

Projection Texturing a Human Head with Open Source.

Disclaimer: This is my first tutorial at this level. There may be more effective methods, but this is to the best of my current knowledge.

We have all seen the realistic digital heads and looked lovingly at them at some point I sure. Maybe there was that one image that inspired you to jump into the world of computer graphics. Wondered if Blender and its open source buddies were up to the task? I believe they have reached a point where they are. I have even ditched Windows XP to proudly bring you this tutorial from Linux. (Note, Blender and Gimp to my knowledge work pretty much the same on any operating system.)

Required Software:

1. Blender (tutorial written at time of 2.45, should work in future versions as concepts are fairly basic at their core.)
2. Gimp (version 2.4.0rc2 used, but any version with iWarp filter) *Photoshop can be used, but the aim is for a complete open source option.
3. A head model. Feel free to use your own, but here is the one I used for this tutorial, complete with finished skin material and 2k textures. [Link to .blend \(packed in 27mb .rar file\)](#)
4. A set of reference textures. Once again, feel free to use your own, but I used free images from <http://www.spectralogue.com/textures/index.php?searchSub=off&keyword=&path=78&startimage=0>
A great reference site is www.3d.sk – not amazingly cheap for hobbyists, but great range and quality for what you do pay.
*****Nudity Warning!***** - Whilst the main set I have referred you to is simply head photos, it is one click away from full body references. The warning is common on the www.blenderartists.org forum - I respect that and maintain it here.

Assumed Knowledge:

This tutorial will contain intermediate level content, but explained as user friendly as I can. If you are new to Blender, try <http://www.math.sunysb.edu/~sorin/online-docs/blender/html/c798.html> or if this link is down, search for "Your first animation in 30 minutes" and complete that tutorial MINIMUM before this one as I will not be covering many shortcuts if any.

In short, I will be mainly using Blenders materials tab and baking function, as well as Gimps basic eraser, brush, iWarp filter and the clone tool. Modelling is NOT covered in this tutorial. Rigging is NOT covered in this tutorial. If I included these topics accurately enough, I could easily fill a book. So texturing it remains.

Required Hardware:

The more grunt your machine has the better. For the record I am using a single core 3.2Ghz machine with 2GB of RAM and a fairly standard video card. The main thing to keep in mind that will chew up your system resources is that I will be going through this tutorial creating very large images. If your system specifications are higher than mine, then good for you, I'm jealous. Potential lack of RAM becomes your main enemy here. That being said, feel free to halve the values I give for each texture to make your images a quarter of the size.

For example – whenever I say 4k (or 4096x4096 pixel) images, just make yours 2k (2048x2048 pixel) images. That is still high enough resolution to make great images for the web, but 4k images will allow

face close-ups in HD movies and stills for print depending on the print quality, but most studios would not go higher than 4k textures for production purposes to my knowledge.

Part1: Looking over our assets

The subject I have chosen is for a few reasons. The reference images are free, which means anyone can participate in this tutorial. They are not perfect, but are high quality. Finally, the subject is a younger female. My reasoning for that is more sound than you might think. Of all the objects to represent, we identify with images of humans the most and scrutinise them for flaws because we are humans. (i.e. You would pick flaws quicker in a computer generated human quicker than a computer generated picnic table.) Secondly, a younger female face has far less defining features such as wrinkles on an older face or stubble on male face. If you watch *Final Fantasy – The Spirits Within*, there is a fair chance you will see the older male doctor as the most convincing character because of the detail that his wrinkled bearded face provides. So in that sense, a younger female face is a very strong candidate for the hardest subject matter to achieve convincingly in 3D – one of the reasons that younger female portrait images flood the galleries of most high end 3D forums and books. The techniques in this tutorial however are not limited to this subject matter though, you could texture anything from a space shuttle to a monkey if you wanted. Now on with the project!

The model I have provided and the reference photo set I have used do not match.


Say *what?*

Why go any further with this tutorial if the assets don't match up? Fair enough question, but method to the madness people. Lets say you wanted to use this tutorial make a CG head of, oh I don't know – Natalie Portman. If you were working in a large production company, you would call her in for a professional photo shoot, take 30 or 40 photos on your extremely good Digital SLR in your well lit studio from any angle you choose. Then she would kindly step in front of your 3D scanner and a 100% accurate head model appears. Use the re-topology tool to get a clean mesh, UV unwrap and you are ready to texture.

You get my point. In reality, you can do a quick search to get high resolution photos of many (famous) people you may want to use as a subject, but to get enough usable angles with matching lighting – not going to happen. If you have a great camera, then you can get someone to take photos of you or you of them, that's fine. But you will probably cause more work for yourself removing shadows and specular lights than using professional images taken in near perfect studio conditions.

Back to our little project and where I am headed with this. We will use images that are close to the model with a skin tone we like and then use the iWarp tool in Gimp to make the photo line up better before we project it on to the mesh. Not a perfect solution, but a practical one for the average CG enthusiast.

The images from the suggested reference set I will be using include:



- DSC_6773.jpg – front projection top of face



- DSC_6758.jpg – front projection bottom of face (Ideally use one reference per image, but the better alignment to start with, the less fiddling later on.)



- DSC_6781.jpg – side projection, both sides (Best direct side angle shot. You could line up another shot of the other side for authenticity, but you will never see both sides at once. Other option is to clone some freckles or something to the other side to avoid symmetry. I will explain later.)



- DSC_6753.jpg – three-quarter projection. (Depending on how the front and side projections turn out, we will probably still need this image, so download it anyway. The only drawback is this image is a little blurry around the eye and upper cheek where we want detail. If you are used to working with 1024x1024 pixel textures then you will think I making a big fuss over nothing. However, I will be doing this tutorial to generate 4096x4096 pixel textures, in which case it will show up. In this case, we would opt for the other two projections where possible when merging together in the Gimp. Again, I will explain more when we get there.)

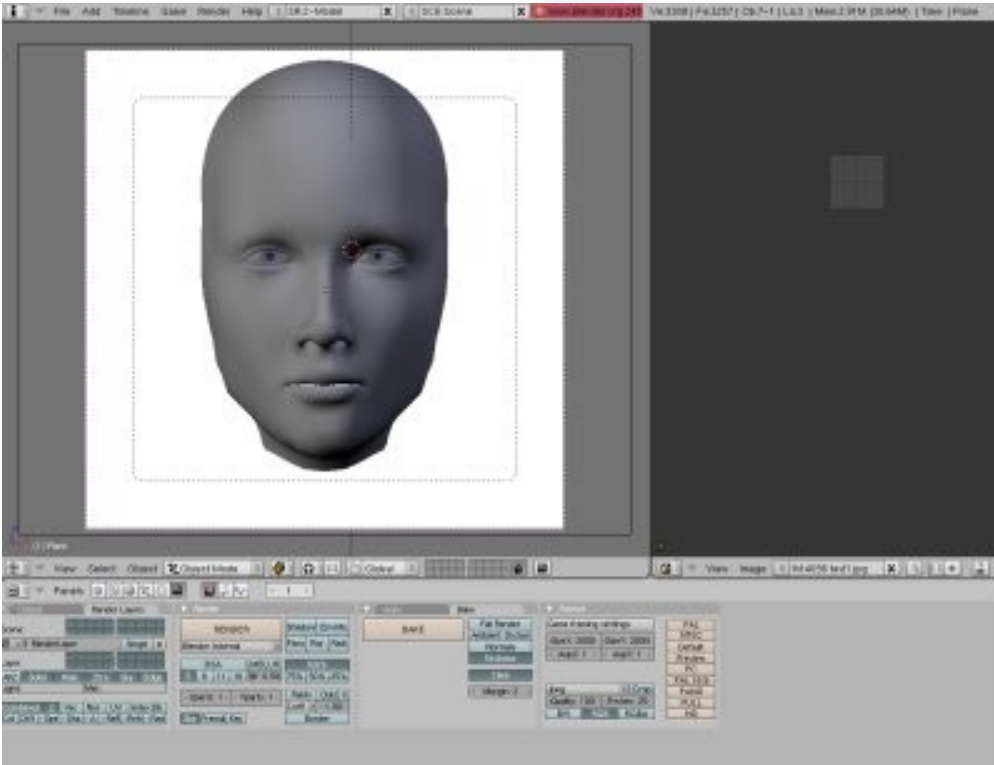


- DSC_6772.jpg - neck projection. (I chose this one over DSC_6771.jpg because even though 71 had more neck showing, 72 showed more than enough and the lighting was more uniform. More uniform light saves you more work later on and helps for more accurate textures in general because you don't have to tweak as much.)

That should be it for now, but we may need to come back. Lazy tutorial writer you say? You are probably right, but it pays to keep an open mind that you might need to go back for more reference/resources during a project.

Part2: Lining up the images for projection

This part here makes all following parts easier. Easier is a good thing. Open the .blend file I have provided for you (or the one you made on your own. Kudos to you!)



This should be roughly what it looks like – A camera, the head mesh (back half hidden by the plane) and the plane. There are a few simple lights in the scene and that is about it. The camera is an orthographic camera (ie. not perspective with depth, more like when you draw a blueprint) with a scale of 1, rendering with a square aspect ratio (same width as height). The plane is funnily enough scaled 1x1 blender unit. By keeping everything simple and uniform, it should take any guess work out of “will my projection fit exactly?”

Render the image from front view. If setup as above, just a matter of hitting F12. I have set the render to 2000x2000 pixels, which is quite possibly the largest render many of you (but not all) have rendered. Never fear, there is no raytracing, shadows, anti-aliasing (OSA) or anything else to slow the progress down meaning the render should be really quick despite the size.




Hooray, I rendered a still image... why did I start this tutorial? Is he taking my learning seriously? - Trust me, it picks up.

Save that image as 'front_mesh.jpg' or similar. (Naming files is important in this tutorial as we will deal with large image files. It is a pain to forget which is which when they are labelled “image1,” “image 2,”

etc and you have to open many 4k images to find what you are after. Too RAM intensive, just easier to name everything.)

Open that image in the Gimp (or Photoshop, etc) and create a new layer. (Top menu bar, Layer – New Layer)




Open `DSC_6773.jpg` and copy the image.

Paste that image into the new layer of the 'front_mesh' image and save it as a .psd file or a file format that can support layers. (I'm still fairly new to both Gimp and Linux, but I am learning to make the switch as I go.) You may have to right click on “Floating selection – Pasted Layer” and select “New Layer” to convert it to a usable layer.

Make the new layer 50% opacity (don't worry about that exact value, just so long as it is enough to see both layers well enough.) Guess where I'm going with this?

Move the pasted layer to match up as best you can – this image was picked because of the straight on top half of the face. Eyes are key, so line them up as best as possible and we'll deal with the rest.



Something like that. Now you might be able to skip this next step if you are using references that match your model, but unless the images line up perfectly and your model matches the images perfectly, it is probably still needed.

There is a big difference in how easy this next part is in Photoshop and the Gimp. Photoshop Liquify filter can zoom and show multiple layers, in which case you just use the main move/distort tool to make the two line up. Main hints – Distort only as much as you need to, and distort larger areas to start with to keep it even, then distort smaller if you have to. Starting with many small distortions has a very high chance of turning out messy.

Gimp's equivalent feature iWarp can only show the one layer preview (at time of writing) and the zoom is only accessible by resizing the window and clicking “reset”. I found out this fairly major setback literally while writing this tutorial, but I really want to see this work in open source, so we shall find a way around.

First we will work on the ever important eyes. They are the 'window to the soul' so to speak. In order to

make this more practical, use rectangle select (top left button in toolkit if you haven't found it) and select a large rectangle around the eyes and eyebrows. Copy and paste as a new layer, drop its opacity and turn off the full face layer to make things easier. This smaller chunk of photo will give us more room to work with in the iWarp tool, (Top menu bar, Filters – Distorts – iWarp). Take a good look at how the eyes do and don't fit. In my case, they are a bit on the large size, mainly vertically. The bottom eyelid needs to come up a little, the top down a fair bit, but without disturbing the eyebrows too much because I like them. (I know you can't see this in the small screenshot, but take my word for it.)



Having done this in Gimp and Photoshop (PS), I have to stress that PS is much simpler if you have it, but if you can set up your screen something like the screen shot above, you can see where the eye level/facial detail was and get a feel of where you are moving it to. Another big hint to make things less frustrating – **don't** do more than three or four moves/edits in one go before committing changes in iWarp. Liquify in PS has undo, iWarp has a 'reset' button. In other words, the eyes could be perfectly fitting, then you accidentally nudge your mouse/tablet pen and you have to start again.

Once the eyes are done, duplicate the whole face layer. (I like to keep a copy of most layers, in case.) Right click (or the Mac equivalent, haven't forgotten you all...) and select "Merge Down" on the Eyes layer. Same function in PS and Gimp.

Use the rectangle select tool again and make a more vertical (taller, not longer) selection around the nose and repeat the last few steps, finishing by merging it down. **CHECK** before merging that you have not changed details around the edges of the nose (or eye, or anything) layer, otherwise there will be a noticeable line in your maps, particularly when it comes to the bump map. This is why I said to select an area larger than you need, but small enough to be effective. Note with the last screen shot that I tried to avoid affecting anything within a "ring" around the eyes about the distance between eye and eyebrow. This leaves plenty of gap for the layers to merge back seamlessly.

Once the nose and eyes are done, that should be it for this reference image. Save the .psd (or similar) file and export as a jpg/tga or a file Blender likes to project with. Remember your easy to understand file names! From here there are two options – Project this image as is onto the front of the face and do a separate projection for the lower face/mouth and merge afterwards as projected images, or merge the mouth on to this image and project once. Personally, I do not think it matters, either way there is work involved, and you are more or less doing the same steps in a slightly different order. How would I do it you ask? More often than not, I would choose the first option and blend them after they had been projected. You already have to blend multiple images in projected form, this way you can choose which parts of which image lines up better with the other (e.g side) projections at that point.

Note: I'm not going to walk through how to line up each and every photo reference angle. *Awww... why not?* For one reason, that would needlessly stretch the tutorial out a lot. The other more important reason,

when you try this again with a different model and a different set of textures, (I'm trying to teach skills here, not a rigid process that works with one strict set of assets) there are going to be photos from many different angles, different areas that need work and I can't cover them all detail by detail. You would get really bored, I'm sure of that much. However, I will give hints:

1. Line up the general position of the image (rough)
2. Rotate the overall image if needed
3. Scale (**only** if needed) the image if needed by referencing a landmark like nose, mouth, eyes, ears or *maybe* even cheek shape, then adjust position to fit as best as possible on that landmark. Also note, you should **only scale down**, otherwise all your detail will stretch and start to become blurry before you even project it.
4. If in Gimp, select a chunk to work with in iWarp, following the advice I gave earlier. If struggling a bit, just keep at it, it does get easier! That is why I suggest having a duplicate layer, in case you need a bit of practice.

Part3: Projecting the images

If everything in Blender lines up and everything in Gimp/PS lined up pretty well, this next part will be very satisfying. Not that the work ends here, but it can be a bit of a shot of inspiration to see it begin to work. Load Blender up if you closed it with the same scene you rendered from. Might seem obvious, but just in case – **DON'T** move the camera between the initial render and the projection of that camera angle (whether it is front/side/whatever) or else all your tweaking in Gimp/PS was a complete waste of time!

I trust you didn't though, so the next step is to go to the texture on the head object and set it to the image you just exported out of Gimp/PS. (I called mine "tweak_front.jpg"), with the texture affecting the colour of the material at 100% (or 1.000)

The parts that make projection are setting the mapping coordinates not to 'UV' or 'Orco', but to 'Object' and giving the name of the plane in the box (see below) and Blender's bake functions.




Enter Blender's UV editor. Versions 2.45 and before, do this by hitting the F key with the head mesh selected. Versions 2.46 onwards (at time of writing) have merged the UV editor with the edit mode. Just enter edit mode (TAB key) and that will put you in UV mode as well. The .blend file I posted has a split window above the button window, one for the 3d viewport and one for the UV editor next to it. Select all

the faces of the head mesh (A key, might have to hit twice if something selected beforehand). The faces should all appear in the UV editor now. With everything still selected, in the UV editor go to the menu labelled 'Image', click 'new' and make a new image that is 4096 pixels by 4096 pixels (or smaller if you are scaling back as mentioned at the start of the tutorial – try and keep to multiples of 512 – 1024, 2048, etc because without getting sidetracked in detail is more efficient to edit/render for your computer.)

Save that image (I know it's blank) as “front_face_project1.jpg” or whatever is appropriate for the angle you are about to do.


Under Render Buttons (where the big “Render” and “Anim” buttons are) click the Bake tab.



From before, the OSA, raytracing and everything else should be turned off. Turn OSA back on (I think it helps, not exactly sure.) Make sure Bake is set to “Textures” – The others won't help us, even “Full Render” will bake on the shadow detail, which we don't want, especially when we already have some shadow detail to paint out later.

Fun Part – Hit Bake! Then watch the first chunk of your texture come alive for you, with skin pores, freckles, the works.

Hopefully that didn't take too long. Save that freshly baked image.



Here is how mine came out, shrunk to 512 pixels for web. Whilst it might not look that usable yet, you have to keep in mind that all we lined up for use was the forehead down to just under the nose. Ignore the rest, look in that region and I think we are off to a good start! Already you should be able to see where the shadow and highlight is an issue- take a look at the nose particularly. Dark edges and a really bright tip on the end. To be fixed later of course.



Here is the first projection texture with most of the unusable parts blacked out just to show you the progress. The main advantage of this technique over the painting from scratch method is that all your facial pore detail is already staring right at you. (Material settings for this preview – shadeless, emit set to 1.000, ref set to 0.000, so all shading comes from the texture at this point.)

Repeat this process, but adjust the angles. For example, rotate the mesh of the Head (not the plane or camera) 90 degrees for side on. (Holding Ctrl key snaps to nearest 5 degrees when rotating.) Pretty simple really. It pays to work on the one angle at a time. All front angle projections, then all side ones, then each angle finish up individually BEFORE rotating to another one. Straight forward way of working, just remember not to move the camera in between in case I haven't stressed that enough.

Also remember to create a new image to bake to **or else it will bake over your previous one!**

My next two projections were the lower mouth and overall side of the face respectively:



The eyes don't match on either projection and there is a smear down the middle!! What is he doing?!

The eyes don't need to match on a projection that is for the lower half of the face. For the side of the face... it's a bit sloppy of me. However, I like the eyes from just the first front on projection so the side isn't really needed in that area. As for the large white smear down the centre – the front projection will be over that. The side of the nostrils and the bridge of the nose however do matter because there is a lot of smearing in the front on projection, you will notice I took a little more care there.


Of course you could line up each angle 100% perfect, but in general that is a waste of time. If you were doing this for a client, you could probably be up to the specular and bump maps with the time you save cutting these little corners.

I am itching to start getting this map together, even though I know we will need to do at least a few more projections – mainly to get rid of all that excess hair!

Part4: Blending the projections together - A

To start of with, this part is alright. Open your first projection and then import/open/copy'n'paste (however you wish) the second one as a new layer. In this case it is the lower part of the face front. Crop out any obvious parts that don't match/aren't needed. (Big smear under the chin, anything to the sides of the face.) Then use the eraser tool to blend the top and bottom halves together. Just check as you erase that the sections you just got rid of weren't as good as the ones underneath on the first projection. Turn the second layer on and off fairly frequently to help check this. Undo buttons are your friend here. If you find it helps, you can crop what you do not need off of the first projection as well, I personally like to leave at least one layer complete and build the rest up. Up to you, whatever confuses you less and gives you freedom to work more efficiently. I recommend you save your progress as a multi-layered document - .psd or whatever you are used to,


Progress Result:



Not bad, right side of the mouth might be drooping a bit, we'll see how that renders in a moment.

Next is adding the side projection. That is a little harder, because you need to keep areas like the side of the nose instead of simply clearing a stripe down the middle. Also check whether or not you like the eyes and mouth corners from just straight on or from the side projection. Generally they should be better from the front, but pays to check.

Progress 2:



Ok, that's a bit more of a dramatic difference. In this progress shot, you will notice the difference in lighter and darker areas a lot more from where the sides of the front photo was in shadow meet the side projection in full light and focus.

More or less just leave it for now. Soften the edges a bit between them, but we have to remove all traces of shadow once the map is all pieced together anyway. Doing it once at the end, we hopefully distort as little detail as possible. Because the main darker lines are in the area where a three-quarter projection would be, you could try one of those to help get rid of the lines, but chances are you will get two more lines either side of the three-quarter projection... Getting away with as few projections as possible without losing quality/detail will make the whole thing easier to keep consistent. Something to note before we progress though – symmetrical small pink blemishes right in the centre of each cheek. Although you can only see one in the following test render, one or both will need to be removed as symmetrical small details can shoot a huge hole in your hard fought effort to capture realism.

Quick Render (texture only, no lighting):



Aside from the dark parts, I'm quite happy with the progress overall. The main thing you are looking for at this stage is details matching geometry, not overall colour shades. The light line between her lips isn't a misprojection, (there is my token invented CG term for the tutorial, deal with it... 🤖) her lips aren't fully closed and the back of the face is showing through. I was (fortunately?) sloppy enough to give an example however – the lower half of the ear probably needs re-doing. The top half looks alright though at a quick glance.

We can't simply leave that hair and the earrings plastered on for a realistic finish. If you were texturing for a game however, (aside from baking much smaller textures) then you may actually want to leave these. Some of you might go to the extent of removing the eyebrows to replace them with particle ones. I won't go to that extent here, but the following step is more or less the same.

The next step to create the bald head and rather than just copy and pasting a small chunk of cheek texture over and over, we'll use another section of the body that has roughly the same sort of look (mainly how obvious the pores are, making the cheek a bad choice to start with). In our case, the reference photos for the side of the neck have the earrings in the way, the back of the neck is covered in hair (as you will probably find to some degree with most human reference.) Good examples here would be a shoulder or a backside. Cheeky I know, but it is a large patch of skin of the same tone. We are effectively skin grafting

like a doctor. Looking at the references we have, the shoulders are either a bit blurry or cropped. This means we need another photo for texturing. I did warn you to get familiar with the idea of going back...

To save you a bit of time, I have looked at the rest of the images in this photo set and the shoulder photos aren't high enough in resolution and her hip had some scratches. What I have used is a hip image from another model on the same site. If you have a 3D.sk account, you can just search straight out for bald heads and you will probably find plenty of images to work with.

Now I did get this reference image from a set containing a lot of (tasteful) nudity. Now whether it be age or whatever personal reason, I have cropped the image to just get a section of usable skin so you don't have to sift through. If you have no objection to nudity, obviously just pick whatever images you like that suit your project.



(thumbnail, right click to save full image.)

The baldness map will need several projections, namely side, front, back and probably top. If your character will have particle hair, don't worry too much on the seams that will be under the hair. Just remember to tidy up the back of the neck because that will probably show through. Depending on how close your bald/skin image matches the rest of your image set, it might pay to tweak the colour and brightness (etc) of the baldness image beforehand to avoid doing as much colour correction and blending afterwards.

This is my progress thus far.



Still a way to go before it is usable.

Part5: Blending the projections together – B

This next step is a manual one and will take some practice to get used to. How easy it is also depends greatly on how well all your photos matched in lighting and skin tone from the start and how neat you were with your projections. People using professional photo shoot reference will have a huge advantage here – but obviously that is part of what you pay for, either by purchasing quality images or taking the time and money to do a quality shoot yourself.

If you have done some photo manipulation before, then you can probably guess where I am going with this and may have your own way of tackling it. Go for it.

Basically we need to match all of the 'patches' into one continuous map. To be perfectly honest, my map doesn't look great. In hindsight, it may have helped to correct the colours of my projections a little before projecting them, remove a shadow or two along the way. Leaving them to the end however, we do the work once and we know that it's over. Once you've practised this whole projection method a few times, you'll pick up where you would rather start correcting the colours and removing shadows. We also need to remove specular highlights (for the sake of this tutorial - shiny bits on the skin), but probably wise to do that last, get all the overall skin working first.

Once again, everyone's map will look different at this point even if you are using the same reference images I am. So what I'm going to do is give a heap of suggestions and examples rather than a step by

step 'fix it' method. (It's generally better to learn the skills behind set steps anyway, that knowledge becomes more flexible and hence useful in the long run.)

The best tools to help you here are:

Eyedropper/Colour picker

Standard Paintbrush – but with **different blending modes**. Standard/Normal blending mode will just paint over and blur out your detail and make the whole technique a waste of time. Other blending modes like lighten, darken, saturation, hue, multiply and screen are probably the most useful as I will demonstrate later.

I find that matching the saturation as a first step should help you. We will look at matching that really obvious point where the bottom of the left cheek meets the neck. Different colours and brightness levels all over the place. Not my proudest moment in texturing, but there it is – lets fix it. (NOTE: I chose this section of skin because it is meant to be pretty much the same. Areas like the eyes nose have different hues/colours **for a reason**. Knock out the differences there and there was no point using photos to begin with!)



Use the eyedropper tool with the max 5px sampling average in Photoshop if that's what you are using. Gimp however I found has a much more flexible range, I used a 20px sampling, (depending on the area) to get a genuine average for that area of texture.



Next step is to match the hue up a bit better. Note that matching the hue and saturation is not a matter of making the entire image the same hue and saturation! Obviously you will knock out the lip colour for starters, but the skin has various differences in colour all over depending on which parts have seen the most sun damage, how close the veins are to the surface – you get the idea. Just make the overall large 'patches' match up, not the individual pixels in the patches. Hopefully that makes sense.



Subtle. Depending on your monitor you may not see a difference, meaning that the projections were very close in hue anyway. Now that the hue and saturation matches up in that general area, it should be easier to see just how much you need to affect the brightness of each patch rather than doing it all at once. Sometimes two patches can be the same hue and brightness, but one looks brighter because there is so much more saturation of that colour. This is where you would use filters like screen, multiply, etc.


Whatever works best for you. (By the way, if you are new to the terms 'hue', 'saturation' and how they differ from just being called 'colour' – congratulations for making it this far! Also, it might pay to pick up a book or google search a bit of colour theory or just the terms to get a better idea. May seem boring, but it will give you more power to create so it's worth it.)



Much better! See how the colour isn't a flat uniform slab, but the overall image blends much better in that section. Now, being careful to not destroy details like pores, scars and the like, repeat that sort of process over the entire image. (Watch out for the seams/edges as well!)

Yes, I know. Time consuming.

My texture is now looking like this:



Part6: The eyes.

Now you might want to leave the eyelashes, but if you have come this far it would be far better to remove the textured ones and then model/use particles for the eyelashes. Many 3D.sk and other professional photoshoots will have close up shots of the eyes closed. The general method then is to make a shape key of the eyes closed (or make a duplicate of your model and just model the eyes shut). Then make a new projection of the closed eyes and paint out with the clone tool and whatever helps to remove any remaining traces of eyelashes.

In this case we get what we paid for. The photo set I've been using does have eyes shut images, but not at great resolution, and there is makeup all over them that make them effectively useless. After searching my usual channels for Creative Commons / free images, I couldn't find much that fit this project and this specific need. Time to put my camera to use. Feel free to use this image, it is 1848x2386 pixels and clear, which even though the lighting isn't professional, is plenty of detail. (Yes, I used the eyes open on in the end.) The other image is of similar detail, but the bottom lid is quite obscured by lashes. Or yet again, try taking your own photo if you like. For reference, mine is a 6 megapixel camera with 40cm minimum focal distance. Also of note, I managed to fluke a pretty clear shot of the iris (coloured bit) of my eye. Study it, use it, whatever.



(thumbnails)



What I did here was render out a full (2000x2000, like with all my others for projection) front image of the model with the current version of the texture - emit 1.0, reflect 0.0 - so just the texture was showing without any shading. I then took that image into Photoshop (whilst I did road test everything and do the bulk in Gimp/Linux, it was just quicker for me to finish in Photoshop/Windows because of familiarity. It will come for Linux as I continue to make the switch I'm sure...) and brought in the eye image as a new layer. I then resized it to roughly the right size, used Liquify (remember - iWarp for Gimp) to match it up better.

I then used the clone tool to remove the eyelashes. I then changed the overall hue of the image slightly to match the rest of the skin. After that I began matching the edges of the eye image to the area surrounding it using the processes we went through in Part 5. Edges only though! You want the changes in colour around the eye the most as the eye is the part that we criticise the most for realism.

One extra thing to note and I'm glad it came up - Specular highlights. The shine on the pores adds detail yes, but you want to remove that otherwise the 'shine' will render out under every lighting condition, which you don't want. Use a standard paintbrush with an eyedropper sampled colour from around that part of the eye - Darken Blending Mode - and opacity on about 50%. We want to dull it down and change to fleshy colour not white, but if we paint it out completely then we can't get the detail for the bump and spec maps.

Project this image into a new 4k (or 2k) texture.

Add that projected image as a new layer over your full texture and erase out anything that isn't around the eye area and that should be it. You already blended the eyes in before that last projection.

This is a 512px version of my final colour map. (Note: I also took the chance to adjust the eyebrows a little when reprojecting the eyes)



Part7: Generating the other texture maps

Yes, that's right. One texture won't cut it. The good news is we are done with projections and the bulk of painting.

Bump Map - Go to the area in Gimp/PS that contains image channels. Separate out the Blue channel. This one should contain the highest contrast to use for the bump. Invert the image. Convert back to rgb 8bit (or greyscale, doesn't really matter). The only trick left is to re-invert the eyebrows, because they will 'dent' the wrong way. Paint over the white hairs with black or however you want to go about that. You may need to play with 'Levels' to get a bit more contrast. Done.



The skin is split into layers in real life, now we can simulate that in Blender with Nodes and Textures. We need to get an Epidermal and a Subdermal map out of the colour map. Fortunately that is even easier than the Bump Map.

Epidermal Map - Drop the saturation of the map by 50%. Done. The idea is that the blood flows in the Subdermal layer underneath. (When you have less blood flowing around or this layer gets thicker (sickness, age, etc) you become pale.)



Subdermal Map - Figured it out yet? Clever you. We are basically going to max out the saturation on this one, making sure there is plenty of red. Take the colour map and turn up the saturation, but not completely full, or else we might lose some colour depth/detail. I then had to darken my map a little as it was quite bright.



Specular Map - This one is a little different, but not hard as I've set most of it up in the .blend file. Basically we put on a shiny blank material to the head model, surround it in even(ish) lighting and bake the map out. This technique works in Blender, but works better/more automatically in Maya, which is where I originally saw it in action. The trick here is I baked one image as you see setup in the .blend file, then moved the hemi light from one side to the other and baked another. Then I merged the two bakes together using screen and dropped the overall brightness of the merged image way down.



Part8: Material and Render Settings

Time to plug in all the maps and get some results! Finally! I was working on my own node setup that was doing alright, but then I found Pixelvore's Skin Node setup at

<http://blenderartists.org/forum/showthread.php?t=121552>

I then removed any textures that were in there already and plugged in mine, modifying a few numbers here and there. (See inside .Blend for values).

The final render was done using a HDRI as the base lighting, then one area light on the side as a fill and a spotlight to bring out the specular highlights a little more. (Note - the specific build of Blender I used was a patched one at the top of page 13 of the above mentioned thread that has support for faster Image Based Lighting. Can be rendered using any version of Blender 2.45 or higher though, best to use 2.46 since it is now officially out (at time of writng).

Warning - expect long render times, this isn't some quick cheat after all...



And there is the Final Render. Something about the eyeballs is off, but this tutorial is about the skin. The eyelashes are just modelled polys, can be substituted pretty easily with particle strands. Something of note - this image is rendered using 2k textures that were sampled down from the original 4k textures. The main reason for this is the bump map was getting cancelled out because of the large amount of oversampling for a texture that was far bigger than needed for this size render. But that gives you some idea about the detail that 4k can be used for - I could basically render out a poster for print or a closeup for a HD film and not be afraid about blurriness.

Thankyous / Reference:

Much of the information that wasn't trial and error was gained from the book **Essence: The Face** by Ballistics Books. <http://www.ballisticpublishing.com/books/essence/face/> I have both the DVD and the book and I recommend them both. Written around Maya, but the book is designed to be as non-software specific as possible and focusses on concepts. It's simply that they had to pick a program to demonstrate with...

Pixelvore for sharing his node setup with the masses.

jrs100000 for pointing out that projection in Blender was quite easy. I was using Zbrush for texturing with this method simply because I didn't think about trying it in Blender before.

And of course, the BlenderArtists community at large! Your comments, critique and support are a large part of how I got to this level of 3D art. As I learn more, I hope to impart that knowledge back to you guys and gals in gratitude.