



THE SONG DECODERS

By breaking music down into its component parts, Pandora
Internet radio tries to figure out what kind of
music you — not your social group, heroes or aspirational self — really like.

By Rob Walker

ILLUSTRATIONS BY WORKBENCH

On first listen, some things grab you for their off-kilter novelty. Like the story of a company that has hired a bunch of "musicologists," who sit at computers and listen to songs, one at a time, rating them element by element, separating out what sometimes comes to hundreds of data points for a three-minute tune. The company, an Internet radio service called Pandora, is convinced that by pouring this information through a computer into an algorithm, it can guide you, the listener, to music that you like. The premise is that your favorite songs can be stripped to parts and reverse-engineered.

Some elements that these musicologists (who, really, are musicians with day jobs) codify are technical, like beats per minute, or the presence of parallel octaves or block chords. Someone taking apart Gnaris Barkley's "Crazy" documents the prevalence of harmony, chordal patterning, swung 16ths and the like. But their analysis goes beyond such objectively observable metrics. To what extent, on a scale of 1 to 5, does melody dominate the composition of "Hey Jude"? How "joyful" are the lyrics? How much does the music reflect a gospel influence? And how "busy" is Stan Getz's solo in his recording of "These Foolish Things"? How emotional? How "motion-inducing"? On the continuum of accessible to avant-garde, where does this particular Getz recording fall?

There are more questions for every voice, every instrument, every intrinsic element of the music. And there are always answers, specific numerical ones. It can take 20 minutes to amass the data for a single tune. This has been done for more than 700,000 songs, by 80,000 artists. "The Music Genome Project," as this undertaking is called, is the back end of Pandora.

Pandora was founded in Oakland a decade ago, and for much of the intervening time has lived a precarious existence (the founders spent one three-year stretch working without salaries while they scrambled for investors). But thanks in part to the popularity of the Pandora iPhone app, its fortunes have lately improved. It has attracted 35 million listeners and claims about 65,000 new sign-ups a day (more than half from mobile-device users). About 75 companies are working Pandora into a variety of gizmos and gadgets and Web platforms. The business model relies largely on advertising, and its founder, Tim Westergren, says Pandora will very likely turn its first profit in the fourth quarter of this year.

However things play out for Pandora as a business, its approach is worth understanding if you're interested in the future of listening. It's the "social" theories of music-liking that get most of the attention these days: systems that connect you with friends with similar tastes, or that rely on "collaborative filtering" strategies that cross-match your music-consumption habits with those of like-minded strangers. These popular approaches marginalize traditional gatekeepers; instead of trusting the talent scout, the radio programmer or the music critic, you trust your friends (actual or virtual), or maybe just "the crowd."

Pandora's approach more or less ignores the crowd. It is indifferent to the possibility that any given piece of music in its system might become a hit. The idea is to figure out what *you* like, not what a market might like. More interesting, the idea is that the taste of your cool friends, your peers, the traditional music critics, big-label talent scouts and the latest influential music blog are all equally irrelevant. That's all cultural information, not musical information. And theoretically at least, Pandora's approach distances music-liking from the cultural information that generally attaches to it.

Which raises interesting questions. Do you *really* love listening to the latest Jack White project? Do you *really* hate the sound of Britney Spears? Or are your music-consumption habits, in fact, not merely guided but partly shaped by the cultural information that Pandora largely screens out

Rob Walker writes the Consumed column for the magazine and is the author of "Buying In: The Secret Dialogue Between What We Buy and Who We Are."

— like what's considered awesome (or insufferable) by your peers, or by music tastemakers, or by anybody else? Is it really possible to separate musical taste from such social factors, online or off, and make it purely about the raw stuff of the music itself?

Tim Westergren is a familiar type: the musician who was not as successful as he might have been and concluded that the system is flawed because it underrates talented people who deserve a bigger audience. He played in bands that never quite took off and for a time worked as a film-score composer. It was that job — a "methodical, calculating form of composition," he says — that led him to dwell on the way music works and forced him to decode the individual taste of whatever director had hired him. He says he was getting pretty good at this. "So I thought I'd try to codify it," he says.

Rangy and bright-eyed at 43, Westergren comes off more like the head of a fan club than an erstwhile rock star. The only time he seems annoyed is when he's talking about how some unpopular musicians are unfairly overlooked — or how some popular ones are unfairly maligned. Pandora is, in effect, a response to both of those problems.

Hé founded his company with two tech-and-business-savvy pals in the start-up-friendly year of 1999. Back then it was called Savage Beast Technologies, and the early (not exactly farsighted) business model involved listening kiosks in record stores. Eventually the company got new financing, beefed up the executive team and landed on using its genome as the engine of an Internet radio service "that plays only music you like."


Pandora went online in 2005 and looked much as it does today. When you arrive at the site, you're invited to type in the name of an artist, or a specific song. Let's say you type in "These Foolish Things," by Stan Getz. The Pandora genome looks for something it judges to have a similar infrastructure — like, when I tried recently, "I Don't Know Why," by Don Byas.

This is Pandora's first guess at a song you will like, based on upon its analysis of the song you picked. You can simply let it play; click a "thumbs down" icon to try another song; or give it a thumbs up if you want Pandora's algorithm to know this was a particularly good choice. You can also click to learn why the song was chosen: you don't get a full breakdown but rather a kind of thumbnail summation. In this case the Byas tune was chosen "because it features swing influences, a leisurely tempo, a tenor-sax head, a tenor-sax solo and acoustic-piano accompaniment."

If you click a lot, the idea is that Pandora's algorithm adjusts, squaring your taste with the genome's database. There are other ways to tweak things — adding more songs to a "station" for the system to scrutinize, creating different stations based on other artists or songs, telling the service not to play a given song for a while. (This happens on a station-specific basis: whatever preferences I express on a station based on "My Sharrona" would not affect the songs on, say, my Yanni station.)

Relying on advertising revenue — visual ads on its site as well as occasional audio ads interspersed between songs on your stations — means that much depends on Pandora's genome doing a good-enough job to keep people listening. (There's also a "premium" ad-free service for \$36 a year, and Pandora makes a small commission if you click through its site to buy a song on iTunes or Amazon.com, but it's primarily an ad-driven business.) Its biggest expense is the licensing fee it pays to publishers and performers; the performance fee, paid to an entity called SoundExchange, which distributes royalties to artists, is equal to something like 50 percent of Pandora's revenue. When you start a station with a specific song, that song isn't the first thing you hear, because this would entail an "on demand" license, which costs even more.

By way of Pandora's Twitter feed, I issued a call for users who not only listened to the service a lot but also felt that it had had some kind of impact on their listening tastes. Summer Sterling, a 21-year-old senior at Washington and Lee University in Lexington, Va., often starts by typing in well-known bands like the Dixie Chicks, and that has led her to music by groups



she had never heard of but now loves, like the Weepies. Stephanie Kessler, a 24-year-old M.B.A. student in St. Louis, started by typing in K T Tunstall and has found her way to Waylon Jennings and David Allan Coe.

Aashay Desai, a 25-year-old computer engineer, has become a “very meticulous” user, building some 30 stations and paying for Pandora’s premium service, which offers better sound quality and more features. Aside from his hard rock/metal station, he has a “metalcore” station that’s “a little more aggressive,” as well as a “polyrhythm metal station” that is probably his “most aggressive.” He has also built an R&B station and a trance station; more recently he discovered Django Reinhardt, whom he used as the basis for a gypsy jazz station.

Others, of course, are not impressed by the genome’s results. Someone passed along to me a harsh assessment by Bob Lefsetz, whose popular Lefsetz Letter critiques pretty much every aspect of the contemporary music business. “I tried and rejected it,” he wrote. “Was flummoxed when a Jackson Browne station I created delivered a Journey song. Huh? ... Jackson is music for the mind, Journey is music for the MINDLESS!”

Jonathan McEuen told me he heard about Pandora a couple of years ago and started using it immediately, “with the goal of breaking whatever algorithm they had.” A devoted music fan and a musician himself, McEuen says he did not believe an online service could understand what sort of music he would like and introduce him to new artists based on some deconstruction of his listening tastes. “You can’t just reduce it to a bunch of numbers,” he recalls thinking. “This is a romantic, emotional thing,” and Pandora’s approach to it “can’t work.”

He has changed his mind. A 28-year-old clinical neuroscience researcher at the University of Pennsylvania, he’s a listener who lacks the time to keep up with music news the way he did while amassing hundreds of CDs as a student. Sometimes he runs Pandora as background music; sometimes he’s more engaged, using it as a way to learn about contemporary classical

and opera — and as a result has become a fan of the music of a young composer named Eric Whitacre. “I don’t know how else I would have found out about it,” he says. “Except through the exhaustive process of making new friends on the Internet. Which is something I’m kind of loath to do.”

What I didn’t hear Pandora users talk about was the Genome Project; many didn’t really know about it. They cared about the music Pandora served up, period. But I wanted to know what was behind that music.

Nolan Gasser was the primary shaper of the lexicon that could reconcile Westergren’s genome metaphor with something a computer could evaluate. Gasser, an actual musicologist, wrote a doctoral thesis that dealt with close analyses of Renaissance composition. “I really needed to know what made that music tick,” he recalls. That systematic study flowed well into his work with Westergren — although they started with 20th-century pop, not Renaissance vocal music. First every piece is broken down into large-scale aspects of music: melody, harmony, rhythm, form, sound (meaning instrumentation and, if necessary, voice), and in many cases the text, meaning lyrics. Each of these broader categories might have 10, 30, 50 elements.

“We have a number of characteristics for vocals,” he continues. “Is it a smooth voice, is it a rough, gravelly voice, is it a nasally voice?” Similar questions are evaluated for every instrument. The upshot was about 250 “genes” for every song in the original pop-rock version of the “genome.”

Gasser also helped develop the training mechanisms to make sure the analysts are consistent about more subjective matters — like how “emotionally intense” that Stan Getz solo is. (It’s a 4 out of 5, in the genome’s view.) The test that candidates take involves being able to pick out, quickly and by ear, harmonic structures, melodic organization and other musical elements. The indoctrination that follows revolves around examples. (You think that vocal gets a 5 on the gravelly scale? Here’s Tom Waits. Is it *that* gravelly?)

Recently I sat in as several of Pandora’s song deconstructors gathered in a small conference room to talk about Indian music. Pandora listeners have been asking for Indian music for a while, but adding it to the service hasn’t been a simple matter. A new genre must arrive in a big batch — about 3,000 pieces of music — because Pandora’s algorithm needs lots of choices to be able to recommend something similar-sounding. And all of it has to be pulled apart first. This entails squaring the very different structures of Indian music with Pandora’s “genome” data points.

Over the previous six weeks or so, the Pandora analysts listened to 650 Indian pieces, and the session I observed was a refresher course. Steve Hogan, who oversees Pandora’s analyst squad, had given a half-dozen of its members the same two songs to analyze. The first was “Raga Ahir Bhairav,” recorded by Bismillah Khan in 1955. But the analysts had not been given this cultural information; all they had for the assignment was the music and their ears. Hogan played a snippet and pointed to Kurt Kotheimer, a bass player who often gigs around the Bay Area.

Kotheimer consulted his listening notes: “Flat second, major third, perfect fourth, perfect fifth, major sixth, flat seventh.” Everybody nodded: that’s the tone set, which helps identify the particular raga, one of 25 new “genes” added to Pandora’s algorithm to accommodate this variety of non-Western music. Based on the beat, everyone agreed that this raga was set in Teentaal, with a 16-beat rhythmic cycle often heard in North Indian classical music; it’s now in the genome too. But that was the easy part, apparently.

They moved on to vocals, and Alan Lin, a violinist, ticked off the scores he came up with for things like rhythmic intensity and the relative exoticism of the melody scale. “I actually put exotic at 3.5,” he said. This prompted Sameer Gupta — a percussionist and an expert on Indian music who was weighing in by speakerphone from New York — to lead a brief discussion of how to think about melody and exoticism in this context. Seven or eight scores related to melody, and then about the same number for harmony. (“A 5 for drone,” one analyst announced.) More scores related to form. Tempo. The timbre of the reeds. When Gupta gave his score for riskiness on the

percussion—a 3.5—Lin did a sort of fist pump: “Yes!” Evidently he’d scored it the same way, meaning progress toward properly fitting Indian music into the Music Genome Project. Things went on like this for a while. “Even if you have a solo violin with a tabla, you’re still going to have monophony,” Gupta remarked at one juncture. “I just wanted to point that out.” It was hard to believe there was a business riding on this kind of conversation.

But while some of the genes involve expert, subjective judgment, they aren’t qualitative in the most traditional sense: there’s no rating that allows an analyst to conclude that a vocal or a sax solo is simply lousy. What Pandora’s system largely ignores is, in a word, taste. The way that Gasser or Westergren might put this is that it minimizes the influence of *other people’s* taste. Music-liking becomes a matter decided by the listener, and the intrinsic elements of what is heard. Early on, Westergren actually pushed for the idea that Pandora would not even reveal who the artist was until the listener asked. He thought maybe that structure would give users a kind of permission to evaluate music without even the most minimal cultural baggage. “We’re so insecure about our tastes,” he says.

While his partners talked him out of that approach, Westergren maintains “a personal aversion” to collaborative filtering or anything like it. “It’s still a popularity contest,” he complains, meaning that for any song to get recommended on a socially driven site, it has to be somewhat known already, by your friends or by other consumers. Westergren is similarly unimpressed by hipster blogs or other theoretically grass-roots influencers of musical taste, for their tendency to turn on artists who commit the crime of being too popular; in his view that’s just snobbery, based on social jockeying that has nothing to do with music. In various conversations, he defended Coldplay and Rob Thomas, among others, as victims of cool-taste prejudice. (When I ran Bob Lefsetz’s dismissal of Pandora by him, he laughed it off, and transitioned to arguing that Journey is, actually, a great band.)

He likes to tell a story about a Pandora user who wrote in to complain that he started a station based on the music of Sarah McLachlan, and the service served up a Celine Dion song. “I wrote back and said, ‘Was the music just wrong?’ Because we sometimes have data errors,” he recounts. “He said, ‘Well, no, it was the right sort of thing—but it was Celine Dion.’ I said, ‘Well, was it the set, did it not flow in the set?’ He said, ‘No, it kind of worked—but it’s Celine Dion.’ We had a couple more back-and-forths, and finally his last e-mail to me was: ‘Oh, my God, I like Celine Dion.’”

This anecdote almost always gets a laugh. “Pandora,” he pointed out, “doesn’t understand why that’s funny.”

By the time the Genome Project got under way, the idea of taking music apart and evaluating it by its acoustic elements was not actually new. “Machine listening” was pioneered in various university settings, often by people who had the exact same problem with collaborative filtering’s reliance on social data that Westergren has. Machine listening basically involves teaching computers to assess sound (or really, waveforms representing sound) into something resembling the way that humans hear it, with the goal of eliminating living, breathing listeners from the evaluation process completely.

Like collaborative filtering, machine listening can deal with a lot of data quickly. And when Westergren was trying to raise a second round of financing after the dot-com bust, most everyone involved in the business of music and technology had come to believe that any recommendation system needed to be able to handle millions of songs, instantly. A bunch of musicians sitting around discussing the finer points of drone and monophony wouldn’t cut it. “Everybody thought it was ridiculous,” Westergren agrees. He gave something like 350 pitches to venture capitalists over three years. “Most investors could not get over this idea that we were using humans.” But to Westergren, there were elements of music that machine listening just

couldn’t capture—like the emotionality of a Getz solo. So yes, he wants listeners to experience new music on the basis of the music and not the influence of other people—but to do it right, people have to analyze the music.

Whatever the algorithmic equation, of course, there’s a listener on the other end who is much harder to decode. What you want to hear can depend on your mood, or whether you’re listening at work or in a nightclub. Context affects any cultural product, but music is different from, say, books or movies. Even a casual listener hears many thousands of songs; and to love a song is to take it in—whether attentively or as background music—over and over. Mick Jagger was once asked what makes a tune a classic, and the co-author of “(I Can’t Get No) Satisfaction” replied, “Repetition.” And yet, even the most conservative listener knows the feeling of hearing a hit single once too often. Maybe because music is so ubiquitous, we respond to it almost like food: sometimes we want to try the new restaurant, sometimes the comfort of a familiar favorite dish.

Still, are all these listener-specific factors really enough to explain what music we like, and why? “Music is an inherently social experience,” argues David Goodman, the president of CBS Interactive Music Group, which includes the popular Last.fm Internet radio service. Last.fm’s social-networking model revolves largely around this idea. “The way in which you experience music by sharing, by storytelling, being part of a community. Last.fm is built on what is organic to music.”


Ali Partovi, the C.E.O. of iLike, makes a related point. Used as an application on Facebook and similar sites, iLike bills itself as a “social music-discovery service” and claims more than 50 million registered users. There’s a huge difference, Partovi argues, between “this computer thinks you’ll like this song” and “your friend thinks you’ll like this song—even if it’s the same song.” The problem with a computer reading waveforms is that it “has no common sense,” summarizes Mike McCready, a founder of a company called Music Xray, a digital-music business for entertainment companies and artists. “It doesn’t take into consideration whether the artist is just starting out or they’re at the pinnacle of their career, it doesn’t take into consideration what they wore to the Grammys or who they’re dating or what they look like or what their age is. You have to factor all of this stuff in.”

And why is that? Surely not one consciously says, “My cool friends like the new Jack White, so I’ll memorize the lyrics and pretend to like it, too, for sociocultural reasons.” Yet the research about how listeners link musical taste (at least at a genre level) and identity is extensive. Surely that’s one reason so much of digital music culture is devoted to opportunities to “share” your taste: the endless options for posting playlists, recommending songs, displaying what you are listening to now, announcing your favorite artists.

Maybe the more vivid illustration of social influence on listening habits isn’t in what we share but in what we obfuscate. Last.fm, for example, publishes a chart listing the songs that its users most frequently delete from their public listening-stream data. The guilty pleasure Top 10 is dominated by the most radio-ready pop artists—Katy Perry’s “I Kissed a Girl,” several tracks by Lady Gaga. The service iLike compiles similar data on the most “suppressed” songs its users listen to in secret; Britney Spears figures prominently. Apparently even listeners who can set aside certain cultural information long enough to enjoy something uncool would just as soon their friends didn’t know. Maybe even in our most private listening moments, what our peers think matters.

Much attention has been focused in the last few years on studying music-liking at the brain level. Daniel Levitin, a neuroscientist (and musician) has been one of the high-profile thinkers in this area, by way of his popular books “This Is Your Brain on Music” and “The World in Six Songs.” One of his central themes is that pretty much all humans are wired to enjoy music, and he says he believes musicality is even important to the evolution of the species.

But when you start talking about individuals, instead of humanity in general, universals are a lot harder to come by. Much depends on culture. The emotions expressed in many of those ragas that Pandora’s experts are pres-



ently decoding, for instance, are lost on the typical Westerner. Just as we're hard-wired to learn a language, but not to speak English or French, our specific musical understanding, and thus taste, depends on context. If a piece of music sounds dissonant to you, it probably has to do with what sort of music you were exposed to growing up, because you were probably an "expert listener" in your culture's music by about age 6, Levitin writes.

The cliché that our musical tastes are generally refined in our teens and solidify by our early 20s seems largely to be true. For better or worse, peers frequently have a lot to do with that. Levitin recalled to me having moved at age 14 and falling in with a new set of friends who listened to music he hadn't heard before. "The reason I like Queen — and I love Queen — is that I was introduced to Queen by my social group," he says. He's not saying that the intrinsic qualities of the music are irrelevant, and he says Pandora has done some very clever and impressive things in its approach. But part of what we like is, in fact, based on cultural information. "To some degree we might say that personality characteristics are associated with, or predictive of, the kind of music that people like," he has written. "But to a large degree it is determined by more or less chance factors: where you went to school, who you hung out with, what music they happened to be listening to."

Pandora's approach to listening violates at least three pieces of conventional digital-music wisdom: it rejects the supremacy of social-data taste communities; it shrugs off the assumption that contemporary listeners must have instant on-demand access to any single song; and, most striking, it rejects what many observers see as a given, which is that music consumers are fundamentally motivated by access to the most massive pool of songs possible. Slacker.com, a rival Internet-radio service, says its library contains about 2.5 million songs. Spotify, the European music streaming service, expected to be available in the U.S. by early next year, is generating enormous buzz because of it offers free, on-demand access to more than 5 million tunes.

Pandora's 700,000-song library sounds puny by comparison. And yet the service has millions of devoted listeners. Why? One answer, perhaps, involves the ways that the genome, quietly, *doesn't* really screen out sociocultural information. For instance, its algorithms are tweaked by genre, and the inclusion of genes for "influence" ("swing" or "gospel," for example) brings in factors that aren't strictly about sound. And Pandora's algorithm does adjust if, for instance, users routinely thumbs-down a particular song under similar circumstances, meaning the genome's acoustic judgment can at times be trumped by crowd taste. But the biggest cultural decision of all may be the one

that also happens to guide Westergren's response to the issue of scale: how, exactly, does a given piece of music get into Pandora's system anyway?

Pandora claims to add about 10,000 songs a month to its library. The "curation" of Pandora, in effect, falls to Michael Zapruder, another musician who has found himself working for a tech company. Zapruder ended up as Pandora's curator because he had a habit of identifying holes in the service's collection. Eventually he was told to fill all the gaps he could. "I had a field day," he recalls; he'd stroll through record stores, buying every single Johnny Cash CD or every tango disc available, plus anything that looked interesting. He paid attention to users' suggestions. Somebody wrote in to say that Pandora needed to improve its jazz-trombone selection; somebody else complained about the dearth of barbershop-quartet music. He took care of it. He has beefed up the Latin-music and the J-pop catalog. The major acquisition project right now is Afrobeat, because by far the biggest failed search is Fela Kuti. Zapruder is in the midst of this research but knows that as this new batch of music comes online, "we're going to get educated by our listeners."

Every Tuesday he looks at the New Music Tipsheet, which lists a few hundred new tracks in a typical week. He scrutinizes the Billboard and CMJ charts. He hears directly from a wide array of distributors, from indie-focused Revolver Records to big shots like Universal Music Group. In addition to what is simply sent to Pandora (by labels, artists, P.R. firms), the company buys hundreds of CDs a month, as well as electronica and hip-hop downloads, acquired from sites like Beatport. Every month, hundreds of bands send songs, and Zapruder does his best to get onto Pandora what he figures his listeners want to hear. Still, the labor-intensive genome simply can't absorb it all.

Westergren maintains that catalog size receded as a problem at around the 300,000-song mark. Since passing that, he says, the number of "missed" searches has declined markedly, so the great majority of people who come to the site and type in an artist or song name get a proper introduction to the Pandora system. But the more surprising part of Westergren's response is his claim that he isn't worried about compiling the biggest possible catalog. "This may seem counterintuitive," he told me, "but we struggle more with making sure we're adding really good stuff." That sounds like a rather subjective, cultural judgment — shouldn't the listener decide what's good, based purely on the genome's intrinsic-of-music guidance? Well, there's no question that Westergren is a champion of the unheard music that gets marginalized by sociocultural judgments. But even he has standards. ♦