

# IEEE Spectrum

---

FEATURE HISTORY OF TECHNOLOGY

## Why JPEGs Still Rule the Web

Thirty years ago, the JPEG became the dominant way we share digital photos on the Internet

BY ERNIE SMITH

17 JUN 2025





NICOLE MILLMAN; ORIGINAL IMAGE: IRINA IRISER

**F**OR ROUGHLY THREE DECADES, THE JPEG HAS been the World Wide Web's primary image format. But it wasn't the one the Web started with. In fact, the first mainstream graphical browser, NCSA Mosaic, didn't initially support inline JPEG files—just inline GIFs, along with a couple of other formats forgotten to history. However, the JPEG had many advantages over the format it quickly usurped.

Despite not appearing together right away—the JPEG first appeared in Netscape in 1994, two years after the image standard was officially published—the JPEG and Web browser fit together naturally. JPEG files degraded more gracefully than GIFs, retaining more of the picture's initial form—and that allowed the format to scale to greater levels of success. While it wasn't capable of animation, it progressively expanded from something a modem could pokily render to a

format that was good enough for high-end professional photography.

For the Internet's purposes, the degradation was the important part. But it wasn't the only thing that made the JPEG immensely valuable to the digital world. An essential part was that it was a documented standard built by numerous stakeholders.

## The GIF was a de facto standard. The JPEG was an actual one

How important is it that JPEG was a standard? Let me tell you a story.

During a 2013 *New York Times* interview conducted just before he received an award honoring his creation, GIF creator Steve Wilhite stepped into a debate he unwittingly created. Simply put, nobody knew how to pronounce the acronym for the image format he had fostered, the Graphics Interchange Format. He used the moment to attempt to set the record straight—it was pronounced like the peanut butter brand: “It is a soft ‘G,’ pronounced ‘jif.’ End of story,” he said.

I posted a quote from Wilhite on my popular Tumblr around that time, a period when the social media site was the center

of the GIF universe. And soon afterward, my post got thousands of reblogs—nearly all of them disagreeing with Wilhite. Soon, Wilhite's quote became a meme.

The situation paints how Wilhite, who died in 2022, did not develop his format by committee. He could say it sounded like “JIF” because he built it himself. He was handed the project as a CompuServe employee in 1987; he produced the object, and that was that. The initial document describing how it works? Dead simple. Thirty-eight years later, we're still using the GIF—but it never rose to the same prevalence of JPEG.

The JPEG, which formally emerged about five years later, was very much *not* that situation. Far from it, in fact—it's the difference between a de facto standard and an actual one. And that proved essential to its eventual ubiquity.

We're going to degrade the quality of this image throughout this article. At its full image size, it's 13.7 megabytes. IRINA IRISER

## How the JPEG format came to life

Built with input from dozens of stakeholders, the Joint Photographic Experts Group ultimately aimed to create a format that fit everyone's needs. (Reflecting its committee-led roots, there would be no confusion about the format's name—an acronym of the organization that designed it.) And when the format was finally unleashed on the world, it was the subject of a book that was more than 600 pages.

*JPEG: Still Image Data Compression Standard*, written by IBM employees and JPEG organization stakeholders William B. Pennebaker and Joan L. Mitchell, describes a landscape of multimedia imagery, held back without a way to balance the need for photorealistic images and immediacy. Standardization, they believed, could fix this.

“The problem was not so much the lack of algorithms for image compression (as there is a long history of technical work in this area),” the authors wrote, “but, rather, the lack of a standard algorithm—one which would allow an interchange of images between diverse applications.”

And they were absolutely right. For more than 30 years, JPEG

has made high-quality, high-resolution photography accessible in operating systems far and wide. Although we no longer need to compress JPEGs to within an inch of their life, having that capability helped enable the modern Internet.

As the book notes, Mitchell and Pennebaker were given IBM's support to follow through this research and work with the JPEG committee, and that support led them to develop many of the JPEG format's foundational patents. Described in patents filed by Mitchell and Pennebaker in 1988, IBM and other members of the JPEG standards committee, such as AT&T and Canon, were developing ways to use compression to make high-quality images easier to deliver in confined settings.

Each member brought their own needs to the process. Canon, obviously, was more focused on printers and photography, while AT&T's interests were tied to data transmission. Together, the companies left behind a standard that has stood the test of time.

All this means, funnily enough, that the first place that a program capable of using JPEG compression appeared was not MacOS or Windows, but OS/2—a fascinating-but-failed graphical operating system created by Pennebaker and

Mitchell's employer, IBM. As early as 1990, OS/2 supported the format through the OS/2 Image Support application.



At 50 percent of its initial quality, the image is down to about 2.6 MB. By dropping half of the image's quality, we brought it down to one-fifth of the original file size. ORIGINAL IMAGE: IRINA IRISER

## What a JPEG does when you heavily compress it

The thing that differentiates a JPEG file from a PNG or a GIF is how the data degrades as you compress it. The goal for a JPEG image is to still look like a photo when all is said and

done, even if some compression is necessary to make it all work at a reasonable size. That way, you can display something that looks close to the original image in fewer bytes.

Or, as Pennebaker and Mitchell put it, “the most effective compression is achieved by approximating the original image (rather than reproducing it exactly).”

Central to this is a compression process called discrete cosine transform (DCT), a lossy form of compression encoding heavily used in all sorts of compressed formats, most notably in digital audio and signal processing. Essentially, it delivers a lower-quality product by removing details, while still keeping the heart of the original product through approximation. The stronger the cosine transformation, the more compressed the final result.

The algorithm, developed by researchers in the 1970s, essentially takes a grid of data and treats it as if you're controlling its frequency with a knob. The data rate is controlled like water from a faucet: The more data you want, the higher the setting. DCT allows a trickle of data to still come out in highly compressed situations, even if it means a slightly compromised result. In other words, you may not keep

all the data when you compress it, but DCT allows you to keep the heart of it.

(See [this video](#) for a more technical but still somewhat easy-to-follow description of DCT.)

DCT is everywhere. If you [have ever seen a streaming video](#) or an online radio stream that degraded in quality because your bandwidth suddenly declined, you've witnessed DCT being utilized in real time.

A JPEG file doesn't have to leverage the DCT with just one method, as *[JPEG: Still Image Data Compression Standard](#)* explains:

*The JPEG standard describes a family of large image compression techniques, rather than a single compression technique. It provides a “tool kit” of compression techniques from which applications can select elements that satisfy their particular requirements.*

The toolkit has four modes:

- **Sequential DCT**, which displays the compressed image in order, like a window shade slowly being rolled down
- **Progressive DCT**, which displays the full image in the

lowest-resolution format, then adds detail as more information rolls in

- **Sequential lossless**, which uses the window-shade format but doesn't compress the image
- **Hierarchical mode**, which combines the prior three modes—so maybe it starts with a progressive mode, then loads sequential lossless compression slowly, but then reaches a lossless final result

At the time the JPEG was being created, modems were extremely common. That meant images loaded slowly, making Progressive DCT the most fitting format for the early Internet. Over time, the progressive DCT mode has become less common, as many computers can simply load the sequential DCT in one fell swoop.

That same forest, saved at 5 percent quality, now down to about 419 kilobytes. ORIGINAL IMAGE: IRINA IRISER

When an image is compressed with DCT, the change tends to

be less noticeable in busier, more textured areas of the picture, like hair or foliage. Those areas are harder to compress, which means they keep their integrity longer. It tends to be more noticeable, however, with solid colors or in areas where the image sharply changes from one color to another—like text on a page. Ever screenshot a social media post, only for it to look noisy? Congratulations, you just made a JPEG file.

Other formats, like PNG, do better with text, because their compression format is intended to be non-lossy. (Side note: PNG's compression format, DEFLATE, was designed by Phil Katz, who also created the ZIP format. The PNG format uses it in part because it was a license-free compression format. So it turns out the brilliant coder with the sad life story improved the Internet in multiple ways before his untimely passing.)

In many ways, the JPEG is one tool in our image-making toolkit. Despite its age and maturity, it remains one of our best options for sharing photos on the Internet. But it is not a tool for every setting—despite the fact that, like a wrench sometimes used as a hammer, we often leverage it that way.

## Forgent Networks claimed to own the JPEG's defining algorithm

The JPEG format gained popularity in the '90s for reasons

beyond the quality of the format. Patents also played a role: Starting in 1994, the tech company Unisys attempted to bill individual users who relied on GIF files, which used a patent the company owned. This made the free-to-use JPEG more popular. (This situation also led to the creation of the patent-free PNG format.)

While the JPEG was standards-based, it could still have faced the same fate as the GIF, thanks to the quirks of the patent system. A few years before the file format came to life, a pair of Compression Labs employees filed a patent application that dealt with the compression of motion graphics. By the time anyone noticed its similarity to JPEG compression, the format was ubiquitous.

Our forest, saved at 1 percent quality.  
This image is only about 239 KB in size, yet it's still easily recognizable as the same photo. That's the power of the JPEG. ORIGINAL IMAGE: IRINA IRISER

Then in 1997, the company owned the patent that made the format

In 1997, a company named Forgent Networks acquired Compression Labs. The company eventually spotted the patent and began filing lawsuits over it, a series of events it saw as a stroke of good luck.

“The patent, in some respects, is a lottery ticket,” Forgent CEO Jay Peterson told CNET in 2005. “If you told me five years ago that ‘You have the patent for JPEG,’ I wouldn’t have believed it.”

While Forgent’s claim of ownership of the JPEG compression algorithm was tenuous, it ultimately saw more success with its legal battles than Unisys did. The company earned more than US \$100 million from digital-camera makers before the patent finally ran out of steam around 2007. The company also attempted to extract licensing fees from the PC industry. Eventually, Forgent agreed to a modest \$8 million settlement.

As the company took an increasingly aggressive approach to its acquired patent, it began to lose battles both in the court of public opinion and in actual courtrooms. Critics pounced on examples of prior art, while courts limited the patent’s use to motion-based uses like video.

By 2006, Forgent’s compression patent expired—and its litigation-heavy approach to business went away. That year,

the company became Asure Software, which now specializes in payroll and HR solutions. Talk about a reboot.

## Why the JPEG won't die

The JPEG file format has served us well. It's been difficult to remove the format from its perch. The JPEG 2000 format, for example, was intended to supplant it by offering more lossless options and better performance. The format is widely used by the Library of Congress and specialized sites like the Internet Archive; however, it is less popular as an end-user format.

See the forest JPEG  
degrade from its  
full resolution to  
1 percent quality  
in this GIF.  
ORIGINAL IMAGE:  
IRINA IRISER

Other image technologies have had somewhat more luck getting past the JPEG format. The Google-supported WebP is popular with website developers (and controversial with end users). Meanwhile, the formats AVIF and HEIC, each

developed by standards bodies, have largely outpaced both JPEG and JPEG 2000.

Still, the JPEG will be difficult to kill at this juncture. These days, the format is similar to MP3 or ZIP files—two legacy formats too popular and widely used to kill. Other formats that compress the files better and do the same things more efficiently are out there, but it's difficult to topple a format with a 30-year head start.

Shaking off the JPEG is easier said than done. I think most people will be fine to keep it around.

*Ernie Smith is the editor of Tedium, a long-running newsletter that hunts for the end of the long tail.*

*This article was updated on 01 July 2025.*

*This article appears in the August 2025 print issue.*

