Universidad de La Laguna

Inferring the 3-Dimensional Shapes of Galaxies Using Artificial Intelligence

Robert Ramón Hidalgo Mahmud

Trabajo de Fin de Grado

Grado en Física - Facultad de Ciencias

Supervisado por:

Dr. Jesús Falcón Barroso

Dr. Marc Huertas-Company

INSTITUTO DE ASTROFISICA DE CANARIAS · IAC ·

Julio 2023

Agradecimientos

Gracias a mis tutores, Jesús y Marc, por tener paciencia conmigo y no darme por perdido. También gracias a todos aquellos profesores, científicos y divulgadores que me inspiraron a emprender este camino. Aunque no ha sido fácil, lo volvería a elegir siempre.

Por supuesto, gracias a mi familia y amigos por su apoyo incondicional. Gracias a mi hermana por motivarme en los días malos y celebrar mis éxitos como los suyos propios. Y, sobre todo, gracias a mi madre por creer en mí más que yo mismo e inspirarme a seguir adelante.

Do not go gentle into that good night,
Old age should burn and rave at close of day;
Rage, rage against the dying of the light.

Dylan Thomas

Resumen

Entender la estructura de las galaxias es crucial en astronomía, ya que proporciona una valiosa información sobre su formación y evolución. Las galaxias en etapas tempranas del Universo ($z \sim 2,5$) eran típicamente ~ 3 veces más pequeñas y $\sim 2,5$ veces más densas que aquellas que vemos actualmente (Trujillo et al. 2006). Además, a $z \sim 2-3$, la mayoría de galaxias presentan una distribución irregular en sus perfiles de luminosidad (Huertas-Company et al. 2015). Sin embargo, hoy en día las galaxias conforman la secuencia de Hubble (1926), que separa entre galaxias elípticas y espirales, además de las irregulares. Para enlazar estas dos poblaciones es necesaria una mayor comprensión de los mecanismos que conducen a la transformación de la forma de las galaxias. En efecto, existe correlación entre la estructura de una galaxia y aspectos como su actividad en formación de estrellas (Conselice 2003), parámetros como su luminosidad total y su radio (Caon et al. 1993), o si procede de una fusión entre dos galaxias (Toomre & Toomre 1972).

Aunque la secuencia de Hubble sigue siendo influyente en la clasificación de galaxias, esta solo atiende a su elipticidad aparente. El problema reside en la manera en la que observamos las galaxias desde la Tierra. Las imágenes obtenidas en los telescopios son proyecciones bidimensionales en el cielo de objetos en tres dimensiones. Por lo tanto, su apariencia depende de la orientación de la galaxia respecto a nuestra línea de visión. La naturaleza del problema da lugar a la existencia de degeneración en las imágenes, lo que implica que hay ocasiones donde la forma original de la galaxia no puede ser totalmente recuperada. Históricamente, obtener las elipticidades intrínsecas de las galaxias, independientes del observador, ha representado un desafío significativo. Los enfoques iniciales se basaban en métodos geométricos (Hubble 1926, Sandage et al. 1970), fotométricos (De Vaucouleurs 1948), estadísticos (Lambas et al. 1992) o cinemáticos (Binney 1985, Weijmans et al. 2014). Métodos como aquellos basados en CAS (concentración, asimetría y suavidad), aunque sean satisfactorios simplificando el problema a unos pocos parámetros, también implican una pérdida de información contenida en los propios píxeles. En este trabajo, hemos adoptado un método de Deep Learning (DL) que nos permite considerar y utilizar cada característica presente en los datos.

Primeramente, hemos generado un conjunto de datos de 10.000 galaxias simuladas utilizando una combinación del perfil $R^{1/n}$ de Sérsic para el plano de la galaxia (x_{gal} - y_{gal}) y una función gaussiana para la componente z_{gal}, siguiendo los pasos de Price et al. (2021). Posteriormente, hemos proyectado cada galaxia con cuatro diferentes inclinaciones, $i = (0^o, 30^o, 60^o, 90^o)$, para simular cómo serían vistas desde los telescopios. El resultado es un conjunto de datos final compuesto por 40.000 imágenes de galaxias. Debido a su construcción, las imágenes finales tienen 64x64 píxeles. Guardar las imágenes en el formato HDF nos ha permitido etiquetar cada imagen con sus cinco correspondientes variables asociadas al modelo: la intensidad (I_e) a un radio efectivo (R_e), que encierra la mitad de la totalidad de la luz del modelo; el índice de Sérsic (n), que configura la

concentración del perfil; la elipticidad intrínseca (q), que mide el achatamiento de la estructura de la galaxia; y la inclinación (i) respecto a nosotros.

A continuación, hemos construido y entrenado una Red Neuronal Convolucional (CNN), una arquitectura específicamente diseñada para el procesamiento de imágenes. El aprendizaje realizado por esta red ha sido supervisado, es decir, le hemos proporcionado las imágenes etiquetadas con su parámetro de achatamiento, q, cuya obtención es el objetivo final del entrenamiento. Más específicamente, la estructura de nuestra CNN comprende una capa de entrada, tres capas convolucionales, dos capas densas (también llamadas totalmente interconectadas) y la capa de salida. Esta última genera una distribución de probabilidad gaussiana, cuya salida es la media y la desviación estándar para cada predicción de q. La intención es cuantificar el grado de incertidumbre en la inferencia de la red.

El primer paso consistió en determinar una configuración óptima para la CNN. Entrenamos la red neuronal utilizando todas las proyecciones y evaluamos la precisión de la predicción obtenida hasta que fuera aceptable. Seguidamente, un aspecto importante fue analizar el comportamiento del modelo según la proyección dada. El estudio concluye que a la red le resulta difícil estimar con precisión los valores de q en galaxias con inclinaciones más bajas (vista de cara), mientras que la precisión mejora a medida que la inclinación aumenta (vista de canto). Esto indica que la CNN tiene complicaciones al lidiar con la degeneración en las imágenes.

No obstante, en la fase final del estudio se probó el modelo proporcionándole una única proyección aleatoria, simulando lo que ocurre en la realidad. Se observó que la CNN puede superar los problemas de degeneración si se le proporciona un conjunto de datos lo suficientemente grande, lo que conduce a predicciones significativamente precisas. Los hallazgos encontrados sugieren que las Redes Neuronales Convolucionales tienen el potencial de recuperar de manera efectiva las formas tridimensionales reales de las galaxias, por lo que es justificable seguir trabajando en esta línea de investigación.

Abstract

Obtaining the intrinsic shapes of galaxies provides insights into their formation and evolution. However, deriving the intrinsic shapes from the 2D data observed from Earth remains an open problem. To tackle this, we propose utilising Artificial Intelligence (AI), specifically a Convolutional Neural Network (CNN), as a promising solution. In our study, we generated a dataset of 40,000 galaxy images for training and testing purposes. The CNN was designed to extract the intrinsic ellipticity value (q) from the 2D projections of each galaxy. The study acknowledges the presence of image degeneracy problems. Nonetheless, the network's predictions were ultimately accurate, indicating the potencial success of implementing this method in the field.

Contents

Introduction

En esta sección inicial, se justifica la importancia del problema a resolver. Expondremos la transformación experimentada por las galaxias a lo largo del tiempo y resaltaremos la relevancia de conocer su estructura para investigar su evolución.

Se mencionarán los métodos tradicionales para calcular la elipticidad intrínseca y cómo las redes neuronales pueden proporcionar un salto cualitativo en este aspecto.

Since galaxies are usually considered as the blocks that form our universe, understanding their formation and evolution is one of the most important problems in astronomy. Trujillo et al. (2006) demonstrated that galaxies in the relatively early stages of the universe ($z \sim 2.5$) were typically $\sim$ 3 times smaller and $\sim$ 2.5 times denser than those we see today. Furthermore, Huertas-Company et al. (2015) showed that at $z \sim 2-3$ the majority of galaxies exhibit an irregular distribution of their light profiles. But on the other hand, nowadays they belong to the Hubble sequence (Figure 1) that separates clearly elliptical and spiral galaxies, besides irregular ones (Hubble 1926). In order to link these two populations, it is essential to obtain a deeper knowledge of the mechanisms that lead to the transformation of galaxy shapes.

1.1 Galactic Structure and Evolution

Galaxies undergoing star formation exhibit diverse morphologies and structures compared to passive galaxies, such as clumpy spiral arms, regions of concentrated star formation and central intense starbursts. Upon observing local galaxies, we find that galaxies undergoing star formation are more irregular and asymmetric compared to those without ongoing star formation. Research has also shown a correlation between not only qualitative but also quantitative measurements of structure and the star formation rate within galaxies (Conselice 2003).

Structure is also an important way of identifying whether two galaxies have recently combined. The resultant structures of major mergers, occurring when the galaxies have a comparable mass, are distorted and Toomre & Toomre (1972) demonstrated that peculiar galaxies arise from such processes. Moreover, the shape of a galaxy light profile is influenced by its scale or mass. As Caon et al. (1993) demonstrated, profiles of ellipticals are well correlated with global galaxy parameters, such as total luminosity and radius.

1.2 Intrinsic Ellipticity and Traditional Methods

The Hubble sequence remains influential in galaxy classification. In this context, ellipticity quantifies the elongation or roundness of a galaxy, and elliptical galaxies are classified from nearly spherical (E0) to highly elongated (E7).

However, the Hubble classification only attends to apparent ellipticity (Figure 3). The problem lies in the way in which we can

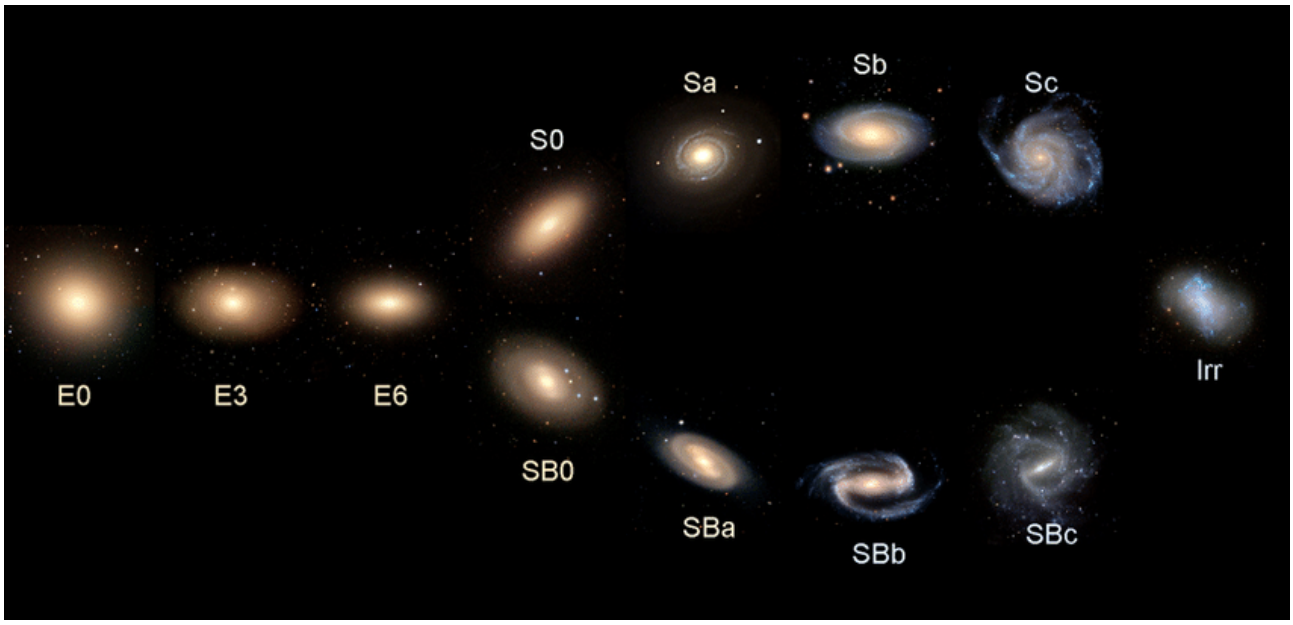

Figure 1: Hubble sequence. 'E' denotes Ellipticals, 'S' and 'Sb' indicate Normal and Barred Spirals, and 'Irr' stands for Irregular.

observe galaxies. The data we gather from telescopes is a 2-dimensional projection in the sky plane of objects in 3 dimensions. Hence the observed appearance depends on the galaxy's orientation, i, relative to our line of sight (Figure 2). The nature of the problem leads to the presence of image degeneracy, which implies that there are occasions when the original shape cannot be fully recovered. We will discuss this issue later in this work.

Our focus centers in the intrinsic ellipticity of galaxies, characterised by the parameter q, independent of the observational effects. The first calculations were derived from the relationship of the measured apparent ellipticity to the inclination angle (Hubble 1926). Later on, Sandage et al. (1970) used these projected axial ratios to estimate the roundness and thickness of the discs in spiral galaxies.

With the advancements in photometry, quantitative measurements of light distributions in galaxies could be obtained. De Vaucouleurs (1948) utilised them to demonstrate that the light profiles of massive elliptical galaxies show approximately the same light distribution, known as the de Vaucouleurs profile. This idea was further developed by other researchers as Sérsic (1963), who demonstrated that a broader range of light distributions could be fitted to galaxy light profiles.

Afterwards, due to the increase in data samples, Lambas et al. (1992) could study galaxies statistically and fit the ellipticities to a Gaussian distribution. Kinematical models have also been used to determine the 3-dimensional shape of galaxies, as Binney (1985) did. Using photometric and kinematic data, Weijmans et al. (2014) related galaxy geometry also to its rotational properties as well. They found out that slow rotators have a rounder shape, whereas fast rotators have a flatter structure and could be associated with the population of spiral galaxies.

1.3 New Approaches and Deep Learning

The advancement of big data and computing resources has facilitated the development of new machine learning techniques for inferring 3D shapes. Traditional methods rely on parameters such as concentration, asymmetry, and smoothness (CAS) to describe galaxies (Conselice 2003). While these parameters simplify the problem by reducing it to a few key characteristics, they also result in a loss of information contained in the pixels themselves.

In contrast, Deep Learning (DL) is a nonlinear learning process that extracts the most pertinent information from the data. It identifies the optimal features that exhibit the strongest correlations with the target prediction (Huertas-Company 2015). The utilisation of this technique, as demonstrated in studies such as Huertas-Company et al. (2015) and Dominguez-Sanchez et al. (2018), has allowed researchers to quantitatively analyse the morphologies and structural characteristics of galaxies with unprecedented precision.

Furthermore, cosmological simulations have made significant progress in accurately reproducing the diverse range of galaxy morphologies observed in the local universe, as demonstrated by Huertas-Company et al. (2019). Combining DL methods with these simulations can provide valuable insights into the recovery of the intrinsic shapes of galaxies.

Objective

En esta sección, presentaremos de manera concisa el objetivo de este trabajo: recuperar la forma tridimensional (la obtención de q) de una galaxia modelada por nosotros, a partir de una imagen bidimensional.

In this project, our aim is to determine whether the advancements of Artificial Intelligence (AI), specifically utilising a Convolutional Neural Network (CNN), can accurately reconstruct the three-dimensional structure of a galaxy from a two-dimensional input.

To achieve this goal, we need to find a proper mathematical model to describe a galaxy's 3D light profile and obtain the 2D images by integrating the flux along a line of sight from Earth. Creating a considerable dataset is important in order to reproduce as many cases as possible.

Afterwards, we will build our network and become familiar with its parameters. We have to find the best way to provide it with the data to train the CNN on recovering the original 3D shape (obtention of q). Finally, we will discuss its predictive power, limitations, and possible enhancements.

Methodology

En esta sección, primero se explican los pasos seguidos para obtener las imágenes: generar el perfil tridimensional y colapsarlo a

(a) NGC 3982 (b) NGC 2008 (c) NGC 4013

Figure 2: A galaxy may appear in any orientation from $i = 0^o$, also called face-on (left), to $i = 90^0$, edge-on (right). Credits: ESA/Hubble and NASA.

lo largo de una línea de visión para obtener la imagen bidimensional, como se vería desde la Tierra. Se detallan las ecuaciones utilizadas y las aproximaciones realizadas.

Finalmente, se expone la arquitectura de nuestra red neuronal, explicando la función de cada capa utilizada y los hiperparámetros con los que se ha configurado. También se describe cómo se han utilizado los datos generados previamente y se explica el proceso de entrenamiento de la red.

3.1 Galaxy Model and Image Generation

In this project, we have followed a similar method to the one proposed by Price et al. (2021) for simulating the galaxy light profiles (Figure 4). Firstly we need to build a sky coordinate system grid $(x_{sky}, y_{sky}, z_{sky})$ and apply a coordinate system change to the galaxy's intrinsic coordinates. The galaxy coordinate system $(x_{gal}, y_{gal}, z_{gal})$ has been defined such that x_{gal} and y_{gal} are the positions within the galaxy midplane and z_{gal} is the vertical position. Here $R_{gal} = \sqrt{x_{gal}^2 + y_{gal}^2}$ is the radial distance from the rotational axis. Being i the inclination angle relative to the line of sight and the major axis (y_{gal}) oriented anticlockwise at an angle PA, the transformation is given by:

$$\begin{aligned} x_{gal} &= x_{sky}\, cos(PA) - y_{sky}\, sin(PA)\, cos(i) \\ &\quad -z_{sky}\, sin(i) \\ y_{gal} &= x_{sky}\, cos(PA) + y_{sky}\, sin(PA) \\ z_{gal} &= x_{sky}\, cos(PA) - y_{sky}\, sin(PA)\, sin(i) \\ &\quad +z_{sky}\, cos(i) \end{aligned} \tag{1}$$

Subsequently, we have used a Sérsic's (1963, 1968) $R^{1/n}$ model as a basis. A common approach is to use a single $R^{1/n}$ profile when there is insufficient resolution in an image to accurately distinguish between the disc and bulge components (Graham & Driver 2005). This approximation serves as a good starting point. Sérsic intensity profile is given by:

$$I(R) = I_e\, exp\left\{-b_n\left[\left(\frac{R}{R_e}\right)^{1/n} - 1\right]\right\} \tag{2}$$

where I_e is the intensity at the effective radius R_e that encloses half of the total light from the model. The parameter n is the Sérsic index that configures the concentration of the profile, and

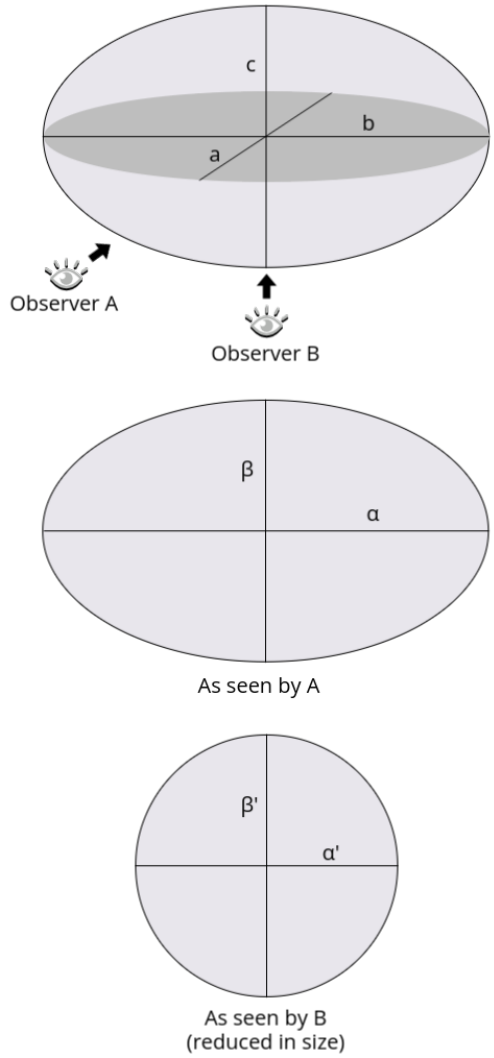

Figure 3: An oblate spheroidal galaxy with axis $a = b$ and $c < a$. Considering $c/a = 0.6$, apparent shape for observer A resembles an E4 galaxy ($\beta/\alpha = 0.6$), meanwhile for observer B it looks like an E0 ($\beta'/\alpha' = 1$).

the constant bn, as explained in Erwin (2015), is the solution of the transcendental equation:

$$\Gamma(2n) = 2\gamma(2n, b_n)\ , \tag{3}$$

where $\Gamma(a)$ is the gamma function and $\gamma(a, x)$ is the incomplete gamma function. In spite of the existence of analytical approximations (Ciotti & Bertin 1999, MacArthur et al. 2003), in this research we have solved it numerically for each n.

We have made the assumption that specific mass galaxy components emit light at a constant mass-to-light ratio, Υ. Otherwise, it would have been necessary to consider the specific characteristics and properties of each component individually. Price et al. (2021) used as a mathematical model a two-dimensional Sérsic flux distribution within the x_{gal} - y_{gal} midplane, combined with a Gaussian profile in the z_{gal} axis:

$$F(R_{gal},z_{gal}) = \tag{4}$$
$$I_e\ exp\left\{-b_n\left[\left(\frac{R_{gal}}{R_e}\right)^{1/n} - 1\right]\right\}$$
$$\times\ exp\left\{-0.5\left(\frac{z_{gal}}{h_z}\right)^2\right\}\ ?$$

where the Gaussian profile has a width related to the flattening q and R_e given by $h_z = q\ R_e/1.177$.

Once we get the 3D profile in the intrinsic coordinate system, we have to collapse the data along a determined line of sight i. The aim is to obtain the integrated flux as it would be seen from Earth. These 2D data are transformed into a 2D magnitude image set by applying the logarithm.

We have made two final considerations. Firstly, we have set $PA = 0^0$, since this parameter represents a 2D rotation that, in principle, does not provide any additional information to the network training process. The final consideration is that the parameters (I_e, R_e, n, q) are uniformly distributed within

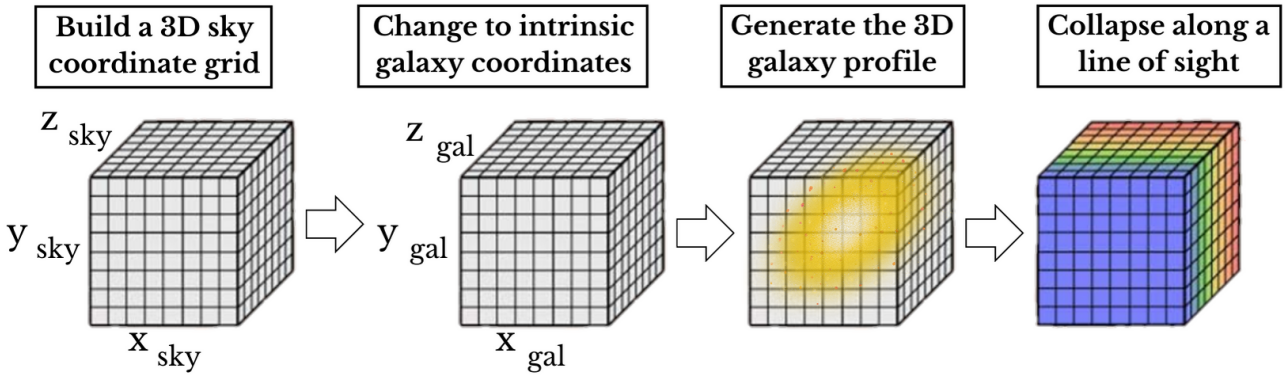

Figure 4: Visual representation of the image generation process.

the given range of values:

$$I_e \in [0.1, 1] \qquad n \in [0.5, 2] \qquad (5)$$
$$R_e \in [5, 30] \qquad q \in [0.2, 1]$$

We are mainly interested in the overall behaviour of both the galaxy model and artificial intelligence. Consequently, we have not given excessive attention to the precise units in this work. However, we have selected a range of values for each parameter that we deem appropiate.

3.2 Network Architecture

Artificial Neural Networks (ANNs) are computational systems that consist of interconnected computational nodes, the neurons, organised into layers. Their purpose is to learn and make predictions from input data, imitating the functioning of biological nervous systems.

The input of an ANN, a multidimensional vector, is fed into the input layer, which then transmits it to the hidden layers. The hidden layers make informed decisions based on the input received from the previous layer and evaluate how internal modifications impact the overall output (Figure 5). This is the process called training. The technique referred to as Deep Learning (DL) involves the combination of multiple hidden layers.

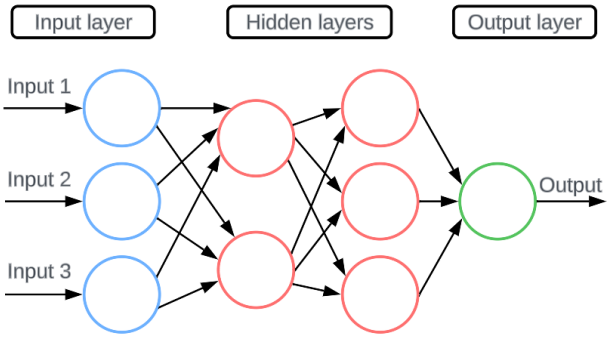

Figure 5: Basic structure of a four-layer ANN.

In this context, Convolutional Neural Networks (CNNs) are specifically designed to process image inputs, using convolutional layers responsible for learning and detecting local patterns. The layers in a CNN consist of neurons organised in three dimensions: height-width as the input spatial dimensionality, and depth, which is an activation volume's third dimension (such as red, green and blue in a RGB image).

The upcoming discussion will provide an overview of the overall design and hyperparameters of our network (Figure 7).

The code for this CNN is implemented in Python, making use of the TensorFlow and Keras libraries.

- Input: Defines the input shape of the model. In our case, it could be 64x64x4 if using the four galaxy projections or 64x64x1 if using only one of them. We will cover this topic later.

- Convolutional layers: Apply a 2D convolution operation with the following main hyperparameters:

  – Filters: The number of filters (or kernels) to apply in the convolution. Each filter learns to extract specific features from the input.

  – Kernel size: Specifies the height and width of the filter window.

  – Activation function: Introduces non-linearity to the network and allows it to learn complex relationships in the data.

  In our CNN, we have sequentially combined three convolutional layers. The initial layer comprises 16 filters, each with a size of 2x2. In the subsequent layer, the number of filters is doubled, and in the third layer, it is tripled.

  Every layer utilises a Rectified Linear Unit (ReLU) activation function (Figure 6). It maps negative values to zero and keeps positive values unchanged, mathematically defined as $f(x) = max(0, x)$. This function is still widely used, despite its simplicity. Deep networks with ReLUs are more easily optimised than networks with another kind of activation function, such as $sigmoid$ or $tanh$ (Ramachandran et al. 2017).

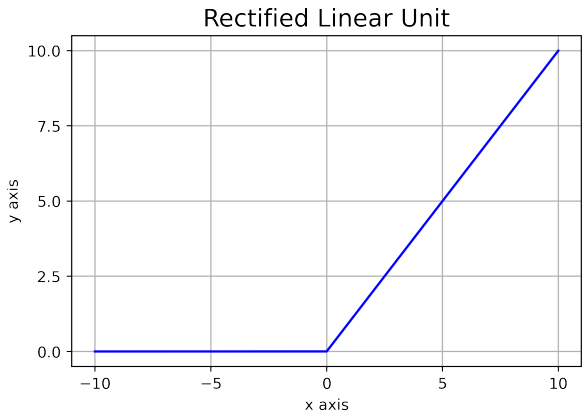

Figure 6: ReLU activation function.

- Pooling layers: Are used to downsample the spatial dimensions (width and height) of the input volume, reducing its size while retaining important information. We have used $max\ pooling$, where the maximum value within each pooling region is retained. In our CNN, every layer divides the input into 2x2 regions, and reduces each region to a single value. Setting a pooling layer after each convolutional layer has allowed us to reduce the number of training parameters from 22.8 million to 300,000.

- Dense (or fully connected) layers: Connect every neuron from the previous layer to each neuron in the posterior layer.

  – Flatten layer: To input data into these layers, it is first necessary to convert the multidimensional feature maps into a single continuous vector.

  – First dense layer: Is composed of 128 neurons and connects the convolutional to a dropout layer.

  – Dropout layer: Applies dropout regularisation to the previous layer. It helps prevent overfitting, which occurs when a network performs extremely

well on the training data but fails to generalise effectively to unseen data, by randomly setting a fraction of the input units to 0 during training time. In this case, it has been set to 0.3 (30% of the input units are set to 0).

  - Second dense layer: Is more specific and determines the number of parameters required for modelling a one-dimensional independent normal distribution.

- Output layer: Is responsible for producing a probability Gaussian distribution, whose output is the mean and standard deviation for each q prediction. The intention is to quantify the degree of uncertainty in the network's inference.

Once the structure is established, the training is carried out. As mentioned previously, we have labelled the images with their parameters. Therefore, during training, we fed the network with each galaxy image along with its q value. This parameter serves as a target value, which means we have followed a supervised learning process. To prevent overfitting, the available data is randomly divided into a training set (used for training the network), a validation set (used to optimise the model during the training process and assess its performance), and a test set (used to evaluate the final performance of the model). From the total dataset, we have set aside 20% of the images for testing; and from the 80% for training, 20% has been allocated for validation. Normalising the data is also crucial to preventing bias within the neural network. When certain features have larger magnitudes than others, there is a risk that the neural network may assign disproportionate importance to those features. By normalising the data, all features contribute more equally to the decision-making process of the model.

During the training process, as the model processes the data through its layers, it compares the outputs to the true target labels using a loss function. In this case, we have employed the $negloglik$ function, representing the negative logarithm of the likelihood. This function measures the disparity between the model's predictions and the actual data. By minimising this function, which is equivalent to maximising the likelihood, the network finds the optimal configuration for the model. The number of epochs, which is a hyperparameter, determines the number of times the model repeats this process and goes through the entire training dataset. In our case, we have specified 100 epochs for each distinct training method, facilitating a comparative analysis.

Results

Aquí mostraremos los resultados en base a la metodología que hemos expuesto anteriormente. Se presentará una muestra del conjunto de imágenes generadas. Asimismo, presentaremos las gráficas obtenidas a partir de las diferentes entradas que hemos proporcionado a la red durante el entrenamiento: cuatro proyecciones, una proyección (para cada inclinación), y una proyección aleatoria.

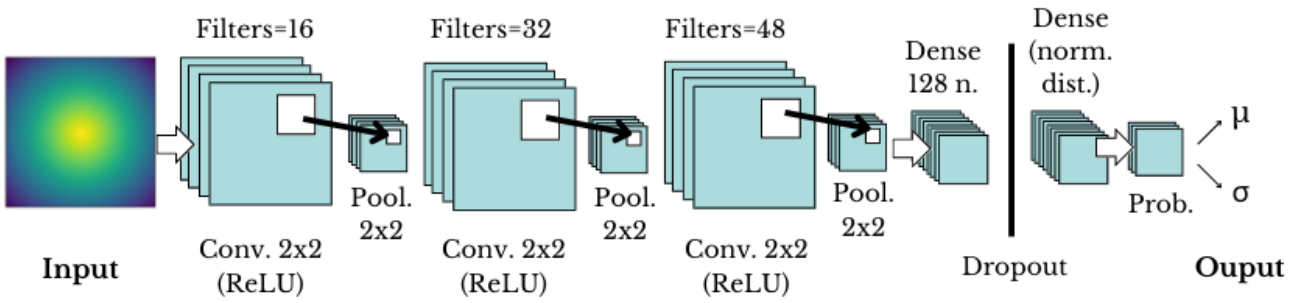

Figure 7: Visual representation of the layers that comprise the CNN architecture implemented in this project.

4.1 Dataset

We have generated a sample of 10,000 galaxies and taken four projections of each one with $i = (0^o, 30^o, 60^o, 90^o)$. The resulting set of 40,000 images was saved in HDF format, enabling the labelling of the images with their corresponding galaxy parameters. The coordinate definition is such that final images have 64x64 pixels.

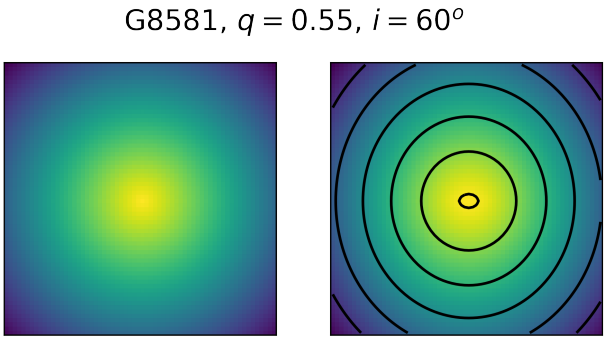

Figure 8: Example galaxy image with and without isophotes lines.

For an easier visualisation of the galaxies, the displayed images will include isophotes (Figure 8), which represent the contours of constant surface brightness.

In Figure 9 we show how a galaxy would appear in the sky, depending on its inclination to our line of sight; and in Figure 10, four randomly selected galaxies. Obtaining a sufficiently large and diverse dataset is important for our CNN, since it enables better generalisation and reduces overfitting.

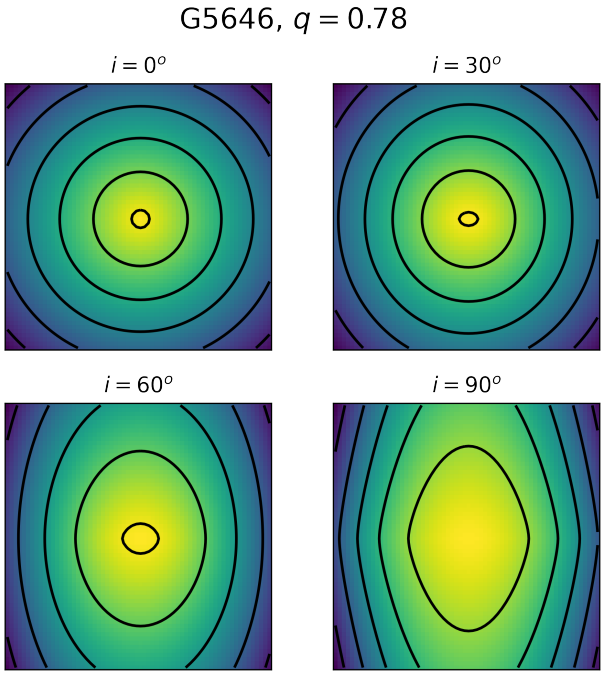

Figure 9: Different projections as seen in the sky of one of the modelled galaxies.

4.2 Network Training

To ensure accurate results, the initial focus was on optimising the CNN. To achieve this, training was conducted using all available projections. Consequently, in this particular phase of the study, the input dimension was set to 64x64x4. Results on the predictions of q

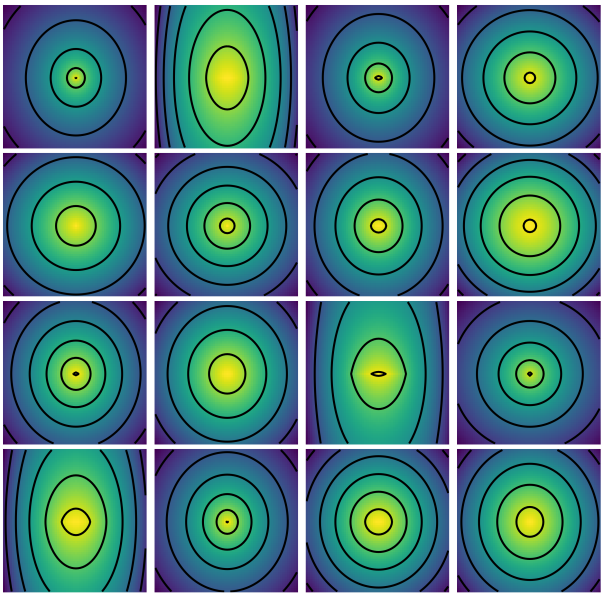

Figure 10: Example of random galaxy images.

(with its respective standard deviation) on the test set are shown in Figure 11.

Following that, we examined the impact of utilising individual projections for each modelled galaxy, resulting in an input dimension of 64x64x1. The initial approach involved considering each of the four projections $i = (0^o, 30^o, 60^o, 90^o)$ independently. Comparative results of the loss function can be observed in Figure 12, while Figure 13 displays the prediction results.

Finally, in an effort to create a more realistic and accurate model that resembles real-life scenarios, we implemented a random projection approach for each image in the CNN (yet with a 64x64x1 input dimension). To achieve this, we assumed that galaxy projections follow a uniform distribution. The training and prediction outcomes are depicted in Figure 14.

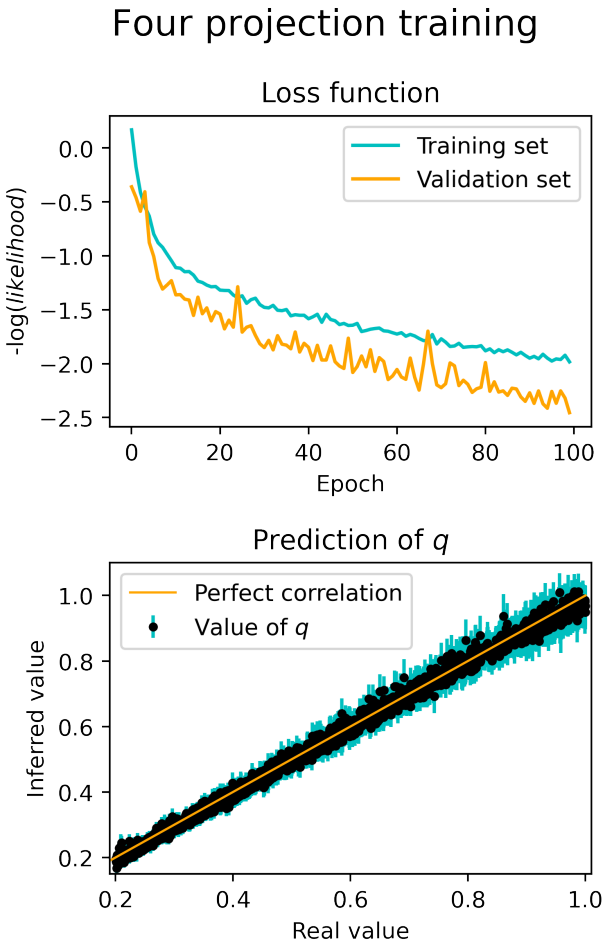

Figure 11: Results of the training using 4 projections.

Discussion

En esta sección, analizaremos en primer lugar la capacidad predictiva de la red. Se expone que la red exhibe un mejor rendimiento en inclinaciones elevadas (vista de canto). Al evaluar la red en un escenario más real, los resultados obtenidos respaldan la efectividad de la red neuronal en su labor de extraer la estructura intrínseca de las galaxias, siempre que se le proporcione una cantidad lo suficientemente grande de imágenes.

Además, se exploran las limitaciones del modelo, incluyendo las complicaciones que

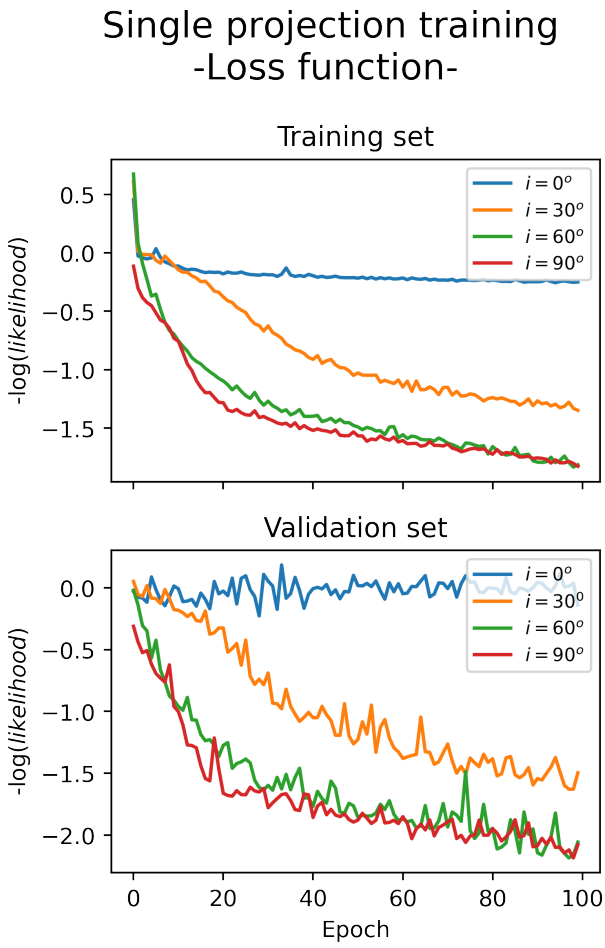

Figure 12: Comparative impact of orientation on loss function during training.

surgen al tratar con imágenes degeneradas y las dificultades que el modelo enfrenta al capturar el flujo de las galaxias muy inclinadas. Se presentan propuestas de soluciones y se sugieren posibles direcciones futuras para el desarrollo del proyecto.

5.1 Predictive Capacity

As illustrated in Figure 11, the initial objectives of refining the model have been successfully achieved. The training results obtained using all projections demonstrate good performance. The inferred values of q exhibit a reasonable level of accuracy, as

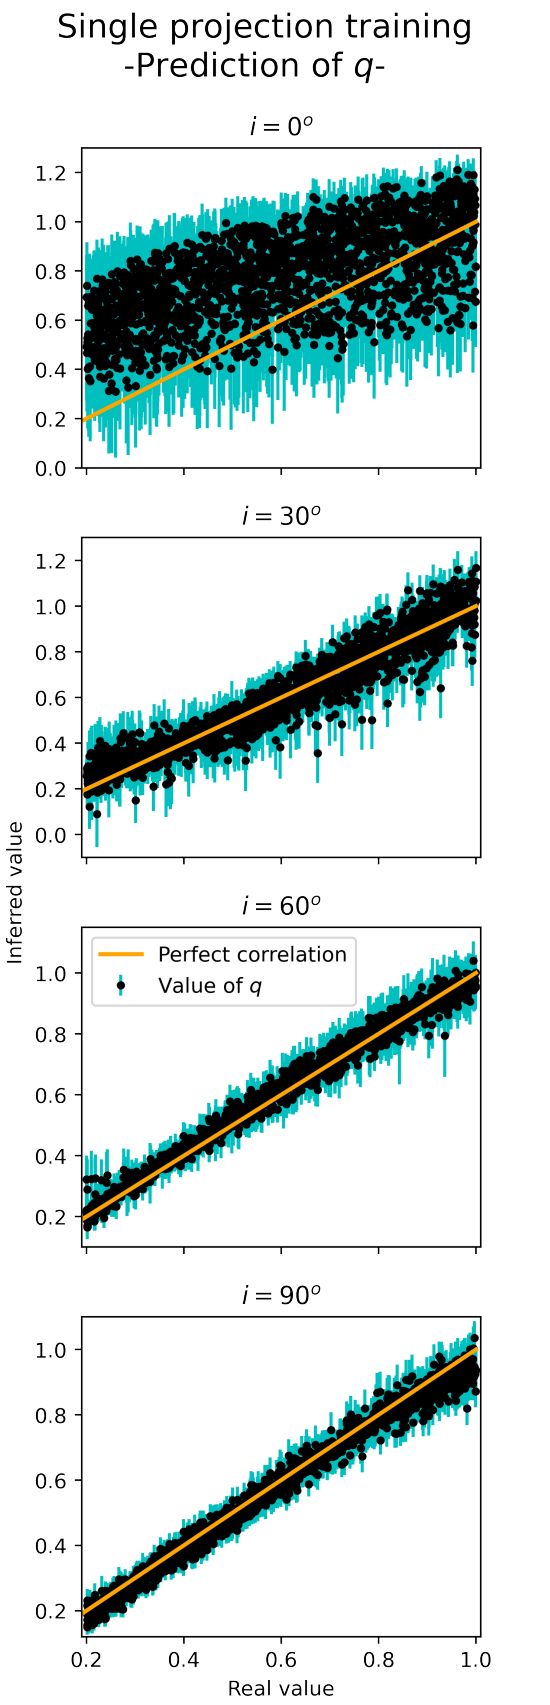

Figure 13: Comparative predictive capacity for each orientation.

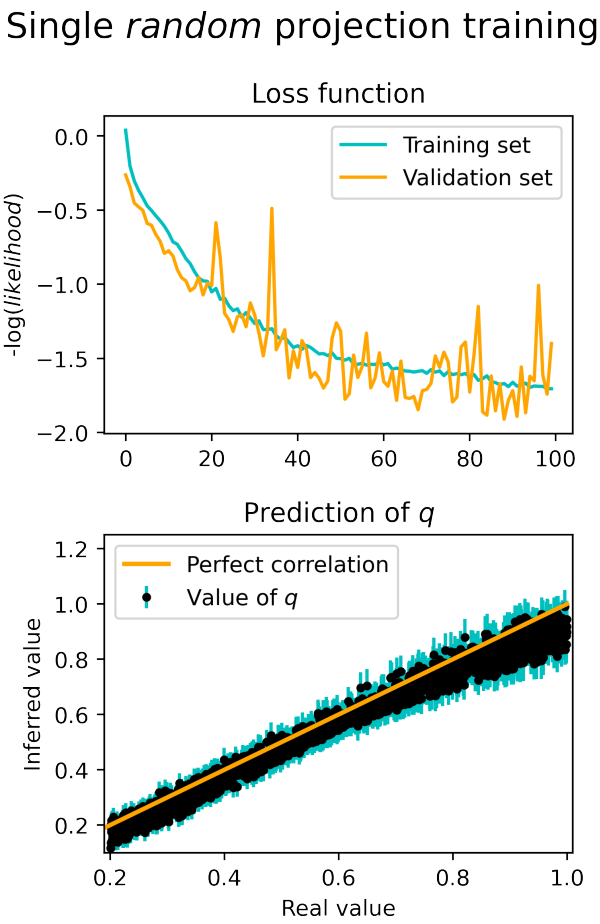

Figure 14: Prediction for each orientation.

they closely align with the ideal correlation line. These findings are in line with our initial expectations when utilising four projections, considering the abundant information available for inferring intrinsic galaxy shapes provided to the network. The successful minimization of the loss function indicates that we have achieved an acceptable network configuration for subsequent work.

The challenges emerge when training with only one projection. With the $i = 0^o$ projection (Figure 12), the loss function of the training set stabilises around a relatively high value, indicating unsatisfactory training performance. Similarly, the validation set loss exhibits a persistent pattern of high values, signalling that the network struggles to make accurate predictions. As depicted in Figure 13, the q values deviate significantly from the ideal results.

As the inclination increases, there is an improvement in the outcomes. At $i = 30^o$, both the validation and training set losses exhibit a more desirable behaviour, characterised by a logarithmic trend. Figure 12 demonstrates a consistent decrease in the loss function with each epoch, indicating the model's progression. Additionally, the predictions are better than those obtained at $i = 0^o$ (Figure 13).

At $i = 60^o$ and $i = 90^o$, the network training demonstrates a significant improvement. The loss function continues to decrease, and the predictions made on the dataset exhibit strong performance, with a similar trend in both cases. The performance variations among different projections will be explored in the next section.

The success of the final training method (Figure 14) can be attributed to the usage of the last two projections ($i = 60^o, 90^o$). Both loss functions exhibit a desirable progression, although it is worth noting that the validation set loss appears somewhat more erratic and fluctuates at certain epochs. This behaviour may arise from the previously mentioned challenges related to the utilisation of low galaxy inclination images during training.

When evaluating the trained model, the CNN provides a test loss function value that serves as an indicator of the final accuracy of the model. The results are presented in Table 1. The single random projection method

yields slightly inferior results when compared to those achieved by training the network with images of higher inclination. Nevertheless, there remains a patent correlation between the actual and network-inferred intrinsic ellipticity q values. Consequently, the results indicate that CNNs have the potential to capture the shapes of galaxies when provided with a sufficiently large dataset.

5.2 Limitations

As mentioned previously, when the network is fed with face-on galaxy images, it struggles to capture the information effectively. Conversely, as the inclination of the galaxy images is augmented to be more similar to edge-on galaxies, the quality of predictions improves. This limitation shares similarities with the image degeneracy encountered in problems involving real galaxy shapes. Upon examining the model definition, Equation 4 reveals that the Gaussian profile on the z-axis captures the information of q within the flux. Therefore, at lower inclinations, the image's underlying parameters are visually indistinguishable. This implies that visual task-oriented neural networks still struggle to overcome the challenges given by image degeneracy.

Furthermore, it is important to acknowledge a significant limitation regarding the galaxy model's construction. The galaxies are built based on a pre-defined sky coordinate grid, which implies that the axes are fixed and do not dynamically adapt to the characteristics of each individual galaxy. As a result, when working with galaxies oriented at high inclinations, the obtained images may exhibit a cropped appearance (Figures 11, 10).

5.3 Future Work

Although acceptable results have been achieved, it is necessary to consider the inherent limitations associated with the chosen method. One possible way to enhance the network's performance at high degeneracy levels is by expanding the input depth, which would offer the model more features to learn from and potentially improve its outcomes. Also, a crucial aspect in avoiding the introduction of any potential bias is the design of a dynamic grid within the galaxy model, ensuring the entire galaxy can be represented.

Understanding the model boundaries is essential for further refinement and advancement in this research line. Further enhancements to the neural network's performance can be achieved through the following approaches:

- Increasing galaxy model complexity: By incorporating separate components such as bulge, disc, and halo into the galaxy model, the network can capture more features and better simulate the complexity observed in real galaxies.
- Dataset augmentation: Augmenting the dataset, in line with what was mentioned above, can significantly improve the network's performance. It enhances data diversity, mitigates overfitting, and facilitates the learning of more generalizable features.

Training method	4 proj.	1 proj.				1 random proj.
		$i = 0^o$	$i = 30^o$	$i = 60^o$	$i = 90^o$	
Test loss function	-2.375	0.680	-1.458	-2.042	-2.040	-1.657

Table 1: Loss function values for testing dataset.

- Incorporating observational and instrumental noise: Introducing factors like the Point Spread Function (PSF) into the training data enables the network to learn how to handle and compensate for such noise in real-world scenarios.
- Real Image Application: This is the ultimate aim of this work. Evaluating the network's performance on real images is crucial to assessing its ability to handle and solve real-life problems.

Conclusions

Se concluye que el conocimiento de las estructuras de las galaxias puede brindarnos información valiosa sobre su formación y evolución. En este sentido, nuevas herramientas como la Inteligencia Artifical pueden resultar particularmente interesantes.

Mediante la generación de 40.000 imágenes de galaxias, hemos logrado entrenar una Red Neuronal Convolucional (CNN) para obtener sus estructuras intrínsecas, las cuales están determinadas por el valor de q. Aunque la red ha enfrentado dificultades al tratar con bajas inclinaciones (vista de cara) debido a la degeneración de las imágenes, es aún así capaz de realizar inferencias precisas si dispone de un conjunto de datos lo suficientemente amplio. Este estudio sugiere que las CNN son capaces de recuperar de manera efectiva la verdadera forma tridimensional de las galaxias.

Understanding galaxy structures is crucial in astronomy as it provides insights into their formation and evolution. There is a relationship between galaxy structure and various aspects of galaxy evolution, including star formation, galaxy mergers, and properties such as luminosity and radius.

Throughout time, obtaining the intrinsic shapes of galaxies has represented a significant challenge. Traditional approaches reduce the problem to a few parameters, which potentially results in a loss of information. In this work, adopting a Deep Learning method has allowed us to consider and utilise every feature present in the data.

Through the generation of a dataset comprising 40,000 images, we successfully trained a Convolutional Neural Network to obtain the intrinsic galaxy shape, determined by the value of q. This highlights the potential of using simulations, in addition to telescope data, to enhance AI techniques in this field.

The employed techniques reveal that the CNN encounters challenges in accurately

estimating the flattening values (q) at lower inclinations (face-on orientation). This indicates the persistence of image degeneracy issues.

Nevertheless, when subjected to a single random projection simulating a real-life scenario, the CNN demonstrates its ability to overcome degeneracy problems when provided with an extensive dataset. Despite the need for further improvements in the methodology, it has resulted in accurate predictions. This study suggests that Convolutional Neural Networks are capable of effectively recovering the true 3-dimensional shapes of galaxies.

References

Binney J., 1985, MNRAS, 212, 767

Binney J. J., Merrifield S. D., 1998, Galactic Astronomy. Princeton Univ. Press, Princeton, NJ

Caon N., Capaccioli M., D'Onofrio M., 1993, MNRAS, 265, 1013

Ciotti L., Bertin G., 1999, A&A, 352, 447

Conselice C.J., 2003, ApJS, 147, 1

de Vaucouleurs G. 1948, AnAp, 11, 247

Domínguez Sánchez H., et al., 2018, MNRAS, 476.3, 3661-3676

Elmegreen D., 1998, Galaxies and Galactic Structure. Prentice Hall

Erwin, P., 2015, ApJ, 799.2, 226

Graham A. W., Driver S. P., 2005, Publications of the Astronomical Society of Australia, 22.2, 118-127

Hubble E., 1926, ApJ, 64, 321

Hubble E., 1930, ApJ, 71, 231

Huertas-Company M., 2015, Proceedings of the International Astronomical Union, 11, (S319), 118

Huertas-Company M., Pérez-González P. G., Mei S., et al., 2015, ApJ, 809, 95

Huertas-Company M., et al., 2019, MNRAS, 489. 2, 1859–1879

Lambas D., Maddox S., Loveday J., 1992, MNRAS, 258, 404

MacArthur, L. A., Courteau, S., Holtzman J. A., 2003, ApJ, 582, 689

O'Shea K., Nash R., 2015, An introduction to convolutional neural networks

Price S. H., et al., 2021 ApJ, 922, 143

Ramachandran P., Zoph B., Le Q.V., 2017, Searching for activation functions

Sandage A., Freeman K., Stokes N. R. 1970, ApJ, 160, 831

Sérsic J.L., 1963, BAAA, 6, 41

Sérsic J.L., 1968, Atlas de galaxias australes. Observatorio Astronomico, Cordoba

Toomre A., Toomre J., 1972, ApJ, 178, 623

Trujillo I., et al., 2006, ApJ, 650, 18

Weijmans A.M., et al., 2014, MNRAS, 444, 3340