Classifying Structures in the Interstellar Medium with Support Vector Machines: The G16.05-0.57 Supernova Remnant

Citation

Beaumont, Christopher N., Jonathan P. Williams, and Alyssa A. Goodman. 2011. "CLASSIFYING STRUCTURES IN THE INTERSTELLAR MEDIUM WITH SUPPORT VECTOR MACHINES: THE G16.05-0.57 SUPERNOVA REMNANT." *The Astrophysical Journal* 741 (1): 14. <https://doi.org/10.1088/0004-637x/741/1/14>.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41397411>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

CLASSIFYING STRUCTURES IN THE INTERSTELLAR MEDIUM WITH SUPPORT VECTOR MACHINES: THE G16.05-0.57 SUPERNOVA REMNANT

CHRISTOPHER N. BEAUMONT^{1,2}, JONATHAN P. WILLIAMS¹, AND ALYSSA A. GOODMAN²

¹ Institute for Astronomy, University of Hawai‘i, 2680 Woodlawn Drive, Honolulu, HI 96822, USA; beaumont@ifa.hawaii.edu

² Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

Received 2011 June 20; accepted 2011 July 24; published 2011 October 11

ABSTRACT

We apply Support Vector Machines (SVMs)—a machine learning algorithm—to the task of classifying structures in the interstellar medium (ISM). As a case study, we present a position–position–velocity (PPV) data cube of ^{12}CO $J = 3-2$ emission toward G16.05-0.57, a supernova remnant that lies behind the M17 molecular cloud. Despite the fact that these two objects partially overlap in PPV space, the two structures can easily be distinguished by eye based on their distinct morphologies. The SVM algorithm is able to infer these morphological distinctions, and associate individual pixels with each object at $>90\%$ accuracy. This case study suggests that similar techniques may be applicable to classifying other structures in the ISM—a task that has thus far proven difficult to automate.

Key words: ISM: individual objects (G16.05-0.57, M17) – ISM: supernova remnants – techniques: image processing

Online-only material: animation, color figures

1. INTRODUCTION

Classifying interesting objects in a data set is an essential early step in most analysis tasks. For many objects in astronomy (stars, galaxies, solar system objects, extragalactic supernovae), algorithms can identify and characterize these objects automatically (Irwin 1985; Bertin & Arnouts 1996; Naylor 1998). Structures in the interstellar medium (ISM), however, have proven difficult to classify in this way. These objects—which include, e.g., molecular clouds, infrared dark clouds, bubbles, jets, radiation-shaped pillars, and filaments—are morphologically complex and heterogeneous. The essential properties of these structures are hard to encode.

To take advantage of increasingly wide-area surveys, previous researchers have mainly relied on manual identification of features in the ISM (Churchwell et al. 2006; Helfand et al. 2006; Curtis et al. 2010; Arce et al. 2010). There are several drawbacks to this approach: it is time consuming, non-repeatable and, when identifying complex structures, affected by difficult-to-quantify selection effects related to how specific people perceive an image. Despite recent advances in the study of specific objects (e.g., infrared dark clouds, Peretto & Fuller 2009; filaments, Men’shchikov et al. 2010), automated feature identification in the ISM is still an open problem.

Machine learning algorithms are designed to infer patterns in data that are otherwise difficult to define explicitly. They are grouped into two classes: supervised (in which the algorithm is “trained” to recognize a pattern via a set of training examples) and unsupervised (in which the algorithm identifies groupings within a data set *a priori*).

These algorithms can mechanize the process of object identification, and hence address many of the shortcomings of manual classification—they scale easily to other similar data, and their results are repeatable. A potential drawback of these methods is that, because the classification is not guided by a physical model, they are susceptible to overfitting and to inheriting selection biases within the training data. However, because the machine approach extends easily to other similar data, these

biases are more readily characterized via classification of test data sets.

In this paper, we explore whether machine learning algorithms can be used to catalog structures in the ISM. We present as a case study a spectral line data cube of ^{12}CO emission toward the M17 star-forming region. Emission from this cloud overlaps with G16.05-0.57, a supernova remnant situated behind the cloud. We use the Support Vector Machine (SVM; Vapnik 1995) algorithm to identify the supernova remnant and, on a pixel-by-pixel basis, classify the original data cube. We are able to use this classification to derive the mass and momentum of the supernova remnant, which would not otherwise be possible with these data.

2. THE DATA

G16.05-0.57 is a supernova remnant in the inner galaxy first discovered via its centimeter synchrotron emission by Brogan et al. (2006). Because it is located behind M17, we serendipitously observed the remnant with the James Clerk Maxwell Telescope (JCMT) during an imaging campaign of the latter region. Our observations of G16.05-0.57 were taken on the nights of 2009 June 23–25, using the Heterodyne Array Receiver Program (Smith et al. 2008). The observations target the ^{12}CO $J = 3-2$ line, which traces moderately excited ($h\nu/k = 16$ K) and dense ($n_{\text{crit}} \sim 10^3 \text{ cm}^{-3}$) gas at a resolution of $15''$. The data were acquired via position-switched raster scans, using a reference position of $(\ell, b) = (15^\circ 5', -2^\circ 4')$. Basket weaving was used to reduce striping artifacts in the final map (see, e.g., Section 2.1 of Davis et al. 2010). Weather conditions were grade 2–3 (225 GHz opacity $\tau_{225} = 0.06\text{--}0.10$). To convert from antenna temperature T_A^* to main beam temperature ($T_{\text{mb}} = T_R^* = T_A^*/\eta$) we assumed an aperture efficiency of $\eta = 0.63$ (Buckle et al. 2009).

G16.05-0.57 is visible as a filamentary arc subtending 20 arcmin on the sky (Figure 1). To our knowledge, these are the first molecular line observations of this object. Features belonging to the supernova have typical line widths of 15 km s^{-1}

Figure 1. G16.05-0.57 supernova remnant. The gray scale shows the intensity of our CO $J = 3-2$ map at several velocities, while the blue contours trace the 20 cm continuum emission presented in Brogan et al. (2006). The contour lines are at 1σ , 1.5σ , and 2σ above the background. The compact, filamentary emission especially noticeable at 49 km s^{-1} is from the supernova remnant. The diffuse emission in the first three panels is from clouds in the Scutum arm, while the bright emission at 19 km s^{-1} is part of the M17 molecular cloud.

(A color version of this figure is available in the online journal.)

FWHM, and are centered on a wide range of velocities from -5 to 90 km s^{-1} . It is impossible to derive a kinematic distance from such broad and scattered lines, but supernova features are absorbed by overlapping cloud emission at 40 km s^{-1} . This corresponds to a near kinematic distance of 3.4 kpc, coincident with the Scutum galactic arm. The fact that G16.05-0.57 emits CO line emission suggests that the remnant is in the process of colliding with another molecular cloud, and we are observing the shocked interface of this collision (Scoville et al. 1977; van Dishoeck et al. 1993).

The morphology of structures in this data cube makes for an appropriate case study of automated feature classification in molecular line data sets. A representative position–velocity slice through the data cube in Figure 2 shows that the remnant overlaps molecular cloud emission from M17 and the Scutum arm. However, the supernova has markedly different structure; features from the remnant are spatially compact and extend over tens of km s^{-1} in velocity space, while the other structures it overlaps with are more spatially diffuse and kinematically narrow. This suggests that a learning algorithm may be able to distinguish emission from G16.05-0.57 based on its unique morphology.

3. SUPPORT VECTOR MACHINES

The SVM algorithm is a supervised learning algorithm that attempts to segregate data points into two categories, based on a representative sample of data belonging to each category. The method has recently been applied to many diverse problems in astronomy, including redshift estimation (Wang et al. 2008), galaxy morphology identification (Huertas-Company et al. 2008), and time series analysis (Kim et al. 2011). Here we provide a basic overview of the algorithm, but refer the reader to Press et al. (2007) for a deeper and more precise derivation, and to Vapnik (1999) for a discussion of the algorithm’s foundation in statistical learning theory. In what follows, we use the SVM^{light} implementation by Joachims (1999). While this code is written in C, we have written a set of wrappers to use these tools within IDL.

During training, the SVM algorithm takes as input a feature vector for each training example—a set of N quantities that describe the discriminating properties of that object. These numbers can be thought of as coordinates for a vector in N -dimensional feature space. Using a training set of pre-classified feature vectors, the algorithm searches for a decision

Figure 2. Position–velocity slice through the data cube, manually labeling some of the emission from the supernova remnant (blue) and foreground clouds (red). The inset shows a position–position view of the cube, and the red line shows the location of the PV slice. The bright feature at 20 km s^{-1} is the M17 cloud, while the emission at 40 km s^{-1} is from the Scutum arm. Individual supernova emission features are distinguished by their narrow spatial extent and broad velocity profiles. Shocked emission from the expanding supernova is detected at all velocities throughout the cube. As a result, it overlaps other structures in several places.

(A color version of this figure is available in the online journal.)

boundary in feature space that optimally separates examples from each category. New data points are then assigned a classification based on the side of the decision boundary on which they fall.

More specifically, SVM seeks a decision boundary B that maximizes a fitness function F given by

$$F = M - C \sum_i \xi_i(B, M). \quad (1)$$

Here M is the margin of the boundary. SVM attempts to separate training data with a large gap between each data point and the boundary, and M defines the size of this gap. $\xi_i(B, M)$ is the degree to which training example i violates this criterion. If example i is separated by more than M from B (and on the correct side), $\xi_i = 0$. Otherwise, ξ_i is the distance that i would have to be moved to satisfy this condition. The adjustable cost parameter C sets a tradeoff between large margins M and poor classifications ξ_i .

For a better understanding of Equation (1), consider first the case where C is very large, the feature space is two-dimensional, and the decision boundary is restricted to a line (Figure 3, where the different symbols denote training examples from two different classes). In this scenario, the algorithm first optimizes Equation (1) over M for a fixed boundary. Since C is very large, even a single training example on the wrong side of the margin

Figure 3. Simple classification task in two dimensions. Dots and Xs represent training examples from two classes. Each training example consists of a two-dimensional feature vector, represented in the figure by the position on the plane. The gray and white areas denote which regions of feature space an SVM classifier assigns to each class. The SVM algorithm seeks a decision boundary that maximizes the distance between training data and the boundary (the margin, shown here as dotted lines).

Figure 4. Illustration of how the cost factor C influences the SVM classification. The symbols and colors are the same as in Figure 3. The left and right figures show classifications with low and high values of C . Decreasing C decreases the penalty for misclassifications near the boundary, increasing robustness against outliers.

will heavily penalize Equation (1). Thus, the optimal M will be the margin that just touches the training example closest to B (i.e., the largest M that satisfies $\sum_i \xi_i(B, M) = 0$). The algorithm repeats this optimization over all boundaries, finding the plane that can accommodate the largest margin. The final classification is illustrated in Figure 3; the background shading shows how the algorithm classifies feature space. The dotted line traces the margin on either side of the boundary.

Next, consider the impact of reducing C . Individual ξ_i terms now penalize Equation (1) less heavily. The optimal boundary in this scenario may be one which misclassifies a small number of outliers, but can afford to partition the remaining data with a larger margin M . This is depicted in Figure 4.

In addition to C , a second set of adjustable parameters characterize a “kernel function.” The kernel function determines the topology of the decision surface. In the simplest case, decision boundaries are hyperplanes in feature space. In this paper, we use the radial basis kernel function (RBF), a popular and effective kernel that allows for non-planar decision boundaries. The RBF kernel has one adjustable parameter, γ , that controls the curvature of the decision surface. For example, Figure 5 shows an SVM classification of a data set with three different values of γ . Low values lead to stiff boundaries, while high values lead

Figure 5. Illustration of how the RBF kernel parameter γ influences the SVM classification. As γ increases, the decision boundary becomes more curved, and conforms more tightly to individual training examples. The margin in the rightmost plot wraps very tightly around the Xs, and is not drawn.

to curved boundaries that may overfit to the training data. In this application, overfitting refers to the case when the decision boundary conforms too tightly to the individual training examples, and the larger-scale organization of the data is ignored.

A small complication arises when different elements of the feature vector have different scales. Components of the feature vector with the very large numerical values will dominate ξ , and the remaining components will have a negligible effect on the classification. To circumvent this, we normalize each element of the feature vectors, such that the dispersions of each element across the data are equal.

From a practical standpoint, then, an SVM-based classification task involves four steps.

1. Manually classify a representative subset of the data (Section 3.1).
2. For each classified example, create a feature vector. This vector should encode the properties of an object that make it identifiable (Section 3.2).
3. Choose a kernel function. In this study, we restrict our attention to the radial basis function.
4. Train the algorithm, and optimize the classification by adjusting free parameters (C and γ in this case; Section 3.3).

In what follows we apply these steps to disentangle overlapping emission from the supernova remnant G16.05-0.57 and foreground molecular clouds.

3.1. Manually Classifying Data

The success of any supervised learning algorithm is limited by how representative the training set is. Because our aim is a pixel-by-pixel classification of the data, our training set consists of a subset of these pixels, manually classified as associated with the supernova remnant (or not). We carried out this manual identification on subsets of four position–velocity slices through the cube. The classification of one of the planes is shown in Figure 2. In total, the training set explicitly labels 0.4% of the pixels in the cube (5% of the pixels above 3σ); however, as we show in Section 4.1, only a small fraction (5%) of these examples are ultimately necessary.

3.2. Creating a Feature Vector

Each pixel in our training set must be assigned a feature vector—a list of numerical attributes that distinguish between the supernova and unassociated foreground objects. When classifying a pixel in the data by eye, it is sufficient to examine the pixels in the immediate vicinity. In particular, a $30 \times 30 \times 100$ pixel sub-cube in position–position–velocity

(PPV) space is sufficient for a human to classify the pixel in the center of that cube. At the presumed 3.4 kpc distance to the supernova, this corresponds to a $3 \text{ pc} \times 3 \text{ pc} \times 25 \text{ km s}^{-1}$ region. In principle, one could use the intensities of these 9×10^4 pixels as a feature vector for the central point. In practice, such a large feature vector is prohibitively slow. We tested three strategies to compress this information. Each of these strategies defines a different feature vector.

Moment. For each pixel p_i , we extract the surrounding $30 \times 30 \times 100$ pixels. We calculate the mean intensity of this cube, and the first and second moments along each direction through the data. These seven numbers constitute the feature vector for p_i . Relative to other cloud emission, supernova features have large velocity dispersions and small spatial dispersions—this information is encoded in the moments of the data.

Derivative. Spatial derivatives are sensitive to edges in images, and such information can be used to identify filamentary structures in astronomical data (Molinari et al. 2011). We generate a feature vector that encodes this information. We approximate the gradient in each direction and pixel location using the Sobel edge detection operator. To generate the feature vector for pixel $p_i = (x_0, y_0, z_0)$, we sample profiles of each derivative along each direction through the pixel:

$$\begin{aligned}
 P_1 &= \{\partial_x(x_0 + \delta, y_0, z_0) \mid -15 \leq \delta < 15\} \\
 P_2 &= \{\partial_x(x_0, y_0 + \delta, z_0) \mid -15 \leq \delta < 15\} \\
 P_3 &= \{\partial_x(x_0, y_0, z_0 + \delta) \mid -50 \leq \delta < 50\} \\
 &\vdots \\
 P_9 &= \{\partial_z(x_0, y_0, z_0 + \delta) \mid -50 \leq \delta < 50\} \\
 P_{\text{final}} &= \{P_1 \cup P_2 \cup \dots \cup P_9\}
 \end{aligned}$$

For convenience, we further downsample P_{final} to 60 elements, which defines the feature vector. We determined the degree of downsampling that was appropriate by examining the widths of typical features in the derivative profiles by eye. Nevertheless, this smoothing may lead to worse performance.

PCA. We approximate the $30 \times 30 \times 100$ sub-cube around each pixel p_i as a linear combination of 15 representative “basis cubes.” We derive the basis cubes using principal component analysis (PCA; Francis & Wills 1999), and these basis cubes capture $\sim 95\%$ of the variance in the data. The 15 weights in the linear combination define our final feature vector for p_i . This is essentially a (lossy) compression of the data and, unlike the first two methods, does not explicitly encode any intuitive, identifying characteristics. Nevertheless, this expression of the data has proved useful in other classification tasks (e.g., asteroid

Figure 6. Performance of each SVM classifier as a function of intensity. We measure the accuracy of the classification by comparing to a manually classified subset of the data (independent from the training data). The lines indicate the fraction of structures correctly classified, with brightness temperatures exceeding the value on the X-axis. The three lines denote SVMs trained using different feature vectors, and the vertical line denotes the 3σ noise level. All three classifiers are more accurate when classifying brighter structures.

(A color version of this figure is available in the online journal.)

taxonomy, Tholen 1984; stellar spectral types, Singh et al. 1998; star/galaxy separation, Cabanac et al. 2002). PCA has also been used to decompose and analyze molecular cloud structure (Heyer & Schloerb 1997; Brunt et al. 2009).

3.3. Training and Optimization

As discussed above, two free parameters influence the training process: γ and C . We use cross-validation to choose optimal values for these parameters. We first partition our classification examples into two independent sets. The first (the training set) is used to train the classifier using a given value for (C, γ) . We then apply the classifier to the second (validation) data set, and measure the accuracy of the identification. We repeat this process for different values of (C, γ) to maximize the performance on the validation set. This approach provides some protection against overfitting, since overfits to the training data will poorly classify the validation set. To maximize the independence of the training and data set, the two samples were drawn from different regions of the cube.

4. RESULTS AND DISCUSSION

4.1. Classification Performance

We evaluate the performance of each classifier by comparing the accuracy with which it classifies the validation data described in Section 3.3. The accuracy is simply the fraction of correctly labeled pixels. We find that the feature vector that encodes the moments of the intensity achieves the highest performance.

The maximum accuracies achieved using the Moment, Derivative, and PCA feature vectors were 83%, 77%, and 75%,

Figure 7. SVM classification of the same data shown in Figure 2. Blue pixels are classified as belonging to the supernova, while red pixels denote foreground cloud emission.

(A color version of this figure is available in the online journal.)

respectively (the y-intercept of Figure 6). However, misclassifications are biased toward low-intensity pixels. This is to be expected since, as Figure 2 shows, pixels corresponding to blank sky have been included in both classes in the classification set. Hence, the proper classification of these faint pixels is not well defined. Misclassifying noise is not problematic, however, because simple thresholding later in the analysis can separate signal from background. For most data analysis purposes, it is more important that emission features be correctly classified.

Figure 6 shows the accuracy at which each classifier identifies pixels above a given intensity threshold. Here, the moment-based classifier has an accuracy of 90% for emission detected at 3σ , and exceeds 95% accuracy for the brightest pixels in the data. Figure 7 shows the classification of the data in Figure 2, using the Moment feature vectors. Figure 8 shows the classification of several position–position channels. Many of the misclassified bright pixels lie near the boundary between supernova and cloud emission. This is perhaps expected, because the moments of the intensity are measured (and implicitly smoothed) over a $30 \times 30 \times 100$ sampling window. Nevertheless, several supernova lines intersect and, presumably, extend behind cloud material. In most cases, the SVM classifier at least partially follows these transitions from supernova to cloud.

There are three reasons why an SVM classifier misclassifies data. First, data from two classes may not segregate perfectly in feature space (due to either noise in the data, errors in the training set, or a poorly designed feature vector). Second, the topology of the decision boundary (determined by the SVM kernel function) may not be able to conform to the distribution of training data in feature space. Finally, the training examples

Figure 8. Same as Figure 7, but for the position–position slices through the cube also shown in Figure 1.

(An animation and color version of this figure are available in the online journal.)

may insufficiently sample how data are distributed in feature space. Each of these possibilities has implications for how to design and improve classification pipelines.

Figures 9 and 10 provide some insight into what limits the performance of our classification using the Moment feature vectors. Using the optimal values of C and γ found above, we measured the learning rate—the accuracy of the classification as a function of training set size. Figure 9 shows this function, and suggests that most of the meaningful information is contained within the first few hundred examples. Remaining examples contain redundant information, and confer little performance gain.

Misclassifications in the validation set may occur because the training set is not representative enough, and the validation set samples a systematically different region of feature space. To test this possibility, Figure 10 shows the classification accuracy when we re-train using the original training set, augmented with a subset of the validation data that were originally misclassified. The solid line re-optimizes C and γ at each step. The dashed line shows the result when we fix these parameters to the values used above.

Note that in this experiment, part of the validation data is now explicitly included in the training set. This increases the risk of overfitting the training data (i.e., devising arbitrary rules that fit the training data but do not generalize well to new

data). As a trivial case of overfitting, when a classifier is re-trained with a correction example in the validation set, it corrects the misclassification for that example only. The dotted line in Figure 10 depicts this scenario. Any meaningful performance gain should fall above this line, since the classifier should ideally use the information in each correction example to correct many additional mistakes.

The figure shows that, even when presented with additional correction examples, the algorithm shows essentially no performance gain. This figure rules out the possibility that misclassifications in the validation set are due to the training and validation sets sampling different regions of feature space. Instead, this strongly suggests that the classification task is limited by the partial overlap of the two classes in feature space. Additional training and correction examples are of little help in this situation, and a better feature vector is needed for further performance gain.

We do not claim that these classifiers will necessarily generalize to other data sets or classification tasks. Other applications likely require re-training the SVM algorithm using the data at hand, or testing new feature vectors. However, Figure 6 suggests that the SVM algorithm is capable of identifying morphological differences in the ISM, and Figure 9 implies that this task can be taught efficiently, with little manual classification. For example, we have started to investigate whether wind-blown bubbles

Figure 9. Accuracy of the classifier using the Moment feature vector, as a function of the training set size. The accuracy is measured using the entire data set (i.e., no filtering based on intensity).

can be identified using *Spitzer* colors and edge information as a feature vector (C. N. Beaumont et al. 2012, in preparation).

4.2. Mass and Momentum of SNR G1605-0.57

As mentioned above, the CO emission from G16.05-0.57 is likely due to the remnant's collision with a molecular cloud, presumably in the Scutum galactic arm at 3.4 kpc. Our pixel-level classification of the data allows us to analyze the properties of this emission in isolation from foreground material. Here we derive an estimate of the mass and momentum associated with the cloud/remnant collision.

In the limit that all material along the line of sight can be described by a single excitation temperature T_{ex} , the equation for the observed radiation temperature T_R is (Ginsburg et al. 2011)

$$T_R = T_0 \left(\frac{1}{e^{T_0/T_{\text{ex}}} - 1} - \frac{1}{e^{T_0/T_{\text{CMB}}} - 1} \right) (1 - e^{-\tau}) f, \quad (2)$$

where $T_0 = h\nu/k = 16.6$ K for CO 3–2, τ is the optical depth, and f is the beam filling factor (which we take to be 1). Furthermore, the column density of the $J = 3$ state is given by

$$N_{J=3} = \frac{8\pi\nu^3}{c^3 A_{32}} \frac{1}{e^{T_0/T_{\text{ex}}} - 1} \int \tau_\nu dv, \quad (3)$$

where A_{32} is the Einstein A coefficient. Equation (2) can be solved for τ , such that Equation (3) explicitly depends only on T_{ex} and $T_R(v)$ (see, e.g., Equation (A8) of Ginsburg et al. 2011).

To constrain the excitation temperature and opacity in the line centers, we obtained supplementary ^{13}CO $J = 3-2$ observations toward three bright knots of emission. Assuming that the filling factor of the gas and excitation temperature of the two CO isotopologues are the same, their intensity ratio gives the gas opacity:

$$\frac{T_{12}}{T_{13}} = \frac{\nu_{12}}{\nu_{13}} \frac{1 - \exp[-\tau_{12}]}{1 - \exp[-\tau_{12}X_{13/12}]}, \quad (4)$$

Figure 10. Accuracy of the classifier using the Moment feature vector, as mistakes from the test data set are added to the training set. The solid line re-optimizes C and γ at each step. The dashed line fixes these parameters at their original optimal values. The dotted line is the trivial case of overfitting, when each correction example causes the algorithm to correctly identify one additional example in the test set.

Figure 11. Normalized histogram of the ^{12}CO opacity inferred from Equation (4), in regions where both ^{13}CO and ^{12}CO are detected.

where $X_{13/12}$ is the abundance ratio of $^{13}\text{CO}/^{12}\text{CO}$, which we take to be 70. Figure 11 shows the inferred opacity for all pixels where we detect emission from both isotopes. The typical optical depth is 3–5.

Plugging $\tau = 4$ into Equation (2) gives an estimate of the excitation temperature along each line of sight, which we find to be $T_{\text{ex}} = 15-30$ K. This in turn allows us to evaluate Equation (3). Finally, we convert from $N(J = 3)$ to $N(\text{CO})$ assuming that the population levels are thermalized, and to $N(\text{H}_2)$ using an abundance ratio $X(\text{CO}) = N(\text{CO})/N(\text{H}_2) = 10^{-4}$. The abundance of CO within shocks is uncertain. In their study of CO and H_2 vibrational lines in the C-type shocks of the Orion KL region, Watson et al. (1985) measure an $X(\text{CO})$

Figure 12. Column density map of G16.05-0.57. The scale bar assumes a distance of 3.4 kpc.

(A color version of this figure is available in the online journal.)

of 1.2×10^{-4} . On theoretical grounds, the value of $X(\text{CO})$ for dissociative shocks may be enhanced by up to a factor of 100 if the re-formation of H_2 in post-shock gas is less efficient than CO (van Dishoeck et al. 1992). Thus, the actual value of $X(\text{CO})$ depends on both the type and strength of shocks in G16.05-0.57, as well as the microphysics of grain catalysis in shocked gas.

Figure 12 presents the column density map of the supernova remnant obtained from this analysis. The angularly integrated column density of the map is $\int N(\text{H}_2) d\Omega = 1.1 \times 10^{16} \text{ cm}^{-2} \text{ sr}$. If we further assume that the remnant is located within the Scutum arm at 3.4 kpc, this implies a total mass of $M = 2300 M_\odot$. We approximate the velocity of the gas at each point by the observed velocity dispersion—this is a lower limit, since it only accounts for motion along our line of sight. Nevertheless, this implies a total momentum of $2.2 \times 10^4 M_\odot \text{ km s}^{-1}$. For comparison, the typical momentum of a supernova explosion is $\sim 1\text{--}10 M_\odot \times 10^4 \text{ km s}^{-1} \sim 10^4\text{--}10^5 M_\odot \text{ km s}^{-1}$.

The characteristic width of filamentary features in the remnant is $30''\text{--}60'' = 0.5\text{--}1 \text{ pc}$. Taking this to be the line-of-sight depth of supernova emission implies a volume density of $n = N/l \sim 500\text{--}1000 \text{ cm}^{-3}$. This is roughly 1–2 orders of magnitude lower than the densities van Dishoeck et al. (1993) measured toward the supernova remnant/molecular cloud collision IC 443. The most likely explanation of this discrepancy is that the filamentary features in our data consist of unresolved filamentary substructure. If this is the case, then the characteristic depth l would be smaller, and the corresponding volume density higher. However, this would not affect our mass and momentum measurements, since the increase in volume density is cancelled out by the decrease in solid angle.

5. CONCLUSION

We have presented a case study of a supervised learning task applied to classifying structures in the ISM. The M17 molecular cloud overlaps shocked CO emission from G16.05-0.57 but, because of the supernova’s distinct morphology in PPV space, the two objects are readily distinguishable by eye. The SVM classification algorithm is able to learn these morphological differences using a representative sample of manually classified

pixels. We emphasize several important characteristics of this approach as follows.

1. Machine-based classification of ISM structures permits a pixel-level classification of data sets. This level of refinement is often prohibitively cumbersome via manual identification.
2. By using an independent set of manually classified validation data, we can characterize the quality of this classification. We can further use this information to refine and improve the algorithm’s performance.
3. Extracting information about the moments of the intensity distribution in our data produced the most effective classification. Other information (the weights in a PCA, spatial derivatives) was less successful.
4. Only a very small fraction of the data ($\sim 0.1\%$) needs to be categorized to train the algorithm. An efficient approach for future work may be to evaluate the classifier’s performance as the training set is assembled, to better understand when the training set is large (and representative) enough.

This case study suggests that automated algorithms are capable of identifying complex structures seen in the ISM. Such an approach may be useful in analyses of current and future surveys of the Milky Way’s ISM, particularly when identifying morphologically distinct structures like bubbles, pillars, and filaments.

We are grateful to C. Brogan for sharing her 20 cm data of G16.05-0.57, and to K. Binsted for discussions regarding this manuscript. We also thank the anonymous referee for his or her careful feedback. J.P.W. acknowledges support from NSF. This material is based upon work supported by the National Science Foundation under Grant No. AST-0908159.

REFERENCES

- Arce, H. G., Borkin, M. A., Goodman, A. A., Pineda, J. E., & Halle, M. W. 2010, *ApJ*, **715**, 1170
- Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393
- Brogan, C. L., Gelfand, J. D., Gaensler, B. M., Kassim, N. E., & Lazio, T. J. W. 2006, *ApJ*, **639**, L25
- Brunt, C. M., Heyer, M. H., & Mac Low, M.-M. 2009, *A&A*, **504**, 883
- Buckle, J. V., Hills, R. E., Smith, H., et al. 2009, *MNRAS*, **399**, 1026
- Cabanac, R. A., de Lapparent, V., & Hickson, P. 2002, *A&A*, **389**, 1090
- Churchwell, E., Povich, M. S., Allen, D., et al. 2006, *ApJ*, **649**, 759
- Curtis, E. I., Richer, J. S., Swift, J. J., & Williams, J. P. 2010, *MNRAS*, **408**, 1516
- Davis, C. J., Chrysostomou, A., Hatchell, J., et al. 2010, *MNRAS*, **405**, 759
- Francis, P. J., & Wills, B. J. 1999, in ASP Conf. Ser. 162, Quasars and Cosmology, ed. G. Ferland & J. Baldwin (San Francisco, CA: ASP), 363
- Ginsburg, A., Bally, J., & Williams, J. P. 2011, *MNRAS*, in press (arXiv:1106.1433)
- Helfand, D. J., Becker, R. H., White, R. L., Fallon, A., & Tuttle, S. 2006, *AJ*, **131**, 2525
- Heyer, M. H., & Schloerb, F. P. 1997, *ApJ*, **475**, 173
- Huertas-Company, M., Rouan, D., Tasca, L., Soucail, G., & Le Fèvre, O. 2008, *A&A*, **478**, 971
- Irwin, M. J. 1985, *MNRAS*, **214**, 575
- Joachims, T. 1999, in Advances in Kernel Methods—Support Vector Learning, ed. B. Schölkopf, C. Burges, & A. Smola (Cambridge, MA: MIT Press), 41
- Kim, D.-W., Protopapas, P., Byun, Y.-I., et al. 2011, *ApJ*, **735**, 68
- Men’shchikov, A., André, Ph., Didelon, P., et al. 2010, *A&A*, **518**, L103
- Molinari, S., Schisano, E., Faustini, F., et al. 2011, *A&A*, **530**, A133
- Naylor, T. 1998, *MNRAS*, **296**, 339
- Peretto, N., & Fuller, G. A. 2009, *A&A*, **505**, 405
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, Numerical Recipes: The Art of Scientific Computing (3rd ed.; Cambridge: Cambridge Univ. Press)
- Scoville, N. Z., Irvine, W. M., Wannier, P. G., & Predmore, C. R. 1977, *ApJ*, **216**, 320

- Singh, H. P., Gulati, R. K., & Gupta, R. 1998, [MNRAS](#), **295**, 312
- Smith, H., Buckle, J., Hills, R., et al. 2008, [Proc. SPIE](#), **7020**, 24
- Tholen, D. J. 1984, PhD thesis, Univ. Arizona
- van Dishoeck, E. F., Glassgold, A. E., Guelin, M., et al. 1992, in IAU Symp. 150, *Astrochemistry of Cosmic Phenomena*, ed. P. D. Singh (Dordrecht: Kluwer Academic Publishers), 285
- van Dishoeck, E. F., Jansen, D. J., & Phillips, T. G. 1993, [A&A](#), **279**, 541
- Vapnik, V. N. 1995, *The Nature of Statistical Learning Theory* (New York: Springer)
- Vapnik, V. N. 1999, [IEEE Trans. Neural Netw.](#), **10**, 988
- Wang, D., Zhang, Y., & Zhao, Y. 2008, in ASP Conf. Ser. 394, *Astronomical Data Analysis Software and Systems XVII*, ed. R. W. Argyle, P. S. Bunclark, & J. R. Lewis (San Francisco, CA: ASP), 509
- Watson, D. M., Genzel, R., Townes, C. H., & Storey, J. W. V. 1985, [ApJ](#), **298**, 316