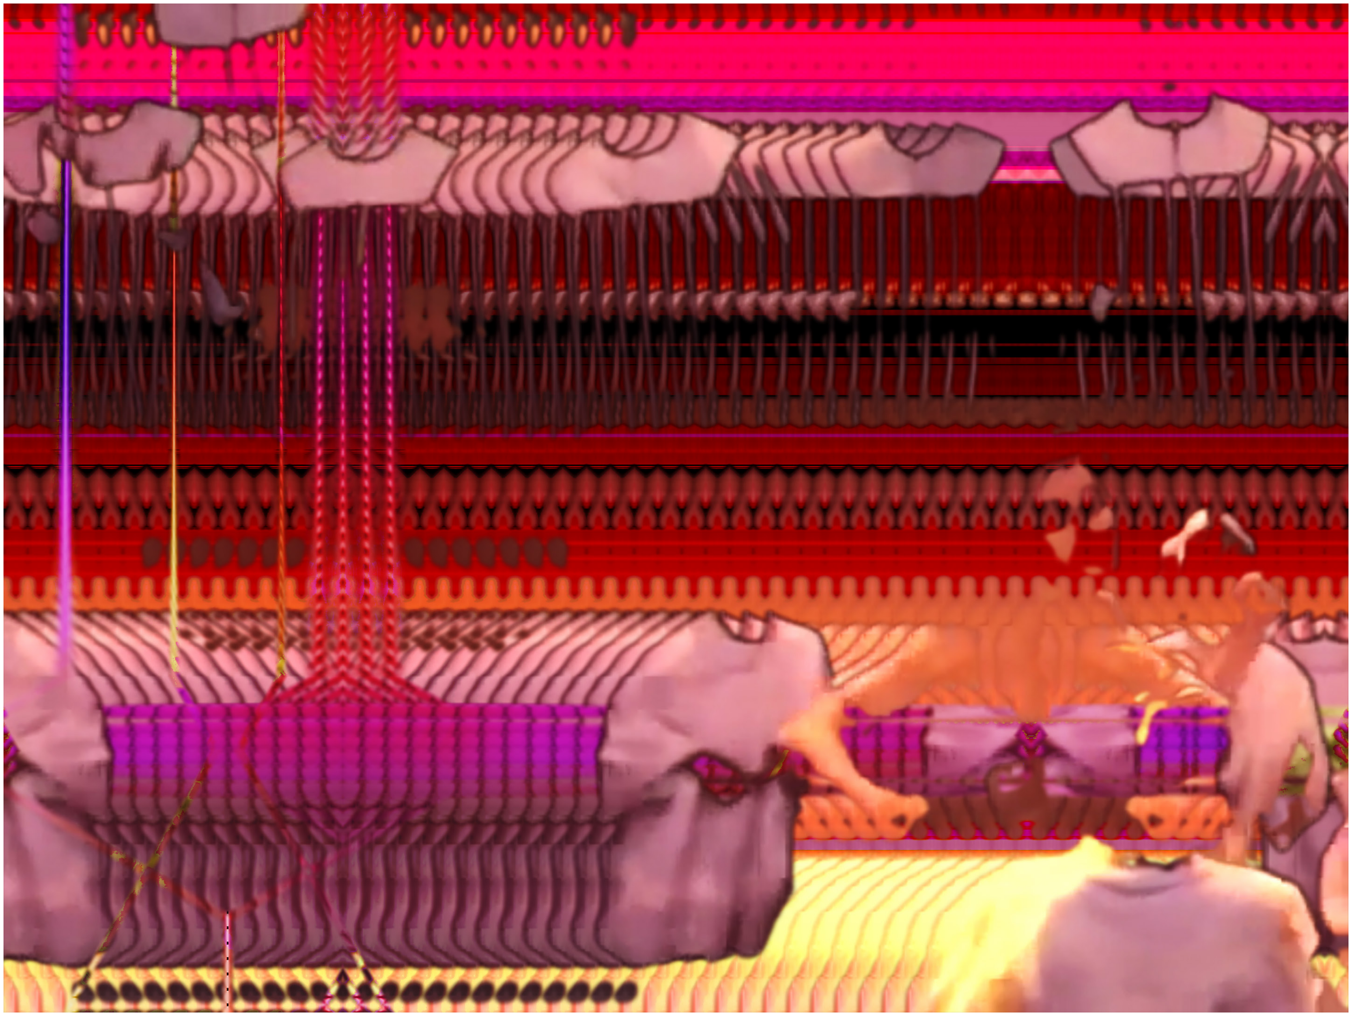

ORGANVM PERCEPTVS

THE LONG ANSWER

I have a confession to make.

Well, it's not a confession, it's more a boast. But boasting would seem so—boastful. So instead I'll call it "journalistic full disclosure": I was once proclaimed the Mayor of Happy Valley.

Perhaps it was because I was at their first and second gig, and then invited them to my program on the radio station WFMU in Jersey City to play their set on air. I'm guessing it's because of this modicum of expertise that I was asked to write liner notes for their debut album.

Little did they know, however, that my knowledge runs deeper than having been at that first gig, on a small stage in a local bar in central Europe during the hot summer of 2011. I have more stories than just that to tell. My people are from Tennessee. I know about Elvis and I know about Patsy. And I'm a born storyteller. This may go on a while.

But first things first. That first gig—it was deep. It was during the biennial Ostrava Days festival in the Czech Republic, a fantastically daring nine days of contemporary composition. Days are long and nights are longer at the festival and one evening that year ended at the Parník Club with performances by some of the festival participants, the Romanian folk band Nea Vasile & Taraful de la Marsa and our heroes, the Happy Valley Band.

The Romanians played first, giving an advantage to the locals in the unfortunate but palpable tension between them and the festival bubble cultural tourists who had descended on their bar. Nobody was likely to throw any punches but there were clearly two audiences there and, perhaps, not many in attendance with the linguistic skills to speak to both of them. Two remarkable things happened during that set. First, the young, Czech composer Lucie Vítková lifted the hems of her skirt and took the stage, giving time-honored footsteps to the traditional rhythms and building at least a bit of a bridge between the factions. Then, later in the set, the band brought up a guest vocalist with thick sideburns and a blue sequined jumpsuit. It would have been a mistake to call him an Elvis impersonator but it also would have been a mistake to assume that he was unfamiliar with the American King of Rock'n'Roll.

Everything that night seemed to lead up to the utterly fantastic, bizarre, unanticipatable performance by the Happy Valley Band. They played songs by Elvis Presley and Patsy Cline. They made a fairly unholy racket. And despite the fact that there seemed to be a Romanian Elvis impersonator in the house, they performed with isolated vocal tracks of Presley himself. It was a glorious shambles.

The population in Happy Valley has pretty much quadrupled since that night in 2011, and there are some fantastically talented people on this record. Along with the core band, which includes mastermind David Kant and the violin duo String Noise (Conrad Harris and Pauline Kim Harris, themselves no strangers to extreme reconfigurings of pop and rock songs), there's the pianist Joseph Kubera, a noted interpreter of John Cage and Morton Feldman, two composers frequently heard at the Ostrava Days festival. There is also the remarkable percussionist Chris Nappi, who (like Kubera) is a regular presence at Ostrava Days. The soprano Charlotte Mundy, who has earned some positive acclaim for her performances of Feldman and Schoenberg, here embodies June Carter and Aretha Franklin. The band's take on the Burt Bacharach tune "This Guy's in Love With You" includes Sam Friedman, whom I once saw in a Wesleyan University student ensemble led by Anthony Braxton. I don't know if it was audacity or arrogance but he played the whole of the concert on mouth harp.

When you listen to this record, if you haven't already, you will be faced with trying to understand just what it is they're doing. As I understand it, they use automatic notation software that generates scores from mp3s but which doesn't work very well, trick it out to work even worse and then play exactly what the computer spits out. But I could be wrong. Look around, they might have put a better explanation here somewhere.

But enough about them. I promised you stories. My mama grew up in Tennessee, just 90 miles from Memphis, and I spent a lot of time on the family cotton and soybean farm when I was growing up. My mama (born just a month after Presley) respected him as a good Tennessee boy, even if she didn't like all that sneering and hip shaking. He bought his own mama a house, after all, and she respected that. But it was songs by Johnny Cash and Hank Williams that she'd sing me at bathtime and bedtime. Nevertheless, in 4th grade I spread a rumor in my school that I was related to the King. Years later, my mama had a boyfriend obsessed with the death sites of country stars and Rebel Army war heroes. He kept albums of photos of the sites he'd visit, and he showed me his pictures of the hole, still unfilled, where Patsy Cline's plane went down. [Cue side A, track 5.] There was a memorial plaque at the bottom of the hole which he didn't think was right, being as some people might want to see it who can't climb down there. He had a point.

As do I, mine being that I know the music from which the Happy Valley blossomed, the fertile ground nourished by the Mississippi River from which they sprung. I wish I could say there's something in what they're doing that reflects the true spirit of the songs they cover, or casts an important light on the region's history, or makes the music relevant for a post-copyright age, or proves some salient point about the influx of YouTube and social media and everyone now being a performer of some sort or another, but I can't. The truth of the matter is I don't really know what the hell they're doing, but I can say that I love this record. And I hope you don't get it either. Understanding things is vastly overrated in contemporary society.

—Kurt Gottschalk
August, 2016

DECOMPOSING MUSIC
MAKING THE HAPPY VALLEY BAND

BY DAVID KANT, COMPOSER AND BANDLEADER

The Happy Valley Band has been called a lot of things: "a shitty MP3 to MIDI converter," "the Shaggs meets Guitar Hero," "James Brown backed by Sun Ra," and "the best executed worst idea." Everyone seems to have their own explanation of what it is and how it works.

The Happy Valley Band is pop music heard by a computer algorithm and (re-)performed by humans. It's about training a machine learning system to hear pop songs: to distinguish between the individual instruments in a band, to determine the notes, pitches, and rhythms played, and to notate it all as standard musical notation. The Happy Valley Band plays what comes out. The machine generating the music has no concern for the limits of human performance. The music is extremely complex, literal, and at times impossible to play. This album, *ORGANVM PERCEPTVS*, is about bringing to life this inhuman music.

I *wish* I could say the Happy Valley Band uses "a shitty MP3 to MIDI converter" that I grabbed online. This would have saved me thousands of lines of code and countless hours of time. The technology I use to make HVB music is all homegrown software that I built, based on ideas from state-of-the-art machine listening and digital signal processing. But it's not *exactly* about getting machines to hear like (we presume) humans do. I became fascinated by a way of hearing that is unique to technology. How does music sound when filtered through a machine process? How can machines help us to hear differently?

The first incarnations of the process hacked together a few pieces of commercially available software, but, as the project continued, I felt the need to build my own tools. I wrote a musical notation transcriber, a polyphonic pitch tracker, onset detectors, a source separation interface, stereo location filters, and a bunch of other junk. I wanted to take responsibility for the music I made and get my hands on the algorithm design. That's where the composition is—HVB music is the artifact of machine hearing algorithms.

The project only took shape when I found people who would play this crazy music. I and four of my closest friends and musical collaborators (Mustafa Walker, Andrew C. Smith, Beau Sievers, and Alexander Dupuis) formed a dedicated band. The music was unlike anything that we had ever played before—compatible senses of humor probably helped. We gave our first performance at a local bar deep in the Czech Republic, an impromptu concert of Ostrava Days 2011 Festival.[1] The band has been growing ever since.

1 See Kurt's liner notes for the full story of our first gig.

HOW IT REALLY WORKS: THE SOFTWARE

SEPARATION

The first step is "source separation." I start with an original recording and isolate the individual voices and instruments from one another. It's like "un-mixing" a record, and for HVB, this usually amounts to separating the bass, drums, guitars, piano, horns, strings, and singing voice. I end up with a bunch of new audio tracks, one for each instrument solo. Or, almost. Sounds bleed together—sometimes in interesting ways—and these artifacts make it all the way into the notes and rhythms that the band eventually plays.

The ability to tell one sound from another is one of the most important things that we do with our ears, whether distinguishing our predator from our prey or navigating city traffic. We do the same with pop music, tracking guitars and drums through a haze of artificial reverb and other modern mixing techniques. We are able to tell instruments apart, follow the individual voices in a band, and even identify new sounds and learn to recognize them. It's not always so easy, however, which makes writing an algorithm to mirror this process even more challenging.

Designing a source separation algorithm raises a number of basic problems about how humans hear. Does perception rely on innate, fundamental cognitive grouping laws? To what extent do we learn to hear from our past experiences? Is hearing influenced by expectations of the sounds that one thing or another might make? These problems are not so easily resolved and implemented in software. There are many theories of human perception, and there are at least as many approaches to machine hearing.

I use a machine learning algorithm.[2] The algorithm decomposes sound (a rock band, for example) into individual pieces called components (which are, hopefully, the instruments the make up that rock band), but the algorithm doesn't know much at all about sound, let alone the human perception of sound.[3] It's not built to model the physiology of human hearing, nor is it hardwired with cognitive mechanisms and perceptual rules. Instead, the algorithm operates solely on statistical features. It treats sound spectra as the outcome of a probabilistic process and finds components by maximizing statistical independence between them. I "teach" the algorithm, in some sense of the word, everything else—how we hear, and how sound works—by example. It learns from data, essentially coming up with a statistical understanding of what things "sound like."

I first find a "basis" for each instrument—a statistical representation of what the instrument should, in general, sound like—and then decompose the entire mix into a collection of "bases" for the band. I often pull example clips from the song itself, which requires some careful and clever use of filters, usually on stereo location and frequency, or I use sample sets. (I all too often end up using the built-in Apple General MIDI Synth as a training set for drums.) I repeat this process for each instrument in the band, and eventually end up with the last remaining element: the voice.

---

2 Paris Smaragdis and Bhiksha Raj's Probabilistic Latent Component Analysis
3 Let alone how much James Brown paid his backing band.

The results are not perfect. Instruments bleed together, parts of instruments get left out, and the process can even introduce its own little sound anomalies. The reasons for these artifacts are suggestive: the inability of the model to fully capture the complexities of how sound mixes acoustically (destructive as well as constructive interference), of mathematical features (like statistical independence) to correspond to human hearing, or of the training process to properly account for how our prior knowledge, experience, and expectations influence our perception, not to mention artifacts of the underlying mathematical representation of sound (Fourier analysis and resynthesis).

Either way, it is fascinating to think that these artifacts are, in *some* sense, the parts of one sound that actually sound more like another, albeit to a machine. At times, the algorithm mistakenly uses the guitar to reconstruct part of the absent singer's voice, or pieces of the piano's attack fuse with a snare drum hit. I could write all of these off as mistakes, except that the artifacts often find structures in the sounds that I have never heard before and now cannot unhear. [4]

For HVB, composing[5] is teaching the algorithm, developing training sets, adjusting priors, and tweaking convergence thresholds. I believe that we learn to hear through experience, and that our hearing is contingent. I like to think of machine learning as a way to build model worlds with unique circumstances, histories, experiences, and expectations. What does it sound like if I decompose a rock band into 50 pieces? How about 500 pieces? How do you make sense of the E Street Band if you've never heard an electric guitar? Or "Suspicious Minds" if you've never heard a Hammond B3 organ? Composing HVB music is like building hypothetical worlds and listening through them.

---

4 Which is, we hope, how you'll feel about these songs after listening to this record.
5 "composing"

TRANSCRIPTION

The next step is to turn these separated audio tracks back into something *the band* can play. This means translating sound into the musical features that musicians are used to reading—notes, chords, melodies, rhythms—and then writing it all down in musical notation.

The question is how to express a perceptual phenomenon—like the experience of pitch, rhythm, or harmony—in quantifiable terms that can be measured in an audio signal, coded in software, and executed by a machine. The interesting part is deciding what these terms mean and how to measure them. We might share a common understanding of pitch, but when we actually get down to defining pitch clearly enough for a machine to understand the concept, there are *lots* of different definitions and *lots* of different algorithms. It all depends on what models of perception you employ.

The real world, however, does not always fit our models so perfectly. Nothing works all of the time. Pitch, for example, is often understood in terms of the period of the audio waveform, but, in reality, waveforms are not exactly periodic, and it is not always so easy to measure the periodicity of a quasiperiodic waveform. My pitch analysis finds the frequencies whose harmonics line up most closely with the spectral peaks of the sound I am analyzing.[6] The process is similar to finding a best-fit line through a dataset, where the slope of the line corresponds to the fundamental frequency as we hear it. The complication is that harmonics of different fundamental frequencies may overlap. Like finding a best-fit line, there usually is not a perfect match, and the solution is more about balancing error than finding an exact solution.

6 The algorithm is similar to Miller Puckette's Pure Data object [fiddle~].

Digital analysis can take these measurements very often—too often, actually. There is a mismatch between the scales of machine analysis and human performance. An unrelenting, continuous stream of 172.27^{th} notes at a tempo of 120 beats per minute (roughly 86 notes per second) that jumps up and down with each and every tiny pitch fluctuation bears no resemblance to the actual music—in most cases—and would be too much for most performers to handle.

The goal is to carve something out of this massive data set that actually looks like music (to a performer). I have to decide where notes begin and end, what pitch to play when the pitch analysis refuses to settle into a stable reading, how loud to play it, which articulations, and what playing techniques to use. This means developing a measurable criteria for how to organize all of this analysis data into musical structures that humans will eventually play, and doing so automatically and algorithmically.

Unfortunately, there isn't any one universal way that works for everything. I often find new pitches by identifying moments where the pitch analysis remains relatively constant, or percussive hits by identifying abrupt increases in amplitude or spectral flux. There are countless other ways to measure change in sound—phase deviation, complex deviation, or even statistical and information theoretical measures such as Kullback-Leibler divergence. Each of these captures slightly different dimensions of what our ears hear (and what they don't). I take a variety of measures and use the most interesting, surprising, or unexpected. I try to find the action in the signal.

Ultimately, writing HVB music comes down to choosing what to notate and what to ignore. I do want the original to be in there—at some level—and I try to keep enough from the original tune such that a listener can still connect the dots.[7] The fun of HVB is transcribing this material in a literal way, with all the chaos that comes out.

7 Okay… *some* dots.

NOTATION

Finally, I write it all down for the band to play. The previous steps don't know anything about music as it exists on the page. Rhythms, pitches, and dynamics are measured to an exacting degree of numerical resolution, with no concern for what can actually be expressed in written notation. Writing this music down means translating the precise decimal values of the computer analysis into a conventional way of writing sound: musical notation.

Notation has its limits. It has its own grids for things like pitch, rhythm, and dynamics, that are circumscribed by aspects of the notation system, such as tempo, meter, divisions of the beat, modulations, and tuning. Musical notation can express a fine degree of resolution, but at the cost of being complex and impractical to read in performance. Translating time-in-seconds (more realistically, time-in-milliseconds) into musical rhythm often requires very small beat divisions and complex tuplets, and the most accurate rendering might be an illegible sea of seventeen-tuplets. Notating this music is a matter of balancing approximation with overwhelming complexity.[8]

HVB transcription plays with the problem of deciding what is intentional in the original recorded performance. Imperfections in the playing do not necessarily mean that the music is written that way, or that the performers even meant to play it that way. It's what gives music feeling and style—the facility with which Patsy Cline sings around the beat, pushing and pulling at the tempo, and how her band responds. It is also what the inhuman precision of a drum machine is missing. When is it necessary to notate these fluctuations and when should I smooth them out?

---

8 Like Brian Ferneyhough interrupted by Kool and the Gang's roadie telling you to get down.

In a way, translating this data into musical notation explores the very limits of the medium, and the inability of written notation to express certain complex nuances of what the ear can hear or what the musicians on these recordings actually play. It triangulates an imagined musical score (never really there in the first place) from a digital analysis of the recorded performance, with all its minute detail and inflection.[9]

9 The E Street Band didn't really play it like that, did they?

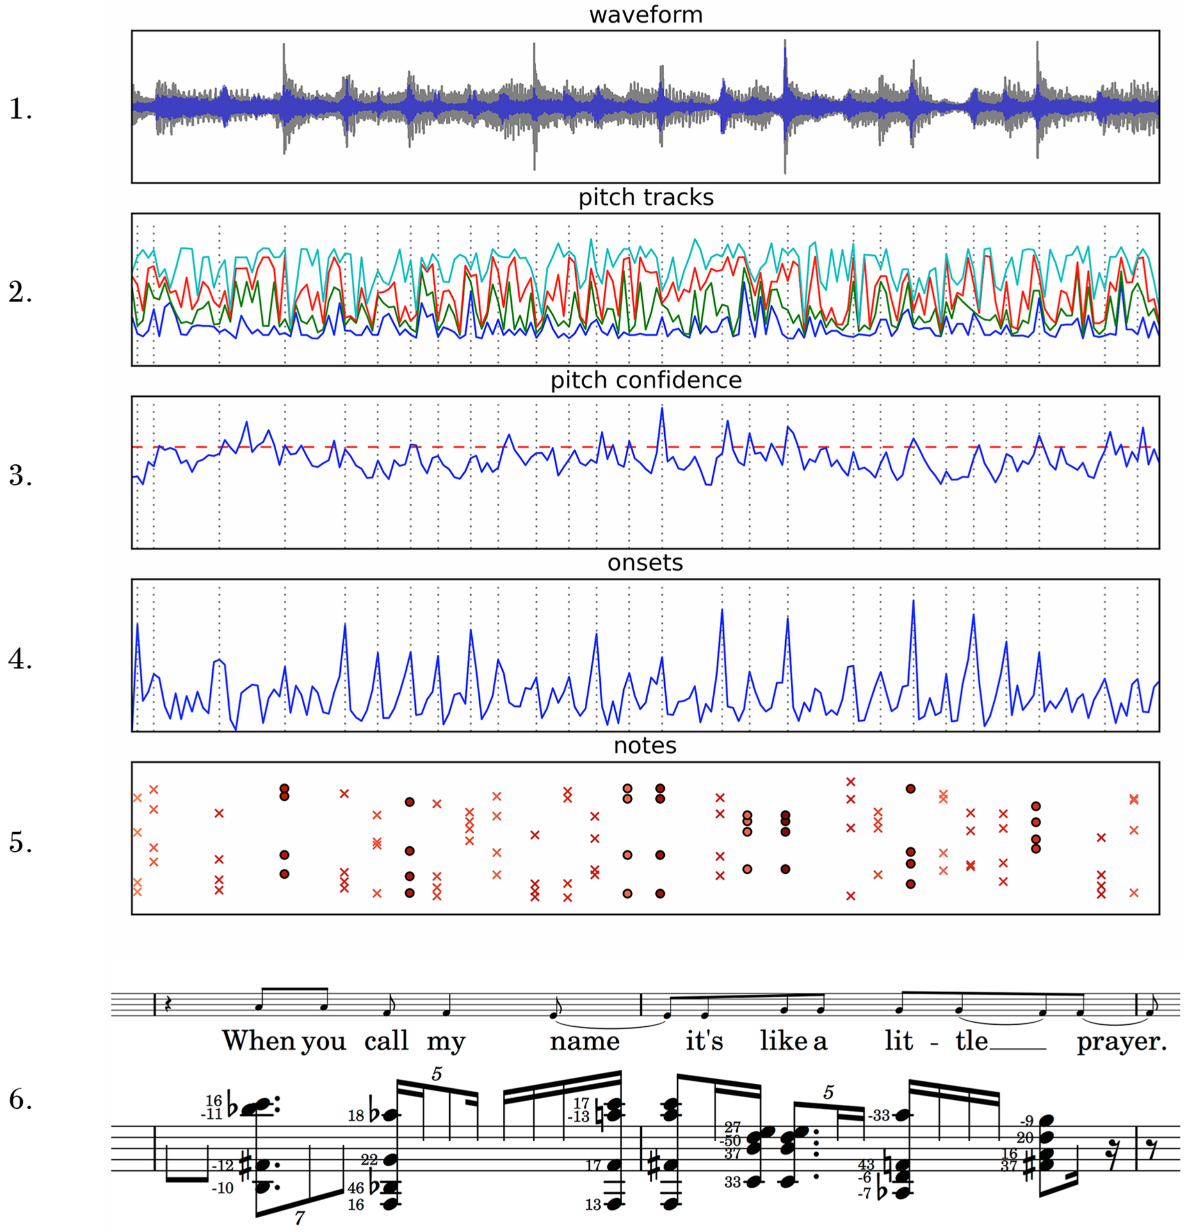
1.
waveform
2.
pitch tracks
3.
pitch confidence
4.
onsets
5.
notes
6.
When you call my name it's like a lit - tle prayer.

Figure: Stages of making a HVB transcription: (1) Separation. The full mix is plotted in grey and the isolated guitar is highlighted in blue. (2-3) Pitch analysis. Multiple pitch tracks (chords) are shown, and the overall pitch confidence is plotted below. Onsets with low pitch confidence become muted strokes. (4) Rhythm analysis. Spectral flux is plotted in blue and peaks indicate note onsets. (5) Piano-roll display of the parsed notes. Loudness (dynamics) is represented by color intensity (from light to dark) and the crosses indicate muted strokes. (6) Rendering into musical notation. The small numbers next to each note are microtonal tuning deviations in cents from twelve-tone equal temperament.


C
26
you make me feel you make me
Sop.
You make me feel You make me
Mez.
You make me feel You make me
Alt.
You make me feel You make me
Vln.
Vln.
Horn.
Pno.
Bass
Kit

Excerpt: *(You Make Me Feel Like) A Natural Woman*, full score. The song as heard and transcribed by the machine hearing process. The music is notated to the pulse of the original song—where the beat actually falls in the recording rather than to an arbitrary pulse. In performance the band plays along to the extracted lead singing voice, which is notated as the top staff in each player's part.

THE BAND: PLAYING THE IMPOSSIBLE

And then we play it. The music is too much information. Pitch is notated to the nearest cent (1/100th of a half-step—imagine 100 extra keys crammed between every note of the piano keyboard) and rhythms are quantized to fine degrees of complex, nested, and incomplete tuplets, each beat surgically dissected into a complicated hierarchy of subdivisions. This complexity gives the music a degree of indeterminacy.[10] It's not always possible to play everything, so performers pick and choose what to focus on. The music is so specific that it is no longer specific, and it ends up being like a musical inkblot: people find what they want in it—a hint of Cecil Taylor or Sun Ra, Ruggles or Ives.

Playing this music has been a process of figuring out how to play what is on the page rather than changing it. The transcriptions are like found objects whose meaning we have to interpret, and there is no consistent criteria from song to song. It emerges from the group. There is something about this overspecificity that gives people room, and the performers find all sorts of interesting things that would have been lost had I filtered them out of the music using my own idea of what I think is possible for *someone else* to play. Sometimes they complain the music is impossible, but I don't change the notes.

I'm often asked how much of the music is made up live. Well, none of it is made up. It's all written down on the music page in some form or another. But like any music, the exact events can change from performance to performance. It's not always easy to tell which notes are artifacts of the process and which are "correct," so the original tune comes in and out of focus.

10 To put it *very* mildly.

WHAT THE FUCK AM I HEARING?

It doesn't always sound right. People like to say the computer is getting it *wrong* because it's bad at hearing. My response is, not quite. It's not wrong; it's *too right*. Machine analysis can be extremely precise, whether in analyzing pitch, measuring loudness, or separating the voice from a full band recording. The analysis algorithms are simply numerical processes that we design. They do their jobs and they do them well. Too well, in fact, and the results of these numerical models do not always correspond to how we actually hear. The "errors" are not mistakes so much as they are the differences between how we expect an algorithm to operate and how it actually operates. We often oversimplify what algorithms really do, abstracting away the details of complex and thorny numerical procedures—which have outliers, exceptions, idiosyncratic artifacts, and caveats—into neater explanations, like pitch or rhythm.

In a way, the production techniques of the original recordings come through in the transcriptions. You can hear a difference between older recordings, in which instruments are hard-panned, and modern mixes with stereo effects, delays, and reverbs. When the recordings are more difficult to separate there's more crosstalk in the HVB music. The voice is the tell; we play along to the extracted vocals, and the quality of its separation hints to the artifacts in the rest of the separated instruments. You don't hear those artifacts directly, but rather as they are captured in the audio analysis, transcribed into music notation, and played by the band.

EXCERPTS:

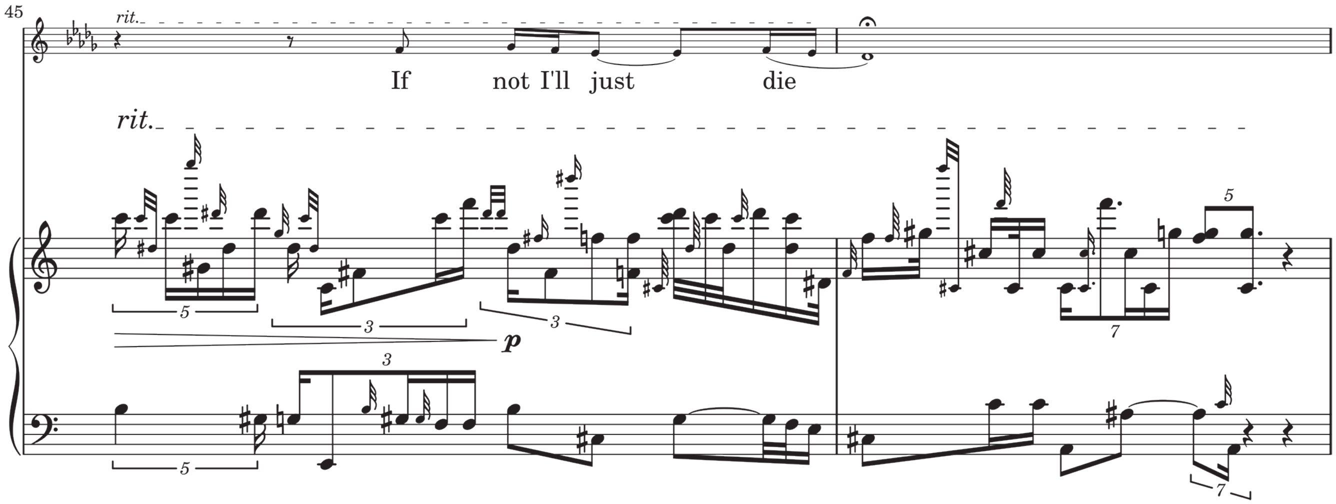

Excerpt: *This Guy's in Love with You*, Piano. For Joseph Kubera's solo piano feature, I left in all of the infinitesimally quick notes and multi-octave leaps that I usually filter out, because he plays piano like a ten-limbed octopus with a separate brain for each digit. The music is riddled with grace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

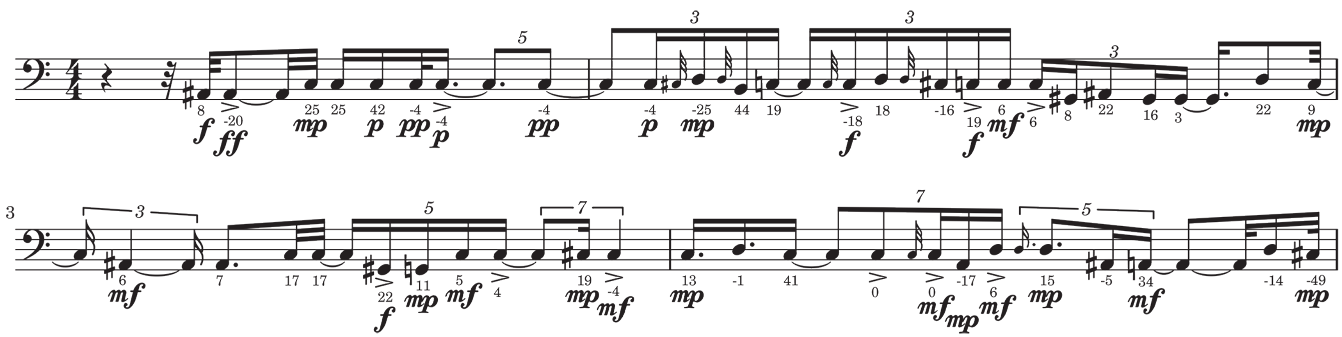

Excerpt: *It's a Man's Man's Man's World*, Timpani. I wanted my transcription to express all of the minute fluctuations of sound (spectra, pitch, dynamics) that are present in a simple timpani roll.[11] The new timpani part is still written for one drum, but with *a lot* of foot pedaling.[12]

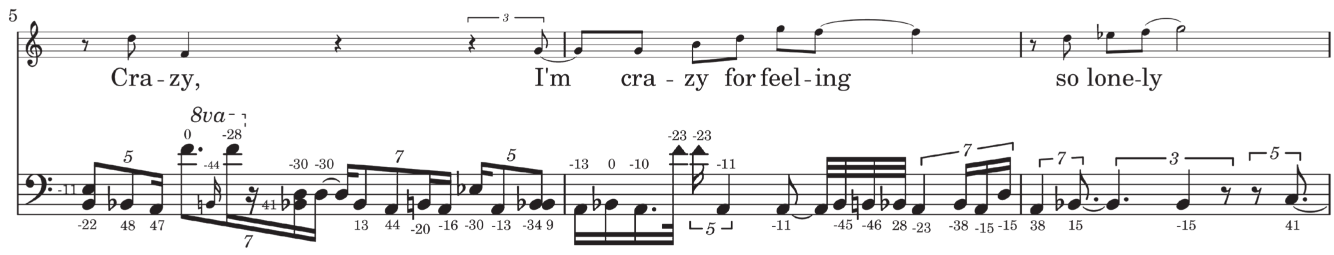

Excerpt: *Crazy*, Upright Bass. This one is a true gem. Sometimes I cannot bring myself to edit a note that is out of playing range. In this case, I write it "as high as possible" (the triangle notehead) and leave it up to bass player Mustafa Walker to deal with. It's now one of my favorite moments in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

11 This is my Grisey moment.
12 Chris Nappi carried the drum up 4 flights up stairs, played 3 takes, and was out the door in under 20 minutes. He's a prince.

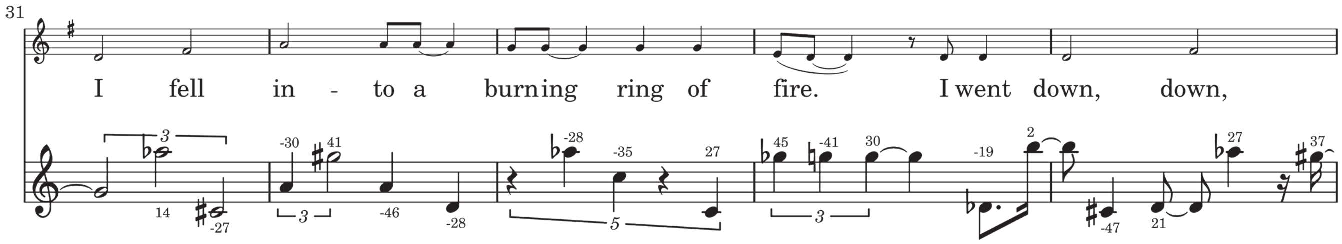

Excerpt: *Ring of Fire*, Electric Guitar. The regularity of the palm-muted single-note electric guitar pulses in the original recording drives the music almost like a percussion instrument, so I tuned my event detection to be more sensitive to rhythmic onsets than to pitched onsets.

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]

13 Seriously, that trumpet entrance might just as well say, "just do something… anything."

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.

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*...

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.

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.

I am not saying that HVB is how *all* machines hear. There is no single answer to that question; machines hear in as many different ways as we design and build them. That's the fun of this project and also its critique. I hope to bring to the surface the kinds of questions that machine hearing systems grapple with, and the decisions that designers of existing systems have made for us.

HVB is the product of my compulsion to overextend and abuse machine hearing models. These algorithms find unexpected things if you look in the right places: analyzing time scales too short or too long to have acoustic-perceptual meaning, or forcing the algorithm to make sense of something—a pitch, a rhythmic onset, a guitar—when that something isn't so clear, or isn't really there at all. HVB facetiously asks the question "What is really there?" as an excuse to negotiate between what these models of human hearing find and what we humans actually hear. If a model describes how things work, then what totally ludicrous crap can I dig out of an audio signal that still fits the mold? (I just hope that, if these computer programs have any sentience, they also have a sense of humor, otherwise I'm guilty of all manner of exploitation.)

Ultimately this project deals with the technology that we increasingly rely on to think, listen, and feel for us. Automation is omnipresent. We offload our brains to nonbiological computation in a sci-fi-esque bio-soft mind-meld. I believe that this technology determines us as much as we determine it. We need to be more responsible with it—i.e., irreverent. This means confronting the assumptions that algorithms rest on, opening up the hood to tease out the space between what these things actually compute and what we tell ourselves they do. I am not satisfied to let capitalist technoculture make these decisions alone, and I am worried what will happen if we do. This notion was only a vague intuition five years ago, but now advertising algorithms and newsbots are harassing people autonomously, and news bubbles have demonstrated the danger of algorithms that reify rather than diversify our beliefs. Algorithmic discrimination is a real and growing problem that we will have to confront. What drives these designs? Who is left out? Who benefits?

It has been more than thirty years since John Oswald released his audio piracy mega-tirade *Plunderphonics*,[14] which demonstrated that the music industry's definition of "property" needed to change. Decades later we only have new questions. What does it mean now that we can reach in and extract pretty much anything from a recording? What does it mean when the process of "sampling" mangles its source and creates something new? What does it mean for an algorithm to be trained to recognize—and eventually replicate—structures and patterns in material *owned* by an individual?[15] Whose property is it then?

If there is one thing that I have learned from all this, it is that our modes of perception are highly mutable. We should use machines to hear differently, not to reinforce our expectations—because whose expectations are they anyway? HVB is about listening through the technology, expressing rather than filtering out its idiosyncrasies and artifacts, and giving voice to its tacit mechanisms.

---

14 ::cheers:: Finally, the *Plunderphonics* reference we've been waiting for!
15 Or–let's be real–a corporation.

EPILOGUE: ABOUT THIS RECORD

When I started writing this music, I never imagined that people would play it. My initial experiments were made with synthesizers, not people. The band and I have spent the past five years making sense of our ludicrous first impulse to play this music. This record contains some of the very first songs—*Crazy* and *It's a Man's Man's Man's World*—as well as a few brand new ones, like *Born to Run* and *Jungle Boogie*.

At this point, HVB is equal parts the music and the band. I choose songs because I want a particular player's take on the music: *(You Make Me Feel Like) A Natural Woman* because there was no way I was going to release this record without Daniel Costello's French horn; *Born to Run* because I wanted to hear a shimmering mess of E Street Band glockenspiel as regurgitated by Chris Nappi; *It's a Man's Man's Man's World* because Pauline Kim Harris shreds the opening string line; and *After the Gold Rush* because Thomas Verchot actually plays all of the flugelhorn microtones. This record features twenty musicians in total, some of whom have been with the band since the beginning and others who are playing with us for the first time. Sometimes these are full solo features, like Joseph Kubera's piano rendition of *This Guy's In Love With You*, and sometimes they just strum a few chords on a fretless electric guitar.[16]

Above all I choose songs because I want to know what will come out the other side of the entire process. The surprises have kept us going. When we stumble upon something that is remarkably musical or gorgeous, we forget all about machine learning and digital analysis—really, we forget about computers altogether. That is what I hope you will hear when you listen to this record.

My deepest gratitude and respect goes to all of the performers who have made this music possible, to the audio research community from whom I have borrowed ideas and source code,[17] and to the original artists (and their lawyers) who hopefully have no idea that this is happening.[18]

—David Kant
Summer, 2016

16 Looking at you, Larry Polansky.
17 https://github.com/davidkant
18 What kind of copy is it anyway after such a bizarre process? © 2016 Only the Errors.

Layout and Design by Mustafa Walker
Handwriting by Rachel Cassandra
Cover Image by Alexander Dupuis

with support from the University of California Santa Cruz, Arts Dean's Fund for Excellence

Indexical