Inferring Galaxy Morphology Through Texture Analysis

Joseph Richards
Andrew Connolly

May 4, 2007

1 Introduction

Galaxies are enormous, gravitationally-bound collections of stars, inter-stellar gas and dust, and dark matter. They generally contain between 10^7 and 10^{12} stars. Galaxies are constantly changing through gravitational interactions with the gas, dust, and other galaxies in their local environment. These encounters cause galaxies to change significantly throughout their lifetimes. The result is a diverse collection of galaxy morphologies throughout the observable universe. Studying galactic morphologies can give astronomers insight into the subtleties of galactic interaction dynamics and can provide details of the long-term trends in galactic evolution.

Galaxies are generally classified into three main groups based on morphology: ellipticals (E's), spirals (S's), and irregulars (Irr's) (Figure 1). Spiral galaxies are further split into regular spirals (S's) and barred spirals (SB's). The main classes of galaxies are in turn divided into sub-groups based on either the ellipticity of E-type galaxies or the tightness of the spiral arms (curvature) of S-type galaxies. This galactic classification scheme was first proposed by Edwin Hubble in 1936 [9].

Hubble believed both that galaxies evolved morphologically in an isolated environment and that the overall trend in galactic evolution was from ellipticals to spirals. Modern theory, however, claims (with a great deal of supporting evidence) that galaxies evolve through gravitational interactions with their surrounding environment. Additionally, astronomers now believe that ellipitcal galaxies are the by-products of interactions between spiral galaxies, which are believed to dominate the early universe.

In an ongoing effort to quickly and precisely characterize the morphological types of potentially millions of galaxies, a collaboration led by astronomer Andrew Connolly (Google, formerly at University of Pittsburgh) has studied methods of texture analysis to extract important morphological features of galaxies. In doing so, the group wishes to answer two main research questions. First, how has the distribution of galactic morphologies in the universe changed with time? In other words, what is the general trend of galactic evolution in the universe? In Astronomy, the further away an object is from us, the earlier in its history we perceive it. Therefore, a deep survey of galaxies provides us an instantaneous snapshot of the entire timeline of galactic evolution in the universe. Secondly, how does the local density of galaxies affect the distribution of morphological types? An answer to this question in turn answers the more scientifically interesting question of how galactic collisions affect galaxy morphology, since local density is a strong indicator of galactic collision rate.

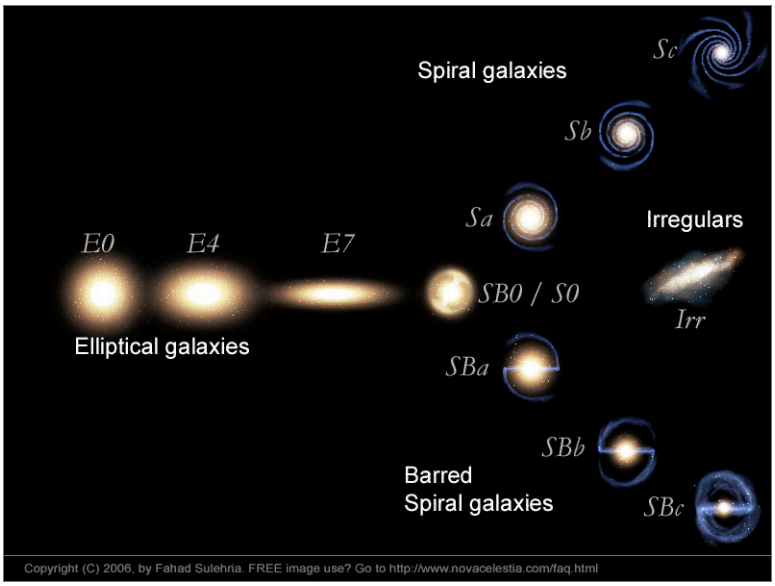

Figure 1: Galactic classification scheme devised by Hubble. The three main classes of galaxy morphology are ellipticals (E's), spirals (S's), and irregulars (Irr's). Spiral galaxies are divided into regular (S's) and barred (SB's) subtypes, and are further classified by the tightness of the spiral arms.

2 The Data

Modern sky surveys have the ability to gather data on millions of galaxies. This sheer amount of data, in the form of mega-pixel digital images, makes it impossible for astronomers to manually study these objects. The classical technique of studying galactic morphologies through visual image inspection is now infeasible. Therefore, we must devise methods that can quantitatively extract important morphological features from galactic images.

One survey in particular, the Sloan Digital Sky Survey (SDSS), promises to obtain images of millions of galaxies [10]. SDSS, when completed, will optically image over a quarter of the entire sky. In June 2005, SDSS completed its first phase, SDSS-I, a five-year campaign that returned data on more than 675,000 galaxies. Now SDSS is in its second phase, SDSS-II, which will last through June 2008.

From SDSS, we have obtained a sample of hundreds of galactic images with which to test quantitative texture analysis methods. Particularly, from SDSS-I we have a sample of over 1400 galaxies in G, R, and I bands plus the corresponding composite images produced by combining the images from these three filters [8]. These galaxies are all from the RC3 catalog [5], and each has diameter between 3 and 30 arcmin (as listed in RC3). To test our methods, we hand-select 105 spiral galaxies with prominent spiral arms from the SDSS-I sample.

3 Methods

Our particular role in this project is to determine the ability of a particular method, the digital curvelet transform, to extract prominent curvilinear features in digital images of spiral galaxies. The curvelet decomposition was developed by a team led by Candes, Donoho, and Starck [6, 2, 3], and has previously been shown to effectively enhance and restore astronomical images with elongated features [11].

In this paper, we aim to show that the curvelet transform can be used to effectively amplify the anisotropic (axis-dependent) features of spiral galaxies while ignoring point-like features such as foreground stars. To test the effectiveness of these methods, we attempt to fit a logarithmic spiral arm model to both the original SDSS and curvelet transform reconstructed images.

3.1 Curvelets

Curvelets were developed [11, 3, 2] to address the problem of sparse representation of objects with lines and edges. Conceptually, the 2-D curvelet transform considers features at several different directions and positions at each length scale and needle-shaped elements at fine scales.

A distinguishing characteristic of the curvelet transform is its parabolic scaling relation that at each scale 2^{-j}, every individual element is contained in a ridge of length $2^{-j/2}$ and width 2^{-j}. This endows the transform with the unique property of being able to easily find lines and edges, which are generally contained in such ridge-like boundaries

We define the continuous-time curvelet transform in $\mathbf{R}^2$ for spatial variable $\mathbf{x} = (x, y)$ at scale 2^{-j}, orientation θ_l, and position $\mathbf{x}_k^{(j,l)} = R_{\theta_l}^{-1}(k_1 \cdot 2^{-j}, k_2 \cdot 2^{-j/2})$ by

$$\varphi_{j,l,k}(\mathbf{x}) = \varphi_j \left(R_{\theta_l}(\mathbf{x} - \mathbf{x}_k^{(j,l)}) \right) \tag{1}$$

where φ_j is a waveform in the domain at scale 2^{-j}, R_θ is the rotation by θ radians, and $k = (k_1, k_2) \in \mathbf{Z}^2$ is a translation parameter. The waveform φ_j is defined by means of its Fourier transform with a basis of cosine functions.

With this construction, we then define a curvelet coefficient of an image $f(x, y)$ at scale 2^{-j}, orientation θ_l, and translation k as the inner product between f and the curvelet $\varphi_{j,l,k}$ defined in Equation (1) as

$$c(j, l, k) = \langle f, \varphi_{j,l,k} \rangle = \int_{\mathbf{R}^2} f(\mathbf{x}) \overline{\varphi_{j,l,k}(\mathbf{x})} d\mathbf{x} \tag{2}$$

Heuristically, the coefficient will be large if there is a prominent linear feature in the image f at the specified scale, orientation, and location. Because of the parabolic scaling relationship, the coefficients will tend to be large only when there are linear features present.

In practice, because we are using pixelated images, we employ a discretized version of the continuous curvelet transform, the digital curvelet transform. Details of two separate implementations of the digital curvelet transform are found in [2].

For any given image, we calculate the curvelet coefficients at every pixel on 6 different scales and for 16 different orientations. After computing these coefficients, we reconstruct

the image f from the m proportion largest coefficients: that is, we set all coefficients to 0 except for the m proportion of largest values, and then invert the original curvelet transform. For an appropriate m, the reconstruction, $\hat{f}_m$ retains the prominent anisotropic features of the original image while thresholding the isotropic features such as points and circles. The output from these routines is a two-dimensional intensity field. We have as our tuning parameter m, the proportion of largest coefficients to retain before inverting the transform.

The approximation error of the curvelet reconstructed image, $\hat{f}_m$, to the original image, f achieves

$$||f - \hat{f}_m||^2_{L^2} \leq C \cdot (\log m * T)^3 \cdot (m * T)^{-2} \tag{3}$$

for a curve that is C^2 except for discontinuities along C^2 curves, where T is the total number of curvelet coefficients calculated. Equation (3) supplies us with an analytic upper bound for the disparity of the reconstructed image to our original image.

3.2 Implementation of curvelet transform

Matlab implementation of the digital curvelet transform is found in the CurveLab package at http://www.curvelet.org. The CurveLab package contains the digital curvelet and inverse digital curvelet transforms, and several scripts to demonstrate their use.

For different choices of the tuning parameter, m, the curvelet transform produces images that vary substantially (Figure 2). Values of m that are too large leave the majority of the image intact, including the isotropic features that are meant to be cut out. Conversely, values of m that are too small begin to cut out many of the important curvilinear features of the image, losing the crucial texture features of the image that we wish to retain.

An important problem that remains in our analysis is the selection of the optimal tuning parameter, m, for each image. Here, we loosely define as optimal the reproduced image that preserves most of the prominent curvilinear features of a given input image while thresholding out most of the point features. Because this problem is somewhat ill-defined, it is difficult to conceive of numerical criteria that will allow us to automatically choose an optimal m. In practice, we find that a value of $m = 0.001$ tends to produce reconstructed images with anisotropic features intact and most point features removed. In our analysis, we use a value of $m = 0.001$.

4 Galactic Model

To analyze how well our methods perform in extracting the spiral features of galaxies while ignoring foreground stars, we now try fitting a parametric model to the curvelet reproduced images. This model will supply us with a quantitative measure of the curviness of each galaxy, which we can compare to different classification schemes in the literature. In this section, we detail the specification of our 8-parameter logarithmic spiral arm model and the fitting of the model to our curvelet transform reconstructed images. The model is often described as the Grand Design spiral model in Astronomy literature. Our parameterization is the same is that in [1].

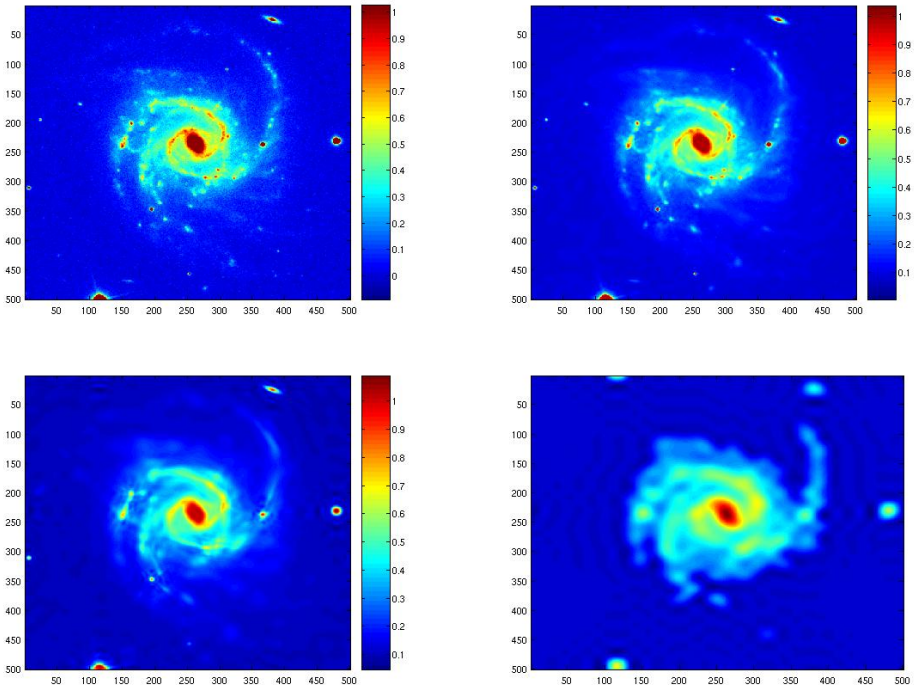

Figure 2: Original image of galaxy NGC 309 plus three images produced by curvelet coefficient thresholding, in order of increasing threshold level (0.01, 0.001, and 0.0001 proportion of curvelet coefficients are used in the image reconstruction, respectively). Observe that isotropic features, such as foreground stars, are filtered by this technique. As more curvelet coefficients are cut, less fine details are observed, and the spiral arms become the prominent features. When too many coefficients are thresholded, the main structures of the galaxy are partially lost.

4.1 Logarithmic Spiral Model

We fit the parametric logarithmic spiral arm model to each galaxy in order to obtain a numerical measure of the curveness of the spiral galaxy. The logarithmic spiral arm model is a four-parameter model. In polar coordinates, the model is

$$\theta = t + k \tag{4}$$

$$r = \beta e^{\alpha t},\ 0 \leq t \leq T, \alpha \geq 0, \tag{5}$$

where $k \in [0, 2\pi)$ is the starting angle of the spiral, $\beta > 0$ is the radial length, $T > 0$ controls the stopping point of the spiral arm, and α controls the curvature of the spiral (a smaller α means higher curvature) [1]. It is the parameter α in which we are most interested.

For each of the 105 galaxies in our reduced dataset we have found both the Hubble and the de Vaucouleurs classification (a more finely-divided modification of the Hubble

scheme, with spiralness from 1-9 [4]). Both of these schemes are based on how tightly the spiral arms are curved. After model-fitting we can compare our fitted values of α from Equations (4) and (5) for each galaxy to each of the two spiral classifications. This allows us a method of determining whether our model-fitting produces results that are consistent with the qualitative-based classifications of both the Hubble and de Vaucouleurs systems.

To give our model more flexibility in finding the optimal fit, we include parameters for the center of the galaxy, (x_c, y_c). This permits the model to search for the center of the galaxy and frees us from having to manually find and fix a center value. Also, since most spiral galaxies have two prominent, nearly-symmetric arms, we model two symmetric logarithmic spiral arms by flipping the first arm over both the x and y axes defined by the center of the galaxy. Our model, in cartesian coordinates, is

$$x = \pm\beta e^{\alpha t}\cos(t+k) + x_c \quad (6)$$

$$y = \pm\beta e^{\alpha t}\sin(t+k) + y_c,\ 0 \le t \le T \quad (7)$$

where the parameters are defined as in (4), (5). An example of this model can be seen in Figure 3.

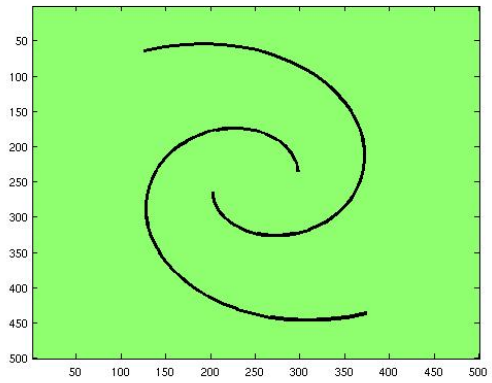

Figure 3: Example of our galactic arm model. The model comprises of two symmetric logarithmic spirals.

4.2 Inclination Angle

The light matter (stars, gas, and dust) of spiral galaxies reside in a plane in space. If this plane is tilted to our line of sight, then we will percieve a different spiral formation than if the plane is perpendicular to us. In the extreme case, a galaxy could be tilted 90 degrees to our line of sight, whereupon we would only detect a thin linear structure and would not be able to discern any spiral features.

In fact, the only case in which our original logarithmic spiral model is valid is if the galaxy of interest is inclined 0 degrees to our sight line (i.e. the plane of the galaxy is perpendicular to us). However, this is a statistically impossible realization. Every galaxy

will have some positive inclination to our line of sight, and therefore we must add inclination to our model.

To model inclination, we need two more parameters: inclination angle, ω, and inclination axis. For inclination axis, we specify the slope, μ of the line that defines the pivot around which the galaxy is rotated. In the model, galaxy inclination results in shrinking the perpendicular distance between each point on the spiral and the inclination axis by a factor of $\cos(\omega)$. Thus, if we call z the perpendicular distance between any point on the spiral and the inclination axis, then we replace z by

$$z' = z\cos(\omega) \tag{8}$$

to obtain our inclined model [1].

Working out the mathematical details, inclination reduces to replacing every point (x, y) on the spiral model by (x_{new}, y_{new}), where the latter is defined by

$$\begin{aligned} x' &= \frac{\mu(y - y_c)}{\mu^2 - 1} + \frac{x_c\mu^2 + x}{\mu^2 + 1} && (9)\\ y' &= y_c - \mu x_c + \mu x' && (10)\\ x_{new} &= \cos(\omega)(x - x') + x' && (11)\\ y_{new} &= \cos(\omega)(y - y') + y' && (12) \end{aligned}$$

Inclination is demonstrated in Figure 4.

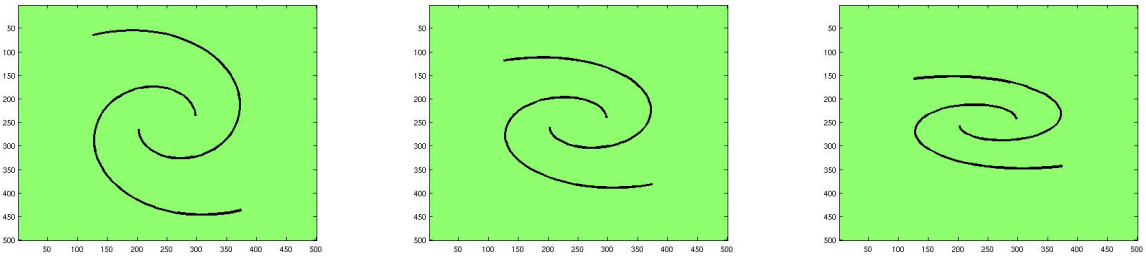

Figure 4: A galactic arm spiral model, and the same model inclined about the $y = 0$ axis by an angle of $\pi/4$ and $\pi/3$.

4.3 Model Fitting

Our complete model has 8 parameters: 4 for the logarithmic spiral, 2 specifying the center of the galaxy, and 2 for the inclination of the galaxy. We now discuss how to fit this model to an image.

4.3.1 Intensity Model

The model we discussed in the previous sections describes an indicator function of whether or not each individual pixel lies exactly on the curve. However, because pixel locations

take on values on $\mathbf{Z}^2$ while our model takes on values in $\mathbf{R}^2$, no pixel will ever fall exactly on the curve. Also, galactic arms are not constrained to narrow bands, but rather are generally dozens of pixels wide in our images. Clearly, we must devise a better model to more accurately mimic the galactic images.

First, instead of considering only pixels that fall exactly on the curve, we define an indicator function which is 1 if a pixel is located within 30 pixels of the curve and zero elsewhere (the choice of 30 is somewhat arbitrary, but was shown to perform well with several images). This gives us a band of intensity one around the spiral arms. In MatLab, this is easily performed by replacing each point on the curve by a circle of radius 30.

However, the galaxy images have higher intensity exactly on the spiral arms and decreasing intensity moving away from the arms. This property is not captured by our current indicator function. To rectify this, we convolve the current function with a two-dimensional Gaussian kernel. Convolving with a wide kernel with high standard deviation, pixels in the center of the indicator function remain near a value of 1, while those toward the outside of the indicator function receive lower values due to the influence of the zero-valued pixels outside the radius of 30.

Finally, we scale down the entire intensity by a factor of 0.4. We choose 0.4 because we would like to pick up spiral arm features down to about an intensity value of 0.2. Anything weaker than 0.2 tends to blend in with the background noise. This model will tend to capture any features with values greater than 0.2 because the difference between our intensity model and a pixel with intensity greater than 0.2 will be smaller than the value of that pixel if it is not captured by the model.

An image representation of our final intensity model can be seen in Figure 5.

4.3.2 Fitting the Model

To actually fit the 8-parameter model, we minimize the total squared difference between the intensity image and our intensity model. This amounts to minimizing the square of the differences summed across pixels. We perform this 8-parameter minimization with an unconstrained nonlinear optimization routine (fminsearch) in MatLab. This minimization routine uses a simplex search method.

Since our function is not necessarily well behaved (i.e. since the images are quite complex, there may be several local minima) it is important to 1) approximate good starting values for as many of the 8 parameters as possible, and 2) run the minimization algorithm several times with different starting values to determine the best fit (i.e. the minimum value of our sum of squared differences function).

Recall that our model has 8 parameters. It turns out that we can quickly and easily obtain good estimates for 6 of those parameters to use as starting values for our minimization routine. First, we can easily visually inspect an image to determine the starting locations of each of the two prominent spiral arms. Since these spiral arms are generally symmetric around the center of the galaxy (recall that symmetry is one of the main assumptions in our model), specifying these locations automatically gives us good estimates for x_c, y_c, β, and k. If we call $s_1 = (s_{1x}, s_{1y})$ and $s_2 = (s_{2x}, s_{2y})$ the starting locations of the two galaxy arms, then:

$$x_{c,0} \quad = \quad (s_{1x} + s_{2x})/2 \tag{13}$$

$$y_{c,0} = (s_{1y} + s_{2y})/2 \quad (14)$$

$$\beta_0 = ||s_1 - s_2||_{L_2}/2 \quad (15)$$

$$k_0 = \pi - \arctan\left(\frac{s_{2y} - s_{1y}}{s_{2x} - s_{1x}}\right) \quad (16)$$

Note that if our estimates are not exact, there is no problem since the parameters are all part of the parameter space that is searched by the minimization algorithm. Supplying close starting values only helps to guide the algorithm to a more accurate fit.

We also have starting values for both the inclination angle, ω, and inclination axis μ from the RC3 catalog [5]. The RC3 catalog lists the log isophotal axis ratio, R_{25}, at the contour corresponding to 25 mag arcsec^{-1} in B, which can be converted to inclination angle by the equation

$$I = \cos^{-1}(10^{-R_{25}}) \quad (17)$$

as specified in [7]. RC3 also supplies an inclination axis for galaxies with significantly non-zero inclination angles (for other cases we set μ_0=0).

The two starting values that are difficult to estimate are α, the curviness, and T, the stopping angle/position of the spiral. The only way to reasonably attempt to estimate these quantities would be to plot spirals with different parameter values. Hence, we choose to randomly generate starting values for these parameters according to:

$$\alpha_0 \sim U[0.1, 0.8] \quad (18)$$

$$T_0 \sim U[1, 8] \quad (19)$$

The bounds on α_0 and T_0 are selected as plausible parameter values for any given galaxy.

We fit the model with 20 different starting values and pick out that model which achieves a minimum sum of squared differences value. Sometimes it occurs that the minimum value is achieved for a spiral that winds tightly around the most intense part of the galaxy. In this case, we select as the optimal fit that which achieves the lowest function value under a parameter constraint on α.

5 Results

We fit the model to a sample of 20 spiral galaxies from the SPSS sample [8]. For 18 of the 20 galaxies, the fits are quite good. The other 2 galaxies, NGC 2595 and NGC 3368 do not achieve plausible fits (judged qualitatively) for any of the 20 random starting values. For NGC 2595, this occurs because the galaxy has a prominent bar, which we have not modeled. Note that we are able to fit the model to other barred galaxies, but in this case the bar is simply too prominent. For galaxy NGC 3368, the spiral arms are not very prominent (the galaxy is classified as Sa/Sb & as de Vaucouleurs spiralness of 2, meaning the arms are very tightly wound), and the fitting algorithm is unable to detect them.

Fits for the 18 galaxies are overplotted on the original image of that galaxy in Figures 6 and 7. Remember that these have all been fit using the curvelet reproduced image of the corresponding galaxy. The model fits well for a wide variety of galaxies: from face-on to inclined and from highly-wound to very straight-armed. Deviations of the model from the actual arms can generally be attributed to deviations of the galaxy arms from actually

following logarithmic spirals. Overall, the fits very closely follow the spiral structure of the galaxies.

To compare the effectiveness of the fitting with the original images (i.e. the images withough curvelet thresholding), we also fit the model to each of the 18 original images. In general, we find that the spirals are easily distracted by bright foreground stars, and result in poorer fits to the galactic arms. We plot the resultant best fits to the original image for both NGC 3184 and NGC 3338 (Figure 8). Obviously, in both of these cases the presence of a bright foreground star has caused the fitting algorithm to fail to appropriately model the galaxy. Compare these fits to those obtained by fitting these two galaxies to the curvelet-reproduced images (Figure 6).

A way to test how well our fitting procedure has performed is to see how closely our parameter for spiralness of each galaxy, α, corellates with the de Vaucouleurs classification of spiralness. As galaxies are less tightly wound, both α and the de Vaucouleurs classification (1-9) should increase. A plot of the de Vaucouleurs specification versus our α parameter from the model fit (Figure 9) shows a moderately strong increasing trend (non-significant p-value in test of zero slope), which is what we would expect. Keep in mind that the de Vaucouleurs classification is performed by visual inspection of each galaxy and not by numeric methods and therefore is not necessarily correct.

One important result we obtain is the inclination angle for each of the 18 galaxies. The standard technique for determining inclination angle is to fit intensity contours to each galaxy and extract the (fairly faint) one that corresponds to a brightness of 25 mag arcsec^{-1}. The log isophotal ratio (log ratio of the major to minor axis of this contour) is computed, and then the inclination is calculated by Equation (17).

However, this technique is not very precise and can also be biased because it assumes that the galaxy has an elliptical halo (outer, less bright region of the galaxy). In the RC3 catalog, measurements for this log ratio are given to the nearest 0.01. Also, standard errors for each estimate are quoted. These standard errors are generally quite large, ranging from 0.02 to 0.07 (while most measurements are in the range of 0.01 to 0.20). For very small inclination angles, small errors in the measurement of the log isophotal ratio can propogate to huge differences in inclination angle. For example, a change in log isophotal ratio from 0.00 to 0.02 changes the derived inclination angle from 0 to 30 degrees! Remember that 90 degrees is the maximum inclination angle possible.

Our techniques circumvent these imprecise methods of inclination angle determination because we simultaneously fit all parameters of our model to capture the spiral arms of the galaxy. We believe that our derived inclination angle values from model fitting are more accurate estimations of the true inclination angles than the methods described above, especially for low inclination galaxies. In Figure 10, we plot the RC3 derived angles versus our galactic fit inclination angles. The line represents the diagonal ($x = y$) line. The error bars on the RC3 values represent one standard error of measurement as quoted in the catalog. We see that our values are generally quite consistent with the RC3 values, with some substantial differences for low-inclination galaxies.

6 Discussion

We have demonstrated a technique that uses a digital curvelet transform to preserve the linear and curvilinear features of spiral galaxies in megapixel images while thresholding point-like features such as foreground stars. These techniques have previously been used for image enhancement and restoration in Astronomy [11]. We threshold an extremely large proportion of coefficients (> 0.99), preserving only the main spiral structure of each galaxy.

To test the effectiveness of these methods, we fit an 8-parameter logarithmic spiral model to each of 20 curvelet-reproduced images of galaxies. Also, we attempt model fits to the original SDSS-I images and show that they are easily distracted and pulled off course by bright foreground stars.

From these models, we obtain two scientifically-interesting parameters: α, the curviness of the galactic arms, and ω, the inclination angle of the galaxy. Comparing α to the de Vaucouleurs classification gives us some sense of how accurately the fits are describing each galaxy, and we see an increasing trend, as expected.

Inclination angle determination is generally a difficult problem in Astronomy. Most estimates are imprecise and many are biased. Our model fits give us an estimate for the inclination of each galaxy that are quite robust and only use the assumption of logarithmic spiral arms in their derivation.

One problem we have neglected to address so far is computing a measure of uncertainty for the logarithmic spiral parameters we have estimated. We are essentially fitting a model of the form data = model + error. When we fit the model we obtain a set of residuals from this fit. One way to obtain standard error estimates on each parameter is to create new images by bootstrapping the residuals from our best fit and adding them to the best-fit model. We could then refit the model to the bootstrap image to obtain new parameter estimates. By repeating this process many times, we can compute empirical standard errors for each parameter which should be good estimates for the standard errors of each parameter. This technique, however, will take a substantial amount of computing time, especially if we choose to use hundreds of bootstrap samples.

With these techniques in place, we wish to analyze more of the 1400 SDSS-I images we have at our disposal. The algorithms we use are fast: the curvelet transform takes only a few seconds for a 500 x 500 image while running 20 model fits takes around two hours. Presently, we have another 85 images identified and ready to be curvelet thresholded and modeled.

References

[1] Kinman Au. *Inferring Galaxy Morphology Through Texture Analysis.* PhD thesis, Carnegie Mellon University, 2006.

[2] Emmanuel Candes, Laurent Denmanet, David Donoho, and Lexing Ying. Fast discrete curvelet transforms. *Multiscale Modeling and Simulation*, 5(3):861–899, 2006.

[3] Emmanuel Candes and David Donoho. New tight frames of curvelets and optimal representations of objects with c^2 singularities. *Comm. Pure Appl. Math.*, 57:219–266, 2002.

[4] G. de Vaucouleurs and A. de Vaucouleurs. Reference catalogue of bright galaxies, 1964.

[5] G. de Vaucouleurs, A. de Vaucouleurs, H.G. Corwin, R.J. Buta, G. Paturel, and P. Fouque. *Third Reference Catalogue of Bright Galaxies (RC3)*. Springer-Verlag, New York, 1991.

[6] David L. Donoho and Mark R. Duncan. *Digital Curvelet Transform: Strategy, Implementation and Experiments*, November 1999.

[7] Z. Frei, P. Guhathakurta, and J.E. Gunn. A catalog of digital images of 113 nearby galaxies. *Astronomical Journal*, 111:174, 01 1996.

[8] David W. Hogg and Michael R. Blanton. Sdss images of selected rc3 galaxies.

[9] E.P. Hubble. *The realm of the nebulae*. Yale University Press, 1936.

[10] Sloan digital sky survey (sdss).

[11] J.L. Starck, D.L. Donoho, and E.J. Candes. Astronomical image representation by the curvelet transform. *Astronomy & Astrophysics*, 2002.

7 Appendix

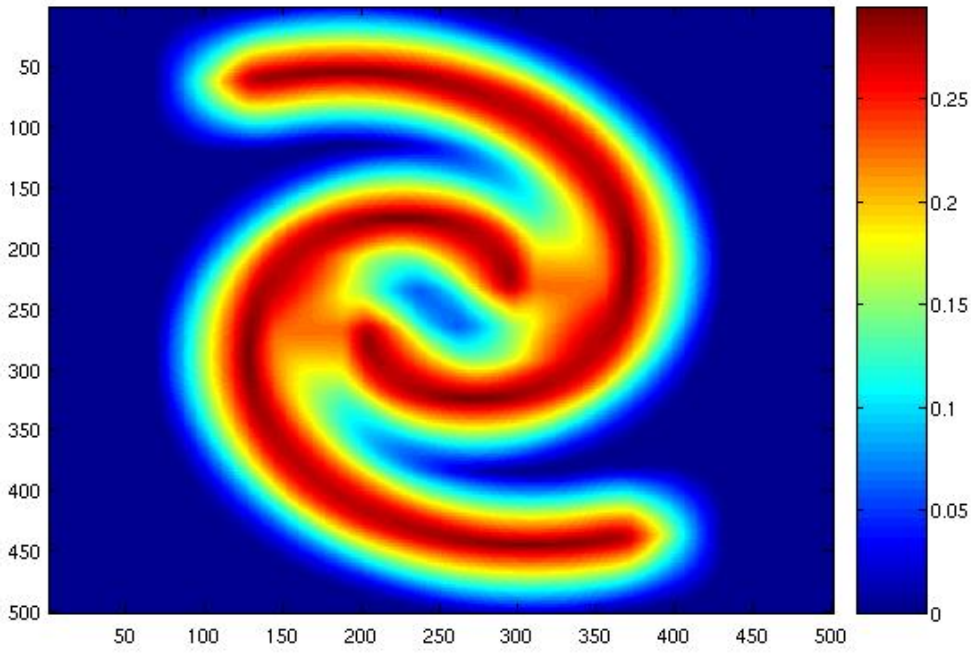

Figure 5: Our final intensity model. Note that pixels close to the logarithmic spirals have large values and that the pixel values decrease as you move away. We fit this model to each galaxy by minimizing the sum of squared differences between the pixel values of the model and the curvelet-reconstructed galactic image.

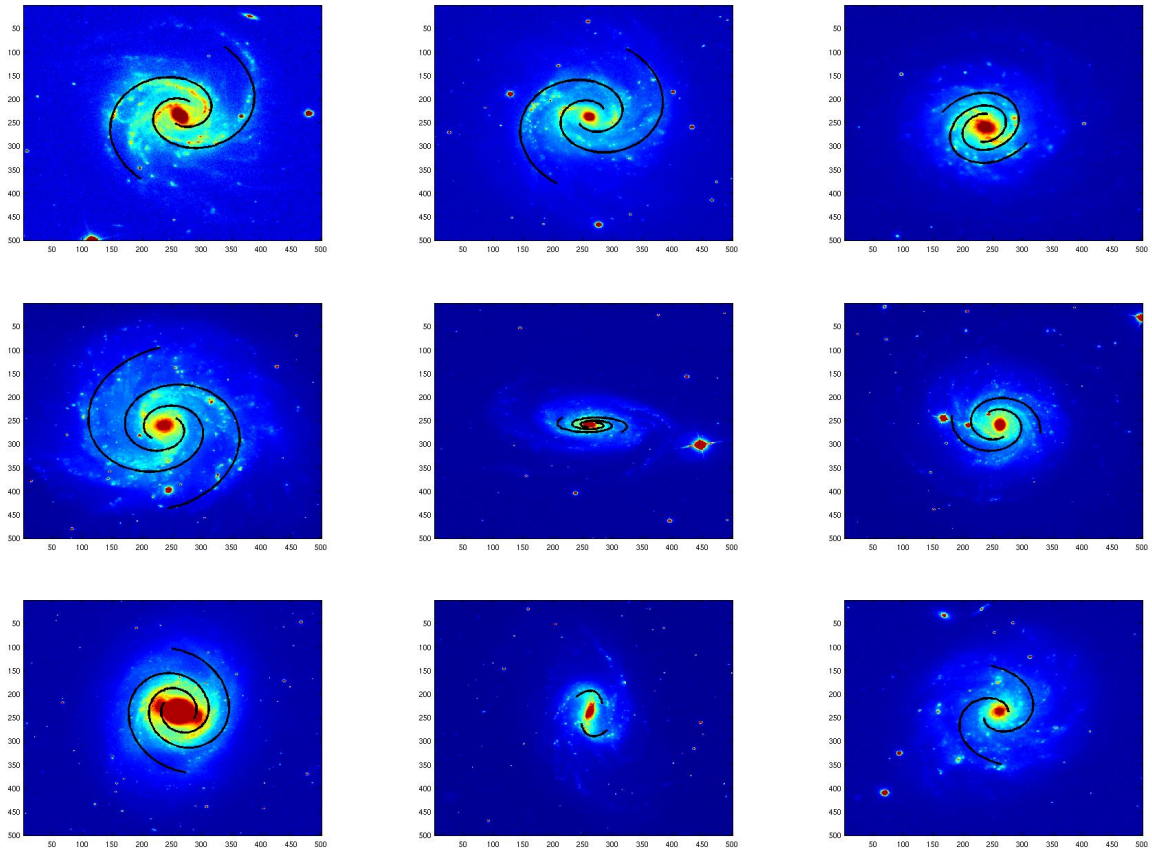

Figure 6: Fits for the first 9 of the 18 galaxies from SDSS-I. Each fit was achieved by minimizing sum of squared differences between the curvelet reproduced image and the intensity model. Galaxies are: NGC 309, NGC 1042, NGC 2967, NGC 3184, NGC 3338, NGC 3344, NGC 3351, NGC 3359, and NGC 3423.

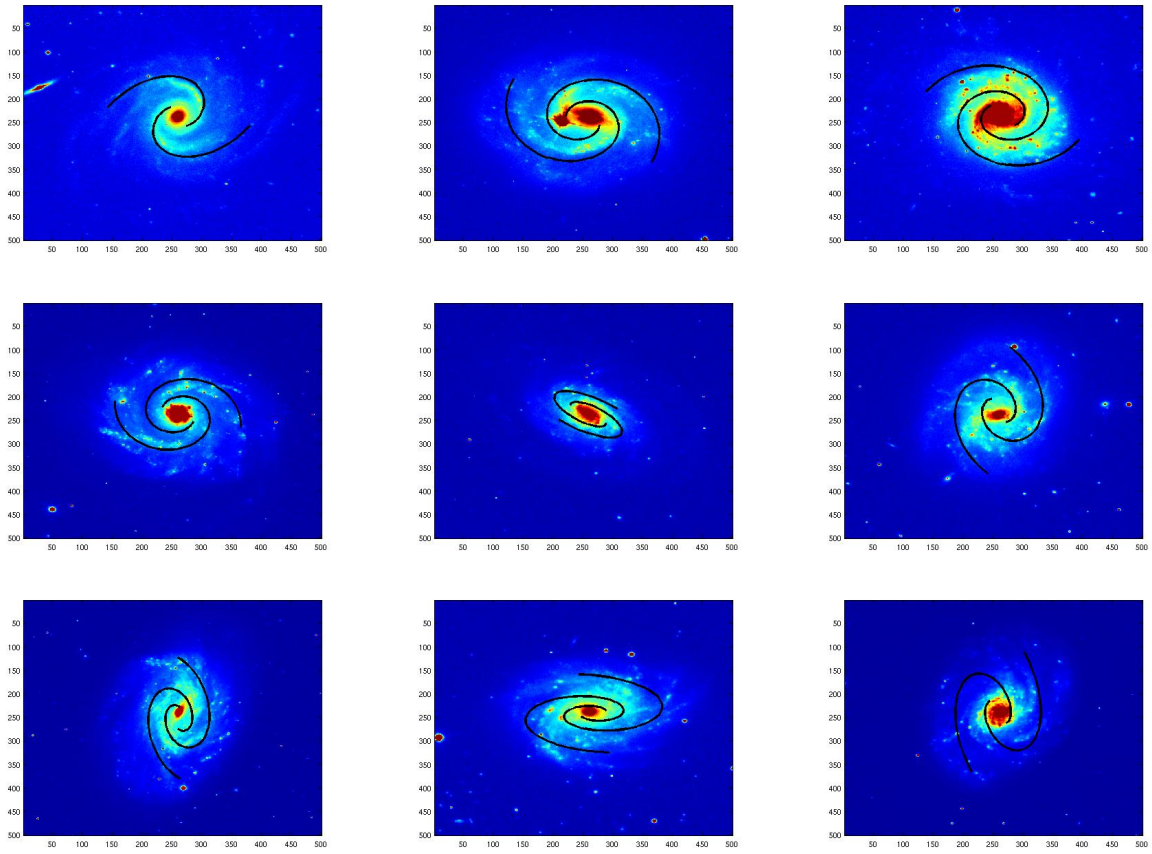

Figure 7: Fits for the second 9 of the 18 galaxies from SDSS-I. Each fit was achieved by minimizing sum of squared differences between the curvelet reproduced image and the intensity model. Galaxies are: NGC 3433, NGC 3507, NGC 3596, NGC 3631, NGC 3684, NGC 3686, NGC 3726, NGC 3780, and NGC 3810.

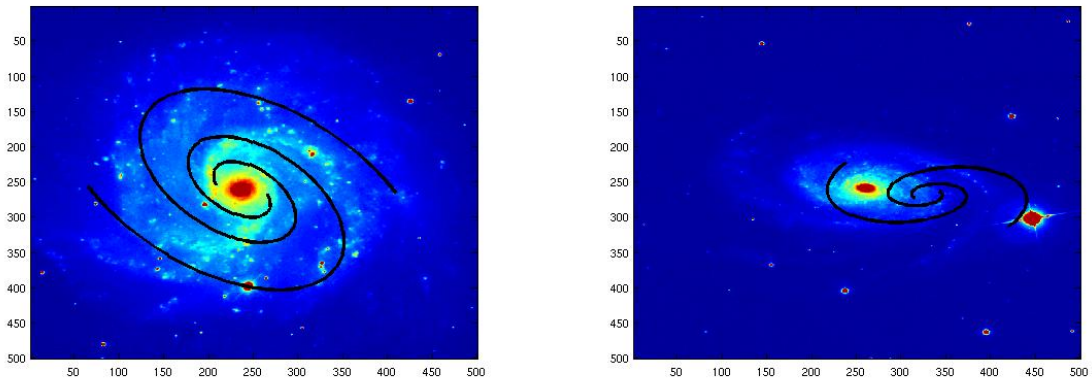

Figure 8: Fits to the original images of NGC 3184 and NGC 3338. Both of these fits are distracted by foreground stars, and result in bad fits.

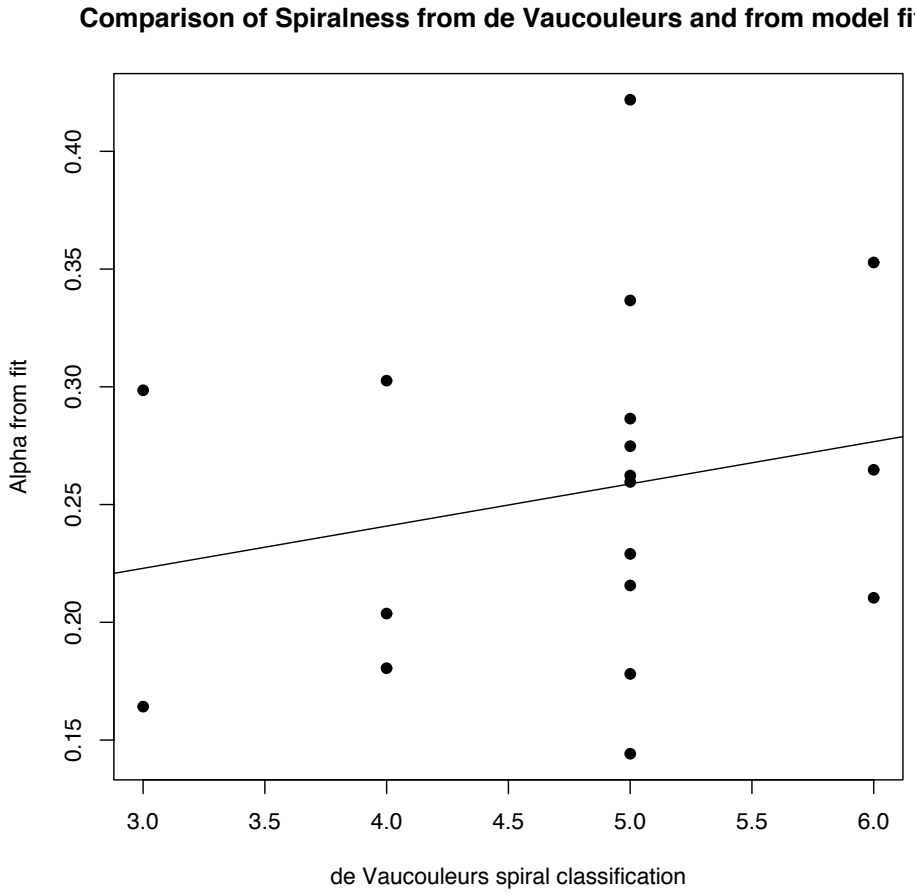

Figure 9: Plot of de Vaucouleurs classification versus α for our model fits. The least squares fit is plotted, and shows a mild increasing trend, as would be expected. The slope of the trend is not significantly non-zero at level .05.

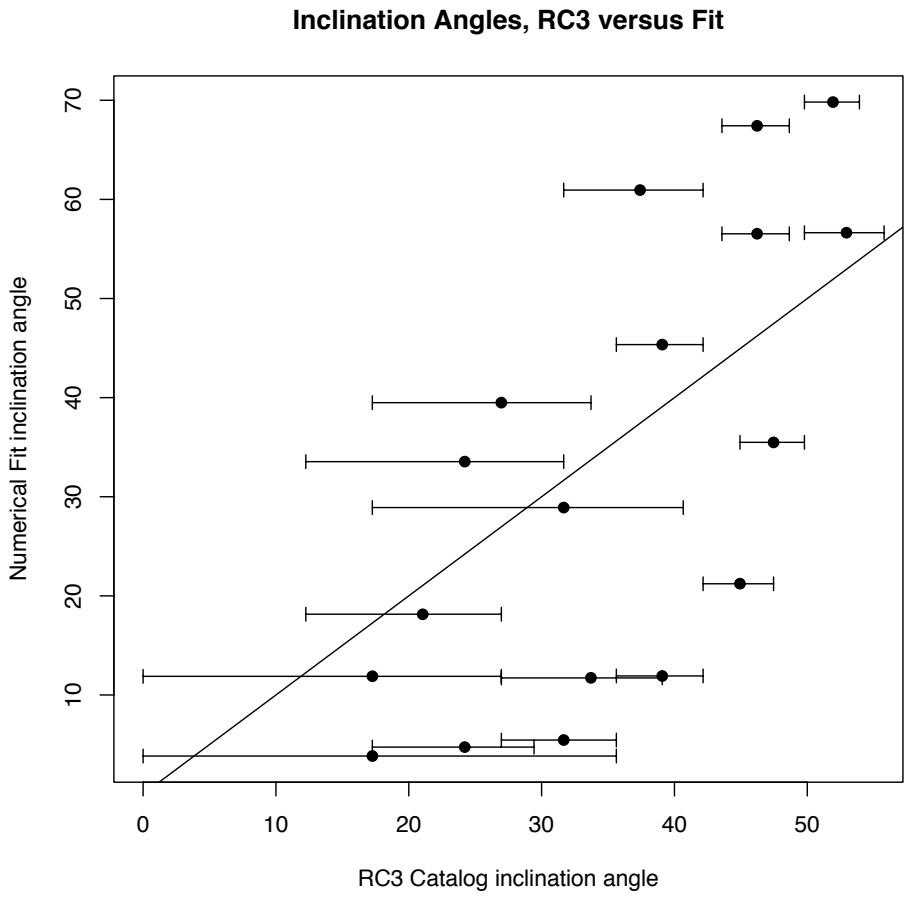

Figure 10: Plot of RC3 inclination angles derived from the log isophotal axis ratios versus the inclination angles determined by our fitting procedure. The line is the $x = y$ diagonal. Error bars represent one standard error of measurement of the RC3 inclination (as quoted in the RC3 catalog).