

[The A.V. Club](#)
[Deadspin](#)
[Gizmodo](#)
[Jalopnik](#)
[Jezebel](#)
[Kotaku](#)
[Lifehacker](#)
[The Root](#)
[The Takeout](#)
[The Onion](#)
[The Inventory](#)

We come from the future
[Shop](#)[Subscribe](#)

[Home](#)[Latest](#)[Reviews](#)[Tech](#)[io9](#)[Earth](#)[er](#)[Science](#)[Field Guide](#)
We come from the future
[Home](#)[Latest](#)[Reviews](#)[Tech](#)[io9](#)[Earth](#)[er](#)[Science](#)[Field Guide](#)

How Shazam Works To Identify (Nearly) Every Song You Throw At It

By
Bryan Jacobs
[9/24/10 10:00PM](#)
[Comments \(230\)](#)

Many of us are prone to using the Shazam music-identification service whenever we encounter unfamiliar songs. After all, it's just so easy to whip out our phones, open an app, and know everything about a mystery song in seconds. But how does Shazam give us all this information so quickly?

Advertisement

There is a cool service called [Shazam](#), which takes a short sample of music, and identifies the song. There are couple ways to use it, but one of the more convenient is to install their free app onto an iPhone. Just hit the "tag now" button, hold the phone's mic up to a speaker, and it will usually identify the song and provide artist information, as well as a link to purchase the album.

What is so remarkable about the service, is that it works on very obscure songs and will do so even with extraneous background noise. I've gotten it to work sitting down in a crowded [coffee shop](#) and [pizzeria](#).

So I was curious how it worked, and luckily there is a [paper written by one of the developers](#) explaining just that. Of course they leave out some of the details, but the basic idea is exactly what you would expect: it relies on fingerprinting music based on the [spectrogram](#).

Related Stories


[Anonymous Claims to Leak Data on the Texas GOP](#)


[What If? Season 2 Promises a Less Apocalyptic Look at the Marvel Multiverse](#)

[You Absolutely Must Watch BBC One's Ghosts Before the American Import](#)

Here are the basic steps:

1. Beforehand, Shazam fingerprints a comprehensive catalog of music, and stores the fingerprints in a database.
2. A user "tags" a song they hear, which fingerprints a 10 second sample of audio.
3. The Shazam app uploads the fingerprint to Shazam's service, which runs a search for a matching fingerprint in their database.
4. If a match is found, the song info is returned to the user, otherwise an error is returned.






G/O Media may get a commission

23% off

[Rotate® DIY Watchmaking Kit](#)

All-in-one kit to build your own mechanical watch.

Founded by watchmakers & artisans dedicated to keeping analog alive.



Here's how the fingerprinting works:

You can think of any piece of music as a time-frequency graph called a spectrogram. On one axis is time, on another is frequency, and on the 3rd is intensity. Each point on the graph represents the intensity of a given frequency at a specific point in time. Assuming time is on the x-axis and frequency is on the y-axis, a horizontal line would represent a continuous [pure tone](#) and a vertical line would represent an instantaneous burst of [white noise](#). Here's one example of how a song might look:

Advertisement

Spectrogram of a song sample with peak intensities marked in red. Wang, Avery Li-Chun. An Industrial-Strength Audio Search Algorithm. Shazam Entertainment, 2003. Fig. 1A,B.

The Shazam algorithm fingerprints a song by generating this 3d graph, and identifying frequencies of "peak intensity." For each of these peak points it keeps track of the frequency and the amount of time from the beginning of the track. Based on the paper's examples, I'm guessing they find about 3 of these points per second. [Update: A commenter below notes that in his own implementation he needed more like 30 points/sec.] So an example of a fingerprint for a 10 seconds sample might be:

Advertisement

Shazam builds their fingerprint catalog out as a [hash table](#), where the key is the frequency. When Shazam receives a fingerprint like the one above, it uses the first key (in this case 823.44), and it searches for all matching songs. Their hash table might look like the following:

Advertisement

[Some extra detail: They do not just mark a single point in the spectrogram, rather they mark a pair of points: the "peak intensity" plus a second "anchor point". So their key is not just a single frequency, it is a [hash](#) of the frequencies of both points. This leads to less [hash collisions](#) which in turn speeds up catalog searching by several orders of magnitude by allowing them to take greater advantage of the table's [constant \(\$O\(1\)\$ \)](#) look-up time. There's many interesting things to say about hashing, but I'm not going to go into them here, so just read around the links in this paragraph if you're interested.]

Advertisement

Top graph: Songs and sample have many frequency matches, but they do not align in time, so there is no match.
Bottom Graph: frequency matches occur at the same time, so the song and sample are a match. Wang, Avery Li-Chun. An Industrial-Strength Audio Search Algorithm. Shazam Entertainment, 2003. Fig. 2B.

If a specific song is hit multiple times (based on examples in the paper I think it needs about 1 frequency hit per second), it then checks to see if these frequencies correspond in time. They actually have a clever way of doing this. They create a 2d plot of frequency hits, on one axis is the time from the beginning of the track those frequencies appear in the song, on the other axis is the time those frequencies appear in the sample. If there is a temporal relation between the sets of points, then the points will align along a diagonal. They use another signal processing method to find this line, and if it exists with some certainty, then they label the song a match.

Advertisement

Top image via [NextWeb](#)

Subscribe to our newsletter!

News from the future, delivered to your present.

Type your email

By subscribing you agree to our [Terms of Use](#) and [Privacy Policy](#).

Bryan Jacobs is a Software Engineer living in San Francisco, CA. He enjoys breaking down complicated topics on [his blog](#) including: the Higgs Boson, the recent financial crisis, the adaptive immune system, and the flow of time. He currently is Director of Engineering at Marin Software, makers of the world-leading paid search management platform.

Advertisement

DISCUSSION


By

petethepanda

I like Soundhound, personally



Featured Videos



[How To Prep Your iPhone or iPad Before Selling](#)

How Shazam Works To Identify (Nearly) Every Song You Throw At It

[Tau Bennett Puppeteer](#)

Friday 9:00AM

[How To Turn Your iPhone Into a Mobile Hotspot](#)

9/14/21 8:00AM

G/O Media may get a commission

23% off

Rotate® DIY Watchmaking Kit

All-in-one kit to build your own mechanical watch.

Founded by watchmakers & artisans dedicated to keeping analog alive.



You may also like

Jalopnik

[Congratulations To Everyone Who Bought The 14 New Lincoln Continentals Sold Last Month](#)

Today 3:02PM

The Root

[Why Are They Like This? Another White Woman, This Time, Katie Couric, Recalls Feeling 'Shaken' After an Interview With Denzel Wash...](#)

Yesterday 10:33AM

The A.V. Club

[Sarah Silverman says Hollywood has a problem with "Jewface"](#)

Today 3:43PM

How Shazam Works To Identify (Nearly) Every Song You Throw At It

How Shazam Works To Identify (Nearly) Every Song You Throw At It

Recommended from G/O Media

[The Root Institute 2021: Unpacking the Attacks on Critical Race Theory](#)

9/21/2021, 9:30 AM



[Here Are the Worst Halloween Costumes of 2021](#)

How Shazam Works To Identify (Nearly) Every Song You Throw At It

9/24/2021, 5:50 PM



11 Netflix Settings Everyone Should Know About

9/30/2021, 10:00 AM

