


## Retinotopy tutorial

Phase-encoded mapping is quite simple really. Surprisingly though, to my knowledge there is no basic description of this method to be found anywhere on the internet (or at least there wasn't when I first wrote this), which is why I put up this tutorial. I also included my Matlab scripts for the stimulus presentation and analysis.

In order to know how to delineate the retinotopic areas a solid understanding of how these areas are organised is required. The point of this tutorial is to explain phase-encoded analysis. However, put briefly, the borders between the early visual areas coincide with the representations of the vertical and horizontal meridians in the visual field. Neighbouring areas contain mirror-reversed maps of visual space. Thus, the meridians usually appear as mirror reversals in the polar angle map.

Moreover, at least the earliest retinotopic regions split lower and upper visual field quadrants into dorsal and ventral regions. That means the ventral portions of V1, V2 and V3 contain maps of the upper visual field quadrant in the contralateral hemifield, while the lower visual field quadrants are represented in the dorsal portions of these areas. The ventral and dorsal halves of V1 are adjacent but the quadrant maps of V2 and V3 are spatially separated. This is illustrated in this schematic drawing:




The foveal representation is difficult to resolve with fMRI. Moreover, there *must* be some kind of gap here because we usually have a fixation dot which remains constant during our scan and thus causes no retinotopic response. Thus, the fovea will appear as non-responsive (this is implied here by the grey oval shapes) and it looks as if V2 and V3 are split into two completely separate regions.


Phase-encoded retinotopic mapping of visual areas in the human brain was described in the classic study by [Sereno \*et al.\*, 1995](#) (among others). Since then it has been used very widely, not only in myriad studies delineating retinotopic areas, but to map all sorts of stimulus attributes. For example, it has been used in optical imaging studies to map orientation columns in V1.

It is based on the assumption that when a stimulus changes cyclically along a particular dimension, the signal of neuronal populations selective for this stimulus dimension will also be cyclically

modulated. This is best illustrated with the example of retinotopy. Assume we have a rotating wedge stimulus like this:




Let's imagine that this wedge cycles all the way around the clock for 5 times during the experiment. Since we know that the early visual areas (and many later ones too...) are retinotopically organised we know that neurons (and thus voxels) with receptive fields in a particular location in space will respond only when the stimulus passes through that location and the response will decay away when the stimulus has moved on. Of course, such a time series is essentially a sinusoid, as in the example below. The green curve is a sinusoid, the grey jagged line is the same sinusoid to which normally distributed noise was added. This line could represent the time series of a particular voxel, in particular it is the time series we would predict for the wedge stimulus shown above, provided that the wedge started cycling from the 3 o'clock position.



Now, if we want to know which retinotopic location this voxel represents or – in other words – at what times in the experimental run the stimulus passed the receptive field of this voxel, all we need to do is work out the phase shift of this sinusoid.

For example, for a voxel representing the location where the stimulus began cycling (3 o'clock) the time series would be a cosine (i.e. starting at a high response). So in theory, we could find out the phase of the underlying sinusoid by creating cosine waves with every possible phase shift and calculating for each wave the correlation coefficient between the wave and the actually observed time series. The phase shift at which we achieve the strongest correlation should then be the one we are after.

This method of works fine and in fact this is pretty much how some analysis packages do it. For the example above, the resulting phase shift comes out at 44 degrees. Since the true phase should be 45 degrees (angle anticlockwise from 3 o'clock) this is a very good match. However, correlating numerous sinusoids is inefficient and also completely unnecessary. Instead, we can take advantage of trigonometry. We generate only two sinusoids as in the graph below, one is a cosine (blue), and the other a sine (red) wave:



We then calculate the correlation coefficient for each of these sinusoids with the observed time series. These coefficients then inform us about the phase of the true underlying sinusoid. As you can guess from the graph, both the blue and the red curve are slightly shifted with respect to the true sinusoid. In fact they will both show the same moderate positive correlation. We can regard these correlation coefficients as the X and Y components of a vector in Cartesian space. If both X and Y are positive and of equal size, the angle of the vector is 45 degree. We can convert these Cartesian coordinates into polar angle by calculating:

$$phase = atan2(Y,X) \text{ (in radians)}$$

Again, the phase in this case came out near 44 degrees, so this method works just as well as the multiple correlations we tried first. However, while this method provides a great estimate for the phase shift of the sinusoid, because the correlation coefficients are in the range from -1 to +1 they do not give us an estimate of the signal *amplitude*. In order to do this, we can employ two strategies. One way would be to conduct a conventional analysis using the general linear model which is typically done in brain imaging studies. We use the cosine and the sine waves we created as regressors in the design matrix. The resulting betas the model fits to the observed data are essentially comparable to the correlation coefficients we calculated above (which we will again call X and Y). However, unlike correlation coefficients they reflect not only the phase but also the amplitude of the sinusoidal modulation (in SPM the betas would be in units of percent global signal change). The signal amplitude can then simply be calculated as the magnitude of the vector described by the two betas:

$$amplitude = \sqrt{X^2 + Y^2}$$

One nice advantage of this method is that since we are using the general linear model, we can add further regressors just as we would in other neuroimaging studies. We can for example model the parameters from motion correction in order to compensate for motion artifacts.


However, most studies use another method. We can also calculate a fast Fourier transform of the observed time series. This performs essentially the same calculation but it does it not only for the fundamental frequency of the stimulus cycle (which was 5 cycles/run in this example), but for the entire frequency spectrum. The transform returns a vector of complex numbers for all the frequencies in the spectrum. The real and imaginary components are equivalents of the two correlation coefficients (or betas) we got in the other analyses (however, because we are interested in the phase of a cosine wave we need to multiply the imaginary component with -1). Then we can use these two components to calculate the phase angle just as we did above. In order to threshold

the data, what people sometimes do is to first divide the power (amplitude) at the fundamental frequency by the power at all other frequencies. This is in essence an F-ratio which can be used to select only significantly active voxels.

The results of these two methods, using the general linear model or the fast Fourier transform, are pretty similar. So in the end it probably comes down to personal taste which one you want to do. One step one might consider when using the general linear model, is to convolve the sinusoidal regressors with the haemodynamic response function. This relates to one aspect of this analysis we haven't addressed yet: the sinusoidal functions assume an instantaneous response to visual stimulation, but we of course know that at least in standard fMRI studies the signal lags seconds behind the neuronal activity that caused it. Using sinusoid regressors convolved with the haemodynamic response function may take care of this problem.

However, another way to address this question is to run two experiments: one in which the stimulus rotates in an anticlockwise direction (as in the example above) and one in which it rotates clockwise. We can then average together the X and Y components calculated by each (after multiplying the Y component from the clockwise run with -1) and then use these average components for calculating the phase. The haemodynamic lags for the two directions should cancel each other out, giving rise to a more exact estimate of the true phase.

But before going to these lengths, you need to ask yourself if this information is actually strictly necessary. In most studies, retinotopic mapping is done in order to delineate the borders of the early visual areas. This can be identified by mapping the phase reversals where upside-down maps turn into upside-up maps. For example, the border between V1 and V2 coincides with the vertical meridian. You can see this in this map:

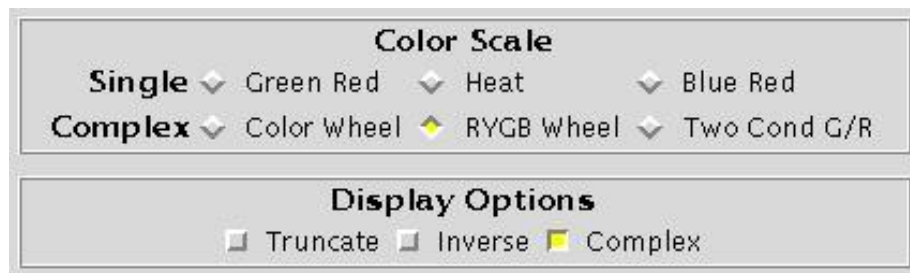


The dorsal borders are quite red, which is the colour associated with the vertical meridian. However, the ventral border only seems to go to blue, or possibly purple. Nonetheless, the mirror reversals in the phase map are immediately clear. Basically, the travelling waves of response in the two areas meet at the meridian. Therefore, an exact knowledge of the phase angle is not actually required for mapping the areal boundaries. But of course, if your experiment depends on knowing precisely where

each polar angle in the visual field is represented in the cortex, you will need to take the necessary steps to make the mapping exact.

Finally, you need to map the borders to delineate the areas. This is generally done by projecting the voxels onto a 3D-reconstruction of the cortical surface and then defining the areas described by the meridians and the outer and inner eccentricities of the stimulus. However, other methods exist as well which can map the areas in the actual volume space ([Dumoulin et al., 2003](#)).

I usually do my mapping on the reconstructed surface in the excellent [FreeSurfer](#). Here is what you need to do to display the maps in that program: You first load the overlays for the real and the imaginary components generated by the Fourier transform into overlays 1 and 2, respectively (obviously, with the appropriate coregistration to the structural brain scan). Then you go into the *View->Configure->Overlay* dialogue, and click the following radio buttons:



Actually, you could also click on the normal colour wheel instead of the RYGB one, but I prefer the latter. The button 'Complex' tells FreeSurfer to interpret the two overlays you loaded as real and imaginary components of the Fourier transform. It then calculates the phase angle and displays them in a colour wheel. Unfortunately, there is no way in FreeSurfer to display the colour wheel in the legend. I made a function for that purpose, which you can download below. In order to threshold the map, you can also load the map of the signal amplitudes into overlay 3 and adjust the threshold in the *View->Configure->Overlay* dialogue just as you would do for a t-map.

Finally, for some kinds of maps, like for instance the polar maps here, it makes sense to let the colour wheel cycle through the clock *twice*, so that all four colours repeat for each visual hemifield (as can be seen in the colour wheel in the figure). You can set this parameter in FreeSurfer in *View->Configure->Phase Encoded Data Display* by setting *Angle Cycles* to 2.

## FREQUENTLY ASKED QUESTIONS

Here is some advice on how to set up your mapping experiments to get good results. Note that these things are of course not set in stone and you may find that other parameters work for you. But these suggestions came in the form of advice from various people (most notably Marty Sereno) and from my own experiences.

*How many cycles & how long should they last?*

I aim to have around 1 minute per stimulus cycle. Between 8-10 cycles per run are usually fine for getting very decent maps of V1-V4/V3A. So altogether that gives around 8-10 minutes per run. However, in the past I have also tried this with 8 cycles of around 30 seconds, which also produced adequate results.

*What stimulus should I use?*

In order to map the earlier visual areas pretty much any stimulus will suffice. Typically, people use high contrast checkerboards that flicker at 2-8Hz. Some people use coloured checkerboards although I am told that there is not much of a difference between coloured and black-and-white stimuli. In a

recent study we used coloured checkerboards in the hope of enhancing the signals in higher ventral areas. In the past I also tried high contrast geometric shapes whose location, width and colour was randomized on every frame (akin to “continuous flash suppression” dichoptic masks). I did not notice much of an improvement over standard black-and-white checkerboards. Marty Sereno has a nice stimulus that contains not only coloured checkerboards but also white random dots moving in various directions (which are randomized over the course of the experiment) and also letters to drive shape-sensitive brain regions.

It probably doesn't matter very much what precise stimulus you choose, however. More importantly, if you are interested in mapping higher visual areas it may be advisable to have a stimulus that covers as much as possible of your field of view. The setup in most fMRI scanners does not allow visual stimulation far beyond 10-15 degrees of eccentricity – in many cases even less. Many of the higher visual areas contain neurons with very large receptive fields, however, and their maps may be biased towards the peripheral visual field (V6 would be a prime example). Such areas will simply not show a very good cyclical signal modulation with a too small stimulus.

*What behavioural task should there be?*

In retinotopic mapping we are not only mapping the response to visual stimulation but also the deployment of visual spatial attention. This is particularly important for higher visual areas that show strong attentional modulation. It may therefore be advisable not to use a fixation task but to use a task that forces subjects to covertly focus attention on the stimulus. In the past I did this by asking them to detect a small grey target that flashes briefly (200ms) on the stimulus at random times throughout the run (obviously, this must not be correlated with the cycling frequency). But other tasks are possible. For example, Marty Sereno asks subjects to monitor for the appearance of digits inside the stimulus (with letters appearing throughout the run as distractors). However, a fixation task has the benefit of stabilising eye movements and even passive viewing can produce adequate results, especially if you care more about the early visual areas.

*What about meridian mapping?*

If you paid attention it may have occurred to you that in order to map retinotopic organisation you really could simply map activations corresponding to the vertical and horizontal meridians. You could have a vertical wedge stimulus and a horizontal and alternate between them. Of course this does work – at least for early visual areas because they are large – and numerous excellent studies used this approach. It however becomes progressively more difficult for the higher regions that are so small that the activations for vertical and horizontal cannot be distinguished very easily. And even for early visual cortex it isn't actually very accurate to define regions based only on the basis of the meridians. A mirror reversal in the phase map is relatively well defined. The exact border of the activation in response to the vertical meridian is not so straightforward. It needn't be (and probably isn't) the middle. Last but not least, there is no real benefit of doing it this way. You still need sufficient data to get accurate maps which isn't all that different from what you would need for phase-encoded or even pRF (see below) approaches. So... don't do that.

*What about pRF mapping?*

Population receptive field mapping ([Dumoulin & Wandell, 2008](#)) is a new way to do retinotopic mapping analysis. It is more accurate and more informative than phase-encoded mapping (e.g. you can infer receptive field profiles for each voxel and it incorporates the haemodynamic lag). I would strongly advise using it if this extra information is useful to you (but this is a whole new tutorial!). It also comes with a large additional computational cost. So if you merely want a quick and dirty way to delineate visual areas, or if you can't create a good model of your stimulus (which is essential for pRF mapping) then phase-encoded mapping is still the ideal way to do it.

*What about a probabilistic atlas or anatomical definitions?*

A lot of people will tell you that retinotopic mapping is always essential because retinotopic organisation is so variable between people. This is not strictly true. The same argument could be made about all fMRI and structural MRI studies and nonetheless group analysis is still standard practice. Of course, the key point here is spatial normalisation to bring the brains from different individuals into alignment. The more sophisticated your normalisation procedure is, the better is your signal-to-noise ratio of your group map. But even with traditional normalisation you can probably already get a relatively decent average map. This could be used to build a probabilistic atlas of where the visual regions are.

Moreover, an increasing number of people are working on purely anatomical definitions of the visual cortex (e.g. based on myelination), or a combination of anatomical and retinotopic properties. For example, the `-label_v1` option in the automatic cortical reconstruction in FreeSurfer does a very good job at delineating V1 based solely on the cortical folding. As some have theorised, the folding morphology of visual cortex may in fact be related to the retinotopic organisation. There may also be other relationships between structure and retinotopic maps. The future may therefore bring new advances in the way we identify and parcellate cortical regions. And even now, in some situations you may be quite justified in not doing any retinotopic mapping – for instance, if you place a small region of interest in the posterior half of the calcarine sulcus you can be pretty confident that this will be within V1, somewhere on the horizontal meridian, and probably within a certain eccentricity range.

But as this example shows, there still remains information to be gained from conducting retinotopic mapping and so far this hasn't been made obsolete by recent advances. Anatomy so far can not really tell you with certainty the precise retinotopic position, receptive field sizes, and other tuning properties of a particular cortical location (but people *are* already working on changing that too!). It also remains very vague about the higher extrastriate regions, most of which are now known to be retinotopically organised but which have only fuzzy structural definitions. In any case, if you want to be sure which particular retinotopic region you are in, functional mapping remains necessary – at least for the time being.

#### *What about other functional definitions?*

Against what some may try to make you believe, it can actually be perfectly justified to omit retinotopic mapping in a visual fMRI experiment if that is consistent with your experiment. People regularly localise the motion-sensitive hMT+ complex (a.k.a. V5) using a simple functional localiser comparing moving and static images. Few people are currently arguing against this practice even though it is abundantly clear that the hMT+ complex consists of a whole party of different retinotopic (and some not-so-retinotopic) regions. This is entirely reasonable as long as you are aware of the fact that your definition is based only on motion sensitivity, and that you aren't claiming that this complex is a functionally homogeneous area. Similarly, there is nothing wrong with localising a cluster in the brain responding preferentially to coloured compared to grey-scale images – but it is difficult to justify calling this V4. Unlike for hMT+/V5 a lot more is known now about the retinotopic organisation of V4 and so to equate V4 with a ventral colour-sensitive area is strictly speaking incorrect.

But this doesn't preclude you from doing meaningful experiments comparing the colour-sensitive cluster to the face-sensitive cluster etc. People talk a lot about functional specialisation of visual cortical areas but the truth is that there will generally be a substantial overlap between the response properties of adjacent visual areas. I can't think of many experiments that showed unequivocally discrete patterns of results between dorsal V3 and V3A, for instance, and it is even rare between V1 and V2. The main differences generally arise between ventral and dorsal regions or between early and higher visual cortices. Some of this overlap in effects may be due to partial voluming (voxels at the areal borders contributing to the response in both regions) but even if you account for that, it is generally the case. It may also be in part due to the limited resolution of fMRI, which is effectively measuring blood flow, not neuronal activity. Disentangling these factors will be an important objective for future research.

## MATLAB FUNCTIONS

*Disclaimer: these are functions I used in the past and that have been successfully used by many of my colleagues. However, they are not supported. Please don't hesitate to contact me if anything is unclear or if you think they aren't working as they should. But I make no commitment that I will help you getting these to run if that requires too much of a time investment. In any case, you may want to use the latest version of Matlab. I have heard from people running into problems with too old versions. For the stimulus scripts you will require a graphics card and drivers that support transparency in Psychtoolbox. Most systems do but there are some setups with problems (such as when the stimulus appears all black). I can't help you if you run into that issue. Checking the Psychtoolbox website or mailing list may be useful in this case. Another alternative would be to change the function so that the overlap of the stimulus aperture and the carrier behind it is calculated without use of transparency. This will inevitably run more slowly though.*

[Retino\\_Fourier4D.m](#) Phase-encoded retinotopic mapping. (Requires SPM & a 4D Nifti image of the whole time series)

[Average\\_Phase\\_Maps.m](#) Script for averaging together several phase maps created by Retino\_Fourier4D. This allows you to combine runs with opposite stimulus directions.

[Colour\\_Wheel.m](#) Script for displaying the RYGB colour wheel FreeSurfer uses (Matlab help included).

[Retinotopic mapping stimulus](#) This zip file contains my mapping localiser (Requires [Psychtoolbox 3](#))