2D–1D Wavelet Reconstruction as a Tool for Source Finding in Spectroscopic Imaging Surveys

L. Flöer[A,C] *and B. Winkel*[B]

[A]Argelander-Institut für Astronomie, Auf dem Hügel 71, 53121 Bonn, Germany
[B]Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany
[C]Corresponding author. Email: lfloeer@astro.uni-bonn.de

Abstract: Today, image denoising by thresholding of wavelet coefficients is a commonly used tool for 2D image enhancement. Since the data product of spectroscopic imaging surveys has two spatial dimensions and one spectral dimension, the techniques for denoising have to be adapted to this change in dimensionality. In this paper we will review the basic method of denoising data by thresholding wavelet coefficients and implement a 2D–1D wavelet decomposition to obtain an efficient way of denoising spectroscopic data cubes. We conduct different simulations to evaluate the usefulness of the algorithm as part of a source finding pipeline.

1 Introduction

The usage of the wavelet transform in astrophysics has become very popular in recent years. Table 1 compiles the number of publications on ADS[1] in a given range of years that have the word 'wavelet' contained in their abstracts. Clearly, the usage of wavelets has gained popularity quickly. Typical applications for wavelet transform–based methods are morphological separation of sources in images and noise removal. The success of wavelet-based methods in astrophysics is in part due to the fact that astrophysical data often contain information on different angular or spectral scales. For example, an optical image of a galaxy contains compact, bright stars as well as extended and faint emissions from the bulge and spiral arms. Multi-scale methods, such as the wavelet transform, allow researchers to investigate the different scales of an image separately (Starck & Bobin 2010).

The most widely used type of wavelet transformation is the so called undecimated or redundant, isotropic wavelet transformation. This is in part due to the algorithmic simplicity of the method but also because undecimated wavelet transforms have proven to be more efficient for noise removal then their decimated counterparts. Apart from that, they also provide a number of computational advantages when reconstructing an image from a subset of its wavelet coefficients (Starck et al. 2010).

In this paper we review the basics of denoising based on the undecimated wavelet transformations in Section 2

[1]NASA Astrophysics Data System, `http://adswww.harvard.edu/`

Table 1. Number of abstracts on ADS containing the word 'wavelet' for given date ranges

Years	Number of abstracts
Until 1995	251
1996–2000	679
2001–2005	1221
Since 2006	1797

and present an extension of the wavelet transform to three-dimensional data as proposed in Starck et al. (2009). Section 3 describes the implementation of the transform in C++ along with a description of where the implementation departs from the original algorithm. A first application of the algorithm is shown in Section 4, where we use the algorithm to implement a source finder and test the performance on simulated HI galaxies. We close the paper with a summary and an outlook on future applications and potential improvements to the algorithm.

2 Wavelet Denoising

The isotropic, undecimated wavelet transform (IUWT) decomposes data $D(x)$ into $J+1$ subbands

$$D(x) = c_J(x) + \sum_{j=1}^{J} w_j(x) \qquad (1)$$

where c_J is a smooth version of the data and the details at position x and scale j are contained in the wavelet

coefficients $w_j(x)$. The IUWT can be efficiently calculated by using the so called 'algorithme à trous' (Holschneider et al. 1989). To calculate the IUWT one needs to convolve the input data with increasingly larger kernels. To calculate the next convolution, the algorithme à trous convolves the previously convolved data again with the same kernel with 2^j zeros inserted between the kernel values. For multidimensional transforms this insertion of zeros is done isotropically in all dimensions. This allows efficient calculation of even the largest scales.

At each step, two consecutive convolved versions of the data $c_j(x)$ and $c_{j+1}(x)$ are used to calculate the wavelet coefficients $w_j(x) = c_j(x) - c_{j+1}(x)$. The number of scales is usually chosen to be $\lfloor \log_2 N_i \rfloor$, where N_i is the number of samples per dimension in a data set, for example an image with $N_1 \times N_2$ pixels. A more detailed description can be found in Starck et al. (2010).

When using wavelets to denoise data, one assumes that the signal in the data, for example the sources, can be described by only a few relevant coefficients in each of the detail subbands $w_j(x)$; that is, that the signal is sparse in a given wavelet representation. Consequently, one can try to detect only the relevant coefficients and reconstruct the image from those.

The detection is usually based on estimating the standard deviation σ_j of the coefficients in a given subband and only take the coefficients with absolute values above a certain threshold $t\sigma_j$ to be significant. t is usually chosen to be between 3 and 5. Then, if one applies Equation 1, with all insignificant coefficients set to zero, one obtains a denoised approximation of the data.

Since this nonlinear denoising benefits greatly if iterated a few times, Murtagh et al. (1995) developed the notion of a multi-resolution support, M, that contains information about whether the data has a significant coefficient at a given location and scale. The multi-resolution support is defined as follows:

$$M(x, j)\begin{cases} 1 & \text{if } w_j(x) \text{ is significant} \\ 0 & \text{else.} \end{cases} \quad (2)$$

Using this multi-resolution support, one can implement the following iterative reconstruction scheme:

1. Detect all significant coefficients $w_j(x)$ and store this information in the multi-resolution support $M(x, j)$.
2. Calculate the IUWT of the data D and reconstruct the image only from the coefficients that belong to M to obtain $\tilde{D}$.
3. Calculate the residual $R = D - \tilde{D}$.
4. Calculate the IUWT of R and again only retain the coefficients that belong to M. Add this reconstruction to $\tilde{D}$.
5. Go to step 3 until the desired number of iterations is reached.

In practice a small number of iterations (<10) is sufficient. Many examples of how iteration improves the denoising process can be found in Starck et al. (2010).

2.1 Extension to 2D–1D Data

The aforementioned decomposition and reconstruction works very well if the relevant signal in the data is isotropic or nearly isotropic. This is true for most 1D and 2D astrophysical data like spectra and images. In the case of imaging spectroscopic surveys like the former Hɪ Parkes All-Sky Survey (HIPASS; Barnes et al. 2001; Koribalski et al. 2004), the ongoing Effelsberg-Bonn Hɪ Survey (EBHIS; Kerp et al. 2011; Winkel et al. 2010), the Arecibo Legacy Fast ALFA Survey (ALFALFA; Giovanelli et al. 2005) and future surveys with WSRT–Apertif (Oosterloo et al. 2009) and the Australian SKA Pathfinder (ASKAP; Johnston et al. 2008), the data is three-dimensional with two angular and one spectral dimension, which is referred to as a data cube. Since spatially unresolved sources can still be resolved spectrally, sources generally do not have the same size among the three different axes of the data cube. This leads to anisotropy, which makes isotropic denoising schemes inefficient, since the wavelet decomposition does not match the natural shape of the sources very well. Nonetheless, the sources can be considered partly isotropic in each individual spectral slice (channel map) and are also approximately isotropic along each line of sight. It is therefore beneficial to split the wavelet decomposition up into a two-dimensional angular and a one-dimensional spectral part.

The theoretical foundation for this 2D–1D transformation is laid out by Starck et al. (2009), who applied a 2D–1D denoising to data from the Fermi LAT (Atwood et al. 2009). Fermi LAT data has either two angular domains and one spectral or two angular domains and one temporal. Even though very different from radio astronomical observations in its noise characteristics, it is similar to data from imaging spectroscopic surveys in terms of the dimensionality.

To calculate a wavelet representation that accounts for this difference in axis type, they first calculate a 2D IUWT of each channel map and subsequently apply a 1D IUWT along each pixel of this wavelet-coefficient data cube. When applying this decomposition with J_1 angular and J_2 spectral scales one arrives at a decomposition of the form

$$\begin{aligned} D(x) = &\, c_{J_1, J_2}(x) \\ &+ \sum_{j_1} w_{j_1, J_2}(x) + \sum_{j_2} w_{J_1, j_2}(x) \\ &+ \sum_{j_1, j_2} w_{j_1, j_2}(x). \end{aligned} \quad (3)$$

Analogous to Equation 1, $c_{J_1, J_2}(x)$ is the smooth version at angular scale J_1 and spectral scale J_2. The second row contains the coefficients that arise from either spectral decomposition of the smooth angular scale or angular decomposition of the smooth spectral scale. The last sum of coefficients w_{j_1, j_2} contains the detail of the data at angular scale j_1 and spectral scale j_2.

To implement an iterative denoising scheme as described above, one again has to construct a five-dimensional (two angular dimensions, one spectral dimension, and two scale indices) multi-resolution support $M(x, j_1, j_2)$. Once all significant coefficients are detected, the iterative reconstruction can be applied as described in the previous section.

3 Implementation

3.1 Scale Selection

In general, the denoising of data is performed by decomposing the image with the maximum number of scales, i.e. $J_i = \lfloor \log_2 N_i \rfloor$. There are however certain advantages in using only a subset of decomposition scales for both the angular and spectral regimes. This is especially true when the denoised image only serves as a mask to find sources in the original data and complete flux reconstruction is not of importance.

The signature of a galaxy in neutral hydrogen surveys is typically small compared to the dimensions of the data cube. It is therefore unlikely that one will miss any sources when leaving out all decomposition scales that belong to larger spectral or angular scales. Especially for single-dish observations, the information contained at very large scales is most likely due to baseline errors or radio-frequency interference (RFI). By neglecting the larger scales, one can suppress those errors in the reconstruction and extract sources even from suboptimal data. This property is investigated in Section 4.3.

For this reason, the reconstruction in our algorithm is done with a 'physical' subset of angular and spectral scales that are likely to contain the signal of sources. Likewise, the coefficients w_{J_1, j_2}, w_{j_1, J_2} and the smooth data c_{J_1, J_2}, which all contain the information at the largest scales, are not taken into account, and are not part of the reconstructed data.

What scales are to be considered physical depends on the type of source one is looking for. In this paper we mainly focus on extragalactic objects, namely HI galaxies, which are typically not larger than a few arcminutes. What scales then contain the desired objects depends on their typical angular and spectral size and the respective sampling of the dimensions on the voxels of the data cube.

Additionally, we only reconstruct the data from the positive wavelet coefficients. This approach is different from the usual iterative approach, where all significant wavelet coefficients are used and the negative values of the reconstruction are set to zero at each iteration. We noticed that, by choosing only the positive coefficients and using mathematical morphology (see next section), we suppress the artifacts that arise during partial reconstruction from wavelet coefficients (Starck et al. 2007). If unsuppressed, these artifacts make the use of the reconstruction as a mask for source finding difficult, as they can also lead to merging of sources. On the other hand, this positivity constraint makes searching for negative features, such as absorption lines, impossible.

3.2 Mathematical Morphology

Another advantage of storing the information of the significant coefficients in the multi-resolution support is that one can perform mathematical morphology on it (Serra 1982). Generally, data cubes are created in such a way that the sampling of the telescope beam fulfills the Nyquist sampling theorem (Nyquist 1928), meaning that it is sampled on at least two pixels in every direction (including the diagonal). This means that real sources are larger than a single pixel in the angular dimension and are most likely also sampled in more than one spectral channel. Furthermore, it is well known that significant structures propagate through the different scales of the IUWT. Sources will therefore be present in multiple adjacent coefficients in the three dimensions of the data as well as adjacent spatial and spectral scales and form connected regions in the five-dimensional multi-resolution support.

To further suppress the noise in the reconstruction we perform a five-dimensional morphological opening of the multi-resolution support. Morphological opening consists of the successive application of an erosion followed by a dilation. The former removes elements from the multi-resolution support if one of its neighbors (in all five dimensions) is 0, and the latter does the opposite, that is adding elements to the multi-resolution support if one of its direct neighbors is 1. This amounts to a successive shrinking and growing of objects in the five-dimensional multi-resolution support.

Objects spanning all five dimensions of the multi-resolution support are not affected by this operation. However, objects that do not span all five dimensions, and are therefore likely to be noise artifacts, are removed. This allows us to use much lower thresholds, of typically $1.5\sigma_j$, during reconstruction. For the purposes of source finding this is an increase in both sensitivity and reliability.

3.3 Memory Layout and Processing

The described 2D–1D denoising has been implemented in C++ using the à trous algorithm in both the two-dimensional angular, as well as the one-dimensional spectral transformation. The complete storage of all wavelet coefficients would take $J_1 \times J_2$ the amount of memory of the original data, which can easily exceed the available computing resources for the typical size of a data cube of several hundreds of MB. Here, we deal with this major issue by performing the reconstruction on the fly. Such a serialized method only needs to store the angular smoothed version c_j, the angular wavelet coefficients w_{j_1}, and the reconstruction $\tilde{D}$. This way, the memory consumption of the algorithm is greatly reduced, being now independent of the number of angular and spectral scales analyzed.

Another memory concern is the size of the multi-resolution support, that has to store $N_1 \times N_2 \times N_3 \times J_1 \times J_1$ Boolean values, where N_i is the size of the data

cube in pixels along the ith axis. For this purpose, the Standard Template Library (STL) for C++ implements a specialized container that is able to store Boolean values as individual bits rather than bytes, which makes the memory footprint of the multi-resolution support acceptable.

The splitting of the different wavelet transformations makes this denoising scheme a prime candidate for parallel computing. Using the OpenMP[2] library, the angular wavelet decomposition of each spectral channel, as well as the spectral wavelet decomposition of each line of sight, was implemented to be computed in parallel.

4 Simulations

To examine the various aspects in the sections below, we created 1000 simple HI galaxy templates using the GIPSY[3] task `galmod`. The galaxies were simple disks with random inclination and maximum rotational velocity, while we kept the overall brightness profile and rotation curve fixed.

Noise was generated according to the specifications of the WALLABY[4] survey (Koribalski & Staveley-Smith 2009; also see Koribalski 2011, this PASA issue). The models were convolved with a Gaussian beam of approximately $30''$ and inserted into data cubes with an rms noise of 1.8 mJy/beam. The exact specifications are, however, not important for the simulations, since all tested quantities are given in terms of signal-to-noise ratios and the algorithm only operates on the pixel grid of the data. Differing beam sizes should therefore yield the same results if the beam is sampled on the same number of pixels.

The algorithm was run on multiple data cubes of 300 by 300 pixels and 600 channels size that each contained 20 random galaxies at random positions.

4.1 Source Scaling

Since the proposed algorithm is sensitive to the complete source signal in the data as opposed to the peak flux, for example, we scaled each of the 1000 galaxies to a fixed set of integrated signal-to-noise ratios. Since the integrated signal-to-noise ratio (ISNR) is dependent on the volume over which it is calculated, we first determined the optimal volume for each of our models. This was done by starting with the brightest voxel of the model and successively adding the next fainter one. This way, both the total flux and ISNR will increase as a function of the number of voxels added. At a certain point, the flux in the added voxels becomes very low, since one adds the faint 'outskirts' of the model. At this point the ISNR will go down since one adds more noise than source flux. This behavior can be seen in Figure 1. The flux at this optimal ISNR is then scaled to yield the desired ISNRs.
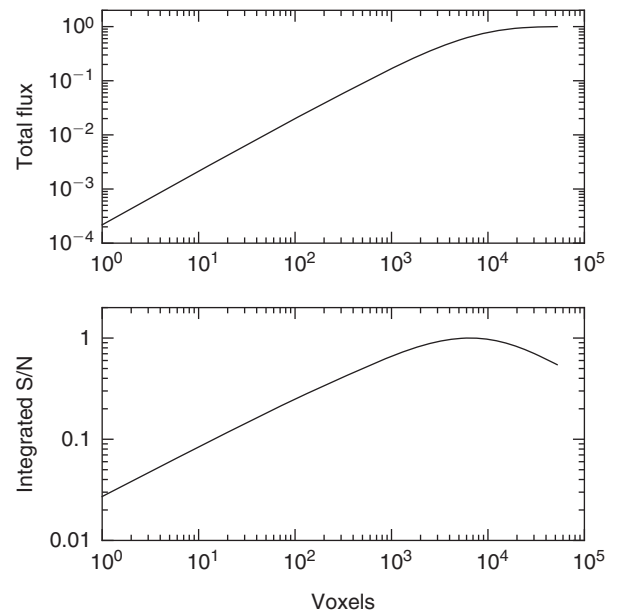
Figure 1 Normalized total flux (top) and normalized integrated signal-to-noise ratio (bottom) as a function of the number of voxels added up, ordered by flux in descending order, for a given source model.
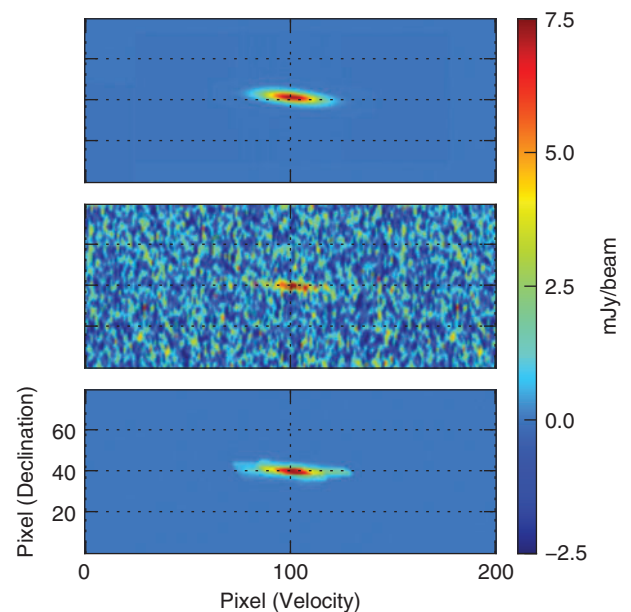

Figure 2 Example of a wavelet reconstruction by the described algorithm. From top to bottom the panels show: source model only; source model with added noise; reconstruction by the algorithm.

4.2 Example Reconstruction

Figure 2 shows the typical result of a reconstruction by the described algorithm. The top panel shows one of our scaled templates. In the middle panel, the simulated noise was added. The bottom panel shows the same data cube after the application of the described denoising algorithm. In general, the reconstruction does not restore the full flux of the inserted model and also has a changed appearance as compared to the model. This is especially true for low
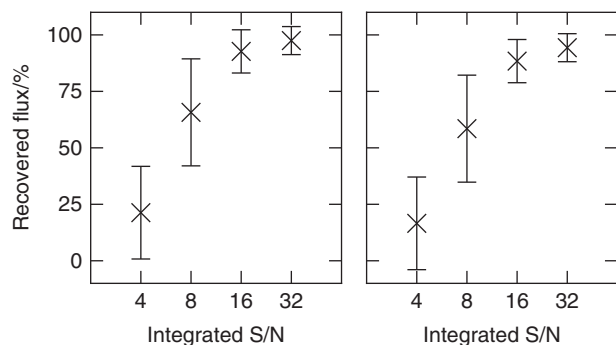
Figure 3 Recovered flux as a function of ISNR. The crosses show the mean of 1000 galaxy models and the error bars indicate the standard deviation in each bin. The left panel shows the recovered flux as obtained from the wavelet reconstruction. The right panel shows the recovered flux when applying the wavelet reconstruction as a mask for data and calculating the flux from the masked data.

Figure 4 Same as Figure 2, but for corrupted data. In addition to the simulated noise, the middle panel shows a sinusoidal varying baseline and added RFI spikes.

signal-to-noise sources, since the reconstruction becomes limited by noise. For more pronounced sources, however, there is a good correlation between model flux and reconstructed flux. This is shown in Figure 3: we plot recovered flux of the reconstruction (left panel) as well as the flux recovered from the data when using the reconstruction as a mask (right panel) as a function of ISNR. Especially for ISNR 32 and 16, the recovered flux matches the model quite well. It is also interesting to note that the flux from the reconstruction seems to be closer to the true flux than the flux calculated from the data. In any case, this evaluation also shows that the masks obtained from the reconstruction should not be used as the final masks without further treatment.

4.3 Robustness

Since real data does not usually contain ideal noise and sources, we evaluated the robustness of the proposed algorithm against two common types of data defects: baseline ripple and RFI.

To simulate these effects, we added a sine wave to one simulated data cube with a varying phase along one angular axis. To simulate the presence of RFI we inserted 30 single-channel spikes in the data and reran the wavelet reconstruction. The result can be seen in Figure 4. Clearly, the wavelet reconstruction is not affected by the rather severe presence of RFI and baseline ripple. This is because both the baseline ripple and RFI are present in scales different from the scales of the sources. By carefully selecting which scales to reconstruct the data from, we can exclude many such defects.

4.4 Completeness and Reliability

The two main measures of the goodness of a source finder are its completeness as a function of source signal and the corresponding reliability. The completeness is expressed as the percentage of sources that have been positively identified by the source finder. The reliability is calculated as the number of true sources divided by the total
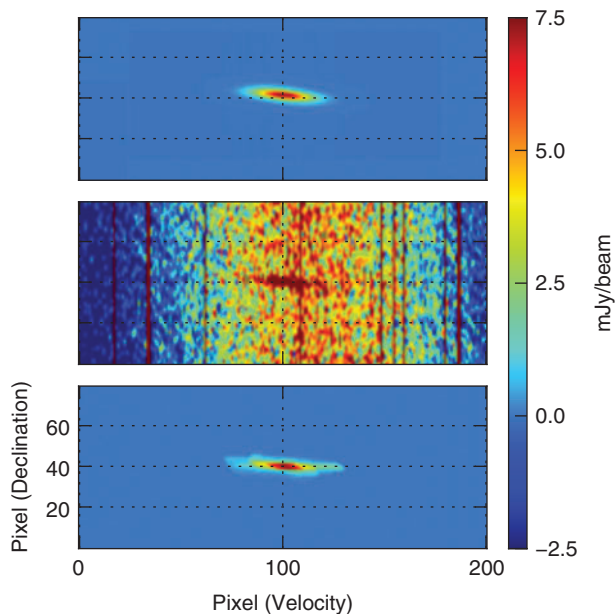
number of objects found by the source finder and gives a measure of the probability that a given object is indeed a source or a false positive.

To test the performance of the algorithm as a source finder, we set up a simple source finding pipeline by using several functions from the SCIPY[5] package `ndimage`. After the wavelet reconstruction, the `ndimage` functions `label` and `find_objects` are used to generate the objects. For this purpose `label` searches the data cube for connected objects, that is, regions where the flux is greater than zero, and marks each region with a unique number. `find_objects` then generates a list of slices that each fully contain one of the labeled objects. Those slices are then used to calculate various parameters like the total flux of the reconstructed object F_R, the total flux in the original data F_D when applying the reconstruction as a mask, and various shape parameters like the size in channels. To check whether a given object is a true detection, we use a noise-free version of the same data set and check for intersections with the noise-free sources above 20% of the peak flux of a galaxy.

4.4.1 False Positives

To achieve a reasonably high completeness even for very faint sources, one has to use very low thresholds which will lead to an increasing number of false positives. After the contrast enhancement, the identification of false positives is a key task of every source finder.

Since the sources of the false positives, noise peaks, are greatly suppressed by the algorithm, they mostly stem from the larger wavelet coefficients where the noise peaks are spread out over a sufficiently large volume to not be
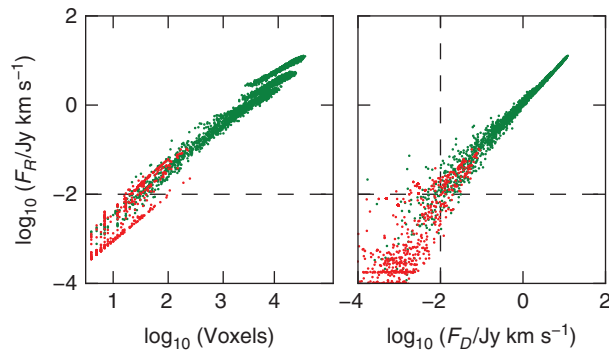
---

[5] http://scipy.org

Figure 5 Distribution of the detection parameters F_R, the flux in the reconstruction, F_D, the flux in the detection as measured on the original data, and the number of voxels a given object occupies. The green points indicate true detections, the red points false detections. The dashed lines indicate $10\,\mathrm{mJy\,km\,s^{-1}}$ for F_R and F_D, respectively.

removed by the morphological opening. This leads to a very low reconstructed total flux F_R and they are therefore easily separated from the real sources.

Figure 5 shows the correlation for three parameters from the simulation with resolved sources (see Section 4.4.3), all directly measured from either the reconstructed or the original data. It is evident that all false positives cluster in one region of the respective plots and that they exhibit very low F_R and F_D. Therefore, by applying a simple cut in the parameter space of the detections, the number of false positives can be greatly reduced without sacrificing much of the completeness. By only taking sources that have both F_R and F_D larger than $10\,\mathrm{mJy\,km\,s^{-1}}$, we exclude 96% of all false positives but only 5% of true positives. The area in which sources fulfill this condition is indicated by the dashed lines in Figure 5.

4.4.2 Point Sources

We tested the completeness as a function of ISNR for both extended sources as well as point sources. To obtain realistic line profiles for the point sources, the extended models were summed in each channel and the resulting spectrum convolved with the beam. The resulting point-source model was then scaled to the desired ISNR.

The results of the run are summarized in Figure 6. Starting from ISNR 0.5, we increased the ISNR by a factor of two from bin to bin. Since the drop in completeness between ISNR 8 and 4 is rather sharp, we ran additional simulations in between those values.

The lower panel in Figure 6 shows that the source finder is indeed sensitive to the extended signal of the sources as we detect sources with a larger line width but lower peak signal-to-noise than the sources we do not detect for smaller line widths.

The reliability for these results is close to 100%. We achieved this by applying the cut discussed in the previous section. Note that Figure 5 was made from the run with resolved sources.
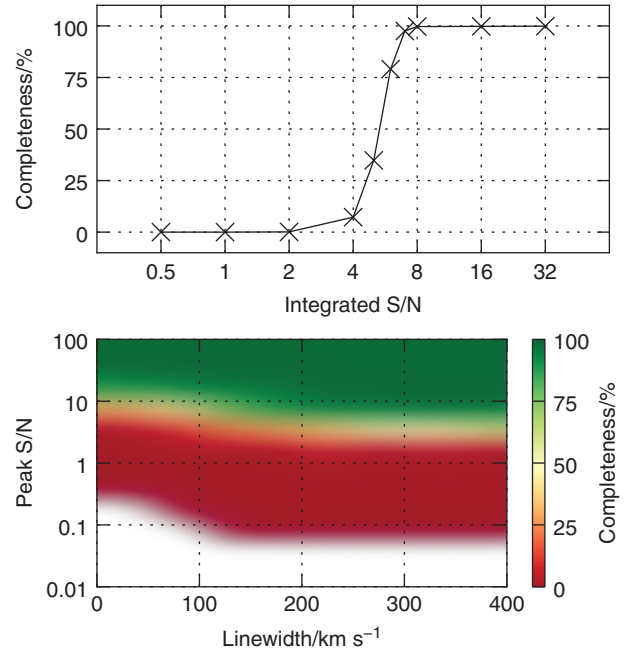

Figure 6 Results from the simulation with spatially unresolved sources. The top panel shows the completeness as a function of the ISNR as probed in our simulations. The lower panel shows the completeness as a function of (logarithmic) peak signal-to-noise ratio and line width of the source. The white areas in the lower plot have not been tested. The reliability for this plot is 99%.
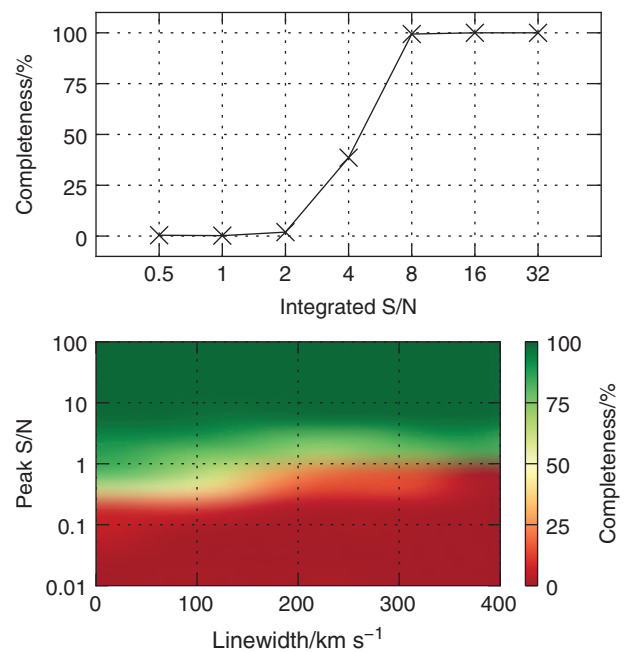

Figure 7 Same as Figure 6 for spatially extended sources. The reliability for this plot is 97% and the fragmentation 3%.

4.4.3 Extended Sources

The second run was made with the extended galaxies which are clearly resolved by the simulated observations. Figure 7 shows the results of this run in a similar fashion to Figure 6. We again cut at F_R and $F_D \lesssim 10\,\mathrm{mJy\,km\,s^{-1}}$ to reach a reliability of 97%.

The boundary from 100% completeness to 0% is substantially smoother than in the case of the point sources. This behavior comes from the fact that extended sources can be extended in the angular domain while at the same time being very narrow in the spectral domain, such as a galaxy seen face-on. This makes it substantially easier to detect galaxies with narrow line widths as long as they are extended in the angular domain. This is also evident from the lower panel in Figure 7, where one can see that narrow line-width galaxies are detected to a lower peak signal-to-noise ratio than large line-width galaxies.

For this simulation, we also encounter a phenomenon usually called fragmentation. We calculate it as the percentage of sources that have been detected two or more times. This can occur when a source with a very large line-width is split into two detections. Furthermore, as mentioned in Section 3.1, wavelet denoising schemes are generally prone to producing artifacts during the denoising process. Because of the simple way we determine whether a source is a true or a false detection, those artifacts can also cause multiple detections of the same source. We are therefore confident that the fragmentation rate will decrease somewhat as the object identification process improves.

5 Summary

We have shown how 2D–1D wavelet denoising schemes can be used for source finding. Even with very simple post-processing of the denoised data, we set up an efficient source-finding pipeline. Particularly because of the algortithm's robustness, it seems promising that it will work well on real data, which is certainly the next test to be passed.

Even though the splitting of the wavelet transformation into a 2D and 1D part avoids some of the difficulties that arise with anisotropic sources, it is far from perfect. A better denoising would involve the use of a full 3D curvelet transformation (Candes et al. 2007; Ying et al. 2005). This transformation is, however, computationally much more difficult and demanding on the available hardware. We therefore think that our approach is a good compromise between sensitivity and computational complexity. But with more powerful hardware or more optimized algorithms, denoising by use of the curvelet transform might become feasible, even for the large data sets we expect from future radio telescopes.

Furthermore we would like to stress that even though this algorithm was developed with Hɪ surveys in mind, it will in principle work for every kind of data that is similar to the data product of such a survey, such as other spectral line surveys.

References

Atwood, W. B., et al., 2009, ApJ, 697, 1071
Barnes, D. G., et al., 2001, MNRAS, 322, 486
Candès, E., Demanet, L., Donoho, D. & Ying, L., 2006, Multiscale Modeling and Simulation, 5, 861
Giovanelli, R., et al., 2005, AJ, 130, 2598
Holschneider, M., Kronland-Martinet, R., Morlet, J. & Tchamitchian, P., 1989, in Wavelets. Time-Frequency Methods and Phase Space, ed. J.-M. Combes, A. Grossmann & P. Tchamitchian (Berlin: Springer-Verlag), 286
Johnston, S., et al., 2008, ExA, 22, 151
Kerp, J., Winkel, B., Ben Bekhti, N., Flöer, L. & Kalberla, P. M. W., 2011, AN, 332, 637
Koribalski, B. S. & Staveley-Smith, L., 2009, ASKAP Survey Science Proposal
Koribalski, B. S., et al., 2004, AJ, 128, 16
Murtagh, F., Starck, J.-L. & Bijaoui, A., 1995, A&AS, 112, 179
Nyquist, H., 1928, Trans AIEE, 47, 617
Oosterloo, T., et al., 2009, in Proceedings of Wide Field Astronomy & Technology for the Square Kilometre Array (SKADS 2009), 4–6 November 2009, Chateau de Limelette, Belgium, available at http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=132
Serra, J., 1982, Image Analysis and Mathematical Morphology (London: Academic Press)
Starck, J.-L. & Bobin, J., 2010, Proceedings of the IEEE, 98, 1021
Starck, J.-L., Fadili, J. & Murtagh, F., 2007, IEEE T Image Process, 16, 297
Starck, J.-L., Fadili, J. M., Digel, S., Zhang, B. & Chiang, J., 2009, A&A, 504, 641
Starck, J.-L., Murtagh, F. & Fadili, J. M., 2010, Sparse Image and Signal Processing: Wavelets, Curvelets, Morphological Diversity (Cambridge: Cambridge University Press)
Winkel, B., Kalberla, P. M. W., Kerp, J. & Flöer, L., 2010, ApJS, 188, 488
Ying, L., Demanet, L. & Candes, E. J., 2005, in Proceedings of SPIE, Volume 5914, Wavelets XI, ed. M. Papadakis, A. F. Laine & M. A. Unser, 351