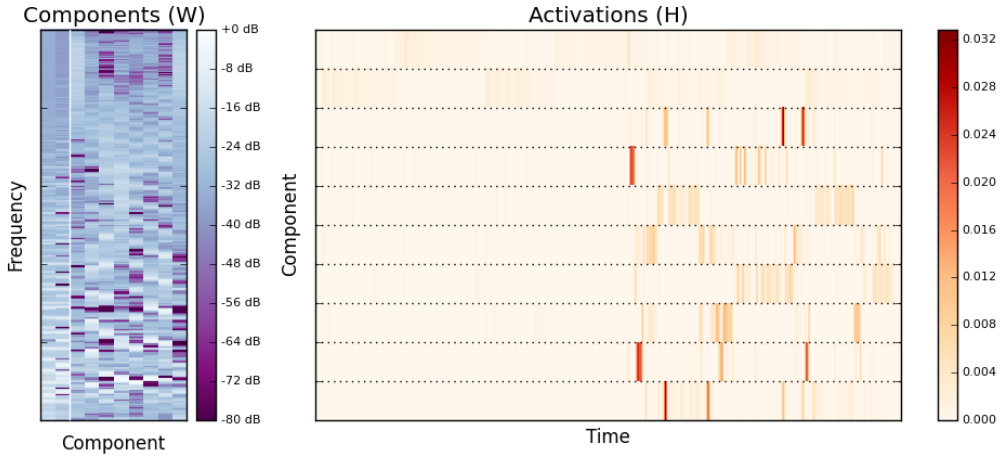

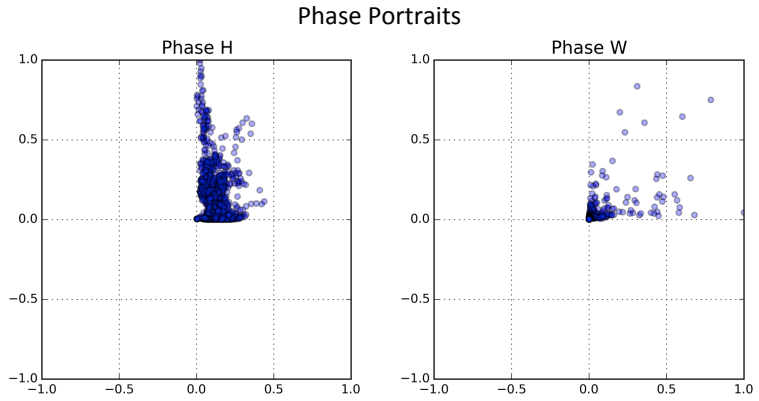

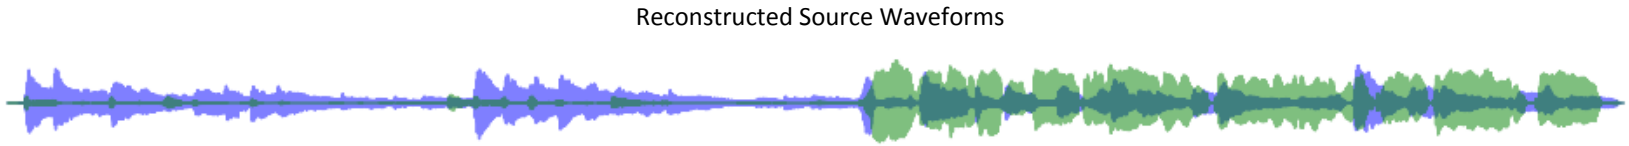

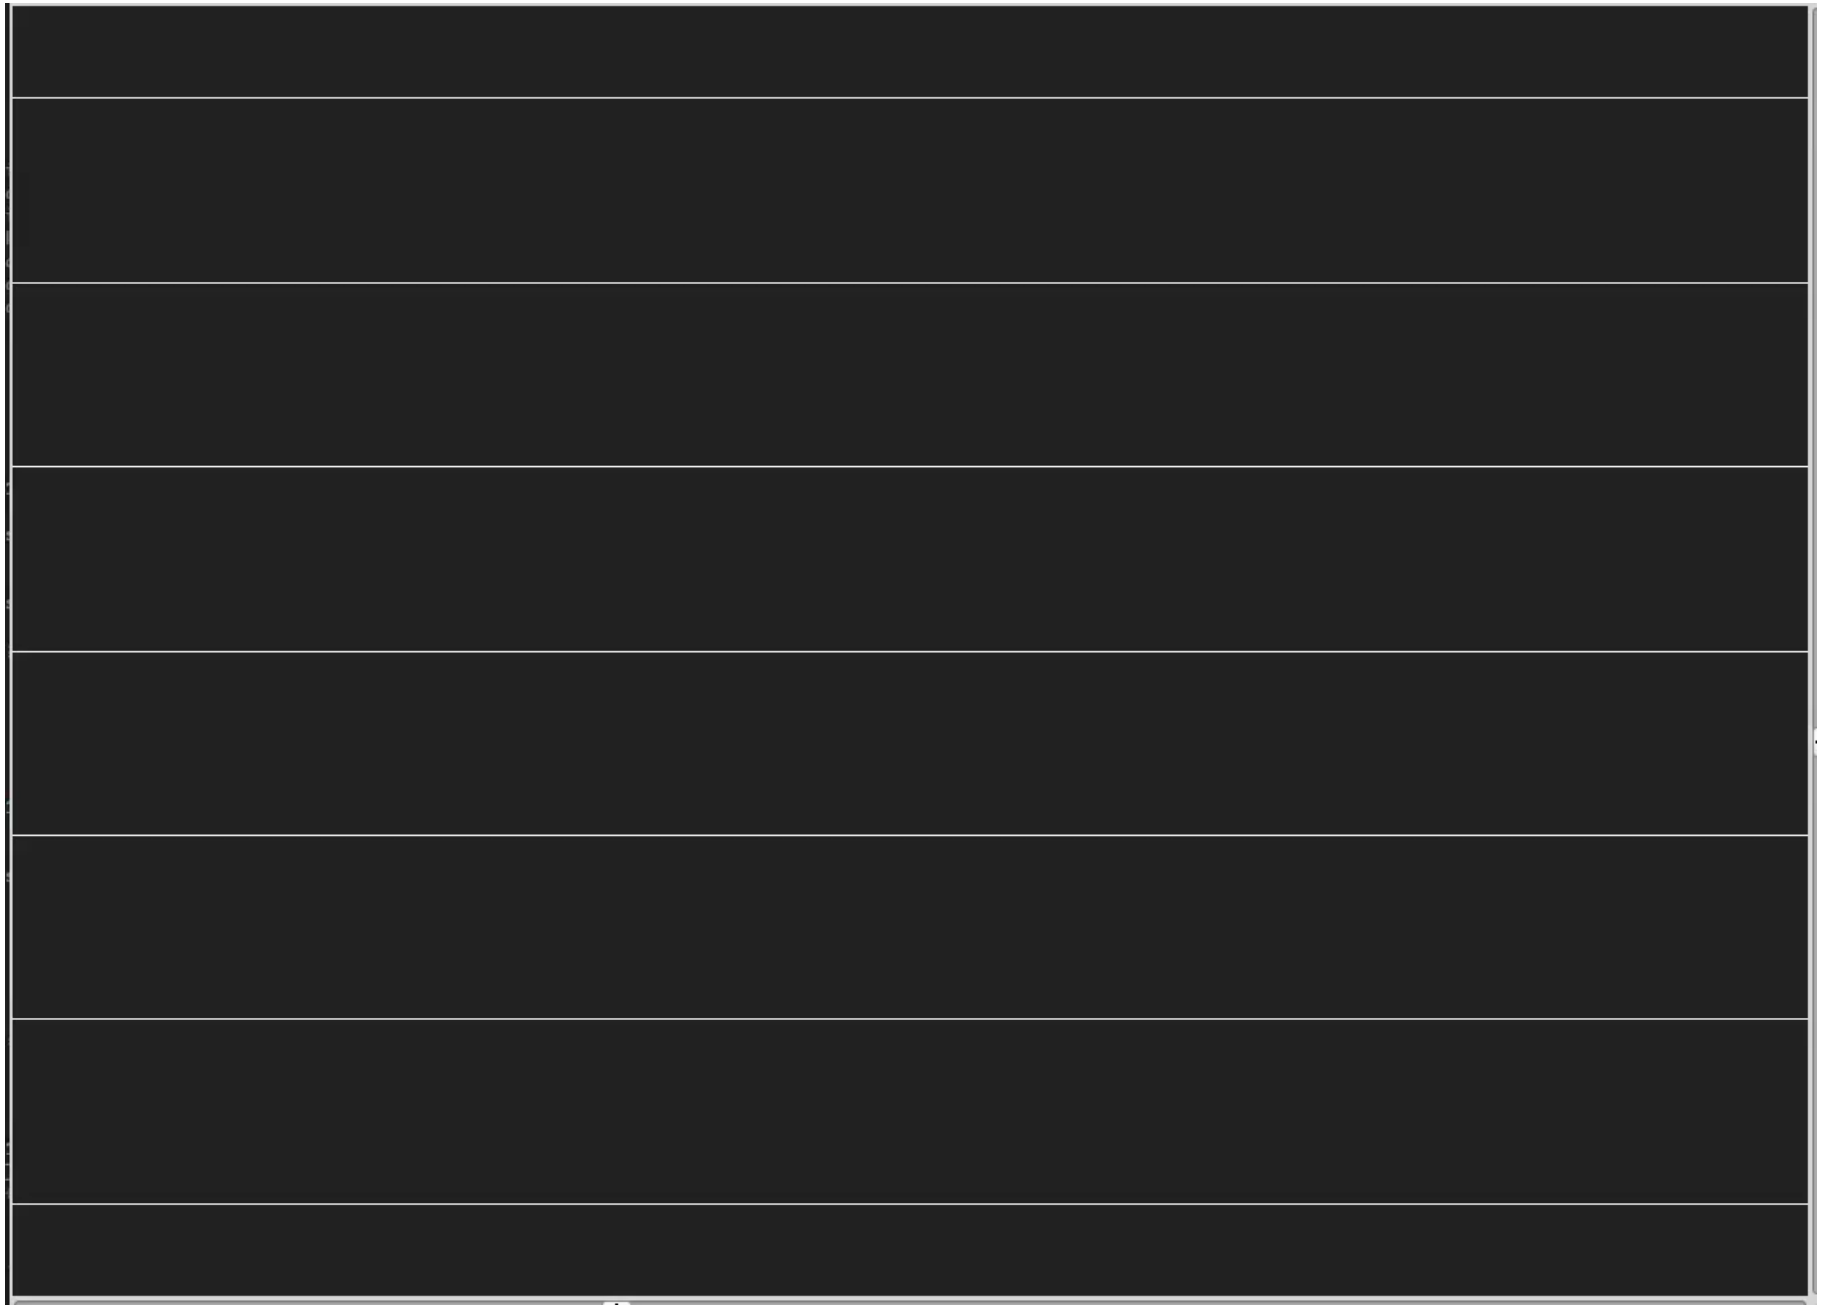

[2] TRANSCRIPTION

1) Pitch Analysis

github.com/davidkant/pyfid

- multiple fundamental frequency estimation based on Miller's Puckette's fiddle~ algorithm
- salience of each frequency depends on the presence of peaks in the spectrum at or near harmonics of *f*
- salience is calculated by a likelihood function *L(f)* where *f* is frequency:

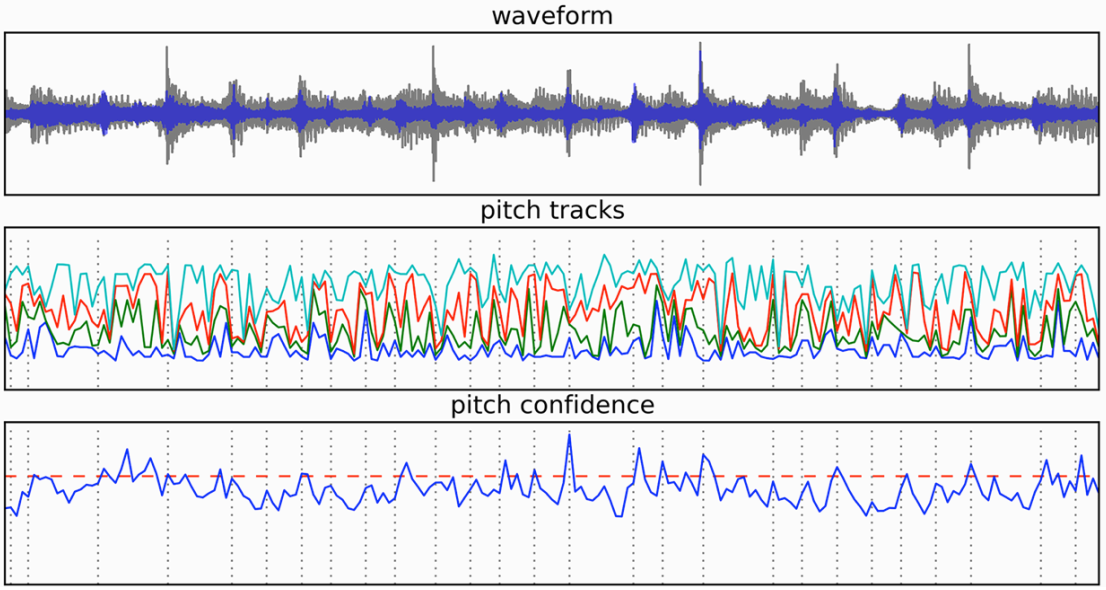

$$L(f) = \sum_{i=0}^{k} a_i t_i n_i$$

- *k* is the number of peaks in the spectrum
- a_i depends on the amplitude of the i^{th} peak
- t_i depends on the distance of the i^{th} peak to the nearest harmonic of *f*
- n_i depends on whether the nearest harmonic is a low or high multiple of *f*

2) Rhythm Analysis

- finding **change** and/or **stability** in an some **audio feature**
- traditionally by finding peaks in an onset detection function (ODF)
- ODF measures change in an audio signal over time, such as amplitude, spectral flux, complex difference, pitch, pitch stability
- peak-picking finds local maxima in the ODF

Finding peaks in amplitude

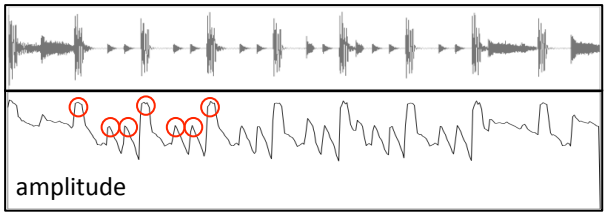

Finding stability in pitch

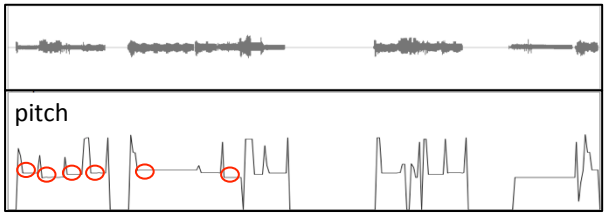

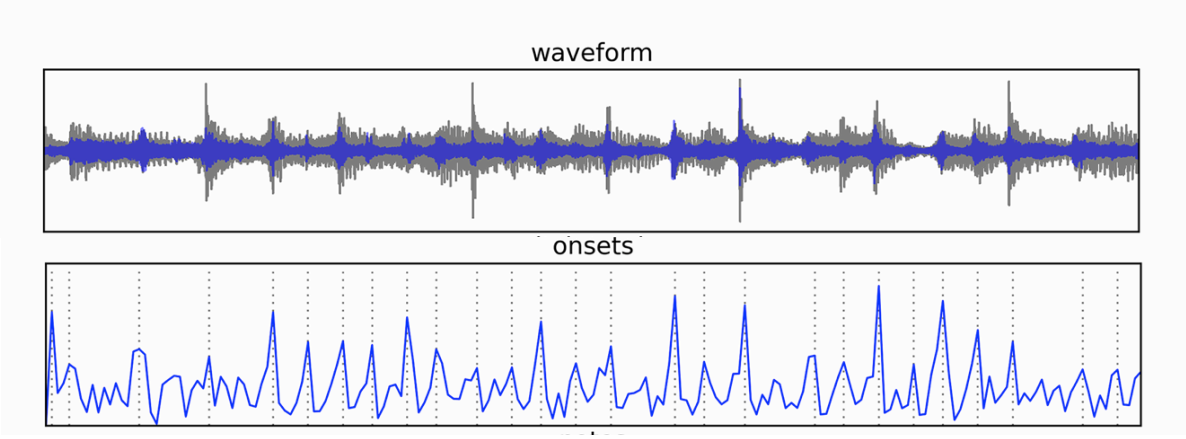

- in HVB, I choose ODFs to match the music, often combining a number of different ODFs, combining pitch with other features

[3] MUSICAL NOTATION

Musical Quantization

github.com/davidkant/trnzscrybe

- convert **time-in-seconds** to **musical rhythms**
- minimize error between actual and quantized time values
- given a scheme of possible subdivisions of the beat

Top: un-quantized time (in seconds); Bottom: error between un-quantized and quantized time

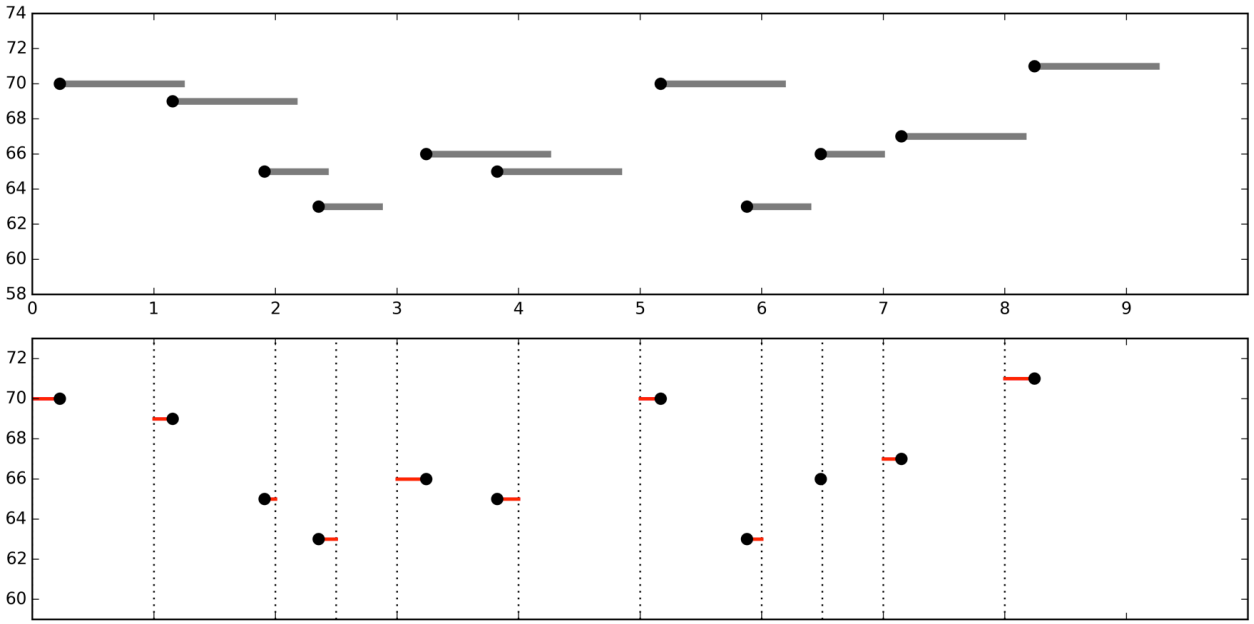

dotted vertical lines shows quantized rhythms

beat division scheme

```
# duration 4 beats
bds.add_tuplet((1,1), 4.0)    # /1
bds.add_tuplet((3,2), 2.0)    # /3
bds.add_tuplet((5,4), 1.0)    # /5
bds.add_tuplet((7,4), 1.0)    # /7

# duration 2 beats
bds.add_tuplet((1,1), 2.0)    # /1
bds.add_tuplet((3,2), 1.0)    # /3
bds.add_tuplet((5,4), 1.0/2)  # /5
bds.add_tuplet((7,4), 1.0/2)  # /7

# duration 1 beat
bds.add_tuplet((1,1), 1.0)    # /1
bds.add_tuplet((2,2), 1.0/2)  # /2
bds.add_tuplet((3,2), 1.0/2)  # /3
bds.add_tuplet((4,4), 1.0/4)  # /4
bds.add_tuplet((5,4), 1.0/4)  # /5
bds.add_tuplet((6,4), 1.0/4)  # /6
bds.add_tuplet((7,4), 1.0/4)  # /7
```

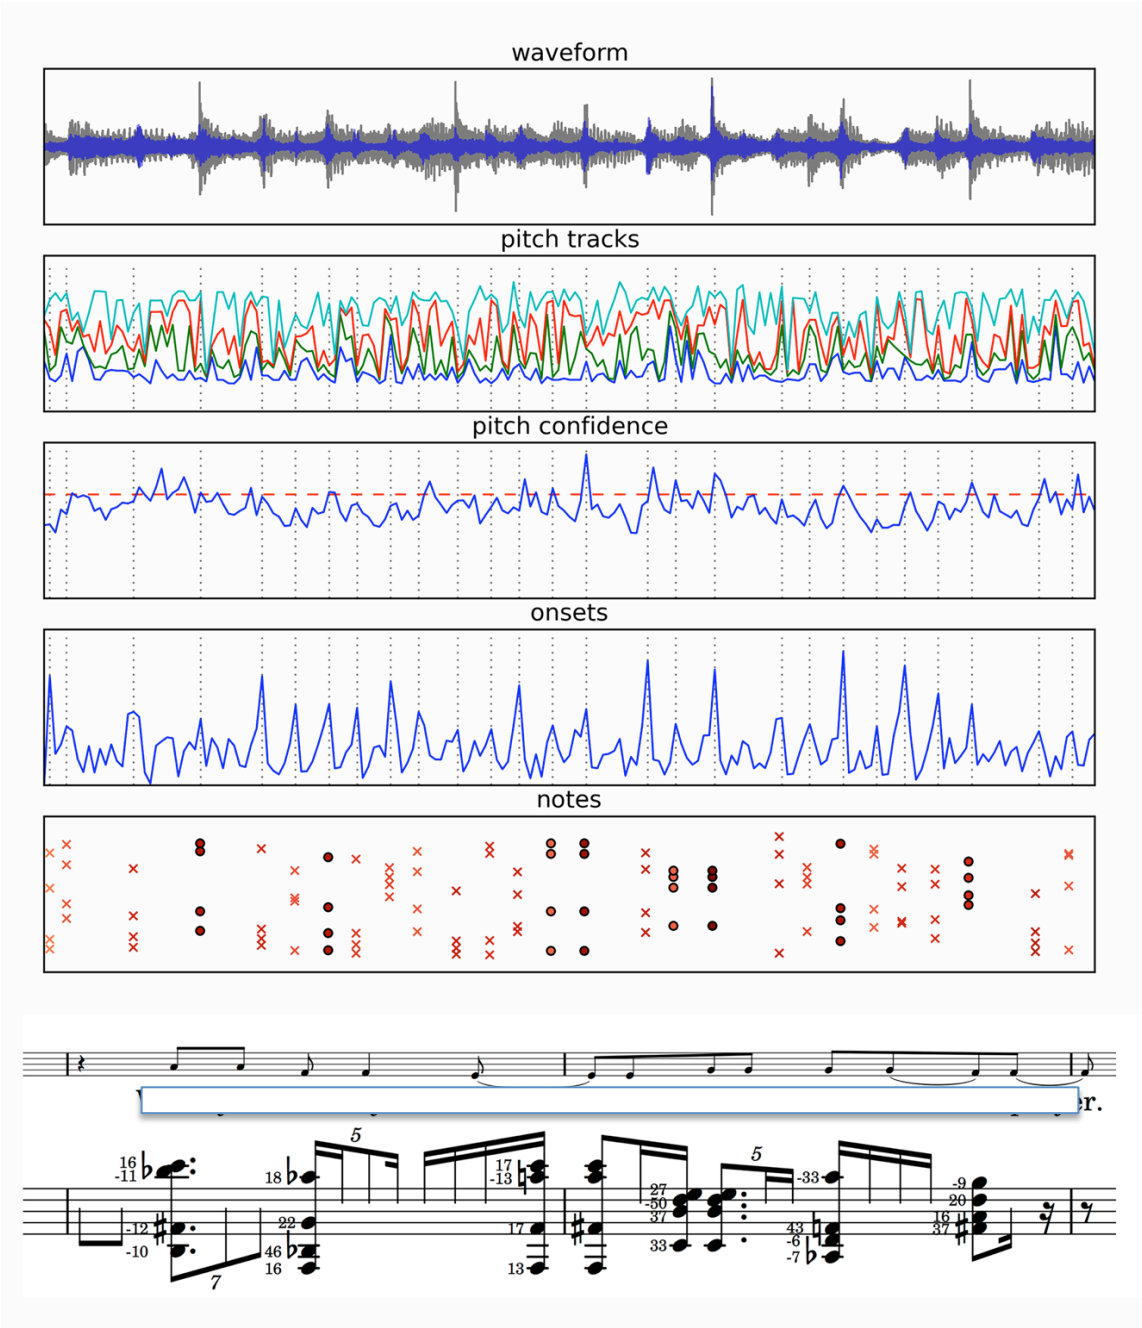

Separation – isolated guitar track is highlighted in blue

Pitch analysis – multiple pitch tracks are shown in different colors

Pitch confidence – low pitch confidences become muted strokes

Rhythm – spectral flux is plotted in blue and peaks indicate note onsets

Piano-roll – pitch versus time display of the parsed notes

Musical notation – rendering into musical notation

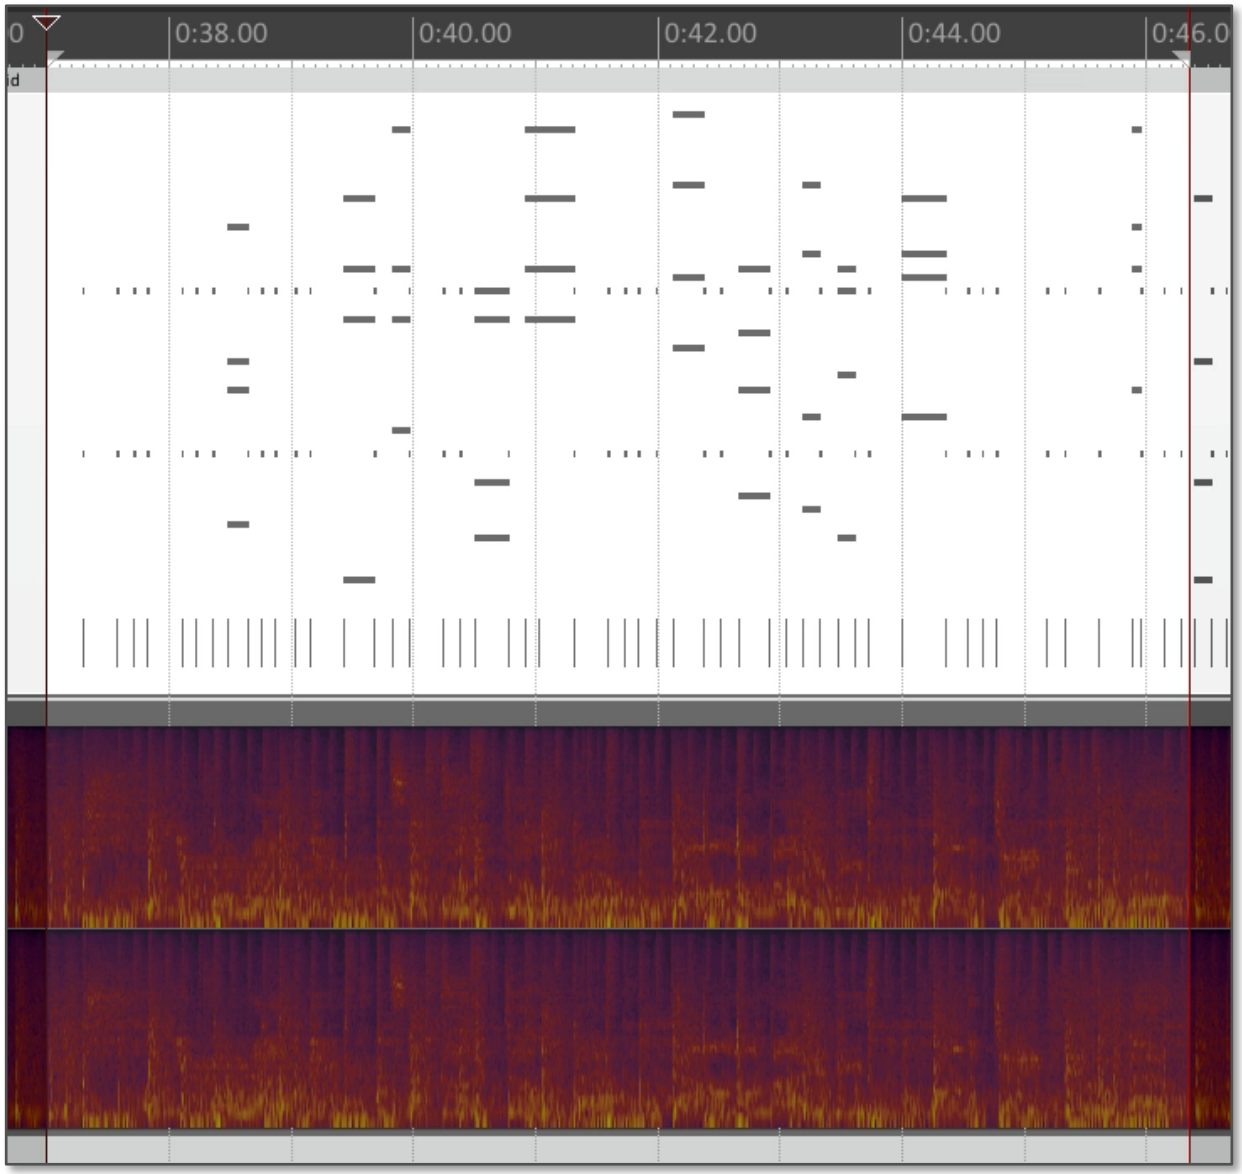
0:38.00
0:40.00
0:42.00
0:44.00
0:46.0

[4] PERFORMANCE

EXCERPTS:

This Guy's in Love With You – Piano

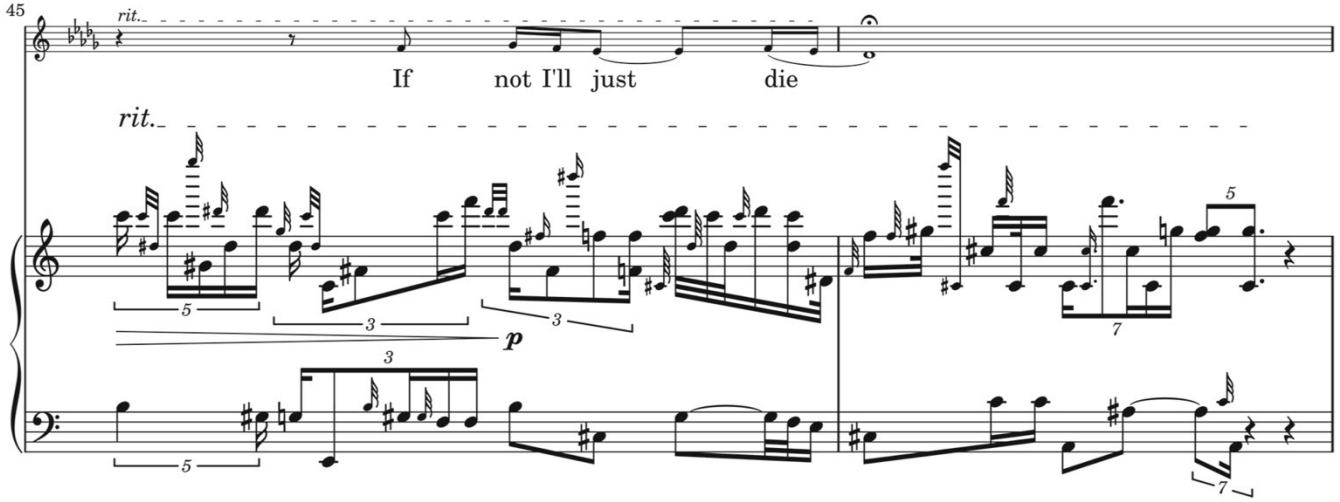

Excerpt: *This Guy's in Love with You*, Piano. For Joseph Kubera's solo piano feature, I left in all of the infinitesimally quick notes and multi-octave leaps that I usually filter out, because he plays piano like a ten-limbed octopus with a separate brain for each digit. The music is riddled with grace notes.

It's a Man's Man's Man's World – Timpani

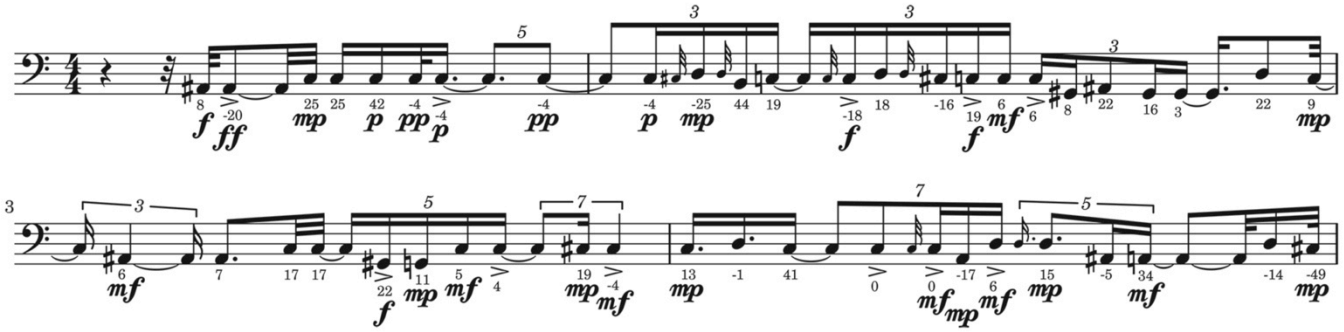

Excerpt: *It's a Man's Man's Man's World*, Timpani. I wanted my transcription to express all of the minute fluctuations of sound (spectra, pitch, dynamics) that are present in a simple timpani roll.[11] The new timpani part is still written for one drum, but with *a lot* of foot pedaling.[12]

Crazy – Upright Bass

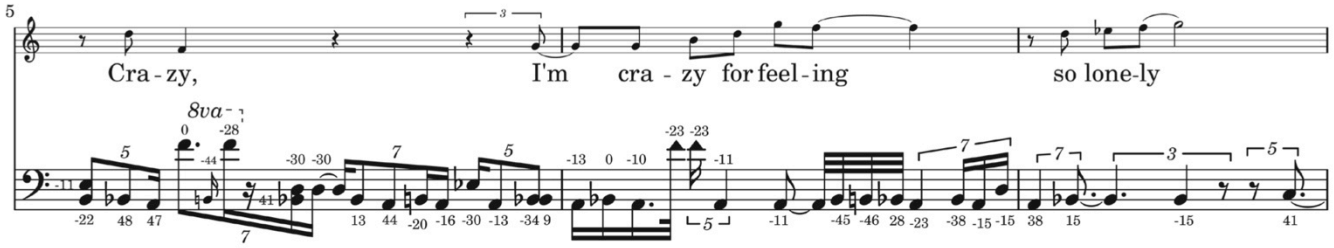

Excerpt: *Crazy*, Upright Bass. This one is a true gem. Sometimes I cannot bring myself to edit a note that is out of playing range. In this case, I write it "as high as possible" (the triangle notehead) and leave it up to bass player Mustafa Walker to deal with. It's now one of my favorite moments in HVB music, ever.

Ring of Fire – Electric Guitar

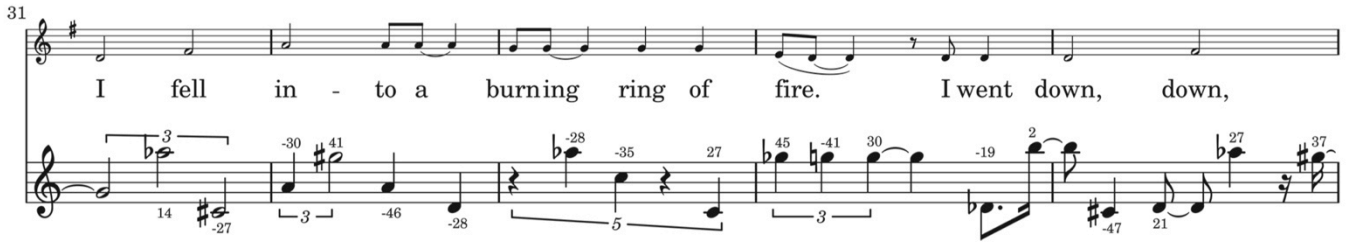

Excerpt: *Ring of Fire*, Electric Guitar. The regularity of the palm-muted single-note electric guitar pulses in the original recording drives the music almost like a percussion instrument, so I tuned my event detection to be more sensitive to rhythmic onsets than to pitched onsets.

Ring of Fire – Trumpets

Excerpt: *Ring of Fire*, Trumpets. Sometimes orchestration is what keeps the song moving, and it doesn't really matter what the band plays, like the alternation between Johnny Cash and the Mariachi trumpets. Timing is more important than the notes.[13]

(You Make Me Feel Like) A Natural Woman – French Horn

Excerpt: *(You Make Me Feel Like) A Natural Woman*, French Horn. I chose this song because I knew Daniel Costello would take his horn part seriously and not object to the characteristically impractical horn leaps. He is the only performer who has ever corrected a microtone.

Like a Prayer – Percussion

Excerpt: *Like a Prayer*, Percussion. I couldn't identify all of the auxiliary percussion instruments by ear, so I used the source separation algorithm instead. I trained it on a large collection of percussion samples and let it decide what was in the recording. The odd thing is I don't remember teaching the computer *Bone Alphabet*...

In the Air Tonight – Electric Guitar

Excerpt: *In the Air Tonight*, Electric Guitar. Another gem in the HVB vault. This transcription of the reverb-soaked, slow guitar bends and sustained lines that open the track is no longer so slow and sustained, but if you listen along with the original, the transcription actually makes some sense.

In the Air Tonight – Violins

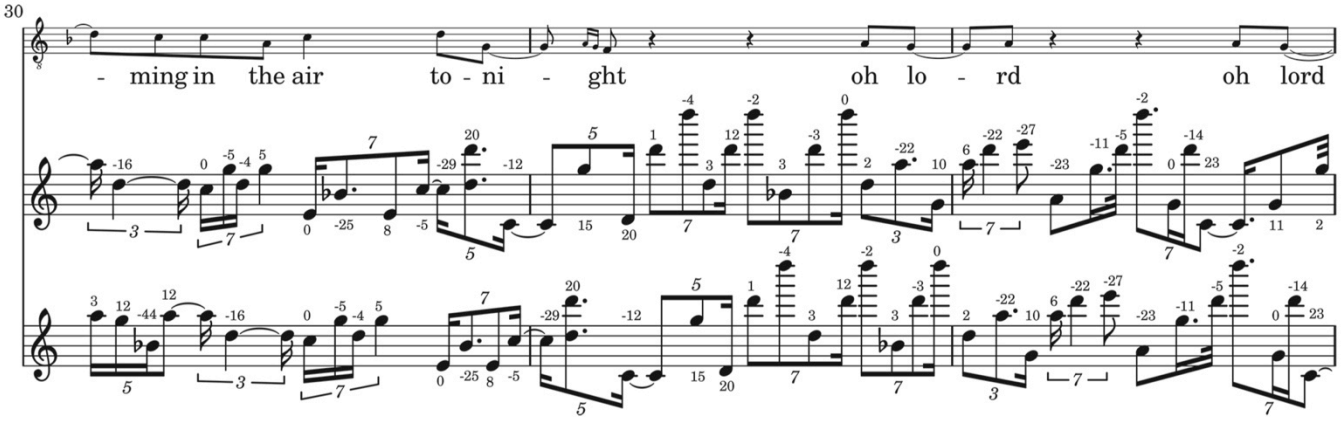

Excerpt: *In the Air Tonight*, Violins. This may be my favorite moment in all of HVB. It's a transcription for two violins of the ethereal, moody, digitally delayed synth pads that accompany Phil Collins' voice in the first section of the song. We do the delay manually by staggering the entrances one beat apart.

Happy Valley Band *ORGANVM PERCEPTVS* | 12” Vinyl LP + Print Booklet

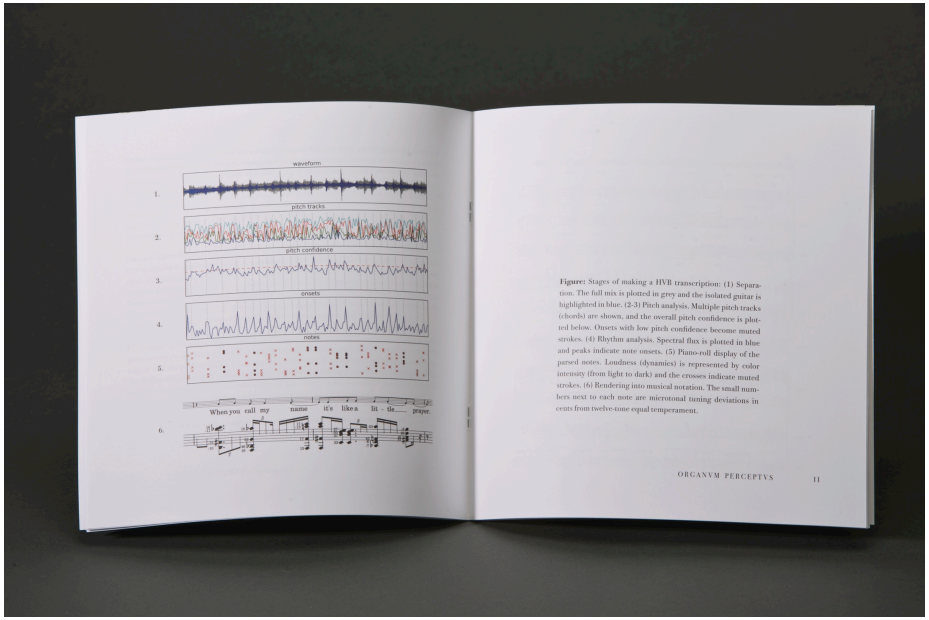

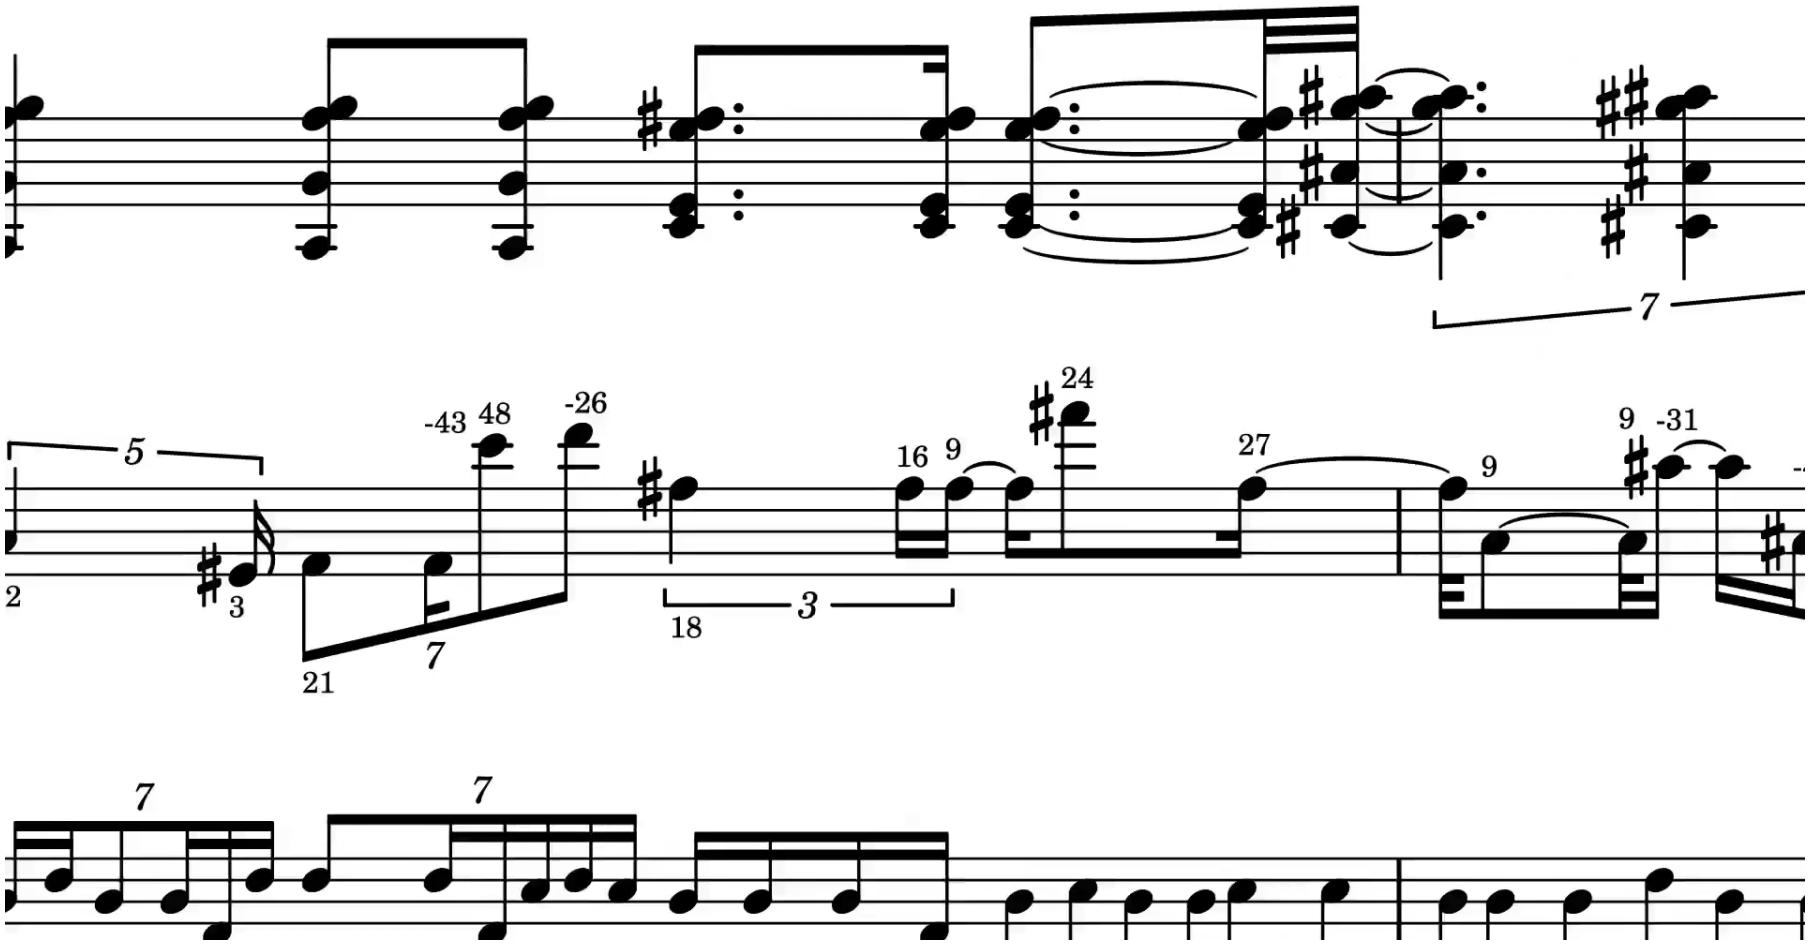
7
5
2
3
-43
48
-26
21
7
18
3
16
9
24
27
9
9
-31
7
7

“conceptually fascinating, but I simply cannot bear to listen to it”

WIRE News In Writing

Album stream and interview: David Kant of Happy Valley Band talks about their ‘machine listening’ album

February 2017

“refracted and amplified through the software in often mystifying ways, resulting in warped interpretations that are unexpected, to say the least” – **Emily Bick, The Wire**

“convinced is an actual recording of my high school’s pep band at a snowy football game when every brass instrument detuned after five minutes out of their cases.” – **Elizabeth Hambleton, New Classic LA**

New Classic LA

MUSIC REVIEWS NEWSFEED FEATURES FILM BLOGS ETC happy valley band

TINY MIX TAPES

NEWS

Happy Valley Band deconstruct pop classics via machine-learning algorithm on debut album ORGANVM PERCEPTVS

by COLIN FITZGERALD · February 10

... Field Items **Happy Valley Band** deconstruct pop classics via machine-learning algorithm on debut ...

"the latest iteration of pop music’s post-human future... delirious and discordant – **Colin Fitzgerald, Tiny Mix Tapes**

KQED Arts

“like nails on a chalkboard. (Actually, nails on a chalkboard sounds much better to me; what does it mean that I prefer pure cacophony to off-time, off-kilter pop music?)” – **Chris Zaldua, KQED Arts**

On February 26, 2017 at 3:06:01 PM, DistroKid (support@distrokid.com) wrote:

DISTROKID

Hi,

We've been notified that one or more of your songs may contain remixes, samples, or other audio that may not be 100% yours.You may only upload audio that you have 100% recorded yourself. Stores won't accept music that contain unauthorized samples, remixes, and so on.

If you have authorization from the original artist or rights-holder, please reply with that evidence. Or, have the original artist/rights-holder contact us (via email, or Twitter DM at @distrokid) stating that you're authorized.

We know this is super inconvenient and we apologize.

UPC:
840095466950

Album title:
Organvm Perceptvs

Sincerely,
DistroKid
http://distrokid.com

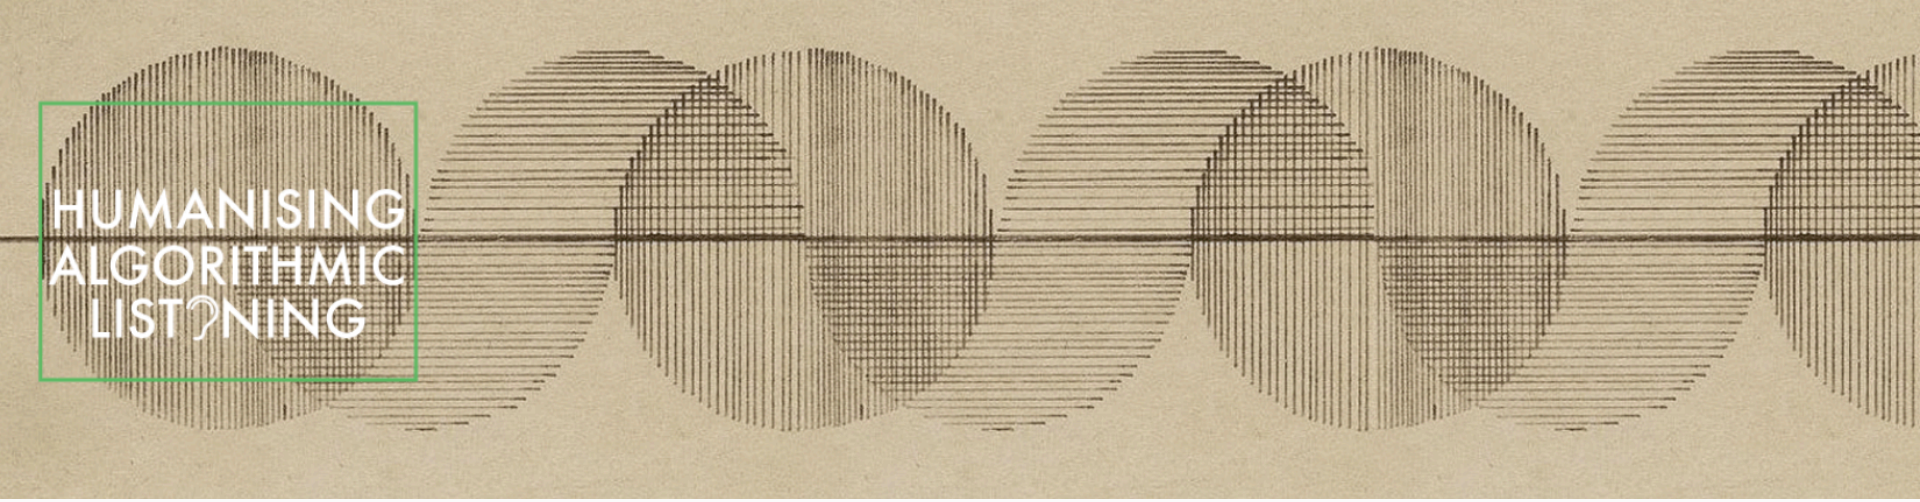

INTRODUCING THE HAPPY VALLEY BAND

Making music through machine ears

By David Kant

Of all the responses to the Happy Valley Band that I have received, this one gives me hope: “as you listen over time and the pieces somehow hold together and get tighter and tighter, you realize that they kind of aren’t strictly errors, there’s some kind of turbo-charged high level thinking going on.” I like to think of this “turbo-charged high level thinking” as a utopian future in which technology extends rather than enslaves our minds, in which we are cyborgs and we are better off for it.

ALL OF THE MUSIC AND NONE OF THE SONGS ©2016 DAVID KANT