Velocity-coherent Filaments in NGC 1333: Evidence for Accretion Flow?

Michael Chun-Yuan Chen[1] , James Di Francesco[1,2] , Erik Rosolowsky[3] , Jared Keown[1] , Jaime E. Pineda[4] ,
Rachel K. Friesen[6] , Paola Caselli[4] , How-Huan Chen[5] , Christopher D. Matzner[6] , Stella S. Offner[5] , Anna Punanova[7] ,
Elena Redaelli[4] , Samantha Scibelli[8] , and Yancy Shirley[9]

[1] Department of Physics and Astronomy, University of Victoria, Victoria, BC, V8P 1A1, Canada; mcychen@uvic.ca
[2] Herzberg Astronomy and Astrophysics, National Research Council of Canada, 5071 West Saanich Rd., Victoria, BC, V9E 2E7, Canada
[3] Department of Physics, University of Alberta, Edmonton, AB, Canada
[4] Max-Planck-Institut für extraterrestrische Physik, Giessenbachstrasse 1, D-85748 Garching, Germany
[5] Department of Astronomy, University of Texas at Austin, Austin, TX 78712, USA
[6] Department of Astronomy & Astrophysics, University of Toronto, 50 St. George St., Toronto, Ontario, M5S 3H4, Canada
[7] Ural Federal University, 620002 Mira st. 19, Yekaterinburg, Russia
[8] Department of Astronomy, University of Arizona, Tucson, AZ 85721, USA
[9] Steward Observatory, 933 North Cherry Ave., Tucson, AZ 85721, USA
Received 2019 October 11; revised 2020 February 5; accepted 2020 February 5; published 2020 March 6

Abstract

Recent observations of global velocity gradients across and along molecular filaments have been interpreted as signs of gas accreting onto and along these filaments, potentially feeding star-forming cores and protoclusters. The behavior of velocity gradients in filaments, however, has not been studied in detail, particularly on small scales (<0.1 pc). In this paper, we present MUFASA, an efficient, robust, and automatic method to fit ammonia lines with multiple velocity components, generalizable to other molecular species. We also present CRISPY, a PYTHON package to identify filament spines in 3D images (e.g., position–position–velocity cubes), along with a complementary technique to sort fitted velocity components into velocity-coherent filaments. In NGC 1333, we find a wealth of velocity gradient structures on a beam-resolved scale of ~0.05 pc. Interestingly, these local velocity gradients are not randomly oriented with respect to filament spines and their perpendicular, i.e., radial, component decreases in magnitude toward the spine for many filaments. Together with remarkably constant velocity gradients on larger scales along many filaments, these results suggest a scenario in which gas falling onto filaments is progressively damped and redirected to flow along these filaments.

1. Introduction

Molecular cloud filaments appear to play a pivotal role in star formation. In addition to being featured prominently in star-forming regions (e.g., Schneider & Elmegreen 1979; Bally et al. 1987) and being ubiquitous in molecular clouds at large (e.g., André et al. 2010), filaments appear to harbor most of the observed dense cores (e.g., Men'shchikov et al. 2010; Könyves et al. 2015), the smallest structure from which stellar systems emerge (see Di Francesco et al. 2007). Theoretically, filaments appear to form naturally from supersonic turbulent motions of a cloud in numerical simulations, both in the absence (e.g., Porter et al. 1994) and in the presence (e.g., Jappsen et al. 2005) of self-gravity. Moreover, filaments appear to be analytically the most favored structure to grow locally and fragment readily under weak perturbations in a finite cloud (Pon et al. 2011). Such properties likely make filaments highly effective at assembling dense cores from a molecular cloud prior to, or even in the absence of, an overwhelming, global cloud collapse.

How dense cores relate to their host filaments is currently not well understood. Gravitationally induced fragmentation along filament lengths, like those found in numerical models (e.g., Bastien et al. 1991; Inutsuka & Miyama 1997), has been suggested to be how supercritical filaments produce cores, as inferred by *Herschel* observations (see André et al. 2014). While Hacar & Tafalla (2011) found that dense structures in the L1517 filament correlate with oscillatory line-of-sight velocities, suggesting filament fragmentation, such behavior has not been generally observed in other filaments (e.g., Tafalla & Hacar 2015). Other core formation mechanisms, such as those that form cores and filaments simultaneously in simulations (e.g., Gong & Ostriker 2011; Chen & Ostriker 2014, 2015; Gómez & Vázquez-Semadeni 2014), may thus play an important role in star formation as well.

In addition to forming cores, filaments in simulations accrete material from their surroundings and transport mass along their lengths to feed dense cores and protoclusters (e.g., Balsara et al. 2001; Smith et al. 2011, 2016; Gómez & Vázquez-Semadeni 2014). Indeed, velocity gradients observed across (e.g., Palmeirim et al. 2013; Fernández-López et al. 2014; Dhabal et al. 2018) and along (e.g., Friesen et al. 2013; Kirk et al. 2013) filaments have been interpreted as evidence for such accretion onto and along these filaments, respectively. Further kinematic studies are needed to understand how such filaments fit within the wide variety of existing models and how they assemble mass in star formation in detail.

A filament that appears to be a single, coherent (i.e., continuous) structure on the sky may not necessarily be truly coherent in three-dimensional, position–position–position (ppp) space. Multiple structures that are distinct in ppp space can appear as a single structure by mere projection along lines of sight. With CO observations, Hacar et al. (2013) showed that a seemingly coherent filament on the sky can in fact contain

multiple velocity-coherent "fibers" when viewed in position–position–velocity (ppv) space.

While some ppv fibers may indeed trace physical ppp subfilaments like those produced in simulations (e.g., Moeckel & Burkert 2015; Smith et al. 2016; Clarke et al. 2017), synthetic CO observations of a simulation showed that ppv fibers do not necessarily map well onto ppp structures, and vice versa (Clarke et al. 2018). Structures that are coherent in ppv space can still suffer from line-of-sight confusion when distinct ppp structures possess similar velocities (e.g., Beaumont et al. 2013). Such a scenario can be common, for example, when multiple ppp structures are swept up by a large-scale flow. Fortunately, denser gas tracers such as NH_3 and N_2H^+ are expected to be less susceptible to these problems due their lower volume-filling fraction in a cloud. This claim seems to be supported by Tafalla & Hacar (2015), who observed only a single N_2H^+ ppv fiber over each line of sight where multiple CO ppv fibers had been detected earlier by Hacar et al. (2013).

Regardless of how well ppv coherent (hereafter velocity-coherent) structures map onto ppp space, multicomponent line modeling is needed to avoid deriving erroneous gas properties that are unphysical. Kinematic analyses that perform multiple-component fits to a large number of spectra, however, are uncommon. This situation is due to the typical need for human intervention in popular least-squares fitting methods, such as the Levenberg–Marquardt (LM; Levenberg 1944; Marquardt 1963; Moré 1978) method, and the inefficiencies associated with many automated approaches, such as the grid-search or Markov Chain Monte Carlo (MCMC) methods.

Recent automated methods for multicomponent fits, such as Behind The Spectrum (BTS; Clarke et al. 2018) and GAUSSPY + (Riener et al. 2019), work only with optically thin, Gaussian lines by design. Other methods that fit hyperfine lines, such as those used by Henshaw et al. (2016, SCOUSE) and Hacar et al. (2017), are semiautomatic and hence are still subject to human biases. An efficient, automated method that fits hyperfine lines is therefore highly desirable for kinematic studies that use species like NH_3 and N_2H^+ to trace dense cores and filaments.

In this paper, we describe an automated, generalizable method that fits two-component NH_3 (1, 1) spectra efficiently using the LM method, without the need for user-provided initial guesses. The fitted models are subsequently used to identify filament spines in ppv space, which are sorted into velocity-coherent filaments accordingly. Moreover, we present a novel approach to study velocity gradients in filaments on beam-resolved scales, where velocity gradients are decomposed into components that are parallel and perpendicular with respect to local filament spines. Such a technique enables us to explore filament kinematics and accretion flow directions on the dense core (<0.1 pc) scale in addition to the filament scale (>0.5 pc).

We apply our new methods to filaments seen in NH_3 (1, 1) data of the NGC 1333 region. Located at a distance of about 295 pc away (Ortiz-León et al. 2018; Zucker et al. 2018), the NGC 1333 star-forming clump in the Perseus molecular cloud is one of the nearest cluster-forming regions. Its properties make NGC 1333 an ideal place to study the interplay between filaments and cores in a cluster-forming environment in detail. NGC 1333 is also one of the most extensively studied star-forming clumps (see Walawender et al. 2008), providing a wealth of context within which our study can be placed.

This paper is laid out as follows. We describe our NH_3 (1, 1) model, synthetic data, and observed data of NGC 1333 in Section 2. Methods behind our analysis, as well as test results of our line-fitting method, are presented in Section 3. The results of our analysis on the NGC 1333 observations are presented in Section 4, followed by a discussion of these results in Section 5. A concluding summary is in Section 6.

2. Models and Data

We used two data sets for our work presented here: one synthetic and one observational. The spectral model behind our line fits is described in Section 2.1 while the synthetic data used to test the accuracy of our line-fitting method is described in Section 2.2. The observations we used for this work are obtained from the Green Bank Ammonia Survey (GAS; Friesen et al. 2017) and are presented in Section 2.3.

2.1. NH₃ Line Models

We modeled observations of the NH_3 (1, 1) inversion transition along a given line of sight with up to two homogeneous bodies of beam-filling gas known as slabs. Each slab in our model corresponds to a kinematic (i.e., velocity) component observed in a spectrum and is assumed to be at a local thermal equilibrium with itself. Furthermore, we assume the emission can be parameterized by excitation temperature (T_{ex}), optical depth (τ_0), velocity dispersion (σ_v), and velocity centroid in the local standard of rest frame (v_{LSR}). The optical depth profile of each slab is described by

$$\tau(v, \sigma_v) = \sum_{i=1}^{n} W_i \tau_0 \exp\left[\frac{-(v - v_{LSR} - \delta v_i)^2}{2\sigma_v^2}\right], \qquad (1)$$

where W_i and δv_i are the relative weight and velocity offset of each of the eighteen NH_3 (1, 1) hyperfine components, respectively. These weights and velocity offsets are tabulated by Mangum & Shirley (2015). We further note that the τ_0 here corresponds to the combined optical depth of all the hyperfine components.

The radiative transfer of our model emission through each slab is governed by:

$$I_\nu = B_\nu(T_{ex})(1 - e^{-\tau_\nu}) + I_\nu^{bg} e^{-\tau_\nu}, \qquad (2)$$

where I_ν^{bg} is the specific intensity of the background radiation and $B_\nu(T)$ is the Planck function at a temperature T. Each slab is assumed to have a constant T_{ex} and we adopt the cosmic microwave background (CMB), as the I_ν^{bg} for our first slab, i.e., the slab farthest from the observer. We then subsequently use the emergent I_ν of the first slab as the I_ν^{bg} of our second slab to complete the calculation. To mimic baseline-removal used in our data reduction, a constant value of $B_\nu(T_{CMB} = 2.73 \text{ K})$ is subtracted from Equation (2) in our final I_ν model.

While we assume the slab farthest from an observer to be the optically thicker slab in our initial guesses, our least-squares fitting routine ultimately decides the order of the modeled slabs along the line of sight. We note that such an ordering of the slabs is unimportant when the two slabs are optically thin with respect to each other, or when two optically thick slabs are not spectrally overlapped due to large v_{LSR} separations. Since the satellite hyperfine lines of NH_3 (1, 1) emissions are optically thin in most cases, they help to constrain our kinematic fits well

even when the main hyperfine lines are optically thick. Such a constraint allows our fitting method to distinguish double spectral peaks resulting from a line absorption profile from those resulting from superpositions. We implicitly explore how the ordering of the slabs affects our fitting accuracy with the performance test described below in Section 2.2.

Since this paper focuses exclusively on understanding the gas kinematics, we analyzed only the NH_3 (1, 1) lines to maximize our spatial coverage for the study. As the (2, 2) line is expected to be weaker than the (1, 1) line, this strategy enabled us to extend our analysis over wider regions where the NH_3 (2, 2) line is not detected. While such an approach does not allow us to derive T_{ex} and τ_0 accurately for purely optically thin slabs, where the two parameters are spectrally degenerate, it does remove the potential bias that comes with assuming a single T_{ex} for both the (1, 1) and (2, 2) lines.

2.2. Synthetic Spectra

To test the accuracy of our line-fitting method, we generated a set of 25,000 synthetic NH_3 (1, 1) spectra, each from either a one- or two-slab model that fills the beam. To motivate our test set physically, we constructed each of our synthetic slabs with the NH_3 model adopted by the GAS DR1 first results paper (Friesen et al. 2017) instead of using our fitting model, described earlier in Section 2.1. The model adopted by the DR1 paper is based on the work by Rosolowsky et al. (2008) and Friesen et al. (2009), and developed from the framework described in Mangum & Shirley (2015).

Each gas slab in our testing model is physically parameterized by the para-ammonia column density, N, and the kinetic temperature, T_k, in addition to σ_v, v_{LSR}, and T_{ex}. We further assumed these slabs are in local thermal equilibrium (LTE), i.e., $T_{ex} = T_k$, and drew parameters behind each instance of synthetic slab randomly from the predefined distributions specified below:

1. N is drawn from a log-uniform distribution in the range of $13 \leqslant \log(N/cm^{-2}) < 14.5$;
2. T_k is drawn from a uniform distribution in the range of $8\ K \leqslant T_k < 25\ K$;
3. σ_v is the quadrature sum of the thermal line width, $\sigma_{v,T} = 0.08\ km\ s^{-1}$, and the nonthermal line width, $\sigma_{v,NT}$, where $\ln[\sigma_{v,NT}/(km\ s^{-1})]$ is drawn from a normal distribution with a mean and a standard deviation of -2.3 and 1.5, respectively. This particular line width distribution is chosen to resemble those found in the GAS DR1 first results, and;
4. v_{LSR} of the first gas slab is drawn from a uniform distribution in the range of $-2.5\ km\ s^{-1} \leqslant v_{LSR} < 2.5\ km\ s^{-1}$ while the v_{LSR} offset of the second slab from the first is drawn independently from the same distribution.

We chose these distributions to represent broadly the typical physical conditions seen toward nearby molecular clouds with NH_3.

We constructed a two-slab spectrum using the same radiative transfer formalism described by Equation (2), with the CMB subtracted as a constant. The final synthetic spectrum is produced by adding random Gaussian noise with an rms value of 0.1 K to the model spectrum. The value of 0.1 K is chosen to mimic the typical noise level found in the GAS DR1 observations (Friesen et al. 2017).

For each instance of randomly generated spectrum, eight additional, spatially correlated spectra are produced. Collectively, these spectra are placed in a 3×3 spatial grid with the original spectrum positioned in the center. The purpose behind creating such a cube is to provide spatial information for fitting methods that utilize them, including the method presented in this paper.

The spatial correlation between pixels of a synthetic cube is achieved by applying spatial gradients to the model parameters, referenced at the central pixel. The gradients used for each parameter are randomly drawn from a Gaussian distribution with 1σ values of 0.2 K, 0.1 km s^{-1}, 0.1 km s^{-1}, and 0.01 for T_k, σ_v, v_{LSR}, and $\log(N)$, respectively, per pixel. Figure 1 (left) shows spectra extracted from such a synthetic cube, displayed spatially on a 3×3 grid.

2.3. NH₃ Observations, Reduction, and Imaging

We observed the Gould Belt molecular clouds with NH_3 (1, 1) and (2, 2) inversion lines as a part of the GAS survey (Friesen et al. 2017). The GAS observations were made with the Robert C. Byrd Green Bank Telescope (GBT) using its 7-beam K-Band Focal Plane Array (KFPA) and its VErsatile GBT Astronomical Spectrometer (VEGAS) backend. The angular and spectral resolutions (FWHMs) of our NH_3 data are $32''$ and $\sim 0.07\ km\ s^{-1}$ (i.e., 5.7 kHz at 23.7 GHz), respectively.

Our targets are observed using the On-The-Fly (OTF) technique, where a $10' \times 10'$ on-sky tile was scanned with a Nyquist-sampled spacing between each row. We reduced these observations with the GBT KFPA data reduction pipeline (Masters et al. 2011) and imaged them with the recipe described by Mangum et al. (2007). Further details of the observations, data reduction, and imaging are available in the GAS first results paper (Friesen et al. 2017). For this work, we will focus on the observations of NGC 1333 star-forming clump, available to the public via Data Release 1 (DR1) of the first results paper. Figure 1 shows example spectra on its right panel extracted from the DR1 data over a 3×3 pixel region in NGC 1333.

3. Analysis Methods

Here, we present our analysis methods in this section below. Details of our spectral fitting method are provided in Section 3.1. We conducted a performance test on our fitting method to quantify its accuracy and completeness, and we present the details and the results of this test in Section 3.2. Methods for identifying velocity-coherent filaments and assigning component memberships to these filaments based on the fits are presented in Section 3.3. Methods behind velocity gradient analysis are presented in Section 3.4.

3.1. Spectral Fitting

We fitted our synthetic and real data automatically using the LM (Levenberg 1944; Marquardt 1963; Moré 1978) least-squares minimization method. We bypassed the need for user-provided initial guesses using an automated approach described in Section 3.1.1 and performed least-squares fits using the PYSPECKIT package (Ginsburg & Mirocha 2011). The Python implementation of the LM method used by the PYSPECKIT is based on the FORTRAN version found in the MINPACK-1 package (Moré et al. 1980), made available via a
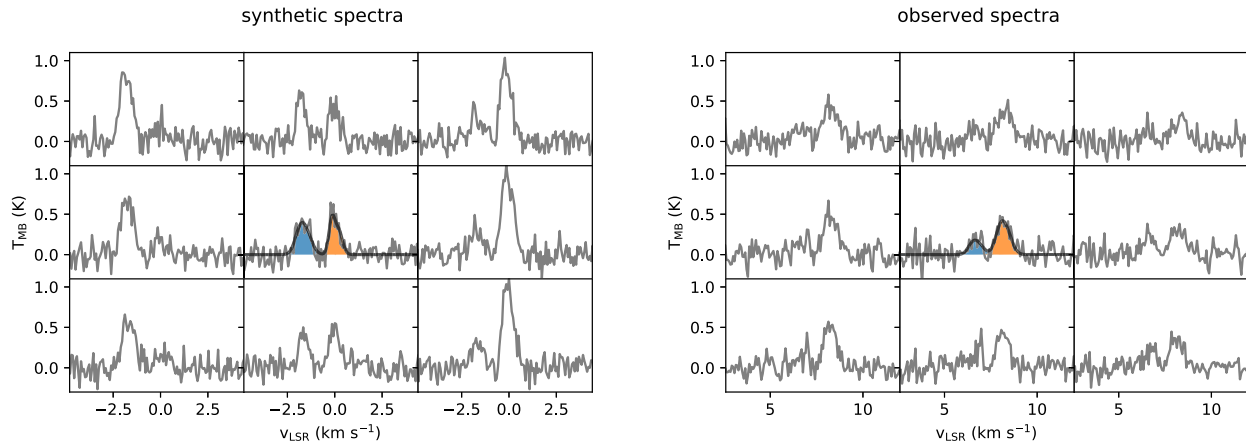
Figure 1. NH_3 (1, 1) spectra (gray) taken from a synthetic spectral cube (left) and the NGC 1333 observation (right) over a grid of 3×3 footprint, zoomed in to focus on the main hyperfine structures. Models fitted to all 18 hyperfine components of the spectra in the center pixel are shown in black, and the individual components of the model are shown in blue and orange.

series of translations (Rivers 2002; Markwardt 2009). Our fits are performed on a pixel-by-pixel basis for all the pixels in our data, excluding noisy regions near the map edges. We use a statistical method described in Section 3.1.2 to discern whether a pixel is better modeled by noise, a one-component model, or a two-component model. Our fitting package is publicly available via *GitHub* under a GNU General Public License as the MUFASA[10] (i.e., MUlti-component Fitter for Astrophysical Spectral Applications). The version we used for this work is archived in Zenodo (Chen 2020a).

3.1.1. Making Initial Guesses

The LM method is an iterative approach to find local minima using a hybrid algorithm of the gradient-ascent and the Gauss-Newton methods (see Lourakis 2005). Due to the nature of these methods, good initial guesses are typically needed with the LM method to find the global minima. Such guesses are particularly required for complex spectral models with many local chi-squared minima and are the reasons why many earlier efforts to fit multiple slab models (e.g., Hacar et al. 2013; Henshaw et al. 2016) require human intervention and are not fully automated. For our method, i.e., MUFASA, we fit spectra automatically using initial guesses generated from a recipe described below. A statistical technique (see Section 3.1.2) is used thereafter to decide which of the one- or two-slab models fitted for each pixel is more appropriate, without human intervention.

The GAS DR1's one-slab fitting method used an effective and automated recipe to make initial guesses. The recipe used the first and second moments of the central, i.e., main, NH_3 (1, 1) hyperfine lines as its initial v_{LSR} and σ_v guesses. Such a calculation excludes the satellite hyperfine lines because their large velocity offsets are not a kinematic feature. The main hyperfine lines were isolated for such a calculation in DR1 via a user-defined spectral window.

For our one-slab model fits, we adopt the DR1 recipe for our initial guesses of v_{LSR} and σ_v, but define our spectral windows automatically instead. We automate such a process by centering a 6 km s^{-1} full-width window on the emission peak of the spatially integrated spectrum of our data. Since the NH_3 (1, 1) emission should be optically thin throughout the majority of the

data, such an emission peak should locate the whereabouts of the main hyperfine lines. Once the window is defined, we follow the DR1 recipe and calculate the zeroth, first, and second moments (μ_0, μ_1, μ_2) of the main hyperfine lines over the window.

To obtain initial guesses for T_{ex} and τ_0 better than assuming fixed values, we use the zeroth moment (μ_0) map as a proxy instead. Specifically, we calculate our guesses by first normalizing the 99.7 percentile value of the μ_0 distribution across the map to one. The initial guesses for T_{ex} and τ_0 are then obtained from the normalized μ_0, i.e., $\widetilde{\mu}_0$, as $\widetilde{\mu}_0 \cdot T_{gmx}$ and $\widetilde{\mu}_0 \cdot \tau_{gmx}$, respectively, where $T_{gmx} = 8$ K and $\tau_{gmx} = 1$. The resulting fits from adopting such guesses do not depend sensitively on T_{gmx} and τ_{gmx} around these chosen values.

We expand the DR1 fitting recipe further for our two-slab fits via the following steps:

1. Adopt the first and second v_{LSR} guesses as $\mu_1 \pm 0.4\mu_2$, respectively.
2. Adopt the first and second σ_v guesses both as $0.5\mu_2$, respectively.
3. Adopt the first and second T_{ex} guesses as $\widetilde{\mu}_0 \, T_{gmx}$ and $0.8 \, \widetilde{\mu}_0 \, T_{gmx}$, respectively.
4. Adopt the first and second τ_0 guesses as $0.75 \, \widetilde{\mu}_0 \, \tau_{gmx}$ and $0.25 \, \widetilde{\mu}_0 \, \tau_{gmx}$, respectively.

As with the guesses used for one-slab fits, our choices for T_{gmx} and τ_{gmx}, and their respective scaling factors, do not affect the fitting outcome sensitively. Our choices for these values are motivated by a hypothetical scenario where the two gas slabs emitting a spectrum have comparable velocity dispersions, densities, and kinetic temperatures. We note that our initial guesses assume the slab further from the observer has a higher optical depth. As we show below in Section 3.2, our recipe for making guesses for the two-slab fit is robust even when the gas slabs emitting the spectrum have dissimilar velocity dispersions, i.e., contrary to this assumption.

To take advantage of spatial information present in our observations, we first fit data that are spatially convolved to an angular resolution twice the size of the original resolution. The parameter maps derived from this initial fit are then median smoothed and interpolated. For T_{ex} and τ_0 guesses, values outside of ranges 3–8 K and 0.2–8, respectively, are removed prior to median-smoothing and interpolation. These post-

---

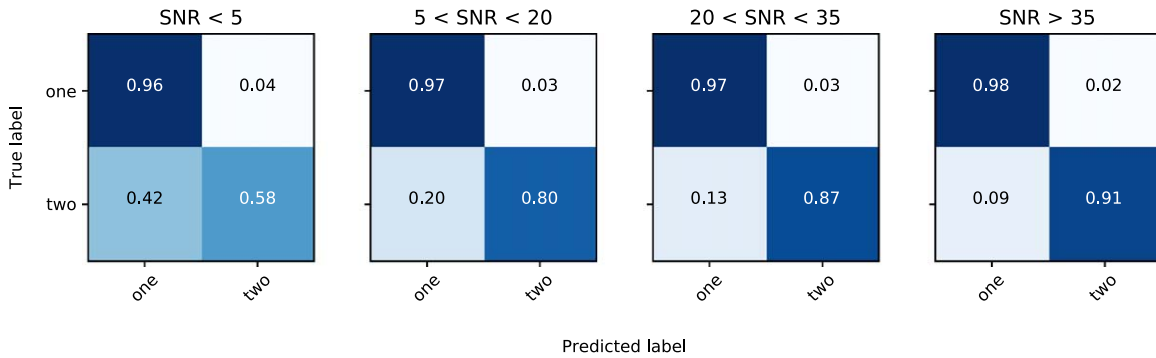[10] MUFASA codebase: https://github.com/mcyc/mufasa.

Figure 2. Confusion matrices quantifying MUFASA's ability to classify the number of velocity slabs behind a synthetic spectrum. Each matrix contains a subset of our samples at a given range of S/N values.

processed parameter maps are then adopted as the initial guesses of our fits to the full-resolution cubes.

By using parameters fitted to the spatially convolved cube as initial guesses for the full-resolution fit, we are able to take advantage of the enhanced signal-to-noise ratio (S/N) in addition to the spatial information present in the convolved cube to improve our fits. Figure 1 shows spectra extracted over a 3×3 pixel region from a synthetic data cube (left) and the DR1 data of NGC 1333 (right), demonstrating the spatial correlation of these spectra between pixels. The two-slab model fitted to the central pixel using this method is overlaid over the central spectrum.

Since moment estimates for making initial guesses may overlook a faint spectral component in the presence of a much brighter one, we further fit one-slab models to our fit residuals in an attempt to recover missing components. The spectral window used to estimate initial guesses for this fit is centered on the $v_{\rm LSR}$ derived from the initial one-slab fit, with a full window width of $7\,{\rm km\,s^{-1}}$.

We use the results of the one-slab residual fit subsequently to assist with the recovery of a missing component. These results are taken in tandem with those of the original one-slab fit as initial guesses for the re-attempt at fitting a two-slab model. We perform these re-attempts over pixels where one-slab models initially fit the full spectrum better than the two-slab models, as determined by the selection criterion described in the next section (i.e., Section 3.1.2). The same criteria is used further to determine whether or not this new two-slab fit is justified over the one-slab fit.

3.1.2. Model Selection

Many previous multicomponent fitting methods selected their best-fit models for each component via an S/N threshold (e.g., Sokolov et al. 2017), a velocity separation threshold, or both (e.g., Hacar et al. 2013, 2017; Chen et al. 2019). While these criteria are effective at avoiding overfitting, they are not necessarily good at picking up all the spectral components present along a given line of sight. We addressed this issue by using the corrected Akaike information criterion (AICc; Akaike 1974; Sugiura 1978) to select the best model between the one- and two-slab fits, on a pixel-by-pixel basis.

The AICc is a second-order corrected estimator, based on the K-L information loss (Kullback & Leibler 1951), for the relative likelihood of one model with respect to another at representing a data set with N samples. Assuming the errors in the data are normally distributed, the AICc can be written in

terms of the χ^2 of the fit as:

$$\mathrm{AICc} = \chi^2 + 2p + \frac{2p(p+1)}{N-p-1}, \qquad (3)$$

where p is the number of parameters used in the model. At each pixel, we accept a two-slab model as the better fit over its one-slab counterpart when their relative likelihood, K_1^2, given by

$$\ln K_1^2 = \exp(\mathrm{AICc}_2 - \mathrm{AICc}_1) \qquad (4)$$

is greater than 5 (Burnham & Anderson 2004). Similarly, we accept a one-slab model as the better fit over a $I_\nu = 0$ model of noise with no free parameter when their relative likelihood, $\ln K_0^1$, is greater than 5.

Pixels with reduced χ^2 values of $\chi_\nu^2 > 1.5$ are further masked from our analysis to ensure spectra which are inadequately modeled by our fits, e.g., those with possibly three or more velocity slabs, are not included in the analyses. In NGC 1333, no pixel is masked out as all pixels best fitted by 2-slab models have $\chi_\nu^2 < 1.5$.

3.2. Performance Tests on Line Fitting

Figure 2 summarizes the performance of our line-fitting method, MUFASA, at identifying the correct number of velocity slabs behind a synthetic spectrum, using confusion matrices. These results are obtained from fits to all 25,000 synthetic test spectra described in Section 2.2 and are binned into separate matrices according to the S/N of their true underlying spectra. Unless stated otherwise, we refer to S/N as the modeled peak brightness to noise ratio of a final, radiatively transferred spectrum rather than that of its individual components. As illustrated in Figure 2, our fitting method identifies true one-slab spectra robustly. The true-positive rate of this identification is greater than 96% even for our lowest S/N ($<$5) bin. For two-slab identification, the true-positive rate correlates with the S/N value of the spectra, reaching values upwards of about 90% at higher S/N. Even at moderate S/Ns between 5 and 20, the true-positive rate is roughly 80% for two-slab spectra.

We next explore the impact of velocity separation between the slabs and intrinsic velocity dispersion (i.e., σ_v) on the success of our fitting method. Figure 3 shows the rate of true-positive identification of a two-slab spectrum as a function of the $v_{\rm LSR}$ separation between the slabs, i.e., $\Delta v_{\rm LSR}$. The left panel shows these rates binned according to the S/N value of the fainter slab, i.e., S/N$_{\rm min}$.

The response curves in Figure 3 show the same shape across various S/N$_{\rm min}$ regimes. Namely, they increase monotonically
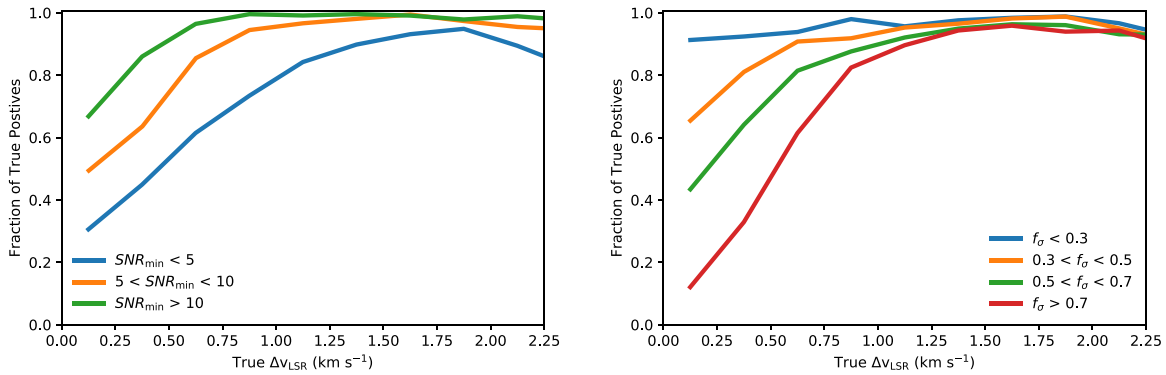
Figure 3. The rate of true-positive identification of a two-slab spectrum plotted as a function of the v_{LSR} separation between the slabs. The left panel shows these test samples binned by the S/N values of the fainter slab while the right panel shows these values binned by velocity dispersion ratio f_σ of the narrow slab with respect to the wide slab. The median σ_v values for the narrow and wide components in these samples are 0.23 km s^{-1} and 0.48 km s^{-1}, respectively.

with Δv_{LSR} until the fraction of true-positive identification plateaus at 100%. Prior to plateauing, these curves shift vertically upwards as the S/N$_{min}$ increases and behave qualitatively the same even when they are binned instead according to the S/N of the brighter slab, i.e., S/N$_{max}$, or the S/N taken from the peak of the combined spectrum (i.e., the S/N).

The only exception to this plateauing trend is when S/N$_{min}$ is low, i.e., at a value less than 5. At this low S/N regime, the true-positive rate reaches a maximum of 95% at a value of about 1.9 km s^{-1} before it turns around instead of plateauing. This turnaround in the true-positive rate likely reflects the limit of our ability to make initial guesses, suggesting that moment maps may have trouble picking up fainter components when these components reside near the edges of the spectral window from which moment maps are calculated from.

The performance of MUFASA decreases with decreasing Δv_{LSR}, likely resulting from a lack of velocity acuity between two slabs when they have similar velocities. As Δv_{LSR} decreases, the spectral profiles of the two slabs start to blend together, making them more difficult to distinguish from a one-slab profile. This lack of acuity is what prompted many studies to adopt a Δv_{LSR} threshold for their model selection to guard against overfitting (e.g., Hacar et al. 2013), and remains a challenge even for advanced machine-learning techniques (e.g., Keown et al. 2019).

The right panel of Figure 3 shows the true-positive identification rates divided into various regimes of velocity dispersion ratio, i.e., the line width of the narrower slab relative to its wider counterpart along a line of sight ($f_\sigma = \sigma_{v,narrow}/\sigma_{v,wide}$). Here, the true-positive rate anticorrelates strongly with f_σ. This true-positive rate is likely enhanced when the spectral profiles from the slabs are less similar, i.e., when f_σ is low, which makes them easier to discern from one another when their amplitudes are different. Indeed, this rate can be as high as 90% for spectra with $f_\sigma < 0.3$ and correlates weakly with Δv_{LSR} in this regime.

Such enhanced identification rates in the low f_σ regime make MUFASA particularly useful for disentangling subsonic gas from supersonic gas along lines of sight. For reference, the median σ_v values for the narrow and wide components in our test samples are 0.23 km s^{-1} and 0.48 km s^{-1}, respectively, whereas the isothermal sound speed at 10 K is $\sim$0.2 km s^{-1}. Like the trend seen in the left panel, the true-positive rate correlates with Δv_{LSR} for most f_σ values prior to plateauing.

The slope of these correlations, however, becomes shallower as f_σ decreases.

Figure 3 reveals that MUFASA is able to recover a large fraction of two-slab spectra that would otherwise be missed by using a Δv_{LSR} threshold for model selection. For a spectral population described by our synthetic spectra, at least 20% and 30% of the second slabs missed by a threshold of 0.25 km s^{-1} (e.g., Hacar et al. 2017) and 0.4 km s^{-1} (e.g., Chen et al. 2019), respectively, are recovered with MUFASA. These recovery rates can be significantly higher, however, depending on the S/N and f_σ.

Figure 2 reveals that MUFASA rarely overfits one-slab spectra, i.e., misidentifying a one-slab spectrum as having two slabs. Our test shows MUFASA only misidentifies one-slab spectra <3% of the time. Even in the lowest S/N regime, such misidentification only occurs 4% of the time. Performing model selection via AICc alone is thus sufficient to guard against overfitting without needing an additional threshold criterion.

To quantify how well MUFASA captures the true v_{LSR} and σ_v, Figure 4 shows the true 1σ errors of these two parameter fits as a function of the following values: the S/N of a two-slab spectrum, the S/N of an individual slab in a two-slab spectrum, and the S/N of a one-slab spectrum. These S/N-error relations are plotted with errors calculated from the median absolute deviation (MAD) of the fitted parameters with respect to the true value behind a synthetic spectrum. As expected, the parameter errors are anticorrelated with the S/N. We note that these errors also account for any potential cases where the spatial order of velocity slabs may not be ordered correctly along a line of sight.

3.3. Identifying Velocity-coherent Filaments

Velocity gradient analyses are only useful for understanding gas flows if the structures being analyzed are velocity coherent, i.e., continuous in velocity. Velocity slabs derived from our fits must therefore be sorted into velocity-coherent structures prior to such analyses. In this subsection, we present methods to reconstruct fits to our data as simple emission line models in position–position–velocity (i.e., ppv) space without hyperfine structures. We further present methods to identify filament spines from these models in ppv space and sort the fitted slabs into velocity-coherent filaments based on these identified spines.
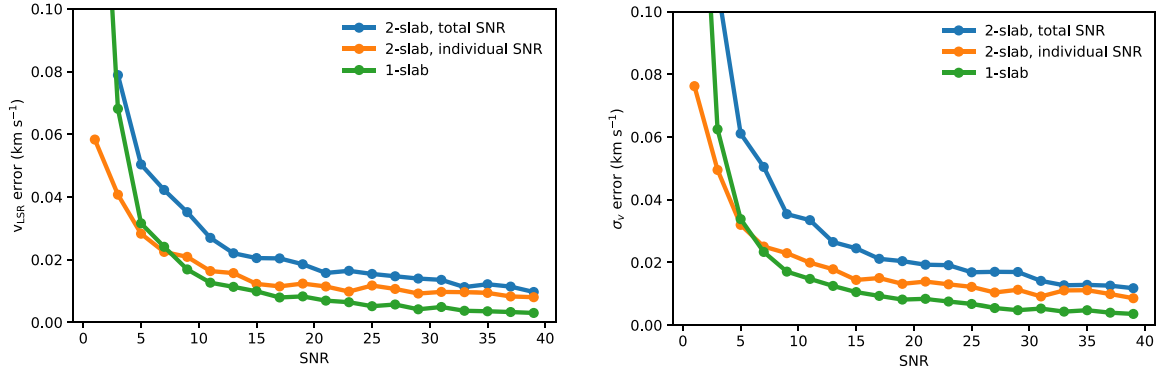
Figure 4. The 1σ error of the fitted v_{LSR} (left) and σ_v (right) as a function of (blue) the total S/N of a 2-slab spectrum, (orange) the S/N of a given slab in a spectrum, and (green) the S/N of a one-slab spectrum, for fits with a correct number of identified components.

3.3.1. Reconstructing Velocity Structures

To help identify velocity-coherent structures in ppv space, we first reconstruct our best-fit models with the hyperfine structures removed to avoid confusion from these nonkinematic features. Such a reconstruction, known as "deblending," is accomplished by computing the spectral profile of each gas slab using a single Gaussian τ_ν profile based on our best-fit model, which accounts for all 18 hyperfine structures (see Section 2.1). In other words, we reconstruct the spectral profile along a line of sight with either one or two Gaussian τ_ν components using our best-fit parameters and number of components as determined by the AICc criterion.

Since our τ_0 value derived from a fit (see Equation (1)) represents the peak optical depth of all the NH_3 (1, 1) hyperfine lines combined, we scale down our fitted τ_0 by a factor of 10 to represent better the actual optical depths of individual hyperfine groups in our reconstruction. For reference, the main and satellite hyperfine groups each contain about 50% and 10% of the optical depth represented by τ_0, respectively. The observed satellite lines are thus typically optically thin even when the main hyperfine lines are not, which enables these thin lines to reveal unobstructed structures along lines of sight.

We further assume each deblended velocity component is optically thin with respect to each other but not with respect to the CMB, individually. Our deblended model, with the CMB subtracted as a constant baseline, can thus be expressed as,

$$I_\nu = \sum_{j=i}^{m} [B_\nu(T_{ex,j}) - B_\nu(T_{CMB})][1 - e^{-\tau_{\nu,j}}], \quad (5)$$

where j designates each velocity component along a line of sight and $\tau_{\nu,j}$ is governed by

$$\tau_{\nu,j} = 0.1 \; \tau_{0,j} \exp\left[\frac{-(v - v_{LSR,j})^2}{2\sigma_{v,j}^2}\right]. \quad (6)$$

Here, $T_{ex,j}$, $\tau_{0,j}$, $v_{LSR,j}$, and $\sigma_{v,j}$ are obtained from previously fitted models with hyperfine structures. To ensure the deblended emission have high spectral acuities for structure identification in ppv space, we further set $\sigma_{v,j}$ to $0.09 \, km \, s^{-1}$ instead of adopting the line widths previously derived from our fits. This constant σ_v value is roughly the minimal line width required to be Nyquist-sampled at our $0.07 \, km \, s^{-1}$ full-width-half-max (FWHM) spectral resolution.

To illustrate structures revealed by deblending, Figure 5 shows the volume-rendered deblended ppv cube of our fits to the NH_3 (1, 1) observations of NGC 1333.

3.3.2. Identifying Filaments

We identify filaments in ppv space by using a multi-dimensional density ridge identification algorithm known as the Subspace Constrained Mean Shift (SCMS); (Ozertem & Erdogmus 2011). The mathematical framework behind SCMS was generalized by Chen et al. (2014) to operate on weighted, particle-like data in addition to their unweighted counterparts, which enables SCMS to run on gridded, multidimensional images. Since the publicly available R code developed by Chen et al. (2015) for cosmological applications only implemented the original framework, we modified the code to reflect the generalized one. We further translated the code to Python, parallelized it for multiprocessing, and made it publicly available on *GitHub* via the CRISPY[11] (i.e., Computational Ridge Identification with SCMS for Python) package under a GNU General Public License. The version of CRISPY we used in this paper is archived in Zenodo (Chen 2020b).

While the DisPerSE algorithm (Sousbie 2011; Sousbie et al. 2011) recently used in star formation studies (e.g., Arzouma-nian et al. 2011) also operates in 3D (e.g., Smith et al. 2016), it requires two more user-defined parameters to run than SCMS. The SCMS' ability to find density ridges consistently is also well established in statistical studies (Chen et al. 2014), making SCMS an attractive option over other methods. Furthermore, SCMS captures local information, such as ridge orientations, better than methods that derive ridges from monolithic filament objects (e.g., Koch & Rosolowsky 2015) or approximate them as line segments (Hacar et al. 2013, 2017).

The SCMS algorithm defines a ridge as a smooth, continuous, one-dimensional object in a multidimensional density field. In addition to this nomenclature, we define a skeleton as a ridge that has been gridded onto an image and a spine as a skeleton with all its branches removed. The SCMS algorithm finds ridges by moving walkers iteratively up the density field using a gradient-ascent method. This approach is subspace constrained (see Chen et al. 2014) to ensure the walkers converge on one-dimensional ridges instead of zero-dimensional peaks. Figure 6 demonstrates how SCMS identifies such a ridge in 2D from an NH_3 integrated intensity map of a source in NGC 1333 using the CRISPY package.

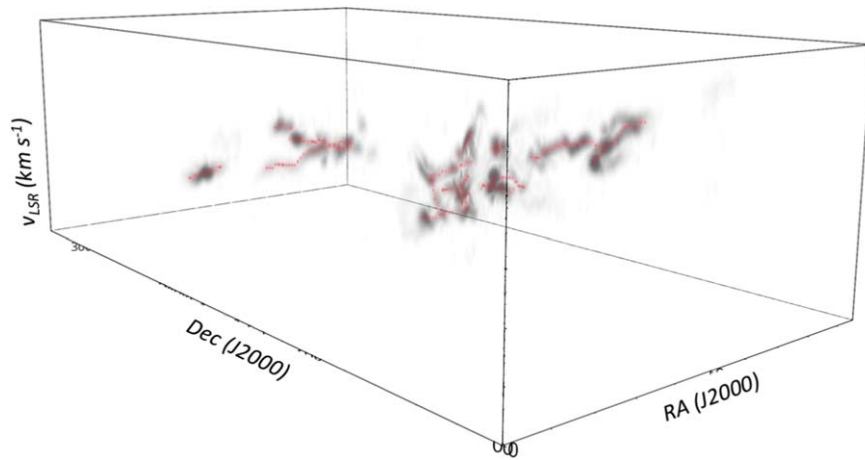[11] CRISPy codebase: https://github.com/mcyc/crispy.

Figure 5. The debelended ppv cube of NGC 1333 reconstructed from the best-fit models (gray), and the filament spines identified from the cube using CRISPY (red).
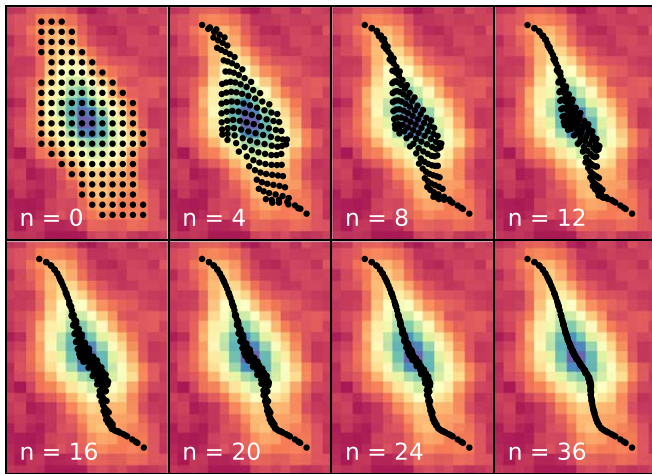

Figure 6. Snapshots of SCMS finding a 2D "density" ridge from the image shown in the background, as carried out by CRISPY. The respective iteration number for each snapshot is labeled in each panel. The black dots represent the SCMS walkers and the background image was taken from the integrated NH_3 emission map of NGC 1333, cropped around a source in its northeast.

Figure 7. A gridded, 3D skeleton being pruned into a spine, projected onto the xy and xv planes on the left and right panels, respectively. Top: the skeleton prior to pruning. Middle: the skeleton decomposed into branches (black), intersections (gray), and "bad" branches (red; see Section 3.3.2). Bottom: the resulting spine, defined by the longest path in the skeleton with "bad" branches excluded.

In general, the SCMS algorithm operates primarily on two user-defined parameters: density (e.g., intensity) threshold and smoothing bandwidth. The density threshold masks out noisy features in the density field while the smoothing bandwidth performs a kernel estimate of the field from particle-like data. Even though a gridded image, e.g., an emission cube, can in principle serve directly as a density field without a kernel estimate, a smoothing kernel is still required by the generalized SCMS to estimate density gradients efficiently and move its walkers accordingly (Chen et al. 2014). A smoothing length comparable to, or greater than, the resolution of the image is thus required still.

For this work, we adopted a density threshold of 0.15 K and a smoothing length of 1.5 pixels for our SCMS run with CRISPY. The deblended cube was spatially convolved to twice its original beamwidth prior to the run. This additional step was performed to suppress noisy features further in the cube without sacrificing spectral resolution, particularly given that SCMS smooths its data indiscriminately in all three dimensions. Since SCMS operates natively in a continuous space, we set our convergence criterion such that the ridges identified are
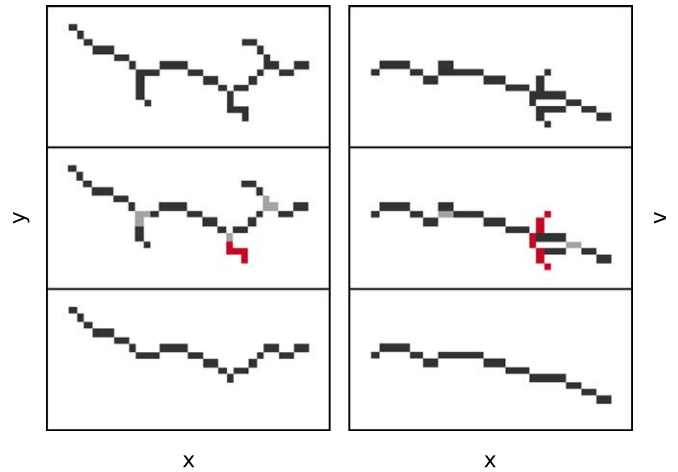
less than one voxel in width prior to regridding. We describe details on our choice of parameters further in Appendix A.

Once CRISPY identified emission ridges in the continuous ppv space, we map these ridges back to the native grid of the deblended cube. These regridded ridges are referred to as skeletons and are subsequently pruned down to branchless structures we call spines. We accomplish such a pruning by using a graph-based technique developed by Koch & Rosolowsky (2015), which we have generalized to operate in 3D.

We prune branches by first decomposing a skeleton into intersection and branch objects known as nodes and edges in graph theory, respectively. We then find the longest path in the graph, measured in Euclidean distance, and subsequently remove all the edges outside of this path. To ensure our spines represent velocity-coherent structures, branches that may bridge velocity discontinuities are further removed. We define these "bad" branches as ones with an on-sky length less than 9 pixels and a velocity projected length greater than that of its on-sky length in pixels (i.e., $\sim$4.8 km s^{-1} pc^{-1} for NGC 1333 with our data). Figure 7 shows a demonstration of our pruning process with "bad" branches shown in red. We note that
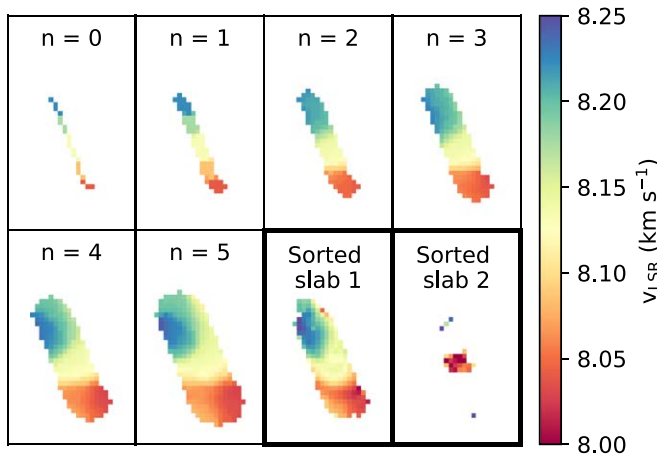
Figure 8. A ppv-footprint during each iteration of its creation, shown in panels labeled with the iteration number n. The v_{LSR} maps of the sorted velocity slabs, which have been assigned to the final ppv-footprint ($n = 5$) as members of an association based on their v_{LSR} proximities, is shown in the last two panels, framed with bold borders. Only the first velocity slab represents a velocity-coherent structure we used for our velocity gradient analyses.

Figure 9. Distance transform of a filament spine overlaid with its corresponding (a) divergence field and (b) parallel field. The vectors on and off the spine are colored white and black, respectively. The color scale corresponds to pixel distances between 0 and 5.

removing "bad" branches does not necessarily impose a maximum velocity gradient limit on a filament. This process merely breaks filaments apart.

3.3.3. Assigning Membership to Filaments

We group our fits-derived velocity slabs obtained with MUFASA into velocity-coherent structures by associating them with filament spines. In brief, we do so hierarchically by first placing velocity slabs into structures we call associations based on each slab's proximity to a spine in the ppv space. Such proximity is calculated using a spatial extension of a spine we call ppv-footprint, a structure from which the v_{LSR} separation between a slab and a spine can be referenced at each pixel. This first step intends to disentangle filaments that overlap in projection into associations.

Associations, which are allowed to have more than one velocity slab at each pixel, are then sorted internally to produce velocity-coherent structures (vc-structures) that contain only a single slab at a pixel. This sortation is carried out based on kinematic similarities between the velocity slabs. The v_{LSR} map resulting from this sortation is subsequently median smoothed and adopted as the new ppv-footprint. This last step serves to grow and update the association iteratively starting from the filament spine, one pixel at a time.

We repeat such an assignation and sortation for five iterations. This number of iterations grows our vc-structures to an extent where S/N values of the new pixels start to drop off below 3. Figure 8 shows ppv-footprints corresponding to each of these iterations in panels labeled with the iteration number n. The v_{LSR} maps of the first and second velocity slabs in the final association are shown in the last two panels of Figure 8. This first velocity slab shown in the figure is representative of those that go into our final vc-structures, which are used in our velocity gradient analyses. Further details of our membership assignment to vc-structures are described in Appendix B.
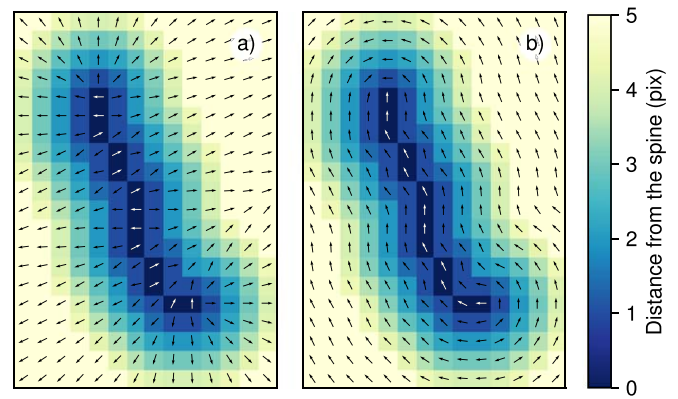
3.4. Velocity Gradient Analysis

3.4.1. Decomposition of Vector Fields

To study gas motions geometrically with respect to filament spines, we devised a technique to decompose a vector field (e.g., the velocity gradient field) into orthogonal components that are either parallel or perpendicular to a filament spine. Such a decomposition is accomplished by first taking a distance transform of a sky-projected spine to map out the shortest Euclidean distance between a given pixel and the spine. In other words, we calculated the radial distance between a pixel and a spine from which radial profiles of filaments can be constructed.

Vector fields that point orthogonally away from the filament spines are then created by taking gradients of our distance transforms over a sampling distance of 1 pixel using the second-order accurate central differences method. We refer to these fields as the divergence fields. Figure 9(a) shows an example of a divergence field superimposed on its corresponding distance transform.

Due to the sampling method of the gradient calculation and symmetry, the divergence field vectors right on the spine often have magnitudes of zero. To avoid a loss of information due to this limitation, we reconstructed an on-spine vector field parallel to the spine by taking gradients, i.e., central differences, of the spine's pixel coordinates. This on-spine field is then rotated by 90° and inserted into the divergence field, as shown in Figure 9(a) in white.

We construct the corresponding parallel field of a filament by rotating the divergence field by 90°. Since the divergence field is discontinuous across the spine, where its vectors point oppositely away from each other, a uniform rotation of the divergence field will result in a discontinuous parallel field, where its vectors are antialigned with respect to each other across the spine. To address this issue, we rotated the divergence field vectors on each side of the spine independently in the directions opposite to each other. Such a rotation is accomplished by exploiting angular degeneracy in the arctan function to differentiate vectors on the two sides of the spine. Figure 9(b) shows an example of a parallel field overlaid on the distance transform of its corresponding spine.
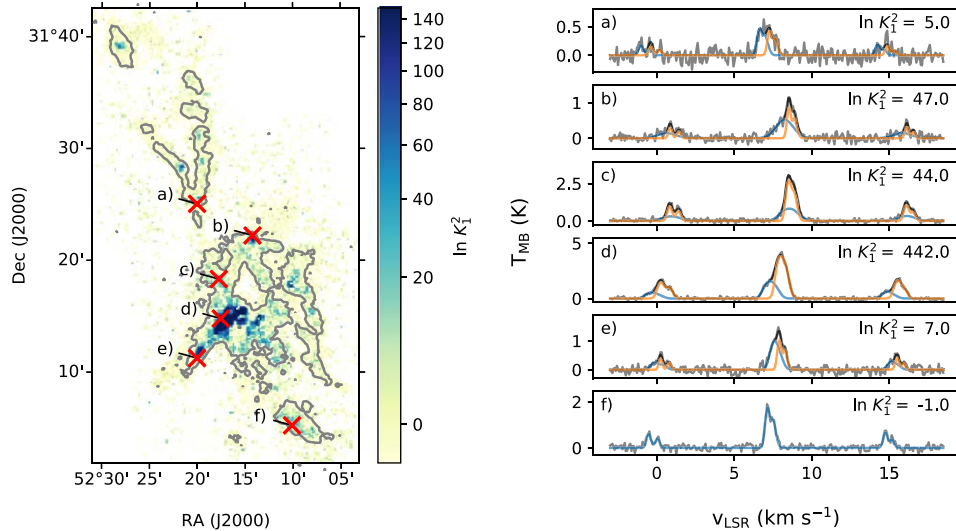
Figure 10. Left panel: relative likelihood of the two-slab fit over its one-slab counterpart in NGC 1333. The gray contour shows the integrated NH_3 (1, 1) emission at the 0.35 K km s^{-1} level. Right panel: the observed NH_3 (1, 1) spectra (gray) extracted from the positions marked with red x's shown in the left panel, cropped to focus on the spectral regions near the main hyperfine lines. The spectra are superimposed with their corresponding two-slab fits (black) and models of their individual components (blue and orange).

3.4.2. Computing Velocity Gradients

We calculate v_{LSR} gradients, i.e., ∇v_{LSR}, from v_{LSR} maps of velocity-coherent structures determined with the methods described in Section 3.3.3. These gradients are calculated on a pixel-by-pixel basis by fitting a plane over pixels within a 6 pixel diameter aperture centered on them. The diameter of the plane-fitting aperture is explicitly chosen to be twice the size of our FWHM beam to ensure the velocity gradients are calculated over resolved structures. To ensure the quality of our calculations, we calculate ∇v_{LSR} only over apertures where v_{LSR} values are available for more than one-third of the pixels.

We further decompose the calculated ∇v_{LSR} fields into components that are perpendicular and parallel with respect to their associated filament spine. This decomposition is accomplished by taking the dot products between the ∇v_{LSR} field and the divergence field, as well as between the ∇v_{LSR} field and the parallel field (see Section 3.4.1).

4. Results

Earlier in Section 3.2, we presented the performance of our fitting method, MUFASA, as characterized by our test fits to synthetic spectra. Here we present our best-fit models to the GAS NH_3 (1, 1) observations of NGC 1333 in Section 4.1, along with the deblended emission reconstructed from these fits. The filament spines identified from the deblended emission by CRISPY and the velocity slabs assigned to these spines are also presented in the same subsection. We further present our velocity gradient analysis on these velocity-coherent filaments in Section 4.2.

4.1. NGC 1333—Fitted Models

Figure 10 shows the relative likelihood of the two-slab fit over the one-slab fit, i.e., K_1^2, in NGC 1333 as determined by the AICc (see Section 2.1). A significant fraction, i.e., 40%, of the pixels in NGC 1333 with S/N > 3 are determined to be better fitted with two-slab models based on the statistically robust threshold of $\ln K_1^2 > 5$ (Burnham & Anderson 2004). This fraction is significantly higher than that suggested by the

GAS DR1 paper (Friesen et al. 2017), where when examined by eye, only 5% or less of the pixels with S/N > 3 appear to be inadequately fitted by a one-component model. No pixel best fitted with our two-slab model has $\chi_\nu > 1.5$, which indicates that our observations of NGC 1333 are indeed well modeled with two or fewer velocity slabs.

The right panel of Figure 10 shows examples of our best fits to the NH_3 (1, 1) emission, superimposed on spectra extracted from positions marked in the left panel. It is qualitatively apparent that these spectra are indeed better fitted by two-component models when $\ln K_1^2 > 5$, even in the limiting case where $\ln K_1^2$ is near 5.

The deblended ppv cube of NGC 1333 derived from our best-fit models is shown earlier in Figure 5, overlaid with their respective spines identified by CRISPY. The left panel of Figure 11 shows these spines projected onto the sky, and overlaid on top of the v_{LSR} maps of selected vc-structures, which are further overlaid on top of the *Herschel* derived $N(H_2)$ map (A. Singh et al. 2020, in preparation). All the spines identified are used to sort fitted slabs into vc-structures, but only vc-structures with spines longer than 15 pixels in length (~6 beam widths) are considered filaments in our analysis. In NGC 1333, we identified 10 velocity-coherent filaments in total, and have labeled them alphabetically from "a" to "j" in Figure 11.

4.2. NGC 1333—Velocity Gradients

The center and right panels of Figure 11 show, respectively, the spatial distribution of perpendicular and parallel components of velocity gradients (∇v_{LSR}) in the NGC 1333 filaments. These filaments display a wealth of ∇v_{LSR} structures within them. A large fraction of these pixels have values of $|\nabla v_{LSR}| > 2$ km s^{-1} pc^{-1} in both components, with many of them exceeding 4 km s^{-1} pc^{-1}.

At our sampling distance of two beam widths (~0.06 pc in NGC 1333), our measured velocity gradients appear consistent with those median values reported by Hacar et al. (2013) in NGC 1333, measured with N_2H^+ in the parallel direction on
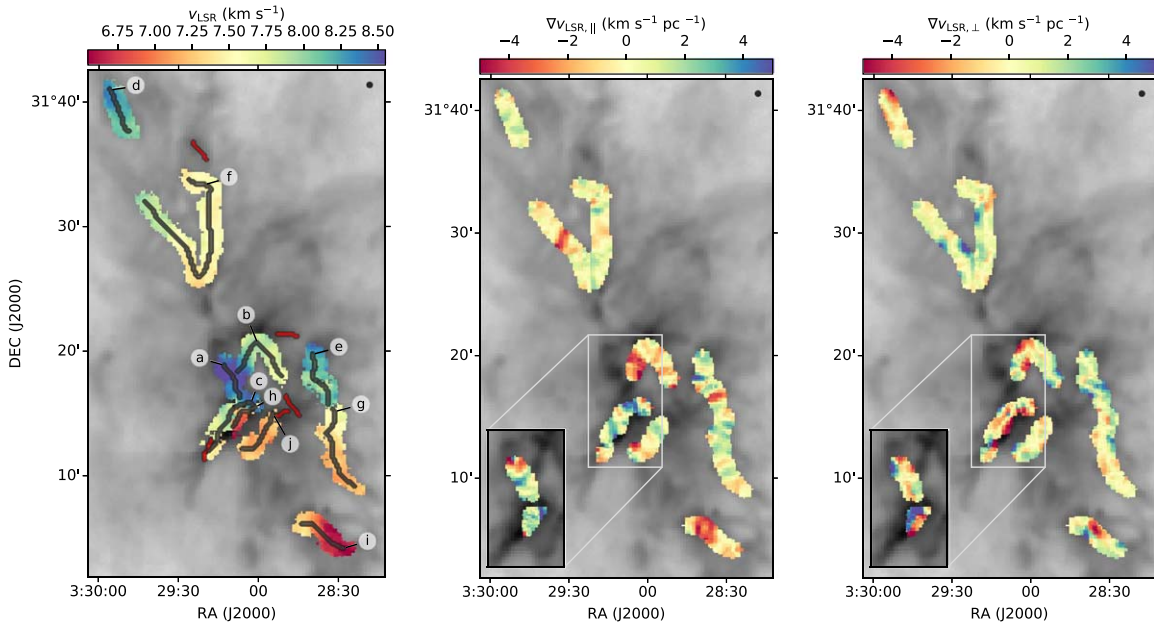
Figure 11. Left panel: projections of filament spines identified in NGC 1333, overlaid on top of the v_{LSR} maps of selected velocity-coherent filaments (color). Spines less than 15 pixels in length are colored in red and their associated filaments are excluded from our analysis. Center and right panels: spatial distribution of the parallel and perpendicular components of ∇v_{LSR}, respectively, relative to their filament spines. The call-out boxes in these panels show the same ∇v_{LSR} components of the additional, overlapping velocity-coherent filaments in the sky. The FWHM beam of the NH_3 (1, 1) data (black circle) and the *Herschel* $N(H_2)$ map (gray; A. Singh et al. 2020, in preparation) are shown on the top right corner and the background of each panel, respectively.

the same spatial scale. Similarly, Lee et al. (2014) also reported comparable values in Serpens Main, measured in the parallel direction for filaments with mass $\sim 4\ M_\odot$. Parallel gradients measured on larger scales (>0.2 pc), however, tend to have smaller values. For example, the Serpens South filaments (Kirk et al. 2013; Fernández-López et al. 2014) and the Serpens Main filaments with masses of $\sim 15\ M_\odot$ (Lee et al. 2014) all have larger-scale parallel gradients $\leqslant 1.5\ \mathrm{km\ s^{-1}\ pc^{-1}}$.

5. Discussion

5.1. Comparing with N_2H^+ Analysis of NGC 1333

Hacar et al. (2017, hereafter H17), conducted a multi-component spectral analysis of NGC 1333 with a dense gas tracer, i.e., N_2H^+ (1−0) lines. Their data have a spatial and spectral resolution (30″ and 0.06 km s^{-1}, respectively) similar to our NH_3 (1, 1) data, and a typical rms noise of 0.15 K, which is about 50% higher than ours. H17 fitted their observations with a semiautomatic method, using either one- or two-component models as determined by eye.

About 15% of the spectra successfully fitted by H17 are fitted with a two-component model. Considering not all of these successfully fitted spectra have S/N > 3, we estimate upwards to about 20% of those spectra with S/N > 3 are fitted with two-component models. This estimate assumes all the successful two-component fits have S/N > 3 in this limit. This 20% fraction is significantly lower than that of our NH_3 (1, 1) fits, where two-component models best fit about 40% of our spectra with S/N > 3. The difference in the model-selection criteria between our method (MUFASA) and that of H17 may contribute predominantly to this reported difference, where the conservative Δv_{LSR} threshold adopted by H17 may have culled out a significant fraction of their two-component fits.

Figure 12 shows the f_σ and Δv_{LSR} values derived from two-slab fits in NGC 1333 plotted against each other. The colored
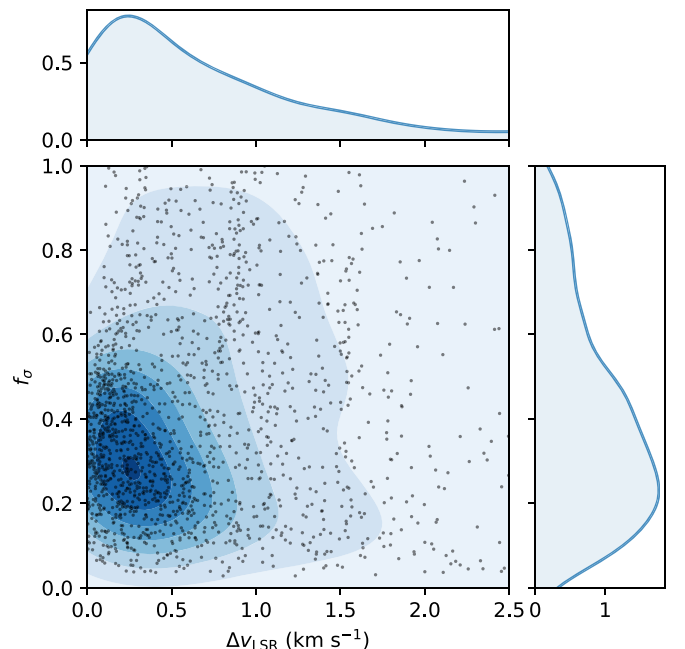

Figure 12. The ratio between the fitted line widths of two velocity slabs (i.e., f_σ), plotted against velocity separation between these slabs (Δv_{LSR}) for successful fits in NGC 1333. The colored background shows the contours of a kernel density estimate (KDE) while the filled curves to the top and the right show the 1D KDE distributions of f_σ and Δv_{LSR}, respectively.

contours in the background show the kernel density estimate of this scatter plot. The filled curves to the right and the top of the plot show the 1D kernel estimated distributions of f_σ and Δv_{LSR}. Most of the points on this plot cluster around f_σ and Δv_{LSR} values of 0.3 and 0.25 km s^{-1}, respectively. This clustering places these values in the regime where the true-positive identification rate for two-slab spectra is in the range of

~70%–90% according to our truth test shown in Figure [3](b), accounting for all S/N values found in our synthetic test set. Given that MUFASA's performance at identifying two-slab spectra decreases toward higher f_σ values and lower Δv_{LSR} values, the true peak of the underlying two-slab population likely sits higher on the f_σ axis and lower on the Δv_{LSR} axis. We reiterate that MUFASA only misidentifies a one-slab spectrum as a two-slab spectrum in <4% of the test cases for all S/N values.

About 40% of our two-slab fits to spectra with S/N > 3 in NGC 1333 have Δv_{LSR} values that are less than 0.25 km s^{-1}, the threshold used by [H17](#) to determine whether or not additional components are justified for their fits to N$_2$H$^+$ (1−0) observations of the same region. If NH$_3$ (1, 1) and N$_2$H$^+$ (1−0) indeed trace the same gas population in this region, then the fraction of dense gas spectra in NGC 1333 with multiple velocity components and S/N > 3 may be significantly underreported by [H17](#) due to their choice of Δv_{LSR} threshold. Given that two-slab identification with MUFASA is successful even with moderate S/N values (i.e., 5–20; see Figure [2](#)), the actual number of two-slab spectra with S/N > 3 is likely higher than those reported in both the [H17](#) study and our study here.

The NH$_3$ (1, 1) and N$_2$H$^+$ (1−0) lines have critical densities of ~2 × 10^3 cm^{-3} and ~5 × 10^4 cm^{-3}, respectively, at gas temperatures $\lesssim$20 K (Shirley [2015](#)). The ratio between these critical densities remains similar even at higher temperatures. If the second velocity component detected in our study tends to trace more diffuse gas, then the difference between the reported fraction of multicomponent spectra between this work and that of [H17](#) may be due to density differences in the tracers themselves in addition to the line-fitting methods used. The sensitivity difference between our data and that of [H17](#) may also play a role as well.

A recent high angular-resolution study of NGC 1333 concluded that NH$_3$ (1, 1) and N$_2$H$^+$ (1−0) trace the same gas population well (Dhabal et al. [2019](#)). Since this study only fits one velocity component along each line of sight, however, it is unclear how robust their conclusion is. Further investigation on how well NH$_3$ (1, 1) and N$_2$H$^+$ (1−0) trace each other in NGC 1333 is thus needed, particularly for diffuse emission to which the data of Dhabal et al. ([2019](#)) are less sensitive.

The filament spines we identified from our NH$_3$ data with CRISPY (see Figure [11](#), left) are morphologically similar to the "filament axes" identified by [H17](#) with N$_2$H$^+$ observations. Some of the longer filaments, however, are "split" differently. Our filament f, for example, is split into filaments *12* and *14* by [H17](#), while our filament g is split into filaments *1* and *2* by [H17](#). Moreover, we identify a kinematically distinct filament (i.e., h) that was not identified earlier by [H17](#), which runs closely parallel to our filament c (i.e., *10*).

Even though the spatial separation between spines of filament c and h is only slightly resolved in our data, the spectral separation of these spines (~0.9 km s^{-1}) is well resolved. When observed at higher spatial resolutions with NH$_3$ and N$_2$H$^+$ (Dhabal et al. [2019](#)), these two filaments can be seen by eye as distinct structures. Filament h is thus likely missed by [H17](#) due to their approach rather than observational biases introduced by the tracers used.
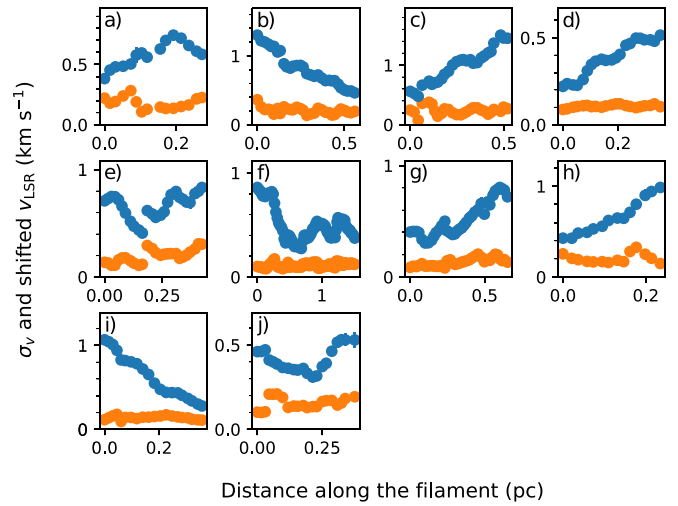

Figure 13. The profiles of σ_v (orange dots) and v_{LSR} (blue dots) along filament spines identified in NGC 1333. The v_{LSR} shown here have been shifted to fit on the same axes as σ_v.

5.2. Velocity Gradients on Large Scales

Figure [13](#) shows 1D σ_v and v_{LSR} profiles of the filaments identified in NGC 1333, taken directly from the pixels on their spines. The v_{LSR} values displayed in Figure [13](#) have been zero-point shifted arbitrarily from the local standard of rest (i.e., LSR) to fit nicely on the same axes as the σ_v. We note that these spine profiles run in the direction that starts on the ends closest to the map origin, i.e., the south-eastern corner.

The v_{LSR} variations along these spines are generally smooth with respect to their estimated errors. Only a few discontinuities are found in these profiles. Such a lack of discontinuity suggests that CRISPY does indeed identify density ridges robustly as continuous structures, which form the basis of our membership assignment to velocity-coherent filaments.

The v_{LSR} values shown in Figure [13](#) do vary significantly along most of the identified filaments. Many of the v_{LSR} profiles are approximately linear and monotonic on scales larger than 0.2 pc, and have velocity gradients typically in the range of 0.8–2.5 km s^{-1} pc^{-1} on those scales. The large-scale (>0.2 pc) velocity gradients are similar to those found with N$_2$H$^+$ observations in NGC 1333 (~0.5–2.5 km s^{-1} pc^{-1}, Hacar et al. [2017](#)) and the Serpens South SF region (1.4 km s^{-1} pc^{-1}, Kirk et al. [2013](#); ~1 km s^{-1} pc^{-1}, Fernández-López et al. [2014](#)). These gradients also fall within the range measured in the Serpens Main SF region (0.7–4.8 km s^{-1} pc^{-1}, Lee et al. [2014](#)).

Velocity gradients along filaments on a large scale have often, but not uniquely, been interpreted as signatures of mass flow toward a star-forming core or cores (e.g., Kirk et al. [2013](#)). While dense structures in NGC 1333 do not necessarily display a classical "hub and spokes" geometry, overdensities of dense cores and Class 0/I YSOs, as estimated by Hacar et al. ([2017](#)), are often found at the ends of our filament spines in projection. Filaments c and h, for example, have very linear v_{LSR} profiles and large-scale ∇v_{LSR} values of roughly 1.9 km s^{-1} pc^{-1} and 2.5 km s^{-1} pc^{-1}, respectively. The fact that the north-western ends of these filaments coincide with the most prominent peak of overdensities in NGC 1333, in the SVS 13 vicinity (see Walawender et al. [2008](#)), suggests these filaments are indeed transporting gas along their lengths toward a small (i.e., $n < 10$) cluster/protocluster.

While filament b has no end that correlates with an overdensity of dense cores and Class 0/I YSOs, it does have an overdensity midway through its length, located at the apex of its sharp turn near the HH 12 IR sources (i.e., VLA 42; see Walawender et al. 2008). Remarkably, this filament has a very linear and continuous v_{LSR} profile despite having such a distinct bend in its middle. Considering that this profile has a large-scale ∇v_{LSR} of $1.4 \, km \, s^{-1} \, pc^{-1}$, filament b is likely a velocity-coherent system of two filaments that are transporting material toward a small hub.

Interestingly, the ∇v_{LSR} seen on the largest scale of our observation ($\sim$4 pc), i.e., at the clump scale, also appears to be fairly ordered along the north–south direction. Nearly all the filaments featuring linear v_{LSR} profiles along their spines have v_{LSR} values that increase northwards. Filament b is the only exception, where half of its western segment prior to its sharp bend has v_{LSR} values that increase southwards instead. Even though filament f does not have an overall linear v_{LSR} profile, its western portion prior to its sharp bend does have a segment, $\sim$0.2 pc in length, with a fairly linear v_{LSR} profile and a ∇v_{LSR} of $\sim$2.5 $km \, s^{-1} \, pc^{-1}$. The v_{LSR} values of this segment increase northwards as well.

In addition to the prevalent trend that v_{LSR} increases northward in most filaments, the median v_{LSR} value of each filament tends to increase northwards across the NGC 1333 clump as well. Considering that NGC 1333 is relatively elongated in the north–south direction on the clump scale ($>$4 pc; see map by Sadavoy et al. 2012 for example), most of these filaments may trace a larger filamentary inflow like those assumed by Matzner & Jumper (2015) in their model, of which these smaller filaments may be a part.

5.3. Velocity Gradients on Small Scales

5.3.1. Parallel Components

In addition to well-organized velocity structures on larger scales, Figures 11 and 13 reveal many quasi-oscillatory v_{LSR} structures that can be found on the 0.1 pc ($\sim$3 beam widths) scale in NGC 1333. This behavior is prominently visible along the spines of many filaments (see Figure 13) and shows up in the parallel $\nabla v_{LSR,\parallel}$ map (see Figure 11, center) as "zebra stripes." These small-scale structures, e.g., gradient peaks and dips, also appear to be somewhat evenly spaced by $\sim$0.1 pc, which suggests a quasi-oscillatory wavelength of $\sim$0.2 pc. Interestingly, this behavior is not confined to the spines of filaments and extends spatially across the width of a filament.

Similar quasi-oscillatory v_{LSR} behaviors have been found in Taurus L1495/B213 (Tafalla & Hacar 2015) and in Serpens South filaments (Fernández-López et al. 2014) based on one-component fits to N_2H^+ (1−0) observations. Filaments and fibers identified from multicomponent fits to $C^{18}O$ (1−0) observations in Taurus-Auriga L1517 (Hacar & Tafalla 2011), and Taurus L1495/B213 (Hacar et al. 2013), respectively, also showed similar results. These quasi-oscillatory v_{LSR} behaviors generally resemble those seen in synthetic $C^{18}O$ observations of simulated filaments, constructed with various degrees of realism (e.g., Moeckel & Burkert 2015; Smith et al. 2016; Clarke et al. 2018).

We find no strong spatial correlations between quasi-oscillatory v_{LSR} and dense structures in NGC 1333. This result is contrary to that found in Taurus-Auriga L1517 (Hacar & Tafalla 2011) using $C^{18}O$ observations but agrees with the behavior found in Taurus L1495/B213 (Tafalla & Hacar 2015) using N_2H^+ (1−0) observations. This agreement extends to synthetic $C^{18}O$ observations of simulations (e.g., Smith et al. 2016). The lack of correlation between quasi-oscillatory v_{LSR} values and dense structures suggests the former is not driven by periodic gravitational instabilities. Alternative mechanisms, such as magnetic waves explored by Tritsis & Tassis (2016, 2018), and Offner & Liu (2018), may be responsible for these quasi-oscillatory behaviors.

5.3.2. Perpendicular Components

Filaments in NGC 1333 also contain a wealth of perpendicular velocity gradients, i.e., $\nabla v_{LSR,\perp}$, structures on smaller scales (see right panel of Figure 11). Regions with high $|\nabla v_{LSR,\perp}|$ values ($>$2 $km \, s^{-1} \, pc^{-1}$) tend to form spatially compact but resolved $\nabla v_{LSR,\perp}$ structures on the outskirts of the filaments, i.e., away from the spine. Similar to interpretations made in the literature (e.g., Palmeirim et al. 2013; Dhabal et al. 2018), these compact $\nabla v_{LSR,\perp}$ structures may be indicative of recent or ongoing accretion of nearby gas onto the filaments.

Freefall accretion in analytic models typically has estimated infall velocities of a few $km \, s^{-1}$ at filament "boundaries" (e.g., Heitsch 2013; Palmeirim et al. 2013). Such an infall velocity will likely result in shocks if the accreting filament is in hydrostatic equilibrium like those described in classic models (e.g., Stodólkiewicz 1963; Ostriker 1964). Even in numerical models for which filaments are not equilibrium substructures, shock-induced discontinuities in velocities are expected from accretion (e.g., Clarke et al. 2018).

We neither saw nor expected velocity discontinuities in our filaments because velocity-coherent structures are continuous in their velocities by definition. Velocity discontinuities, however, can be inferred from the Δv_{LSR} observed between velocity slabs along a line of sight. With the exception of filaments c and h, which have typical Δv_{LSR} values of $\sim$0.8 $km \, s^{-1}$ between them, we did not find filaments with overlapping velocity slabs that had Δv_{LSR} values greater than 0.4 $km \, s^{-1}$, i.e., about twice the isothermal sound speed at 10 K. Interestingly, the region where filaments c and h overlap along lines of sight is also where some of the most prominent $\nabla v_{LSR,\perp}$ structures are found in NGC 1333. We note that this observed $\nabla v_{LSR,\perp}$ structure is unlikely driven by the highly collimated outflow originated from IRAS 4 (e.g., Blake et al. 1995) given that it poorly aligns with the orientation of the outflow.

Except for where filaments overlap in projection, we typically only detect two-slab spectra near filament spines rather than the edges. This lack of detection near the edges is likely limited by the sensitivity of our data. Not much information is therefore available on Δv_{LSR} over filament edges to infer the nature of spatially compact $\nabla v_{LSR,\perp}$ structures that reside there. Furthermore, despite having a critical density of $10^3 \, cm^{-3}$, it remains unclear how effective NH_3 is at tracing accretion flows, which themselves are likely more diffuse than the dense filaments. Further investigation with NH_3 synthetic observations, similar to that conducted by Clarke et al. (2018) with $C^{18}O$ transitions, would be highly valuable.

To search for potential sources of accretion flows, we looked for structures around our filaments in the column density, i.e., $N(H_2)$, map of NGC 1333 derived from *Herschel* observations

(A. Singh et al. 2020, in preparation). We find no strong spatial correlation, however, between ambient *Herschel* structures (e.g., subfilaments) and the observed $\nabla v_{LSR,\perp}$ structures. This lack of correlation suggests that accretion from subfilaments, such as those seen in Taurus B211/B213 (Goldsmith et al. 2008; Palmeirim et al. 2013), is unlikely to explain the origin of the compact $\nabla v_{LSR,\perp}$ structures seen near filament edges.

Nevertheless, the lack of visible, interconnected ambient structures does not necessarily rule out $\nabla v_{LSR,\perp}$ as a sign of accretion flows onto dense filaments in NGC 1333. According to models where a postshock layer of a converging flow produces filaments (e.g., Chen & Ostriker 2014; Chen & Ostriker 2015), a subfilamentary network only arises in a strong magnetic field. In these models, gravity drives the accretion flows in a postshock layer. The resulting flows move predominately along the field lines and may not necessarily contain substructures with densities high enough to be distinguished from the rest of the planar, accretion flow. Without visible substructures, these flows may appear as a large-scale background to *Herschel* due to their planar geometry, making them difficult to discern.

Postshock accretion models, such as those developed by Chen & Ostriker (2014), have been proposed by Dhabal et al. (2019) as an explanation for the observed large $\nabla v_{LSR,\perp}$ along the south-western edge of filament *c*. This particular $\nabla v_{LSR,\perp}$ structure has been found in both the high-resolution NH$_3$ and N$_2$H$^+$ observations by Dhabal et al. (2019) as well as our NH$_3$ observations. The filament *h* we identify with two-slab fits, which runs parallel to filament *c*, also displays similar $\nabla v_{LSR,\perp}$ over the same region. In a postshock accretion interpretation, such a similarity suggests that filament *h* belongs to the same planar flow as filament *c*. Interestingly, filament *h* is spatially well resolved in the high-resolution observation by Dhabal et al. (2019) as a filament distinct from *c*.

It is worth noting that Walsh et al. (2006) measured infall velocities of $\sim 1\,\mathrm{km\,s}^{-1}$ toward the south-western edge of filament *h* with HCO$^+$ (1−0) observations, modeled as self-absorbed lines. The infall velocities measured with HCO$^+$ (1−0), which were interpreted as a sign of large-scale (~ 0.2 pc) infall, are similar to the observed v_{LSR} separation ($\sim 0.8\,\mathrm{km\,s}^{-1}$) between filaments *c* and *h* in NH$_3$ along lines of sight. Given that this infall region spatially correlates with filaments *c* and *h*, the HCO$^+$ (1−0) observed there may indeed trace the same planar accretion flow as that suggested by the large NH$_3$ $\nabla v_{LSR,\perp}$ we see toward filaments *c* and *h*.

Not all the observed $\nabla v_{LSR,\perp}$ structures in our filaments can be well explained by models of accretion flow along a postshock layer. While the compact nature of $\nabla v_{LSR,\perp}$ may be explained by clumpy, inhomogeneous accretion, the sign (i.e., direction) alternation of $\nabla v_{LSR,\perp}$ along filament edges, however, does not conform well to the planar geometry naively expected from a postshock layer. Some, if not all, of these observed $\nabla v_{LSR,\perp}$ features, may thus be driven by a different physical process.

While rotation of small bodies, such as dense cores, may produce compact $\nabla v_{LSR,\perp}$ signatures, no clear correlation exists between cores and many of these compact $\nabla v_{LSR,\perp}$ regions in the majority of the cases. Inhomogeneous accretion flows, on the other hand, similar to those seen in the nonmagnetized simulation by Clarke et al. (2017) may explain the sign-changing behavior of these $\nabla v_{LSR,\perp}$ along filament edges. Indeed, the shocked regions bordering dense structures
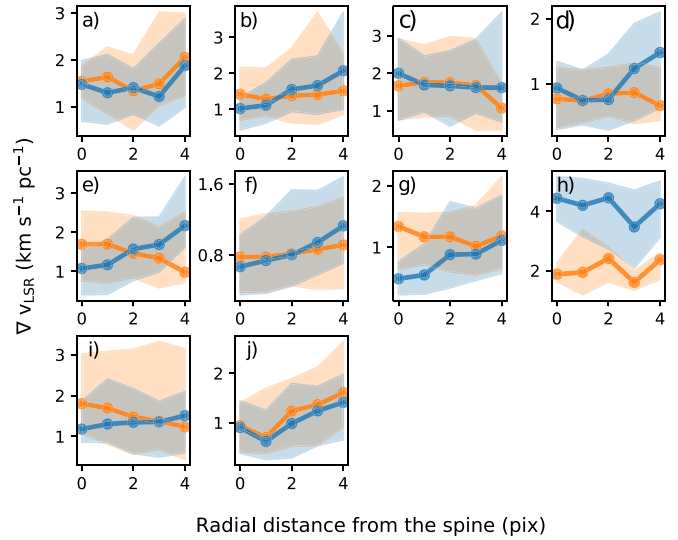

Figure 14. Median magnitudes of the perpendicular (blue) and parallel (orange) velocity gradients plotted as a function of their distance to the filament spines. The median values are marked by the connected dots while the 15th and 85th percentile values mark the lower and upper bound of the shaded regions, respectively.

in their simulation, as traced by local velocity divergence, morphologically resemble the $\nabla v_{LSR,\perp}$ structures seen in NGC 1333. Further investigation with synthetic observations of simulations is needed to see if such a resemblance holds and whether or not magnetic fields play an important role in such an accretion.

5.3.3. Radial Dependencies

Figure 14 shows the magnitudes of the perpendicular and parallel velocity gradients plotted as functions of their respective distances from the filament spines. Their median, 15th, and 85th percentile values are marked. While the parallel velocity gradients show little correlation with their distances from filament spines, the perpendicular velocity gradients tend to decrease as one moves toward the spine in many filaments.

Specifically, filaments *b*, *e*, *f*, *g*, and *j* clearly show such a trend. Indeed, the Wald Test (see Fahrmeir et al. 2013) revealed that these filaments all have p-values <0.01 for a null hypothesis where the data are consistent with a zero-slope linear trend. The linear least-squares regression slopes of these trends are 0.40, 0.41, 0.17, 0.21, and 0.16, respectively, in units of km s^{-1} pc^{-2}. Moreover, the Pearson correlation coefficients, i.e., *r*-values (see Cohen 1988), of these regressions all fall in the range of 0.2–0.4, suggesting that these positive correlations are indeed significant but relatively scattered.

The decreasing trends in these filaments persist when gradients are calculated with smaller sampling distances ($r < 3$ pixels). Indeed, the regression slopes for these filaments actually increases with smaller sampling distances.

A decrease in $|\nabla v_{LSR,\perp}|$ toward a filament spine contradicts the behavior of free-falling gas. Such a freefall behavior is often assumed for gas accretion onto a filament in simple analytical models (e.g., Heitsch et al. 2009; Palmeirim et al. 2013). While filaments themselves are not typically expected to behave like pressureless systems, except for maybe the very massive ones (e.g., $M \sim 600\,M_\odot$; Gómez & Vázquez-Semadeni 2014), a simple analytical examination of such an
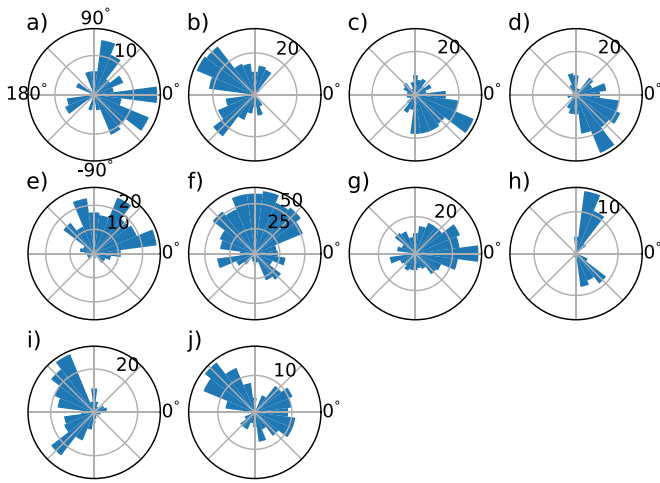
Figure 15. Polar histograms of the $\nabla v_{\rm LSR}$ orientations shown for each filament.

assumption can still provide valuable perspectives in light of our results.

Consider a parcel of gas falling onto an infinitely long cylinder axially centered at $r = 0$. Such a parcel would have a velocity profile of $v_{\rm ff} \propto [\ln(r_0/r)]^{1/2}$ if the parcel was initially at rest at r_0 (see Heitsch et al. 2009). The radial derivative of this profile, i.e., $\nabla v_{\rm LSR,\perp}$, thus would be $dv_{\rm ff}/dr \propto [r^2 \ln(r_0/r)]^{-1/2}$ and would increase monotonically in magnitude toward the filament spine for r values less than $\sim r_0/2$. The observed $\nabla v_{\rm LSR,\perp}$ for that parcel of gas should thus increase toward a filament spine if the parcel's emission dominates over others along the line of sight, and its trajectory is not parallel to the plane of the sky.

While such a geometric assumption does not describe a simple, axially symmetric accretion of a filament, it reasonably approximates inhomogeneous (e.g., Clarke et al. 2017) or anisotropic (e.g., Chen & Ostriker 2014) accretion flows found in realistic simulations. In fact, the former model is unlikely to produce observable $\nabla v_{\rm LSR,\perp}$ in the first place, contrary to what we have observed. Thus, if these nonsymmetric assumptions hold true for our observations, then the observed decrease in $|\nabla v_{\rm LSR,\perp}|$ toward the filament spine indeed suggests these filaments do not behave like a pressureless system under the accretion flow interpretation. The $\nabla v_{\rm LSR,\perp}$ we observed may thus suggest ongoing accretion that is being damped by the higher density material as the accreting gas moves closer to the filament spine.

5.3.4. Orientations

Figure 15 shows the orientation angles of the measured velocity gradients, θ, binned into polar histograms. The direction along a filament spine away from the end closest to the image origin, i.e., the bottom left corner, defines the zero-point reference of our angle θ. The convention is set such that vectors with $-180° < \theta < 0°$ point away from the spine and vectors with $0° < \theta < 180°$ point toward the spine.

Most of the $\nabla v_{\rm LSR}$ vectors within an NGC 1333 filament are not randomly oriented and often display unimodal or bimodal behaviors. A circular statistics analysis conducted with a Rayleigh test (Wilkie 1983) shows that the θ values found in 9 out of the 10 filaments are very unlikely to have been drawn from a random distribution ($p < 0.01$).

Coherent $\nabla v_{\rm LSR}$ orientations may seem perplexing at first considering how complex the $\nabla v_{\rm LSR}$ structures appear on small scales. A clearer picture emerges, however, when the $\nabla v_{\rm LSR,\parallel}$ on large scales and the radial dependency of $|\nabla v_{\rm LSR,\perp}|$ on small scales are considered together. For example, mass flows along filaments combined with perpendicular (i.e., radial) accretion onto filament edges may have caused the preferential $\nabla v_{\rm LSR}$ orientations observed here. After all, physics that are likely important in forming filaments, e.g., gravity and magnetic fields, do tend to impose order. If the observed $\nabla v_{\rm LSR}$ values are indeed indicative of mass flows, then filaments in NGC 1333 may be viewed as loci of collapsing flows where radially accreted gas changes direction to flow along filaments and into cores.

6. Summary

In this paper, we devised an efficient and robust method (i.e., MUFASA) to fit two-slab NH_3 (1, 1) spectra automatically, one that is generalizable for other molecular species. We tested the performance of MUFASA on synthetic spectra and found it to be particularly robust at identifying two-component spectra with dissimilar σ_v between their components. This behavior makes MUFASA effective at disentangling subsonic gas from super-sonic gas along lines of sight. By selecting our best-fit model via a statistical approach, we recovered $\sim 40\%$ of the two-component NH_3 spectra from our synthetic data with $\Delta v_{\rm LSR}$ values below the $0.25\ {\rm km\ s^{-1}}$ culling threshold used by Hacar et al. (2017) for their study of NGC 1333 with N_2H^+. Based on our fit residuals, we find no strong evidence that three-component fits are warranted in NGC 1333.

We identified 3D filament spines in ppv space from MUFASA line-fitted models using our implementation of the generalized (i.e., weighted) SCMS algorithm (see Chen et al. 2014), CRISPY. The generalized SCMS operates on a well-established statistical framework where the orientations of the density ridges (e.g., spines) are well defined locally. We sorted the fitted models into velocity-coherent filaments with these spines and measured the velocity gradients of their filaments on a beam-resolved scale (~ 0.05 pc). We further decomposed these velocity gradients into parallel and perpendicular components with respect to the local spine.

By applying our analysis techniques to the observation of NGC 1333, we found the following:

1. Many filaments have remarkably linear changes in $v_{\rm LSR}$ along their spines on larger scales (>0.2 pc). The $\nabla v_{\rm LSR}$ corresponding to these changes are in the range of $0.8–2.5\ {\rm km\ s^{-1}\ pc^{-1}}$, similar to those found in previous works (e.g., Kirk et al. 2013). Several of these filaments have ends or sharp bends that spatially correlate with the kernel estimated overdensities of dense cores and embedded YSOs derived by Hacar et al. (2017). This result suggests the observed $\nabla v_{\rm LSR}$ may indeed be tracing mass flow along filaments toward star-forming cores.

2. Most filaments with a linear $v_{\rm LSR}$ profile along their spines have values that increase northwards. Considering that the median $v_{\rm LSR}$ of these filaments tends to increase northwards as well, these filaments may trace a larger-scale (>4 pc) filamentary accretion flow, similar to that assumed by Matzner & Jumper (2015) in their model, from which the NGC 1333 clump may be fed.

3. The $\nabla v_{\mathrm{LSR},\parallel}$ measured on small scales show quasi-oscillatory v_{LSR} along filaments. These quasi-oscillations, however, do not correlate well with dense structures, a result similar to those found by Tafalla & Hacar (2015) in real N_2H^+ observation and by Smith et al. (2016) in synthetic $C^{18}O$ observations of a simulation. This lack of correlation suggests periodic gravitational instabilities are not responsible for such an observed behavior. Alternative mechanisms, such as magnetic waves (e.g., Tritsis & Tassis 2016) may be responsible instead.

4. The $\nabla v_{\mathrm{LSR},\perp}$ found on small scales tend to form compact structures near the filament edges, potentially indicating perpendicular accretion flows. The compact nature of these $\nabla v_{\mathrm{LSR},\perp}$ structures combined with an apparent lack of ambient subfilaments suggest these accretion flows are likely clumpy, i.e., inhomogeneous. Alternations in the direction of these $\nabla v_{\mathrm{LSR},\perp}$ structures along filament edges also suggest these accretion flows may not be purely planar like those found in simulated magnetized postshock layers (e.g., Chen & Ostriker 2014), except possibly those first reported by Dhabal et al. (2019) for filaments c and h.

5. The magnitudes of the measured $\nabla v_{\mathrm{LSR},\perp}$ decrease prominently toward filament spines in half (i.e., 5) of our filaments. Assuming our observations trace gas flows that are inhomogeneous or anisotropic with respect to the filament spines, such a trend is inconsistent with freefall accretion models and suggests that these filaments do not behave like pressureless structures. Such an observed behavior may thus indicate the infall of accretion flows being damped by the denser, pressure-supported gas within filaments.

6. The ∇v_{LSR} vectors measured on small scales are not randomly oriented within a filament. Their orientations tend to be unimodally or bimodally distributed. This global trend within filaments conforms to a scenario by which the gas falling onto a filament is redirected to flow along the filament as it approaches the spine.

For our interpretations, we assumed the observed velocity gradients are indeed signs of accelerating gas seen along lines of sight. We plan to apply our analysis to synthetic NH_3 observations in future work to understand better the nature of these velocity gradients.

Appendix A
Parameter Choices for CRISPY

As mentioned in Section 3.3.2, the operation of SCMS, including our implementation in CRISPY, primarily depends on two user-defined parameters: density (i.e., intensity) threshold and smoothing bandwidth. We adopted a density threshold of 0.15 K to capture most of our emission in the model while avoiding going near the typical rms noise level of our data ($\sim$0.1 K). We further adopted a smoothing length of 1.5 pixels, which corresponds to about the 1σ sampling width of our data, i.e., $\sim$3 pixels across the FWHM beam. Since our deblended cube at its native resolution is too noisy for SCMS even after a density cutoff, we further smoothed our deblended cube spatially to twice its original beamwidth prior to running SCMS. We smoothed the data only spatially and not spectrally to avoid further loss in our spectral resolution.

In addition to density threshold and smoothing bandwidth parameters, SCMS requires a few additional parameters to run in practice: convergence criterion, the maximum number of iterations, and walker placement. Convergence criterion and the maximum number of iterations are used to decide when to stop running SCMS further. We set our convergence criterion to 10^{-3} to ensure the ridges represented by the converged walkers have scatters that are smaller than the equivalent width of a voxel in the deblended cube. We set the maximum number of iterations to 1000 and the walker placement such that a walker is placed at each voxel in the deblended cube above a density threshold. We adopted a walker placement threshold of 0.16 K to sample the density field well without placing walkers near the edges of the field defined by our cutoff threshold (0.15 K).

The CRISPY implementation of SCMS also includes scaling parameters for which each dimension of the deblended emission can be rescaled. The purpose of such a rescaling is to renormalize the distance between each particle in the field and consequently the density field. Such a renormalization is essential for structure identification in a parameter space with dimensions that are not necessarily physically related, e.g., a ppv space.

Since the two spatial distances in a ppv space are physically related, only the velocity axis needs to be rescaled, provided that the smoothing bandwidth was already chosen appropriately based on the spatial sampling. For our runs, we kept velocity scaling the same as the one native to our deblended cube. We made this choice deliberately to avoid elongating spatially compact structures along the velocity axis such that they are misidentified as filaments. Shorter scaling was avoided to prevent further loss in our spectral resolution from bandwidth smoothing.

Appendix B
Membership Assignment to Velocity-coherent Structures

As briefly described in Section 3.3.3, velocity slabs are assigned memberships to velocity-coherent structures (vc-structures) based on their proximity to filament spines in the ppv space. This process is performed iteratively, starting with

pixels spatially closest to the filament spine. At each iteration, velocity slabs are assigned to spine-derived ppv-footprints nearest to them in v_{LSR} along a line of sight, which forms what we call an association. The slabs in each association are subsequently sorted internally, based on their kinematic similarities, into a vc-structure that contains only a single velocity slab along a line of sight. Here, we describe the process of assigning membership to vc-structures in more detail.

We first create a ppv-footprint, a spatial extension of a spine, to serve as a reference from which the v_{LSR} proximities between a slab and a spine are calculated along a line of sight, i.e., a pixel. We construct a ppv-footprint by dilating, i.e., expanding, filament spines by one pixel in the two spatial dimensions but not in the velocity dimension. This expansion is initially accomplished by taking the first-moment map of a filament spine, dilating the map's on-sky footprint by one pixel, and adopting the median v_{LSR} value of the moment map within a 3 pixel radius of each pixel as its new value.

Once the ppv-footprint is created, velocity slabs with v_{LSR} values closest to the ppv-footprint at each pixel are then assigned to that ppv-footprint as a member of the association. Only slabs with velocity separations $<0.21\,km\,s^{-1}$ from the ppv-footprint, i.e., about three spectral channel widths, are accepted to ensure the assigned members are reasonably velocity coherent with respect to the ppv-footprint. We note that this threshold implicitly imposes an upper limit to velocity gradients of $0.42\,km\,s^{-1}\,pix^{-1}$ (i.e., $\sim 29\,km\,s^{-1}\,pc^{-1}$) for a given association.

Following these assignments, member slabs within each association are further sorted into a vc-structure containing only a single slab along each line of sight, based on their similarities in v_{LSR}, σ_v, and δv_{LSR}, i.e., the Jacobian estimated error of v_{LSR} from the fits. We used δv_{LSR} similarity as our additional proxy for spatial coherence assuming that spectral components which are spatially similar in their properties, such as their brightness, will produce fitted δv_{LSR} that are spatially similar as well.

At each iteration, we sort these velocity slabs as follows:

1. Assign velocity slabs with the smallest δv_{LSR} values at each pixel in a given association into a new, single slab structure we call a vc-structure.
2. Median smooth the δv_{LSR} map of the vc-structure with a circular aperture, 1 pixel in radius, to serve as a reference map.
3. Reassign velocity slabs at each pixel with the most similar δv_{LSR} values, i.e., the least difference between the smoothed δv_{LSR} map and their respective δv_{LSR} values, to the vc-structure.
4. Create reference maps of v_{LSR} and σ_v by employing the same median-smoothing technique used in step 2 on property maps of the vc-structure.
5. Reassign velocity slabs with the minimum quadrature sum difference between the reference maps and their respective v_{LSR} and σ_v values, i.e., $[(v_{LSR} - v_{LSR,\,ref})^2 + (\sigma_v - \sigma_{v,ref})^2]^{0.5}$, to the vc-structure.

Following each sortation above, we median smooth the final v_{LSR} map in each vc-structure with a 3 pixel radius aperture and adopt the resulting v_{LSR} map as our new ppv-footprint for the subsequent iteration of membership assignment and sortation. This procedure is carried out for five iterations in total, growing ppv-footprints and their respective associations radially one pixel at a time. This five pixel radial distance typically marks the spatial extent for which the S/N values of our pixels start to fall below 3.

To illustrates such a process, Figure 8 shows a ppv-footprint at each iteration in panels labeled with the iteration number n. The last two panels of Figure 8 show v_{LSR} maps of the first and second velocity slabs in an association, taken from the final iteration. The first slab shown here defines the final vc-structure used in our filament analyses.

ORCID iDs

Michael Chun-Yuan Chen ⓘ https://orcid.org/0000-0003-4242-973X
James Di Francesco ⓘ https://orcid.org/0000-0002-9289-2450
Erik Rosolowsky ⓘ https://orcid.org/0000-0002-5204-2259
Jared Keown ⓘ https://orcid.org/0000-0003-2628-0250
Jaime E. Pineda ⓘ https://orcid.org/0000-0002-3972-1978
Rachel K. Friesen ⓘ https://orcid.org/0000-0001-7594-8128
Paola Caselli ⓘ https://orcid.org/0000-0003-1481-7911
How-Huan Chen ⓘ https://orcid.org/0000-0001-6222-1712
Christopher D. Matzner ⓘ https://orcid.org/0000-0001-9732-2281
Stella S. Offner ⓘ https://orcid.org/0000-0003-1252-9916
Anna Punanova ⓘ https://orcid.org/0000-0001-6004-875X
Elena Redaelli ⓘ https://orcid.org/0000-0002-0528-8125
Samantha Scibelli ⓘ https://orcid.org/0000-0002-9485-4394

References

Akaike, H. 1974, ITAC, 19, 716
André, P., Di Francesco, J., Ward-Thompson, D., et al. 2014, in Protostars and Planets VI, ed. H. Beuther et al. (Tucson, AZ: Univ. Arizona Press), 27
André, P., Men'shchikov, A., Bontemps, S., et al. 2010, A&A, 518, L102
Arzoumanian, D., André, P., Didelon, P., et al. 2011, A&A, 529, L6
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Bally, J., Langer, W. D., Stark, A. A., & Wilson, R. W. 1987, ApJL, 312, L45
Balsara, D., Ward-Thompson, D., & Crutcher, R. M. 2001, MNRAS, 327, 715
Bastien, P., Arcoragi, J.-P., Benz, W., Bonnell, I., & Martel, H. 1991, ApJ, 378, 255
Beaumont, C. N., Offner, S. S. R., Shetty, R., Glover, S. C. O., & Goodman, A. A. 2013, ApJ, 777, 173
Blake, G. A., Sandell, G., van Dishoeck, E. F., et al. 1995, ApJ, 441, 689
Burnham, K. P., & Anderson, D. R. 2004, Sociological Methods Research, 33, 261
Chen, C.-Y., & Ostriker, E. C. 2014, ApJ, 785, 69
Chen, C.-Y., & Ostriker, E. C. 2015, ApJ, 810, 126
Chen, H.-R. V., Zhang, Q., Wright, M. C. H., et al. 2019, ApJ, 875, 24
Chen, M. C.-Y. 2020a, MUlti-component Fitter for Astrophysical Spectral Applications, 0.1, Zenodo, doi:10.5281/ZENODO.3637660
Chen, M. C.-Y. 2020b, Computational Ridge Identification with SCMS for Python, 0.1, Zenodo, doi:10.5281/ZENODO.3637613
Chen, Y.-C., Genovese, C. R., & Wasserman, L. 2014, arXiv:1406.1803
Chen, Y.-C., Ho, S., Freeman, P. E., Genovese, C. R., & Wasserman, L. 2015, MNRAS, 454, 1140
Clarke, S. D., Whitworth, A. P., Duarte-Cabral, A., & Hubber, D. A. 2017, MNRAS, 468, 2489
Clarke, S. D., Whitworth, A. P., Spowage, R. L., et al. 2018, MNRAS, 479, 1722
Cohen, J. 1988, Statistical Power Analysis for the Behavioral Sciences (Mahwah, NJ: Lawrence Erlbaum Associates)
Dhabal, A., Mundy, L. G., Chen, C.-y., Teuben, P., & Storm, S. 2019, ApJ, 876, 108
Dhabal, A., Mundy, L. G., Rizzo, M. J., Storm, S., & Teuben, P. 2018, ApJ, 853, 169
Di Francesco, J., Evans, N. J., II, Caselli, P., et al. 2007, in Protostars and Planets V, ed. B. Reipurth et al. (Tucson, AZ: Univ. Arizona Press), 17

Fahrmeir, L., Kneib, T., Lang, S., & Marx, B. 2013, Regression: Models, Methods and Applications (Berlin: Springer),

Fernández-López, M., Arce, H. G., Looney, L., et al. 2014, ApJL, 790, L19

Friesen, R. K., Di Francesco, J., Shirley, Y. L., & Myers, P. C. 2009, ApJ, 697, 1457

Friesen, R. K., Medeiros, L., Schnee, S., et al. 2013, MNRAS, 436, 1513

Friesen, R. K., Pineda, J. E., co-PIs, et al. 2017, ApJ, 843, 63

Ginsburg, A., & Mirocha, J. 2011, PySpecKit: Python Spectroscopic Toolkit (v0.1.21), Astrophysics Source Code Library, ascl:1109.001

Goldsmith, P. F., Heyer, M., Narayanan, G., et al. 2008, ApJ, 680, 428

Gómez, G. C., & Vázquez-Semadeni, E. 2014, ApJ, 791, 124

Gong, H., & Ostriker, E. C. 2011, ApJ, 729, 120

Hacar, A., & Tafalla, M. 2011, A&A, 533, A34

Hacar, A., Tafalla, M., & Alves, J. 2017, A&A, 606, A123

Hacar, A., Tafalla, M., Kauffmann, J., & Kovács, A. 2013, A&A, 554, A55

Heitsch, F. 2013, ApJ, 769, 115

Heitsch, F., Ballesteros-Paredes, J., & Hartmann, L. 2009, ApJ, 704, 1735

Henshaw, J. D., Longmore, S. N., Kruijssen, J. M. D., et al. 2016, MNRAS, 457, 2675

Hunter, J. D. 2007, CSE, 9, 90

Inutsuka, S.-i., & Miyama, S. M. 1997, ApJ, 480, 681

Jappsen, A. K., Klessen, R. S., Larson, R. B., Li, Y., & Mac Low, M. M. 2005, A&A, 435, 611

Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python (v1.2.0), http://www.scipy.org/

Keown, J., Di Francesco, J., Teimoorinia, H., Rosolowsky, E., & Chen, M. C.-Y. 2019, ApJ, 885, 32

Kirk, H., Myers, P. C., Bourke, T. L., et al. 2013, ApJ, 766, 115

Koch, E. W., & Rosolowsky, E. W. 2015, MNRAS, 452, 3435

Könyves, V., André, P., Men'shchikov, A., et al. 2015, A&A, 584, A91

Kullback, S., & Leibler, R. A. 1951, Ann. Math. Statist, 22, 79

Lee, K. I., Fernández-López, M., Storm, S., et al. 2014, ApJ, 797, 76

Levenberg, K. 1944, QApMa, 2, 164

Lourakis, M. I. A. 2005, A Brief Description of the Levenberg–Marquardt Algorithm Implemened by Levmar, Tech. Rep., Institute of Computer Science, Foundation for Research and Technology http://users.ics.forth.gr/~lourakis/levmar/levmar.pdf

Mangum, J. G., Emerson, D. T., & Greisen, E. W. 2007, A&A, 474, 679

Mangum, J. G., & Shirley, Y. L. 2015, PASP, 127, 266

Markwardt, C. B. 2009, in ASP Conf. Ser. 411, Astronomical Data Analysis Software and Systems XVIII, ed. D. A. Bohlender, D. Durand, & P. Dowler (San Francisco, CA: ASP)

Marquardt, D. W. 1963, SJAM, 11, 431

Masters, J., Garwood, B., Langston, G., & Shelton, A. 2011, in ASP Conf. Ser. 442, Astronomical Data Analysis Software and Systems XX, ed. I. N. Evans et al. (San Francisco, CA: ASP), 127

Matzner, C. D., & Jumper, P. H. 2015, ApJ, 815, 68

Men'shchikov, A., André, P., Didelon, P., et al. 2010, A&A, 518, L103

Moeckel, N., & Burkert, A. 2015, ApJ, 807, 67

Moré, J. J. 1978, Lecture Notes in Mathematics, Vol. 630 (Berlin: Springer), 105

Moré, J. J., Garbow, B. S., & Hillstrom, K. E. 1980, User Guide for MINPACK-1, Tech. Rep. ANL-80-7, Argonne Nat. Lab., Argonne, IL http://cds.cern.ch/record/126569

Offner, S. S. R., & Liu, Y. 2018, NatAs, 2, 896

Ortiz-León, G. N., Loinard, L., Dzib, S. A., et al. 2018, ApJ, 865, 73

Ostriker, J. 1964, ApJ, 140, 1056

Ozertem, U., & Erdogmus, D. 2011, The Journal of Machine Learning Research, 12, 1249

Palmeirim, P., André, P., Kirk, J., et al. 2013, A&A, 550, A38

Pon, A., Johnstone, D., & Heitsch, F. 2011, ApJ, 740, 88

Porter, D. H., Pouquet, A., & Woodward, P. R. 1994, PhFl, 6, 2133

Riener, M., Kainulainen, J., Henshaw, J. D., et al. 2019, A&A, 628, A78

Rivers, M. 2002, Perform Levenberg–Marquardt Least-squares Minimization, http://cars9.uchicago.edu/software/python/mpfit.html

Robitaille, T., & Bressert, E. 2012, APLpy: Astronomical Plotting Library in Python (v1.1.1), Astrophysics Source Code Library, ascl:1208.017

Rosolowsky, E. W., Pineda, J. E., Foster, J. B., et al. 2008, ApJS, 175, 509

Sadavoy, S. I., di Francesco, J., André, P., et al. 2012, A&A, 540, A10

Schneider, S., & Elmegreen, B. G. 1979, ApJS, 41, 87

Shirley, Y. L. 2015, PASP, 127, 299

Smith, R. J., Glover, S. C. O., Bonnell, I. A., Clark, P. C., & Klessen, R. S. 2011, MNRAS, 411, 1354

Smith, R. J., Glover, S. C. O., Klessen, R. S., & Fuller, G. A. 2016, MNRAS, 455, 3640

Sokolov, V., Wang, K., Pineda, J. E., et al. 2017, A&A, 606, A133

Sousbie, T. 2011, MNRAS, 414, 350

Sousbie, T., Pichon, C., & Kawahara, H. 2011, MNRAS, 414, 384

Stodólkiewicz, J. S. 1963, AcA, 13, 30

Sugiura, N. 1978, Communications in Statistics—Theory and Methods, 7, 13

Tafalla, M., & Hacar, A. 2015, A&A, 574, A104

Tritsis, A., & Tassis, K. 2016, MNRAS, 462, 3602

Tritsis, A., & Tassis, K. 2018, Sci, 360, 635

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22

van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., et al. 2014, PeerJ, 2, e453

Walawender, J., Bally, J., Francesco, J. D., Jørgensen, J., & Getman, K. 2008, in Handbook of Star Forming Regions, ed. B. Reipurth (San Francisco, CA: ASP), 346

Walsh, A. J., Bourke, T. L., & Myers, P. C. 2006, ApJ, 637, 860

Wilkie, D. 1983, Journal of the Royal Statistical Society Series C, 32, 311

Zucker, C., Schlafly, E. F., Speagle, J. S., et al. 2018, ApJ, 869, 83