Mapping the Invisible: Gravity as a Lens to Reveal Dark Matter

Avin Mohebi

École polytechnique fédérale de Lausanne

`avin.mohebi@epfl.ch`

Supervisors: David Harvey and Felix Vecchi

LASTRO, EPFL

September 30, 2025

1 Introduction

Cosmology is the science that studies the Universe as a whole: its origin, its structure, and how it evolves over time. The goal is to build a coherent picture of how everything we see today came to be. One of the most surprising discoveries of modern cosmology is that most of the matter in the Universe is not the ordinary matter we are familiar with and can see, but something invisible known as dark matter.

Dark matter does not emit, absorb, or reflect light. We cannot see it directly, but we can see the effects of its gravity. Observations ranging from the way galaxies rotate to the way matter is spread out across the Universe, show that dark matter must exist. In fact, it outweighs normal matter by about five to one. Understanding it is crucial, because dark matter plays a central role in shaping the Universe.

In order to understand dark matter, we must look in places where dark matter concentrations are high and easy to study. The way structures form in the Universe is hierarchical. Small objects form first, and then merge to make larger ones. Stars come together to form galaxies, galaxies collect into groups, and groups grow into clusters. Galaxy clusters are the largest structures held together by gravity. They contain thousands of galaxies, hot gas that shines in X-rays, and enormous amounts of dark matter. These are the ideal places to study dark matter.

Gravity does more than hold clusters together: it also bends light. According to Einstein's theory of general relativity, any massive object can bend the path of light that passes near it. When the alignment is just right, the bending can be dramatic, stretching background galaxies into arcs or even full rings, known as strong lensing. More often, the effect is subtle: background galaxies are stretched just a little so that their shapes look slightly elongated. This is weak gravitational lensing.

As galaxy clusters are massive objects, they bend light from the background galaxies too. Since their mass is dominated by dark matter, most of the bending is also done by dark matter. Therefore, by studying gravitational lensing, this hidden matter can be detected and mapped.

The stretching of background galaxies in gravitational lensing is called shear. For one galaxy, the change is too small to notice, because galaxies are not perfectly round to begin with. But if we look at thousands of galaxies behind a cluster, the small distortions add up in a consistent pattern. Without lensing, galaxy shapes would average out to circular, since there is no preferred direction in the Universe. With lensing, the pattern of shear reveals how mass, including dark matter, is distributed.

The power of weak lensing is that it measures mass directly. It does not matter whether the mass is in stars, gas, or dark matter. For this reason, weak lensing is one of the most reliable ways to map the total mass in galaxy clusters.

My research focuses on improving how we select the background galaxies that are used for weak lensing. Not all galaxies are useful: small galaxies with indistinguishable shapes, and faint or noisy detections can all reduce the signal. Traditionally, astronomers made these selections by hand, using visual checks or simple thresholds. My project aims to make this process automated and optimized, so that the galaxies providing the clearest lensing signal are selected consistently. This allows the analysis of many clusters quickly and reliably.

My work contributes to making better maps of dark matter. These methods will also be essential for future surveys such as ESA's Euclid mission, which will measure weak lensing for tens of thousands of clusters. Automated methods like the one I am developing will ensure accuracy, efficiency, and scalability.

2 Data

The data used in this project came from Hubble Space Telescope images of the Bullet Cluster. The cluster is a well-known system of two colliding clusters, in which dark matter can be studied through the gravitational lensing of background galaxies. The Hubble image can be seen in Fig. 1.

Figure 1: RGB image of a galaxy cluster taken by the Hubble Space Telescope. Credit: NASA/ESA/Hubble.

The Bullet Cluster is an ideal case study, as it provides one of the clearest pieces of evidence for dark matter, showing the separation of baryonic (visible) matter and dark matter after a collision. The filters used in this study were the F814W and the F606W, which isolate light in the near-infrared/red (814 nm) and optical/yellow (606 nm) parts of the spectrum, respectively, allowing for the comparison of how objects appear in different wavelengths.

3 Methods

The software used for creating the mass maps was PyRRG: a program developed by David Harvey to calculate shear measurements. PyRRG extracts sources, provides tools to separate galaxies from stars and noise, excludes cluster members, and applies cuts on parameters such as galaxy size and brightness. Each of these processes will be discussed in more detail in the following sections, where the methodology for source extraction, star–galaxy separation, and parameter cuts will be outlined.

3.1 Shape Measurement

As a first step, I familiarized myself with PyRRG and how it worked. This included visually understanding galaxy–star separation, learning how a cluster member catalog is made, and checking what each step does by overlaying the selected sources on the Hubble image of the Bullet Cluster.

3.1.1 Source Extraction

The first step in PyRRG was to extract all sources from the Hubble image. This was done with SExtractor (Source extractor), a tool which identifies objects above a brightness threshold and returns a catalog along with their properties: total flux, peak surface brightness, effective diameter, and shape estimates. This catalog cannot be used right away, as it does not only include the background galaxies we are looking for. It includes everything in the image: stars, cluster galaxies, background galaxies, and noise detections. This raw catalog was the foundation of the analysis.

3.1.2 Separating Stars

Stars, being point sources in our own galaxy, do not provide lensing information. Their shapes are unaffected by the Bullet Cluster's mass, and including them only contaminates the signal. Therefore, they had to be removed. In order to do this, PyRRG produces plots such as size versus peak magnitude. Magnitude is a measure of the brightness of a source. The higher the magnitude, the less bright the source is. Stars and galaxies occupy different regions in these plots: stars appear along a tight sequence because they are point sources, while galaxies spread out as extended sources.

Noisy detections also tend to appear at small sizes, but unlike stars, they do not follow the magnitude–peak magnitude relation. This makes them distinguishable and prevents them from contaminating the stellar sequence.

By drawing a separation line, stars and noise were excluded. This selection was repeated multiple times until the results looked coherent when using other methods to verify the selection. For example, this separation was confirmed by overlaying the galaxies on the Hubble image: the bright stars should disappear from the selection, leaving only galaxies.

3.1.3 Removing Cluster Members

Cluster members are not useful for weak lensing, since they are not behind the cluster and hence their shapes are not distorted by its mass. If left in, they dilute the signal. To remove them, a cluster member catalog was matched to the shear catalog, and identified cluster members were excluded. In practice, this was done by plotting the color (F814W–F606W) against magnitude and removing galaxies along the red sequence, which we assume to be cluster members. These galaxies appear bright because they are nearby, and red because they are composed of older stellar populations that have stopped forming stars. Cooler, older stars emit stronger at redder wavelengths. A visual check by overlaying the selected sources on the Hubble Image of the Bullet Cluster confirmed this worked: the removed galaxies were bright, red cluster members, while the remaining galaxies were fainter and more scattered in the background. An example of this visual check and overlaying the selections on the Hubble image using the ds9 software can be seen in Fig. 2.

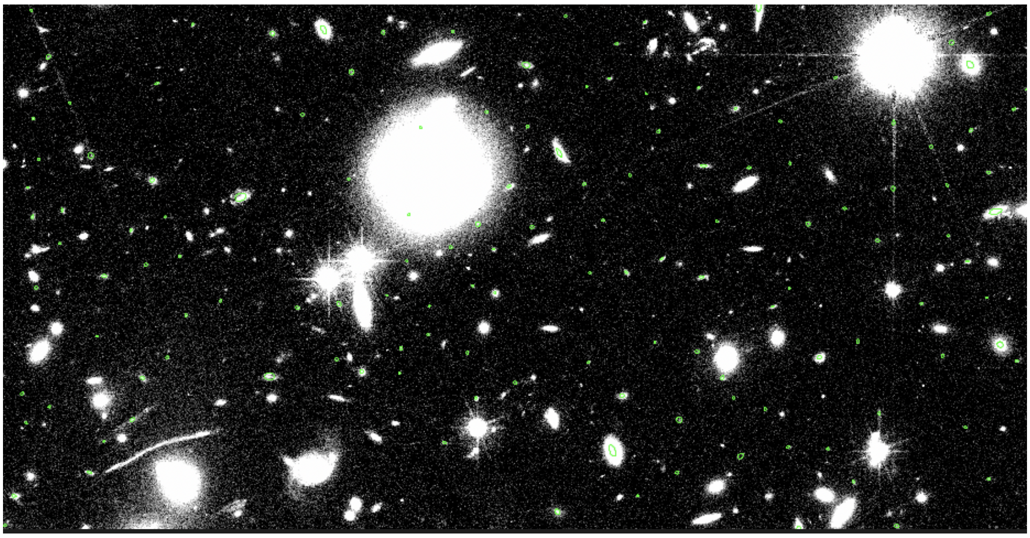

Figure 2: Selection of galaxies in the F814W filter using `ds9`, a visualisation tool, shown here for sources of a minimum radius of 4 pixels. Green contours mark the objects included after exclusions of stars and cluster members, highlighting the background galaxies used for lensing analysis.

3.1.4 Shape Measurement and Shear Catalogs

With a catalog of background galaxies, PyRRG measured galaxy ellipticities and calculated shear. The output was a shear catalog listing the distortion for each galaxy, with uncertainties. However, the results are still contaminated if all background galaxies are used. Not all background galaxies are equally useful. Some are too bright and likely cluster members, or foreground galaxies, or saturated objects. Some are too faint, with shapes dominated by noise. Very small galaxies make the shape measurement hard and unreliable. This demands applying a selection on the background galaxies. This selection is usually done manually based on educated guess and previous experience. However, this is not the most reliable way and does not necessarily produce optimized final results.

3.2 Optimisation

3.2.1 Parameter Cuts

In order to make a good selection of useful background galaxies, cuts have to be made on certain parameters. Exclusions should be applied with respect to the magnitude of the selected sources in order to consider a certain brightness range. Also, a selection of certain radii of background galaxies in order to consider their size is necessary. For this reason, cuts had to be applied on:

- **Minimum magnitude** to remove very bright galaxies.

- **Maximum magnitude** to remove faint galaxies.
- **Minimum radius** to exclude very small galaxies.

I then implemented functions to apply these cuts efficiently. Different parameter cuts affect the measurements differently. The minimum radius parameter directly influences the shear calculation itself, and therefore, the cuts have to be made before the calculations are done. To handle this, I created base shear catalogs with different minimum radius values but no magnitude cuts. The minimum and maximum magnitude parameters do not get involved in the process of shear calculations; therefore, cuts on these parameters can be applied after the shear catalog has been made. I implemented functions that apply minimum and maximum magnitude cuts to the base catalogs that were previously created. This saved time and made it easier to test many parameter combinations. It also made the process of creating new shear catalogs faster, more efficient, and systematic. The results obtained using different parameter cuts then had to be compared.

3.2.2 Shear and SNR Maps

For each parameter set, two kinds of maps were produced: shear maps and SNR maps. Shear maps show how the shapes of background galaxies are distorted by the cluster's gravity, while SNR (signal-to-noise ratio) maps show how strong and reliable that signal is compared to random noise. To make the SNR maps, the shear field is converted into a convergence map (which traces the mass) using a mathematical tool called a Fourier transform. This step introduces some uncertainty, so dividing the mass signal by its error gives the SNR. If the selection of galaxies was poor, the resulting maps showed fuzzy mass peaks and weak SNR. For example, Fig. 3 shows an SNR map with cuts that are not optimized, including noise and not revealing the mass patterns. With good selections, however, the maps reveal clear mass peaks, matching where the cluster's matter is located (Later seen in results, Fig 6).

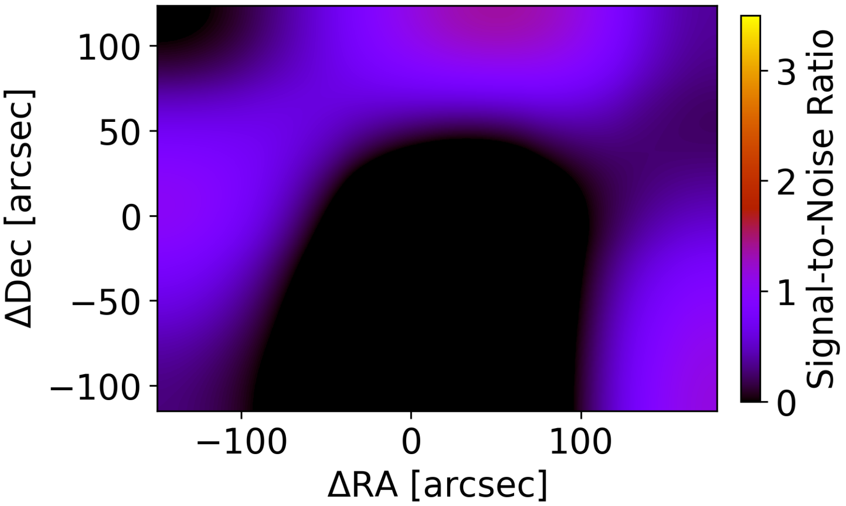

Figure 3: SNR map for the Bullet Cluster using non-optimized cuts (minimum radius = 2 pixels, minimum magnitude = 21.0, maximum magnitude = 24.5)

3.2.3 Choosing the Right Metric

At first, maps were compared using the maximum pixel value in the SNR map. But this was unstable, since one noisy pixel could dominate the results. Percentiles provided a better solution. For example, the 95th percentile measures the value above which the top 5% of pixels lie. This was less sensitive to outliers and provided a stable comparison metric.

3.2.4 Randomized Maps

The Fourier transform used to calculate the shear measurements and plot the SNR maps, tends to create some noise, especially in the corners of the maps. These are called boundary effects of the Fourier transform. To minimize the influence of such noise on the comparison of results, randomized shear catalogs were generated. For this, I created a function that randomly rotates the shear, removing the lensing signal but keeping the boundary effects. Averaging many random maps gave a baseline of what pure noise looks like. The SNR was then calculated in four ways:

- Shear signal divided by shear noise.
- Shear signal minus the average of random maps, divided by shear noise.
- Shear signal minus the average of random maps, divided by shear noise and the random maps' standard deviation.
- Shear signal divided by the random maps' standard deviation.

The most reliable was the second method: subtracting the random baseline and dividing by shear noise. This approach accounts for the influence of the randomized maps without over-correcting. When the different methods were compared visually, the maps evaluated with the second method not only scored better but also looked cleaner and more physically meaningful, with clearer mass peaks and less noise. By combining the 95th percentile metric and the subtraction of the mean of random maps, the results became more reliable and less influenced by outliers and noise.

4 Results

4.1 Parameter Exploration

With the metric in place, the 95th percentile value was first plotted by only varying one parameter at a time, for a specific range. These ranges were determined by relying upon the parameter cuts that had been determined manually and based on experience. The initial values of parameters were 3 to 4 pixels for minimum radius depending on the filter based on educated guess, 21 for the minimum magnitude, and 27 for the maximum magnitude. The fixed parameters were kept at this initial value every time.

- **Minimum magnitude (range: 18–21):** little effect, since very bright galaxies are rare.

- **Maximum magnitude (range: 24–29):** strong effect, as including faint galaxies increased the amount of source galaxies, but also noisy measurements. Therefore it's a trade-off.

- **Minimum radius (range: 1–6 pixels):** strong effect, since it directly influences shear calculations. This is because the minimum radius parameter does not only filter the catalog afterwards. To calculate the shear measurements, Gaussian weightings are applied to measure moments. That radius also controls how the Gaussian weighting is applied in the moment calculations, and therefore is used in the process and the calculations have to be redone if this parameter is changed.

Only integer values were initially explored. Afterwards, based on results, some ranges were studied in more detail. For example, since the maximum magnitude showed more sensitivity, and since it is a logarithmic value, a 0.1 step was chosen to investigate a smaller range of values around the peaks. Example plots can be seen in Fig. 4.

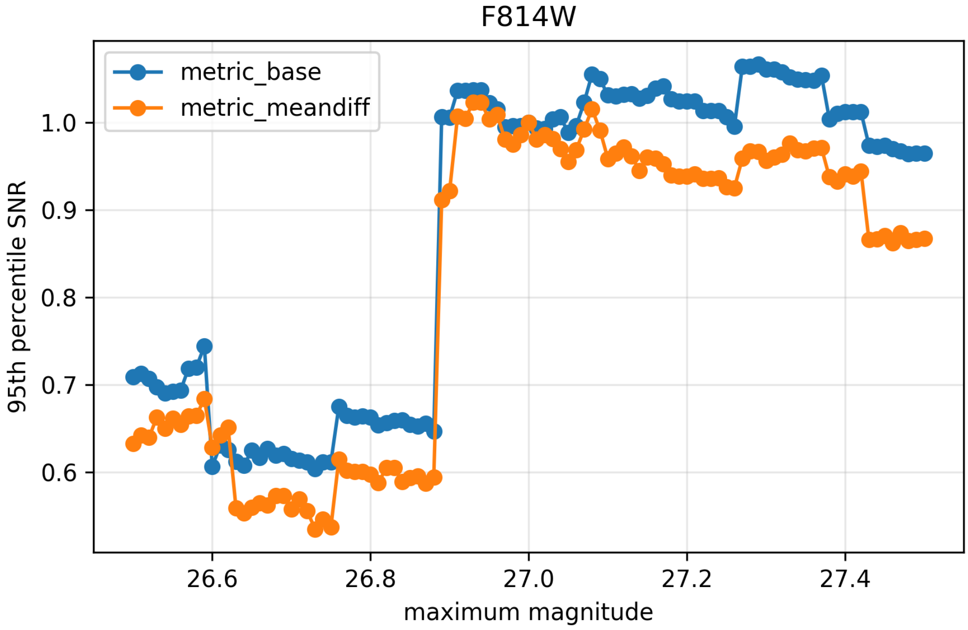

Figure 4: Comparison of two metrics for evaluating map quality in the F814W filter. The blue line shows the base 95th percentile SNR metric, while the orange line includes subtraction of the mean of randomized maps. The latter provides a more stable baseline and avoids fluctuations from noise.

4.2 Grid Search Optimisation

Considering that the maximum magnitude and the minimum radius have the strongest impact, a 2D grid search across minimum radius and maximum magnitude identified the best parameter combination. This plot provided us with the optimized value for the cut parameters, using the developed metrics. These optimized cuts were determined to be at:

- F606W: min radius = 4 pixels, min mag = 21, max mag = 25.5

- F814W: min radius = 4 pixels, min mag = 21, max mag = 26

These improved the SNR maps by ∼65% for the F606W filter and by 50% for the F814W filter.

The grid search done for the F606W filter and the F814W filter can be seen in Fig.5, where the maximum magnitude cuts are explored from 24.0 to 28.0 and the minimum radius cuts are explored from 2 pixels to 5 pixels. Each square represents the results using the 95th percentile and the mean difference metric.

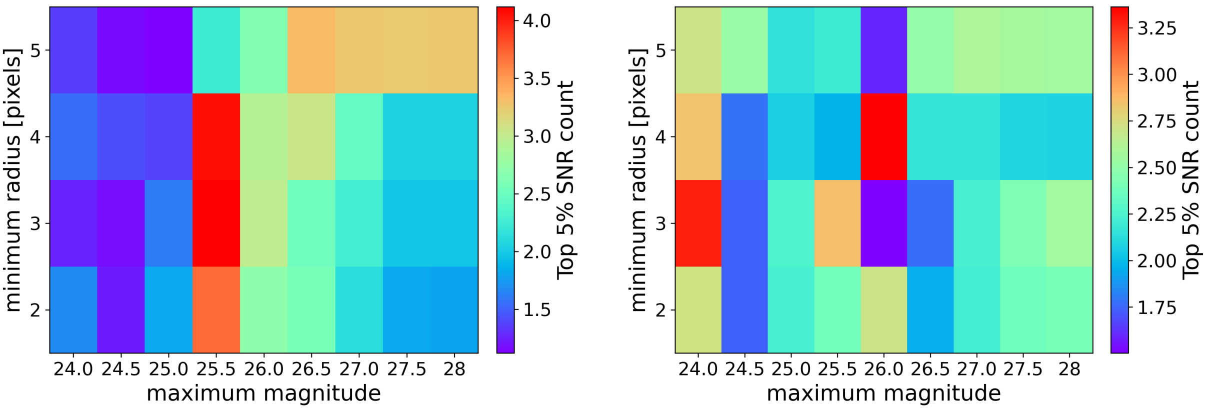

(a) F606W filter

(b) F814W filter

Figure 5: Heatmaps showing the effect of varying minimum radius and maximum magnitude on the 95th percentile SNR for both filters. Redder colors correspond to parameter sets with higher signal-to-noise, while bluer colors indicate poorer performance.

Fig. 6a shows the dark matter mass map that was generated with the best educated guess, manually and taking more time, while Fig. 6b shows the dark matter mass map generated with this method, which looks sharper and took less time to generate.

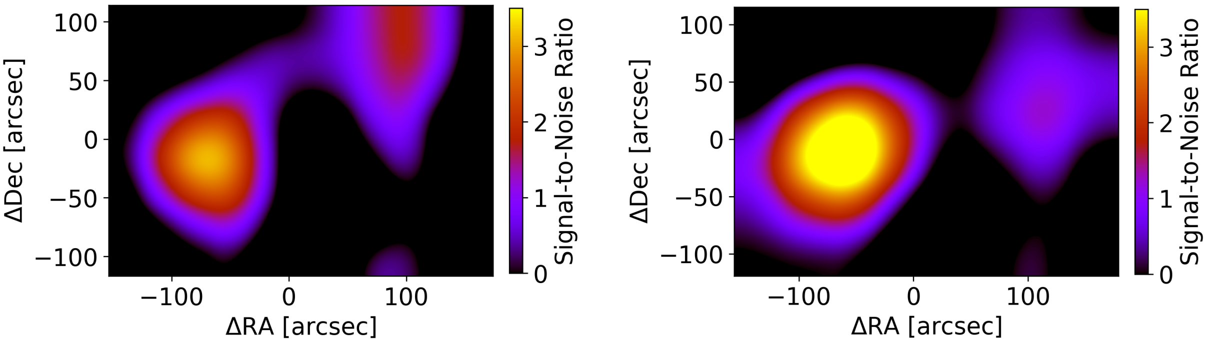

(a) Fiducial cuts (minimum radius = 4 pixels, maximum magnitude = 27)

(b) Optimized cuts (minimum radius = 4 pixels, maximum magnitude = 25.5)

Figure 6: Comparison of SNR maps for the Bullet Cluster with different parameter cuts. Subfigure (a) shows fiducial cuts based on educated guess, while (b) shows optimized cuts, where the mass peaks appear sharper and more reliable.

5 Discussion

The optimized cuts found through the 2D grid search improved the SNR maps by about 65% for the F606W filter and by about 50% for the F814W filter compared to the initial parameter guesses. This improvement was not only seen in the numerical metrics but was also clear in the visual quality of the maps (Fig.6). The optimization process demonstrated the benefits of a more automated approach: catalogs could be generated more quickly because the cuts were applied after catalog creation, reducing the number of full PyRRG runs required. This made parameter optimization substantially more efficient. At the same time, the use of randomized maps and percentile-based metrics allowed for objective comparisons, reducing noise bias and preventing individual outliers from dominating the results. Finally, excluding faint, small, or poorly measured galaxies sharpened the contours of the mass reconstructions, producing cleaner maps with more distinct mass peaks.

There are, however, limitations. The method was only tested on a single system, the Bullet Cluster, and the results may depend on the choice of filters (F814W and F606W). Moreover, the influence of other parameters, such as the noise cut, was not fully explored in the grid search. By adding more dimensions to the grid search, such other parameters can be taken into account more directly. Also, the lower magnitude threshold was determined to have little influence on the final results, however by adding this parameter in the grid search as well, the final outcome can be still improved, even if the influence is not as big as the studied parameters. This leaves open the question of how robust the optimized cuts are across different clusters and observational conditions. A more systematic study, covering a variety of clusters and filter combinations, would be needed to assess the generality of these results.

Looking ahead, this project could be extended by moving beyond grid searches to fully automated optimization schemes. Techniques such as gradient descent or other machine learning–based methods could explore parameter space more efficiently and adaptively, learning optimal cuts without the need for predefined grids. This would not only speed up the process further but could also uncover parameter combinations that are not easily found through manual tuning. Ultimately, such automation could streamline weak lensing analyses across many clusters, making studies of dark matter distributions more practical and consistent.

6 Conclusion

Dark matter is invisible, but its gravity allows us to detect it through weak lensing. The accuracy of lensing maps depends critically on selecting the right background galaxies. This research developed and tested an automated method for background galaxy selection and introduced reliable and robust metrics for evaluating the final SNR maps. The approach improved the quality of the final maps by 65% for the F606W filter and by 50% for the F814W filter, demonstrating clear advantages over manual parameter selection. This method was also faster than the manual method used in the past, by eliminating unnecessary calculations using a systematic method to create shear catalogs. This method can be applied to many clusters and scaled to large surveys such as Euclid, where automation will be essential. Having faster, more objective, and robust methods and tools for this selection, enables uniform and coherent mapping when dealing with several clusters. In the future, more parameters can be taken into consideration at the same time, making the grid search multi-dimensional. Also, by replacing the grid search with smarter algorithms, a faster approach can be implemented, by exploring the parameter space in a way that it focuses on the desired area automatically. For example, using a gradient descent, or an MCMC algorithm. This will allow an even faster, more reliable, and more exact approach when automatically choosing the parameters, allowing the least amount of manual interference made when mapping dark matter on a larger amount of data.

Availability of data and software code

Software: `https://github.com/DavidHarvey1986/PyRRG`
Data: Hubble Space Telescope Bullet Cluster archive