

Analyzing the face on the Shroud of Turin with a three-dimensional morphable model

Game and Media Technology Master Thesis

Utrecht University



Author

Nicola Chinellato
4220633

Supervisor

Prof. dr. Remco C. Veltkamp

Abstract

In this thesis I study the claim that the face that is visible on the Shroud of Turin contains some three-dimensional information. I do so with the aid of the computer vision technique called morphable models. I analyze how an implementation of such a model acts when given three-dimensional meshes obtained from real faces, from noise functions and finally from the Shroud face image, trying to understand whether the outputs of the model can be used to classify the inputs.

Table of Contents

1. Introduction	4
2. Related work	5
2.1. Dating the Shroud	5
2.2. Shroud image formation hypotheses	6
2.3. Three-dimensional morphable models	7
3. Objectives	9
4. Resources and materials	11
4.1. Shroud pictures	11
4.2. Three-dimensional morphable model	14
4.3. Other resources	15
5. Methodology	18
5.1. Literature research	18
5.2. Image manipulation	18
5.3. Three-dimensional mesh	20
5.4. Evaluation	23
5.5. Experiments definition	24
6. Results	26
6.1. Preliminary phase	26
6.2. Experiment 1	27
6.3. Experiment 2	28
6.4. Experiment 3	29
7. Evaluation	32
7.1. Fitting real faces	32
7.2. Fitting non-face meshes	32
7.4. Robustness of the fitting process	34
7.5. Choosing the significant parameters	34
7.6. Comparing the fitting results	37
8. Conclusions	41
8.1 Answers to the research questions	42
8.2 Future work	43
9. References	44
A. Appendix	47

1. Introduction




Figure 1.1: The Shroud of Turin. The darker lines are burn marks from a fire erupted in the chapel where the Shroud was kept. The body image is fainter, but still quite clear, in between the burn marks. It is possible to see both the front and the back of the body since the body was laid on the Shroud, which was then folded on top of it.

The Shroud of Turin (Figure 1.1) is the linen cloth that for many years has been regarded as the one in which Jesus was buried. As such it has always been considered of great value and, in more recent years, it caught the attention of the scientific community. This interest started at the very end of the nineteenth century and is still vivid today. The image of a body is impressed on the Shroud and, despite the numerous studies that have been made, its origin is still debated. Some of the properties of this image are particularly interesting from a scientific perspective. In this thesis we will analyze the claim that it is possible to obtain some three-dimensional information from the part of the image containing the face and we will do so with the aid of the computer vision technique known as morphable models. Our objectives are to test if it is possible to create a three-dimensional mesh with the information that the image is supposed to carry and use the properties of the morphable models to decide whether such mesh can be classified as an actual face. The approach we take is to study the theories about the Shroud image and a particular implementation of a morphable model and then to use the morphable model fitting function to process the three-dimensional meshes obtained from the Shroud, comparing the results to ones obtained with meshes representing actual faces. After this introduction, the text will continue with the description of what has been achieved so far researching both the Shroud of Turin and the morphable models (Section 2), followed by a more detailed explanation of what the objectives of this study are and which research questions I am trying to answer (Section 3). Section 4 contains a description of the materials that I had at my disposal and that I used, in particular the Shroud pictures and the morphable model implementation. After a description of how I planned my experiments and their definition (Section 5), in Section 6 I will explain how I implemented such experiments in practice and present the results obtained from each. The evaluation of those results is presented in Section 7. Finally Section 8 contains the conclusions and a discussion on possible future improvements to my work. At the end of the document there are a list of references (Section 9) and an appendix with the images of the results and the data obtained from them (Section A).

2. Related work

The Shroud of Turin has become one of the most studied archaeological artifacts since Secondo Pia took a picture of it in 1898 and noticed that its photographic negative showed very clearly the shape of the face (Figure 2.1) and the rest of the body.

Pia's findings arose the scientific community's interest and led to the formation of the Shroud of Turin Research Project (STURP), a group of people, scholars or not, interested in a scientific analysis of the Shroud. The group was granted access to the relic in 1978 and performed a wide variety of studies of which the results are summarized in [Schwalbe 1982].

The STURP's analysis led to two major debates that contributed to keep alive the interest in the item until today: one on the Shroud's age and one on its image formation.




Figure 2.1: a comparison of the detail of the face on the Shroud and its photographic negative (this is the image referenced as 330px-Shroud_positive_negative_compare.jpg in the list of Section 4.1).

2.1. Dating the Shroud

The STURP scientists were allowed to collect a single piece of linen from a corner of the Shroud, which was then cut in two: one piece was reserved as a backup and the other was furtherly cut in three and split between three laboratories. The laboratories had the task of independently performing carbon dating on the pieces assigned to them. The result of these analyses, published in [Hedgesi 1989] collocates the Shroud's origin between AD 1260 and 1390 with 95% confidence.

This result arrived after a big debate on the test's methodology, with some scholars challenging the validity of the results even before they were published [Meacham 1987], a debate that continued after the results were published (see for example [Rogers 2005] and [Freer-Waters 2010]). The main argument concerns the pieces of linen on which the carbon dating was performed: since they came all from the same thread, cut from the very angle of the linen cloth, they might be contaminated and not representative of the whole item.

Recently the skeptics of the carbon dating have been trying some different dating methods (e.g. [Fanti 2015]).

The dating issue, however, is not much relevant for my project so I will not analyze it any further. A recent compendium of all the scientific studies, dating included, can be found in [Siefker 2015].

2.2. Shroud image formation hypotheses

The STURP analyses were not able to identify an image formation explanation that would satisfy all the members of the team, so the question of how the image was formed stayed open, as it still remains. Throughout the years, a lot of theories have emerged on this matter and they have been collected in [Fanti 2011], where the author gives an overview, followed by his opinion, of all the theories concerning the Shroud image formation. He proceeds to group the theories thusly:

- **Contact:** these theories see the image as formed through contact between the body and the linen cloth and mostly derive by [Volckringer 1991]. In [Fanti 2004], though, the authors deny their validity by showing that the Shroud presents color even in places where it would not have touched the body.
- **Gas diffusion:** these theories see the image as formed by some gases, mostly created by the body's processes of decomposition, which were absorbed by the cloth. The main argument for this thesis was published in [Rogers 2008], however in [Fanti 2011] it is argued that this theory fails to explain most of the characteristics studied on the Shroud.
- **Radiation:** these theories see the image as a result of some type of radiation, emitted by the body or by the surrounding environment, which interacted with the cloth changing its color. This is the most accredited theory between those who do not trust the carbon dating. Being the most accredited it is also the most studied. Numerous theories emerged regarding the type of radiation that could have caused the image to form. The theories range from ultraviolet radiation [Baldacchini 2008] to a radiation emitted by the body enveloped in the Shroud [Rinaudo 1998], however the most followed one currently is the corona discharge theory [Fanti 2005] [Fazio 2015].
- **Artistic forgery:** these theories see the image as a forgery, produced by some artist in medieval times. These theories are supported mostly by those who regard the carbon dating as trustworthy and are backed by an analysis that reportedly found traces of iron oxide and other pigments which would be evidence of tempera paint [McCrone 1997]. The detractors of this theory argue that nobody has been able to identify a technique that would produce physically similar results.
- **Mixed:** these theories combine one or more of the ones above. For example in [Fazio 2014] the authors analyze the stains and the body image on the Shroud separately and assert that they were generated in different times. In particular the formers are actual blood stains, result of the cloth touching a dead body and the latters are the result of some energy released by some chemical reaction in

the body throughout time. They conclude that the blood is proof of the Shroud being an actual burial cloth (not necessarily the one of Christ though). These theories are often dismissed as not very scientific by other scholars (e.g. [Latendresse 2014]).

Some other theories, not grouped in the categories above and with less consideration by the scientific community, include some type of proto-photography [Allen 1995]. The debate on the image formation has led to the experimentation on techniques to reproduce the Shroud as similar as possible to the original. A wide variety of techniques has been tried: in [Craig 1994] a dust transfer technique with materials available in medieval times is used to create an image of a face with similar characteristics to the Shroud, in [De Liso 2010] the author uses seismic activity to create an image with properties resembling the Shroud's ones with a phenomenon similar to the corona discharge; some people have tried different bas relief techniques (e.g. [Delfino 2000] and [Nickell 1987]). In particular in [Garlaschelli 2010] a convincing replica is created using a combination of bas relief and frottage; in [Di Lazzaro 2012] the authors use a laser to color a piece of linen cloth and give it physical properties very close to the one of the Shroud, however they do not produce an image, but just a spot. Although a laser would not be available in the 14th century, the authors argue that some kind of radiation could have had the same effect.

2.3. Three-dimensional morphable models




Figure 2.2: an example of how a morphable model works, taken from [Blaiz 1999]. In the paper the average face is subdivided in different segments that can be independently modeled in order to increase the degrees of freedom of the model. The picture shows how the resulting face varies by increasing or decreasing the deviation from the average of each segment.

Morphable models are a computer vision technique that has been introduced in 1999 to provide a framework for easily synthesizing new three-dimensional faces that would look realistic [Blaiz 1999]. The model was built by recording the shape of $m = 200$ real faces (half males and half females), each face being described by its shape and its texture, represented by the two vectors:

$$S_i = (X_1, Y_1, Z_1, X_2, Y_2, Z_2, \dots, X_n, Y_n, Z_n) \text{ and } T_i = (R_1, G_1, B_1, R_2, G_2, B_2, \dots, R_n, G_n, B_n)$$

This allows new faces to be defined in barycentric coordinates as a linear combination of shapes and textures:

$$S_{mod} = \sum_{i=1}^m a_i S_i \text{ and } T_{mod} = \sum_{i=1}^m b_i T_i$$

After computing the average shape (\underline{S}) and texture (\underline{T}), principal component analysis was then used to perform a basis transformation to an orthogonal coordinate system formed by the eigenvectors s_i and t_i , so that a new face could be expressed as:

$$S_{mod} = \underline{S} + \sum_{i=1}^m \alpha_i s_i \text{ and } T_{mod} = \underline{T} + \sum_{i=1}^m \beta_i t_i$$

With this solution each new face is defined just by the shape and texture coefficients α_i and β_i which describe how much the new face is distant from the mean face, which has all the coefficients equal to zero (Figure 2.2).

In the paper the authors proposed morphable models as a way of synthesizing new faces, but they also discussed their possible utility in face recognition tasks. In fact their properties and in particular their capability of abstracting from pose and illumination render them quite versatile for many computer vision applications. In particular morphable models have been widely used for three-dimensional reconstruction, both from two-dimensional and three-dimensional sources. A survey on the first approach can be read in [Abate 2007] and for the second approach in [Bowyer 2006]. In [Doshi 2013] both approaches were studied for four different morphable model implementations.

In this project I used the 3D to 3D approach with the aid of a morphable model developed at Utrecht University [ter Haar 2009], which I will describe more in depth in Section 4.2.

3. Objectives

The objective of this thesis is to understand whether it is possible to use morphable models to test the claim that three-dimensional information can be extracted from the Shroud of Turin's face image. In order to achieve this broad objective it is easier to divide it into smaller ones, which I have done as follows.

The first thing to understand is if between all the theories that try to explain how the Shroud image formed and thus how the three-dimensional information was encoded in it, there is one (or more) that provides a mathematical way of extracting such information. This translates into finding an interpretation of the image that can lead to the construction of a three-dimensional mesh.

The second thing would be to test if a morphable model can be consistently fitted to the three-dimensional mesh obtained from the Shroud's face image and what can be said about the possible outcome of such fitting.

The third and final thing would be to understand if there is a way to measure, using the results from the fitting, how likely it is for the three-dimensional mesh to be the representation of an actual face. In other words this final step is about using the results from the fitting of a morphable model to distinguish between three-dimensional meshes that represent faces and three-dimensional meshes that do not.

These objectives can be summarized in the following research questions:

- **Is there a plausible interpretation of the Shroud of Turin's image formation that leads to the construction of a three-dimensional mesh?**
To answer this question, after some literature research (of which the outcome has already been described in Section 2), I selected the one theory about the Shroud's image formation that have not been dismissed and that is not just outlined, but gives some mathematical explanation. I then implemented that interpretation to produce a three-dimensional mesh.
- **Can a morphable model be consistently fitted to the three-dimensional mesh obtained from the Shroud of Turin's image?**
To answer this question, after making sure that the previous one had a positive outcome, I used the morphable model implementation at my disposal (described in the following section of this document), fitting it to the three-dimensional meshes previously obtained.
- **Is it possible to use the results of the morphable model fitting to distinguish between meshes that represent faces and meshes that do not?**
To answer this question I compared the results of the morphable model fitting to meshes that represent faces and meshes that do not. I did this by analyzing the fittings parameters distributions of both categories.
- **Can the meshes obtained from the Shroud be classified as faces after having analyzed the results of the morphable model fitting to them?**
To answer this question I compared the results of the morphable model fitting to the meshes obtained from the Shroud with both the results of the fitting to face meshes and to non-face meshes.

Below each question I summarized the method that I used to answer them, Section 5 will go more in depth on this subject, describing what was my plan and how I planned to carry it out.

4. Resources and materials

4.1. Shroud pictures

What follows is the list of the pictures that I had available. The pictures depict various parts of the Shroud and some even its entirety, however for this project we were interested only in the head, so as a preliminary step I discarded some pictures and cropped some others in order to keep just the head. Following all the pictures that were at my disposal are listed with their original filename and resolution, a description of their content, whether I did some changes to them and finally the final name given to them. All the pictures are grayscale images and for all the used pictures I initially produced a high resolution version, with the original resolution, and a low resolution one, where the height has been reduced to 1000 pixels at maximum. Later, as I will explain, I did some more selection and uniformed the selected pictures to just one resolution, but initially I thought it would be handy to have a lighter version of the pictures to be used for quick tests and similar cases where I would prefer to have a quick result rather than a precise one. I also converted the images that were in the TIFF file format to JPEG since the extra information contained in that specific format was not particularly useful for my purposes. JPEG compression alone allowed to shrink down the files to a quarter of their size.

The pictures came from three different sets, depending on how and by whom they were obtained, and are grouped accordingly in the following subsections.

Enrie images




Figure 4.1: In order Enrie_01.jpg, Enrie_02.jpg, Enrie_03.jpg and Enrie_04.jpg

These pictures are a contact copy of the original negatives, taken in 1931 by Giuseppe Enrie and later digitized:

- **05.tif** (7903x9959): positive of the full picture of the Shroud. This picture has been cropped to contain just the head. Final name: *Enrie_01.jpg*;
- **06.tif** (7927x9897): negative of the full picture of the Shroud. This picture has been cropped to contain just the head. Final name: *Enrie_02.jpg*;

- **07.tif** (7059x9322): negative of the center part of the Shroud, depicting the head and shoulder from the front and the back. This picture has been cropped to contain just the head and rotated 180°. Final name: *Enrie_03.jpg*;
- **08.tif** (6671x9541): negative of the bottom part of the Shroud, depicting the legs and torso from the front. This picture has been discarded since it didn't contain the head;
- **09.tif** (7897x10021): positive of a detail of the Shroud, depicting the whole frontal image, the head and feet are partially cut off. This picture has been discarded since it didn't fully contain the head;
- **10.tif** (7668x9905): positive of a detail of the Shroud, depicting the head from the front. This picture has been cropped to remove the space around the head that was included in it. Final name: *Enrie_04.jpg*;

DHL images




Figure 4.2: In order DHL_01.jpg, DHL_02.jpg and DHL_03.jpg

These are pictures of the face that were used by the Dutch Holographic Laboratory¹ (hence the name DHL) to produce an hologram of the Shroud²:

- **DHL_face_final.tif** (3904x5205): the original picture. This picture has been kept as is. Final name: *DHL_01.jpg*;
- **DHL_face_final_kwart_res.tif** (1952x2603): the same picture as above, but scaled down by a factor of 2. This picture has been kept as is. Final name: *DHL_02.jpg*;
- **DHL_face_final_kwart_res_levels.tif** (1952x2603): the same scaled down picture, but digitally adjusted in order to enhance the contrast. This picture has been kept as is. Final name: *DHL_03.jpg*.

¹ <http://www.holoprint.com/>

² <http://shroud3d.com/>

Miller images




Figure 4.3: In order Miller_01.jpg, Miller_02.jpg, Miller_03.jpg and Miller_04.jpg

These are ultraviolet pictures taken by Vernon Miller, member of the Shroud of Turin Research Project:

- **Untitled-4.jpg** (9282x13923): negative of the face. The bottom part of the picture seems damaged; The damage seems due to the file being corrupted, the picture has been discarded since that made it unopenable;
- **Untitled-5.tif** (7800x10046): detail of the face. This picture has been discarded since it contained only a detail;
- **Untitled-13.tif** (4206x9309): full picture of the Shroud. In here the image of the body is quite outstanding in respect to the other pictures and it is clear that it has undergone some sort of manipulation or enhancement. This picture has been cropped to contain just the head. Final name: *Miller_01.jpg*;
- **Untitled-14.tif** (8640x10800): positive of the face and the upper torso. This picture has been cropped to contain just the head. Final name: *Miller_02.jpg*;
- **Untitled-15.tif** (8400x11594): very bright positive of the middle part of the Shroud, depicting the head and upper torso from the front and the back. This picture has been cropped to contain just the head. Final name: *Miller_03.jpg*;
- **UV Photo by Vern Miller-19.tif** (7300x7125): positive of the head and full torso from the front. This picture has been cropped to contain just the head. Final name: *Miller_04.jpg*.

Other images



Figure 4.4: Other_01.jpg

- **330px-Shroud_negative.png** (169x272): low resolution negative depicting the face. This picture has been kept as is and converted to the JPEG format for conformity with the other pictures. Final name: *Other_01.jpg*;
- **330px-Shroud_positive_negative_compare.jpg** (330x272): a side by side comparison of the picture above with the positive of the same part. This picture has been discarded since it is basically a repetition of the one above, but with its positive on the side (Figure 2.1).

4.2. Three-dimensional morphable model

I had at my disposal a morphable model implementation developed by a researcher at Utrecht University as a PhD project [ter Haar 2009]. The model is built using a set of 100 three-dimensional face scans provided by the University of South Florida, that originally collected them to develop its own morphable model [Sarkar 2005]. Each face was described as a set of three-dimensional points, their mean face was calculated and then principal component analysis was used to obtain 99 eigenvectors and the respective eigenvalues from the 100 faces.

With this set up it is possible to define new faces, different from the 100 original ones, using the following formula:

$$S = \underline{S} + \sum_{i=1}^{99} \omega_i \sigma_i s_i$$

Where:

- \underline{S} is the mean face;
- ω_i is a weight;
- $\sigma_i = \lambda_i^2$ and λ_i is the i th eigenvalue;
- s_i is the i th eigenvector.

The eigenvectors and eigenvalues are defined by the model, which means that they can be considered constants while operating with the same model. A face can thus be described just by the set of 99 weight parameters (ω).

The implementation that I had, provided a large set of functionalities relative to morphable models, from model construction and initialization to simple mesh visualization, the ones that I used for this project are the following:

- *Pose normalization*: a function that aligns a three dimensional mesh with the mean face of the morphable model, using the nose tip as center point. Each three dimensional mesh has to go through this function before being able to be used as input for the fitting function.
- *Morphfit*: a function that performs the fitting of the morphable model to a three-dimensional mesh. To perform the fitting, the software starts from the mean face (where all weights are 0) and adjusts it iteratively, selecting for each of the 99 weights the one in a set of ten that gives the best result before moving to the next one. The best result is considered the one that gives the minimum distance (measured as the root mean square) between the resulting face and the input mesh. The *morphfit* function outputs, among other files, a PLY file with the three-dimensional mesh that results from the fitting and a text file containing the list of 99 weights that were used to build it.
- Multiple sets of three dimensional meshes in the PLY format, that were obtained from laser range scans of real faces and that were used in the past to test the software. Of those I used a set of 234 faces with neutral expressions.

4.3. Other resources

Helping tool

Before beginning with the experiments I created a graphic tool that would help me visualize what I was doing and that I kept improving while conducting the experiments. The tool is written in C++, uses the libraries OpenGL³ and OpenCV⁴ and is capable of:

- Loading a grayscale image and creating a three-dimensional mesh from it:
 - The mesh creation process is described in Section 5.3 and mathematically defined in Section 5.5;
 - The influence of the pixel intensity on the mesh can be changed online;
 - The mesh is scaled and rotated to the same scale and orientation of the mean face of the morphable model (the orientation is approximative since the mesh has to go through the pose normalization function of the morphable model anyway).
- Applying one or more OpenCV filters to the image and updating the mesh consequently:
 - The available filters are blur, gaussian blur, median blur, bilateral and non-local means;
 - The filters can be applied in any order;
 - The image can be reset to its original state.

³ <https://www.opengl.org/>

⁴ <http://opencv.org/>

- Applying a mask to the image so that the masked pixels will not appear in the mesh.
- Loading a cloth depth map to influence the position of the vertices in the mesh:
 - This process is described in Section 5.3 and mathematically defined in Section 5.5;
 - The influence of the cloth depth map on the mesh can be changed online.
- Visualizing the resulting mesh (see Figure 5.2 for an example visualization):
 - The camera can rotate around the mesh and zoom in and out.
- Exporting the mesh in the PLY, OBJ or CSV formats.

MeshLab

MeshLab⁵ is a free and open-source tool for visualizing and processing three-dimensional meshes. In this project I mostly used it to visualize the outputs of my tool and the results of the morphable model fittings. All the images of three-dimensional meshes that appear in this document have been produced using MeshLab.

Face scan modeling tool

This is a software, developed by a student at Utrecht University, that loads the average face of the morphable model just described in section 4.2 and provides an user interface that allows to increase and decrease each one of the weights in order to create arbitrary faces. I used this tool to inspect what effect each of the weights had on the final face while I was trying to select which weights were more relevant, as will be detailed in section 7.5.

Morphable Model Fitter

This is a graphical user interface created as a MSc project by a student at Utrecht University [van der Weerd 2011]. The software uses the Basel Face Model to perform 2D to 3D face reconstruction. The user inputs one or more pictures depicting the same person from different angles, manually positions some landmarks (e.g. nose tip, pupils, upper and lower lips and so on) and then the software uses the model to create a three-dimensional mesh as similar as possible to the one described by the inputs. I initially tried to use this tool to include 2D to 3D reconstruction in the project in addition to the 3D to 3D reconstruction that would be done with the morphable model described in section 4.2.

The initial results were not very satisfying (Figure 4.5), probably due to the fact that the software works better with more than one angle as input, but the Shroud only contains the frontal view so each fitting had to be performed on just one image. Also some of the needed landmarks are very specific and cannot be defined with absolute certainty on the Shroud images.

⁵ <http://www.meshlab.net/>



Figure 4.5: two three-dimensional meshes created with the Morphable Model Fitter tool from two images of the Shroud (the noise-reduced versions of DHL_03.jpg and Miller_01.jpg). The tool outputs, other than a PLY file containing the mesh, also a JPEG image with the mesh aligned with the face in the original picture and superimposed on it. These are the images resulting from the usage of the two Shroud pictures.

Given that the results were not promising, that the software only outputs a PLY file and some JPEGs, but no measurements that could be used to compare different outputs, and given that the underlying morphable model was different from the one that I was using for the rest of the project, I decided to drop the 2D to 3D reconstruction and the usage of this tool and to proceed with just the 3D to 3D reconstruction.

5. Methodology

This section contains a detailed description of the different steps that I went through to answer the research questions outlined in Section 3 and is concluded by a definition of the experiments, of which the results are presented in Section 6.

5.1. Literature research

The first step was to find more information on the Shroud of Turin, since I previously had only some general knowledge about it. After reading some historical information on the Shroud of Turin website⁶, the Shroud of Turin Wikipedia page and most of the critical summary written by the Shroud Center of Colorado [Siefker 2015], I focused on the different interpretations given to the Shroud image and the data supporting or discrediting them. The main resource for this task was the compendium by Fanti [Fanti 2011] that presents all the information regarding the image formation theories.

To learn about morphable models, I started from a compendium written by a student at Utrecht University, who analyzed different implementations of morphable models both for 2D to 3D and 3D to 3D face reconstruction [Doshi 2013], reading also all the relevant papers referenced by it. I then studied the relevant part of the PhD thesis that described the work on the morphable model implementation that I had at my disposal [ter Haar 2009] (see Section 4.2 for a description of the model).

The results of this first step were already presented in Section 2.

5.2. Image manipulation

The available images needed some more adjustments, other than the ones already described in Section 4.1, before being ready to be used. As mentioned, some of the pictures had already been manipulated with image processing softwares, however some more manipulation was needed to remove the noise. In particular the cause of most of the noise was imputable to the linen fabric pattern and to the blood stains. Digital noise reduction in order to provide a better 3D interpretation of the Shroud image had already been performed in the past [Tamburelli 1981] however I was not able to obtain the images resulting from that particular study. It is also worth noting that too much or inadequate image manipulation have been shown to lead to some false positives, as explained in [Di Lazzaro 2013], where the authors discuss how image processing can bring to misinterpretations of the images due to effects such as pareidolia and the Gestalt principles⁷.

⁶ <http://shroud.com/>

⁷ Pareidolia is a phenomenon for which we tend to see familiar patterns even where they do not exist. A famous example is the “Face on Mars” (Figure 5.1).

The Gestalt principles are some rules developed by psychologists in the 1920s to describe how the human mind tends to process visual perception by creating groups of similar elements.




Figure 5.1: the “Face on Mars”, a famous example of pareidolia.

Some of the noise, in particular the one deriving from the blood, had to be removed manually, due to the size of the stains that makes them hard to automatically replace. For this process I used the image processing software Photoshop⁸ in its version CS6 which also provides some noise reduction utilities. These utilities are useful for removing speckles and small amounts of noise so I used them to improve the quality of the images after manually removing the blood stains and the most apparent lines caused by the cloth structure.

I used the helping tool described in Section 4.3 to load those images and apply the OpenCV’s noise reduction functions that I had included in the software. Since there were multiple filters that could be applied in multiple orders, I went through a process of trial and error to select the best combination of filters. More so, given that there is not an objective way of defining what “best” means in this case, I decided basing on which combination gave the less noisy looking results. However I did not choose which filters to use randomly, but I proceeded by considering what their effect is:

- The box filter simply assigns to each pixel the average of its surroundings. This is the most basic way of filtering and it simply blurs the image.
- The gaussian blur filter uses a Gaussian function to blur the image, it is often used in computer vision to reduce noise, however it is also known to reduce details as well as edges.
- The median filter uses the median of the surrounding pixels, it is used often for noise reduction as well, especially when pre-processing images for further manipulation. It is known to preserve the edges better than the previous two filters when the noise is not too conspicuous.
- The bilateral filter uses a weighted average of the surrounding pixels. The weights are assigned in such a way that the process is very good at preserving the existing edges. The downside is that it is subject to the so-called staircase effect, where some flat intensity regions appear on the image and thus new edges are created, transforming gradients in “staircases”.
- The non-local means filter [Buades 2005] instead of calculating the mean in a surrounding area it does it for the whole image (hence non-local), using the difference between the active pixel and the other pixels as a weight. This filter is very good at maintaining details while increasing the image clarity, the downside is that it takes quite a long time to process.

Finally, since I was planning to use a morphable model and morphable models are concerned only with faces, it was necessary to select only the part of the images containing the face and exclude all the surroundings. To do so the easiest way was to create a black and white mask that defines which pixels in each image are part of the

⁸ <http://www.adobe.com/products/photoshop.html>

face and which aren't. Since all the images contain the same face, albeit with the aesthetic differences caused by the different ways of achieving them, it made sense to use the same mask for all of them. So I shrunk and cropped the images until they all fitted inside the same mask that I created. This mask can be loaded by my tool, which at the moment of creating the three-dimensional mesh disregards any pixel that is covered by the mask. The results of this step and the following one are presented in Section 6.

5.3. Three-dimensional mesh

The image of the Shroud reportedly carries some three-dimensional information [Jackson 1984]. If that is the case it should be possible to generate from it a three-dimensional mesh representing a face. In order to do so it is necessary to understand what the interpretation of this information is. The interpretation strongly depends on the means by which the image was formed. Considering the different theories listed in Section 2.2 and the current beliefs within the scientific community, there are only two image formation hypotheses that are worth to investigate, that is the two that are still being studied and have not been dismissed by the scholars: the image as a result of radiation and the image as a forgery.

In the first case the interpretation given by the scholars is that the intensity of the color on the image is representative of the distance between the body and the linen cloth that was wrapped around it. This idea originated from an experiment, narrated in [Schumacher 1999], where the author produced a three-dimensional model with the VP-8 Image Analyzer, a machine that interprets the white levels in an image as depth values and produces a depth map from them. In the second case it is difficult, if not impossible to give an interpretation. This is because not even the method of fabrication of the image has been settled upon, also it would be hard for a medieval artist to willingly include such information. It is thus more likely that in this case the presence of three-dimensional information is coincidental. It is also possible that the three-dimensional information is void of the meaning that some scholars try to give it and that the results of the VP-8 Image Analyzer experiment and those alike are being given more significance than they actually carry.

With all of this in mind, and given that the color intensity as depth information is the only interpretation of the image's possible three-dimensional information found in previous literature that has a mathematical interpretation, it is also in my opinion the only theory explorable as well as worth exploring.

To generate a three-dimensional mesh from that interpretation the first step is to create a depth map similarly to how it was done with the VP-8 Image Analyzer. In a preliminary experiment I did exactly that, obtaining the result depicted in Figure 5.2. Despite the result appearing promising at a first glance, that image suffers of a big bias. Without thinking of the consequences I left on each three-dimensional point the same color that the corresponding pixel had on the original image. This causes our eyes to see a shape even where there is actually none. In fact, the same mesh visualized without any texture is too full of noise to distinguish any actual shape, if not barely. It is important to avoid this kind of biases, or at least to recognize them soon, and not to draw conclusions from them.




Figure 5.2: Result of a preliminary experiment where I tried to replicate the VP-8 Image Analyzer process (with and without colors). The image used was Miller_01.jpg.

This is just a preliminary result obtained without any noise reduction from the image named *Miller_01.jpg* (Figure 4.3). It looks more noisy than the result obtained by the VP-8 Image Analyzer (Figure 5.4), however there are some considerations to be made. First, the VP-8 Image Analyzer, being a machine built around the 1970s could not have had the same resolution as the pictures in my possession, with so many more pixels there is space for a lot more noise. Second, there is not much information about the machine⁹ and its way of creating the depth map, so it is hard to say if it limited itself to plainly taking the color values and outputting them as depth the way I did or if it did some kind of smoothing in the process.



Figure 5.3: The depth map created by the VP-8 Image Analyzer.

To improve the result from Figure 5.2 I applied some noise reduction to the image, as described in the previous section, and reduced the resolution (*i.e.* the number of final 3D points that form the mesh).

Another concern with the VP-8 Image Analyzer experiment and my preliminary experiment alike is that in them the Shroud was represented as a flat surface (since the photos were taken while the linen cloth was straightened), however when the image

⁹ I was able to find a brochure with the VP-8 Image Analyzer technical specifications at: <http://shroudnm.com/docs/1977-08-VP8ImageAnalyzerBrochure.pdf>

was formed the Shroud was supposed to be wrapped around a body. This means that by flattening it and then creating a depth map, the information about the cloth shape is not considered and this is why the results obtained in that way look more like bas reliefs than actual three-dimensional faces. If there was a way of knowing the Shroud's shape it would be possible to have a more accurate result, unfortunately we cannot know for sure the exact shape that the Shroud would have had, since we do not know for sure the shape of the body that was wrapped in it. Nonetheless it is possible to recreate a similar scenario and thus obtain an approximation of the shape. In [Jackson 1984] the authors tackled this issue by laying an actual linen cloth on a volunteer and then recording its shape and digitizing it. They then modified the depth map created by the VP-8 Image Analyzer accordingly, as shown in Figure 5.4. As it is possible to see in the picture, the authors considered two possible ways of mapping the cloth shape to the depth map: vertically and perpendicularly to the cloth. In the paper they eventually considered the first approach as more realistic. It is also the the approach that conforms better to the image formation by radiation hypotheses, where the radiation is imagined as being a vertical field.




Figure 5.4: Visual explanation of how the depth map is modified according to the cloth shape in [Jackson 1984]. The author proposed two ways of considering the projection: vertical and perpendicular to the cloth.

To obtain a similar model of a cloth, using the same method (*i.e.* using a real person and a real cloth and then digitizing the result) could prove to be cumbersome. Nowadays, however, computers are much more powerful in respect to when the paper was written and physics simulations can be very accurate. In fact I was able to simulate a cloth being laid on a person using the computer graphics software Blender¹⁰, which contains an implementation of the Bullet Physics library¹¹. After obtaining the model of the cloth I replicated the same process described in the paper to obtain the final version of the mesh.

¹⁰ <https://www.blender.org/>

¹¹ <http://bulletphysics.org/>

5.4. Evaluation

In the morphable model a face is described by a set of 99 parameters (as detailed in Section 4.2). The parameters of each one of the original faces that served to build the model are distributed around the mean face ones. A way of measuring how good (*i.e.* realistic) the fitting of a mesh obtained from the Shroud image is, could be to see where its parameters fall in such a distribution.

Analyzing all the 99 parameters together in a single distribution could cause some meaningless results due to what is known as the curse of dimensionality¹². Moreover the parameters do not have all the same impact on the characteristics of the faces they describe and thus some parameters might be more important than others to analyze. For example the first 10 parameters are used for what is called the coarse fitting of the morphable model, that is defining the general shape of the face that is being created (*e.g.* height, width and depth), and, considering that I will be using a single mask to define where the face is in the Shroud images, it is likely that all the faces created will have a similar general shape and thus that the corresponding parameters will not vary much. This is why it would probably be wiser to analyze each parameter individually, giving more importance to the more meaningful ones. Unfortunately, however, there is no way of knowing with precision what each parameter does, since the parameters are the result of a mathematical process and are not designed one by one. I had at my disposal a visual tool, developed by another student, that allows to see the mean face changing shape as you manually increase or decrease the parameters values, but even with that it was really hard to understand what was really going on and to assign a specific role to each parameter. This is why I chose instead to look at how the values of each parameter are distributed in a sample of 234 faces that I knew for sure were created by fitting the model to real faces. From this analysis I noticed that only some of the values were normally distributed while others were more evenly distributed.

Categorizing the results of the morphable model fittings based on the 99 parameters, dividing them between fittings of faces and fittings of non-faces, is a way of deciding whether the meshes created from the Shroud images belong to the faces or the non-faces. To do so I fitted the morphable model to some meshes created from Shroud images and some meshes that were not faces. Initially I tried using three-dimensional models of objects (*e.g.* a teapot), but due to some technical issues with the fitting software I had to switch to meshes created with some noise functions. Once obtained these meshes I analyzed the distributions of those parameters that in the 234 real faces are normally distributed and calculated whether those resulting from the Shroud fittings could be distinguished from those resulting from the noise fittings. The results of this evaluation, which were obtained using MATLAB¹³ are presented in Section 7.

Another interesting form of evaluation that I initially thought of, this time qualitative, was a human evaluation, where the results are shown to some people and their opinion is asked. This is because even if face recognition algorithms exist, humans are still the best at recognizing other humans. Humans are also known to experience the so called uncanny valley [Mori 2012] where a face that closely resembles a human one but has

¹² The curse of dimensionality refers to the fact that analyzing data in high-dimensional spaces often results in the data being too sparse to convey any meaning at all.

¹³ <https://www.mathworks.com/products/matlab.html>

some slight differences causes some uneasiness in human observers. Unfortunately, for lack of time I was not able to proceed with this second evaluation, that was however less important than the quantitative one and that I had planned only as a nice-to-have feature for this project.

5.5. Experiments definition

To put the described methodologies into practice I devised three experiments that differ on how the images are treated and how the three-dimensional mesh is created. For all the experiments I chose to use one picture from each of the sets described in Section 4.1. This is because some of the pictures are very similar to each other and particularly to those in the same set (e.g. the images in the DHL set, shown in Figure 4.2 are actually the same image, just with different resolution or contrast), rendering the usage of all of them unlikely to give different results and thus redundant. The pictures from different sets however have been created in different ways, or have undergone different kinds of manipulations, so choosing only a single picture would be too selective.

In each set the chosen picture was the one where the face image is more clear and defined. The resulting list is:

- *Enrie_03.jpg*
- *DHL_03.jpg*
- *Miller_01.jpg*

Each experiment was conducted using the tool that I created and following these steps, which were repeated for every picture:

1. Load the image in the tool. The loaded image can be:
 - a. The raw version of the image;
 - b. The filtered version of the image, that went through the image manipulation process described in Section 5.2.
2. Use the tool to generate and export a three-dimensional mesh. The way to pass from the image to the mesh, described in Section 5.3, is to create a vertex for each pixel included by the mask, such that $v = (x, y, z)$, where: v is the vertex; x and y are the pixel's indices, shifted and scaled to be aligned with the mean face of the morphable model (since all the pictures have the same resolution and the face is in the same position in all of them, the shifting and scaling is consequently the same as well); z can be calculated using one of the following formulae:
 - a. $z = \alpha I_s$ where I_s is the image intensity at that pixel and α is a scaling factor;
 - b. $z = \beta I_c - \alpha I_s$ where I_s and α are the same as before (α has the same role, not necessarily the same value), I_c is the intensity of the cloth height map and β is another scaling factor.
3. Use the morphable model software to create a fitting from the mesh;

4. Collect and analyze the parameters resulting from the fitting, comparing them to those coming from real faces and non-faces.

The way of conducting the first two steps, that is the way of obtaining the three-dimensional mesh, is what differentiates the experiments from each other. In conclusion the three experiments are defined as such:

- **Experiment 1** is similar to the VP-8 Image Analyzer test, the main differences being the usage of the mask and the higher resolution of the images. It follows steps 1.a and 2.a.
- **Experiment 2** improves on the previous one by adding image manipulation before the three-dimensional mesh creation. It follows steps 1.b and 2.a.
- **Experiment 3** maintains the usage of polished images and adds the cloth shape into account. It follows steps 1.b and 2.b.

The experiments were conducted in sequence up until step 3, however step 4, that is the final evaluation, was done after having collected the parameters from all three of them.

6. Results

6.1. Preliminary phase




Figure 6.1: In order dhl.jpg, enrie.jpg, miller.jpg and mask.jpg. These pictures have all the same dimensions (573x763px) and in each one the face image is in the same position (described by the white pixels in mask.jpg).

During the preliminary phase I scaled each picture to the same resolution and cropped and rotated them so that the face would be in the same position in each picture. I finally renamed them *dhl.jpg*, *enrie.jpg*, *miller.jpg*. I then created a mask (*mask.jpg*) that defines the position of the face. The results can be seen in Figure 6.1.

One issue that I was able to identify after fitting the morphable model to the first meshes I created is that the model is built using faces without any facial hair and thus it is not capable of creating any face with facial hair. The face on the Shroud however appears to have a beard, so the point cloud obtained from its image will have a beard too (*i.e.* the vertices corresponding to the beard pixels will be more prominent) and the morphable model will not be able to interpret it during the fitting, producing instead a face with a more prominent chin (Figure 6.2).




Figure 6.2: In order: the mean face of the morphable model, the fitting of a mesh created from the image miller.jpg not masking the beard and masking the beard and finally the mask used to remove the beard. See how in the first fitting the chin is much more prominent due to the beard being misinterpreted and in the second fitting the face is shorter due to the hole not being filled in completely.

The morphable model is built in such a way that it should be able to fill blank holes, so I tried to create another mask, that would cut out the beard as well, and to use it in place of *mask.jpg*, in the hopes that the model would be able to fill in the space by itself. Unfortunately it turned out that the morphable model cannot cover holes as big as this one and would instead create shorter faces as is evident from Figure 6.2. Creating a whole new morphable model that could handle beards would take too much time and would go beyond the scope of my project, so I decided to discard the new mask and proceed with the old one. This is because I think it is better to have a face with a more prominent chin than necessary rather than a face that is shorter and thus has wrong proportions.

6.2. Experiment 1

The first experiment was very similar to the VP-8 Image Analyzer experiment and to the preliminary experiment described in Section 3.3. The data is obtained in a similar fashion, but the images are masked and the result is passed through the morphable model fitting.

As was already clear from the preliminary experiment, the images are too much noisy and thus the resulting meshes do not look like faces at all, but more like forests of spikes. Only the mesh obtained from *miller.jpg* is clear enough that you can distinguish some details of the Shroud from some angles (Figure 6.3).




Figure 6.3: the meshes created from (in order) *dhl.jpg*, *enrie.jpg* and *miller.jpg*.
First a frontal view and then a side view.

The fact that only in the last mesh some features are visible was expected since the Miller scans had already been manipulated before entering my possession. In fact, as it is possible to see in Figure 6.1, *miller.jpg* has considerably less noise than the other two selected images.

Figure 6.4 shows the results of the fitting of the morphable model to the meshes.

Despite the noise, the model was able to generate some credible faces, although some features look a bit unconventional. In particular all the faces are very flat, due to the fact that the original mesh is derived from a flat surface, and there seem to be some noise around the eyes, making them appear strange and slightly uncanny. Finally, the chin, as mentioned before, is quite prominent due to the misinterpretation of the beard. In these fittings it is almost as prominent as the nose.




Figure 6.4: the faces resulting from the meshes in Figure 6.3

6.3. Experiment 2




Figure 6.5: a comparison between the original version of the images (on the left) and their polished version (on the right). The images are in order: dhl.jpg, enrie.jpg and miller.jpg

In experiment 2, before loading the images in my tool, I polished them with Photoshop, manually removing the blood stains and the more evident lines created by the folds of the cloth and its fabric and then using the noise reduction tool of the program. Even if I am using a mask and thus only the face part is relevant, I polished the images in their entirety because it did not take too much extra time and the images might be useful for something else in the future.

The result, which is visible in Figure 6.5, was then loaded in my tool, where it got through some of the noise reduction functions of OpenCV. The meshes created consequently are shown in Figure 6.6.




Figure 6.6: the meshes created with the polished images.

Polishing the images made it so that some details of the face can be seen on the resulting meshes, especially on the one created from *miller.jpg*. All the meshes are still quite flat, with the tip of the nose on the same level as the cheekbones, the eyebrows and the chin.

The faces created by the morphable model (Figure 6.7) do not look much better than the ones obtained in the previous experiment, just slightly less flat. The roughness around the eyes is still there and in some cases spreads even to the nose.




Figure 6.7: the faces resulting from the meshes in Figure 6.6

6.4. Experiment 3

For this experiment I first simulated a cloth folding over a face using Blender. The face used was the mean face of the model (see Figure 6.2). I then used the resulting mesh to create an height map, which I aligned to *mask.jpg*. The result is shown in Figure 6.8.




Figure 6.8: a mesh with a piece of cloth lying over a face simulated with Blender and the height map created from it. The height map is aligned with *mask.jpg*

The interpretation of the Shroud image that I am considering considers it as a representation of the distance between the body and the cloth. The cloth height map represents the distance between the cloth and a plane. The distance of the face from the plane can thus be obtained by subtracting the first from the second with a formula like:

$$z = \beta I_c - \alpha I_s$$

Where:

- I_c is the intensity of the cloth height map;
- I_s is the intensity of the Shroud image;
- α and β are two weights used because although both intensities can range from 0 to 1, they represent different distances that likely have different ranges.

A schema of this interpretation is shown in Figure 6.9.




Figure 6.9: a schema of the interpretation that I used to create the three-dimensional mesh in this experiment. The cloth height map gives the distance between the cloth and a plane and the Shroud image gives the distance between the cloth and the body. To get the distance between the body and the plane (i.e. the z coordinate of the points in the mesh) it is sufficient to subtract the two values for each point. In the picture the point that is being considered is the central point.

The meshes resulting from the formula above and the polished images created in the previous experiment are noticeably less flat than the ones obtained before. More so the nose is more prominent than the other features (Figure 6.10). Unfortunately there is still some noise and, in particular, the first mesh is still the one with less distinguishable details.




Figure 6.10: the meshes obtained by applying the cloth depth map to the images created in experiment 2.

Figure 6.11 shows the fitting of the meshes. It is clear that the cloth depth map had a noticeable effect on the fitting process, the faces have more curvature and the nose is no longer flattened out. The weird shape of the eyes and the general uncanniness however are still there.



Figure 6.11: the faces resulting from the meshes in Figure 6.10.

The appendix contains a printout of the CSV file with all the parameters resulting from the fittings (Section A.1) and a table with pictures of all the three-dimensional meshes and their fittings (Section A.2).

6.5 Other results

While performing the experiments I tried to use the tool that I implemented with some face images that were not taken from the Shroud. Specifically I used two paintings that represent the face of Jesus Christ¹⁴. The first painting is a more realistic depiction, whereas the second one is more sketched and artistic. The aim of this test was to see if the same approach used with the Shroud images would lead to similar results with some artistic productions. The resulting meshes are different from the Shroud ones and from each other, however the fitting results are similar, in that the fitting parameters are not comparable to those resulting from real faces, as the next section will detail.




Table 6.1: The meshes and fittings obtained from two paintings.

¹⁴ The two paintings are respectively “Prince of Peace” by Akiane Kramarik and the “Face of Christ” by Larry Cole.

7. Evaluation

Each one of the faces created with the morphable model can be described by a set of 99 parameters, that the morphable model fitting software saves as a text file. I wrote a script to gather the parameters of the faces created during the three experiments and group them into a single CSV file. There are a total of nine faces created from Shroud images (three for each experiment).

In order to evaluate these results I compared them with some fitting of real faces and some fittings of non-faces. The reasoning that led to this decision is that if it is possible to distinguish between the two categories from the fitting outputs, the Shroud fittings could be assigned to one of the two categories, it would then be possible to say whether the meshes obtained from the Shroud images can be considered representations of faces or not.

7.1. Fitting real faces

I had at my disposal 234 meshes created from three-dimensional laser scans of real faces. These were obtained from the work of other students on the same morphable model that I am using. An example of a face mesh and the corresponding fitting can be seen in Figure 7.1.




Figure 7.1: a mesh obtained from a laser scan of a real face (left) and the morphable model fitting obtained from it (right).

After performing the model fitting for all the 234 faces, I used the same script as before to gather the resulting parameters and group them into another single CSV file.

7.2. Fitting non-face meshes

To generate some non-face fittings I first tried using two of the classical computer graphics meshes: the Utah teapot and the Stanford bunny (Figure 7.2). Unfortunately the fitting software crashed without giving any meaningful error message. I then tried cutting the bottom part of the teapot, because I thought that since the faces are not supposed to be closed surfaces maybe that was causing some issues, but the program still crashed. It turned out that if there are two overlapping vertices the software cannot handle the three-dimensional mesh.




Figure 7.2: the Stanford bunny and the Utah teapot

I then decided to try a different approach, I created three images with the same resolution as the Shroud images described in section 4.1 (573x763px), the first image is completely black, the second is filled with random noise and the third with a gradient noise similar to Perlin noise [Perlin 1985] (Figure 7.3).




Figure 7.3: flat_plane.jpg, random_noise.jpg and perlin_noise.jpg

I loaded all this images in my tool and created three meshes for each one of them, so that I would end with nine non-face meshes:

- For the first triplet I just extruded the z value as described before;
- For the second triplet I first applied the mask;
- For the third triplet I applied the mask and the cloth height map.

The resulting meshes are displayed in Figure 7.4.




Figure 7.4: the nine meshes created from the images in Figure 7.3.

This method guarantees that the generated meshes do not have overlapping vertices and in fact the software did not crash, even with the rectangular shapes (it is somewhat noteworthy that for the three random noise meshes it took way more time than for all the other fittings that I did during this project). From the results of the fittings, which can be seen in the Appendix, I again grouped all the parameters into a third CSV file with the same script I used previously.

7.4. Robustness of the fitting process

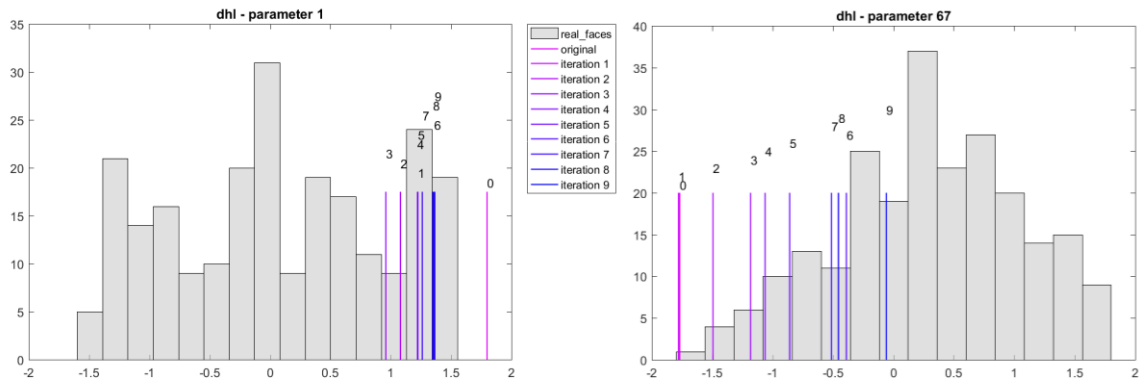


Figure 7.5: two plots that show the different values of parameters 1 and 67 for 9 fitting iterations on the image dhl.jpg. The first graph shows the values converging, which is the more common case, whereas the second one shows the less common case of the values diverging.

Before starting to compare the results of the fitting process, we wanted to understand if the process itself was robust. To do so I iterated the fitting process nine times, using the output of an iteration as the input for the following one. I did so for each of the nine Shroud meshes and the nine non-face meshes, collecting the output parameters in a CSV file. I then created a plot for each parameter of each iterative process where the values of the parameter for each iteration is visualized as a vertical line. For almost all of these plots the parameters seem to converge towards a certain value after some oscillation, except for some rare cases. Figure 7.5 shows an example of a plot where the parameters values converge and one where they do not.

In the appendix there is the complete list of plots from which I took the ones that appear in this and in the following sections (Section A.3).

7.5. Choosing the significant parameters

The devised way to compare the fitting results with each other is to compare their parameters, however as mentioned before using all 99 of them could create some confusion. In fact different parameters give different contribution to the fitting and some are more relevant than others. More so some parameters tend to vary a lot, while others remain more or less similar throughout all the fittings.

First of all I measured the mean, median, standard deviation and skewness of all 99 parameters for the sample constituted by the 234 fittings of real faces. This can be considered a sample from the population of all the possible fittings of real faces and thus it is good for describing the characteristics that the fittings of real faces have. I

plotted these measures into four graphs to see if they followed any trend (for example the first parameters could have higher importance and the last one a lower one or vice versa), however this does not seem to be the case (Figure 7.6).

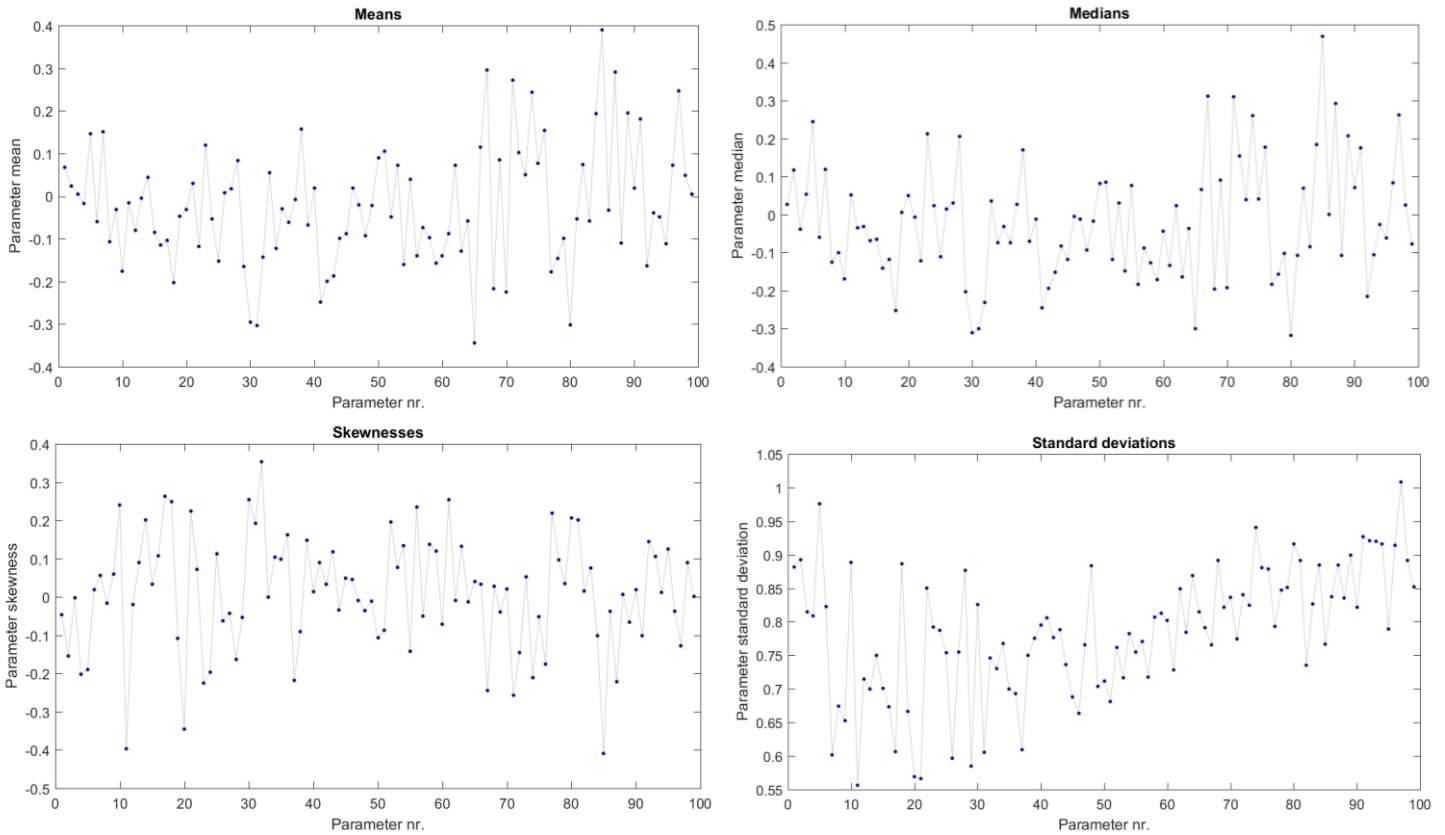


Figure 7.6: The graphs displaying mean, median, skewness and standard deviation for the 99 parameters of the 234 real faces. It is immediately apparent that these measures do not follow any particular trend based on the parameter number.

To understand more about every parameter I created 99 histograms plots, one for each parameter, using again the list of 234 real face fittings. From those it became apparent that some of them were similar to normal distributions, whereas others were more randomly distributed (Figure 7.7 shows one example of both cases). To have a measure of the fitness to a normal distribution that is more quantitative than just plotting a normal distribution over the histograms I run the Anderson-Darling test for every parameter sample.

With this normality test I was able to identify all the parameters for which the sample of real faces fittings resembles a normal distribution. I then discarded those that failed the test from my analysis. This is because I noticed that the Shroud fittings tend to have more parameters far from the mean and thus, if that is the case even for those parameters that should follow a normal distribution, it would be possible to assert that the Shroud fittings do not come from the same distribution of faces as the real face fittings.

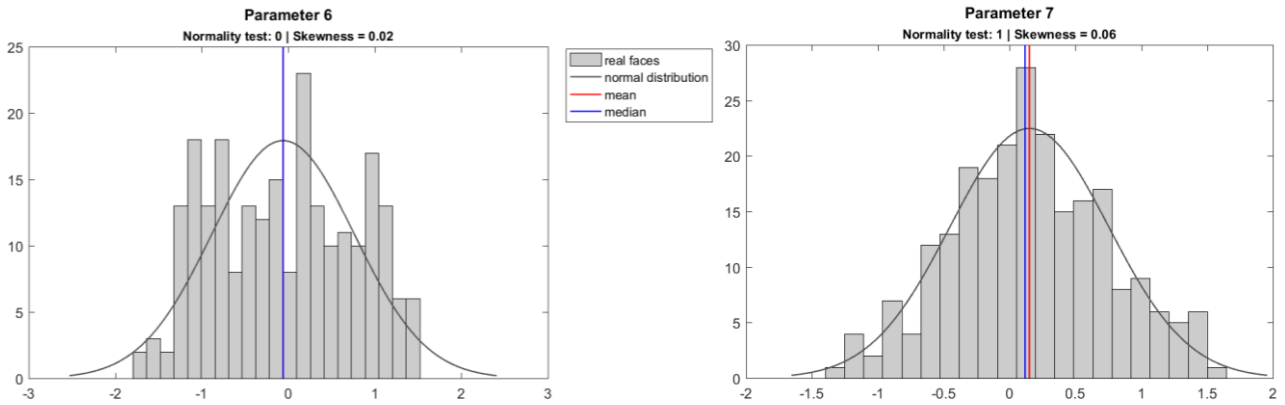


Figure 7.7: histogram plots with 21 bins for the parameters 6 and 7. The first fails the normality test whereas the second passes it.

Figure 7.8 shows an example of this situation, where the parameters follow a normal distribution and a random sample of nine real faces is much closer to the mean than the nine Shroud fittings and the nine noise fittings. Of course it is not possible to immediately draw a conclusion from just one of the 99 parameters, however this example illustrates how the parameters that follow a normal distribution can help in reaching such conclusion. It would be harder to assert that those fittings come from a different population if the underlying histograms would not follow a normal distribution, as can be seen from the second graph in Figure 7.8. The inspection of those parameters that do not follow a normal distribution therefore cannot give many useful insights, which means that the parameters can be safely discarded from further analysis.

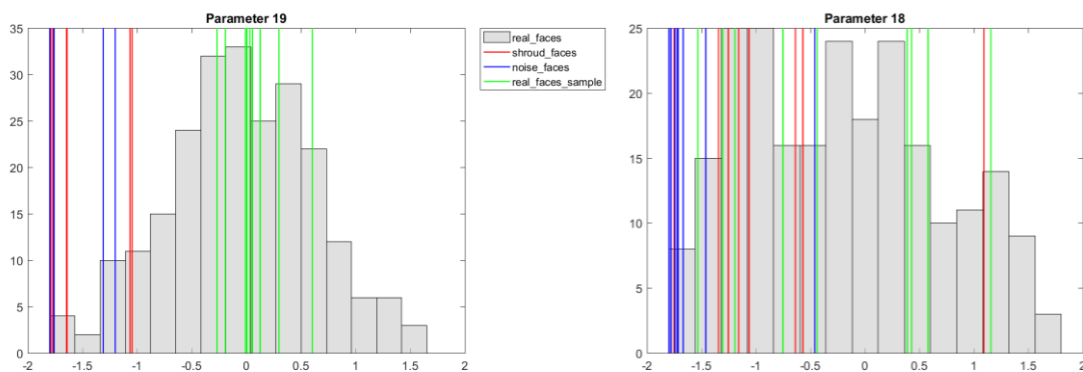


Figure 7.8: the histogram plot for parameter 19 and 18 (this time with 15 bins). The vertical lines are the 9 Shroud fittings (red), the 9 non-face fittings (blue) and 9 random samples from the 234 real faces fittings (green). The first is an example of how for some parameters the Shroud and noise fittings tend to be more far off from the mean compared to the real face fittings. In the second, even if the Shroud and non-face fittings are still more decentralized than the real face fittings it is harder to assert that they come from different distributions since the underlying histograms have taller bins even far from the mean.

More so, as mentioned before, the first ten parameters are used to model the general shape of the face (*i.e.* width, height, slantiness and so forth). Given that I am using a mask to define most of those features it is safe to assume that those parameters will not be very meaningful for my experiments, even if they were to follow a normal

distribution, and indeed in most of the Shroud and noise fittings they tend to have a similar value (Figure 7.9).

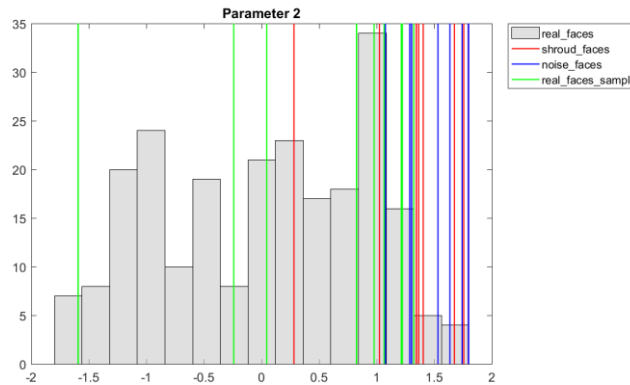


Figure 7.9: the histogram plot for parameter 1. The line colors are the same as those in the other plots. It is noteworthy that, for the Shroud and non face fittings, the values tend to be bundled around the same point. This is because the first ten parameters are used to describe the shape of the face, which in my experiments is always defined by the same mask.

Eventually I indeed discarded the parameters that describe the general shape of the face and the parameters that are not distributed as a normal distribution, ending up with the following list of valid parameters: 12, 13, 15, 16, 17, 19, 21, 22, 25, 26, 27, 29, 31, 33, 34, 35, 36, 37, 38, 39, 40, 42, 44, 45, 46, 47, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 61, 62, 63, 66, 67, 69, 71, 72, 73, 75, 76, 77, 78, 79, 82, 83, 86, 88, 90, 95, 98. In total I discarded 41 parameters and kept 58.

7.6. Comparing the fitting results

The most common way to measure where a data point is located in a distribution is the z-score:

$$z = \frac{x - \underline{X}}{S}$$

Where x is the data point, \underline{X} is the sample mean and S is the sample standard deviation. In other words the z-score is a measure of how many standard deviations the data point is far from the mean. For a generic distribution Chebyshev's rule guarantees that at least $1 - 1/k^2$ data points are within k standard deviations from the mean (*i.e.* have a z-score for which the absolute value is lower than k). Since the values of all the valid parameters are normally distributed, however, it is possible to use the more strict 68-95-99.7 rule, which says that in a normal distribution at least 68% of the data points are within one standard deviation from the mean (*i.e.* their z-score is contained in the $[-1 \ 1]$ range), at least 95% of them are within 2 standard deviations from the mean (*i.e.* their z-score is in the $[-2 \ 2]$ range) and at least 99.7% of them are within 3 standard deviations from the mean (*i.e.* their z-score is in the $[-3 \ 3]$ range).




Figure 7.10: the parameters z-scores distribution for the fittings of two different real faces. The first one is more representative of the other real faces, whereas the second one is one of the few exceptions where an outer bin is significantly taller than an inner bin.

To verify this, for each of the 234 real faces fittings, I measured the z-scores of their 58 valid parameters and plotted them in a histogram chart with 8 bins one standard deviation wide. As expected most of the real faces respect those proportions, having most of the data points within one standard deviation from the mean and almost all of them within 3 standard deviations. Among the 234 histogram plots obtained there are only a few that do not respect such proportions. Figure 7.10 shows an example of A plot that respect the proportions and one that does not. It is also worth noting that only 2 out of the 13572 parameters described by these graphs (234 faces with 58 valid parameters each) have a z-score lower than -3 or higher than 3, which is the 0.015%.

Not all of the histograms have high enough scores to fit the 68-95-99.7 rule, with the lowest scores being 34.5% for the [-1 1] range (the second plot in Figure 7.10), 81.0% for the [-2 2] range and 98.3% for the [-3 3] range, however most of them have a score above 60% in the [-1 1] range, above 90% in the [-2 2] range and 100% in the [-3 3] range. It is also worth pointing out that the scores are not supposed to adhere strictly to the 68-95-99.7 rule since the rule is applied to normal distributions and I am counting together samples from different populations that are normally distributed (the populations being the 58 ranges of valid samples possible values that were selected in the previous section). What the rule tells is the probability of a sample of falling in the three ranges, which means that with enough samples their distribution would follow the rule. With just 58 samples (from different populations) exceptions can happen.

I repeated the same process for the Shroud and non-face fittings parameters. The non-faces z-scores histograms look quite different from the real faces ones, as expected. They have much taller bins further from 0 and the [-1 0] and [0 1] bins are generally short. This is true especially for the fittings of rectangular meshes (the ones created by simply extruding the values from *plane.jpg*, *noise.jpg* and *perlin.jpg*), for the other meshes the difference between the bins heights is not as big. However they do not respect the distribution of the parameters z-values for the fitting of real faces, although in some cases they come close to the borderline. There is one notable exception to what just said, in fact the histograms plot for the fitting to the flat plane with the mask is similar enough to the plots for the real faces (Figure 7.11).




Figure 7.11: the parameters z-scores distribution for the fittings of the three meshes obtained from the flat plane. The first plot is for the unaltered flat plane, the second for the plane with the mask and the third for the plane with the mask and the cloth depth map. The second plot is the only exception in the set of non-faces z-scores distributions that resembles a distribution of real faces z-scores.

Even with this resemblance, however, its score in the $[-2, 2]$ range is only 82.8% (the 68-95-99.7 rule requires a 95%) and of the 234 real faces distribution only 2 have this exact same score and only 1 has a lower one (81%). It is also important noting that this exception appeared only once, even with previous non-face fittings that I have generated while testing the scripts that I wrote and that all resulted more similar to the other two configurations visible in Figure 7.11. I have not included those graphs in this document, as they were not part of the experiments, but they are present in the material that accompanies it. Finally, there are 7 out of 522 parameters that have a z-score below -3 or above 3, which is the 1.34%, around 90 times more than the set of real faces and enough for some distributions to fail the 68-95-99.7 rule for the $[-3, 3]$ range. In fact none of the distributions comes close to respecting the rule in all three ranges at once (not even in two at once), except for the one discussed above, that is close but still fails both in the $[-1, 1]$ and in the $[-2, 2]$ ranges.




Figure 7.12: the parameters z-scores distribution for the fittings of the three meshes obtained from miller.jpg. The first plot is for the unaltered image, the second for the filtered one and the third for the filtered image with the cloth depth map. The second plot is the one that comes closest to respecting the 68-95-99.7 rule in the Shroud set.

The Shroud faces z-scores histograms do not have such extreme cases, with very tall bins far from the mean, as the non-faces ones, however they still look similar to the non-faces z-scores histograms that do not come from fittings of rectangular meshes and they are not very similar to the real faces ones. The z-scores are more evenly distributed between the bins, which means that the outer bins are taller and the inner bins are shorter than they should be in order for these histograms to be considered close to the real faces z-scores histograms (Figure 7.12).

The second histogram in Figure 7.12 has a score of 39.7% in the [-1 1] range, 77.6% in the [-2 2] range and 98.3% in the [-3 3] range. It is the best scoring histogram for the first two ranges and still is quite far from respecting the 68-95-99.7 rule. Between these histograms there is no exception like the one that appeared in the non-faces set of histograms. Finally, there are 3 out of 522 parameters that have a z-score below -3 or above 3, which is the 0.57%, less than the non-faces ones but still almost 40 times more than the real faces ones.

Table 7.1 is a summary of all the z-scores from the three sets of fittings that shows their minimum, maximum and average for the three different ranges. Overall, it is clear that the real faces score better than the other two sets, especially considering the averages (66.4%, 96.2% and 100%) which are very similar to the boundaries set by the 68-95-99.7 rule. This was expected because 13572 z-scores are quite a big sample that, being taken from normally distributed populations, should adhere to the rule. In conclusion the Shroud fittings seem to have a slightly better score than the non-face fittings, not good enough to be considered fittings of real faces though. Also the set of non-face fittings contains the borderline cases of the rectangular-shaped meshes, which have contributed to lowering their average scores.

Range	Measure	Real faces	Non-faces	Shroud
[-1 1]	Min	34.5%	5.2%	20.7%
	Max	93.1%	63.8%	37.9%
	Avg	66.4%	22.6%	29.3%
[-2 2]	Min	81.0%	27.6%	56.9%
	Max	100.0%	82.2%	77.6%
	Avg	96.2%	51.3%	66.1%
[-3 3]	Min	98.3%	94.8%	98.3%
	Max	100.0%	100.0%	100.0%
	Avg	100.0%	98.1%	99.4%

Table 7.1: the minimum, maximum and average scores in the three sets of fittings for the three z-scores ranges.

8. Conclusions

In this thesis I studied the claim that the image on the Shroud of Turin contains some three-dimensional information using a morphable model. The results show that it can be possible to distinguish between three-dimensional meshes that represent faces and three-dimensional meshes that do not by fitting a morphable model to the meshes and analyzing the resulting parameters. The fitting parameters of a mesh that represents a face tend in fact to be normally distributed around the mean face of the morphable model, whereas the fitting parameters of a mesh that does not represent a face vary from a distribution that is close to uniform to more extreme distributions where the bulk of the values is more than one standard deviation away from the mean.

With this metric, it would seem that the three-dimensional meshes created from the Shroud using the image intensity as a measure of the distance between the body and the cloth, could not be defined real faces. There are, however, some caveats to be added. First of all, the faces population described by a particular implementation of a morphable model is not the population of all the possible existing faces, it is in fact a population limited by the capabilities of the model itself. A morphable model is in fact bound by the original set of faces with which it was built, it can create faces that are different from the original ones, but it cannot stray too much from them. In this particular case the morphable model implementation that I had at my disposal did not include bearded faces, whereas the face on the Shroud appears to have one. The second caveat is that the statistics that I used to distinguish between faces and non-faces is not infallible, it does not define without doubt whether a three-dimensional mesh is a face or not, but it defines the likelihood of that mesh to be a face or not. In some cases, like the parameter distribution of the fitting to a flat plane, the likelihood of not being a face is very high, in other cases it is lower. For example Figure 8.1 shows two examples of distributions that look similar, one from a real face and the other from a mesh obtained from the Shroud. If it was not known that the first histogram originated from a real face it could be mistakenly classified it as a non-face.




Figure 8.1: two histograms resulting one from a real face and the other from a Shroud mesh, but both with a similar distribution. This shows that the distinction between faces and non-faces is not infallible and there can be gray situations.

This shows that it is possible, albeit unlikely, for the parameters of the fitting to a real face to be distributed in such a way that it could not be classified as a real face with

confidence. It also means that the second histogram could be the result of a fitting to a real face, as it is not much different from the first one. Since the first one is one of the worst scoring in the set of real faces and the second one is the best scoring in the set of the Shroud meshes it is highly unlikely that the Shroud meshes can be considered faces, but it is not impossible. Similarly the exception in the set of non-faces described in section 7.6 shows that it is possible, albeit unlikely, for the parameter distribution of a fitting to a non-face to be close enough to a normal distribution that the original mesh could be erroneously classified as a real face.

Finally, the fact that the meshes obtained from the Shroud are not likely to be faces does not necessarily mean that the Shroud image does not contain three-dimensional information. I do not think that the results of this study are enough to dismiss that claim, but maybe they can bring some more attention to it. In fact, being a novice to the subject, I found it quite hard to find the relevant literature about the Shroud and the interpretation of its image. I found a lot of papers that were not peer reviewed (also those that are peer reviewed mostly date back to the STURP analyses), or that were self-published or self-referential and the different compendiums that I read would not distinguish between the former and the latter categories. I think that a good way to test the claim that the Shroud image contains three-dimensional information would be to first define better what these information are. These are the reasons that lead me to abstain from commenting on the claim itself and stick to saying that the meshes that I obtained are unlikely to be faces, but there could be a different way of generating them if the claim were to be defined more in detail by the sindonologists.

With all of this in mind, I will now present the answers to the four research questions posed in section 3.1 and I will conclude with a discussion on possible future improvements to my study.

8.1 Answers to the research questions

- **Is there a plausible interpretation of the Shroud of Turin's image formation that leads to the construction of a three-dimensional mesh?**
Yes, the interpretation that sees the Shroud image intensity as a measure of the distance between the body and the cloth can be mathematically defined in such a way that it is possible to create three-dimensional meshes. However it is also currently the only interpretation that can lead to the creation of a three-dimensional mesh.
- **Can a morphable model be consistently fitted to the three-dimensional mesh obtained from the Shroud of Turin's image?**
Yes, in fact with the proper construction and alignment any three-dimensional mesh can be used as input of the morphable model fitting function to obtain a three-dimensional face as an output. The face-like characteristics of the output however depend on the quality of the input. A morphable model is also limited by its implementation. The one that I had available for example could not create bearded faces.
- **Is it possible to use the results of the morphable model fittings to distinguish between meshes that represent faces and meshes that do not?**
Yes, more precisely it is possible to distinguish between three-dimensional meshes that represent faces which are part of the population of the faces that

the model can create and meshes that do not. However there can be some rare occurrences of false positives.

- **Can the meshes obtained from the Shroud be classified as faces after having analyzed the results of the morphable model fitting to them?**

No, the meshes obtained from the Shroud are more likely to be representing non-faces than to be representing faces, with the caveats discussed above.

8.2 Future work

In this project I used a morphable model which was not implemented with the Shroud of Turin in mind, but with a more generic purpose. I think that the results, especially the fittings, could be consistently improved by creating an ad hoc morphable model. Knowing that Jesus is believed to have been a bearded man from Galilee, it could be interesting to limit the capabilities of the morphable model by using, during its initialization, a set of faces obtained from middle-eastern bearded men. Not being able to recreate the beard was in fact a big limit, in that it substituted the beards with larger chins, creating less realistic faces.

The evaluation method could be improved as well. A more in depth study on the role of the different parameters in the fitting process, could give a better understanding on their importance, leading to a more precise way of comparing different fittings. In this project I simply excluded the parameters that were used for the coarse fitting and those that were not normally distributed in the real faces sample, however a finer distinction could be put in place. For example instead of binarily excluding or including the parameters, they could be assigned a weight, so that the parameters that have more influence during the fitting process would have more influence in the evaluation process as well. For this purpose, a study that did not have its main focus on the Shroud of Turin, but on the morphable model itself, could be an interesting option.

9. References

- Abate, Andrea F., et al. "2D and 3D face recognition: A survey." *Pattern recognition letters* 28.14 (2007): 1885-1906.
- Allen, Nicholas PL. "Verification of the nature and causes of the photo-negative images on the Shroud of Lirey-Chambery-Turin." *De Arte* (Pretoria) 51 (1995): 21-35.
- Baldacchini, Giuseppe, et al. "Coloring linens with excimer lasers to simulate the body image of the Turin Shroud." *Applied Optics* 47.9 (2008): 1278-1285.
- Blanz, Volker, and Thomas Vetter. "A morphable model for the synthesis of 3D faces." *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 1999.
- Bowyer, Kevin W., Kyong Chang, and Patrick Flynn. "A survey of approaches and challenges in 3D and multi-modal 3D+ 2D face recognition." *Computer vision and image understanding* 101.1 (2006): 1-15.
- Buades, Antoni, Bartomeu Coll, and J-M. Morel. "A non-local algorithm for image denoising." *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. Vol. 2. IEEE, 2005.
- Craig, Emily A., and Randall R. Bresee. "Image formation and the Shroud of Turin." *Journal of Imaging Science and Technology* 38.1 (1994): 59-67.
- De Liso, G. "Shroud-like experimental image formation during seismic activity." *Proceeding of the international workshop on the scientific approach to the acheiropoietos images*, ENEA, Frascati. 2010.
- Di Lazzaro, Paolo, Daniele Murra, and Barrie Schwartz. "Pattern recognition after image processing of low-contrast images, the case of the Shroud of Turin." *Pattern Recognition* 46.7 (2013): 1964-1970.
- Delfino, Vittorio Pesce. *E l'uomo creò la Sindone*. Vol. 233. Edizioni Dedalo, 2000.
- Doshi, Vikram. "3D Faces - A comparative study on 2D to 3D and 3D to 3D reconstruction of faces." *Utrecht University*, 2013.
- Fanti, Giulio, and Roberto Maggiolo. "The double superficiality of the frontal image of the Turin Shroud." *Journal of Optics A: Pure and Applied Optics* 6.6 (2004): 491.
- Fanti, Giulio, Francesco Lattarulo, and Oswald Scheuermann. "Body image formation hypotheses based on corona discharge." *Third Dallas International Conference on the Shroud of Turin*: Dallas, Texas. 2005.
- Fanti, Giulio. "Hypotheses regarding the formation of the body image on the Turin Shroud. A critical compendium." *Journal of Imaging Science and Technology* 55.6 (2011): 60507-1.
- Fanti, Giulio, et al. "Non-destructive dating of ancient flax textiles by means of vibrational spectroscopy." *Vibrational Spectroscopy* 67 (2013): 61-70.
- Fazio, G., Y. Clement, and G. Mandaglio. "The mysterious coexistence of bloodstains and body image on the Shroud of Turin explained by a stochastic process." *Med Archaeol and Archaeometry* 14 (2014): 20-23.
- Fazio, Giovanni, and Giuseppe Mandaglio. "Stochastic effects and corona discharge for the Shroud body image generation." *The Journal of The Textile Institute* 106.8 (2015): 904-906.

- Freer-Waters, Rachel A., and A. J. Timothy Jull. "Investigating a dated piece of the Shroud of Turin." *Radiocarbon* 52.4 (2010): 1521.
- Garlaschelli, Luigi. "Life-size Reproduction of the Shroud of Turin and its Image." *Journal of Imaging Science and Technology* 54.4 (2010): 40301-1.
- Hedgesi, R. Housle, et al. "Radiocarbon dating of the Shroud of Turin." *Nature* 337 (1989): 611.
- Jackson, John P., Eric J. Jumper, and William R. Ercoline. "Correlation of image intensity on the Turin Shroud with the 3-D structure of a human body shape." *Applied Optics* 23.14 (1984): 2244-2270.
- Jackson, John. "The vertical alignment of the frontal image." *Shroud Spectrum International* 32.33 (1989): 3-26.
- Latendresse, Mario. "Critical Evaluation of the Paper by Fazio et al., 'The mysterious coexistence of bloodstains and body image on the Shroud of Turin explained by a stochastic process'." *Med Archaeol and Archaeometry* 14 (2014): 367-373
- Levine, Martin D., and Yingfeng Chris Yu. "State-of-the-art of 3D facial reconstruction methods for face recognition based on a single 2D training image per person." *Pattern Recognition Letters* 30.10 (2009): 908-913.
- McCrone, Walter C. *Judgement day for the Turin Shroud*. Microscope Publications, 1997.
- Meacham, William. "Radiocarbon measurement and the age of the Turin Shroud: possibilities and uncertainties." *Shroud Spectrum International* 19 (1986).
- Mori, Masahiro, Karl F. MacDorman, and Norri Kageki. "The uncanny valley [from the field]." *Robotics & Automation Magazine, IEEE* 19.2 (2012): 98-100.
- Nickell, Joe. *Inquest on the Shroud of Turin*. Prometheus Books, 1987.
- Paysan, Pascal, et al. "A 3D face model for pose and illumination invariant face recognition." *Advanced video and signal based surveillance, 2009. AVSS'09. Sixth IEEE International Conference on. IEEE, 2009.*
- Perlin, Ken. "An image synthesizer". *Proceedings of the 12th annual conference on Computer graphics and interactive techniques (SIGGRAPH '85)*. ACM, New York, NY, USA (1985): 287-296.
- Rinaudo, J. B. "Image formation on the Shroud of Turin explained by a protonic model affecting radiocarbon dating." *III Congresso internazionale di studi sulla Sindone, Torino*. 1998.
- Rogers, Raymond N. "Studies on the radiocarbon sample from the Shroud of Turin." *Thermochimica Acta* 425.1 (2005): 189-194.
- Rogers, Raymond N. "A Chemist's Perspective on the Shroud of Turin." Lulu. com, Lulu Enterprises, Inc (2008).
- Sarkar, Sudeep. "USF HumanID 3D Face Database". University of South Florida. 2005.
- Schwalbe, Larry A., and R. N. Rogers. "Physics and chemistry of the shroud of turin: A summary of the 1978 investigation" *Analytica Chimica Acta* 135.1 (1982): 3-49.
- Schumacher, Peter M. - "Photogrammetric Responses from the Shroud of Turin" retrieved from <http://www.shroud.com/pdfs/schumchr.pdf> (1999).
- Siefker, Bob, Keith Propp, Dave Fornof, Ares Koumis, Rebecca Jackson, John Jackson. "The Shroud: A Critical Summary of Observations, Data and Hypothesis." *The Shroud Center of Colorado*. 2015.

Tamburelli, Giovanni. "Some results in the processing of the holy Shroud of Turin." *Pattern Analysis and Machine Intelligence*, IEEE Transactions on 6 (1981): 670-676.

ter Haar, Frank B. "Reconstruction and analysis of shapes from 3d scans". Utrecht University. Doctoral dissertation, 2009.

van der Weerd, Jacob C. "Fitting 3D Morphable Models on 2D face images using the Graphics Processing Unit." Utrecht University. Master thesis, 2011.

Volckringer, Jean, and Rex Morgan. "The Holy Shroud: science confronts the imprints." Runciman Press, 1991.

A. Appendix

A.1. Fittings parameters

<i>shroud_faces_parameters.csv</i>												
dhl_filter_cloth	1.3834	0.281836	-0.699023	0.618164	-1.44199	1.43027	-0.517383	1.1877	0.970899	1.0752	-0.29707	0.0216798
dhl_filter	-0.674414	1.34473	1.40215	-1.79707	0.817383	1.35293	1.50996	0.440039	1.61895	-0.488086	-1.79121	0.999023
enrie_filter_cloth	0.967383	1.36465	-0.0380859	-0.738867	-1.52051	0.560742	-0.120117	1.37285	0.808008	0.536133	-0.113086	1.0377
enrie_filter	-1.38809	1.6752	0.0931641	-1.79707	0.123633	1.79707	1.7209	1.66348	1.46074	-0.941602	-1.7959	1.01426

0.217383, -0.80918, -1.79824, -1.70918, 1.28262, 1.6248, -0.941601, 1.66816, 0.843164, -0.638086, -1.79824, 0.556055, -0.335742, 0.236133, 1.61543, -1.76074, 1.60254, -1.79824, -1.57324, 1.56035, -1.31074, 0.995508, 1.37754, -1.60137, -1.5123, -1.62949, -1.30137, -0.264258, -1.73965, 0.358008, 1.43496, 1.64707, 1.1291, 1.00137, 0.614649, -0.609961, -0.95332, 1.21582, 0.802149, -0.853711, 1.74434, -0.648633, 0.111914, 1.4373, 0.258399, -1.78652, 1.63535, -1.6459, -1.19824, -1.60137, -1.00605, -1.41621, 0.0732423, -1.66582, -0.864258, -0.223242, -1.75137, -0.423633, -1.32012, -0.77168, 1.0541, 1.20293, -0.0298827, 1.16074, -1.5334, 0.197461, 1.62949, -1.27559, 0.0919923, 0.508008, -0.0287108, 0.288867, 1.76074, -1.29785, -1.68457, 0.795117, -1.0377, -0.0240233, -0.700195, 0.697852, 1.77949, -1.09277, -0.832617, -1.60605, -0.655664, 0.94043, 0.42832, -0.0896483, -1.56152, 0.710742, 0.961524
miller_filter_cloth, 1.29902, 1.0248, -0.181055, 0.316992, -0.442383, -0.17168, 0.445899, 1.64355, 1.14316, 1.29902, 0.267774, -0.689648, -0.373242, 0.916992, -1.30137, -0.593555, 0.573633, -1.1584, -1.04473, 0.17168, 0.0873047, -0.39668, -0.836133, 0.241992, -0.901758, -1.25801, -0.575976, -0.72832, -1.15137, 1.6834, -0.20918, 1.23574, 1.66699, 0.829102, 0.676758, -0.257226, -1.14902, -0.722461, 1.40684, -0.311133, -1.04824, 1.38574, -0.799804, -1.41387, 1.5334, -1.23691, 1.79824, -1.75957, -1.67402, 1.64473, 1.02949, 1.75137, 1.7127, 1.44785, -0.266602, -1.0248, -1.08223, -0.970898, 1.63887, -1.79824, -1.61074, -0.444726, 0.260742, -0.182226, -1.71855, 0.898242, -0.823242, -0.363867, 0.275977, -0.710742, -0.898242, 0.595899, -1.64824, -1.49824, 0.260742, -1.79824, 1.12324, 1.46074, -1.46074, -1.34941, 1.55215, 0.981446, -1.56152, 1.49824, 1.64824, 0.818555, -1.68574, 1.6752, 1.1209, -1.42207, 0.749414, 0.49043, 1.79824, -0.0662109, 1.20645, 1.77012, 1.34824, -1.6002, 1.79355
miller_filter, 1.22988, 1.40449, -1.44785, -1.76074, -0.612305, 1.2041, 1.7959, -1.2627, -0.556055, -1.68574, -1.79824, -0.954492, 1.19707, -0.451758, -1.45371, 1.1748, 1.42324, 1.09277, -1.79824, 1.49824, 1.56152, -1.2416, 1.49824, -1.1584, -0.679101, -1.7877, -0.979101, 1.26387, -1.79707, 1.75137, 0.0017579, 0.55957, -1.72324, -1.34824, -1.62949, -1.42324, -1.6459, -1.4748, 1.70918, 0.999024, 1.1502, 0.973242, -0.764648, -1.18535, 0.341602, -1.79707, 1.49824, -1.64824, 0.949805, 1.49824, -0.354492, 0.823242, 0.598242, 0.308789, -1.08105, -0.298242, 0.0580079, -0.298242, -0.695508, -1.74902, -0.226758, -1.77949, 0.97793, 1.10098, -1.18301, -0.244336, -0.407227, -1.1373, 1.62949, 0.505664, -0.663867, 0.138867, -0.598242, 0.0369141, 1.49824, -0.823242, 0.324024, 1.49707, 0.151758, 0.450586, -0.388476, 0.0779298, 0.499805, -0.170508, -0.0720702, -0.199805, -1.4209, 0.438867, 1.26387, -0.143555, -1.79824, -0.176367, -1.47949, 0.161133, 1.49355, -1.57793, -0.0134765, -1.59668, 1.61543
miller, 1.30488, 1.75605, -1.78535, -1.77949, -0.497461, -0.700195, 0.445899, 0.598242, -0.00175781, 1.43965, -1.79824, -1.75605, 0.704883, 1.19121, -0.983789, 0.360352, 1.46074, -1.72207, -1.6459, 1.32949, 1.37168, -0.993164, 0.0298829, -1.79355, 1.79355, -1.40449, -1.67637, -0.0521483, -1.76074, -0.0955078, -0.12832, -1.1877, -0.751758, -1.67871, -1.7373, -1.65293, -1.64824, -1.39395, 0.592383, 1.71855, -0.257226, 1.26504, -0.632226, -0.223242, 0.500977, -0.48457, 1.79824, -1.70449, 1.76074, 0.984961, -0.72832, 1.40566, 1.6834, -1.64824, 0.635742, -1.47949, -0.95332, -1.45957, -0.110742, -1.26152, -0.45293, -0.963867, 1.77246, -1.61074, -1.79824, 0.611133, -0.288867, -1.14316, 0.689649, 0.437695, 0.403711, 1.72324, -1.79824, -1.71621, 1.56387, -0.895898, 1.64824, 1.6459, -1.79707, -1.34121, 0.748242, -0.0732421, -0.729492, 0.722461, 0.682617, -1.4748, 0.448242, 1.19824, 1.34824, -1.78887, -1.79824, 0.891211, -0.921679, -1.79824, 0.475195, 0.199805, 0.202149, -0.176367, 1.46895

noise_faces_parameters.csv

noise, 1.79824, 1.79824, -1.79824, -1.79824, -1.66934, 1.04941, 1.79824, -0.881836, 1.79121, -1.58613, -1.79824, -1.79824, 0.0591798, -0.305273, 0.824414, 1.7959, 1.79824, -1.77832, -1.79824, 1.79824, -0.901758, 1.10566, 1.77012, -1.79824, 1.79824, -1.79824, -1.79824, -0.0978516, -1.79824, 0.958008, 1.79824, -1.79824, -1.75254, -1.79121, 0.837305, 1.79824, -1.79824, -1.32012, 1.79824, 1.79824, -1.19707, 1.12793, -1.7959, 0.911133, -1.79355, -1.7959, 1.79824, -1.59902, 1.19004, 1.74082, 1.0166, -1.59668, 1.7959, -1.79824, 1.78066, -1.79824, -1.79004, -1.79355, -1.79121, -1.79824, 1.78066, -1.78887, 1.76543, -1.79824, -1.79824, 1.74902, -1.78887, 1.53926, 0.819727, 1.78418, 1.7877, 1.79824, -1.79121, 1.20293, 1.79238, 1.7959, 1.14199, -1.47832, -1.79824, -0.907617, 1.7877, 1.7959, 1.76426, -0.466992, -1.47012, -1.50996, -1.71855, 1.79004, 1.61191, -1.78418, -1.52637, -1.44902, -1.64355, -1.79707, -0.813867, -1.76074, -1.78535, 1.4502, -1.63418
noise_mask_cloth, 1.37871, 1.53223, -1.74668, -1.44551, -1.6166, -1.56738, 1.5873, 1.12207, -0.360352, -0.319336, -1.42207, -1.05879, -0.106055, 0.288867, 1.59668, 0.997852, 0.629883, -1.46074, -1.19941, 1.10801, -


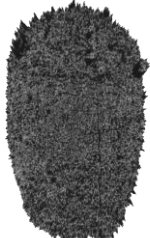
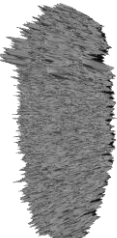

1.24043, -1.1959, 1.66582, 1.38926, 1.41387, -0.989648, -1.16191, 0.84668, -
1.7373, 0.440039, 1.17012, -0.319336, -0.329883, -1.24863, 0.0626954, -
0.198633, -1.31074, -1.37754, 1.6166, -0.652148, -0.368555, 0.0345704, -
0.440039, 1.68457, -0.237305, 0.135352, 0.270117, 0.13418, 1.1373, 0.10957,
-0.0474609, 0.307617, 0.433008, -0.557227, -0.893555, -0.157617, 0.118945, -
0.363867, 0.0134767, -0.516211, 0.463477, -0.558398, -0.518555, -0.145898, -
0.127148, -0.499805, -1.30488, -0.305273, -0.219727, 0.577149, 1.47949, -
0.0205077, -0.45293, -0.300586, 0.0462892, -0.0333983, 0.220899, -0.670898,
0.19043, 0.274805, 0.558398, 0.143555, -0.291211, -0.205664, -0.172851, -
0.0697265, -0.54082, 0.12832, 0.406055, 0.24082, 0.168164, 0.0521485, -
0.0263671, -0.25957, -0.185742, -0.182227, 1.65762, -0.000585848, -0.156445
noise_mask, 1.78887, 1.74199, -1.62949, -1.68105, -1.3248, 0.172852,
1.65762, 1.53574, 1.77949, -0.489258, -1.77012, -1.78887, -0.813867,
0.729492, 0.223242, -0.308789, 1.47949, -1.79824, -1.79824, 1.16426, -
0.860742, -0.373242, 1.53105, -1.61074, 1.61074, -1.72324, -1.49824,
1.70449, -1.79355, 0.710742, 1.28965, -1.73613, 0.418945, 0.331055,
0.151758, 1.1959, -1.79824, -1.69746, -1.15488, 1.49355, -1.78887, 0.208008,
-0.513867, 1.12324, -1.40801, 0.826758, 1.63887, -0.0474609, 0.40957,
0.295898, 0.251367, -0.495117, 1.79824, -1.77949, 1.08574, -1.57324, -
1.7584, -1.45137, -1.29082, -1.64707, -0.474023, -1.60137, 1.24395, -
1.64824, -1.64824, 1.1959, -1.48301, 1.09277, -1.23105, 0.327539, 0.279492,
-0.245508, -1.64824, -0.100195, 0.066211, 1.7584, 1.62949, -1.46777, -
1.79355, -1.5709, 1.60723, 0.281836, 0.0638673, -0.301758, -0.188086, -
0.436523, 0.43418, 0.456445, -0.598242, -1.60605, 0.311133, 0.0837892,
0.0720704, -1.22637, 0.430664, -0.832617, 0.59707, -0.226758, -0.37207
perlin, 1.77012, 1.79824, -1.79824, -1.79824, 0.376758, 1.7291, 1.79824,
1.55449, 1.79824, 1.79824, -1.79824, -1.79824, 0.162305, 1.79824, 1.64824,
1.39512, 1.79824, -0.459961, -1.79824, 1.65293, -0.298242, 0.43418, 1.79824,
-1.79824, 1.79824, -1.79824, -1.79824, 0.451758, -1.79824, -1.04473, 1.7959,
-1.79824, -1.70918, -1.78887, -0.598242, 0.598242, -1.79824, 1.78652, -
0.941602, 1.79824, -1.41504, -1.19238, 1.64824, -1.79824, -1.74082, -
1.78418, 1.79824, 1.35176, 1.79824, 1.77949, 1.7748, -1.15605, 1.79824, -
1.79824, 1.79824, -1.79824, -1.79238, -1.79824, -1.34824, -1.64824, 1.78184,
-1.49824, 1.77246, -1.79824, -1.79824, 1.77246, -1.56621, 1.14668, 1.76543,
1.7959, -1.65293, 1.78535, -1.79824, -1.79824, 1.49121, -1.18887, 1.49707,
1.79355, -1.79824, -1.67637, 1.79824, -0.0146484, 1.72793, -1.42324, -
1.78184, -1.79824, -1.79824, 1.78887, 1.79473, -1.25449, 1.79824, -1.78887,
-1.17715, 1.77129, 1.36582, -1.28496, -1.75605, -1.79355, -1.46426
perlin_mask_cloth, -0.336914, 1.63652, 0.564258, -1.13965, -0.966211, -
0.15293, 1.61426, -0.62168, 1.7291, -1.36582, -1.45137, 0.656836, -0.827929,
1.69043, -1.6916, -1.77832, -1.72207, -1.79473, -1.75957, -0.650976,
0.73418, -1.65293, 0.723633, 1.45957, 1.79707, -1.38105, 0.18457, -0.982617,
0.393164, 0.995508, -1.79824, -1.1502, 1.00371, -0.579492, 1.79824, -
1.79707, -1.74199, -1.79824, -1.73965, -1.51934, -1.59199, 1.79824, -
1.79824, -1.79355, 1.79824, 1.7209, 1.75605, -1.79824, 0.301758, 1.79707, -
1.76074, 0.376758, -1.7209, -0.60293, -1.0377, -1.72324, -1.76074, -1.61543,
-1.79707, -1.79824, -1.79824, -1.79707, -1.79824, -1.79707, -1.79824, -
1.71738, -1.79824, 1.7959, -1.42324, -0.76582, -1.79824, 1.79824, -1.79824,
1.79824, -1.79473, 1.79355, -1.15254, 0.236133, 1.77832, -1.69863, 1.26152,
-1.79355, -1.75137, -1.32129, 1.5334, -1.79824, -1.7666, 1.79355, -1.79824,
-0.912305, -1.79355, -1.79824, -1.79121, 0.667383, 1.79473, 1.71973,
1.79824, 1.79824, 1.79824
perlin_mask, -1.6002, 1.30371, -1.40449, -1.79824, 0.742383, 1.37637,
1.54512, -1.03301, 1.64824, -1.48887, -1.77949, -1.72324, 1.56387, -1.29551,
-1.46426, 0.373242, -1.38926, -1.74199, -1.76074, -1.61074, -1.32949, -
0.0580077, 1.40449, -1.79824, -1.27324, -1.49824, 1.17949, 0.264258,
1.37051, 1.64824, 0.376758, -1.27324, -1.62949, 0.129492, 0.954492, -
0.973242, -1.14199, 0.151758, 1.38574, 1.53574, -1.32949, -0.770508,
1.25918, 1.29551, 1.47363, 0.0580079, 1.72324, -1.79824, 1.27676, 0.279492,
-1.79824, -1.23926, 0.804492, -0.298242, 1.08691, -1.32949, 0.751758, -
0.293555, -1.3459, 0.110742, 0.898242, -1.72324, -1.0834, 1.64238, -1.79824,
1.29668, -1.79824, 1.06582, 0.676758, -1.79824, -1.79824, 1.72324, -
0.223242, 0.523242, 1.6459, 1.79824, -1.79824, 1.72324, -1.79824, 1.79355,
1.19824, -1.19824, 0.901758, 0.785742, -1.64355, -1.6459, -1.79824,
0.579492, -1.79824, 0.756446, -1.64824, -0.748242, 0.0955079, -0.0041015,
0.148242, 0.973242, 0.541992, -0.595898, 0.138867
plane, 1.78066, 1.79707, -1.79824, -1.79824, 0.0837891, 1.79824, 1.79824, -









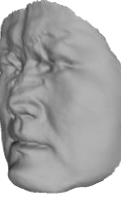



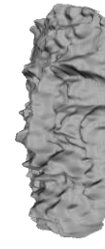


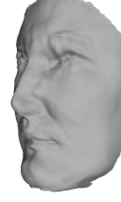





```






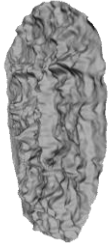






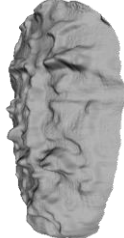



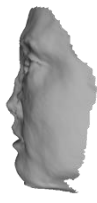



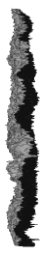
1.61191, 1.79824, -1.5627, -1.79824, -1.79824, 1.72324, -1.05176, -1.08926,
1.78887, 1.79824, -1.72324, -1.79824, 1.61074, 1.23926, 1.79824, 1.79824, -
1.79824, 1.49824, -1.49824, -1.79824, 1.64824, -1.64824, 1.48652, 1.79824, -
1.76074, -1.79824, -1.79824, 1.17598, 1.79824, -1.79824, 1.7959, 1.79824,
1.64824, 0.751758, 1.07051, 1.64824, -1.79824, -1.72324, -1.12324, 1.79824,
-1.79121, 1.79824, 1.19824, -0.863086, 1.79824, 1.79824, -1.79824, 1.79824,
-1.79824, 0.0017579, -1.79824, -1.79824, -1.79824, 1.79824, -1.79824,
1.79824, -1.79824, -1.79824, 1.19824, -1.65176, 1.79824, 1.49824, 1.49824, -
1.49824, 1.79824, -1.79824, -1.77949, 1.76074, -0.654492, 0.919336, 1.68457,
-1.70449, 1.79824, 1.79824, 0.812696, 1.68223, -1.49824, -1.79824, -1.79824,
-1.79824, -0.00175781, 1.79824, -1.49824, 1.79707, -1.79824, -1.79824, -
0.898242, -1.47598, -1.79824, 0.301758, 0.607617, 0.145899
plane_mask_cloth, 0.611133, 1.2873, 0.529102, -1.43848, -1.67051, 1.77949,
1.42441, 0.631055, 1.29785, -0.95332, -0.913477, 0.607617, -0.594727,
1.64707, -1.56387, -0.729492, -1.71035, -1.66816, -1.30957, 0.382617,
0.0498048, -1.74902, -0.804492, 1.40566, -1.10332, 0.342774, 1.36465, -
1.04941, -0.881836, 1.78652, -0.893555, -0.0345702, 1.79824, 0.970899,
1.7959, 0.962696, -1.78887, -0.945117, -1.75137, -0.00527344, -1.79824,
1.7373, -1.79355, -1.58848, 1.50645, -1.79824, 1.79238, -1.79824, -0.586523,
1.70332, -1.7877, 1.79355, 1.78887, 1.37988, -1.50176, 0.254883, -0.97793, -
1.64238, 0.382617, -1.79824, -1.72324, -1.32832, 1.60957, -1.79355, -
1.76074, 1.79824, -1.79824, -1.46895, 1.09512, -1.63418, -0.97207, 1.79355,
-1.79121, -1.78418, 0.00996104, 0.850195, -1.10918, -0.218555, -1.79355, -
1.38926, -0.641601, 0.579492, -0.204492, -1.47832, 1.79824, 1.11855, -
1.79707, 1.79824, -1.16777, -1.42559, -1.69277, 1.7373, 0.929883, 0.811524,
1.78418, -1.2123, 1.78887, -0.816211, 1.64824
plane_mask, -1.75723, 1.07637, -1.44316, -1.79824, 0.0404298, 1.60488,
1.7584, 1.50762, 1.63887, -1.62363, -1.78887, -1.71152, 0.404883, -1.03535,
-1.38926, -1.57324, 1.76074, -1.71387, -1.79824, 0.287695, -1.79121, -
0.877148, 1.05527, -1.78184, -1.42324, -1.62949, -1.14551, -1.31074, -
1.79824, 1.78887, 0.736524, -1.79824, -1.79824, 1.74434, 1.79824, -1.52637,
-1.79824, -1.7748, 1.74199, -1.72324, -1.6459, -1.21348, 1.79824, -
0.0486327, 1.64824, -1.19824, 1.79824, -1.79824, -1.73262, 1.05176, -
1.11855, -0.388476, 1.79824, -1.71387, -0.898242, 0.995508, -0.532617, -
1.57324, -1.79824, -1.42324, 1.61074, -1.79824, 0.911133, 0.301758, -
1.05645, 1.7959, -1.79824, 0.901758, 0.320508, -0.820898, -1.28145, 1.79824,
-1.46074, -0.639258, 0.0158204, 1.79824, -1.55449, 1.79824, -1.79824, -
1.35176, 1.79355, -0.543164, 1.64824, -1.79824, 1.79824, 0.901758, 1.46074,
-1.49824, -1.6459, -0.523242, -0.782226, -0.117773, 0.0767579, -0.595898,
1.79824, -1.34355, 1.31074, 1.33652, 0.223242













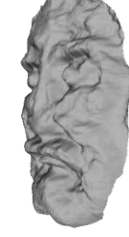




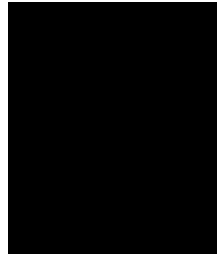



```




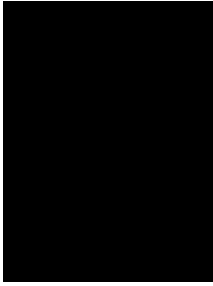
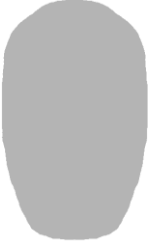



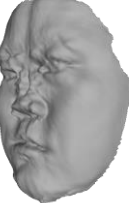

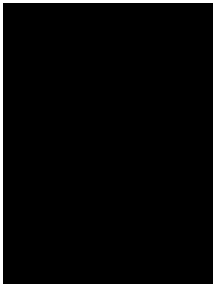
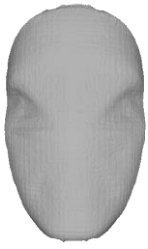





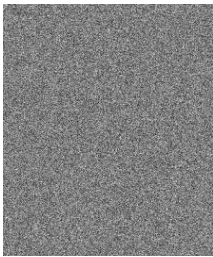
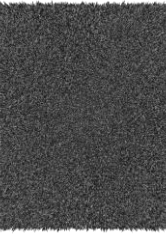
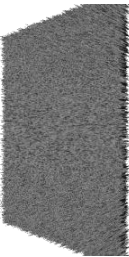

A.2. Fittings table


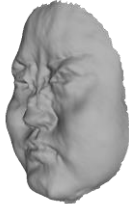

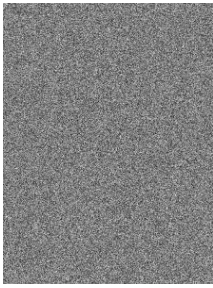

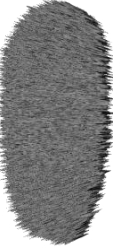




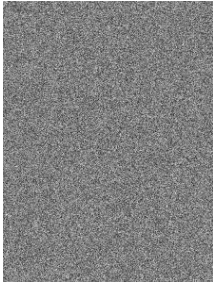

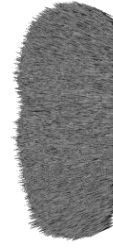




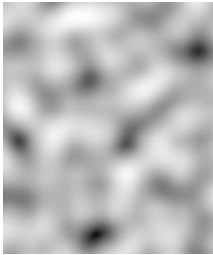
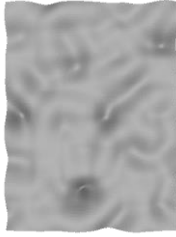


Name	Image	Mesh and fitting results		
<i>dhl</i>				




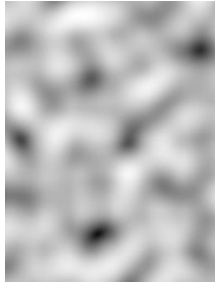
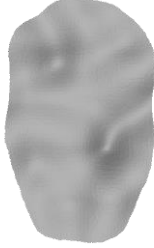




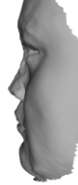
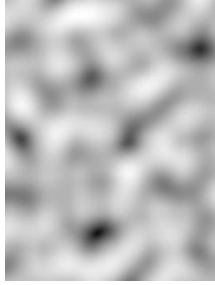
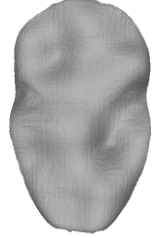
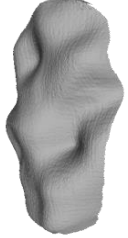


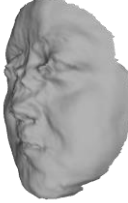
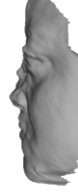
				
<i>dhl_filter</i>				
				
<i>dhl_filter_cloth</i>				
				
<i>enrie</i>				

				
<i>enrie_filter</i>				
				
<i>enrie_filter_cloth</i>				
				
<i>miller</i>				

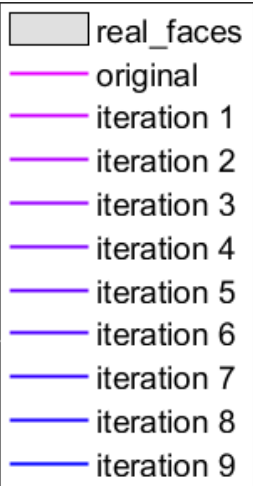
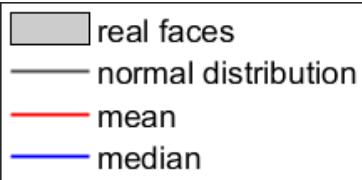
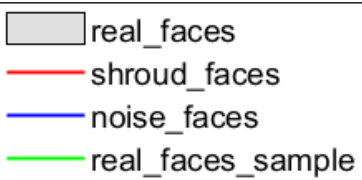
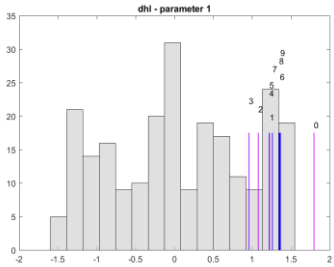
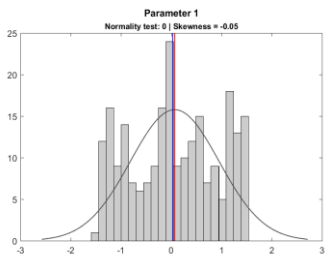
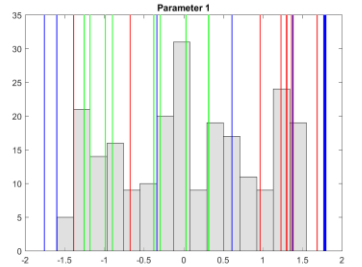
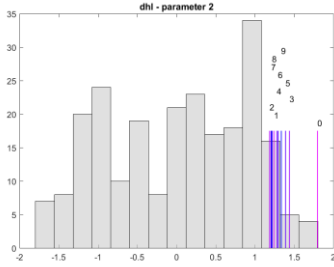
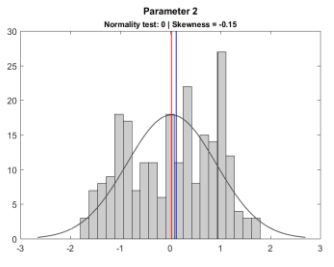
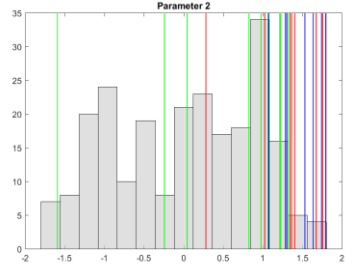
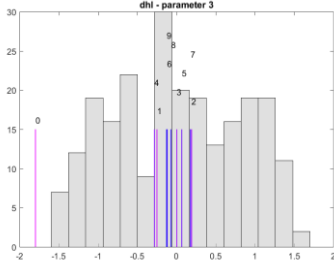
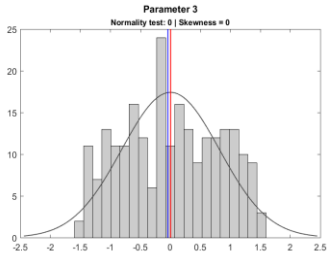
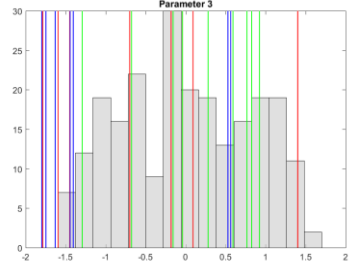
				
<i>miller_filter</i>				
				
<i>miller_filter_cloth h</i>				
				
<i>plane</i>				

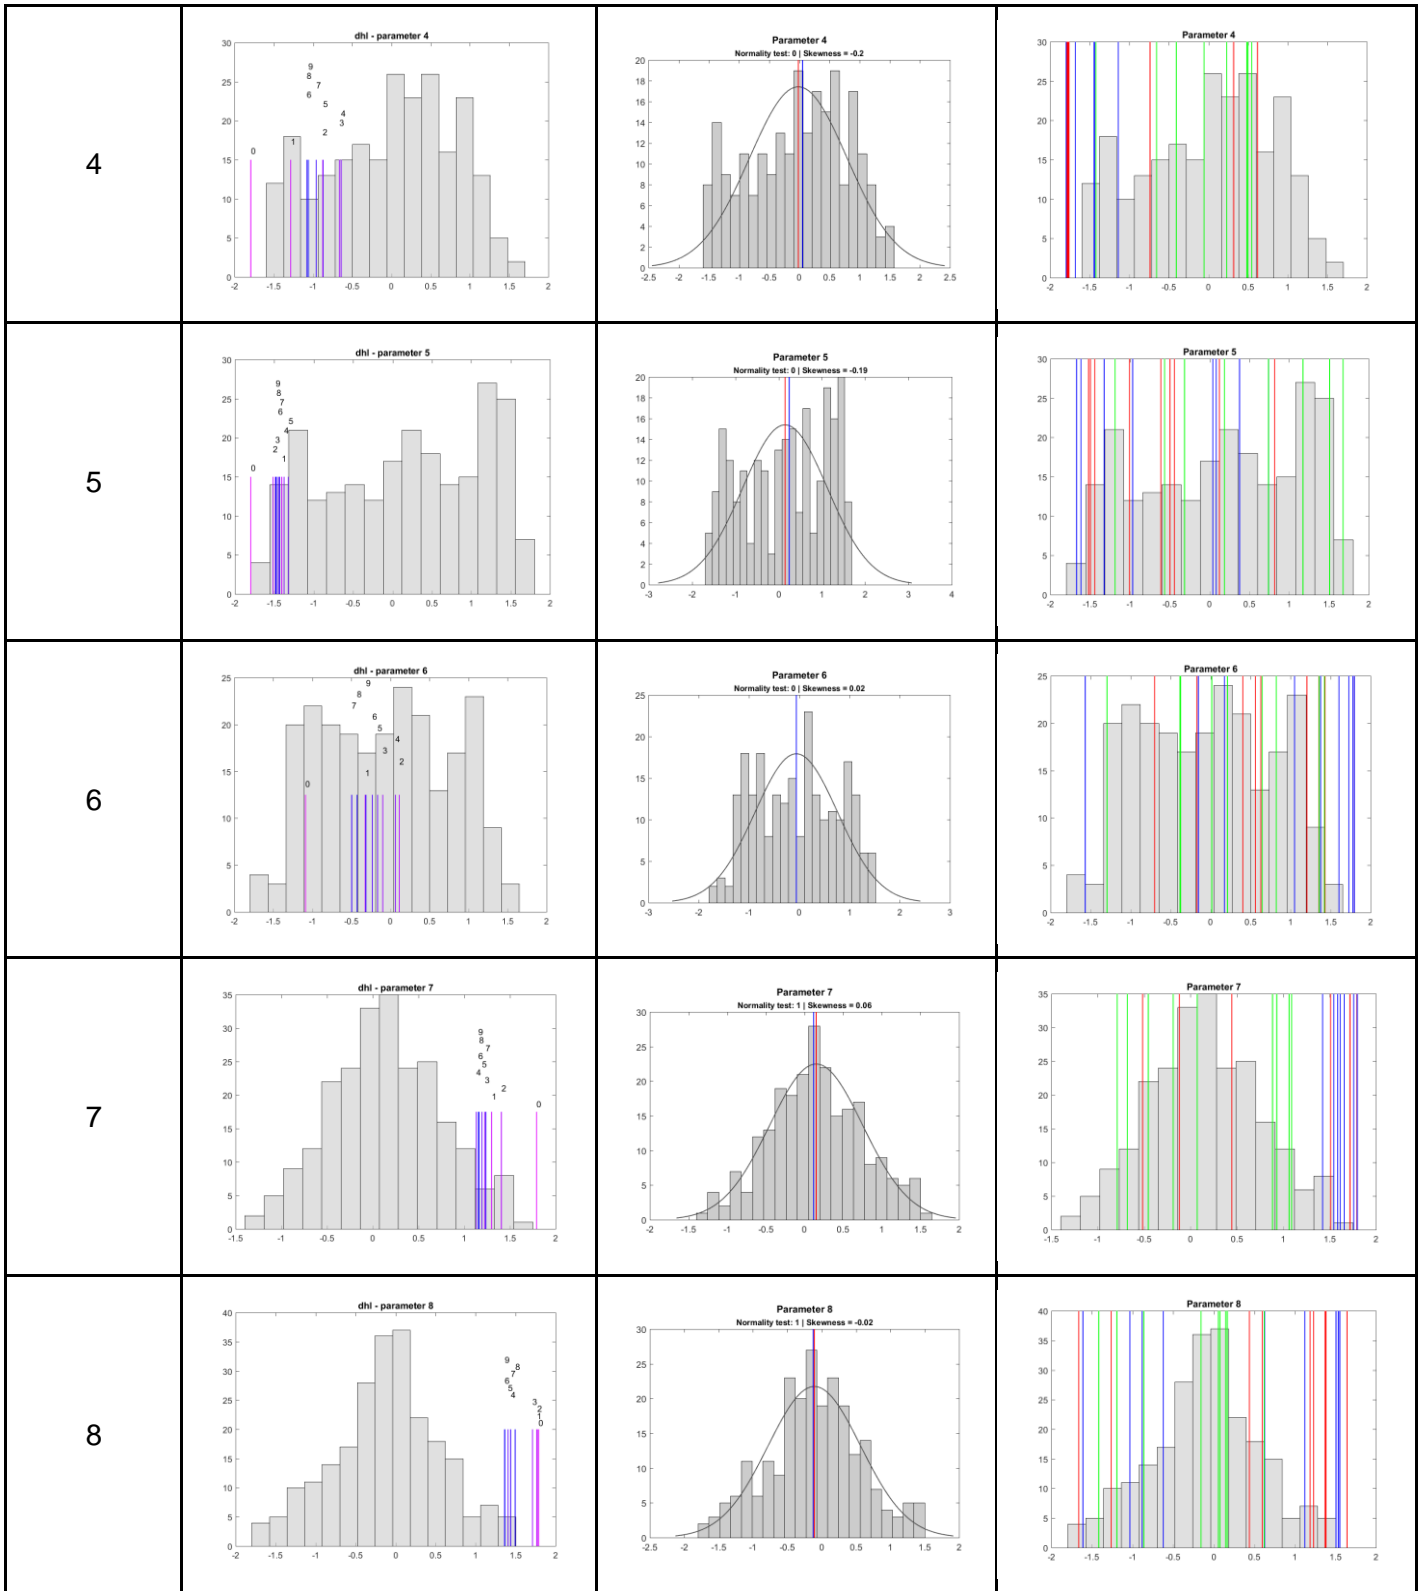
				
<i>plane_mask</i>				
				
<i>plane_mask_cloth</i>				
				
<i>noise</i>				

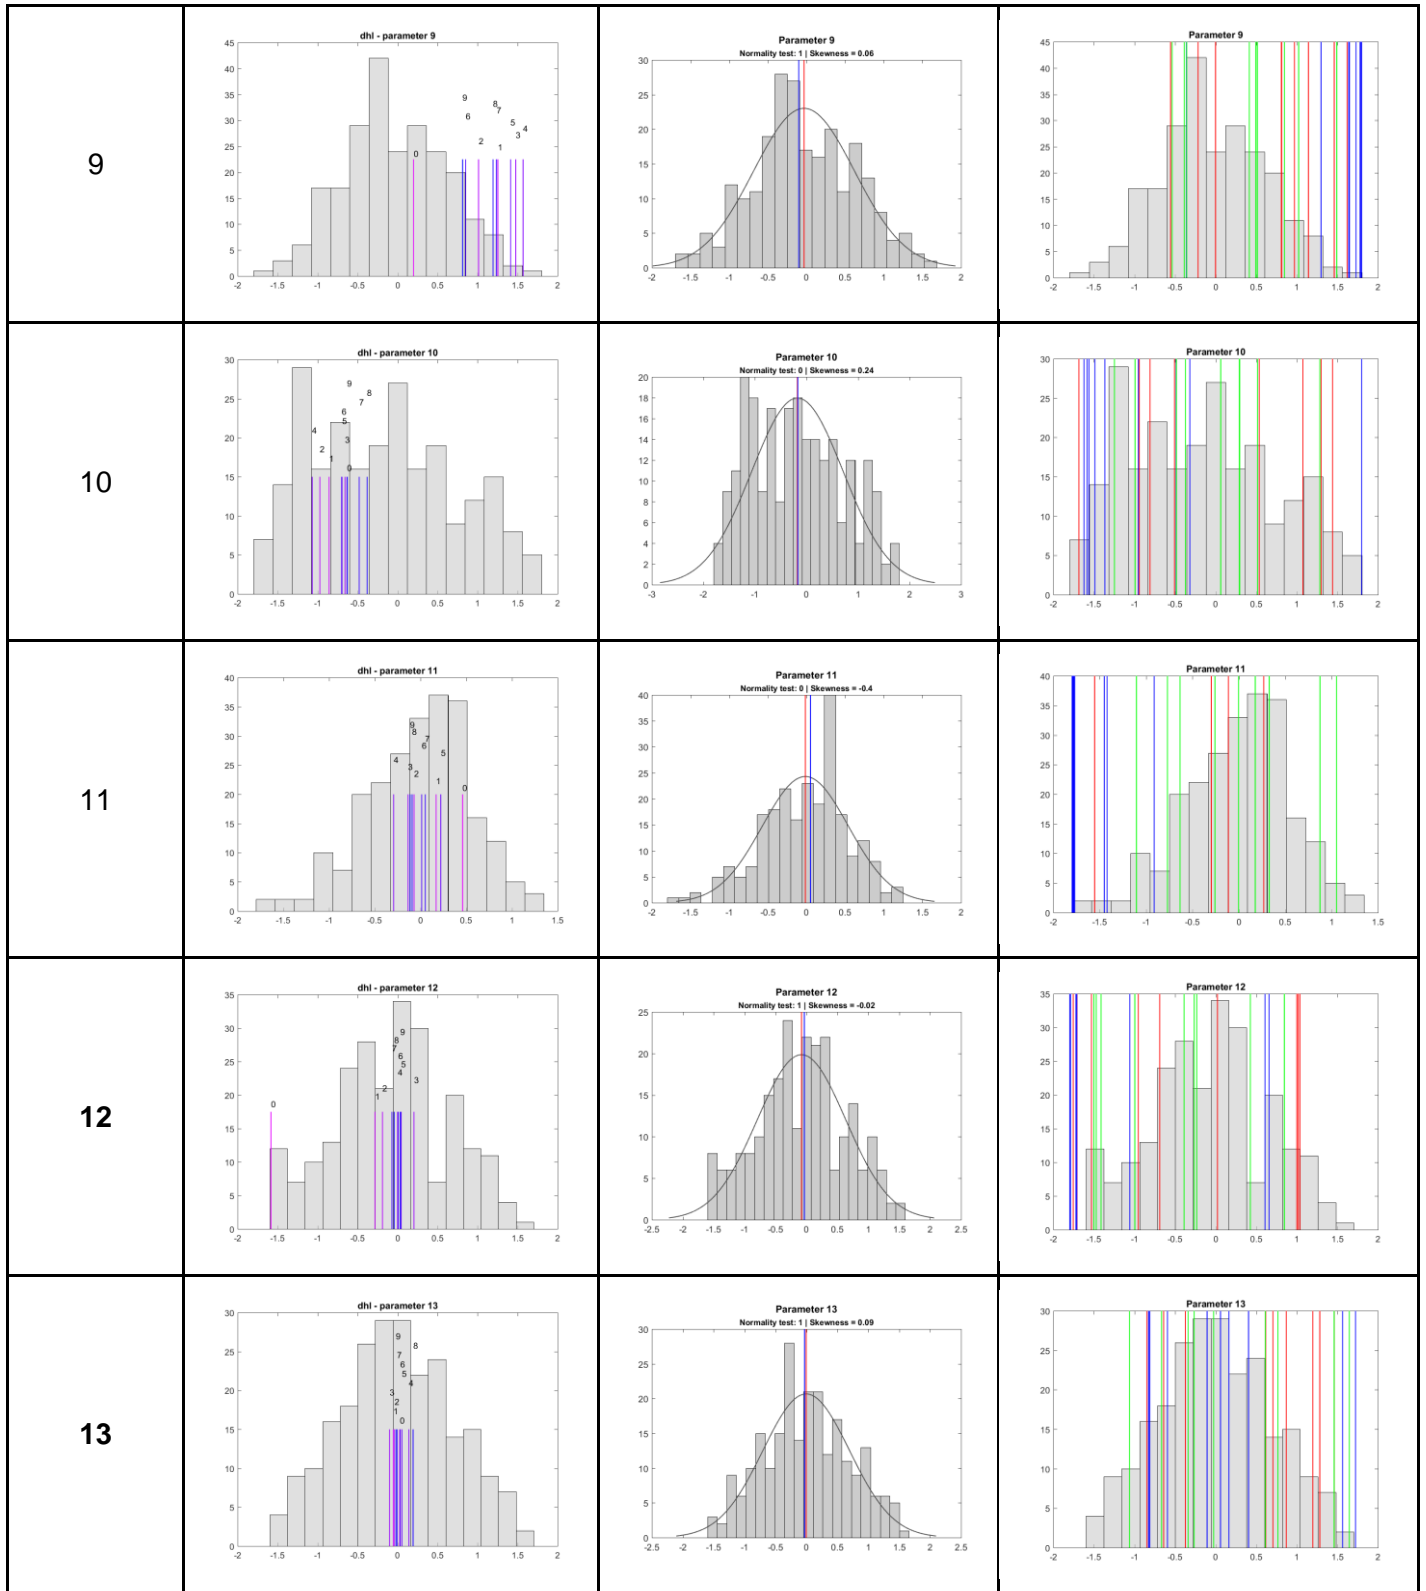
				
<i>noise_mask</i>				
				
<i>noise_mask_clot h</i>				
				
<i>perlin</i>				

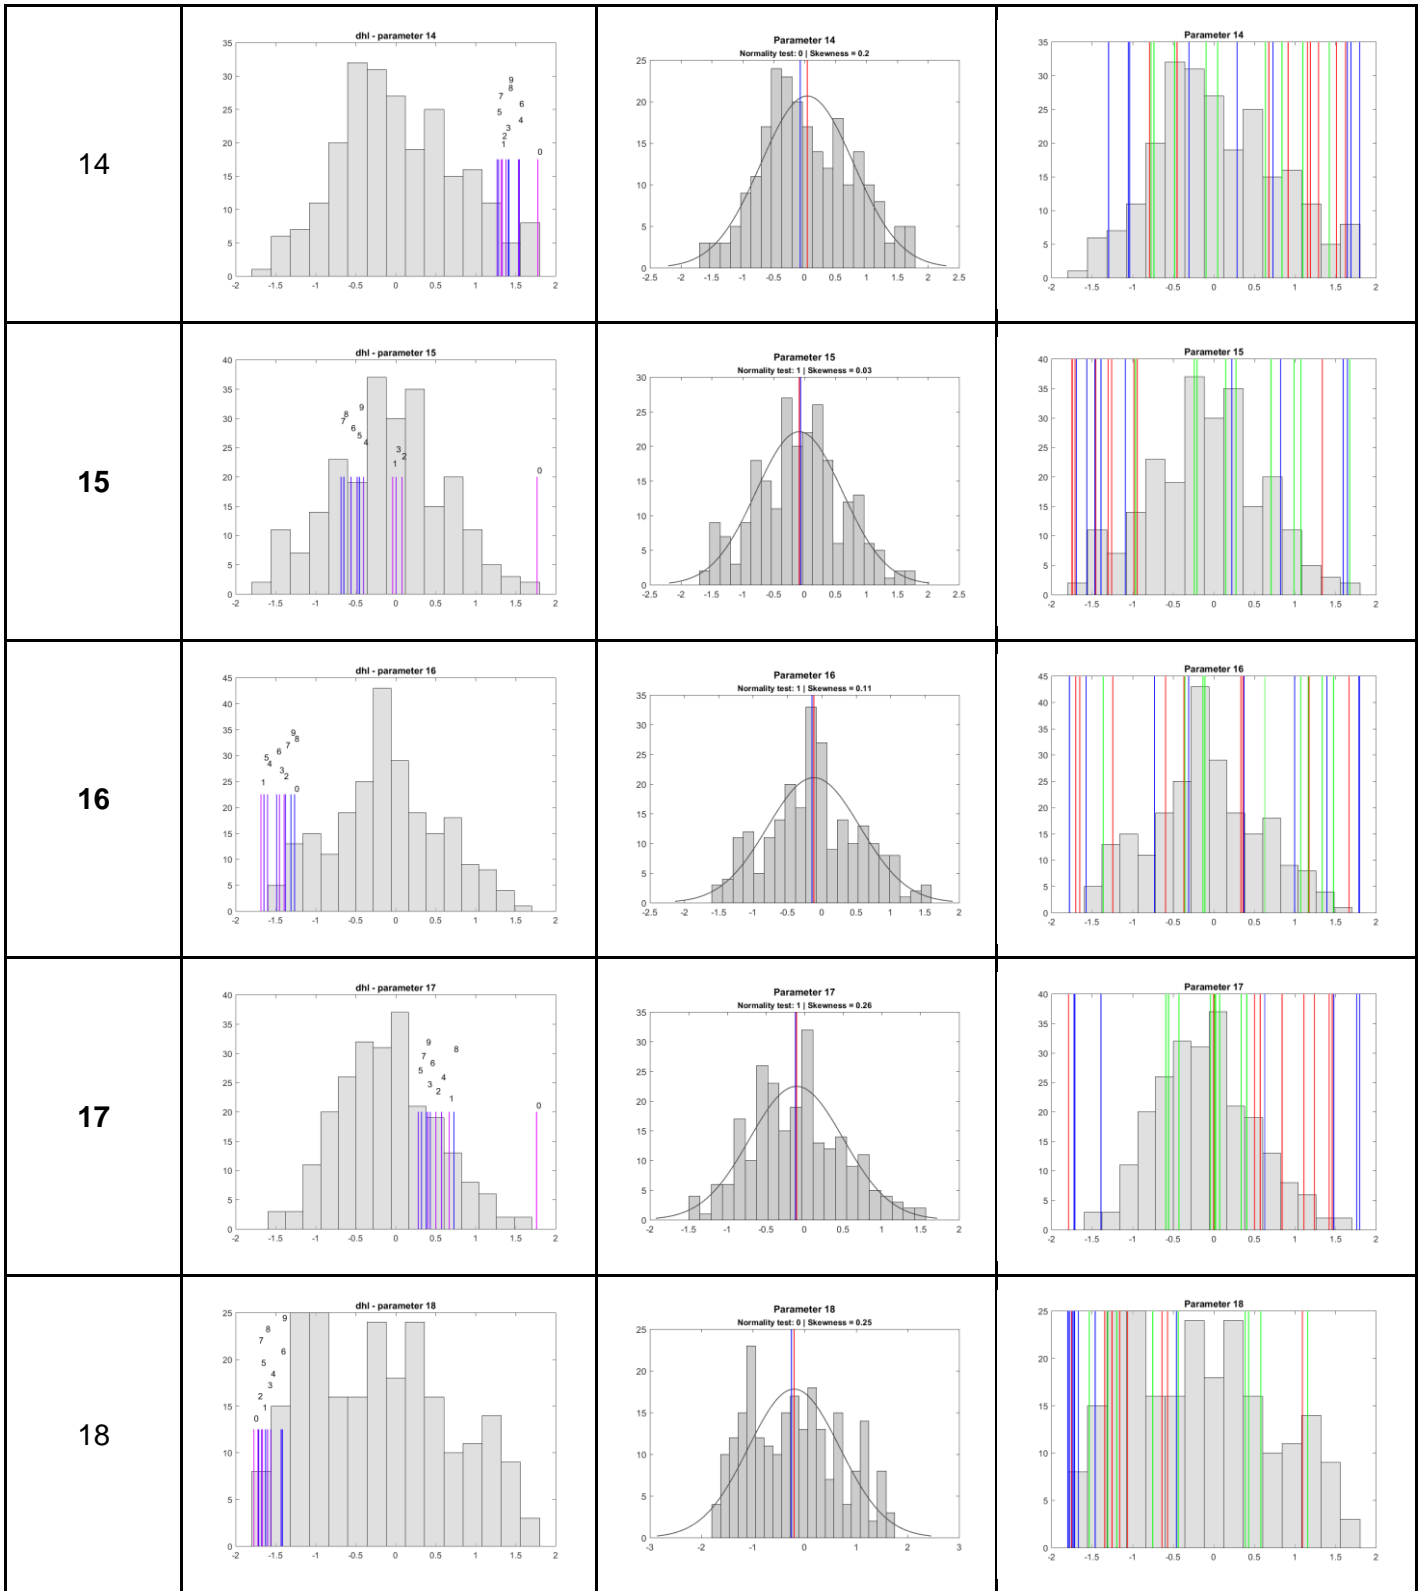
				
<i>perlin_mask</i>				
				
<i>perlin_mask_cloth</i>				
				

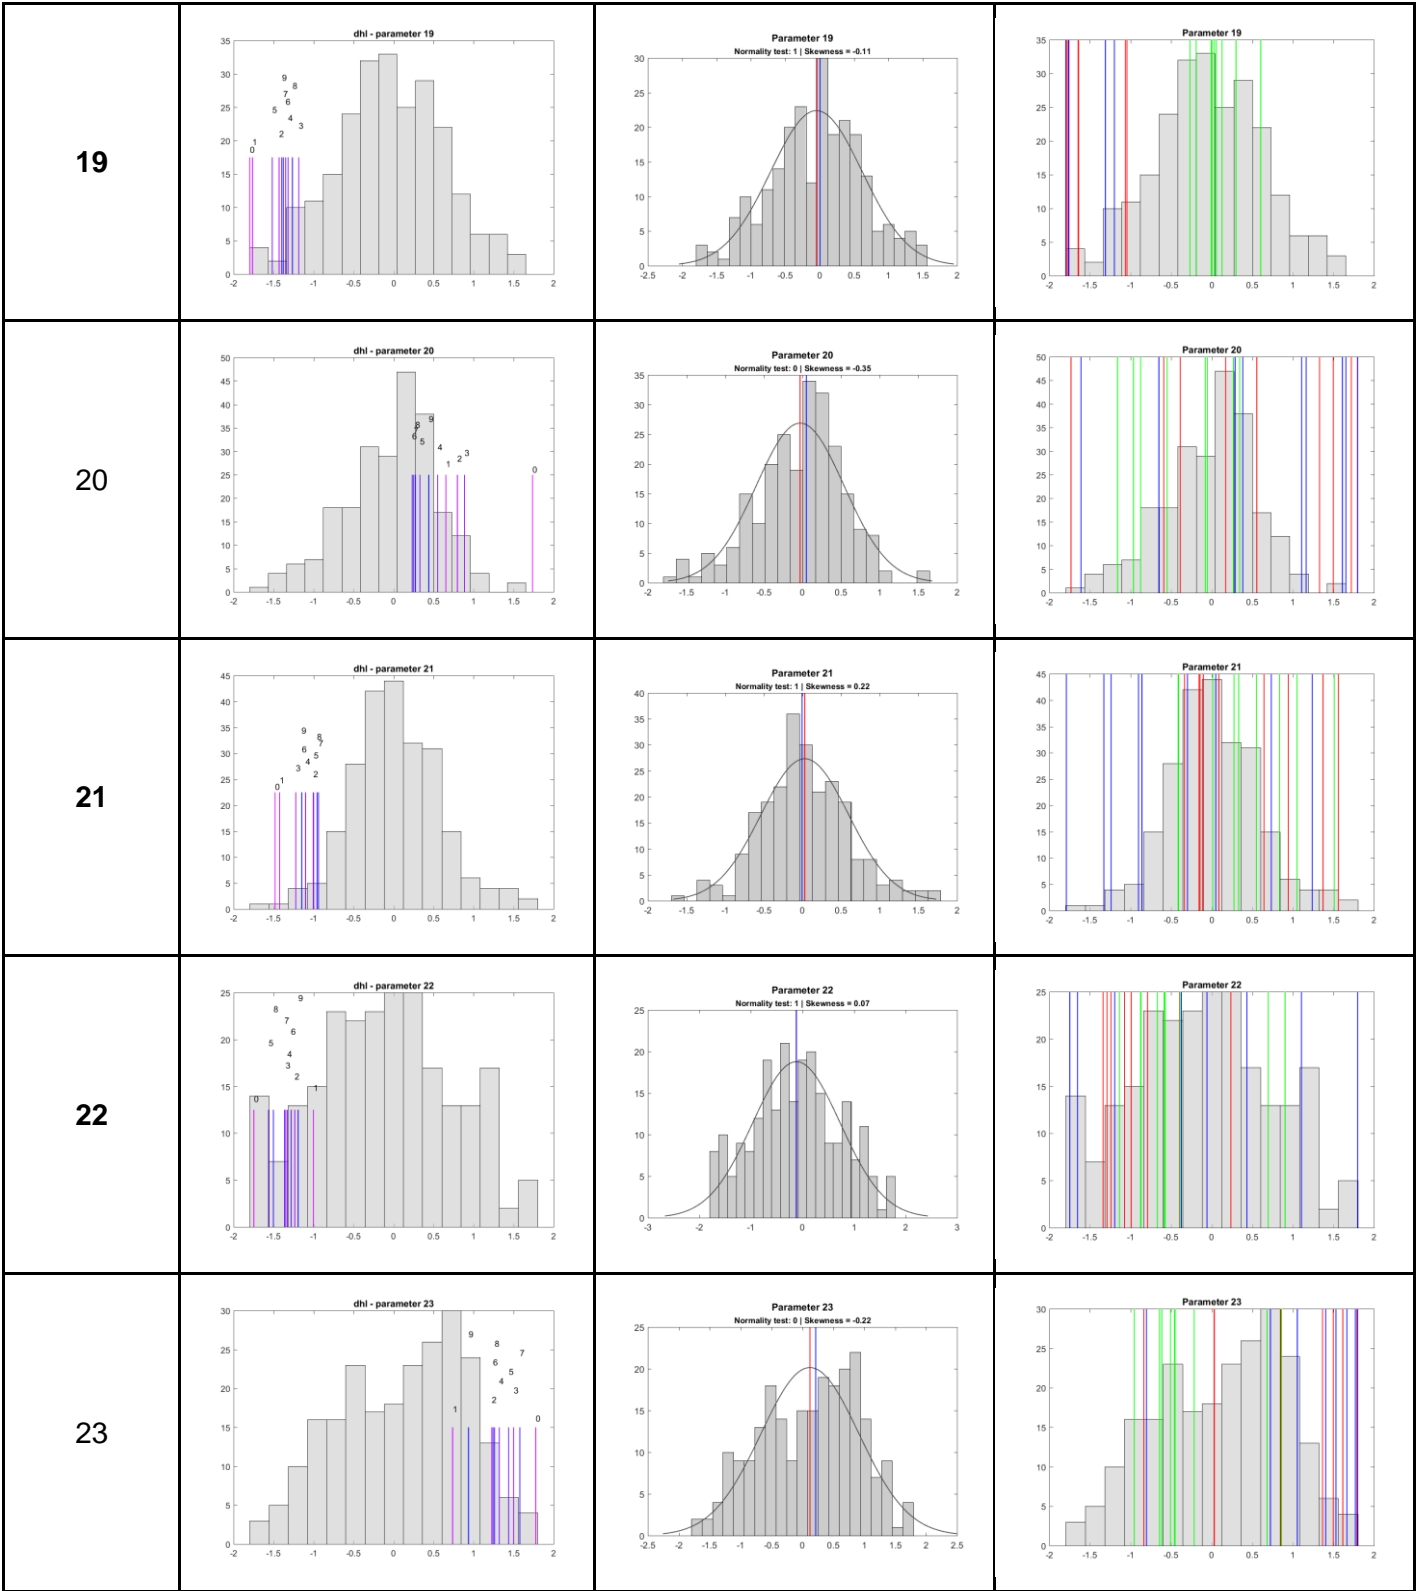
A.2. Parameters plots

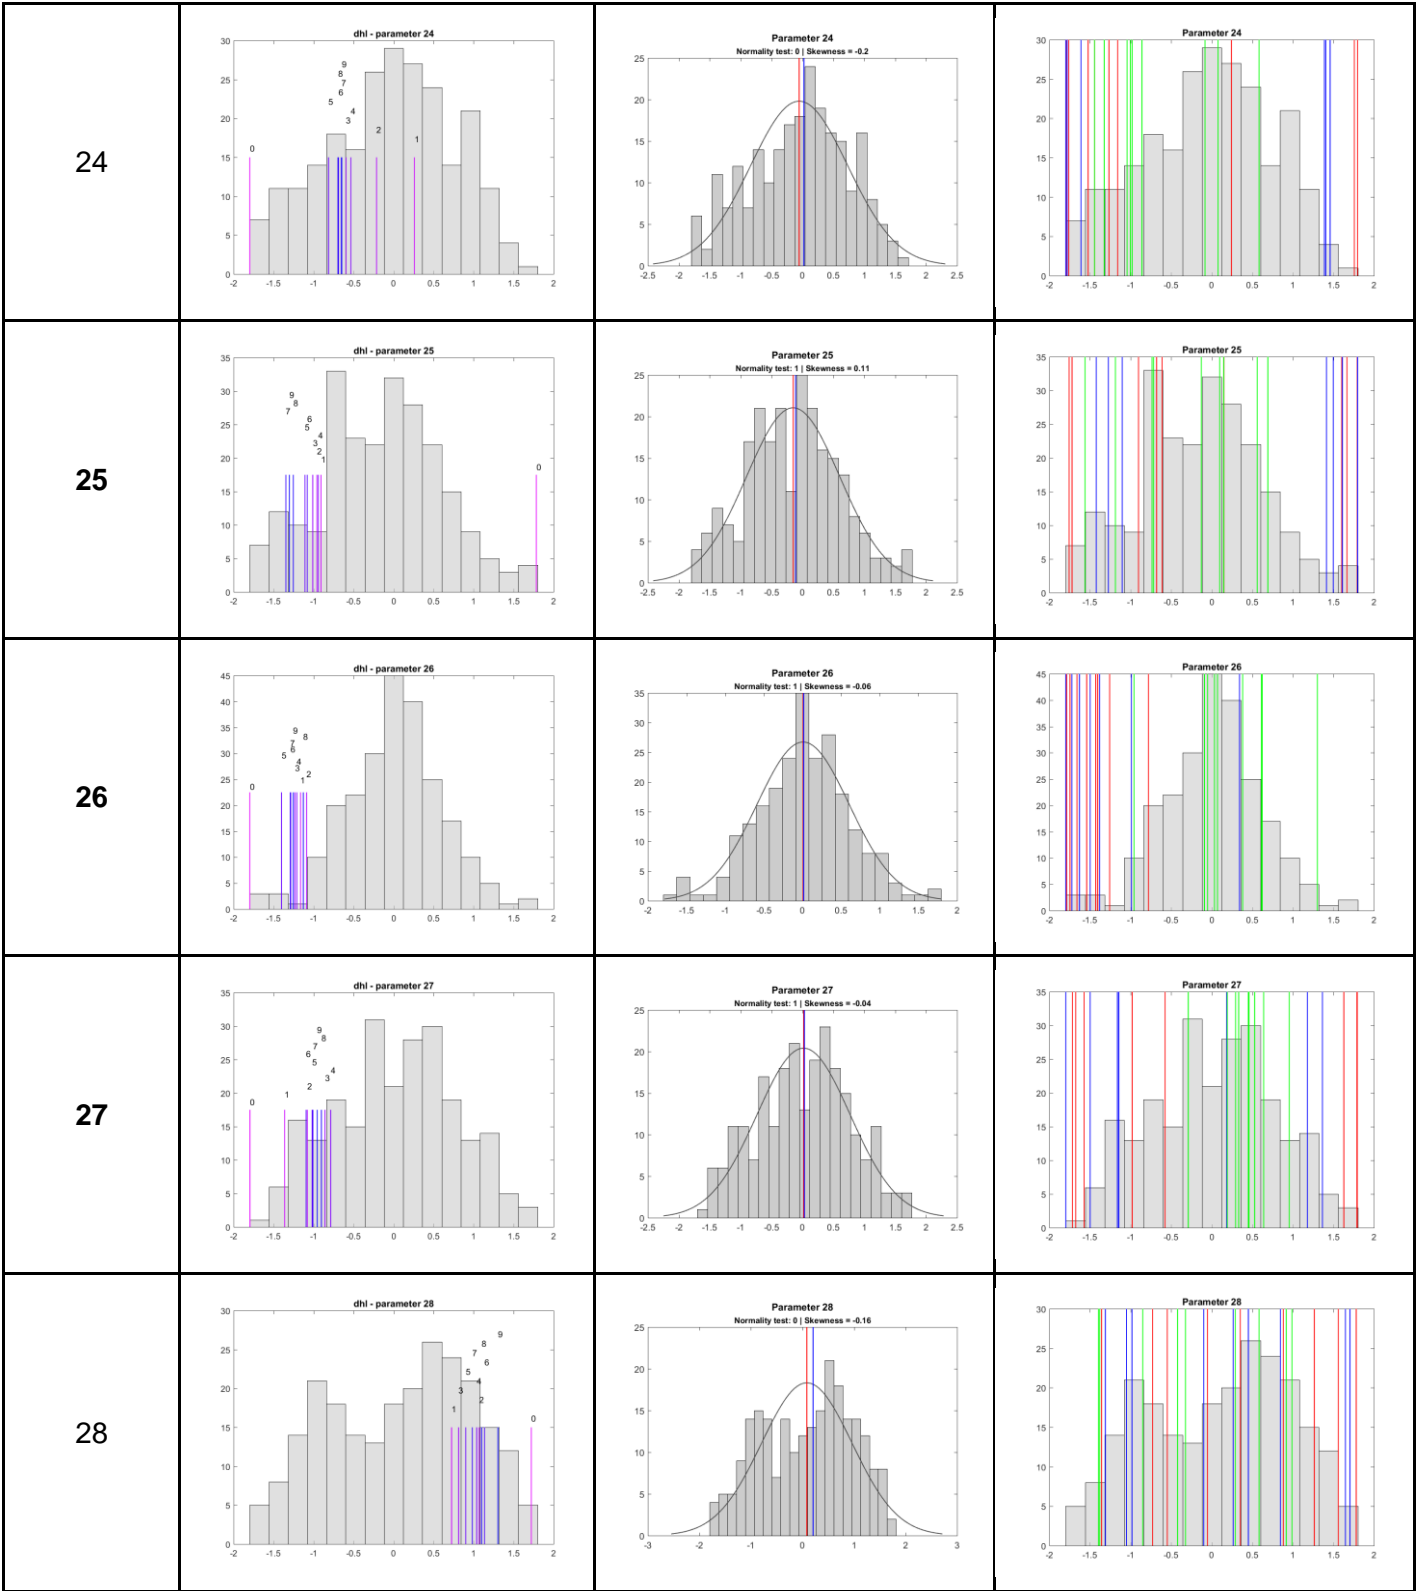
Parameter nr. Bold indicates a valid parameter	Convergence  The original fitting uses <i>dhl.jpg</i> as starting image	Normality 	Comparison 
1			
2			
3			

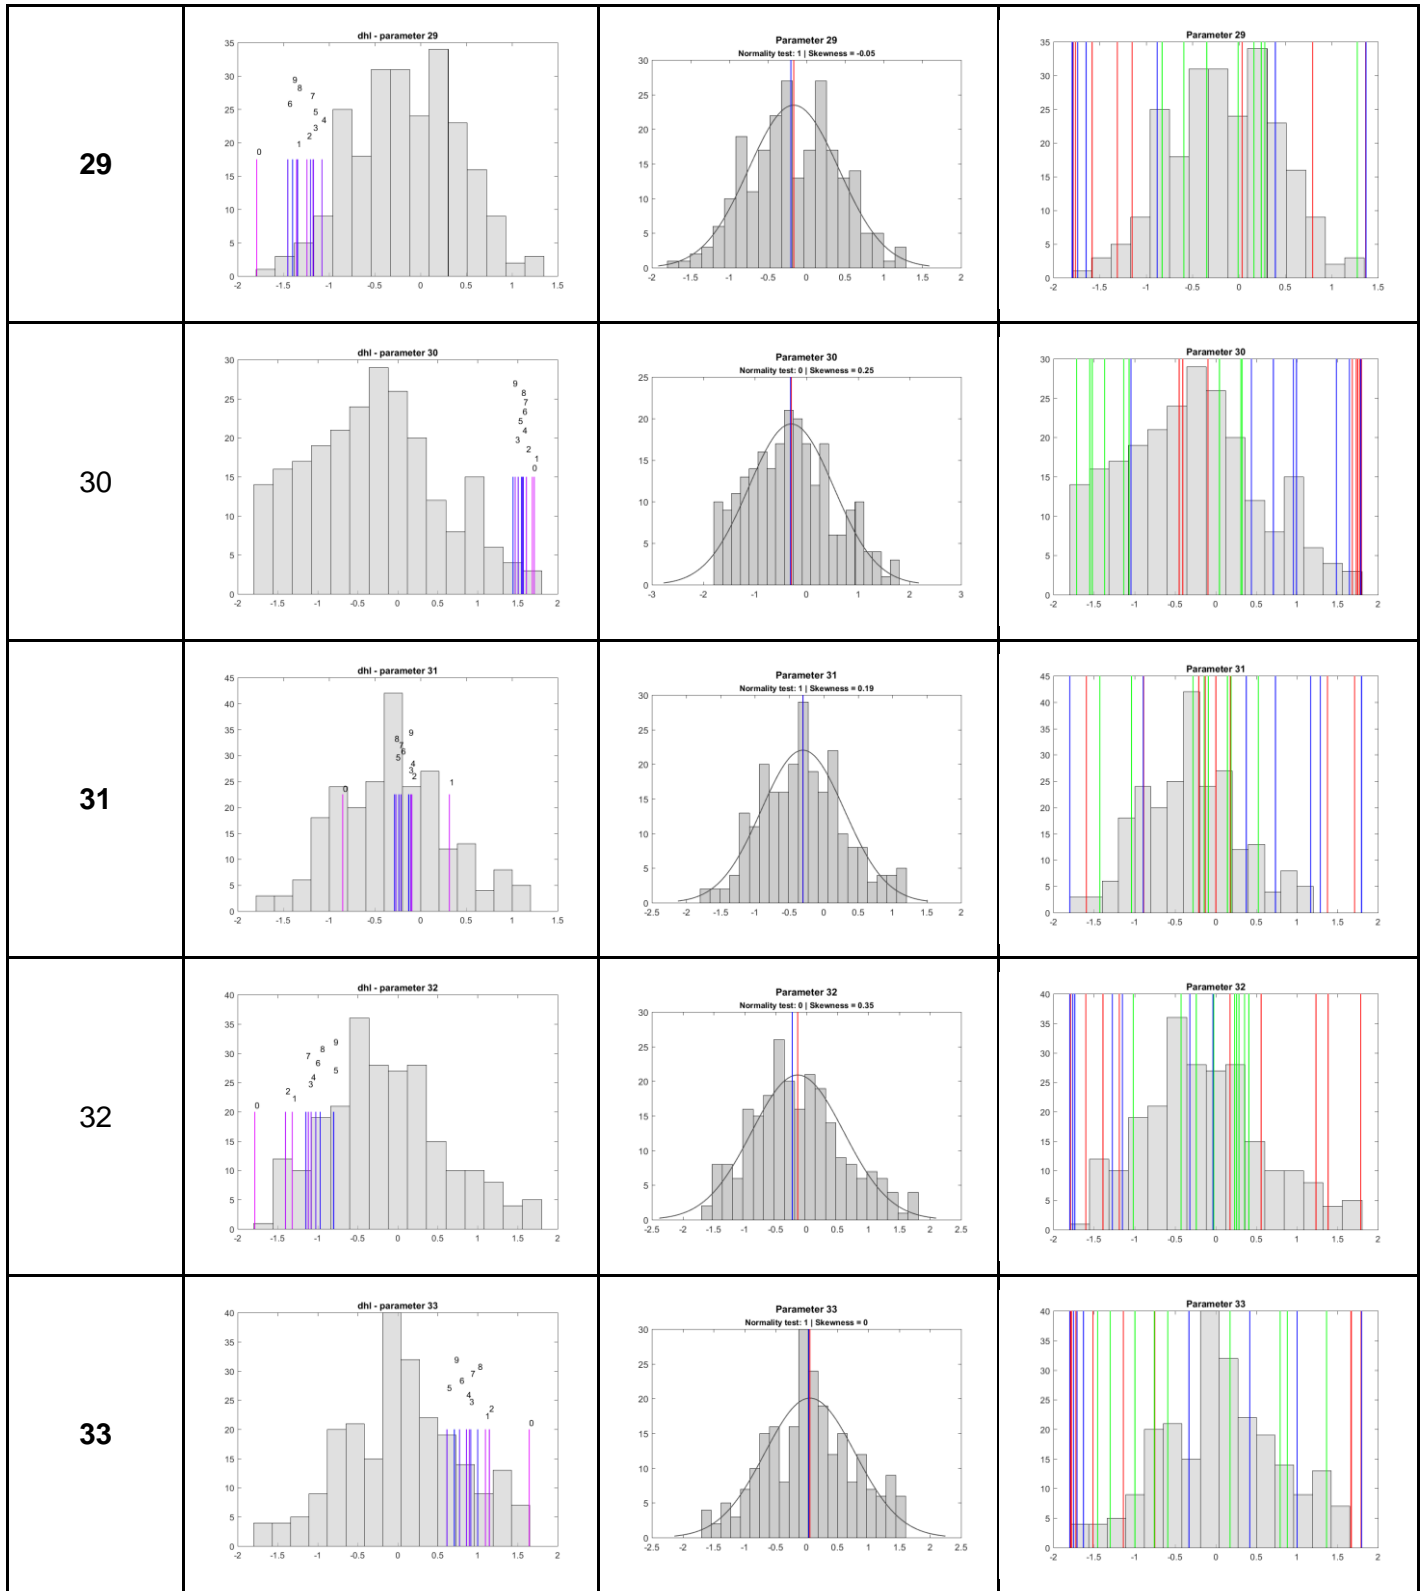


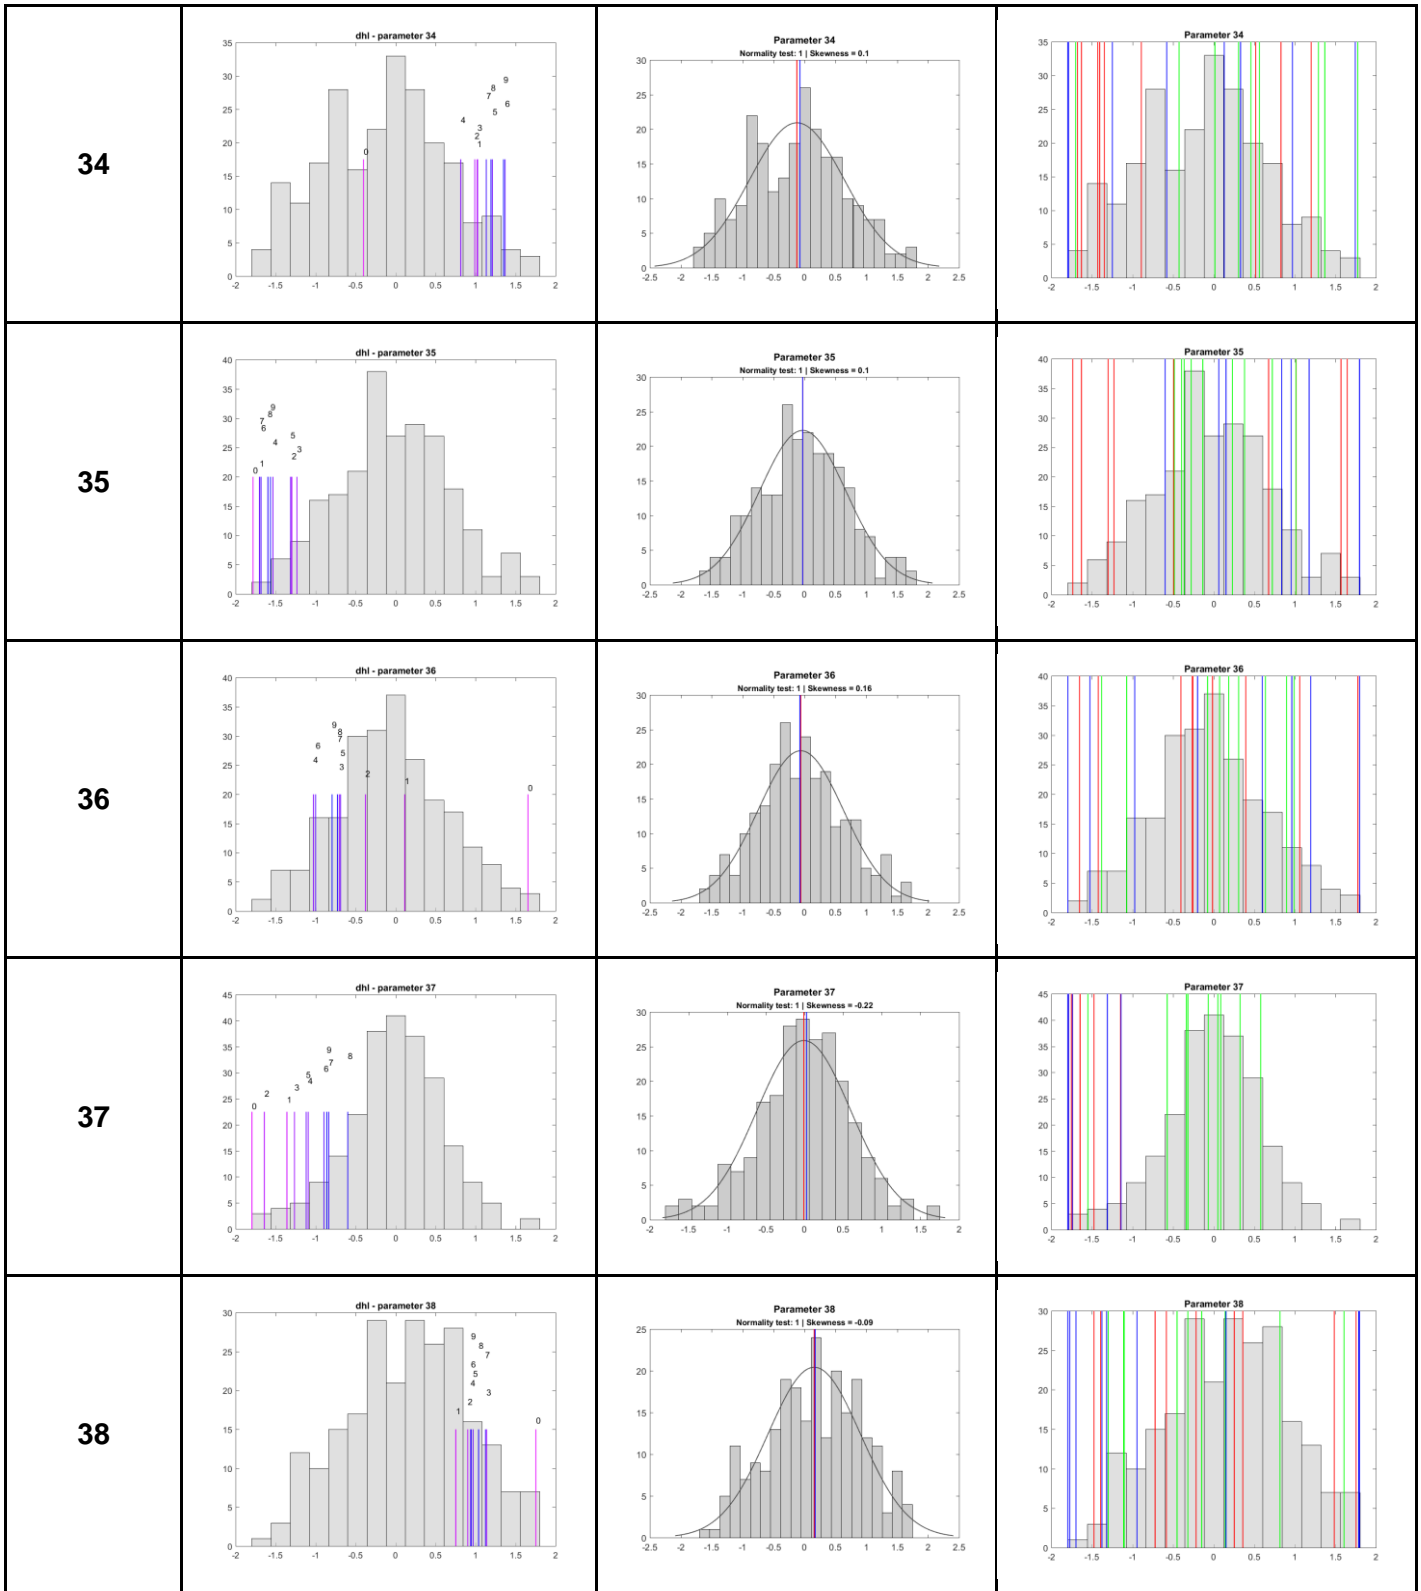


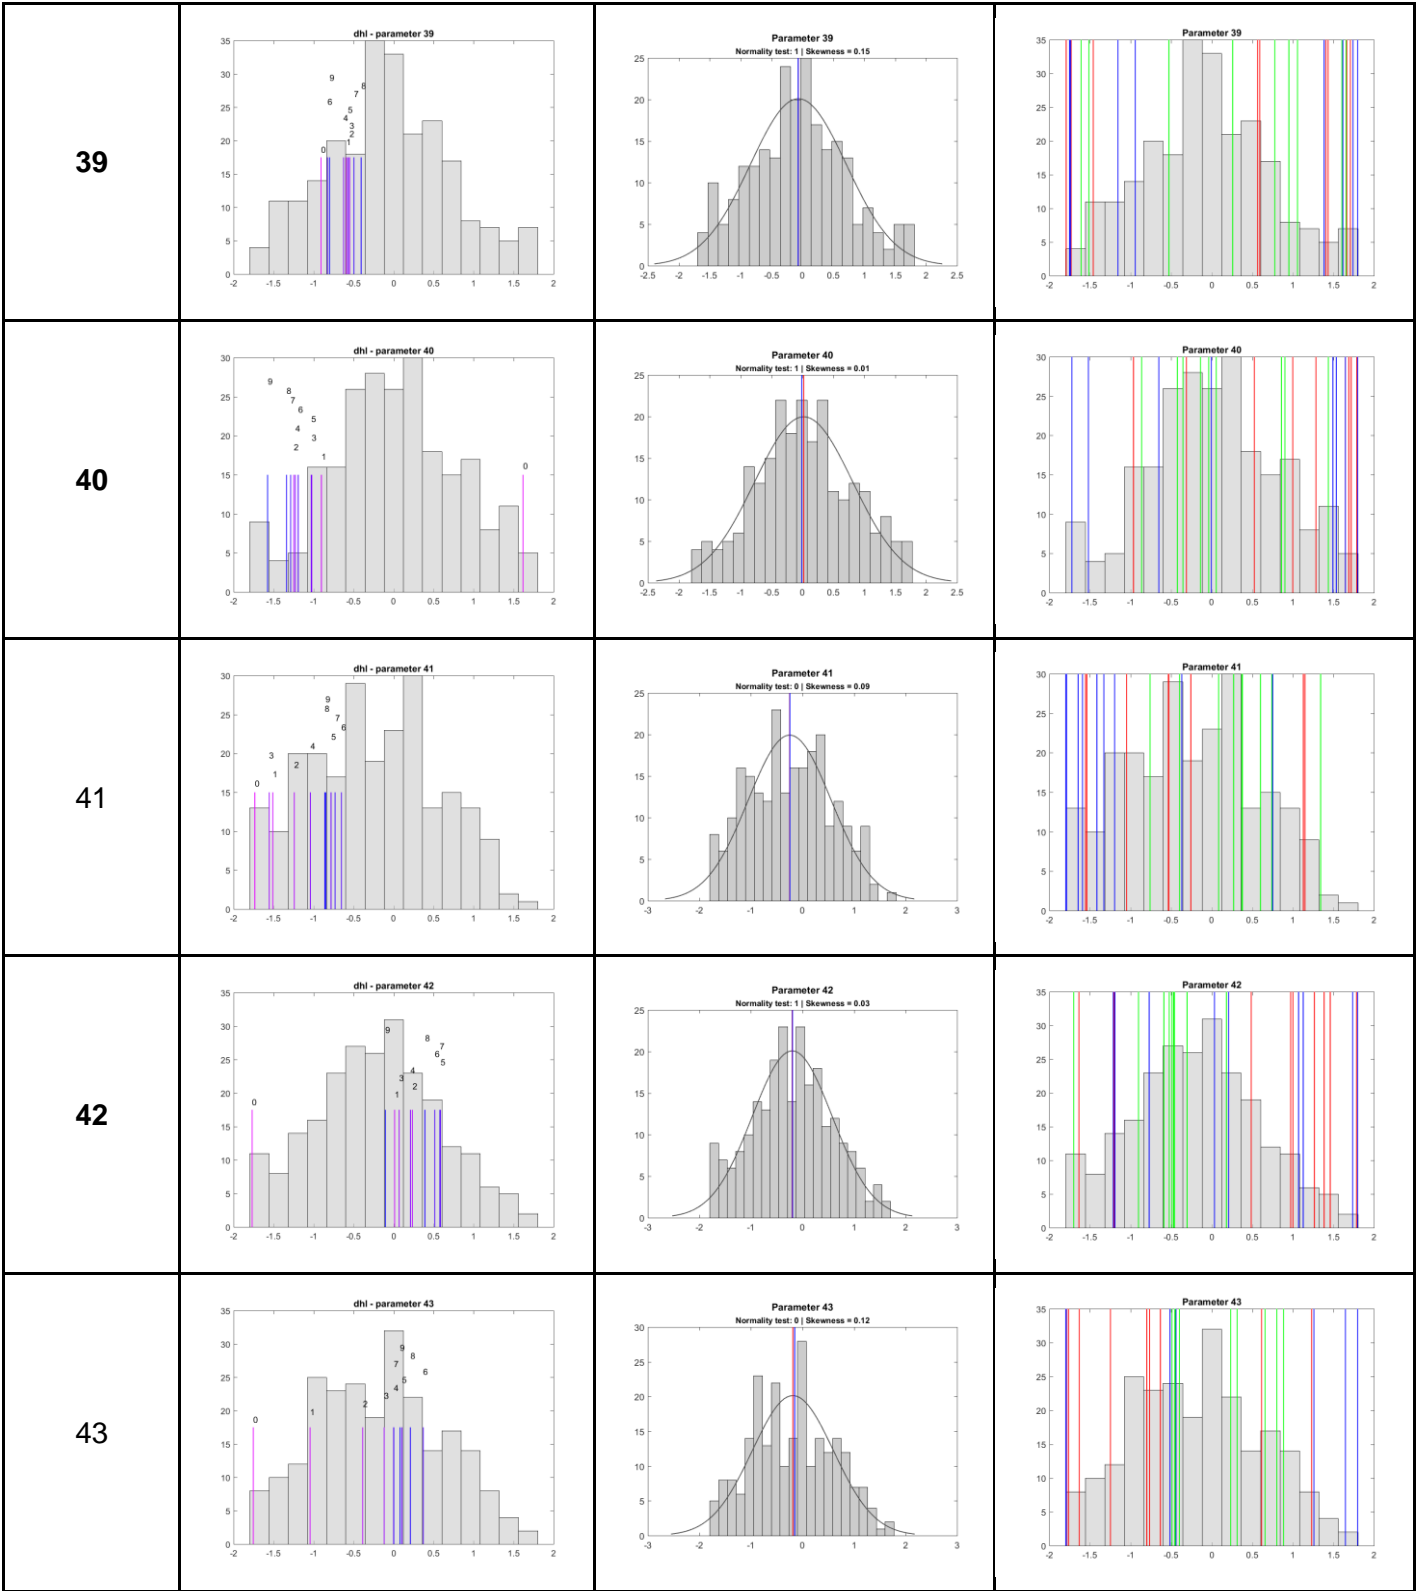


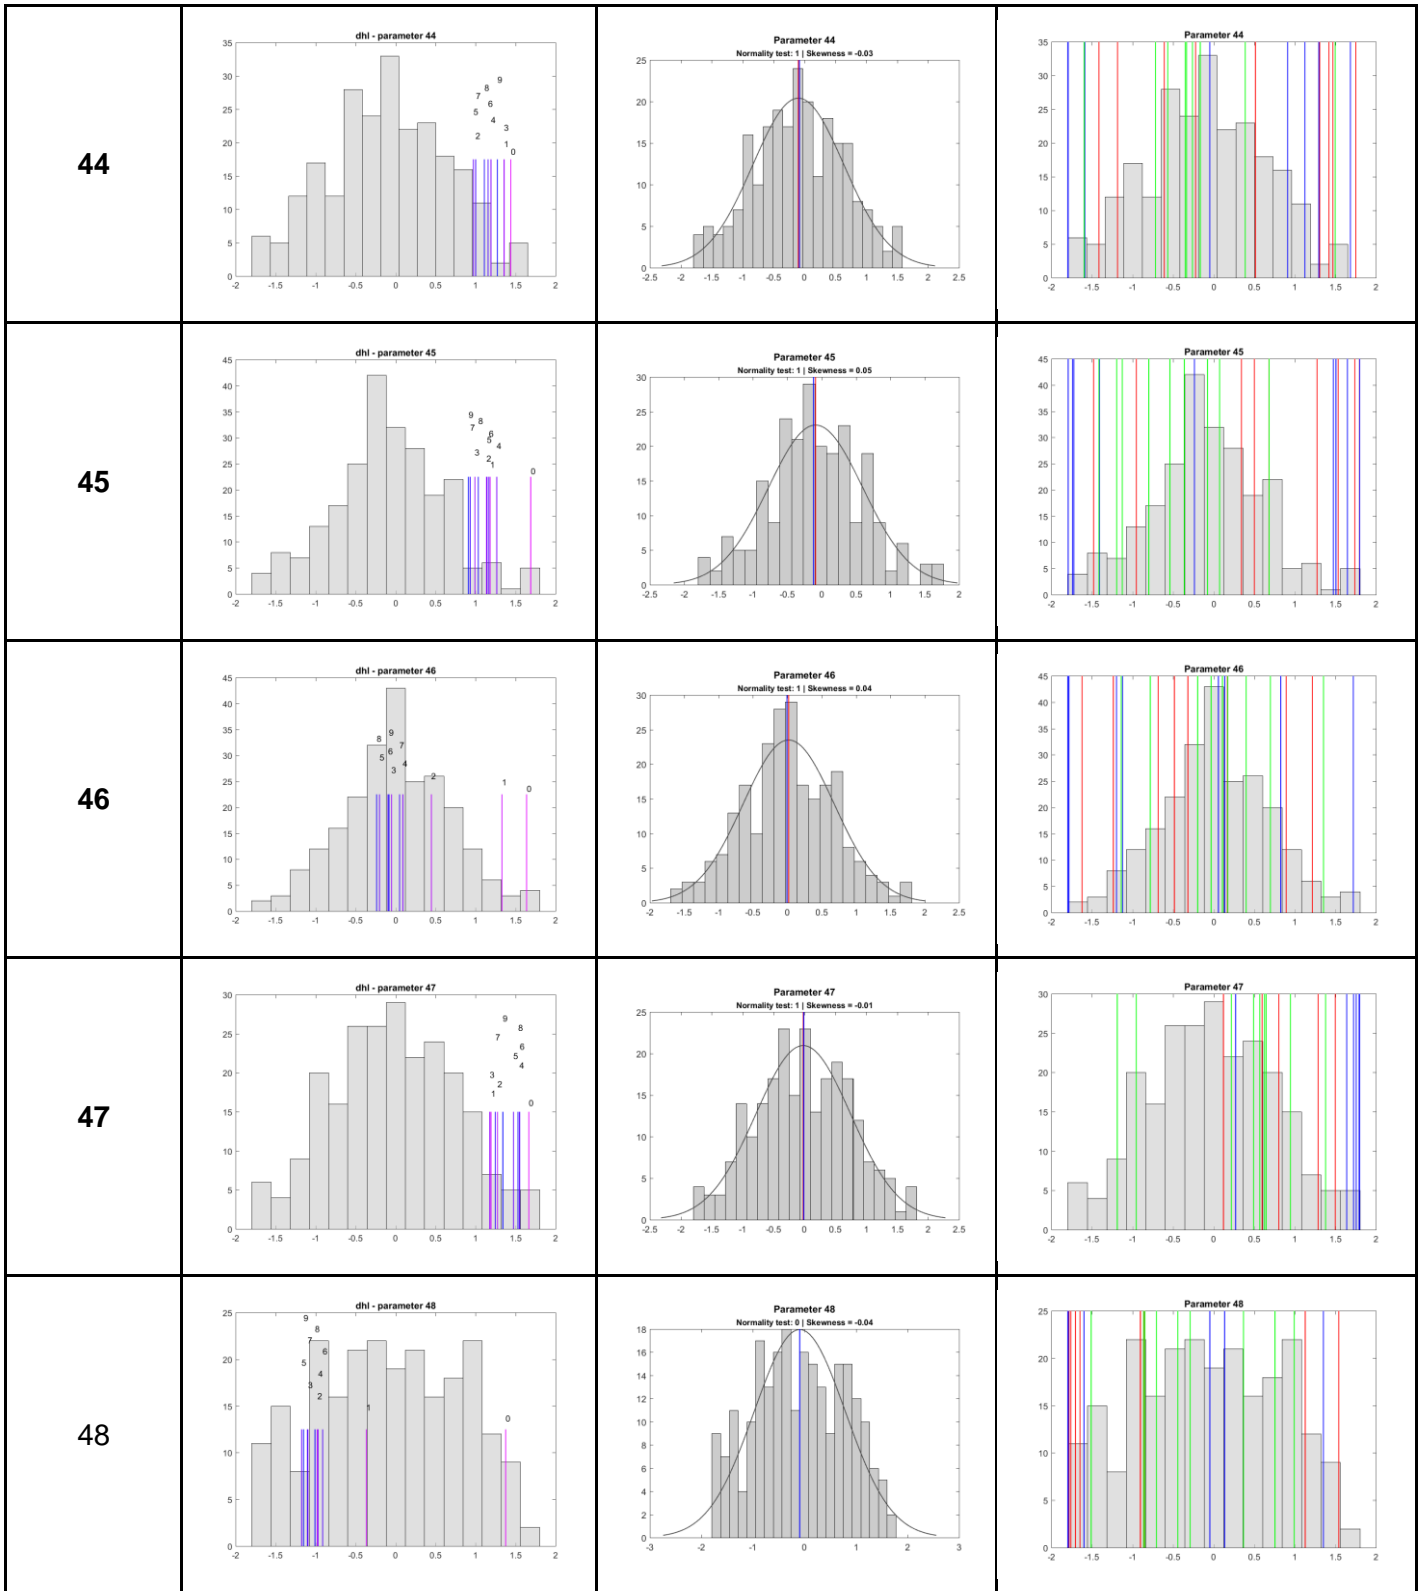


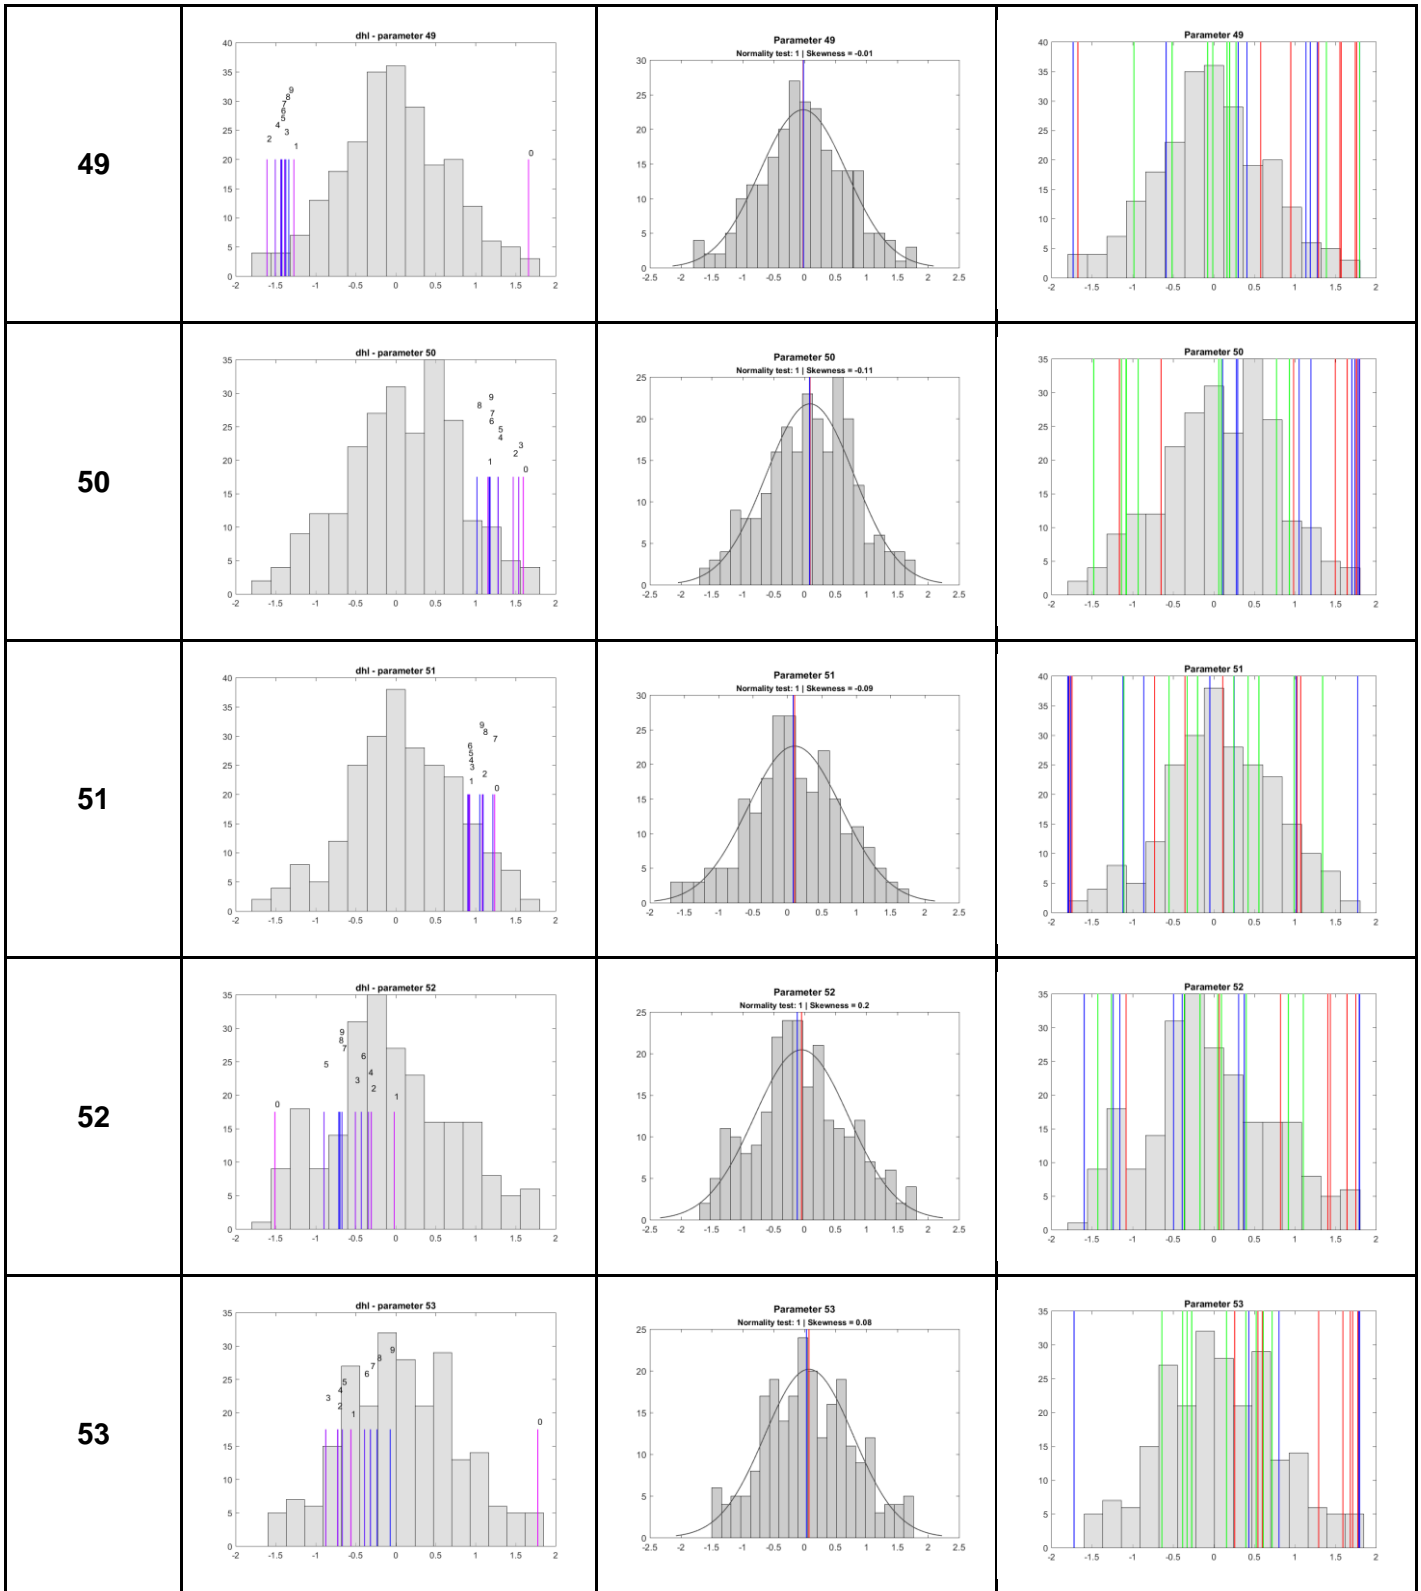


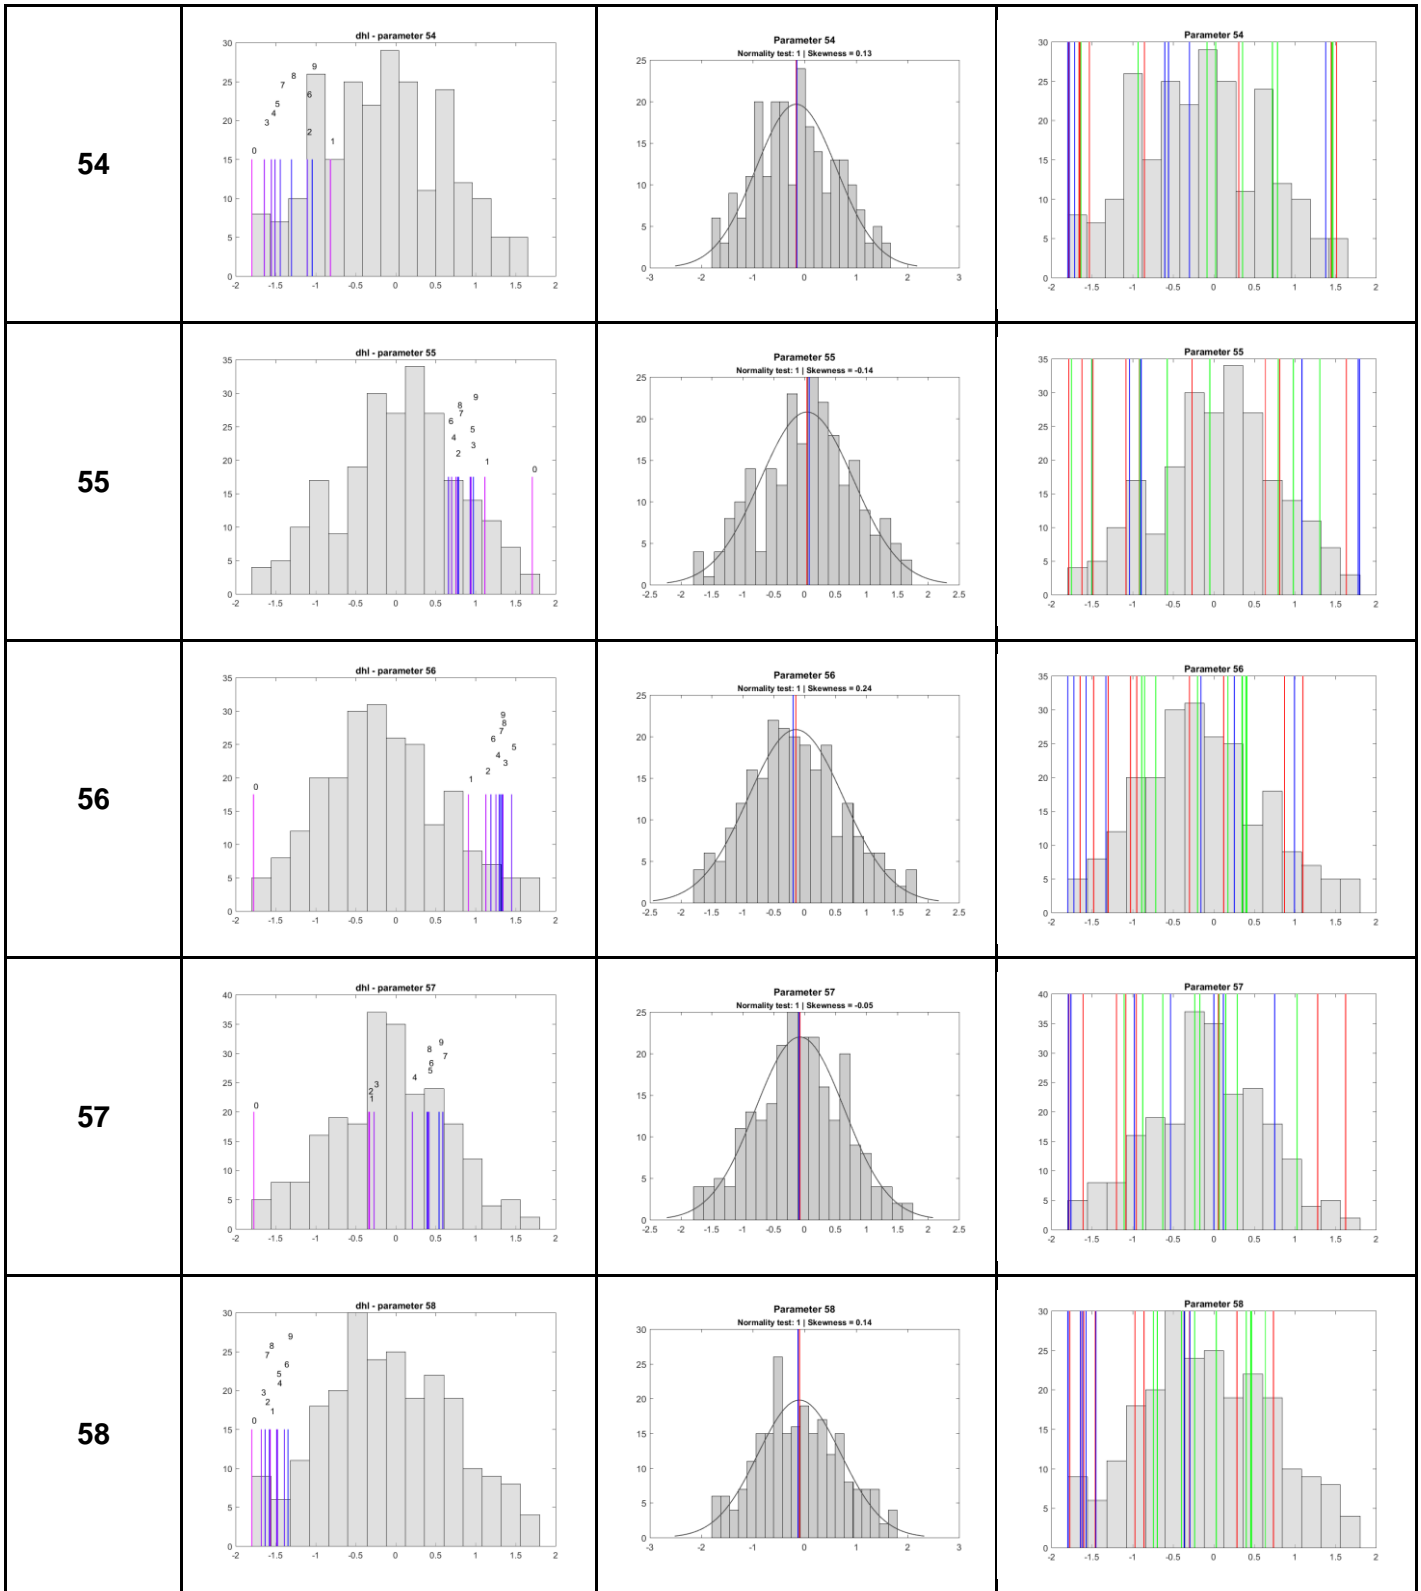


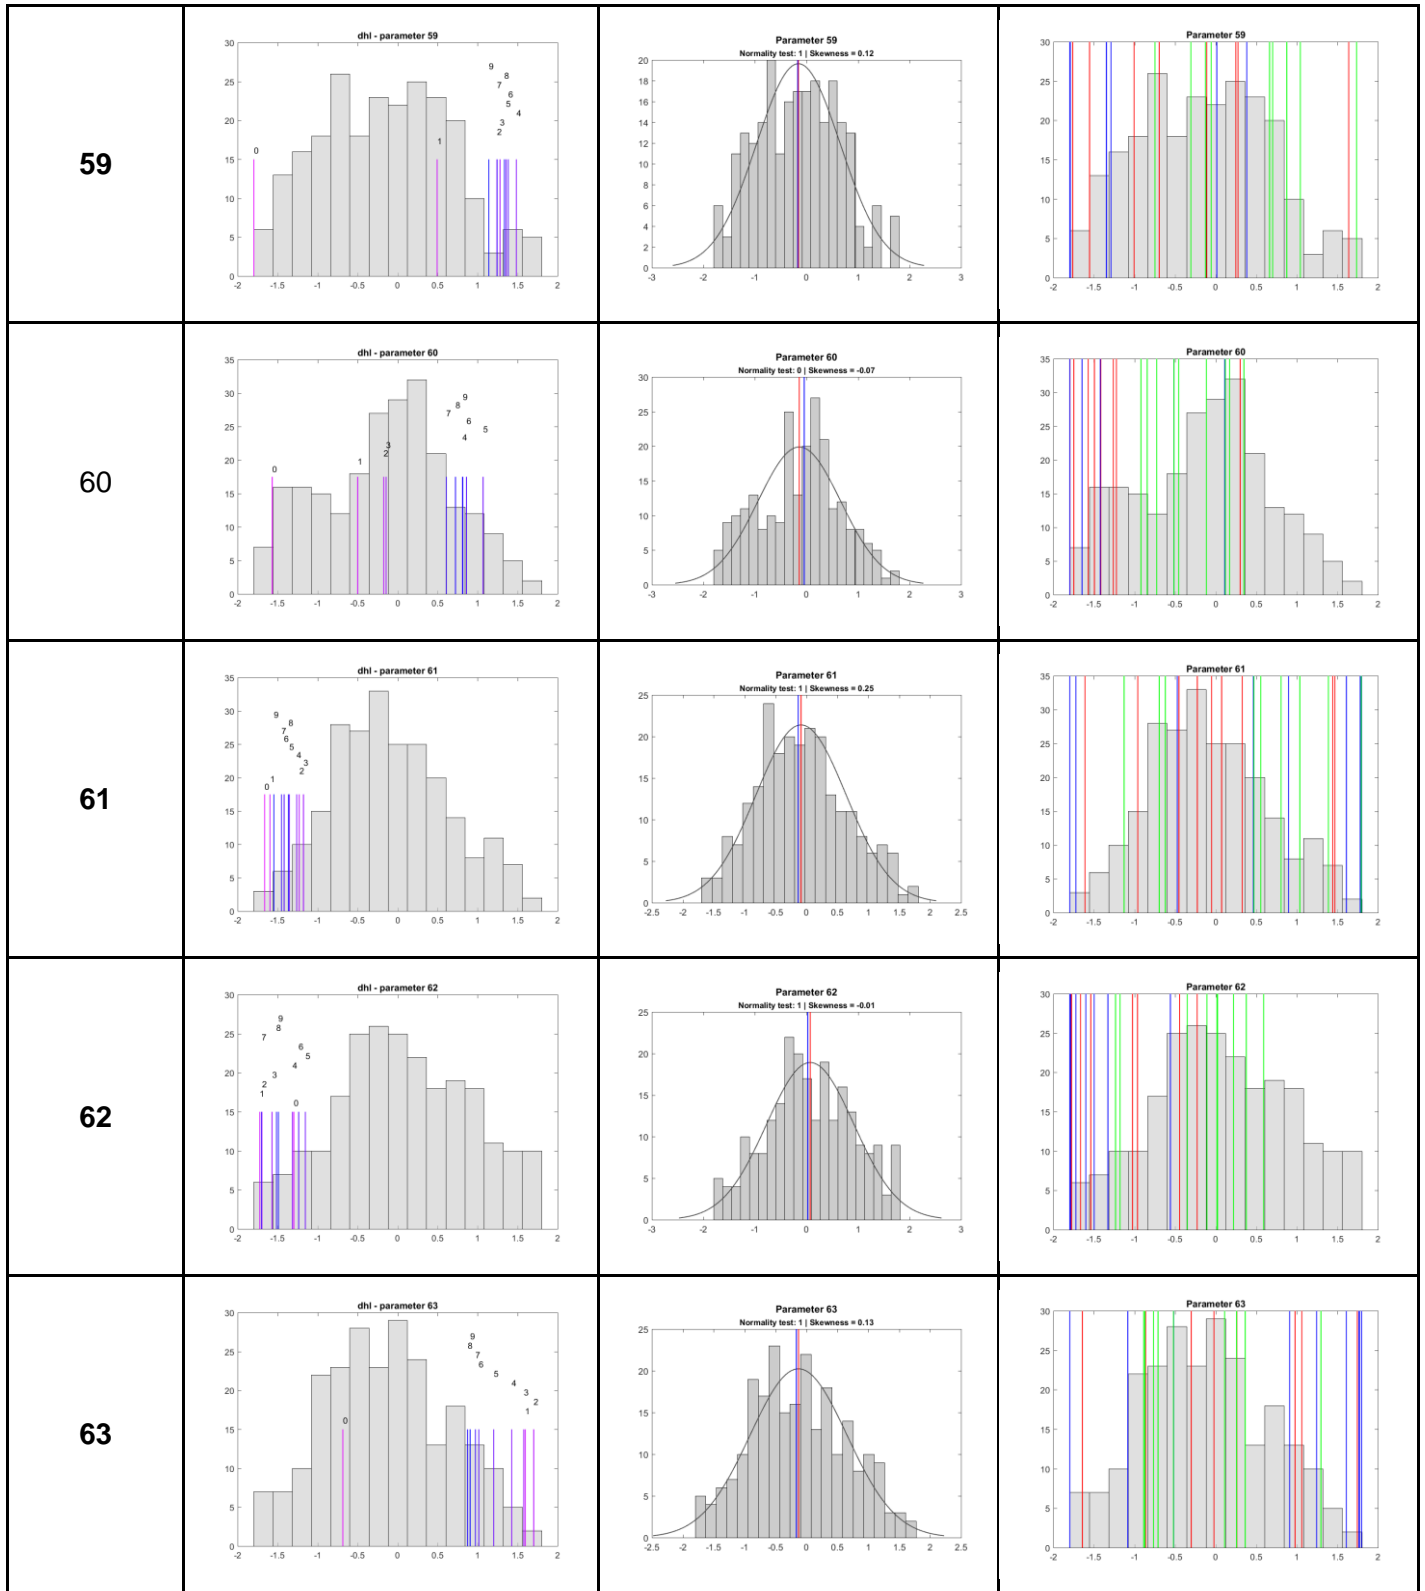


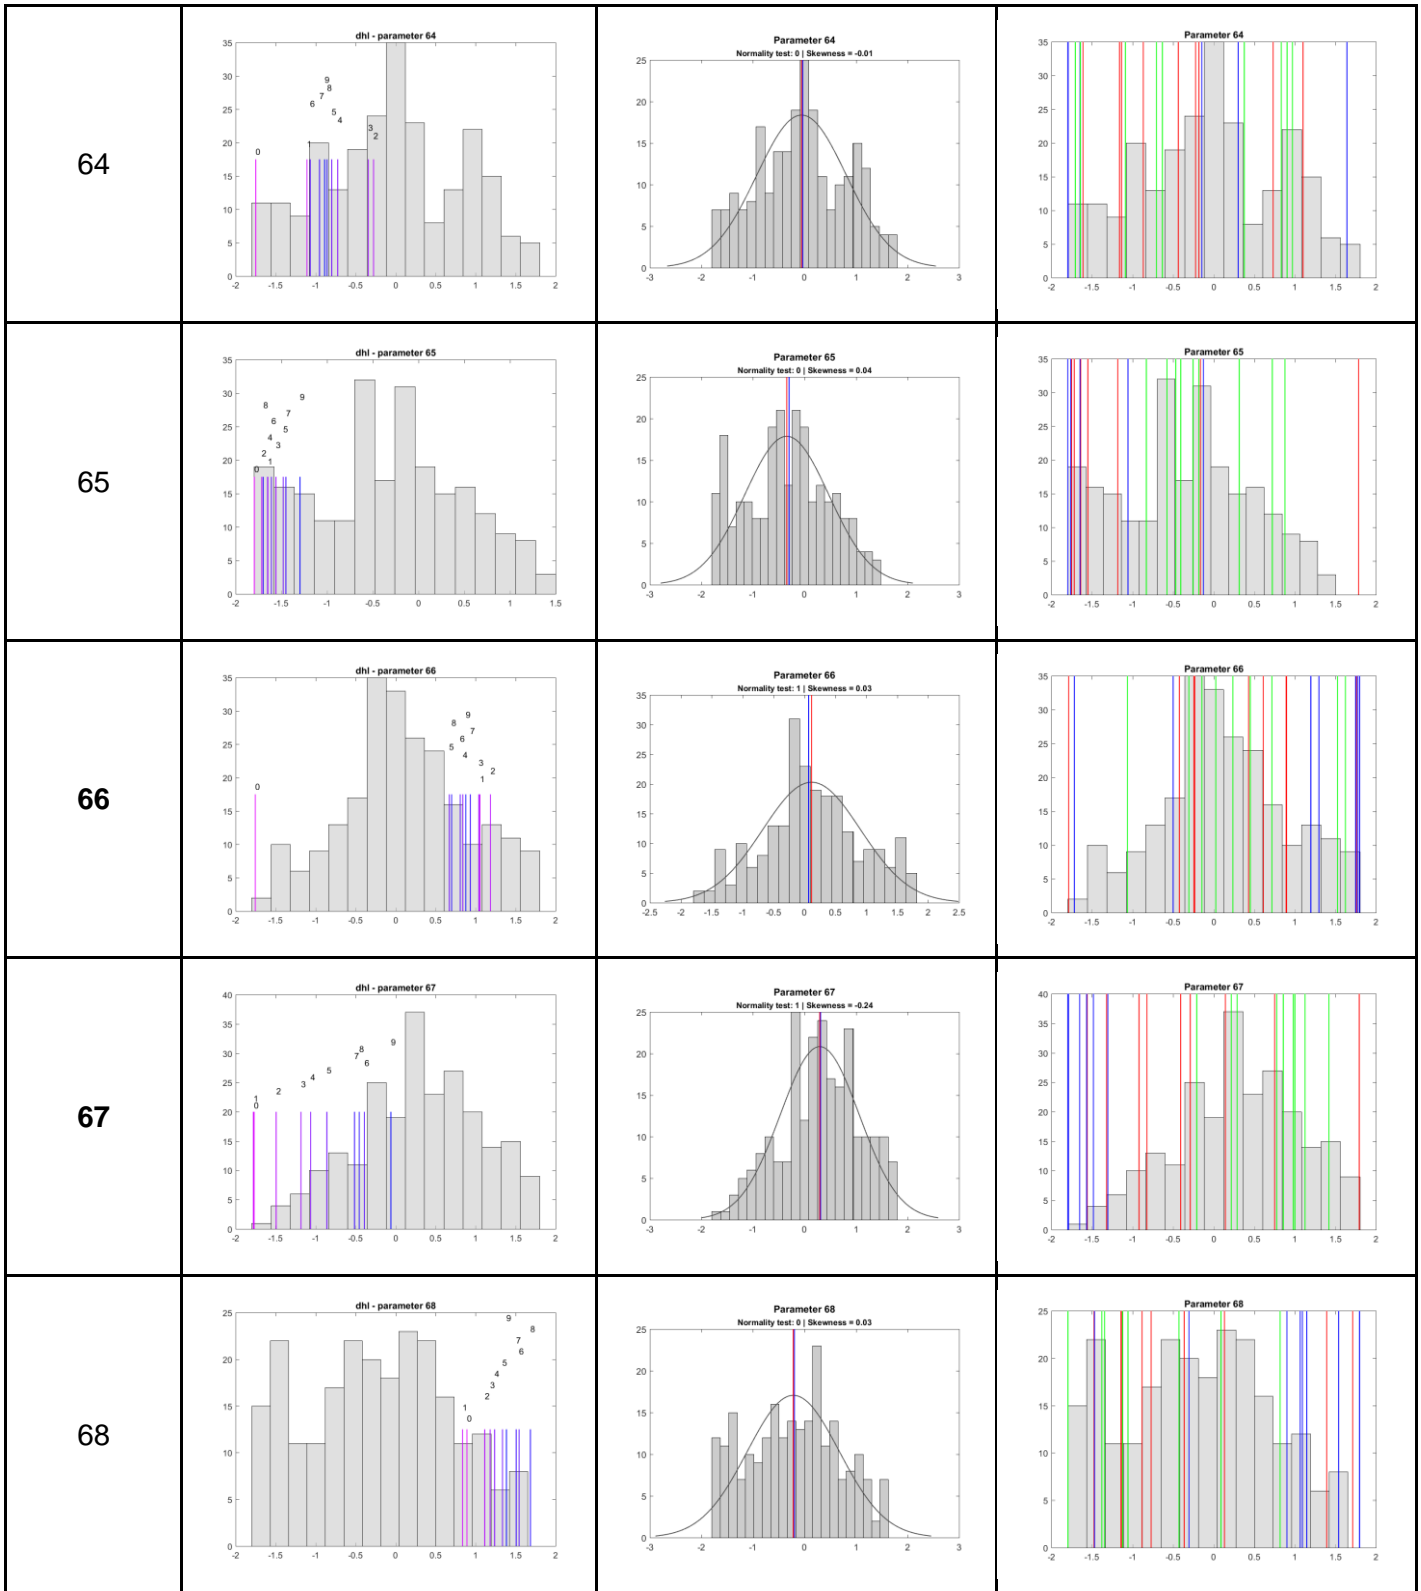


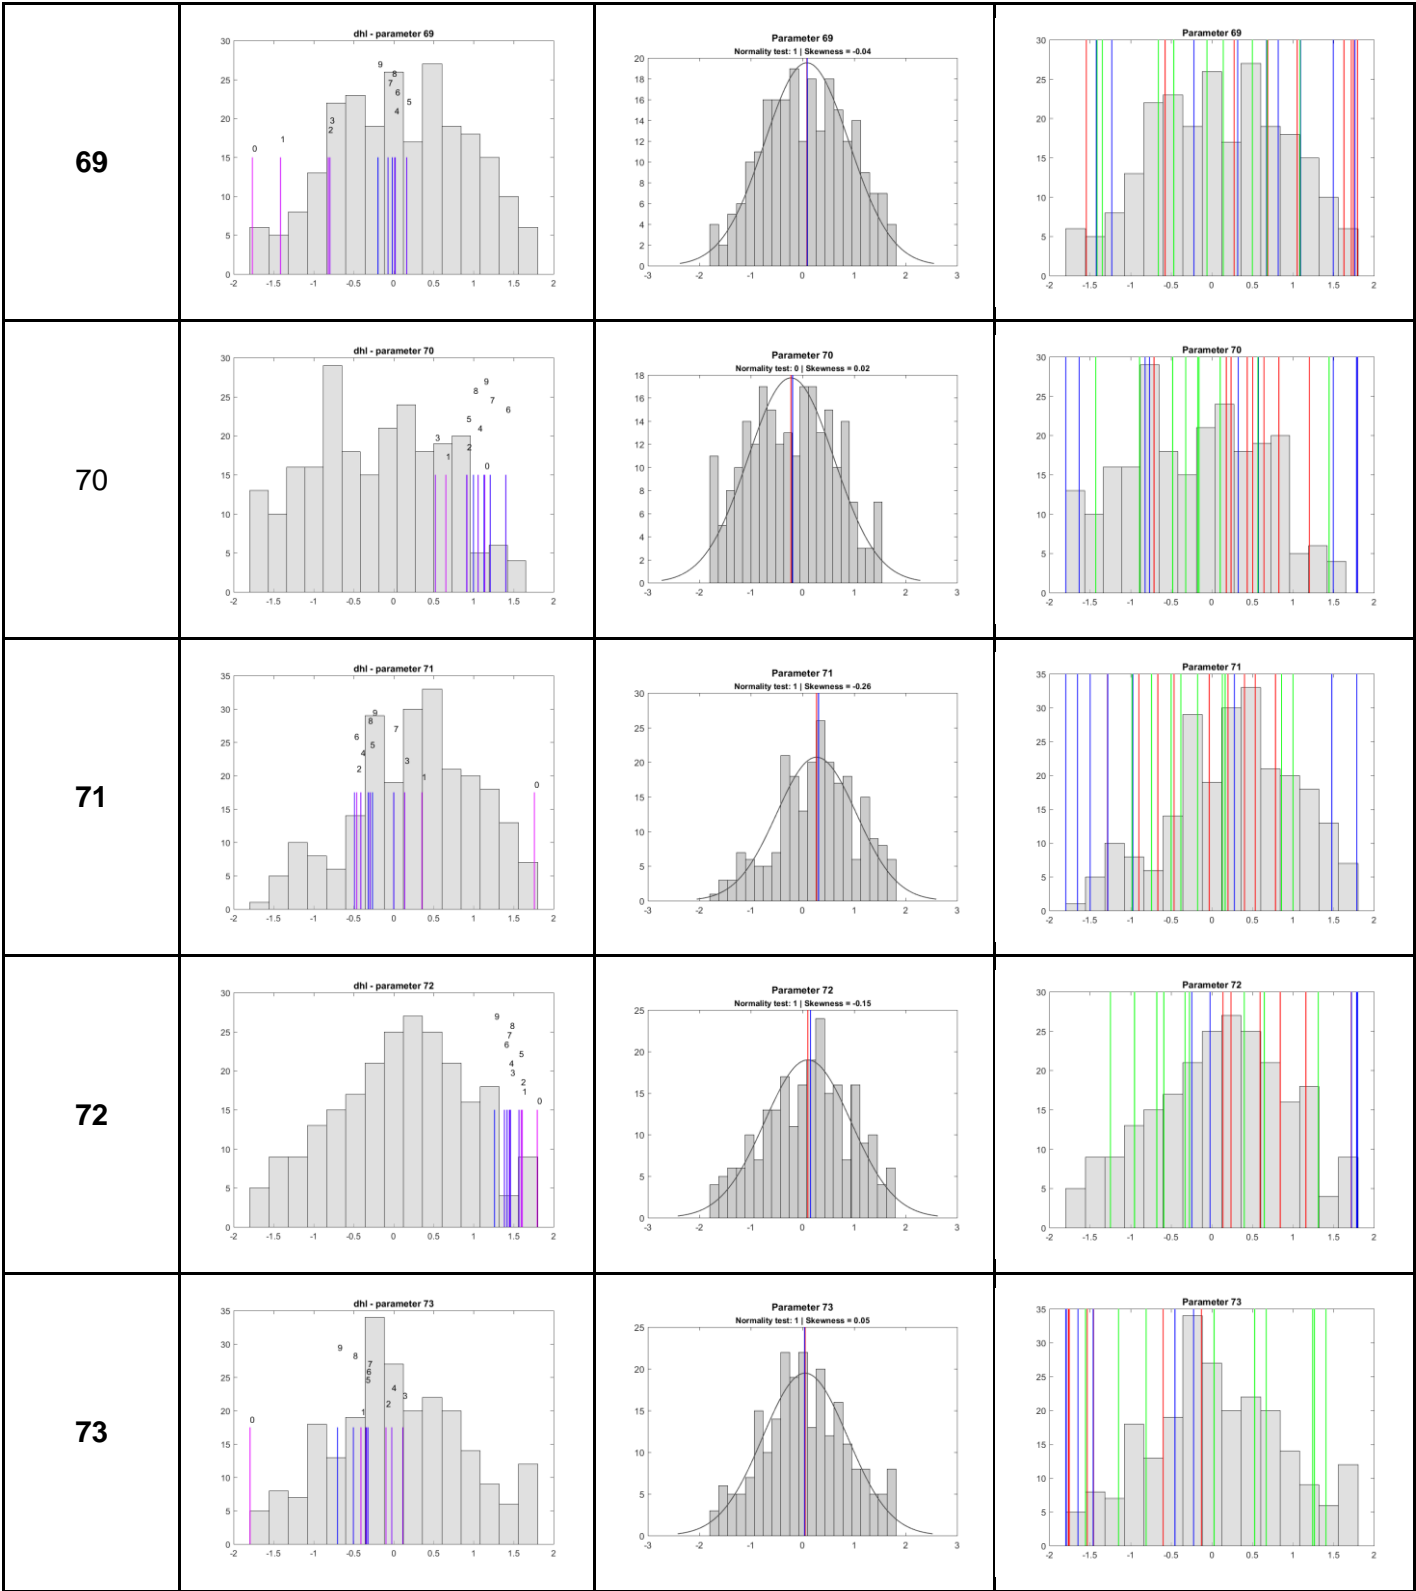





















74






75






76

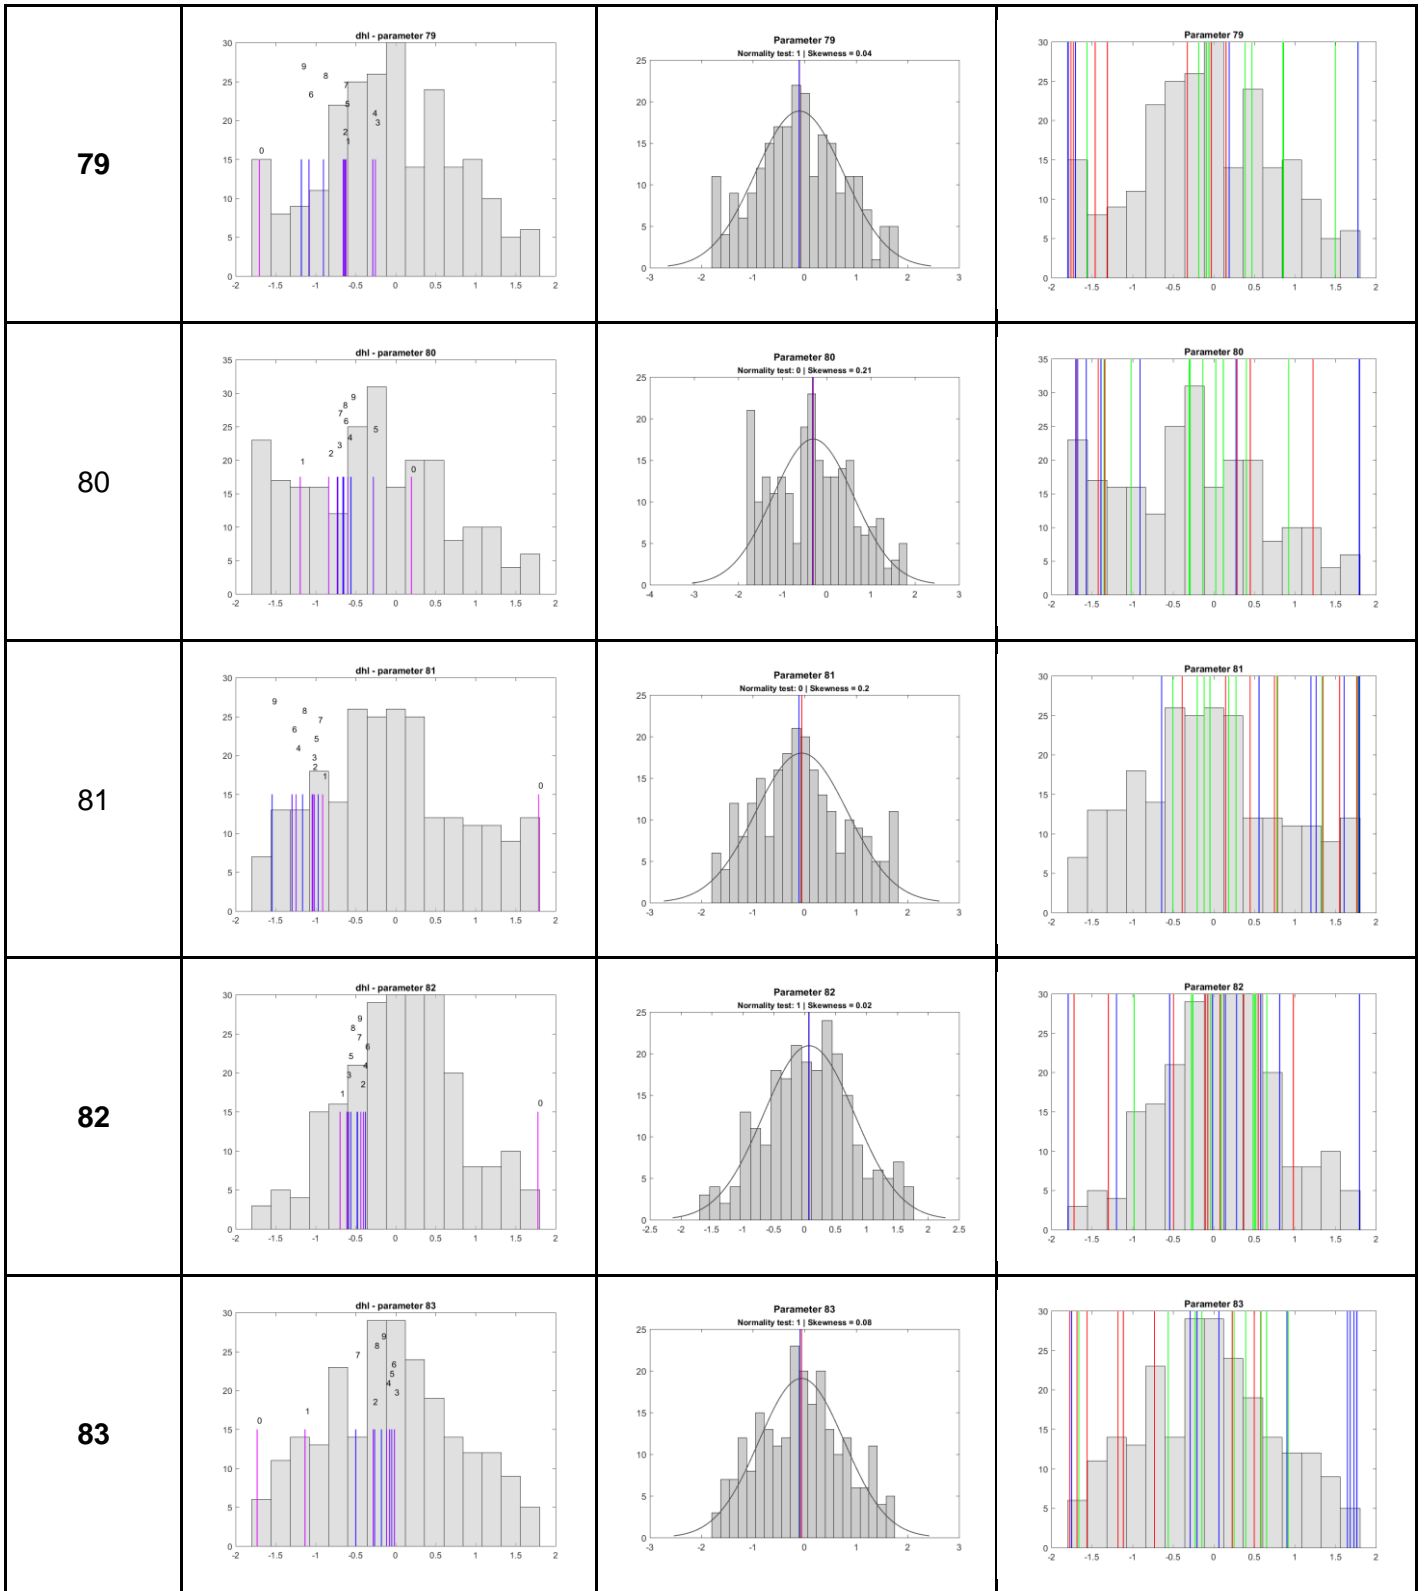


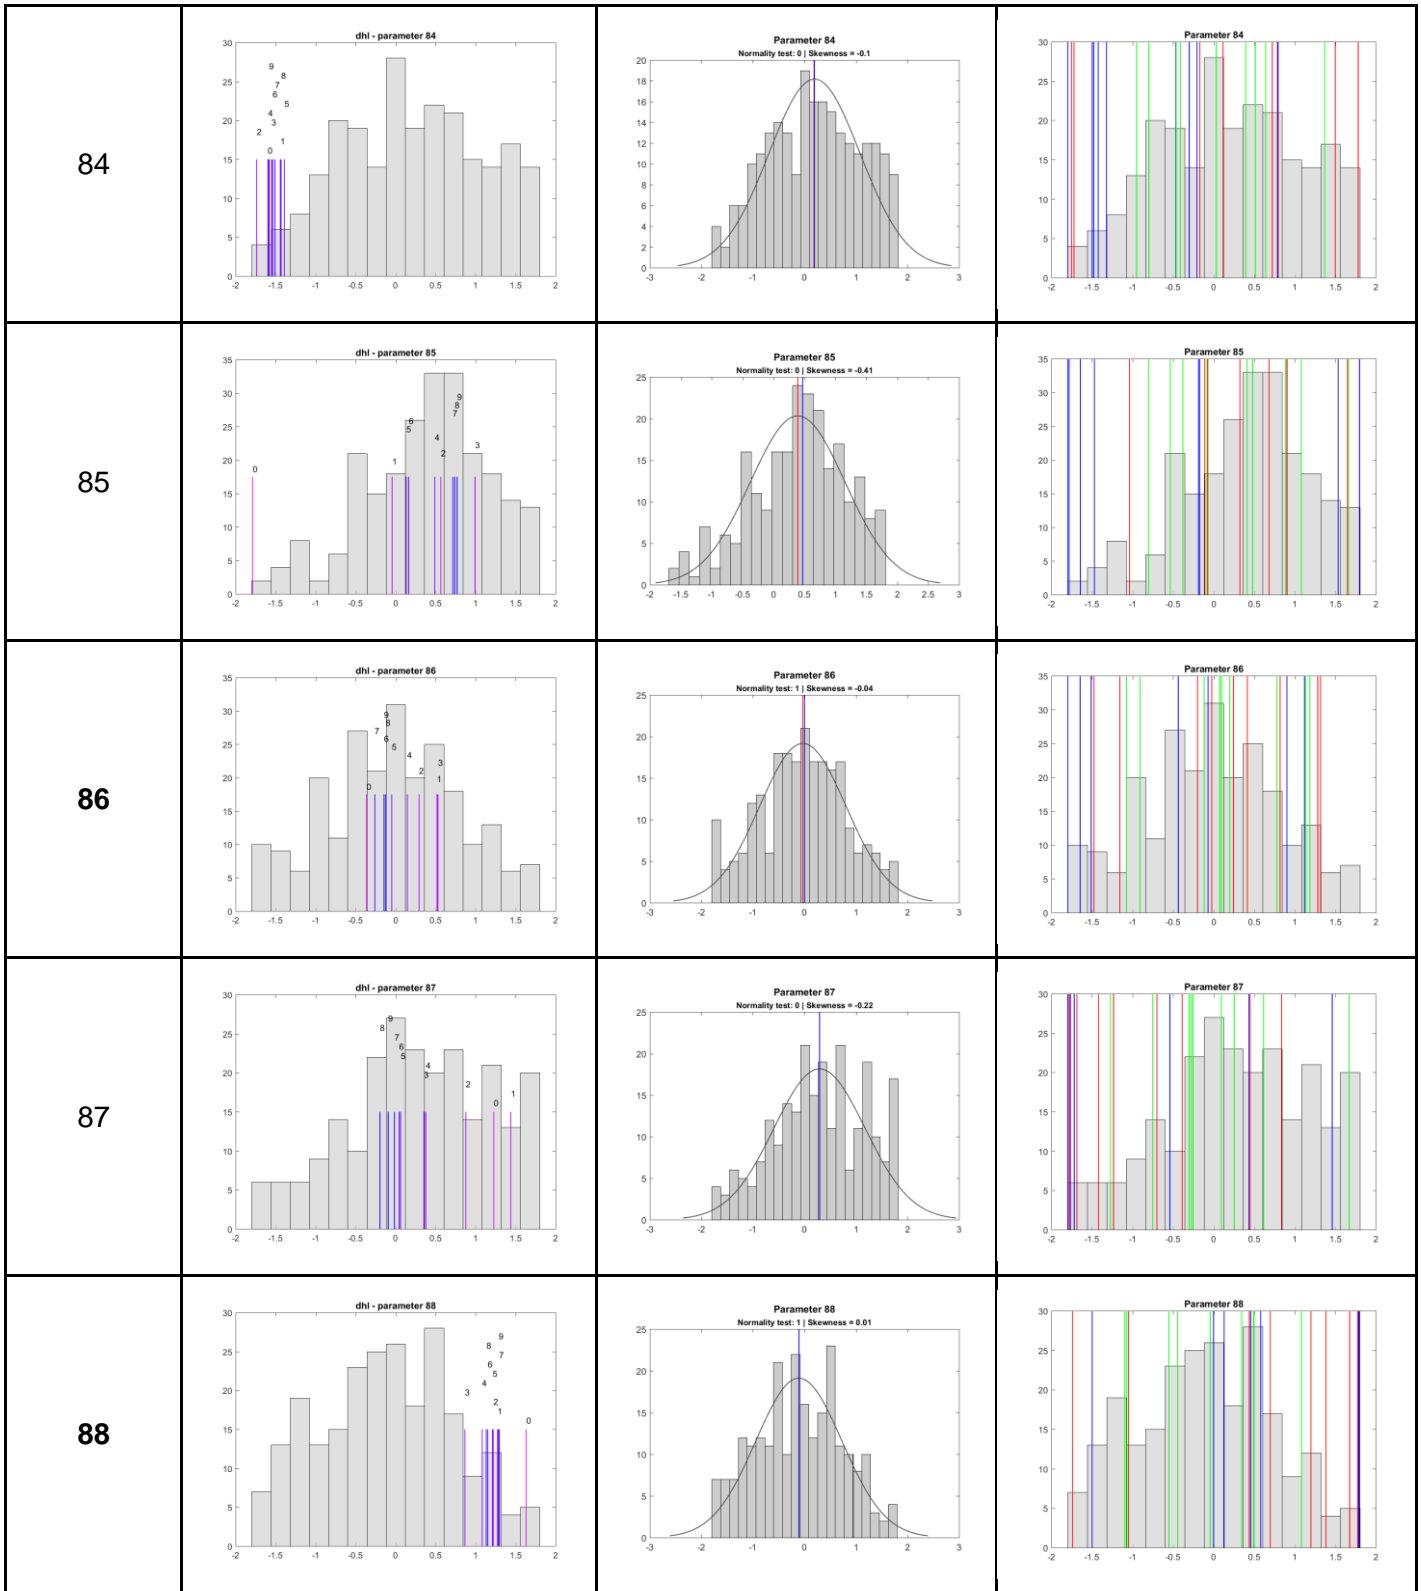
77

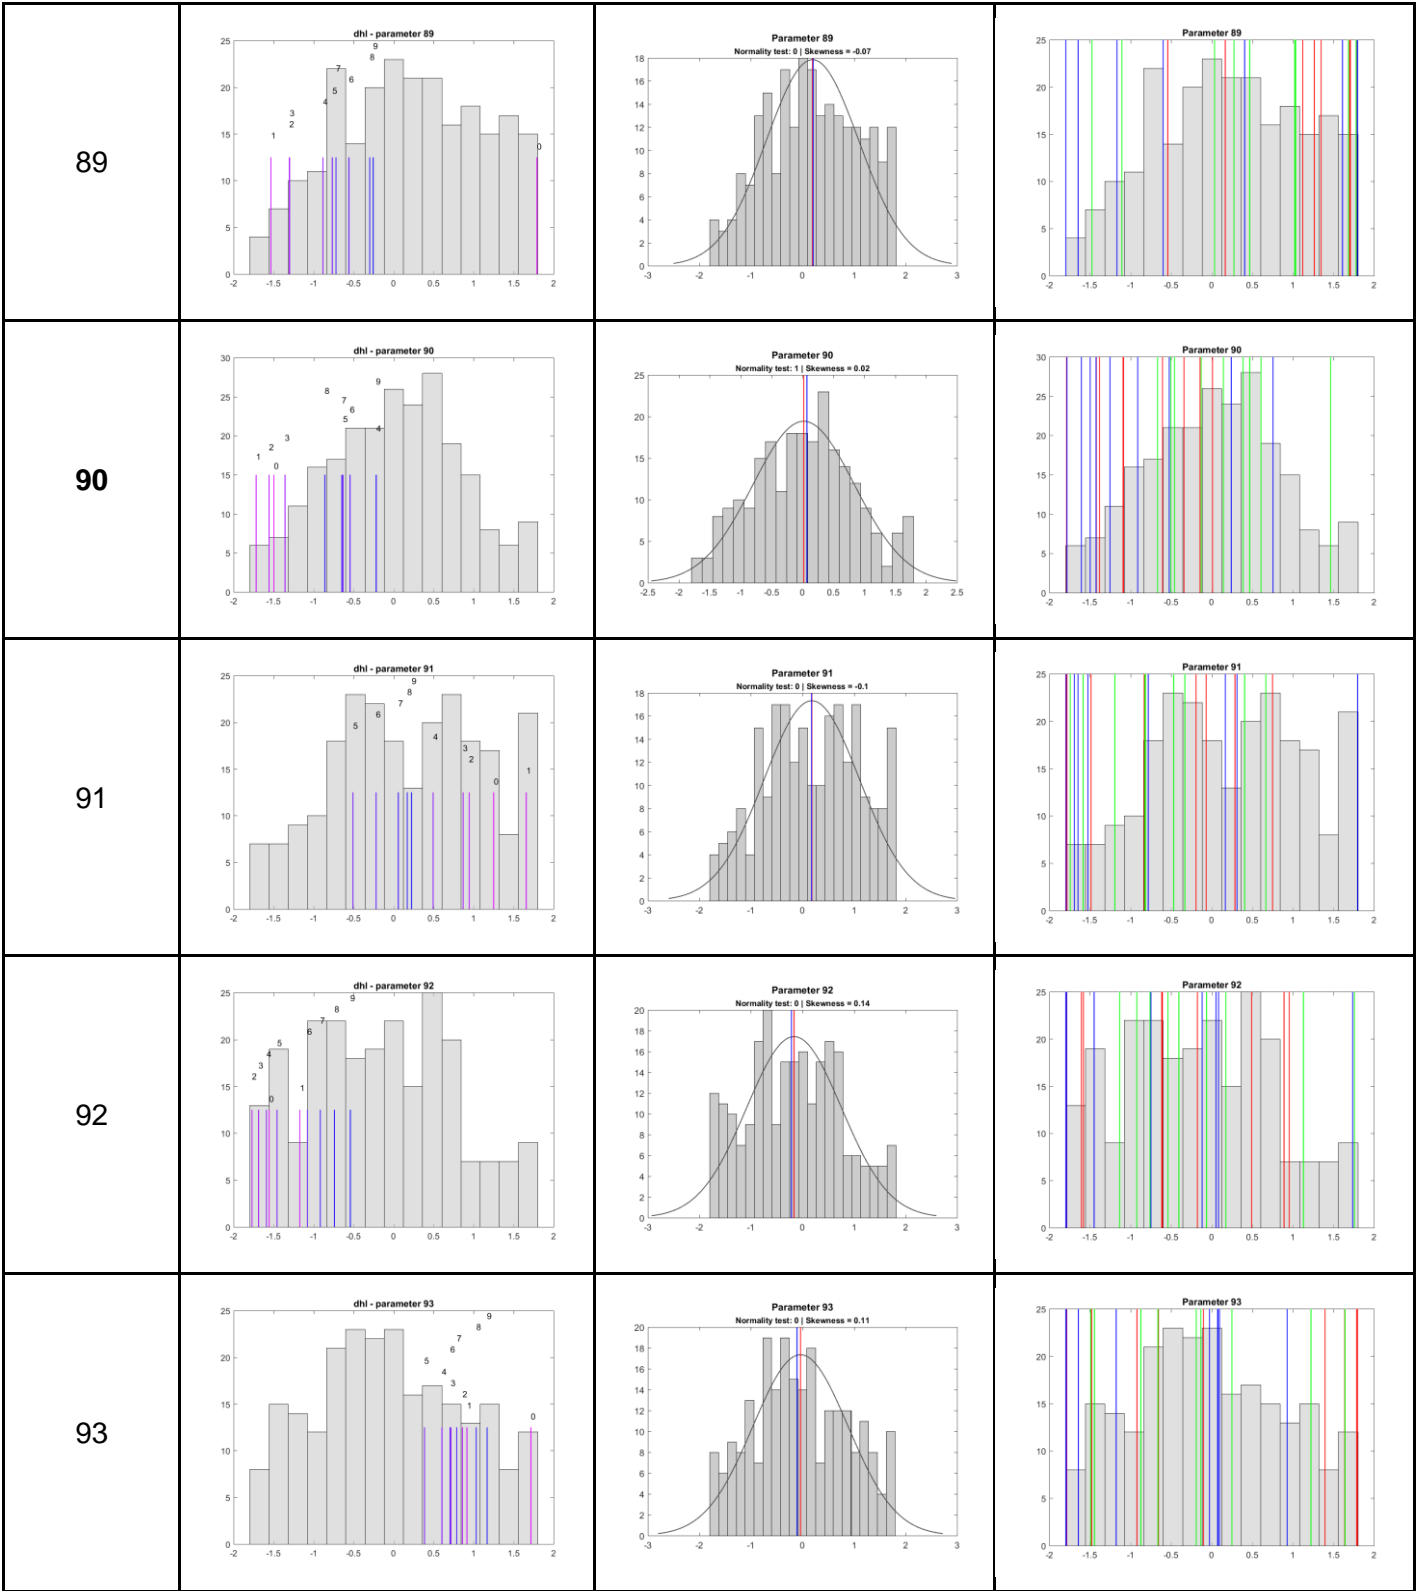


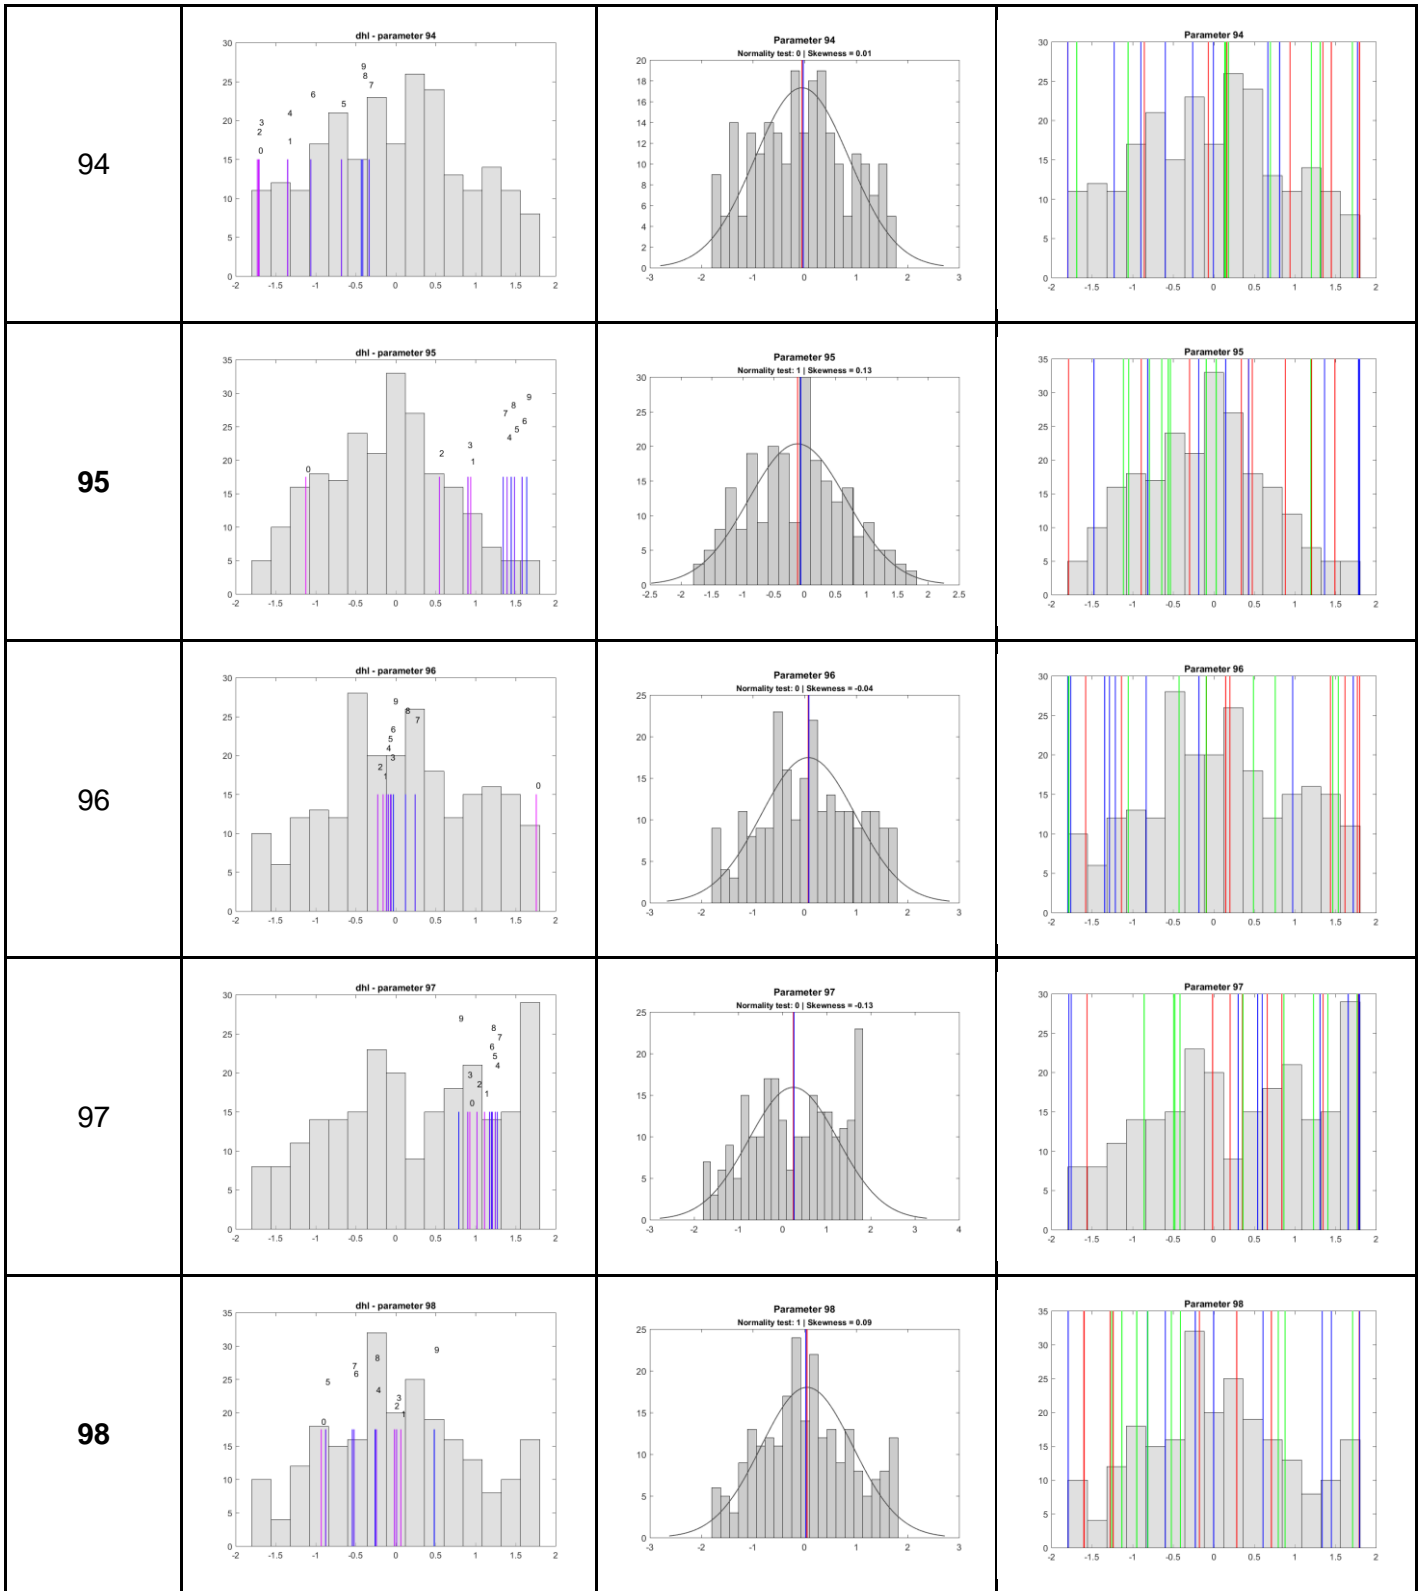
78














A.4. Parameters z-scores

