

ImageVice™ 2.0 Documentation

Index

1: Overview	2
2: Why you need to use ImageVice™	3
3: When to use ImageVice™	5
4: When not to use ImageVice™	7
5: Getting started with ImageVice™	8
6: Understanding the ImageVice™ Controls	10
7: Preparing graphics with ImageVice™	17
8: Using ImageVice™ for GIF animations	30
9: Contact information	37

ImageVice™ 2.0 Documentation

Overview

ImageVice™ is an editor's choice award-winning color reduction filter for Photoshop 3.0 or better presenting designers with a new technology that enables common multimedia format files, such as GIF, PICT, BMP, and PNG, to be made up to 70% smaller than previously possible.

ImageVice™ 2.0 Documentation

Why do you need to use ImageVice™

There are two main reasons to use ImageVice™ – smaller files and better looking images. The third reason, of course, would be to get both of those at once.

Smaller files:

ImageVice™ is capable of creating images that are more compressible than they would be using any other color reduction tool. Smaller files are critical for image applications such as the World Wide Web, where larger graphics files mean longer load times for viewers and create limits for the amount of image content that can be delivered on any single Web page at once. Multimedia applications are also subject to limited storage space, and smaller files allow more content.

ImageVice™ works as an image filter. It is not a true compression tool. It does not create files or actually compress image data in any way; what it does is change the image data so that it will compress much more efficiently. Because ImageVice™ is removed from any particular file format, it can easily be used to make files in any file format that uses lossless compression, such as GIF, PNG, BMP, and PICT formats, smaller. This allows great flexibility and broad uses for ImageVice™ to create more compressible files in many formats for many different applications, ranging from Web pages to game development.

Better image quality:

The second main reason to use ImageVice™ is image quality. The color reduction in ImageVice™ is not only capable of producing much more compressible images, it is also capable of producing the best visual quality images. Sparing the complex jargon required to explain how ImageVice™ actually accomplishes what it does, it is simplest to say ImageVice™ closely models the human visual system and the way the eye sees using proprietary technology.

ImageVice™ 2.0 Documentation

ImageVice™ provides complete flexibility – it has already been called the most flexible color reduction tool ever – and can be used to make files that are smaller, images that look better, or both at one time to meet your exact needs.

For Web graphics work in particular, where issues of image quality and file size must be carefully juggled, ImageVice™ is invaluable. It can be used to make what was impossible to do before, because of file size issues, possible.

Our own Web site design would not have been possible without the help of ImageVice™ to make the intense graphics and animations an acceptable size.

ImageVice™ 2.0 Documentation

When to use ImageVice™

It isn't possible to discuss all the possible application of ImageVice™, so this subject is limited to using ImageVice™ for preparing graphics for the Web that will be saved in GIF or PNG formats.

There are several situations in Web graphics production where the best tool you could possibly have is ImageVice™. These are:

- When you have a complex image, such as one that is photographic in nature, that must be saved in GIF or PNG format for any reason. In this case, ImageVice™ can provide both superior image quality as well as smaller file size.
- When you have images that combine complex areas with large flat areas of color that must not dither. In this case ImageVice™ will ensure that flat areas of color are not dithered, which will improve both image quality and compressibility.
- When you have very simple images, where the only areas of color that are not flat are antialiasing along edges, such as type graphics. In this case, ImageVice™ can be used to color reduce to extremely low bit-depths without any visual loss in image quality and improve file size by allowing lower bit depths to be used.
- When you need to lower the bit depth or number of colors used in images that have already been saved in indexed format. In this case ImageVice™ can be used to lower the bit-depth without noticeable loss in image quality, which will decrease file size significantly.
- When you must have the absolute best image quality possible. In this case, ImageVice™ can be used without compression settings for the best image quality possible. Though file sizes at settings for maximum quality may be slightly larger than those produced from Photoshop's color reduction in comparison, the images will look much better.

ImageVice™ 2.0 Documentation

- **When your image would benefit from different areas of the image, such as foreground and background objects, being color reduced with different settings. ImageVice™ is a true filter that can be applied to any selection made. It can be used with more aggressive settings for compression on background areas and settings for higher quality on foreground areas that are more important to the visual appearance of the image, which is a flexibility and benefit that no other color reduction tool can offer.**

ImageVice™ 2.0 Documentation

When not to use ImageVice™

ImageVice™ should only be used for images that will be saved in a file format that employs any form of lossless, meaning the image data is compressed and decompressed to be exactly what it started as with no "loss" of image data taking place in that process, compression such as GIF, PNG, BMP, and PICT. ImageVice™ should not be used for the following:

- ImageVice™ should not be used for images that will be saved in lossy compression formats such as JPEG or any of the proprietary formats using wavelet or fractal compression. Using ImageVice™ when saving in a lossy file format will only degrade the final image quality.
- You should not use ImageVice™ on any image that could normally be saved as a JPEG for Web graphics, either. This is a common question that we have been asked.
- ImageVice™ color reduction is not meant to be a replacement for JPEG. As good as ImageVice™ is, the compression ratios that can be achieved still do not exceed those of lossy formats. When you would normally save an image as JPEG, such as a photographic image with no need of transparency, you should still be saving those images as JPEG, and if you are comparing the results of ImageVice™, you should compare the results to other color reduction tools. Making a comparison to JPEG is a fallacy. The two are not comparable. They are different things with different uses.

ImageVice™ 2.0 Documentation

Getting started with ImageVice™

Macintosh installation

When you first download ImageVice™ for Macintosh, it is in a Binhex encoded Stuffit archive. The first step is to decode and decompress the archive. This may or may not be automatically handled by your software, depending on configuration. If it is not automatically done, you should first use Stuffit Expander to decode and decompress the file you downloaded.

Once you have decoded and decompressed the archive file containing ImageVice™, all that needs to be done to install ImageVice™ is to drag the plug-in from its folder to your plug-ins folder. It is normal to also place filter plug-ins, like ImageVice™, in the "Filters" folder inside Photoshop's "Plug-ins" folder, but this is not necessary for Photoshop to load plug-ins. It is only for organizational purposes.

After you have placed ImageVice™ in your plug-ins folder, ImageVice™ will be available the next time you launch Photoshop. If you aren't using Photoshop, you may need to refer to your image-editing application's documentation for instructions on installing plug-in filters.

The speed menu provides a way to see approximate download times for your file at different transmission speeds. The speed menu does not have any real effect on the compression. It is simply a previewing mechanism. However, all figures for download time are only estimated. Under real world conditions on your Web site, they will vary based on many factors such as server load, line noise, and packet latency that cannot be accounted for during estimation.

Windows installation

When you first download ImageVice™ for Windows it is in a zip file. The first step is to decompress this file. This may or may not be automatically handled by your software depending on configuration. If it is not automatically done, you should use WinZip or compatible utility to decompress the file you downloaded first.

After decompressing the zip file, copy the plug-in from its directory to your plug-ins directory.

ImageVice™ 2.0 Documentation

Within Photoshop it should be located in Photoshop\Plug-ins\Filters, but in other image-editors, the location may vary slightly. If needed, consult your image editor's manual on installing filter plug-ins.

Locating the plug-in in Photoshop

Once you have properly installed ImageVice™, it will appear under the Filters->BoxTop Web Tools->ImageVice... menu in Photoshop.

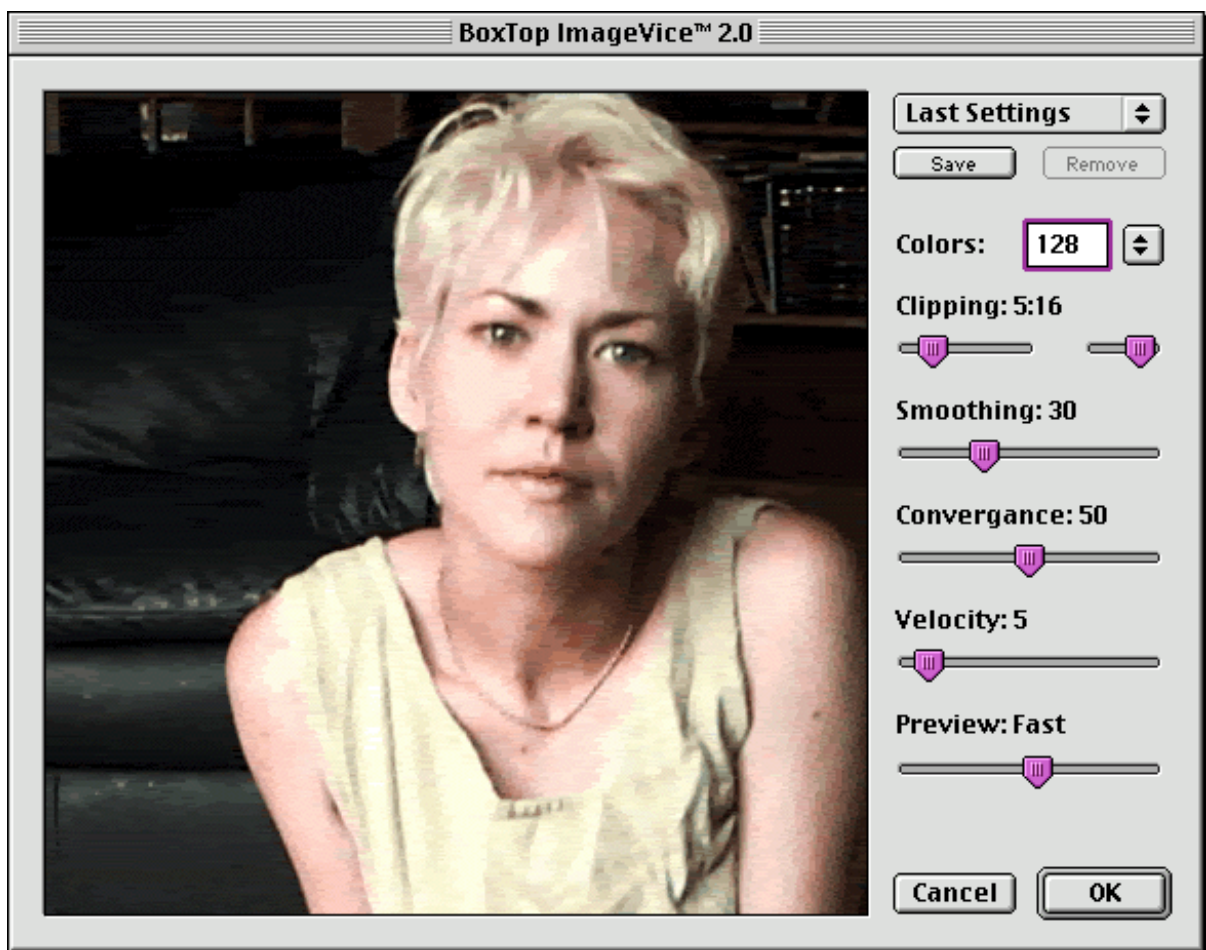
Other applications may locate filter plug-ins under an "Effects" menu, but in all applications, ImageVice™ will be under the "BoxTop Web Tools" sub-menu.

ImageVice™ 2.0 Documentation

Understanding the ImageVice™ controls

ImageVice™ has a single, relatively simple, dialog for specifying color reduction parameters. Though, there are only six individual parameters, which can be specified for color reduction, not all are immediately obvious in their function.

Below is a screen shot of the Macintosh version. All controls are common between Macintosh and Windows versions.



ImageVice™ 2.0 Documentation

The only controls in ImageVice™ that are comparable to controls found in other color reduction tools are the depth menu and colors text box. Both of these provide a means to set the number of colors to be used in the final image.

All other controls are unique to the ImageVice™ plug-in, and require a basic level of understanding before using the ImageVice™ plug-in.

More detailed explanation for controls are provided below.

Depth menu

The depth menu shows the lowest possible bit-depth for the image given the current number of colors specified, and can also be used to set the number of colors specified to the maximum number of colors possible for a given bit-depth.

Lowering the bit-depth will improve an image's compressibility and reduce the final file size. Very few images need to have more than 128 colors, which is the maximum for a 7 bit image, to adequately represent them and many images require even fewer colors.

Colors text box

The colors text box provides another means to set the number of colors to be used. Odd numbers of colors can be specified using the colors text box, whereas only even bit depths can be specified using the depth menu.

Even if lowering the set number of colors will not allow a lower bit depth, it will still improve compressibility. An image using 120 colors will normally be smaller than an image using 128 colors, though both images are 7 bit.

You should set the colors to the minimum number that will produce satisfactory visual results for the individual image. If you are unsure of what this number is, starting at 128 colors or 7 bit, is a good idea. Then you can undo the filter and try a higher or lower number, if it looks like fewer colors will do or more colors are needed.

ImageVice™ 2.0 Documentation

Clipping sliders

The lower clipping slider controls the threshold level for sampling dark colors from the image. The range of the lower clipping slider corresponds directly to gray component values 0 – 50. If the lower clipping slider was set to 50, no colors would be sampled from the image that had a gray value lower or darker than 50 during color reduction.

The upper clipping slider controls the threshold level for sampling light colors from the image. The range of the upper clipping slider is scaled to correspond to gray component values 236 – 256 inversely, which is best thought of as the number of high or light gray levels not to sample. Setting the upper clipping slider to 20, has the effect of no light colors with a gray values of 236 or above being sampled during color reduction.

ImageVice™ color reduction of a gradient with no clipping settings applied.



ImageVice™ color reduction of a gradient with maximum dark clipping of 50 applied with no light clipping set. All colors with a luminosity value of 50 or lower are automatically mapped to black.



ImageVice™ color reduction of a gradient with maximum light clipping of 20 applied with no dark clipping set. All colors with a luminosity value of 235 or greater are automatically mapped to white.



ImageVice™ 2.0 Documentation

ImageVice™ color reduction of a gradient with maximum dark and light clipping settings applied.



Mid-range values for clipping will improve image compressibility by mapping all colors that fall below the thresholding values to a single color in the ranges that subtleties in color are most difficult to distinguish and reduce entropy. Image quality will also be improved with mid-range clipping settings because more palette entries will be available for mid-range colors that are more easily distinguishable.

You should reduce the lower or dark clipping setting for images that are predominantly very dark colors to prevent harsh borders where thresholding occurs, and leave more palette entries available, when they are actually needed in the image.

Correspondingly, you should also reduce the upper clipping setting in images that are predominantly very light colors for the same factors.

In images that have few dark or light areas, you can safely increase the settings for clipping to very high values to improve compression.

Smoothing slider

The smoothing slider controls the tolerance for when very similar adjacent pixel color values are merged.

This is a complex space variant process with only one controllable parameter, so the action taken will vary significantly depending on the actual image, unlike the upper and lower clipping sliders which have predictable results on all images.

Increasing the value for smoothing will increase the compressibility of the image and normally has little apparent visual impact on the final image. However, if the original image has predominant areas that are very smooth linear color blends, high smoothing settings will become more obvious and possibly undesirable.

ImageVice™ 2.0 Documentation

Convergence and Velocity sliders

The two sliders for convergence and velocity work together to give a ratio for an internal mechanism that there is no sane way to explain.

The convergence slider is an overall tolerance value for this internal mechanism, and the velocity slider provides relational amplitude for the mechanism. Velocity is an easier to understand term, though, as the visual effect is apparent in terms of distance. Though, the process is not easily explained in words, its visual effects are apparent, and the controls are easily understood by example.



The two images above show ImageVice™ applied with convergence and velocity settings off in the top image compared to the default settings used in the bottom image. All other settings remain at default values.

ImageVice™ 2.0 Documentation

At normal size, there is no easily apparent difference between the images, but at a magnified view, the visual effects of both convergence and velocity becomes apparent. They will impart horizontal tendency to images, which will greatly improve compressibility. The top image is 13,107 bytes but the bottom image, using default convergence and velocity settings, is just 9,642 bytes.

To better understand the visual effects of both convergence and velocity settings, you should compare images filtered with both convergence and velocity sliders set to zero to an image filtered with the default values and then one filtered with both convergence and velocity at maximum settings.

Convergence and velocity settings are primarily responsible for incredible compression levels that can be achieved, and when at the highest settings imparting a "painterly" look to images.

The default settings for these work very well with almost all images, and produce very good quality to compressibility ratios.

Increasing the convergence slider value will have less visually apparent effect than the velocity slider, which should rarely be used at its maximum setting.

When maximum quality is desired the convergence and velocity sliders can be set to values lower than the defaults, but it is not recommended for compressibility's sake to turn convergence and velocity off completely. Very low settings will still significantly improve compressibility without any noticeable visual effect on most images.

The preview slider

The preview slider controls the sampling rate for and speed of the preview in ImageVice™. It has no effect on final results, but setting this for a faster preview can get you to the final results quicker. It can be said ImageVice™ is great, but it can't really be said it's fast. Speed is always a relative thing, though.

Compared to other color reduction methods capable of producing similar quality levels, few of which have ever found use in commercial software, ImageVice™ is a speed demon.

ImageVice™ 2.0 Documentation

Compared to the color reduction methods that are commonly used in other commercial software, ImageVice™ is somewhat slower.

It can take from several seconds to several minutes to run on an image depending on the size, and regenerating previews can get agonizingly monotonous unless some brilliant steps are taken.

This is what the preview slider is. You can have a faster preview with less precision, or a slower preview with more precision. The process of generating the preview is interruptible, so there is never any unnecessary wait.

You can confirm the settings as soon as you want to apply them to the image, without having to wait for the preview to regenerate.

Image preview

The image preview provides full-size feedback on given settings, but with variable precision, so final outcome will be greater quality than in the preview unless the preview slider is set to true.

OK and Cancel

Both do the usual thing.

ImageVice™ 2.0 Documentation

Preparing graphics with ImageVice™

- A: Differences in color reduction methods
- B: Color reduction goals for optimization
- C: Picking a target bit-depth for color reduction
- D: Analyzing the image for compressibility
- E: Removing little used colors for better compression
- F: Making the image even smaller with your brain
- G: Post production issues

The frog image used in this example is typical of many transparent GIFs used on the Web. The image was created from a digital photograph of a tree frog, who haplessly wandered onto a window. Our goal is to make him hop as fast as possible downloading.



The final frog from this example is only 4278 bytes in size and has significantly better image quality than a comparable size file from Photoshop.


ImageVice™ 2.0 Documentation

A: Differences in color reduction methods

Before going through the process of how to best optimize the image, it is beneficial to look at some key differences between color reduction methods and capabilities. ImageVice™, Photoshop, and DeBabelizer are compared here. ImageVice™ is very different from both in that it does not use traditional dithering; instead, it uses proprietary shape-form shading color mapping techniques. The advantages to these will soon become apparent in the illustrations that follow. Even at the settings for maximum image quality in ImageVice™, where file size is not the priority, it can produce smaller files than Photoshop and DeBabelizer just because of the different color mapping techniques employed.




ImageVice produces a file that is 12,300 bytes at 8-bits using best quality settings. Note the difference in appearance in the magnified sections.



ImageVice™ 2.0 Documentation



Photoshop produces a file that is 12,589 bytes at 8-bits using diffusion dithering for best quality.



DeBabelizer doesn't get its own illustration. It came in last place with a file that was 13,901 bytes big.

If you look at the magnified sections the difference in color mapping can be seen. The image color reduced with Photoshop uses dithering, which causes the grainy appearance and is responsible for making the file compress less efficiently.

The differences in size are not dramatic, though. 289 bytes is only a small difference in size. The problem with other color reduction tools is they are not very flexible. Photoshop is only capable of producing two file sizes for any target number of colors setting. It only allows you to turn dithering on or off. DeBabelizer is equally limited.

12,589 bytes or 11,774 bytes are the only possibilities in Photoshop, in this example.

ImageVice™ 2.0 Documentation

Unlike Photoshop and DeBabelizer, ImageVice™ has ultimate flexibility in comparison with nearly unlimited possibilities. It can produce any number or file sizes ranging from 12,300 bytes to 8088 bytes, in this example, by varying clipping, smoothing, and convergence settings.

4501 bytes smaller is dramatic.

B: Color reduction goals for optimization

The goal, however, is not just to be dramatic, but to achieve the very best possible relationship between image quality and file size, taking into consideration the usage of the image – the perfect optimization.

For this example, we will assume that our frog is a plain and ordinary transparent GIF needing average image quality – enough to convey to the viewer that it's a frog and enough not to embarrass creative types, but not as much as would be wanted for portfolio artwork or a scientific imaging example. The frog is just "typical".

C: Picking a target bit–depth for color reduction

The first step in preparing an image is deciding on a target bit–depth. Does the image need 256 colors or 128 colors or 64 colors, etc. The number of colors and bit–depth that you choose before you even begin preparing the image will have the greatest effect on the file size of your GIF.

The number of colors needed will always vary by image. You should examine your image, and make the best guess you can at the least number of colors and lowest bit–depth that will be able to achieve acceptable image quality first. Never just hit the OK button without thinking about this. Your files will be bigger than they need to be almost every time.

Rarely does an image need more than 128 colors with ImageVice™, which is 7–bits and the default depth in ImageVice™.

The frog image has only one significant range of colors, greens, with small areas of yellows and gray values. There are not many separate hues in the image. Images that have a single predominant color range can almost always be represented well with just 32 colors or 5–bits, which will be our target number of colors for this example.

ImageVice™ 2.0 Documentation

If in doubt about how many colors it will take to adequately represent an image, you should start at 4 or 5 bits and work upwards until a suitable depth is reached. If the results are not acceptable at a particular depth, just undo ImageVice™ and then reapply it at the next higher depth.

You should not spend time tweaking the controls in ImageVice™ for compression, until you have settled on a bit depth that is suited to the image first.



This image shows the results of ImageVice™ at 32 colors and default settings for all other controls. The file is 5015 bytes, and if that is small enough for you, then you could stop here.

It is 1810 bytes smaller than the comparable settings of 32 colors and diffusion dithering in Photoshop, but it is not as small as it can be made using ImageVice™.



ImageVice™ 2.0 Documentation

Differences in appearance are noticeable between ImageVice™, above, and Photoshop, below, at lower bit depths, but the matter of image quality is subjective.

Both are more than acceptable for the use of the frog image, but more quality may be desired for some images, in which case the number of colors used should be increased or the compression settings lowered.

ImageVice™ provides plenty of room for individual interpretation, unlike other tools.

D: Analyzing the image for compressibility

Since ImageVice™ is a true image filter and can be applied to different areas of the image, quality and compression can be greatly improved by separating areas of the image with different properties and using settings in ImageVice™ that are best suited to each individual area.

This is actually a very simple process that can be quickly done on all images using Photoshop's layers and QuickMask capabilities with the following step-by-step procedure.

Step 1:

First, you should begin by flattening your working document, if it is not already on a single background layer. The frog image is already flattened, so this step is not needed with it.

Step 2:

Next, duplicate the background layer twice, so that your document contains three identical layers containing the entire image.

Step 3:

On the top layer, apply ImageVice™ at settings of 32 colors and all other controls set to their maximum values. These are the settings that will result in maximum compression. They are listed to the right of the illustration showing the visual results.

ImageVice™ 2.0 Documentation



Colors: 32
Clipping: 16-15
Smoothing: 40
Convergence: 50
Velocity: 5

Step 5:

On the base layer, apply ImageVice™ at settings of 32 colors and all other controls except clipping set to zero. Clipping should be set to the mid-range defaults for best quality on all images that are not primarily very dark or very light colors. These are the settings that will produce best image quality, and are listed to the right of the illustration showing the visual results.



Colors: 32
Clipping: 16-15
Smoothing: 0
Convergence: 0
Velocity: 0

Step 6:

Starting with the top layer to which you applied the most aggressive settings in ImageVice™ for compression, use Photoshop's QuickMask feature to mask off the sections of the image that are acceptable image quality. The mask you create need not be exacting. The image below shows the mask created very quickly using a large brush size and the pencil tool in Photoshop.

ImageVice™ 2.0 Documentation



After you have created the mask, load the selection from it, and delete the portions of the top layer that are not acceptable quality. This will cause those areas to be filled-in by the layer below, where higher settings for quality were applied to create a composite image.



In this example, there was no need to continue this process with the middle layer, because all areas of it were acceptable image quality for the purposes of the frog image. However, if needed the same process should be repeated with the middle layer, so the final result will be a composite image formed by three layers in Photoshop.

Step 7:

Flatten the image to merge the three layers, then apply ImageVice™ at the same settings for best image quality, as first used on the base layer to consolidate subtle variations in colors.

ImageVice™ 2.0 Documentation



Colors: 32
Clipping: 16-15
Smoothing: 0
Convergence: 0
Velocity: 0

The file produced by this process is only 4888 bytes in size and somewhat smaller than could be achieved by a single application of ImageVice™ to the entire image for the same image quality.

E: Removing little-used colors for better compression

After primary color reduction, you can also use ImageVice™ to automatically remove little used colors to create an even smaller file, with no visible loss in quality from your current image.

Simply apply ImageVice™ to your entire image one more time, with best settings for image quality and a number 2 less colors set. In this case, ImageVice™ was reapplied 30 colors and the same settings used last on all other controls, which are listed to the right of the illustration.

This step reduced the size of the previous file by another 547 bytes, for a file size of 4341 bytes. To make the file even smaller, manual tweaking is needed, which you can easily do if you choose, by following a few simple rules of thumb laid out in the next section.

However, compared to what Photoshop and other color reduction tools are capable of doing, ImageVice™ has already far exceeded their capabilities in both quality and compression. A comparable sized file of 4356 bytes produced by Photoshop is shown below, under the file produced by ImageVice™.

ImageVice™ 2.0 Documentation



ImageVice - 4341 bytes
Colors: 30
Clipping: 16-15
Smoothing: 0
Convergence: 0
Velocity: 0



Photoshop - 4356 bytes
Colors: 13
Dithering: Off

No other color reduction tool is capable of matching the quality to compression ratios that ImageVice™ is capable of producing. At the same file size, there is no question as to which of the two images above looks better.

F: Making the image even smaller with your brain

The nature of ImageVice™ color reduction makes it extremely easy to perform some simple manual tweaks to the final image in Photoshop to obtain an even smaller final file size. For this step you don't need to use ImageVice™, just your brain.

The simple rules of thumb for final manual tweaking are:

- Remove infrequent colors altogether if there are even remotely similar colors you can use instead in the palette. Infrequent colors greatly increase entropy.

ImageVice™ 2.0 Documentation

- Merge colors that are visibly indistinguishable from one another. Change the lesser common of the two colors to the more common one.
 - Get rid of stray pixels of a different color that break up otherwise flat areas of colors, whenever possible.
 - Any place you can set a single pixel to the color of it's horizontal neighbor it will make the image compress better.
- If you followed all the steps in this example, you've already taken care of the first two rules of thumb using ImageVice™, but you can still continue with the last two rules of thumb to make a significant difference in file size.



Before tweaking the file is 4341 bytes.



Only 57 pixels were changed for this example, but that made the final file size 63 bytes smaller. One pixel can make more than a bytes worth of difference in a GIF image.

ImageVice™ 2.0 Documentation



After tweaking, the file is 4278 bytes, without any visible difference in appearance.

For the purpose of this example, only minimal tweaking was done just to illustrate the possibility. Much more could be done following the rules of thumb given for an even smaller file.

When file size issues are most critical, your own two eyes are the most valuable tools you have to squeeze extra bytes off of your images.

G: Post production issues with ImageVice™

Since ImageVice™ is an image filter and does not actually create any files, there are post production issues that you need to be aware of, else you may inadvertently negate any advantage that would have been achieved by ImageVice™.

Optimally you should use ImageVice™ in conjunction with PhotoGIF™ to produce GIF files. They are designed to work very well together, and can save several post production steps for you.

In the absence of PhotoGIF™ you should follow the procedure described below for producing GIF files in Photoshop.

1: Change the document mode to "Indexed" first in Photoshop, without changing any of the settings in the mode change dialog. These should be "Exact" palette and no dithering. Selecting any other settings in the mode change dialog will cause Photoshop to color reduce the image again and negate all benefits from ImageVice™.

ImageVice™ 2.0 Documentation

Do not save your document directly from RGB mode using the GIF89a Export plug-in. This will cause your final file size to be larger than it would be otherwise, and can result in unpredictable results because of the way the GIF89a Export plug-in handles transparency when saving from RGB mode images.

2: Save the indexed mode document using the Compuserve GIF plug-in. The "GIF" format option in Photoshop's standard save dialog is the Compuserve GIF plug-in.

However, if you are using the Compuserve GIF file format plug-ins on a Macintosh, be aware that their files will contain information in the Macintosh resource fork – even if you have custom icons and previews turned off in Photoshop's preferences – which will make the files appear larger in the Finder's Get Info dialog.

To accurately tell how large your GIF files are if using a Macintosh, you should download the free GIF Prep utility, which will strip resource forks from your GIFs, so the Finder will report an accurate size for the GIF file.

3: In host applications besides Photoshop, the proper procedure may vary, and there is no way to address all possible host applications that ImageVice™ may be used with. If you aren't using Photoshop, you should make a point to use ImageVice™ with PhotoGIF™ to ensure you are receiving the full benefit of ImageVice – not to mention the benefits PhotoGIF™ provides for GIF production.

ImageVice™ 2.0 Documentation

Using ImageVice™ for GIF animations

This is an article written by Travis Anton, so you know who the 'I' in it refers to, that used to be in the now defunct Tips and Tricks section of our Web site. It is included as an extra bonus to the manual now.

After months of working with MWB Interactive and invisible, ink. on the new look for our Web-site, one last thing remained – making the pages load fast enough that people would wait around to see them. It was a job I had saved for myself, and not an easy one – or at least it wouldn't have been an easy one without ImageVice™, PhotoGIF™, and GIFmation™. The web site as it is, wouldn't have been possible without them.

The primary designer at invisible, ink., Ed Sultan, had been given the full set of BoxTop's tools to work with at the beginning of the project, PhotoGIF™, ProJPEG™, GIFmation, ColorSafe, WebClipper, and ImageVice™, which he used to create the original graphics and animations for the site. He did a pretty good job at optimizing files, but not good enough to really show off.

>Re: Squished...

>Ed, > >Two down. 41,474 bytes to 14,615 bytes on the projector. >More to go...

>You're just gonna keep rubbing my nose in it, aren't >you?! You're >a FREAK!

I used the same tools to make the final set, but as you can gather the ones I made were many times smaller than Ed's because I knew how to use the tools just a little better. (OK, I do have an unfair advantage in that department having been the one to design the software, I am the expert of experts on how to use it best.) That knowledge is what I'm about to share in this article – how I crunched the alligator-gun and other odd-ball gadget animations on our site down to size using ImageVice™, PhotoGIF™, and GIFmation™.

ImageVice™ 2.0 Documentation

I got the original artwork for the alligator-gun as a Photoshop layers file from Ed. There were seventeen layers of artwork for five frames of animation, and they had not been created in such a manner that I could directly import the layers file into GIFmation™ as it was. I didn't want to take that approach anyway, though.

My plan was to use the new ImageVice™ plug-in for color reduction to achieve better optimization results than could be achieved otherwise with only GIFmation™. ImageVice™ uses new color reduction techniques that make indexed files up to 70% more compressible than other color reduction methods.

ImageVice™ even beats the color reduction in PhotoGIF™ and GIFmation™, which now rate only second best in the compression category, though still the more versatile overall since ImageVice™ can't dither if needed or use fixed palettes or create super palettes like both PhotoGIF™ and GIFmation™ can.

Step one – Preflight planning.

It never works best to just jump head-first into things without some time to think about what needs to be done. I took several minutes to look at the original art and see how the motion in the frames worked taking note of which parts of the image "moved" by switching layers in Photoshop. I also compared what I was seeing in Photoshop to the original animation Ed had created in GIFmation™. Then I took careful note of colors in the layers.

The primary color scheme was a metallic blue-gray, with only four other distinct colors that didn't fall in the blue-gray range. The first critical decision to be made was bit depth and how many colors to use for the best compromise between overall visual quality and the final file size for the animation.

Color reduction to 32 colors would leave 26 colors for the blue-gray range – plenty to be realistic, but color reduction to 16 colors would only leave room for 12 colors in the blue-gray range. That wouldn't be enough, so I set my target at 5 bits and a maximum of 32 colors.

ImageVice™ 2.0 Documentation

The next thing I did was think about how to get the best results from color reduction. ImageVice™ works as a true filter in Photoshop, so one option that allows is the ability to use different settings in ImageVice™ for separate areas of the same image. Quality and compression can be varied by image area for best effect.

In the layers, the "gun" and its shadow were the high entropy areas. I knew they would be responsible for the majority of the resulting file size, while other background areas were very low entropy areas with little continuous tone.

I saved a copy of the original layers file to begin modifying.

Step two – color reduction

I selected a layer containing only the "gun" image for the first frame and pulled down the menu for ImageVice™ then typed in 28 colors – saving room for 4 colors latter for the colors in the background area of the animation – leaving all the other settings at their defaults and clicked OK. The default settings in ImageVice™ are moderately high on the compression scale. They work extremely well on most images to achieve better compression and image quality at the same time, but on this image the defaults for convergence were a little too much. It had more smooth gradient areas than most images with the metallic surface, and convergence settings become more obvious in areas of smooth gradients. I undid ImageVice™ then ran it again, turning down the convergence settings to about 20:3, which worked beautifully.

I quickly ran ImageVice™ with the last settings used on all the other gun layers, and the first step in the color reduction process was done.

I needed to composite the five frames before the second step in the color reduction process, so I duplicated and flattened with the right layers visible for each frame until I had five separate frame images.

Then I created a cell-strip format document – one long enough for all the frame images to fit side by side in horizontally – in Photoshop to paste my frame images into. I needed them all in one document for ImageVice™ to create a single palette for them all. The current version of ImageVice™ can't create super palettes from separate images like can be done in PhotoGIF™ or GIFmation™.

ImageVice™ 2.0 Documentation

The results were almost perfect – to the human eye they were, but since the animation had to line up and match perfectly in color with other graphics I needed to make sure they were also mathematically perfect as well and correct color drift in the background areas of the frames.

Since Ed had used colors from the Web-safe aka. Netscape palette for the background colors on our site and in the animation, this was easy as pie. I just ran the WebClipper plug-in (included free with ColorSafe), which automatically forces close colors to exact matches from the Web-safe palette to make sure the colors would match the surrounding graphics perfectly.

It was time to start optimizing, now that I had achieved the optimal color reduction possible for compressibility.

Step three – optimization

Optimization, as it means for GIF animation, is simply the elimination of redundant data. Areas of the image that do not change from frame to frame are made transparent so that the image will compress more efficiently. The process is similar to more complex delta compression methods used in real video formats, but it isn't quite the same thing.

An important distinction needs to be made here between visual and mathematical change. Software has an inherent limitation in that it deals only in numbers and absolute values. Human vision is not that discerning. Two colors that are visually indistinguishable in the same surroundings are not always mathematically the same color even if there are only 32 possible colors they may be, as in the case of this animation.

You can achieve better optimization with a little extra manual effort and your own two eyes every time. GIFmation™ has the best one step, automatic optimization around, but it can't compare to the results if you put a little human judgment into the process. GIFmation™ is also the only GIF animation tool that provides the ability to make manual optimizations quickly and easily. It is the perfect optimization tool, but like all tools the actual results depend a bit on well it's used.

ImageVice™ 2.0 Documentation

I didn't have an animation for GIFmation™, though, at this point. I just had an image document with the five frames positioned side by side horizontally in Photoshop – a format that could not go directly to GIF animation through PhotoGIF™ or GIFmation™.

My plan was to open Ed's original GIF animation file with PhotoGIF™'s cell-strip mode and paste in my frame document, which was why I had created it that way. PhotoGIF™ can only save GIF animations this way if it opened a GIF animation in cell-strip mode, otherwise only a single image could be saved with PhotoGIF™ from a single document.

If I hadn't had Ed's original animation, I would have needed to make a dummy five frame GIF animation in GIFmation™ quickly to create a template for PhotoGIF™ to use. All that would have been needed were five frames of anything and the correct logical screen size set.

Before I saved my frames as a GIF animation using the above trick I wanted to take the opportunity to rough out the transparency for the frames while they were visible side by side in Photoshop. This could have been left to GIFmation™, but since I was already in Photoshop, it was easy to use its QuickMask feature to create an alpha channel mask that PhotoGIF™ could use to set transparency with.

The advantage of doing this initial rough step in Photoshop was that I could see all of the frames at once and easily tell which parts of the frames never changed. The same could have been done in GIFmation™ alone, but it may have taken a little longer because GIFmation™ hasn't got the sophisticated brush and selection tools of Photoshop – just the basics.

After I masked the first frame, I copied and pasted to quickly finish masking all the frames.

With that done, I opened Ed's original alligator-gun animation using PhotoGIF™, then changed the image mode from Indexed to RGB before pasting my frames document, and the alpha channel I had created for transparency, into the original animation document PhotoGIF™ opened.

ImageVice™ 2.0 Documentation

I changed the image mode back to Indexed after pasting with Exact palette selected and dithering turned off, which was the default in Photoshop's mode change dialog since after using ImageVice™ on the new frames there were less than 32 colors used in the document.

All that was left was to save the file with PhotoGIF™ so it would then be a GIF animation document that I could open in GIFmation™.

I saved it to a new file and set the transparency in PhotoGIF™ from the alpha mask I had created, then switched to GIFmation™.

After opening the partially optimized animation in GIFmation™, the first thing I did was export it as optimized to let GIFmation™'s magic work before I spent more time working my own.

I closed the partially optimized animation and opened the one that had just been exported using GIFmation™'s automatic optimization.

Now it was time to go through the individual frames and see what could still be done to make the file even smaller.

I removed stray pixels and visually unchanging areas from each frame with GIFmation™'s transparency brush, stopping to preview the animation in between frames to make sure that I hadn't made any visible error. When I had satisfactorily tweaked the four optimized frames, then the last step needed was to use GIFmation™'s automatic optimization one more time to take care of automatic cropping where my tweaks had made it possible again.

A last preview to confirm all was well, and I was done. In all it took about half an hour for me to complete final production from the art to the finished animation.

Compared to the days it took for Daniel Thomas at MWB Interactive to conceive, model, and render the alligator-gun and the time it took Ed Sultan at invisible, ink. to make it a working part of our web site, production time was trivial in perspective, even going whole hog on tweaking to make the animation as small as possible.

ImageVice™ 2.0 Documentation

The net result of my efforts – an animation much less than half the size of the first, 36,893 bytes down to 14,234 bytes for the alligator-gun. I was able to make other tool animations as small as 8k each with the same techniques.

Spending an extra twenty minutes on production for the alligator-gun animation saved 22k of file size off a graphic that will be accessed an average of 5000 times a day – 22k off one animation means around 108MB of data transfer saved every day for our site, and a page that loads twice as fast for our viewers.

Think of what the same savings would mean to you.

ImageVice™ 2.0 Documentation

Contact Information

BoxTop Software, Inc.
PO Box 2347
Starkville, MS 39760

Voice: +1-662-263-5410
Fax: +1-662-263-5412

info@boxtopsoft.com
<http://www.boxtopsoft.com>
<ftp://ftp.boxtopsoft.com/pub/>