



Probabilistic Detection of Spectral Line Components

Vlas Sokolov¹ , Jaime E. Pineda¹ , Johannes Buchner^{1,2,3} , and Paola Caselli¹

¹ Max Planck Institute for Extraterrestrial Physics, Gießenbachstraße 1, D-85748 Garching bei München, Germany; vlas-sokolov@mpg-alumni.de

² Instituto de Astrofísica, Facultad de Física, Pontificia Universidad Católica de Chile, Casilla 306, Santiago 22, Chile

³ Millennium Institute of Astrophysics, Vicuña, MacKenna 4860, 7820436 Macul, Santiago, Chile

Received 2020 January 24; revised 2020 March 5; accepted 2020 March 16; published 2020 April 3

Abstract

Resolved kinematical information, such as from molecular gas in star-forming regions, is obtained from spectral line observations. However, these observations often contain multiple line-of-sight components, making estimates harder to obtain and interpret. We present a fully automatic method that determines the number of components along the line of sight, or the spectral multiplicity, through Bayesian model selection. The underlying open-source framework, based on nested sampling and conventional spectral line modeling, is tested using the large area ammonia maps of NGC 1333 in the Perseus molecular cloud obtained by the Green Bank Ammonia Survey (GAS). Compared to classic approaches, the presented method constrains velocities and velocity dispersions in a larger area. In addition, we find that the velocity dispersion distribution among multiple components did not change substantially from that of a single-fit component analysis of the GAS data. These results showcase the power and relative ease of the fitting and model selection method, which makes it a unique tool to extract maximum information from complex spectral data.

Unified Astronomy Thesaurus concepts: [Computational methods \(1965\)](#); [Star forming regions \(1565\)](#)

1. Introduction

Spectral line observations provide a unique insight into the kinematics of astrophysical sources. However, the underlying physical complexity of the latter often results in observed spectra bearing a signature of multiple independent components with distinct radial velocities. In particular, a line-of-sight superposition of material with different radial velocities emitting in optically thin regime will naturally result in a spectral profile with multiple peaks, as seen in both Galactic (e.g., Busquet et al. 2013; Tanaka et al. 2013) and extragalactic (e.g., Koch et al. 2018) sources, as well as in molecular cloud simulations (Clarke et al. 2018). The observational signature of the multiple components, as well as the underlying number of the line-of-sight velocity components, shall be referred to as *spectral multiplicity* henceforth, and is the topic addressed in this work.

Despite the widespread occurrence of spectral multiplicity, the issue has yet to be addressed in a statistically sound manner. Traditionally, spectral multiplicity is often assessed by eye (e.g., Rosolowsky et al. 2008; Pineda et al. 2010, 2011, 2015; Hacar & Tafalla 2011; Beuther et al. 2015; Hacar et al. 2016; Pon et al. 2016; Barnes et al. 2018; Monsch et al. 2018). Recently, the drastically increasing number of spectra that can be delivered by the contemporary instruments has led to the development of a semiautomatic approach, where a by-eye fit to an averaged spectrum is propagated into individual spectra within the averaged area (Hacar et al. 2013, 2017, 2018; Henshaw et al. 2013, 2016), and the individual components are judged against heuristic criteria. More recent works have used information criteria to prevent overfitting, supplementing the conventional methods above (such as AIC in Henshaw et al. 2016; Chen et al. 2020).

However, currently employed methods to determine spectral multiplicity have their shortcomings. The commonly used approach of evaluating the number of spectral components by eye or by imposing detection heuristics does not have a statistically sound foundation. Furthermore, forward-fitting

multiple components in this fashion is invalidated by the sheer number of spectra in contemporary data sets, where tens or hundreds of thousand spectra can comprise a single observational result. Moreover, as most current approaches employ local, and not global, regression methods, the multitude of spectra analyzed in contemporary data sets creates yet another complication where the output optimal parameters are highly dependent on the initial conditions of the fit. More sophisticated frameworks such as GaussPy+ (Riener et al. 2019), while providing a fully automated way to separate Gaussian spectral components, come with limitations of their own, such as not supporting complex hyperfine structures and absorption line profiles in the case of GaussPy+. It is for these reasons that an automated, global minimization framework, complete with a statistically robust selection of spectral multiplicity, is sorely needed.

In this study, we present the results of a Bayesian model selection for measuring the spectral multiplicity on the Green Bank Ammonia Survey (GAS; Friesen et al. 2017) data for NGC1333 in the Perseus molecular cloud, through the use of the nested sampling algorithm (Skilling 2004). The framework we present requires minimal human interaction, delivers full sampling over the free parameter space, and allows well-defined detection criteria. By treating each spectrum in the spectral cube independently, we avoid biasing neighbors.

2. Method Summary

Here we explain the classical method for line fitting and our new method.

2.1. Classical Line Fitting

To model spectral lines, an observed spectrum is assumed to contain signal from an astrophysical source and a known Gaussian noise contribution with known amplitude σ . Fitting

minimizes the χ^2 statistic, defined as

$$\chi^2(\theta) = \sum_{i=1}^N \frac{(y_i - \mathcal{M}(\theta, i))^2}{\sigma^2},$$

where y_i is the observed signal amplitude at spectral channel i (of N) and $\mathcal{M}(\theta, i)$ is the model of the astrophysical source dependent on some source parameters θ (following Friesen et al. 2017) that are varied during the fitting process. The square root of the Gaussian noise variance, σ , is taken from the GAS DR1 for our ammonia application.

The χ^2 statistic originates from assuming a Gaussian likelihood valid across the spectral range, which can be written as

$$\mathcal{L}(D|\theta) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} \times \exp\left[-\frac{1}{2\sigma^2}(y_i - \mathcal{M}(\theta, i))^2\right] \quad (1)$$

or, dropping constants, as

$$\log \mathcal{L}(D|\theta) = \frac{1}{2}\chi^2. \quad (2)$$

Simple minimization algorithms have difficulties identifying globally best parameters in complex models $\mathcal{M}(\theta, i)$ when the signal is faint. There may be multiple local optima, or no well-defined optimum. Even if the fit succeeds, quantifying whether a more complex model is better than a simpler model can be difficult (Protassov et al. 2002). In the situation of deciding how many k components are justified by the data, the model of k components contains the $(k-1)$ component model at the border of the parameter space (component amplitude is zero). However, just such border situations are not permitted by model comparison methods based on F-tests, likelihood ratio tests, or more generally those relying on Wilks's theorem.

2.2. Bayesian Framework

Bayesian model selection with nested sampling solves these problems elegantly. Nested sampling (Skilling 2004) is a global parameter space exploration algorithm (as opposed to local optimization in classical frameworks), which allows both parameter estimation and model comparison, even in low-signal data. *Parameter estimation* gives the ranges of parameters θ that are probable given the data. To obtain all the posterior probability densities (and not just a global maximum of the likelihood function), we need to define prior densities over the parameter space. With these, we can use the likelihood to reweigh the prior probability density $\pi(\theta)$ to a posterior probability density $P(\theta|D)$ with Bayes's theorem:

$$P(\theta|D) \propto \pi(\theta) \times \mathcal{L}(D|\theta),$$

which can be used both for point estimation (i.e., finding the best-fit parameter values) as well as for uncertainty analysis (e.g., based on the posterior spread around the point estimate).

In our application we use uninformative (flat) priors. The prior $\pi(\theta)$ defines the free parameters of the line modeling, which are listed below. Unless specified otherwise, the value ranges indicate uniform priors.

1. Gas kinetic temperature, ranging from cosmic microwave background (CMB) temperature to 25 K.
2. Excitation temperature, ranging from CMB temperature to 25 K.

3. Velocity dispersion of the line, ranging from 0.05 to 2 km s⁻¹.
4. Logarithm of the ammonia column density, ranging from 12 to 15, probing H₂ column densities of 10²⁰–10²³ cm⁻² for typical ammonia abundances of 10⁻⁸.
5. The mean v_{lsr} of all velocity components, ranging from 3 to 10 km s⁻¹.
6. Velocity separation between the closest components, ranging from 0.2 to 3 km s⁻¹.

All the parameter priors, except for the mean v_{lsr} , are treated independently for each velocity component. The parameters 5 and 6 can be shown to linearly transform into a set of centroid velocities for each velocity component. The priors above were chosen to be broad enough to encompass all typical physical conditions of the low-mass molecular clouds (see Friesen et al. 2017 for overall distributions of the free parameters).

Model comparison evaluates whether to prefer one model over another. We first estimate the likelihood integral Z , called the Bayesian evidence, marginalized (integrated) over the entire parameter space,

$$Z = \int p(\theta)\mathcal{L}(\theta)d\theta.$$

If we have a set of equally probable models under consideration, the Bayes factor K_j^i gives the odds ratio of \mathcal{M}_i model over model \mathcal{M}_j :

$$K_j^i = \frac{P(\mathcal{M}_j)Z_i}{P(\mathcal{M}_i)Z_j} = Z_i/Z_j, \quad (3)$$

where we implicitly assume that the competing models are equally likely a priori, i.e., $P(\mathcal{M}_i) = P(\mathcal{M}_j)$.

Model comparison can also be used to quantify how strongly the data support the presence of a spectral line in the observed spectrum. A special case occurs when only noise is considered (i.e., $\mathcal{M} = 0$ in Equation (1)). If the noise amplitude is known, the corresponding evidence Z_0 can be derived analytically (see, e.g., Buchner 2017 for a similar approach).

For our model selection purposes, we consider three models: one yielding a noise-only spectrum (with analytical evidence Z_0 above), the model where the data are a sum of noise and one spectral component (its evidence Z_1 , sampled by MultiNest), and the two-component model (with Bayesian evidence denoted as Z_2). Consequently, we will refer to two Bayes factors of interest in this study: K_0^1 , denoting the odds of a spectrum being present in the data, and K_1^2 , for the odds of the two-component model prevailing over the single-component one.

We use the popular MultiNest (Feroz & Hobson 2008; Feroz et al. 2009) implementation of nested sampling through the PyMultiNest (Buchner et al. 2014) Python interface. This has been widely used in the past to perform model selection for astrophysical spectra (e.g., Bernardi et al. 2016; Feldmeier-Krause et al. 2017; Lavie et al. 2017; Baronchelli et al. 2018). The likelihood function was sampled through a wrapper to pyspeckit, specifically designed to perform nested sampling of spectral cubes.⁴ The results of this study, as well as the code to reproduce them fully, are publicly available.⁵

⁴ <https://github.com/vlas-sokolov/pyspeckit>

⁵ <https://github.com/vlas-sokolov/bayesian-ngcl333>

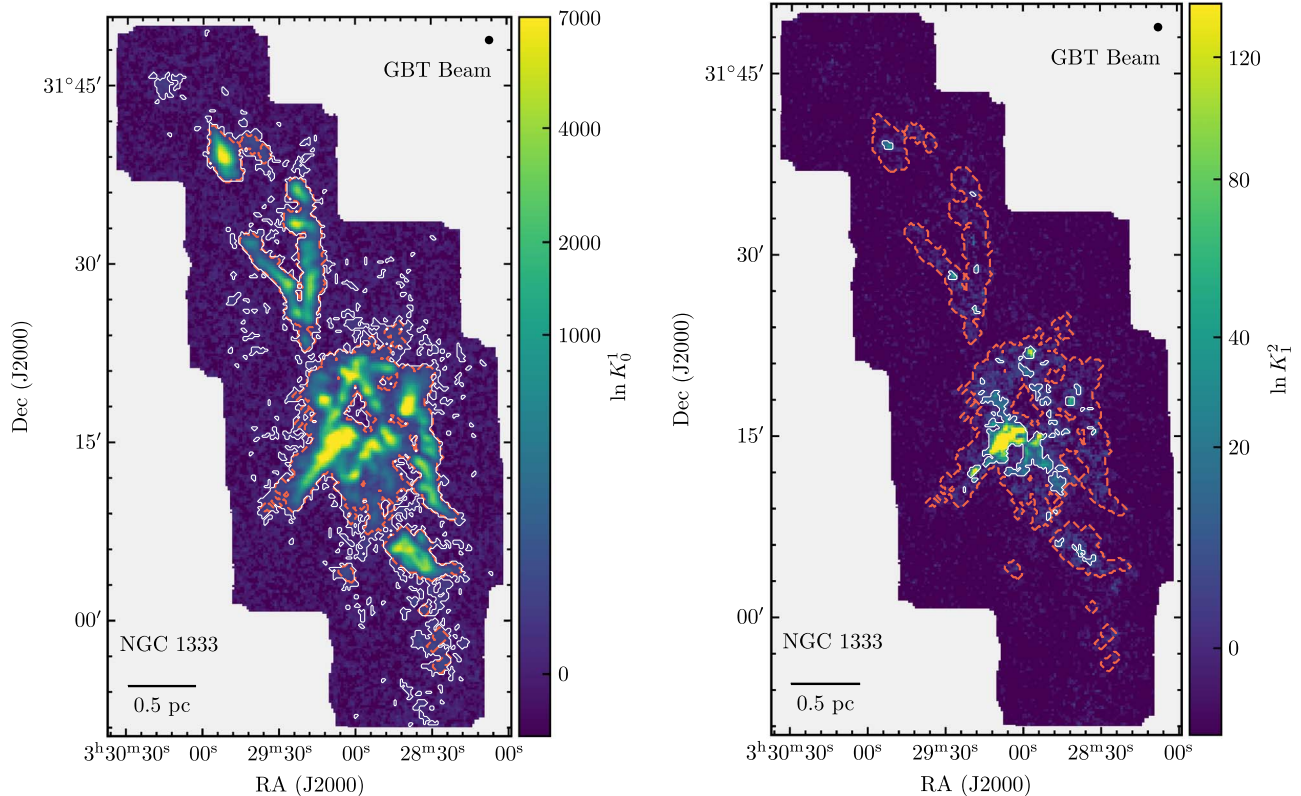


Figure 1. The odds ratios (Bayes factors) for detecting one (left, K_0^1) and two (right, K_1^2) spectral line components. Both maps are overlaid with the heuristical decision boundary of $\ln K_i^j = 5$ (solid white lines) and the previous quality assessment for significant single-spectrum detection (dotted red lines; Friesen et al. 2017). To highlight only spatially continuous spectral component detections, we have applied small feature removal on to the contours, only considering continuous regions of nine or more pixels. Scale bar and beam are shown in the bottom left and top right corners, respectively.

3. Data

We demonstrate our technique on GAS DR1 (Friesen et al. 2017) data in the NGC1333 region. These data cover ≈ 1222 arcmin² in the sky of the young embedded cluster in the nearby Perseus molecular cloud. These data were observed using the On-The-Fly mapping technique with the K -Band Focal Plane Array at the Green Bank Observatory. The spectral resolution of these data is 5.7 kHz, or ≈ 0.07 km s⁻¹ at the frequency of these lines, and a beam FWHM of 32". The typical rms of the cube is ≈ 0.11 K in the Main Beam scale. We focus on NH₃ (1, 1) and (2, 2), at rest frequencies of 23.6944955 and 23.7226336 GHz, respectively.

The numerous hyperfine subcomponents per line, high spectral resolution, and sensitivity of the observations allow an accurate determination of the gas kinematics. The initial GAS DR1 (Friesen et al. 2017) fitted a single-component fit using pyspecfit, however, acknowledged the presence of multiple components in a small fraction of the data. In the same region, other works have shown a small fraction of multiple components along the line of sight and fit then multiple components by hand (e.g., Hacar et al. 2017).

4. Results and Discussion

We fitted each pixel spectrum separately with zero-, one-, and two-component models. We present the odds of detecting one and two components in Figure 1. We emphasize that the sampling routine does not use knowledge about nearby pixels, except for the inherent correlation between nearby spectra due

to pixel sizes smaller than beam size, and that we have used the same sampling setup for every spectrum sampled.

The odds ratios for detecting a single component are shown in the left panel, overlaid with white contours of a conventional ($\ln K_0^1 = 5$; see Section 4.2) cut. The red contours, marking the stringent detection criteria imposed by the GAS DR1 quality assessment (Friesen et al. 2017), are fully enclosed within the signal detected in our method. Additional regions emitting ammonia extend beyond the DR1 detection range, indicating that our probabilistic framework is sensitive to fainter emission than the heuristic quality cuts imposed in Friesen et al. (2017).

The right panel of Figure 1 shows the same Bayes factor cut ($\ln K_1^2 = 5$) for model selection of detecting two components in white. The strongest evidence for the secondary component is found toward the eastern side of the main group of filaments, but isolated traces of secondary components can be found throughout the NGC 1333 region, beyond the regions where multiple components have been fit by hand in previous studies.

4.1. Comparisons with Previous Work

In Figure 2 we plot the maximum likelihood estimate (MLE) values of the centroid velocities at each pixel passing the decision boundary criterion above. The apparent lack of scatter in nearby pixels' values, even in the regions with faint emission (beyond the DR1 detection boundary), illustrates the relative ease in constraining the cloud kinematics.

The original GAS-DR1 paper (Friesen et al. 2017) shows the velocity dispersion distribution for all pixels in the regions studied. The distribution is composed of a substantial number




Figure 2. Map of MLEs for the centroid velocities V_{lsr} , overlaid with $K_0^1 = 1$ contours. The extended spatial region (relative to $K_0^1 \approx 100$ contours in Figure 1) shows the relative ease of constraining the centroid velocity parameter.

of pixels with a narrow (subsonic) velocity dispersion and a wide distribution of broad velocity dispersion (corresponding to the supersonic value). However, the degree to which the spectral multiplicity affected this broad velocity dispersion population was unclear. In the top panel of Figure 3 we plot the same KDE⁶ estimated velocity dispersion distributions with and without including the pixels where multiple components were found, and it shows that the DR1 reported values were biased toward larger values in places where we resolve two components. Meanwhile, the bottom panel of Figure 3 compares the velocity dispersion distributions from a common region where only one component is present. This shows good correspondence between the DR1 results and this study (Pearson’s $r = 0.95$). Spectral multiplicity of $\text{N}_2\text{H}^+ (1-0)$ has previously been reported in NGC 1333 (see Figure 12 in Hacar et al. 2017). While for all the regions reported to have a secondary component in Hacar et al. we report at least one significant spectrum with a secondary component, our method recovers more regions with two components. We attribute the discrepancy to a mixture of difference in critical densities of the two transitions, unequal signal-to-noise coverage of the two data sets, and method-specific differences. Future studies that would perform nested sampling analysis on Hacar et al. data should be able to pinpoint the reason for this apparent inconsistency.

⁶ SciPy implementation of Gaussian kernels with a bandwidth determined using Scott’s rule.

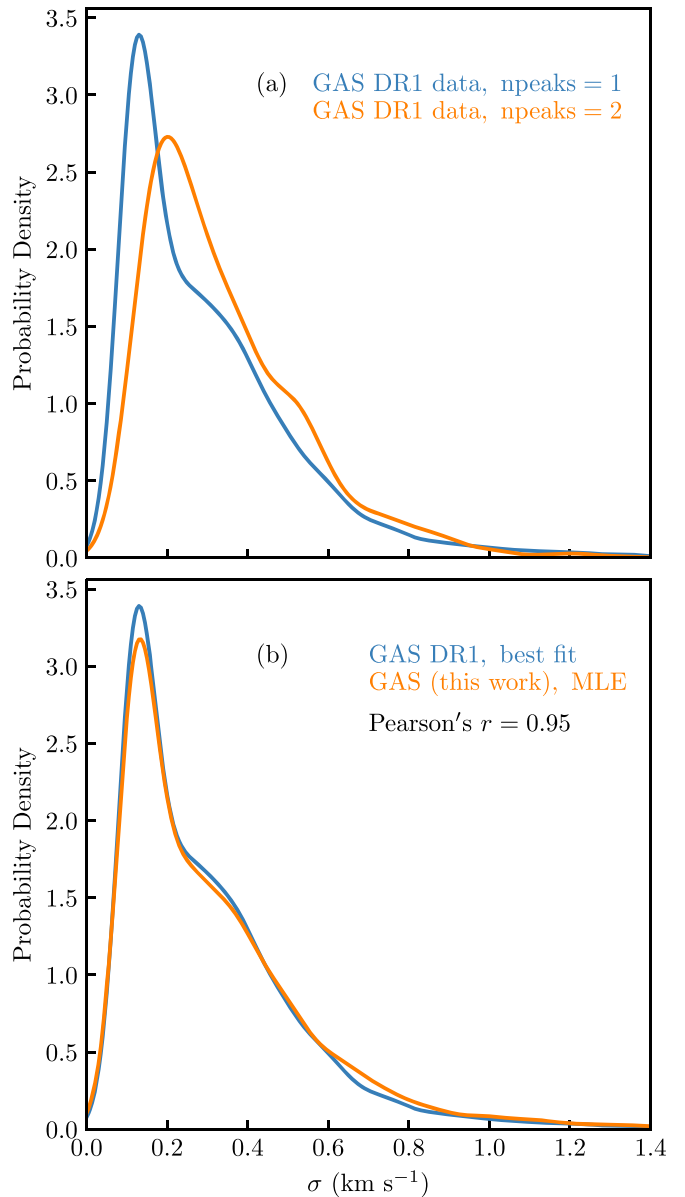


Figure 3. (a) The distribution of the GAS-DR1 best-fit velocity dispersion values, colored in blue and orange where this study finds one and two line-of-sight components, respectively. (b) Best-fit velocity dispersion distributions in GAS-DR1 result in blue and the corresponding MLE values from this work are shown in orange.

4.2. Heuristical Decisions on Spectral Multiplicity

The probabilistic approach in Bayesian model comparison does not implicitly require us to make a decision on the number of line-of-sight components, but, while K_j^i values alone can be used as a measure of model comparison, it is sometimes necessary to make an explicit decision for the preferred model. We adopt a decision threshold of $\ln K_j^i = 5$, with higher K_j^i -values indicating the \mathcal{M}_i model is considered to be true. The chosen threshold roughly corresponds to a “decisive” evidence strength on the Jeffreys’s scale (Jeffreys 1939). To further justify our decision threshold, we have simulated a 64×64 synthetic spectral cube containing only the white-noise component, and ran the same inference routine on it as on the actual GAS data. In total, 4096 spectra were generated with the noise amplitude randomly sampled from the DR1 rms

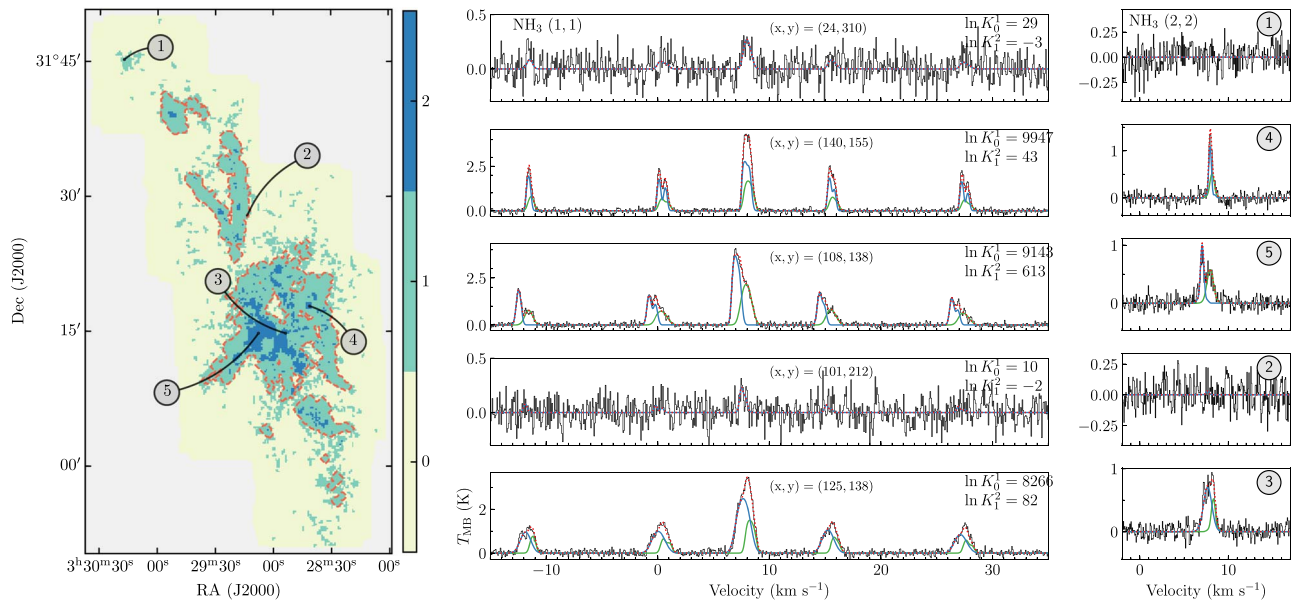


Figure 4. Selected spectra and the spectral multiplicity map. The map of spectral multiplicity (left) was constructed with a $\ln K_0^j = 5$ threshold for selecting the more complex model. The numbered annotations point to selected spectra displayed on the right side. The NH_3 (1, 1) and (2, 2) spectra are shown in the middle and right columns, respectively. The spectra are overlaid with their best MLE fits—solid colors for individual components, dotted red line for their total.

maps. We have found no false positives to be generated, with the maximum $\ln K_0^1$ value of 3.16 and 99.9th percentile of 2.13. Furthermore, a nested sampling analysis of M. Chen et al. (2020, in preparation) on a control sample of 10,000 two-component synthetic spectra, generated to resemble the GAS data, finds no cases where two components are identified where only one was generated, and 12% of cases where ones with two components were misidentified as one. We note that the 12% fraction is dependent on the signal-to-noise ratio of the components and is dropping to zero when their separation becomes sufficiently large (M. Chen et al. 2020, in preparation). Following the approach above, we derive a map of spectral multiplicity for GAS results on NGC 1333.

Illustrating both the parameter estimation and model selection, Figure 4 shows the spectral multiplicity with selected spectra overlaid with their best-fit (MLE) profiles. The Bayes factors corresponding to the detection of a single spectrum (K_0^1) and detection of a double component (K_1^2) are annotated alongside each spectrum. While the full Bayes factor and MLE maps are available in DataVerse at doi:10.7910/DVN/PDH6VT, we briefly describe a few representative spectra below. Ammonia spectra labeled as (1) are taken from a position where only one faint line component can be constrained, but no detection was reported in Friesen et al. (2017). Note that the odds for finding a secondary component (K_1^2) are against the two-component model. The spectral lines denoted as (2) lie at the detection limit of the GAS DR1 results, but are significantly detected as one component ($\ln K_0^1 = 10$) in our results. The map pointers (3) and (5) demonstrate a confident detection of secondary line components. Finally, a secondary component that would have normally missed the authors’ by-eye inspection yet is unambiguously present with a high certainty ($\ln K_1^2 = 43$) is labeled as (3).

4.3. Limitations and Future Prospects

Despite the relative ease of the setup compared to conventional methods of modeling additional spectral

components, the framework presented here must not be viewed as without limitations. First and foremost, we emphasize that the choice of priors must be educated, and future empirical studies should be undertaken to fully validate the decision boundaries and the impact of prior volume on them. Nevertheless, we see the Bayesian model selection as an improvement over the state-of-the-art heuristical quality control and by-eye human decisions for determining spectral multiplicity. Second, the model selection framework we follow is only meaningful if the true model is included among the competing ones. As the only alternative to astrophysical ammonia emission included in the model space is Gaussian noise, the non-LTE effects, imaging artifacts, skewed spectral baselines, and other features not described by the emission models will often manifest as additional velocity components. To illustrate this, we have sampled a three-component model on all five spectra shown in Figure 4. We find that a third component is ruled out ($\ln K_2^3 < 2$) in all the spectra, except for (3), where the three-component model is deemed the best ($\ln K_2^3 = 10.2$; $\ln K_3^4 = -9.0$). We note that at 5 km s⁻¹, this extra feature is unlikely to be physical as its velocity is atypical for the cloud kinematics in the central NGC 1333 region (see Figure 2). Third, in the analysis above we operate under the assumption that the nearby spectra are not spatially aware. We expect that an improved framework, where both the correlation between nearby spectra and the continuity of the physical properties are taken into the account, would greatly improve the performance of the Bayesian inference on spectral cube data. Finally, the framework presented here is more limited by the computational time available than other approaches we compared it to. While sampling a moderately large spectral cube for up to two ammonia components is certainly feasible on a small cluster, sampling highly dimensional models (e.g., 15 or more free parameters) for large Atacama Large Millimeter/submillimeter Array spectral cubes would require considerably more computational power.

V.S. is grateful to Sebastian Grandis, Suhail Dhawan, and Linda Baronchelli for useful discussions on Bayesian statistics. The authors would like to thank Brian Svoboda for his insightful comments during the preparation of this manuscript.

Facility: GBT.

Software: MultiNest (Feroz & Hobson 2008; Feroz et al. 2009), PyMultiNest (Buchner et al. 2014), scipy (Jones et al. 2001), astropy (Astropy Collaboration et al. 2013, 2018), pypspeckit (Ginsburg & Mirocha 2011), APLpy (Robitaille & Bressert 2012).

ORCID iDs

Vlas Sokolov  <https://orcid.org/0000-0002-5327-4289>
 Jaime E. Pineda  <https://orcid.org/0000-0002-3972-1978>
 Johannes Buchner  <https://orcid.org/0000-0003-0426-6634>
 Paola Caselli  <https://orcid.org/0000-0003-1481-7911>

References

- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Barnes, A. T., Henshaw, J. D., Caselli, P., et al. 2018, *MNRAS*, **475**, 5268
- Baronchelli, L., Nandra, K., & Buchner, J. 2018, *MNRAS*, **480**, 2377
- Bernardi, G., Zwart, J. T. L., Price, D., et al. 2016, *MNRAS*, **461**, 2847
- Beuther, H., Ragan, S. E., Johnston, K., et al. 2015, *A&A*, **584**, A67
- Buchner, J. 2017, arXiv:1707.04476
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
- Busquet, G., Zhang, Q., Palau, A., et al. 2013, *ApJL*, **764**, L26
- Chen, M. C.-Y., Di Francesco, J., Rosolowsky, E., et al. 2020, arXiv:2002.11736
- Clarke, S. D., Whitworth, A. P., Spowage, R. L., et al. 2018, *MNRAS*, **479**, 1722
- Feldmeier-Krause, A., Kerzendorf, W., Neumayer, N., et al. 2017, *MNRAS*, **464**, 194
- Feroz, F., & Hobson, M. P. 2008, *MNRAS*, **384**, 449
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
- Friesen, R. K., Pineda, J. E. & co-PIs et al. 2017, *ApJ*, **843**, 63
- Ginsburg, A., & Mirocha, J. 2011, PySpecKit: Python Spectroscopic Toolkit, version 0.1.21, Astrophysics Source Code Library, ascl:1109.001
- Hacar, A., Kainulainen, J., Tafalla, M., Beuther, H., & Alves, J. 2016, *A&A*, **587**, A97
- Hacar, A., & Tafalla, M. 2011, *A&A*, **533**, A34
- Hacar, A., Tafalla, M., & Alves, J. 2017, *A&A*, **606**, A123
- Hacar, A., Tafalla, M., Forbrich, J., et al. 2018, *A&A*, **610**, A77
- Hacar, A., Tafalla, M., Kauffmann, J., & Kovács, A. 2013, *A&A*, **554**, A55
- Henshaw, J. D., Caselli, P., Fontani, F., et al. 2013, *MNRAS*, **428**, 3425
- Henshaw, J. D., Longmore, S. N., Kruijssen, J. M. D., et al. 2016, *MNRAS*, **457**, 2675
- Jeffreys, H. 1939, The Theory of Probability, Oxford Classic Texts in the Physical Sciences (Oxford: Oxford Univ. Press) <https://books.google.de/books?id=vh9Act9rtzQC>
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, version 1.4.1 <http://www.scipy.org/>
- Koch, E. W., Rosolowsky, E. W., Lockman, F. J., et al. 2018, *MNRAS*, **479**, 2505
- Lavie, B., Mendonça, J. M., Mordasini, C., et al. 2017, *AJ*, **154**, 91
- Monsch, K., Pineda, J. E., Liu, H. B., et al. 2018, *ApJ*, **861**, 77
- Pineda, J. E., Goodman, A. A., Arce, H. G., et al. 2010, *ApJL*, **712**, L116
- Pineda, J. E., Goodman, A. A., Arce, H. G., et al. 2011, *ApJL*, **739**, L2
- Pineda, J. E., Offner, S. S. R., Parker, R. J., et al. 2015, *Natur*, **518**, 213
- Pon, A., Johnstone, D., Caselli, P., et al. 2016, *A&A*, **587**, A96
- Protassov, R., van Dyk, D. A., Connors, A., Kashyap, V. L., & Siemiginowska, A. 2002, *ApJ*, **571**, 545
- Riener, M., Kainulainen, J., Henshaw, J. D., et al. 2019, *A&A*, **628**, A78
- Robitaille, T., & Bressert, E. 2012, APLpy: Astronomical Plotting Library in Python, version 1.1.1, Astrophysics Source Code Library, ascl:1208.017
- Rosolowsky, E. W., Pineda, J. E., Foster, J. B., et al. 2008, *ApJS*, **175**, 509
- Skilling, J. 2004, in AIP Conf. Ser. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. R. Fischer, R. Preuss, & U. V. Toussaint (Melville, NY: AIP), 395
- Tanaka, T., Nakamura, F., Awazu, Y., et al. 2013, *ApJ*, **778**, 34