

# What Makes a Good HDRI and How to Use It Correctly

HDRIs are everywhere these days. If you've got a half-decent camera, a tripod and some software you can even make them yourself.

But just like creating art in Blender, being able to do it at all is not the same as being able to do it **well**.


So, after I created my first crappy HDRI and discovered how challenging it could be, I decided to embark on a quest. I wanted to create the perfect high dynamic range environment map that would give you perfectly accurate and realistic lighting as if you had teleported your CG scene to the actual location of the photo itself.

In truth, this is an unending quest, but I've made some fair progress over the years. So without further ado, let me explain...

## What Makes a Good HDRI

Just like art, the quality of an HDRI can be a subjective thing, but I think we can all agree that there are a few fundamental attributes that define (although not exclusively) how useful or accurate an HDRI is.

## Dynamic Range



Let's begin with what is, to me, the most important aspect of any HDR image that you intend to use for lighting.

Dynamic range is an indication of the **contrast in an image**. Typically you'll see something like "12 EVs" which means that when they took the photos, for each camera angle they captured multiple photos at different exposures that encompassed a 12 EV range.

**But what the heck is an EV in the first place?** Well it stands for Exposure Value. As you might know, *exposure* refers to the amount of time that the camera shutter is open, and thus the amount of light that gets captured on the sensor. But an *exposure value* (in terms of HDRIs) is a measurement of the *range* of exposure. So if you capture two images, one twice as bright as the other (double the exposure time), and merge them into a single image, that new image would have a dynamic range of **one EV**.

*(Side note for you techies... This is a pretty silly way to measure it in my opinion, because a single shot already encompasses a certain range of brightness (dynamic range), and depending on the camera you use this can vary between 10 and 14 EVs. But it seems to be the norm for HDRI dynamic range measurements to ignore this fact and just include the **range between the shutter speeds.**)*


Each EV means double the lighting information, so a dynamic range of 2 EVs means you captured three photos, each one doubling the last exposure time. You might start with an exposure of 1/16th of a second, then do 1/8th and then 1/4th.



*Side note: Cameras can't do 1/16th exposures, instead they do 1/15th because they want to make our lives difficult*

But most environments have an extremely wide range of brightness. The filament of a light bulb might only be properly exposed (not over-bright or “blown out”) at 1/4000th of a second, and in the same view you can see the corner of the room cast in shadow, which might require an entire second of exposure before it’s bright enough to see clearly.

So to capture all of this you would shoot 5 photos with a 3 EV gap between each one. The exposure times would be 1/4000th, 1/500th, 1/60th, 1/8th of a second and finally 1 whole second.



*Both the light bulb and the darkest shadows are properly exposed*

You might think that 5 photos with 3 EVs between them would mean a dynamic range of 15 EVs ( $5 \times 3$ ), but in fact there are only 4 gaps between the 5 photos, which means the dynamic range is really 12 EVs ( $4 \times 3$ ).

This is all quite easy to do, in fact the camera does all the work of calculating the right values and capturing them in sequence, all you have to do is point and shoot.

The challenge comes in when the range of brightness you need to capture is so high that it's difficult to get it all in one go. For example, any time you need to capture the sun. The sun is so bright that even  $1/4000$ th of a second is too long and captures too much light. And usually the fastest any typical commercially available camera will go is  $1/4000$  or  $1/8000$  (which is still not enough).

Since this is a difficult problem to solve, most folks just don't bother at all and thus we're left with **HDRIs that do not contain the entire brightness of the sun.**



*Flashing red area is a solid colour, indicating the brightness is clipped and therefore incorrect.*

But it's not an impossible problem. You can buy **Neutral Density filters**, which essentially block a certain amount of light, allowing you to shoot darker photos. Thus you can shoot a larger exposure range and, with a strong enough filter, capture the entire brightness of the sun.

My point is this: It's quite difficult to capture the whole dynamic range of a scene in many cases. For an interior shot, 12 EVs is usually enough, but for an outdoor shot where the sun is visible, you might need around **22 EVs** to capture everything depending on the time of day and the weather.

If an HDRI does not contain the complete dynamic range possible for that particular shot, the lighting it creates in your 3D scene will be incorrect.



## Resolution

If you only care about the lighting that an HDRI gives you, resolution isn't actually all that important. A 2000×1000 HDRI will give you lighting almost identical to a 16000×8000 one. If you go smaller than 2k, you might start to see blurry shadows, but even that is not usually a big deal.

But lighting isn't everything, and in many cases you'll have very shiny surfaces that reflect your HDRI. Depending on the focal length of the camera and how flat the surface is, I find an HDRI needs to be at least 4000×2000 in order to get crispy reflections on most surfaces.

Additionally, if the HDRI is actually visible in the background (e.g. the sky), resolution is obviously very important as you don't want any of that ugly pixelation nonsense distracting the viewer from your awesome artwork.



*1k vs 16k HDRI in a 1080p render, 35mm focal length*



*100% crop*

No surprises there.

Now let's say you're rendering an image in Blender and want to know what the minimum resolution of the HDRI you'll use in the background should be...

The horizontal field of view of the 35mm (default) camera in Blender is about  $49^\circ$ , which is about 14% of  $360^\circ$ . If the resolution of the render is  $1920 \times 1080$ , then the HDRI's horizontal resolution needs to be at least  $1920 * (100/14)$ , which is just shy of 14000 pixels.



But of course you may want to zoom in a little, so I'd say that **14000×7000 is the bare minimum** for any HDRI, though you should rather try and find something bigger.

The problem with large HDRIs is the amount of memory they consume. An HDRI is stored in memory as a 32-bit image, meaning each channel for each pixel needs 32 bits of memory during render time. This means that a 14000×7000 HDRI will consume about **1 GB of RAM**. If you're rendering on your CPU that's not such a big deal, but if you're rendering on your GPU, RAM is suddenly much more precious. Add to that the geometry of your scene and all the other textures you're using, and chances are your GPU simply doesn't have enough memory and you won't be able to render with it.

The formula for working out how much memory an image will use is:  $\frac{((\text{bit\_depth} * \text{number\_of\_channels} * \text{x\_res} * \text{y\_res}) / 8) / 1024}{1024}$  to get it in megabytes.

## Aesthetics

Finally we reach the matter of beauty and taste.

There are certain things like giant stitching artefacts, lens flares, chromatic aberration, wobbly horizons and the shadows of cameramen that we can probably all agree are bad things, but what about the beauty of the lighting that an HDRI gives? Does it even matter? Surely someone can find a use for even the most boring overcast sky?

I suppose.

But as an artist, it matters to me. I don't want to make boring things, even if they are useful to someone way out there in the outer reaches of the internets. I want to make things that are pretty, things that can turn your grey, lifeless model into a beautifully presented work of art.



To me, lighting is the **single most important part of the CG pipeline** because it's what creates mood and emotion in an image. Obviously you couldn't have an image at all without modeling, texturing, composition and all the rest, but lighting is responsible for the sole purpose of why we create art in the first place: to express our thoughts and feelings, to **put our emotions into something that will outlast us**.

Without good lighting, all the art we could ever make would be little more than a form of factual essay without any adjectives.

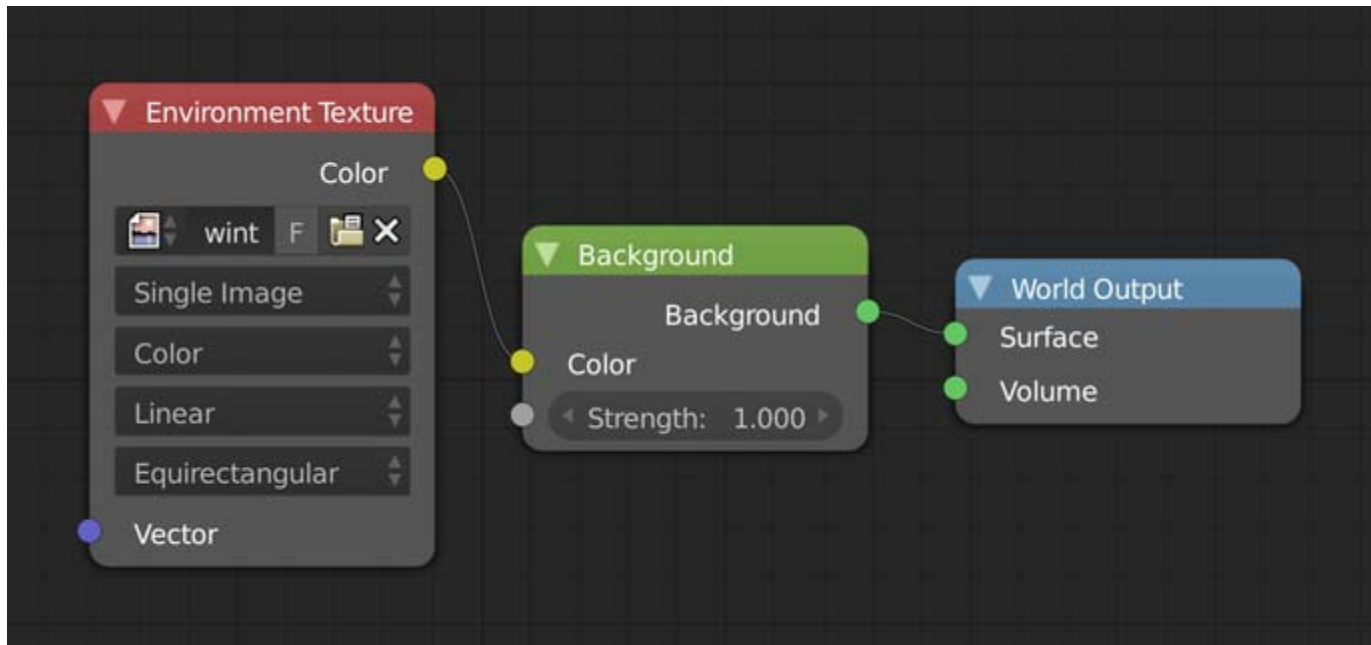
OK, maybe I'm being a bit melodramatic, but you get the idea. In my opinion at least, a good HDRI is one that makes all the rest of your hard work look good.

## The Right Way to use an HDRI

So there have been many tutorials and articles that claim to show you the proper way to do HDR lighting in Blender. I've done some incorrect ones myself in fact.

But to be honest, there's really no special trick, and practically no way to guarantee that your lighting will be 100% accurate to real life due to the simple fact that colour management is a bitch.

So in the case that an HDRI was shot with the full dynamic range possible (including the full brightness of the sun), then the most correct way to use such an HDRI is simply this:



Yes, that's all. **No fancy node setups or expensive plugins.** What you have right there is the most accurate lighting you could ever get, and it takes just 15 seconds to set up.

In the past I have said that you're supposed to plug the HDRI into the strength input of the background shader too, but **that is wrong.** The only reason you might want to do that is if your HDRI's dynamic range is not high enough and the sun's (or any other light source's) brightness is actually clipped. Sadly, this is true for most HDRIs out there, even the ones you have to pay for.

Plugging it into the strength input is basically the same as multiplying it by itself, which is like increasing the contrast (the brighter pixels get brighter, and the darker pixels get darker). If the HDRI *does* contain the full dynamic range, then **doing any sort of manipulation to it will potentially ruin the accuracy.**