



CHALMERS
UNIVERSITY OF TECHNOLOGY

Neural Network Constraints on the Cosmic-Ray Ionization Rate and Other Physical Conditions in NGC 253 with ALCHEMI Measurements of HCN

Downloaded from: <https://research.chalmers.se>, 2024-12-20 20:34 UTC

Citation for the original published paper (version of record):

Behrens, E., Mangum, J., Viti, S. et al (2024). Neural Network Constraints on the Cosmic-Ray Ionization Rate and Other Physical Conditions in NGC 253 with ALCHEMI Measurements of HCN and HNC. *Astrophysical Journal*, 977(1).
<http://dx.doi.org/10.3847/1538-4357/ad85db>

N.B. When citing this work, cite the original published paper.



Neural Network Constraints on the Cosmic-Ray Ionization Rate and Other Physical Conditions in NGC 253 with ALCHEMI Measurements of HCN and HNC

Erica Behrens^{1,2} , Jeffrey G. Mangum² , Serena Viti^{3,4,5} , Jonathan Holdship^{3,5} , Ko-Yun Huang³ , Mathilde Bouvier³ , Joshua Butterworth³ , Cosima Eibensteiner² , Nanase Harada^{6,7} , Sergio Martín^{8,9} , Kazushi Sakamoto¹⁰ , Sebastien Muller¹¹ , Kunihiko Tanaka¹² , Laura Colzi¹³ , Christian Henkel^{14,15} , David S. Meier^{16,17} , Víctor M. Rivilla¹³ , and Paul P. van der Werf³

(ALMA Comprehensive High-resolution Extragalactic Molecular Inventory (ALCHEMI) collaboration)

¹ Department of Astronomy, University of Virginia, P.O. Box 400325, 530 McCormick Road, Charlottesville, VA 22904-4325, USA; eb7he@virginia.edu

² National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903-2475, USA

³ Leiden Observatory, Leiden University, P.O. Box 9513, NL-2300 RA Leiden, The Netherlands

⁴ Argelander Institut für Astronomie der Universität Bonn, Auf dem Hügel 71, 53121 Bonn, Germany

⁵ Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

⁶ National Astronomical Observatory of Japan, 2-21-1 Osawa, Mitaka, Tokyo 181-8588, Japan

⁷ Department of Astronomy, School of Science, The Graduate University of Advanced Studies (SOKENDAI), 2-21-1 Osawa, Mitaka, Tokyo, 181-1855, Japan

⁸ European Southern Observatory, Alonso de Córdova, 3107, Vitacura, Santiago 763-0355, Chile

⁹ Joint ALMA Observatory, Alonso de Córdova, 3107, Vitacura, Santiago 763-0355, Chile

¹⁰ Institute of Astronomy and Astrophysics, Academia Sinica, 11F of AS/NTU Astronomy-Mathematics Building, No.1, Sec. 4, Roosevelt Rd, Taipei 10617, Taiwan

¹¹ Department of Space, Earth and Environment, Chalmers University of Technology, Onsala Space Observatory, SE-43992 Onsala, Sweden

¹² Department of Physics, Faculty of Science and Technology, Keio University, 3-14-1 Hiyoshi, Yokohama, Kanagawa 223-8522, Japan

¹³ Centro de Astrobiología (CAB), CSIC-INTA, Ctra. de Ajalvir km 4, E28850, Torrejón de Ardoz, Spain

¹⁴ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

¹⁵ Xinjiang Astronomical Observatory, Chinese Academy of Sciences, 830011, Urumqi, People's Republic of China

¹⁶ New Mexico Institute of Mining and Technology, 801 Leroy Place, Socorro, NM 87801, USA

¹⁷ National Radio Astronomy Observatory, PO Box O, 1003 Lopezville Road, Socorro, NM 87801, USA

Received 2024 August 14; revised 2024 September 19; accepted 2024 September 20; published 2024 November 29

Abstract

We use a neural network model and Atacama Large Millimeter/submillimeter Array (ALMA) observations of HCN and HNC to constrain the physical conditions, most notably the cosmic-ray ionization rate (CRIR, ζ), in the Central Molecular Zone (CMZ) of the starburst galaxy NGC 253. Using output from the chemical code UCLCHEM, we train a neural network model to emulate UCLCHEM and derive HCN and HNC molecular abundances from a given set of physical conditions. We combine the neural network with radiative transfer modeling to generate modeled integrated intensities, which we compare to measurements of HCN and HNC from the ALMA Large Program ALCHEMI. Using a Bayesian nested sampling framework, we constrain the CRIR, molecular gas volume and column densities, kinetic temperature, and beam-filling factor across NGC 253's CMZ. The neural network model successfully recovers UCLCHEM molecular abundances with $\sim 3\%$ error and, when used with our Bayesian inference algorithm, increases the parameter-inference speed tenfold. We create images of these physical parameters across NGC 253's CMZ at 50 pc resolution and find that the CRIR, in addition to the other gas parameters, is spatially variable with $\zeta \sim$ a few $\times 10^{-14} \text{ s}^{-1}$ at $r \gtrsim 100$ pc from the nucleus, increasing to $\zeta > 10^{-13} \text{ s}^{-1}$ at its center. These inferred CRIRs are consistent within 1 dex with theoretical predictions based on nonthermal emission. Additionally, the high CRIRs estimated in NGC 253's CMZ can be explained by the large number of cosmic-ray-producing sources as well as a potential suppression of cosmic-ray diffusion near their injection sites.

Unified Astronomy Thesaurus concepts: Cosmic ray astronomy (324); Interstellar medium (847); Molecular gas (1073); Neural networks (1933); Starburst galaxies (1570); Star forming regions (1565); Astrochemistry (75); Galaxy nuclei (609); Stellar feedback (1602)

Materials only available in the online version of record: machine-readable table

1. Introduction

Understanding the physical processes associated with star formation is pivotal for identifying how star formation proceeds in different environments. Star-forming regions vary in many ways, including, but not limited to, their temperatures, densities,

turbulence, radiation strength, and molecular complexity (e.g., D. Downes & P. M. Solomon 1998; P. Padoan & A. Nordlund 2002; M. H. Heyer & C. M. Brunt 2004; M. Gong et al. 2020). All of these factors combine to determine the nature and lifetime of star-forming regions and the galaxies that host them (see, for example, E. Corbelli et al. 2017; M. Chevance et al. 2020; J. Kim et al. 2021; V. A. Semenov et al. 2021). Starburst galaxies feature some of the most extreme examples of star-forming environments, demonstrated by the high surface brightness of their emission as well as enhanced star formation rates (e.g.,

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

E. Schinnerer et al. 2007; A. K. Leroy et al. 2013; E. J. Murphy et al. 2015; N. Brunetti et al. 2021; D. Callanan et al. 2021; C. Eibensteiner et al. 2022). Naturally, we expect that the physical processes occurring in starburst galaxies must differ from those in more quiescent regions, such as in the Milky Way, to account for their increased star formation output.

One important avenue of study is the feedback processes that occur as a result of star formation activity while also affecting the gas and dust that will form future generations of stars. Recently formed stars produce large quantities of energy in the form of ultraviolet (UV) and X-ray radiation (S. Lepp & A. Dalgarno 1996; D. J. Hollenbach & A. G. G. M. Tielens 1999), turbulence from stellar winds and gravitational collapse (e.g., A. F. Loenen et al. 2008), and ionizing cosmic rays (F. Acero et al. 2009; P. P. Papadopoulos 2010; E. Bayet et al. 2011; M. Padovani et al. 2022). All of these mechanisms contribute energy and pressure to the surrounding medium that can have a marked effect on the molecular gas chemistry (J. Holdship et al. 2021; E. Behrens et al. 2022; J. Holdship et al. 2022; K. Y. Huang et al. 2023). However, the relative dominance of each of these feedback mechanisms can vary based on environment (e.g., E. Cenci et al. 2024; E. Schinnerer & A. K. Leroy 2024), so it is important to understand the role of each process in different star-forming contexts.

NGC 253 is one of the nearest ($d \sim 3.5 \pm 0.2$ Mpc; R. Rekola et al. 2005) starburst galaxies; thus, it has been the subject of many studies investigating its intense and complex star formation (e.g., K. Sakamoto et al. 2011; M. J. F. Rosenberg et al. 2013; A. K. Leroy et al. 2015, 2018; J. G. Mangum et al. 2019; N. Krieger et al. 2020; R. C. Levy et al. 2022). Though only featuring a modest total star formation rate (SFR) of $\sim 5 M_{\odot} \text{ yr}^{-1}$, nearly half of the galaxy's star formation is occurring in its central kiloparsec, making NGC 253's Central Molecular Zone (CMZ) a very active star-forming environment (A. K. Leroy et al. 2015). For comparison, the SFR in NGC 253's CMZ is larger than the total SFR across the entire Milky Way galaxy ($1.65\text{--}1.9 M_{\odot} \text{ yr}^{-1}$; L. Chomiuk & M. S. Povich 2011; T. C. Licquia & J. A. Newman 2015). Recent observations and analyses from the Atacama Large Millimeter/submillimeter Array (ALMA) Comprehensive High-resolution Extragalactic Molecular Inventory (ALCHEMI) Large Program (S. Martín et al. 2021) have revealed a wealth of molecular complexity, making it possible to study shocks, heating processes, masers, and molecular distributions at ~ 28 pc resolution (N. Harada et al. 2021; J. Holdship et al. 2021; E. Behrens et al. 2022; D. Haasler et al. 2022; N. Harada et al. 2022; J. Holdship et al. 2022; P. K. Humire et al. 2022; K. Y. Huang et al. 2023; M. Bao et al. 2024; M. Bouvier et al. 2024; J. Butterworth et al. 2024; N. Harada et al. 2024; K. Tanaka et al. 2024).

Previous studies by J. Holdship et al. (2021, 2022) and E. Behrens et al. (2022) have investigated the heating processes associated with the intense star formation occurring in NGC 253's nucleus. Their results point to high cosmic-ray ionization rates (CRIRs; $\sim 10^{-14}\text{--}10^{-12} \text{ s}^{-1}$) across the NGC 253 CMZ, as well as show a spatially varying CRIR distribution with higher ionization rates in the center that decrease toward the edge of the CMZ. For comparison, S. Ravikularaman et al. (2024) found a CRIR of $2 \times 10^{-14} \text{ s}^{-1}$ in the Milky Way CMZ, which is 2–3 orders of magnitude larger than the CRIR in other parts of the Galaxy. J. Holdship et al. (2021, 2022) and E. Behrens et al. (2022) also report that these high CRIRs are the main

contributors to the high kinetic temperatures (≥ 300 K) measured by J. G. Mangum et al. (2019), making them the dominant heating mechanism in the CMZ. Additionally, E. Behrens et al. (2022) found a positive correlation between the CRIR and the location and number of radio continuum sources (primarily H II regions and supernova remnants) from recent star formation. However, due to computational limitations, these studies limited their focus to 5 (J. Holdship et al. 2022) and 10 (E. Behrens et al. 2022) Giant Molecular Cloud (GMC) regions in the CMZ that were identified by A. K. Leroy et al. (2015) based on peaks in the dense gas emission. Thus, these investigations were unable to fully spatially sample the CRIR and other physical conditions across the entire CMZ.

Chemical modeling codes, such as UCLCHEM (J. Holdship et al. 2017),¹⁸ which was used in J. Holdship et al. (2022) and E. Behrens et al. (2022), are powerful tools for understanding the chemistry of star-forming regions. By providing chemical models with sets of physical conditions, such as kinetic temperature, volume density, and CRIR, we can predict the abundances with respect to total hydrogen of various species within a region of interest. For instance, E. Behrens et al. (2022) found that the abundances of HCN and HNC are greatly influenced by cosmic-ray chemistry in the UV-shielded environments of NGC 253. While previous studies suggested that HCN and HNC were sensitive to changes in kinetic temperature for $T_{\text{K}} \lesssim 50$ K (L. Colzi et al. 2018; A. Hacar et al. 2020), E. Behrens et al. (2022) found that, at high cosmic-ray ionization rates, the formation of HCN and HNC is dominated by reactions of HCNH^+ with electrons. Since both HCN and HNC are destroyed at relatively equal rates via reactions with ions, their cosmic-ray-influenced formation is primarily responsible for their relative abundances (for more details on the HCN and HNC chemistry, see E. Behrens et al. 2022).

When we combine information derived from these chemical pathways with radiative transfer calculations, we can directly compare model output to measured molecular intensities, thereby allowing us to constrain the gas conditions in specific star-forming regions and make inferences about the physical processes occurring there. However, chemical models are computationally expensive and time-consuming. Previous parameter-inference studies in the NGC 253 CMZ were able to provide strong constraints on the conditions in only a handful of regions. In order to perform a comprehensive study of the gas conditions across NGC 253's CMZ, we must construct a method that allows us to thoroughly sample the CMZ's spatially varying chemistry in a reasonable time frame.

In this paper, we present a neural network that can emulate and adopt the role of a chemical model in our physical parameter-inference algorithm, allowing us to ascertain the gas conditions in a given region 10 times faster than with the chemical model alone. This new implementation to our algorithm allows us to infer the physical properties across the entire NGC 253 CMZ at 50 pc resolution and will pave the way toward extending this analysis to other galaxies, as well as images with higher resolution or larger fields of view. Ultimately, the ability to constrain molecular gas conditions and make inferences about the physical processes associated with starburst activity will help us better understand the relationship between star formation and galaxy evolution in starburst galaxies.

¹⁸ <https://uclchem.github.io/>

Table 1
Example Average Integrated Intensity and rms Values for Regions in the Outer (Region 26) and Inner (Region 53) CMZ

Transition	Rest Frequency (GHz)	E_u (K)	Region 26		Region 53		
			$\langle \int S_i dv \rangle$ (Jy km s ⁻¹ beam ⁻¹)	rms (Jy km s ⁻¹ beam ⁻¹)	$\langle \int S_i dv \rangle$ (Jy km s ⁻¹ beam ⁻¹)	rms (Jy km s ⁻¹ beam ⁻¹)	
HCN	1–0	88.6316	4.25	4.7	0.7	8.8	1.3
	2–1	177.2611	12.76	13.3	2.0	33.5	5.0
	3–2	265.8864	25.52	13.9	2.1	56.1	8.4
	4–3	354.5055	42.53	11.0	1.7	68.6	10.3
HNC	1–0	90.6636	4.35	2.2	0.3	7.9	1.2
	2–1	181.3248	13.05	2.2	0.7	19.9	3.0
	3–2	271.9811	26.11	4.3	0.7	37.9	5.7
	4–3	362.6303	43.51	2.4	0.4	40.9	6.1

(This table is available in its entirety in machine-readable form in the [online article](#).)

We present our data and identification of regions of interest in Section 2. In Section 3, we describe our neural network architecture and its implementation into our Bayesian nested sampling algorithm. We summarize our results in Section 4 and discuss the implications of our work on star formation and galaxy evolution in Section 5.

2. Observational Data

2.1. ALCHEMI Observations

We use data from the ALMA Large Program ALCHEMI: the ALMA Comprehensive High-resolution Extragalactic Molecular Inventory (S. Martín et al. 2021), which imaged the NGC 253 CMZ in ALMA Bands 3–7 (84–373 GHz) at 1.^{''}6 (~28 pc) resolution during ALMA Cycles 5 (2017.1.00161.L) and 6 (2018.1.00162.S). Details of the observations and data reduction can be found in S. Martín et al. (2021) and E. Behrens et al. (2022). As noted in E. Behrens et al. (2022), we extract the HCN and HNC $J=1-0$, 2–1, 3–2, and 4–3 transition integrated intensities using the CubeLineMoment code.¹⁹ The frequencies and upper-state energies for our transitions of interest are shown in Table 1.

2.2. Region Selection and Moment 0 Maps

In order to sample the entire nuclear region in NGC 253, we divide the CMZ into adjoining hexagonal regions 50 pc in height using the integrated emission averaging code HERA (E. Behrens 2024), the HEXagonal Region Averager code.²⁰ This code divides integrated intensity maps for each transition of interest into either adjacent circular or adjoining hexagonal regions of a user-specified size, averages the integrated emission over the region, and calculates noise statistics associated with the averaged emission. We choose a common region size of 50 pc, which is the typical size of a GMC structure in the NGC 253 CMZ (A. K. Leroy et al. 2015). The code-generated regions are pruned such that only regions that meet a specified signal-to-noise ratio (S/N) threshold (in our case, $S/N \geq 3$) are retained. This information is written to a file from which we can extract and analyze it. Figure 1 shows the final 94 regions we selected to study, where each of the hexagonal regions has $S/N \geq 3$ in at least one transition of interest.

However, unlike the regions studied in J. Holdship et al. (2022) and E. Behrens et al. (2022), not all hexagonal regions

contain significant emission for every studied transition (see the bottom panel in Figure 1). In some cases, a portion of the image pixels within a given hexagon are masked due to low S/N. In other cases, entire hexagonal regions are masked. Rather than removing regions that do not feature a full set of eight detected transitions for HCN and HNC or, alternatively, constraining models of those regions with fewer than eight transitions, we fill in low S/N image pixels with 3σ upper-limit integrated intensities. These limits allow us to constrain our models with an equal number of transitions per region while also testing the efficacy of this algorithm in regions with low levels of emission. We derive these 3σ limits using our single-channel rms σ_{noise} , flux-calibration uncertainty σ_{fluxcal} , and an assumed line width Δv . Based on an inspection of the range of line widths seen in the CMZ for HCN and HNC near the edge of our S/N threshold, we assume a random Δv between 20 and 100 km s⁻¹ for each low S/N image pixel. We integrate over a Gaussian with width $\Delta v / \sqrt{8 \ln 2}$ and a peak value equal to 3 times the noise value, or $3 \times \sqrt{\sigma_{\text{noise}}^2 + \sigma_{\text{fluxcal}}^2}$, where σ_{fluxcal} is equal to 15% of the measured flux. Using these upper-limit values in place of masked pixels, we then obtain a unique average integrated intensity and rms for each hexagonal region, which we use to constrain our physical and chemical models, as discussed in Section 3. Average integrated intensity measurements and rms values for example outer and inner CMZ regions (regions 26 and 53, respectively) are shown in Table 1. Average integrated intensities and the associated rms values for each of the 94 regions is provided as a machine-readable table.

3. Bayesian Inference with Neural Networks

To constrain the gas properties in each of our designated regions, we consider the chemical pathways and molecular emission resulting from a given set of input gas conditions. Previous studies (E. Behrens et al. 2022; J. Holdship et al. 2022; K. Y. Huang et al. 2023) used the chemical modeling code UCLCHEM (J. Holdship et al. 2017)²¹ to calculate molecular abundances given a set of input conditions. Using the radiative transfer code SpectralRadex,²² these studies associated modeled abundances, in combination with the input gas conditions, with modeled integrated intensities for each transition of interest, which could be directly compared to observations. The aforementioned studies successfully

¹⁹ <https://github.com/keflavich/cube-line-extractor>

²⁰ <https://github.com/ebehrens97/HERA>

²¹ <https://uclchem.github.io/>

²² <https://spectralradex.readthedocs.io/en/latest/>

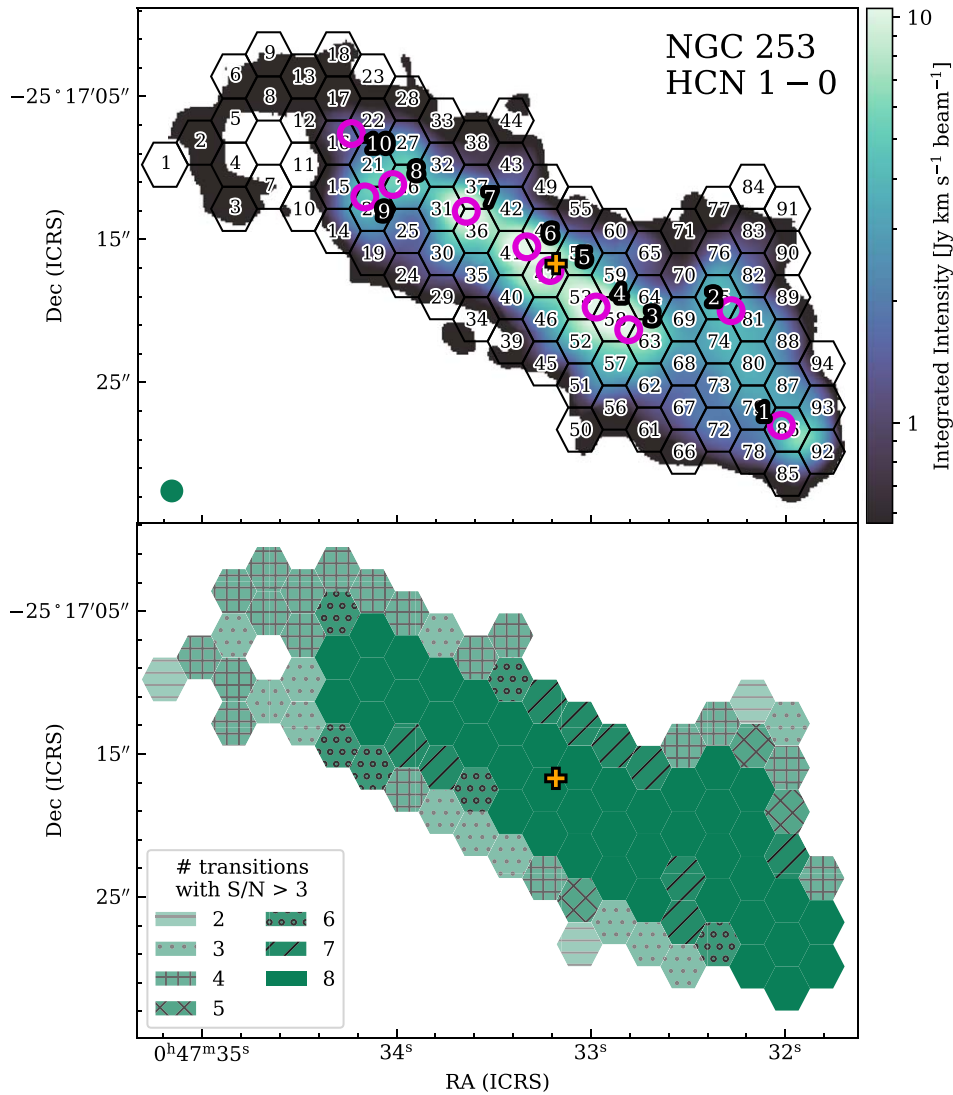


Figure 1. Top: map of NGC 253 HCN 1–0 integrated emission with $S/N > 3$ overlaid with the regions analyzed in this article. Purple circles indicate locations of GMCs studied in J. Holdship et al. (2022) and E. Behrens et al. (2022); note J. Holdship et al. (2022) only investigated GMCs 3–7). The circle in the bottom-left corner indicates the size of the $1.''6$ (~ 28 pc) ALCHEMI beam. Orange crosses in both panels indicate the location of the galaxy’s kinematic center (J. L. Turner & P. T. P. Ho 1985). Bottom: map of the hexagonal regions used in this analysis with hatching to indicate how many total HCN and HNC transitions have been detected with $S/N > 3$.

demonstrated the validity of these methods by estimating certain gas properties, such as volume density and CRIR, for a small number of regions (5–12, from A. K. Leroy et al. 2015; K. Y. Huang et al. 2023) in the NGC 253 CMZ. However, the use of a chemical modeling code significantly prolongs this process, making this strategy untenable for modeling a larger number of regions. We address this issue by developing a neural network model that can replace and fulfill the role of a chemical model in our Bayesian inference algorithm (see Figure 2). Further details regarding our Bayesian inference analysis are described in J. Holdship et al. (2022) and E. Behrens et al. (2022).

3.1. Neural Network Architecture

We use the Python package TensorFlow²³ to build and train our neural network model. We find that a neural network with fully connected layers and 12,020,002 trainable

parameters produces a model that best reproduces the UCLCHEM abundances (note that these trainable parameters—neural network weights and biases; see Appendix A—are separate from the physical parameters listed in Table 2). We start with an input layer of four nodes, where each node corresponds to one of our chemical model parameters (kinetic temperature T_K , H_2 volume density n , CRIR ζ , and H_2 column density N).²⁴ After values from the input layer are passed through the hidden layers, the values calculated in the neural network’s output layer represent the molecular abundance predictions of that epoch for each physical parameter combination in the training set. An epoch is defined as the time it takes for the neural network to be trained on the entire training set one time. See Appendix A for more details on the neural network training process.

²⁴ Note that the beam-filling factor is not an input for the neural network, as it is applied after the integrated intensities are calculated with SpectralRadex (see Section 3.4).

²³ <https://www.tensorflow.org/>

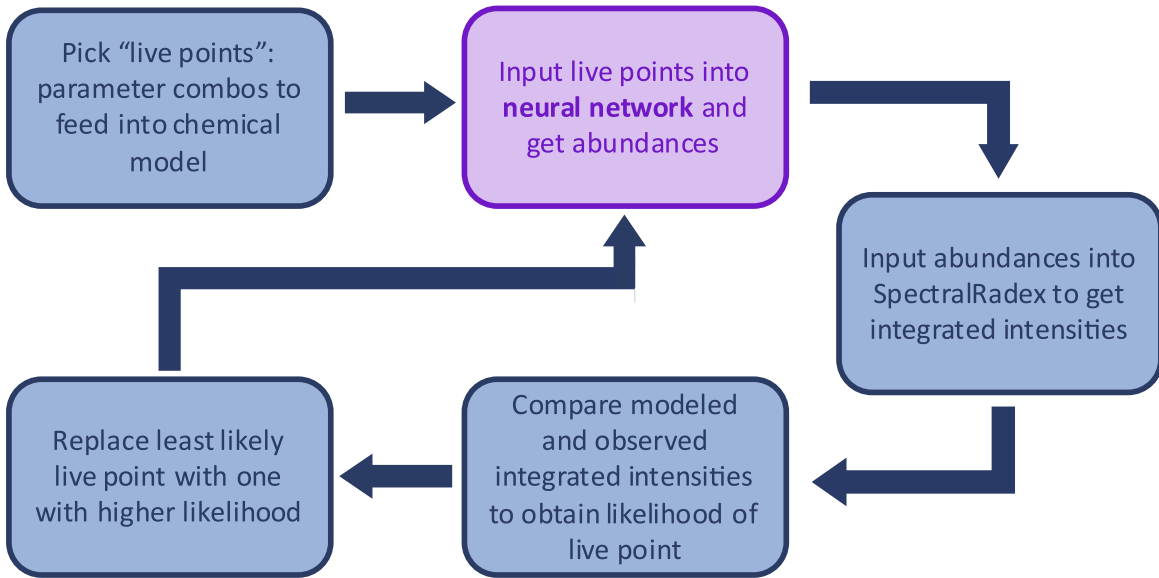


Figure 2. Flow chart describing our Bayesian nested sampling algorithm, where the highlighted purple rectangle indicates the neural network’s position in our modeling process.

Table 2
Prior Distributions and Training Data

Parameter	Range (Buffer Range)	Distribution Type	No. Points (No. with Buffer)	
T_K	Temperature (K)	50–300	Linear	15
n_{H_2}	Volume Density (cm^{-3})	10^3 – 10^7 ($10^{2.6}$ – $10^{7.4}$)	Log	20 (24)
ζ	Cosmic-ray Ionization Rate (ζ_0) ^a	10 – 10^7	Log	25
N_{H_2}	H_2 Column Density (cm^{-2})	10^{22} – 10^{25} ($10^{21.6}$ – $10^{25.4}$)	Log	15 (19)
η_{ff}	Beam-filling factor	0–1	Linear	N/A ^b

Notes.

^a $\zeta_0 = 1.36 \times 10^{-17} \text{ s}^{-1}$.

^b Parameter η_{ff} is not a neural network parameter, as it is applied after the intensities are calculated by the neural network and SpectralRadex.

Periodically throughout each training epoch, our training algorithm compares the node values for a given layer to the actual values provided by the training data using a loss function. We use a mean-squared-error loss function, and the calculated loss is used to tweak the neural network in order to improve its performance in future epochs. At the end of each epoch, the loss for a separate data set, the validation data, is calculated (see Section 3.2 for a breakdown of the training and validation sets). The neural network does not “see” the validation data until the end of an epoch, and thus the validation set is used to assess the neural network’s efficacy so the model can be improved for the next round of training. This process continues over several epochs until the neural network reaches its specified stopping criterion. In this case, we cease neural network training once the calculated loss has not improved over 20 consecutive epochs, and the neural network then reverts back to using the weights from the epoch with the lowest validation loss. We employ this criterion to avoid overfitting, which occurs when the model simply memorizes the abundances associated with each parameter combination without learning how to predict abundances from new, unseen data.

3.2. Training Data Set

We train one neural network model to predict the abundance of both HCN and HNC since previous ALCHEMI studies have

shown that these molecules have similar spatial distributions (see E. Behrens et al. 2022). We derive our training data set from a grid of UCLCHEM models for HCN and HNC using input parameters that span our parameter space (see Table 2). We use a two-phase UCLCHEM model structure²⁵ that takes only the assumed volume density as input in phase 1 and models the freefall collapse of a gas cloud with an initial density of 100 cm^{-3} down to the specified density at a fixed temperature of 10 K, transforming the gas from atomic/ionic to molecular form. This stage is followed by phase 2, which uses the calculated abundances inherited from phase 1 and four of our five input parameters (T_K , n_{H_2} , ζ , and N_{H_2}) to chemically model the cloud itself over a timescale of 6 Myr. The output from phase 2 makes up the training data for our neural network. We note that the abundances used at the beginning of stage 2 may be sensitive to the initial conditions (e.g., UV field, elemental abundances, mode of collapse)—varying these parameters would require reconstructing the training data set, which is not a case considered here. Here, we assume that we are probing the inner, shielded part of the cloud (visual extinction A_v of 10 mag) such that UV radiation (set to 1 Habing in both UCLCHEM phases) is negligible. X-ray ionization is not included in UCLCHEM, but other studies have

²⁵ Note that this two-phase approach differs from that used in E. Behrens et al. (2022), which only included the second phase of this modeling process. Further discussion on this subject can be found in Section 5.3.

shown that the ionization contribution from X-rays in the regions we study is likely several orders of magnitude lower than the ζ values we infer ($\zeta_{\text{x-ray}} \lesssim 10^{-17} \text{ s}^{-1}$ for $N(\text{H}_2) \gtrsim 10^{23} \text{ cm}^{-2}$; V. H. M. Phan et al. 2024), thus making X-ray ionization a negligible contributor to the heating budget within the NGC 253 CMZ. See J. Holdship et al. (2017) for more details on UCLCHEM and its capabilities. Table 2 shows the parameter space over which we run these models and also indicates the types of prior distributions we assume for each parameter, as well as the number of values we sample for the training set for each parameter within our desired range. Including buffer points, which will be discussed later, each initial data set contains $\sim 170,000$ parameter combinations.

Any parameter combinations that result in very low abundances are likely to be dominated by numerical error rather than chemistry (J. Heyl et al. 2023). Hence, we retain only parameter combinations that result in abundances for both HCN and HNC above 10^{-12} (a conservative value corresponding to typically the lowest observed molecular abundances), leaving $\sim 115,000$ data points.

We divide the UCLCHEM output into a training set (60% of data, $\sim 69,000$ points), validation set (15% of data, $\sim 17,000$ points), and a test set (25% of data, $\sim 28,000$ points). We use the parameter values in our training set as input for the neural network’s first layer (using the log of the parameter value for n_{H_2} , ζ , and N_{H_2}), and these values are then propagated through the neural network (see Appendix A for more details), ultimately resulting in molecular abundance predictions for both HCN and HNC. We use the log of the UCLCHEM fractional abundance values when training the neural network in order to increase the magnitude of differences between abundances. We calculate the mean-squared-error loss between the neural network result and the training data continuously throughout each epoch, but we calculate the validation loss, or loss derived from the validation data, only once at the end of each epoch (see Figure A2). Once we reach our stopping criterion (see Section 3.1), we test the newly trained neural network model on the test set, a subset of the UCLCHEM grid that was not included in the training process. We then compare the model’s predicted abundances to the actual abundances produced by UCLCHEM in order to assess its accuracy.

3.3. Training Results

Figure 3 shows the results from our neural network training process, including the results for the best model (panel (C)), which we later use in our Bayesian inference algorithm. We compare the neural network model-predicted values against those calculated with UCLCHEM and find that, on average, the neural network model is able to recover molecular abundances within 3% of the chemical model’s values.

However, Figure 3(a) shows there are outliers, making up $\sim 1\%$ of the test set where the neural network was unable to recover an abundance within 1 order of magnitude of the UCLCHEM value (green and black Xs in Figure 3(a)). Initial testing revealed that the neural network was unable to predict abundances for parameter combinations that included the minimum and maximum column and volume density values that we probed (10^{22} cm^{-2} , 10^{25} cm^{-2} and 10^3 cm^{-3} , 10^7 cm^{-3} , respectively). Excluding the edge points, all other column and volume density values within our parameter ranges are padded above and below by values of that parameter that the neural network can use to predict an abundance. However,

values on the edge of our parameter space are not padded by additional parameter values on each side, resulting in poorer abundance predictions for parameter combinations that contain these edge points. To ensure that the neural network is equally trained and sensitive to all values that fall within our desired parameter space, we include additional column and volume density buffer points outside of our desired parameter ranges (see Table 2) when training our neural network model, though values outside of this original range will not be considered later in our Bayesian analysis. We found no evidence that the neural network was poorly predicting abundances for other parameter combinations including temperature and CRIR values on the edges of our parameter space, so we use the original parameter ranges for T_{K} and ζ that are listed in Table 2.

Figure 3(b) shows the predictions of our neural network model when including the buffer points shown in Table 2 in our training set. These buffer points become the new edge points (green and black Xs in Figure 3(b)) in this expanded parameter space and are thus subject to the same scatter as the edge points in Figure 3(a). However, when plotting the results of this same model but removing these buffer points (Figure 3(c)), we see that the abundance predictions from parameters with values in our original parameter space have significantly less scatter. By including buffer points while training our neural network model, we ensure that our model can robustly predict abundances across our desired parameter space.

3.4. Implementation with Bayesian Nested Sampling

We implement our fully trained neural network model for HCN and HNC as a replacement for the chemical model’s role in our inference procedure. The original implementations of this algorithm (see E. Behrens et al. 2022; J. Holdship et al. 2022) used the nested sampling code UltraNest (J. Buchner 2016, 2019, 2021)²⁶ to feed parameter combinations into our chemical model, UCLCHEM, which would calculate and return abundance estimates for our molecules of interest. In our updated algorithm, our neural network model fills this role instead (see Figure 2). The HCN and HNC model predicts abundances that we then feed into our radiative transfer code, SpectralRadex, in order to obtain integrated intensities for these molecular species. Note that SpectralRadex does a radiative transfer calculation that accounts for optical depth. We use collisional excitation rates from M. Hernández Vera et al. (2017), who calculated rates for HCN and HNC collisions with ortho- and para- H_2 (see Appendix C for a detailed analysis of these new rates as compared to those used in E. Behrens et al. 2022). We adopt a 3:1 ortho:para hydrogen ratio. After the integrated intensities are calculated, we apply a beam-filling factor, which is a free parameter in our algorithm ranging from 0 to 1. We then compare these modeled integrated intensities F_r , which are calculated for some set of parameters θ , to our observed intensities F_d and their associated uncertainties σ_F using the following log-likelihood function:

$$P(F_d|\theta) = \exp\left(-\frac{1}{2}\sum_i^N \frac{(F_{d,i} - F_{r,i})^2}{\sigma_{F,i}^2}\right), \quad (1)$$

where we sum the values for each transition i .

²⁶ <https://johannesbuchner.github.io/UltraNest/>

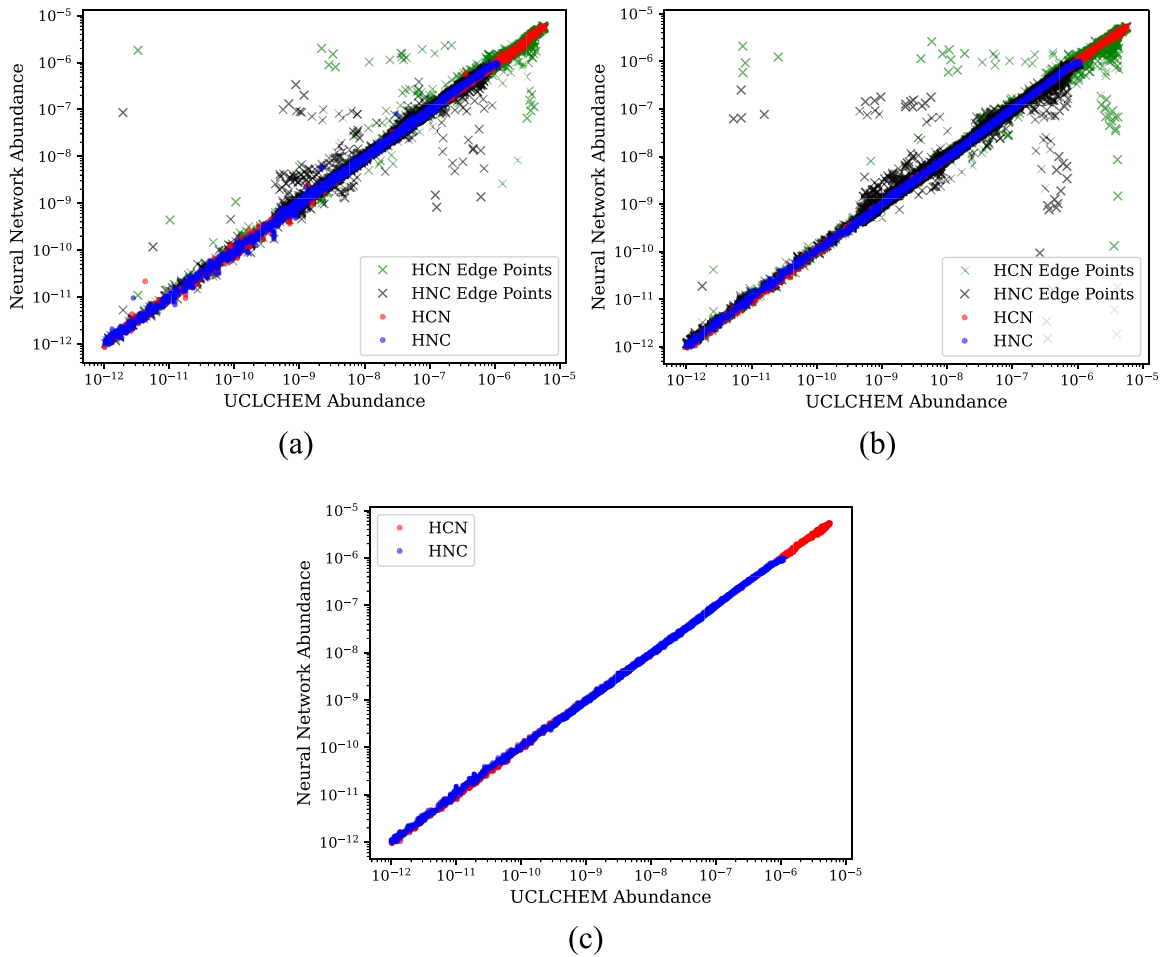


Figure 3. (a) Neural network predictions of test set parameter combinations where the model was trained using the original parameter ranges in Table 2 without buffer points. Predictions for parameter combinations that used maximum or minimum values for column or volume density are indicated with X's. (b) Same as (a), but neural network training set included additional buffer points for volume and column density. (c) Neural network predictions using the same model as in (b), but buffer points are not plotted.

UltraNest uses the results from each log-likelihood comparison to choose additional parameter combinations to sample within our designated parameter space (see Figure 2 and Table 2). This process concludes when UltraNest has sampled the majority of the probability density.

4. Results

4.1. Neural Network versus Chemical Model

We compare the nested sampling results using our neural network model to those obtained with UCLCHEM to confirm that the neural network can reliably replace the role of a chemical model in our nested sampling algorithm. Figure 4 compares the inference algorithm results using UCLCHEM and the HCN/HNC neural network for regions 26 and 53 (see Figure 1 for region numbers), which were chosen to represent regions in the outer and inner CMZ, respectively. The neural network clearly replicates the results we obtain using UCLCHEM. The final inferred parameter values and their uncertainties derived from UCLCHEM and our neural network are nearly identical. Additionally, both the 1D and 2D posterior distributions for the two methods feature very similar shapes and characteristics, demonstrating that the neural network can reliably reproduce UCLCHEM results in this context in a fraction of the time. We note that the 2D posteriors generated when

including the neural network model are somewhat smoother than those produced when our algorithm includes UCLCHEM. This difference may result from the fact that the neural network models are simply an approximation of the UCLCHEM calculations. This approximation may smooth over the functional relationships between model parameters and output abundances and could be further related to the spacing between parameter points in our training grid. Further testing is required to determine whether a more finely spaced grid would more closely resemble the structure seen in the 2D UCLCHEM posterior distributions. Nevertheless, the parameter values inferred by our neural network algorithm are well within the uncertainties obtained from the same analysis with UCLCHEM. The neural network+UltraNest algorithm completed its explorations of the parameter space in $\sim 7\%$ and $\sim 15\%$ of the time it took for UCLCHEM+UltraNest to do the same for regions 26 and 53, respectively.

4.2. Parameter Maps

We show maps of the inferred parameter values derived from our combined Bayesian nested sampling and neural network algorithm in Figure 5 and their associated uncertainties in Figure 6. In general, we see notable differences in the inferred parameter values in the central 100 pc versus the rest of the

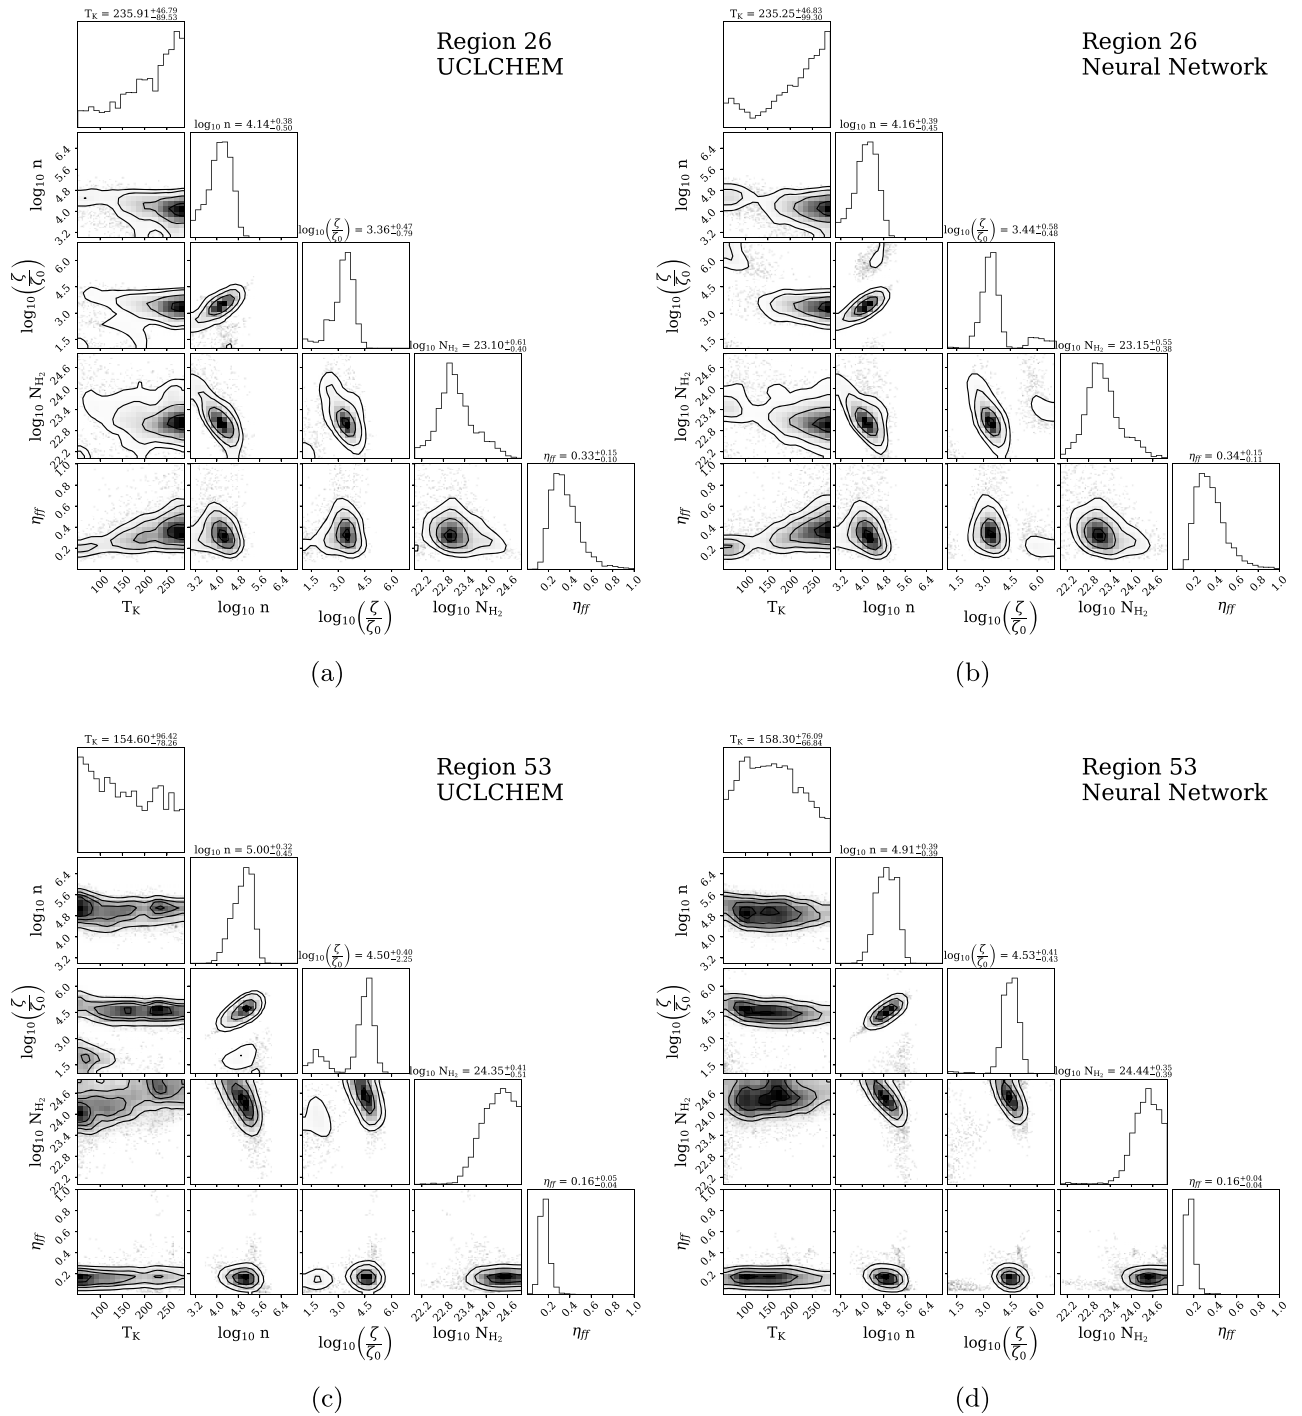


Figure 4. (a) Corner plot for region 26 using UCLCHEM. (b) Corner plot for region 26 using the HCN/HNC neural network. (c) Corner plot for region 53 using UCLCHEM. (d) Corner plot for region 53 using the HCN/HNC neural network.

CMZ. H_2 volume and column density are enhanced by 1–2 orders of magnitude in the nucleus of the CMZ with values reaching as high as $8 \times 10^4 \text{ cm}^{-3}$ and $5 \times 10^{24} \text{ cm}^{-2}$, respectively, compared to outer CMZ values of $\sim 5\text{--}6 \times 10^3 \text{ cm}^{-3}$ and $\sim 3\text{--}4 \times 10^{22} \text{ cm}^{-2}$. We will discuss comparisons of our measurements with those from other NGC 253 studies in Section 5.1.

The CRIR is also an order of magnitude larger in the central $\sim 10''$ ($\sim 4 \times 10^{-13} \text{ s}^{-1}$) than in regions $\gtrsim 15''$ away from the nucleus ($\sim 2 \times 10^{-14} \text{ s}^{-1}$). However, the inferred CRIR map also indicates a ring of high-CRIR ($\gtrsim 5 \times 10^{-13} \text{ s}^{-1}$) regions

exists around the edge of the CMZ. We do not believe these parameter estimates to be physically motivated and will discuss this issue in more detail in Section 5.2. Comparison to other CRIR estimates, both observational and theoretical, will be addressed in Sections 5.3 and 5.4.

Figure 5 also shows a decrease in the beam-filling factor in NGC 253’s nucleus, which is consistent with the gas emission originating from smaller, higher-density regions and is supported by our estimation of higher volume and column densities in these regions. Finally, we see very little spatial correlation between the kinetic temperature and our other fitted




Figure 5. Parameter-inference results from our Bayesian nested sampling plus neural network algorithm for kinetic temperature (upper left), H_2 volume density (upper right), CRIR (middle left), H_2 column density (middle right), and beam-filling factor (bottom). Gray contours correspond to the integrated HCN 1–0 emission at levels of 1, 3, and 5 $\text{Jy km s}^{-1} \text{beam}^{-1}$. White and black crosses indicate the galaxy’s kinematic center (J. L. Turner & P. T. P. Ho 1985). Green circle in the bottom-right corner of each map represents ALCHEMI’s $1''.6$ (~ 28 pc) beam.

parameters, with no consistent increase or decrease in temperature in the central versus outer CMZ regions. In general, we do not constrain the kinetic temperature well with our HCN and HNC measurements (see, for example, the kinetic temperature panels in Figures 4, 5, and 6), which is consistent with the results from E. Behrens et al. (2022). This finding is reinforced by the kinetic temperature uncertainty map in Figure 6. While the rest of our parameter maps show uncertainties are lowest in the center of the CMZ, where we have the highest S/N, the kinetic temperature uncertainty map demonstrates no such correlation.

5. Discussion

5.1. Column Density, Volume Density, Filling Factor, and Kinetic Temperature

In Section 4.2 we found that there are significant gradients in the inferred values of hydrogen column density N_{H_2} and volume density n_{H_2} (Figures 5 and 6). Estimates for volume and column density from this work and other ALCHEMI studies are shown in Table 3 for the inner ($r \lesssim 100$ pc) and outer

($r \gtrsim 100$ pc) CMZ. Note that with the exception of this work, all of the studies listed in Table 3 estimated the volume and column densities on scales of ~ 28 pc, equivalent to the ALCHEMI beam, whereas we estimate these parameters on scales of 50 pc. Therefore, our estimates of the volume and column density are largely consistent with the lower end of the ranges provided for these other studies. This slight discrepancy likely results from averaging these parameters over larger areas that may include more diffuse gas components, especially in the outer CMZ.

We also see a clear distribution across the CMZ in the dense gas beam-filling factor η_{ff} , with the lowest values found in the center. In the inner CMZ ($r \lesssim 100$ pc), $\eta_{\text{ff}} \lesssim 0.1$, while in the outer CMZ ($r \gtrsim 100$ pc), $\eta_{\text{ff}} \gtrsim 0.3$. The gradients in n_{H_2} , N_{H_2} , and η_{ff} are consistent with gas originating from smaller, higher volume and column density regions as one moves from the outer portions of the CMZ toward the nucleus of NGC 253. Kinetic temperature, by contrast, is poorly constrained by our models, exposing no systematic changes as a function of position in the CMZ.

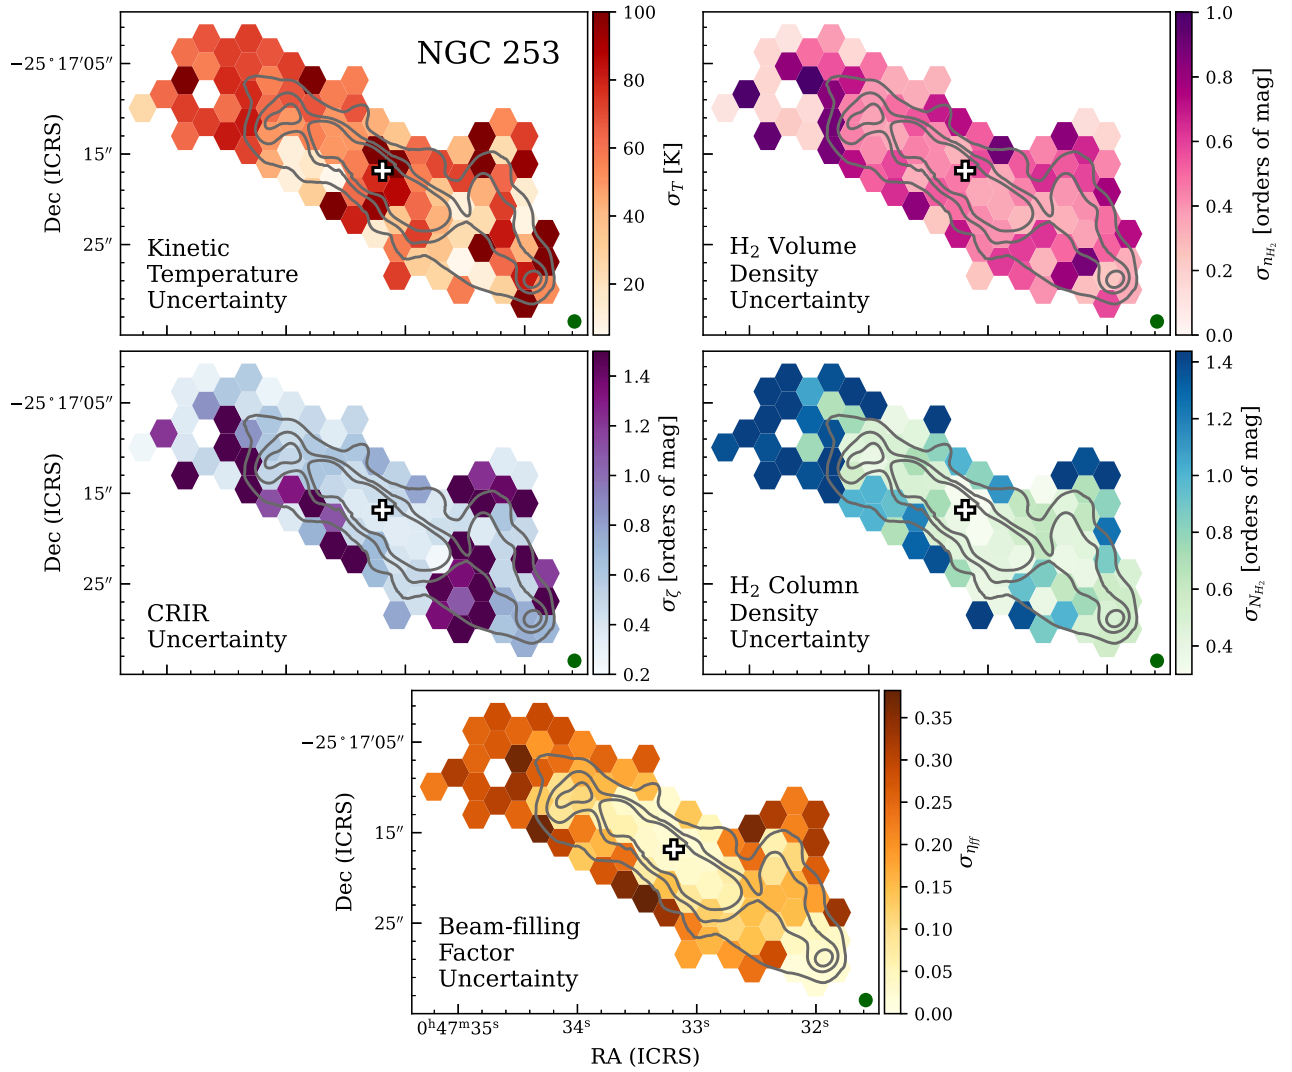


Figure 6. Uncertainties on our five inferred parameters. For the parameters over which we consider a log-uniform distribution (n_{H_2} , ζ , and N_{H_2}), we show their uncertainties as orders of magnitude from the inferred value, as shown above the corner plot 1D distributions in Figure 4.

Table 3

Comparison of NGC 253 Volume and Column Density with the Literature

ALCHEMI Study	$r \lesssim 100$ pc		$r \gtrsim 100$ pc	
	$\log n_{\text{H}_2}$ (cm^{-3})	$\log N_{\text{H}_2}$ (cm^{-2})	$\log n_{\text{H}_2}$ (cm^{-3})	$\log N_{\text{H}_2}$ (cm^{-2})
J. Holdship et al. (2022)	5.3	24.3	4.5	22.6
E. Behrens et al. (2022)	5.0	24.0	3.8–4.7	23.0
K. Tanaka et al. (2024)	5.0–5.4	23.7	3.8–4.5	23.0
This work	4.7–5.0	23.8–24.7	3.7–4.3	22.5–23.5

The centrally condensed nature of the NGC 253 CMZ may reflect the process that catalyzed the starburst in this galaxy. Mergers are a common mechanism for generating starbursts in galaxies (L. Armus et al. 1987; F. Renaud et al. 2022). As NGC 253 shows no signs of having participated in a merger, the trigger for the burst of star formation in the NGC 253 CMZ must be associated with its high concentration of dense gas. Starburst galaxy evolution models can reproduce such a process. E. Cenci et al. (2024) have modeled starburst galaxy evolution using the FIREBOX simulation (R. Feldmann et al. 2023). They find that global gravitational instabilities that drive

central gas compaction are the main catalyst for starbursts in galaxies, independent of a galaxy’s merger status.

5.2. Apparent High- ζ Regions

As noted above in Section 4.2 and in Figure 5, the CRIR maps feature regions primarily along the edge of the CMZ with CRIRs seemingly higher than those found in the nucleus ($\zeta \gtrsim 5 \times 10^{-13} \text{ s}^{-1}$). E. Behrens et al. (2022) noted a positive correlation between the number of cosmic-ray sources (e.g., supernova remnants, which are primarily located in the center of the CMZ) and the CRIR in NGC 253. As such, the lack of cosmic-ray sources around the edge of the CMZ begs questions about the legitimacy of these high inferred ζ values. We will explore two possible explanations for this behavior.

5.2.1. High $X_{\text{HCN}}/X_{\text{HNC}}$

First, we will consider the HCN/HNC abundance ratio, $X_{\text{HCN}}/X_{\text{HNC}}$. E. Behrens et al. (2022) found a negative correlation between $X_{\text{HCN}}/X_{\text{HNC}}$ and the CRIR, with $X_{\text{HCN}}/X_{\text{HNC}} < 4$ for $\zeta > 10^{-13} \text{ s}^{-1}$ and decreasing toward unity for even higher CRIRs. They explained that at high CRIRs both HCN and HNC are expected to be formed and destroyed at similar rates, resulting




Figure 7. Probability density for HCN/HNC abundance ratio (a) and cosmic-ray ionization rate map (b) with dotted hatching in both plots corresponding to $X_{\text{HCN}}/X_{\text{HNC}} > 4$.

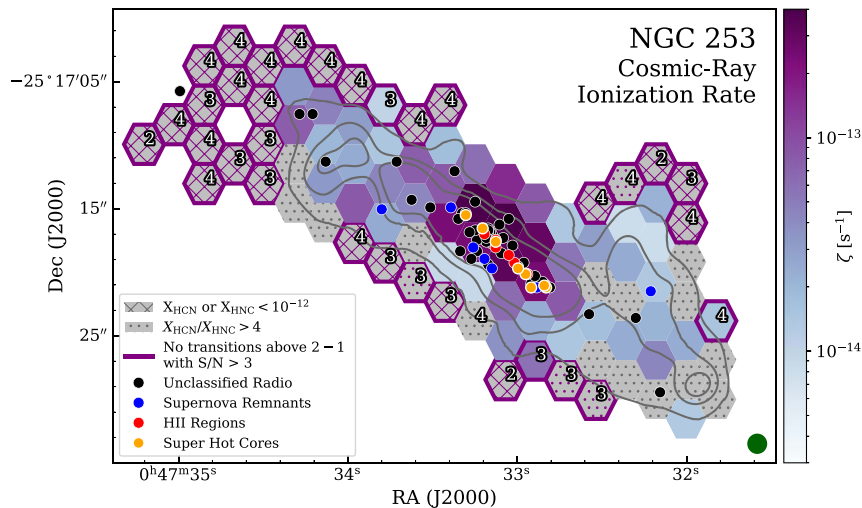


Figure 8. CRIR map estimated from our parameter-inference algorithm. All grayed-out, hatched regions are estimated by our algorithm to have $\zeta > 5 \times 10^{-13} \text{ s}^{-1}$ and one or more of the following issues. The “X” hatching indicates regions where the estimated HCN or HNC abundance was below 10^{-12} with respect to hydrogen, and the dotted regions indicate locations where $X_{\text{HCN}}/X_{\text{HNC}} > 4$, with the two appearing together when both conditions are true. White and black numbers indicate the number of transitions that are detected with $S/N > 3$ for a given region, and purple outlines indicate when there are no transitions above the $J = 2-1$ that meet this threshold. Gray contours signify HCN 1-0 emission at 1, 3, and 5 $\text{Jy km s}^{-1} \text{beam}^{-1}$. Colored dots indicate the locations of recent star formation via radio continuum sources (supernova remnants and H II regions; J. S. Ulvestad & R. R. J. Antonucci 1997) and super hot cores (F. Rico-Villas et al. 2020). Note half of the unclassified sources are expected to be H II regions, and half are thought to be supernova remnants. The green circle in the bottom-right corner represents the ALCHEMI 1''/6 (~ 28 pc) beam.

in $X_{\text{HCN}}/X_{\text{HNC}}$ tending toward 1. However, as shown in Figure 7, we find $X_{\text{HCN}}/X_{\text{HNC}} > 4$ in many regions with high ζ . Figure 7(a) shows a clear bimodal behavior in $X_{\text{HCN}}/X_{\text{HNC}}$, where all 24 regions with $X_{\text{HCN}}/X_{\text{HNC}} > 4$ also feature $\zeta > 5 \times 10^{-13} \text{ s}^{-1}$. These regions are not located within the nucleus’s innermost $5''$ and generally do not contain cosmic-ray-producing sources (see Figure 8).²⁷

Additionally, we find that most of the regions identified in Figure 7 pair high CRIRs with very low kinetic temperatures such that T_{K} is often pinned against the lower end of our prior

distribution (see Figure B1 for examples). Though we do not constrain T_{K} well with our models, the high CRIRs that we estimate in these edge regions are not compatible with such low kinetic temperatures. Cosmic-ray ionization is expected to contribute significant energy to interstellar gas, raising the kinetic temperature to several hundreds of Kelvin (E. Bayet et al. 2011; E. Behrens et al. 2022), assuming the gas column density is not so high as to shield the inner layers from ionization (M. Padovani et al. 2022). The unphysical combination of high ζ and very low T_{K} , in addition to anomalously high $X_{\text{HCN}}/X_{\text{HNC}}$ values, indicate that the chemistry included in our model is not well suited to the conditions in these regions.

As mentioned in Section 3.2, we assume in our model that the gas we are probing is dense enough such that it is shielded

²⁷ Note that, of the “Unclassified Radio” sources labeled in Figure 8 (identified in J. S. Ulvestad & R. R. J. Antonucci 1997), half are expected to be supernova remnants, and half are likely H II regions.

from UV radiation and therefore dominated by cosmic-ray ionization. While this was true for the dense gas in the GMCs studied in E. Behrens et al. (2022), it is likely not the case for the regions in more diffuse areas of the NGC 253 CMZ. As a result, UV radiation is likely impacting the HCN and HNC chemistry. M. G. Santa-Maria et al. (2023) found that HCN was enhanced by far-UV radiation, increasing the HCN/HNC ratio as the incident radiation field increased. Thus, one explanation for the high-CRIR values outside of the inner 5'' is that they are not shielded against UV radiation and are therefore not well described by the chemistry in our model.

Another possibility is that some of these regions are affected by shocks. R. Meijerink et al. (2011) discussed the influence of mechanical heating from shocks on the HCN/HNC ratio, showing that it is enhanced by shock activity. K. Tanaka et al. (2024) constrained the kinetic temperature in NGC 253 and noted four high-temperature regions around the outskirts of the CMZ (approximately regions 34, 49, 51, and 71 from this work). They hypothesized that the gas in these regions is emitting from above the plane of the CMZ and could therefore be interacting with the large-scale outflow (A. D. Bolatto et al. 2013). In the southwest, region 86 (GMC 1) is another location potentially featuring shocks. M. Gorski et al. (2017) showed that this region is close to an expanding shell of dense gas, and K. Y. Huang et al. (2023) also found evidence indicative of shocks here. Since the chemistry in shocked regions is likely not dominated by cosmic-ray ionization, the model presented here may struggle to constrain the parameters in those areas.

5.2.2. Low X_{HCN} and X_{HNC}

After considering the regions marked by high $X_{\text{HCN}}/X_{\text{HNC}}$ that may involve chemistry not included in our model, we see in Figure 7 that there are still several regions, primarily in the northeast and northwest, that feature uncharacteristically high CRIRs. To address these regions, we then look beyond abundance ratios and consider the individual values of X_{HCN} and X_{HNC} . As mentioned in Section 3.2, UCLCHEM abundance estimates, upon which our neural network is based, are dominated by numerical errors when considering abundances below 10^{-12} . As a result, we remove all parameter combinations from our training set that yield HCN or HNC abundances below this threshold. However, the neural network is able to characterize the functional relationships that exist between the parameters and molecular abundances during training. Since these functions are smooth around the abundance threshold we have imposed, it can still extrapolate and produce abundances below the 10^{-12} threshold if provided with the relevant parameter values by our nested sampling algorithm.

We find that in many of the high- ζ regions along the outskirts of the CMZ, the most likely range of parameters yields X_{HCN} and X_{HNC} values below 10^{-12} (see the “X”-hatched regions in Figure 8). Interestingly, despite not being trained on parameter combinations that result in such low abundances, the neural network does an excellent job at replicating the results that would be produced with UCLCHEM. Figure B2 in Appendix B shows the posterior distributions achieved for region 1 using UCLCHEM and the neural network model, demonstrating nearly identical results in this low-abundance region. It is thus clear that our high- ζ result is not a product of the neural network, but of the chemical model’s inefficacy at such low abundances. We see in Figure 8 that all low-abundance regions feature four or fewer combined HCN




Figure 9. Comparison of our CRIR estimates with those from E. Behrens et al. (2022), J. Holdship et al. (2022), and V. H. M. Phan et al. (2024). As proxies for GMCs 2–10, which were not studied in this work, we provide values for regions 75, 58, 53, 47, 36, 26, 20, and 16, which are the regions that overlap most significantly with the plotted GMCs. The widths of the shaded regions at each data point indicate the most likely 66% of each ζ posterior distribution.

and HNC transitions that meet our S/N criteria (as compared to the full set of eight that are available) and no higher-energy transitions detected above $J = 2-1$. The lack of HCN and HNC emission that spans the excitation ladder in these regions corroborates the low abundances estimated by UCLCHEM and the neural network. Since chemistry is largely irrelevant in this low-abundance regime, we determine that parameter results for any regions that correspond to such low abundances are unreliable.

5.3. Comparison of Cosmic-Ray Ionization Rate to Other Measurements

We compare our CRIR results to those predicted using similar methods in J. Holdship et al. (2022) and E. Behrens et al. (2022). While J. Holdship et al. (2022) studied H_3O^+ and SO and E. Behrens et al. (2022) used HCN and HNC, both studies employed Bayesian nested sampling with a combination of chemical and radiative transfer models in order to infer the dense gas physical parameters in NGC 253’s CMZ. As mentioned in Section 1, both studies investigated a smaller number of GMC-sized regions (5 and 10, respectively) than in the study presented here. We compare our parameter estimations to those inferred in J. Holdship et al. (2022) and E. Behrens et al. (2022), which are plotted as colored-shaded regions in Figure 9. Since GMCs 1–10 were not explicitly studied in this work, we associate each GMC with the hexagonal region with which it shares the most overlap (see Figure 1). Note that both of the aforementioned studies used smaller regions (~ 28 pc circles, equal to the size of the ALCHEMI beam) than the 50 pc hexagonal regions we study here. We are unable to plot a value for GMC 1, as region 86 (its closest proxy) was removed due to possible UV or shock contamination. Though E. Behrens et al. (2022) was able to obtain a reasonable CRIR estimate for GMC 1, which corresponded to a smaller region associated with denser gas, the larger area of our hexagonal regions results in the inclusion of more diffuse gas that is more likely to be influenced by UV radiation. Thus we are unable to obtain a reliable CRIR estimate for this region using our current chemical framework.

As shown in Figure 9, this work predicts slightly lower CRIR values than those estimated previously with HCN and

HNC in E. Behrens et al. (2022), suggesting a CRIR range of 5×10^{-15} – 10^{-12} s^{-1} , whereas E. Behrens et al. (2022) predicted a range of 6×10^{-13} – $5 \times 10^{-12} \text{ s}^{-1}$. This lower range agrees better with the values reported in J. Holdship et al. (2022), who estimated a CRIR range of 10^{-14} – 10^{-12} s^{-1} . The difference we see between our study’s CRIR estimates and those estimated in E. Behrens et al. (2022) is likely a result of updates we implemented in our modeling procedure and not the effect of our neural network use. The model differences are as follows:

1. *Multi-phase modeling.* In this work, we use UCLCHEM’s two-phase model that includes the freefall collapse of a cloud in order for the gas to start the second phase of the model in molecular form, rather than atomic or ionic. This first phase was not used in E. Behrens et al. (2022) but was indeed implemented in J. Holdship et al. (2022). In cases where the CRIR was well constrained using only phase 2, the addition of phase 1 had little effect on our posterior distributions. In regions where our posterior distributions were less well constrained, introducing phase 1 into our modeling had a more significant effect on the shape of the posterior distributions but did not greatly alter the median values derived from each distribution. Figure B3 demonstrates the differences in the posterior distributions with and without phase 1 modeling for sample regions 26 and 53.
2. *New collisional excitation rates.* We implement more recent collisional excitation rates in our radiative transfer modeling with SpectralRadex. These rates are derived from calculations done by M. Hernández Vera et al. (2017) that included collisions with ortho- and para- H_2 , whereas E. Behrens et al. (2022) used older rates that were based on scaled calculations from collisions with helium (F. Dumouchel et al. 2010). These new rates did not result in significant changes in our posterior distributions for the majority of tested regions. A detailed analysis of the new collisional rates and their impact on each of our SpectralRadex parameters can be found in Appendix C.
3. *Beam-filling factor.* Recent work by J. Butterworth et al. (2024) suggests that HCN emission in NGC 253 originates from sources significantly smaller than the ~ 28 pc ALCHEMI beam, indicating that a beam-filling factor is a necessary addition to an algorithm that models HCN, and likely also HNC. Thus, we also implement a beam-filling factor as one of our free parameters. The addition of this parameter did not greatly change the median CRIR values inferred for any of our test regions—small variations were noted, but any differences were within parameter uncertainties. However, including a beam-filling factor did slightly widen the posterior distributions derived for each parameter. This difference is demonstrated in Figure B4.

For the most part, none of these changes individually have a large impact on the CRIR estimates in NGC 253’s CMZ. However, when applied together, we see a slight decrease (about half an order of magnitude) in the predicted CRIR estimates as compared to E. Behrens et al. (2022). With the exceptions of GMCs 1 (discussed above) and 5, all CRIR uncertainties overlap with those from both J. Holdship et al. (2022) and E. Behrens et al. (2022). It is important to

remember, however, that the regions we study here are nearly twice as large in diameter and contain ~ 3.5 times the area as those investigated in the aforementioned studies. Therefore, a decrease in the estimated CRIR could also be attributed to the gas parameters being averaged over a larger area, thus likely including both diffuse and dense components, whereas the smaller GMC structures, as defined in A. K. Leroy et al. (2015), are centered on denser gas components.

5.4. Comparison of Cosmic-Ray Ionization Rate to Theoretical Predictions

To confirm the validity of the high CRIRs derived from molecular emission in E. Behrens et al. (2022), J. Holdship et al. (2022), and this work, we consider CRIRs derived from theoretical calculations and fits to nonthermal radio, X-ray, and gamma-ray emission data from V. H. M. Phan et al. (2024). This study considers a theoretical transport model for cosmic rays from which they derive cosmic-ray spectra (informed by observations of nonthermal emission) that are used to derive CRIRs as a function of molecular cloud H_2 column density in NGC 253, M82, and Arp 220. In developing the transport model, they consider cosmic-ray protons and electrons injected by supernova remnants, as well as secondary and tertiary electrons injected from proton–proton interactions and low-energy photon interactions with cosmic-ray-induced gamma-rays, respectively. V. H. M. Phan et al. (2024) include four processes in their model to produce cosmic-ray-induced nonthermal emission: π_0 decay from proton–proton interactions, bremsstrahlung radiation, inverse Compton scattering, and synchrotron radiation. Using this model and parameters derived from fits to nonthermal emission from Fermi-LAT (A. A. Abdo et al. 2010) and H.E.S.S. (H. E. S. S. Collaboration et al. 2018), they model cosmic-ray spectra and derive ζ as a function of H_2 column density, obtaining a value of $1.5 \times 10^{-14} \text{ s}^{-1}$ for an H_2 column density of 10^{23} cm^{-2} . However, the dependence of the CRIR on column density is weak, with the CRIR varying by less than 0.5 dex for $10^{22} < N(\text{H}_2) < 10^{25} \text{ cm}^{-2}$.

This value is included as a reference in Figure 9 and is consistent with the lowest end of our modeled CRIR range. V. H. M. Phan et al. (2024) consider possible sources of discrepancy between their calculated values of ζ and those derived through molecular emission modeling from J. Holdship et al. (2022) and E. Behrens et al. (2022). They note that the gamma-ray data, from which key parameters such as the supernova rate and interstellar medium (ISM) density are derived, have uncertainties that could alter the gamma-ray flux by a factor of 2, resulting in proportional variations of the CRIR. Additionally, V. H. M. Phan et al. (2024) note they are limited by the point-spread functions of the nonthermal emission data, which are on the order of a few arcminutes (1 arcminute $\simeq 1020$ pc in NGC 253). Thus, any CRIR values derived from Fermi-LAT and H. E. S. S. measurements must be treated as an average over the entire CMZ. As shown in Figures 5 and 8, the distributions of the gas parameters and locations of supernova remnants in NGC 253 vary spatially, with a higher density of both molecular gas and supernova remnants in the center of the CMZ, where the CRIR is also enhanced. Thus, we would expect any calculation of an average CRIR to be more consistent with values derived from molecular emission for the outer CMZ.

Finally, V. H. M. Phan et al. (2024) also consider the possibility of ionization contributions from sources that accelerate cosmic rays with an energy $E \lesssim 280$ MeV. At GeV and TeV energies, the cosmic-ray spectrum is dominated by π_0 decay, which drops off sharply for $E \lesssim 280$ MeV, as π_0 cannot be produced at energies below this threshold. However, V. H. M. Phan et al. (2024) suggest there may be a population of sources that accelerate cosmic rays with MeV energies. These sources, which may include protostellar jets (M. Padovani et al. 2015, 2016; B. A. L. Gaches & S. S. R. Offner 2018), H II regions (F. Meng et al. 2019; M. Padovani et al. 2019), and stellar wind termination shocks (K. Scherer et al. 2008), are likely deeply embedded within the cloud and unable to penetrate into the more diffuse ISM due to the high gas density. If so, these MeV sources could contribute to the CRIR but might not be observable to gamma-ray facilities.

5.5. Cosmic-Ray Enhancement in Star-forming Regions

While the CRIR estimates presented by V. H. M. Phan et al. (2024) are consistent with the lowest values we infer in the outer CMZ, work by V. A. Semenov et al. (2021) could further explain the enhancement in CRIR values that we see in NGC 253's nucleus. Many cosmic-ray transport models assume isotropic diffusion and therefore estimate CRIRs far from the sources that produce cosmic rays. However, as Figure 8 shows, many of the central regions in NGC 253's CMZ contain sources that will produce and accelerate cosmic rays, such as supernova remnants and H II regions. Therefore, it is important to consider how the CRIR and its effects on star formation would manifest in regions near these sources. V. A. Semenov et al. (2021) use simulations to test the effect of varying the cosmic-ray diffusion coefficient κ_{CR} in star-forming regions. In environments with strong shocks from supernova explosions, the cosmic-ray current is much higher, which results in plasma instabilities and the excitation of nonresonant modes. The background magnetic field is therefore amplified, which can significantly suppress κ_{CR} . In the Milky Way, κ_{CR} is assumed to be $\sim 10^{28} \text{ cm}^2 \text{ s}^{-1}$. However, studies of nearby supernova remnants and molecular clouds have shown that within ~ 50 pc of cosmic-ray injection sites, κ_{CR} can be 10–100 times smaller (see, for example, H. Li & Y. Chen 2010; M. Ajello et al. 2012; Y. Hanabata et al. 2014), and theory predicts it could be suppressed by a factor of up to 10^6 (P. Blasi et al. 2007). This suppression of cosmic-ray diffusion would result in a buildup of cosmic rays and increased ionization rates at injection sites, which has important implications for star formation and galaxy evolution.

V. A. Semenov et al. (2021) test this theory by running simulations (i) without cosmic-ray feedback, (ii) with cosmic-ray feedback and a constant diffusivity of $\kappa_{\text{CR}} = 10^{28} \text{ cm}^2 \text{ s}^{-1}$, and (iii) with cosmic-ray feedback where κ_{CR} is suppressed in star-forming regions. Their simulations show that, in models where κ_{CR} is suppressed, the cosmic-ray pressure is significantly higher in the galaxy centers, and cosmic-ray pressure makes up a much larger fraction ($>75\%$) of the total pressure budget. They also find that, in these same models, fewer dense, star-forming clumps are allowed to form and the overall SFR is reduced by a factor of ~ 3 – 4 . These effects are magnified in galaxies that were simulated to have a higher gas-mass fraction f_{g} of 40% versus those with $f_{\text{g}} = 20\%$. Additionally, V. A. Semenov et al. (2021) conclude that, after the last supernova in the simulation explodes, cosmic rays are finally

able to escape their injection sites and permeate the ISM, resulting in an extended vertical pressure gradient that can stabilize the disk against fragmentation and drive galactic winds. It is clear that cosmic rays play a crucial role in the regulation of star formation and galaxy evolution.

Though V. A. Semenov et al. (2021) show that cosmic-ray feedback has the power to disrupt, and even halt, star formation, NGC 253 may be at too early a stage of its starburst phase to show symptoms of star formation suppressed by cosmic-ray feedback. Recent analyses have identified 14 super star clusters (SSCs) in NGC 253 with a likely age range of 0.01–3 Myr and an age gradient whereby the older SSCs are located in the center of the NGC 253 CMZ (A. K. Leroy et al. 2018; E. A. C. Mills et al. 2021; R. C. Levy et al. 2022; J. Butterworth et al. 2024). This gradient may account for the greater number of supernova remnants and H II regions in the center of the NGC 253 CMZ, as the younger SSCs farther from the center may not have existed long enough for massive stars to form and die. It is unclear, though, at exactly what point from the beginning of a starburst phase cosmic-ray feedback is expected to noticeably affect star formation. In the V. A. Semenov et al. (2021) simulations, supernova explosions occur between 3 and 43 Myr after the start of the simulation, and supernova remnants are thought to have lifetimes $\sim 10^5$ – 10^6 yr (S. K. Sarbadhicary et al. 2017; A. Bamba & B. J. Williams 2022). In noninteracting galaxies, such as NGC 253, the starburst phase is thought to last ~ 70 Myr (E. Cenci et al. 2024). However, V. A. Semenov et al. (2021) show that the effects of cosmic-ray feedback (e.g., decrease in SFR and number of dense gas clumps) are able to be seen 300–800 Myr after the beginning of the simulation. Thus, the possible quenching effects of cosmic-ray feedback may not yet be visible in NGC 253. Further investigations of starburst galaxies at later stages of evolution are necessary in order to better understand the full scope of cosmic-ray feedback's role in star formation and galaxy evolution.

6. Conclusions

We present a neural network model that we have trained to replace the function of a chemical modeling code in a gas parameter-inference algorithm. We use ALMA measurements of the first four rotational transitions of HCN and HNC to constrain our neural network and radiative transfer models in a Bayesian nested sampling framework in order to infer the gas conditions across the NGC 253 CMZ at 50 pc resolution. We find that the neural network model can very effectively reproduce the same molecular abundances and gas parameter values as would be found using a chemical modeling code but in $\sim 10\%$ of the time. Using this algorithm, we draw the following conclusions regarding the molecular gas conditions in NGC 253:

1. Our results show clear spatial gradients across the CMZ in the gas volume and column density as well as the beam-filling factor. Volume and column densities are higher by an order of magnitude in the center ($r \lesssim 100$ pc) of NGC 253's CMZ ($n_{\text{H}_2} \sim 5 \times 10^5 \text{ cm}^{-3}$, $N_{\text{H}_2} \sim 10^{24} \text{ cm}^{-2}$) than they are farther away from the nucleus. The beam-filling factor reaches its minimum ($\eta_{\text{ff}} < 0.1$) in NGC 253's nucleus, which is consistent with the emission in this area originating from small, dense gas components.

2. We estimate CRIR values of $\zeta \sim 4 \times 10^{-13} \text{ s}^{-1}$ in the nucleus of the NGC 253 CMZ and $\zeta \sim 2 \times 10^{-14} \text{ s}^{-1}$ in its outer regions. These derived CRIRs are consistent with those derived for specific GMCs in the NGC 253 CMZ by J. Holdship et al. (2022) and E. Behrens et al. (2022).
3. We are unable to derive CRIR estimates for regions with low HCN and HNC abundance as well as regions likely featuring UV-dominated or shock chemistry ($X_{\text{HCN}}/X_{\text{HNC}} \gtrsim 4$), which is not included in our models. Chemical modeling of HCN and HNC in photodissociation regions and shock environments will be addressed in a future study of the sources of energy in the NGC 253 CMZ.
4. The CRIR estimates in the outer CMZ are consistent within ~ 0.2 dex with theoretical predictions of the CRIR V. H. M. Phan et al. (2024) derived from CMZ-averaged nonthermal emission.
5. The high-CRIR estimates in the inner CMZ are likely a result of their proximity to supernova remnants and other cosmic-ray-producing sources, as well as potential cosmic-ray diffusion suppression that prevents cosmic rays from escaping into the outer CMZ.

This study demonstrates that neural networks can be an effective and efficient tool for replacing chemical models in larger parameter-inference algorithms. This method has the potential to significantly decrease the time required for inferring physical parameters across sources in large samples or in individual images with large numbers of pixels. As discussed in Section 5.2, this model is not appropriate for regions where the chemistry is very sensitive to small changes in temperature, i.e., photodissociation regions. Future possibilities include incorporating additional molecular tracers that trace the same gas component into a single neural network model in order to constrain the gas parameters even further. Additional future work will include utilizing the neural network’s speed to investigate the limitations associated with using molecular emission to infer gas parameters, particularly focusing on how the number of constraining transitions affects the parameter-inference outcomes.

Acknowledgments

We thank Tommaso Grassi for his very helpful and constructive feedback in reviewing this article. We also thank Phoebe Behrens for her crucial support to the authors throughout this process. This work is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research

and innovation program MOPPEX 833460. V.M.R. and L.C. acknowledge support from the grant PID2022-136814NB-I00 by the Spanish Ministry of Science, Innovation and Universities/State Agency of Research MICIU/AEI/10.13039/501100011033 and by ERDF, UE. V.M.R. also acknowledges the grant RYC2020-029387-I funded by MICIU/AEI/10.13039/501100011033 and by “ESF, Investing in your future,” and from the Consejo Superior de Investigaciones Científicas (CSIC) and the Centro de Astrobiología (CAB) through the project 20225AT015 (Proyectos intramurales especiales del CSIC); and from the grant CNS2023-144464 funded by MICIU/AEI/10.13039/501100011033 and by “European Union NextGenerationEU/PRTR.” N.H. acknowledges support from JSPS KAKENHI grant No. JP21K03634. P.H. is a member of and received financial support for this research from the International Max Planck Research School (IMPRS) for Astronomy and Astrophysics at the Universities of Bonn and Cologne. K.S. acknowledges the grant MOST 111-2112-M-001-039 from the Ministry of Science and Technology in Taiwan.

This article makes use of the following ALMA data: ADS/JAO.ALMA#2017.1.00161.L and ADS/JAO.ALMA#2018.1.00162. S. ALMA is a partnership of ESO (representing its member states), NSF (USA) and NINS (Japan), together with NRC (Canada), MOST and ASIAA (Taiwan), and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO and NAOJ. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

Facility: ALMA.

Software: CASA (CASA Team et al. 2022), Astropy (Astropy Collaboration et al. 2013, 2018, 2022), MLFriends (J. Buchner 2016, 2019), UltraNest (J. Buchner 2021), SciPy (P. Virtanen et al. 2020).

Appendix A Neural Network Architecture

Neural networks consist of multiple layers of nodes, where each node is connected to all nodes in the layers immediately preceding and following it (see Figure A1). The number of layers and nodes needed is typically unique to the problem at hand and is found primarily through trial and error—in our case, we found a neural network with four layers each containing 2000 nodes produced results that aligned best with chemical model calculations. Every node has a unique value and bias, and each connection between nodes in different layers is given a weight (see Equation (A1)). Altogether, these

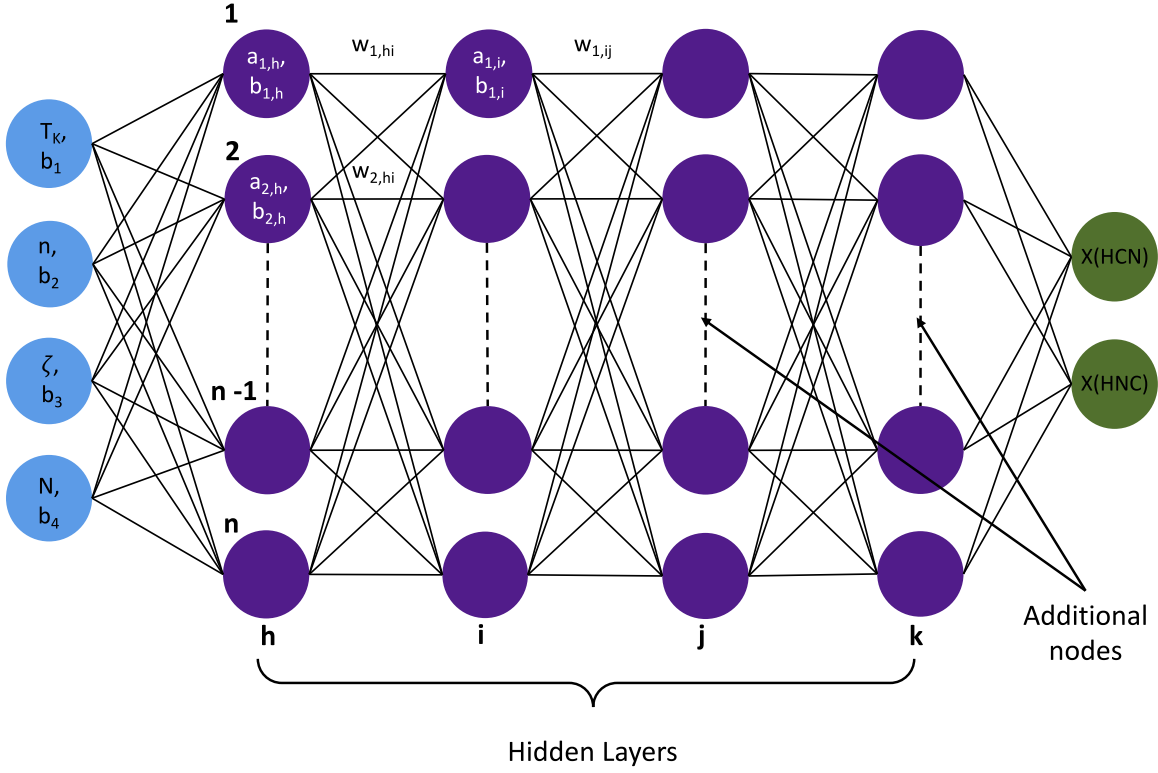


Figure A1. Schematic of our neural network architecture. Blue circles represent the nodes in the input layer of the neural network that contain our input parameters, purple circles represent the nodes that make up the neural network’s inner hidden layers, and the green circles represent the output nodes in the final layer, which contain the abundances for HCN and HNC, respectively.




Figure A2. Exponential moving average of the mean-squared-error loss as a function of training epoch calculated for both the training and validation data sets. Here, the neural network was trained for many epochs past our early stopping criteria (described in Section 3.1) to demonstrate the divergence between the training and validation set losses. The dashed gray line indicates the epoch from which the final weights would have been obtained, as the validation loss does not improve for more than 20 epochs after this point.

quantities are used to calculate the values of nodes in the following layer. The values of the nodes in the first layer are the values of our input parameters (kinetic temperature T_K , H_2 volume density n , CRIR ζ , and H_2 column density N), and each node is assigned an arbitrary starting bias and each node connection an arbitrary starting weight. During the neural network training process, the node values, biases, and weights in each layer are successively manipulated to produce

molecular abundances, which are stored in the final layer of the neural network. To calculate the value $a_{1,j}$ of a node in row 1 and layer j , we use the following prescription:

$$a_{1,j} = \Phi \left(\sum_{n=1}^{2000} a_{n,i} w_{n,ij} + b_{1,j} \right), \quad (\text{A1})$$

where $a_{n,i}$ is the value of a node in the previous layer i and some row n , $w_{n,ij}$ is the weight associated with the unique connection between that node and our target node $a_{1,j}$, $b_{1,j}$ is the bias of our target node in layer j , and Φ is an activation function through which the sum of the combinations for each node in the previous layer are passed. Activation functions, such as the rectified linear unit (ReLU) that we employ, allow neural networks to consider nonlinear solutions.

As discussed in Section 3.1, we assess the difference between the UCLCHEM and neural network abundances using a mean-squared-error loss function, where the loss is calculated on the training set throughout each epoch and calculated on a subset of the validation set just once at the end of each epoch. We calculate the validation loss using mini batches with a batch size of 64. The exponential moving average of the loss associated with our neural network training is shown in Figure A2.

Appendix B Additional Figures

This Appendix includes a number of figures provided in support of the discussion of Section 5.2 (Figures B1, B2, B3, and B4).

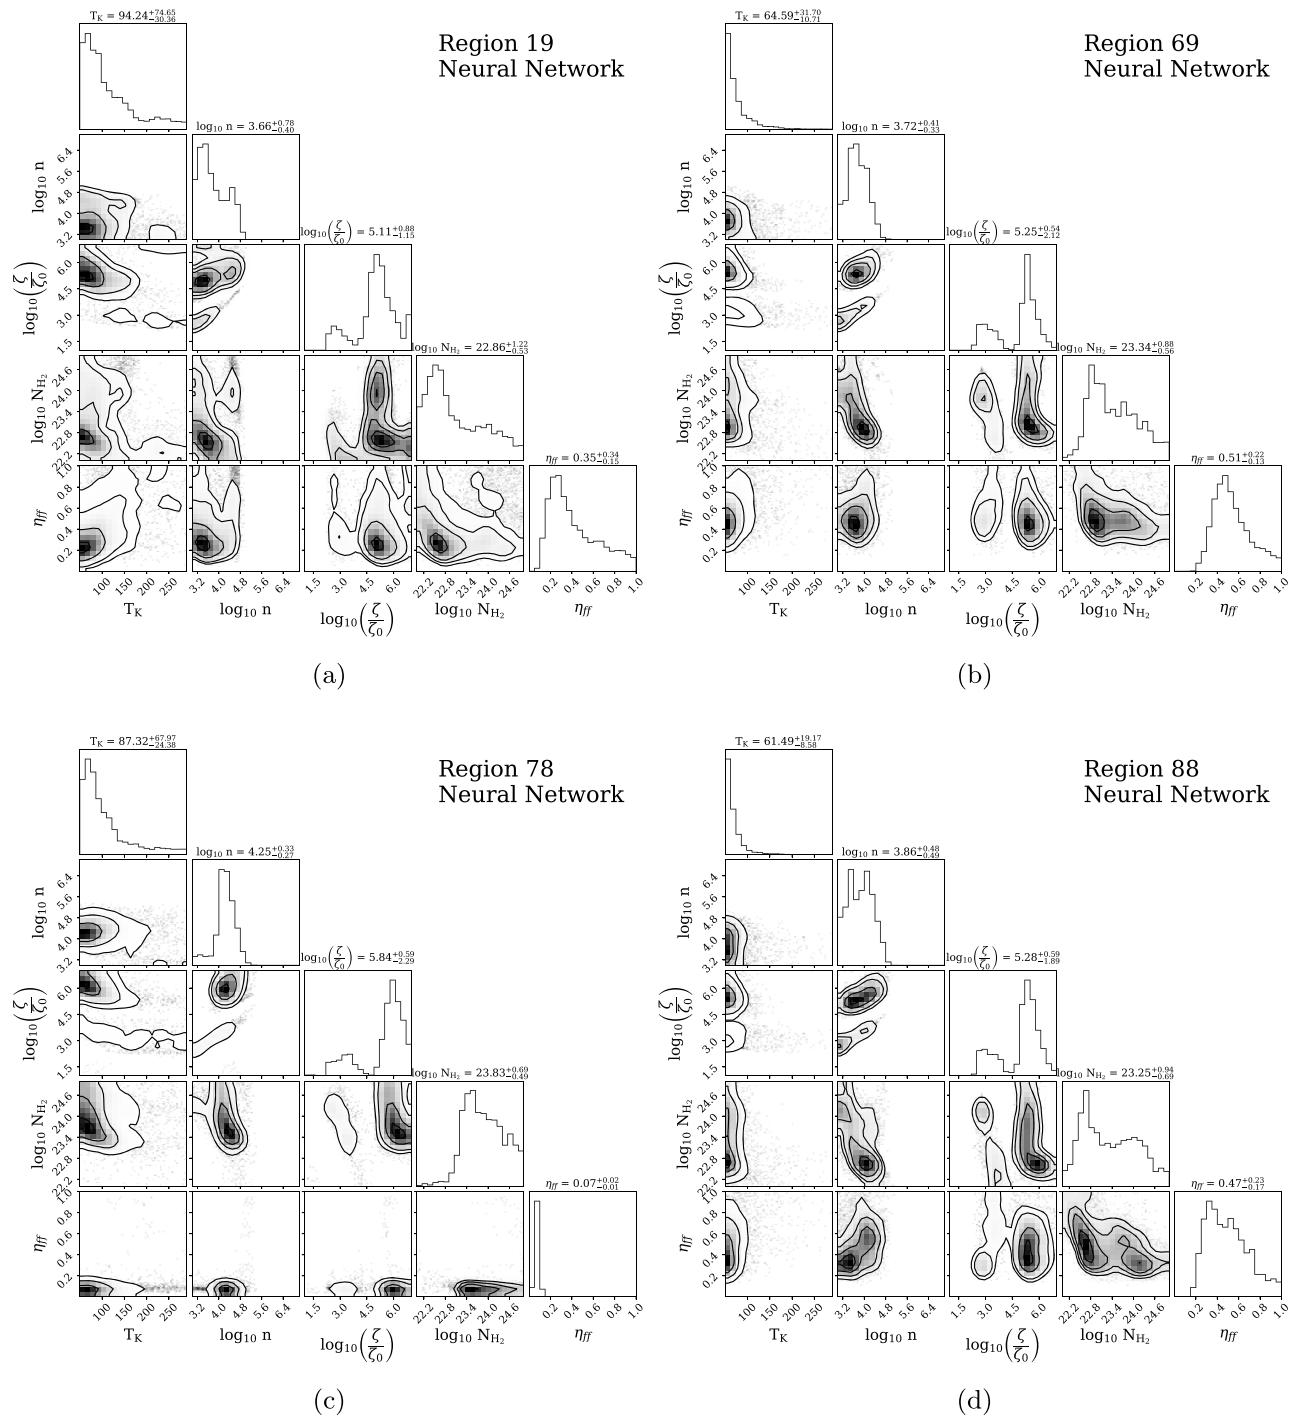


Figure B1. Posterior distributions for regions 19 (a), 69 (b), 78 (c), and 88 (d), which all feature high CRIRs paired with low kinetic temperatures.

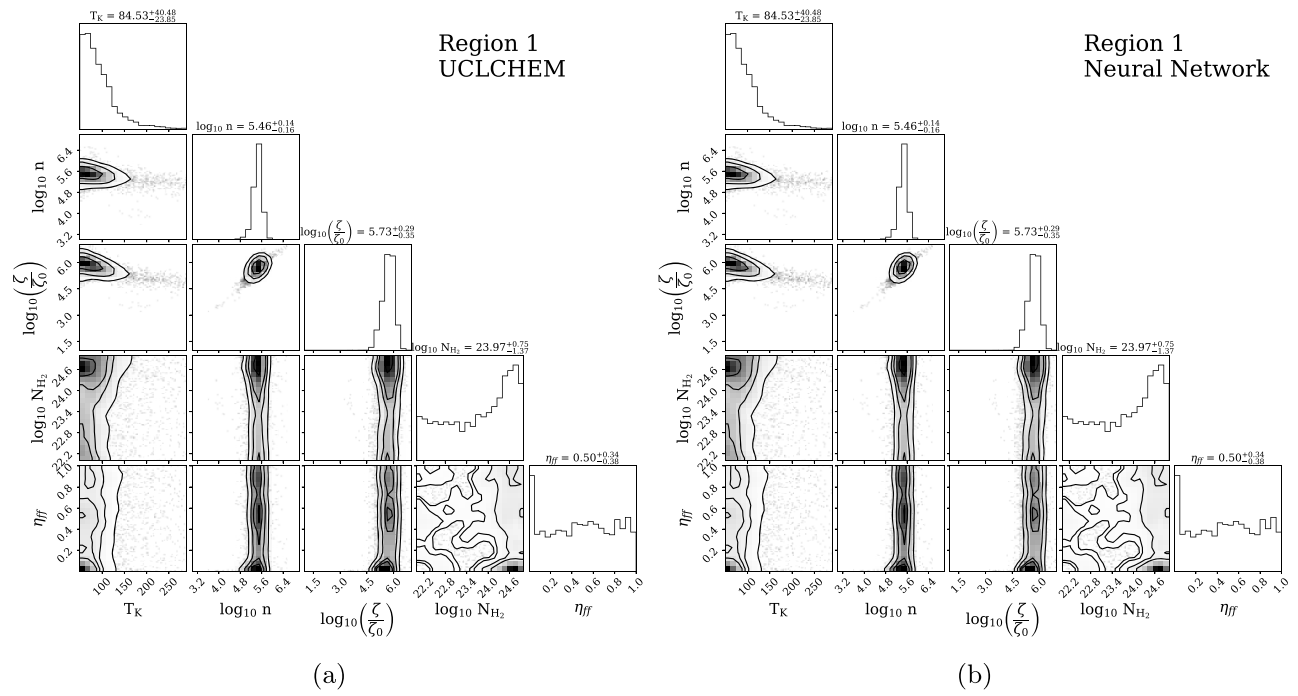


Figure B2. Posterior distributions for region 1 using UCLCHEM (a) vs. the neural network (b).

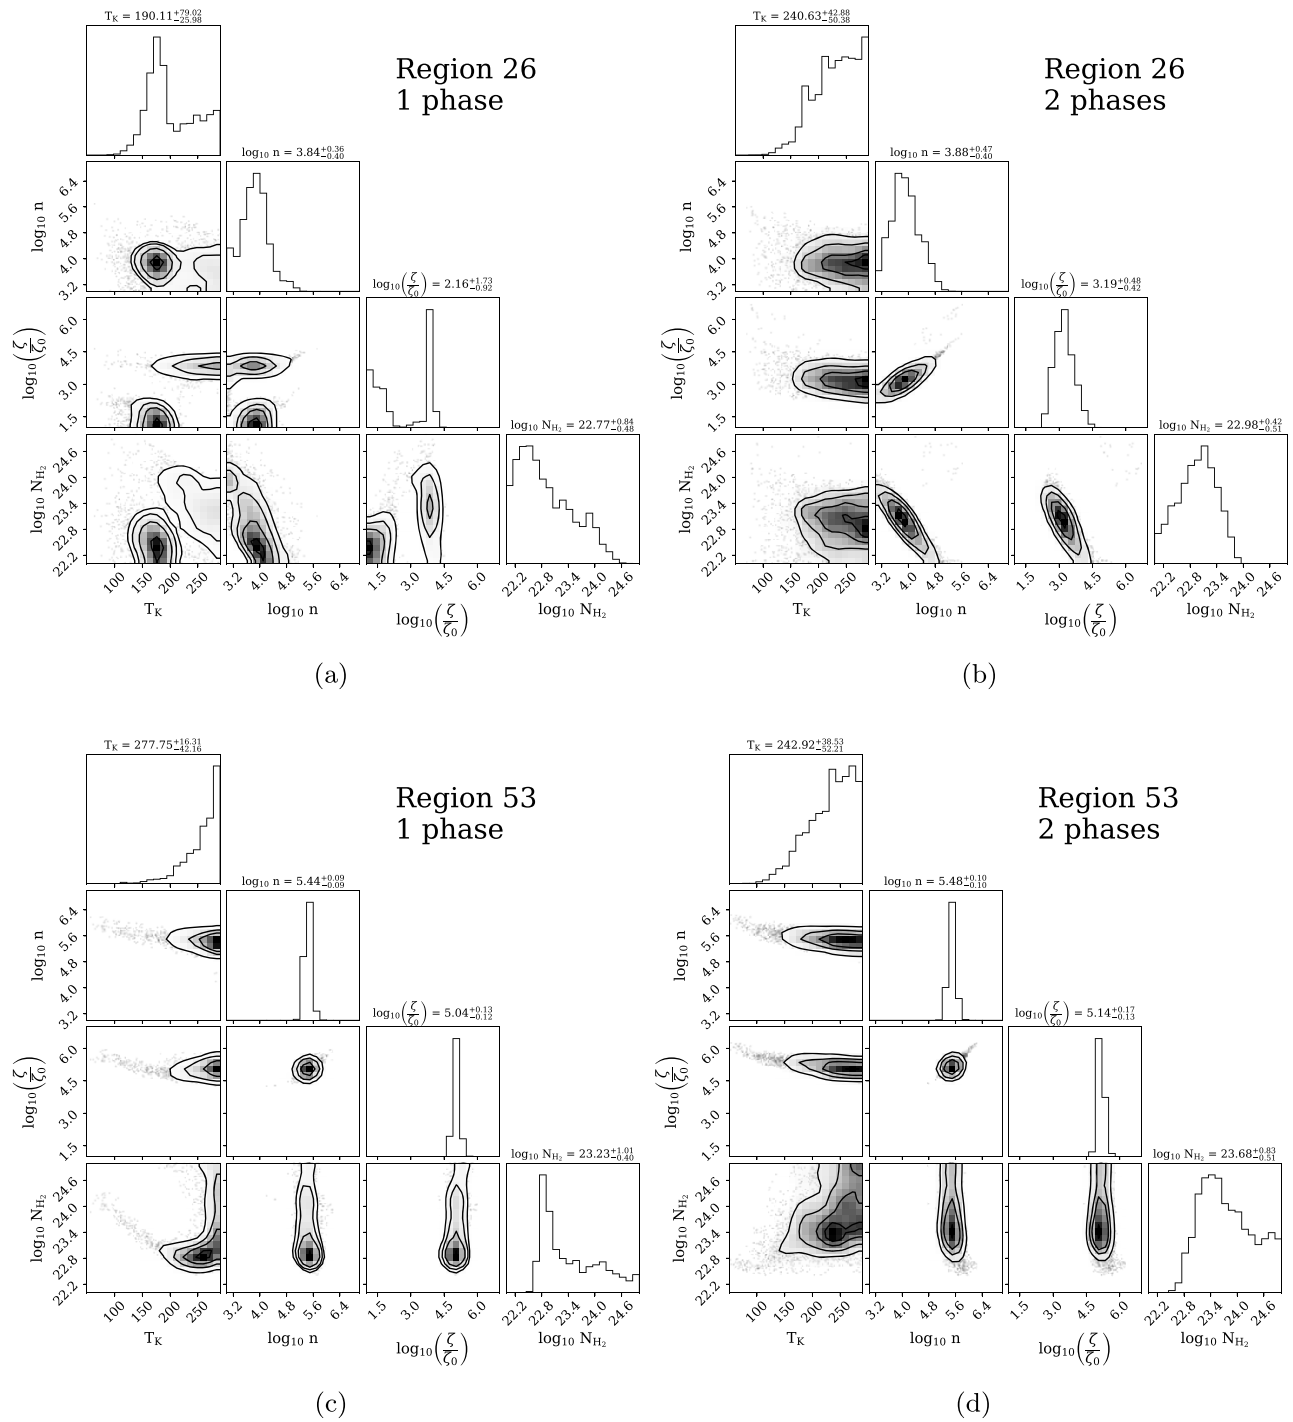


Figure B3. Posterior distributions for regions 26 (top) and 53 (bottom) using only one phase of chemical modeling (left column) vs. two phases (right column).

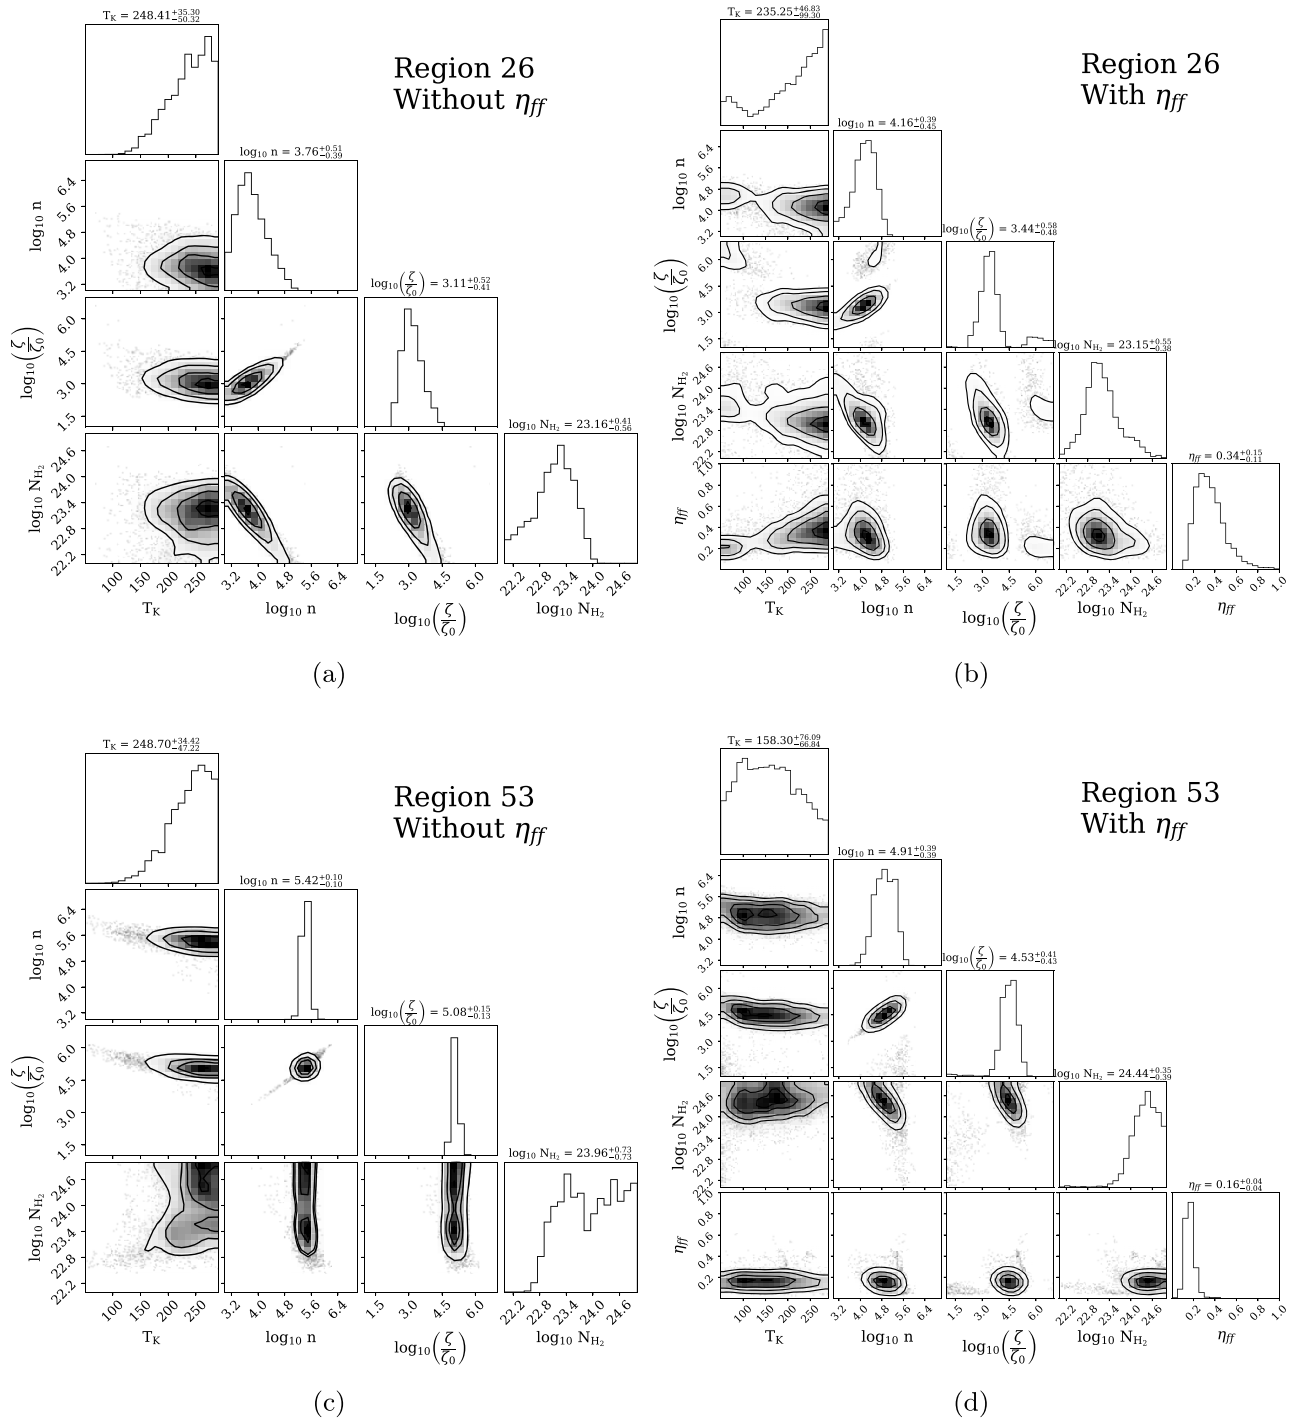


Figure B4. Posterior distributions for regions 26 (top) and 53 (bottom) before (left) and after (right) including a beam-filling factor as a free parameter in our modeling algorithm.

Appendix C New Collisional Rates Analysis

In this Appendix we compare published HCN and HNC collisional excitation rates used in radiative transfer and statistical equilibrium calculations to predict HCN and HNC abundances. The intent of this comparison is to highlight differences between excitation rates used for each isomer that might impact derived physical and chemical properties. In the following we highlight the salient features of the published HCN and HNC excitation rates (Table C1). We also provide a direct comparison to the two HCN and HNC excitation rate sets predominant in the astrophysical literature.

The He excitation rates published by F. Dumouchel et al. (2010) were considered the default HCN and HNC excitation rates before subsequent excitation rates involving H_2 as the collision partner were available. The F. Dumouchel et al. (2010) rates are available in the Leiden Atomic and Molecular Database (LAMDA; F. L. Schöier et al. 2005) where a scaling factor of 1.363 is used to correct these He-based collisional excitation rates by the difference in mass between He and H_2 .²⁸ This approximation thus makes these rates appropriate to use for modeling the interstellar medium where ground-state

²⁸ <https://home.strw.leidenuniv.nl/~moldata/>

Table C1
Published HCN and HNC Excitation Rates

Reference	Isomer(s)	Col. Partner	No. Energy Levels	T_K Range (K)
F. Dumouchel et al. (2010)	HCN/HNC	He	26	5–500
F. Dumouchel et al. (2011)	HNC	para-/ortho-H ₂	16	5–100
M. H. Vera et al. (2014)	HCN	para-/ortho-H ₂	13	5–100
M. Hernández Vera et al. (2017)	HCN/HNC	para-/ortho-H ₂	26	10–500

($j = 0$) para-H₂ is a common collision partner. F. Dumouchel et al. (2010) noted the following propensity rules for collisional de-excitation:

1. HCN–He favor even ΔJ transitions.
2. HNC–He favor odd ΔJ transitions.
3. Propensity favoring even/odd transitions in HCN and HNC decreases with increasing temperature.

M. H. Vera et al. (2014) and F. Dumouchel et al. (2011) published some of the first excitation rates for HCN and HNC, respectively, which used para- ($j = 0$ and 2) and ortho- ($j = 1$) H₂ as the collision partner. Both studies noted that the rate coefficients for collisions with different H₂ rotational states are very different. Collisions with para-H₂ ($j = 0$) are significantly smaller than those for collisions with ortho-H₂ ($j = 1$) and with excited para-H₂ ($j = 2$). For the HNC excitation rates, F. Dumouchel et al. (2011) noted that propensity rules in favor of $\Delta J = 1$ rates are larger than those involving $\Delta J = 2$ for all H₂ rotational states. For the HCN excitation rates, M. H. Vera et al. (2014) noted that propensity rules in favor of: (1) even ΔJ transitions for HCN in collisions with para-H₂ ($j = 0$) and (2) odd ΔJ transitions for HCN in collisions with excited para- and ortho-H₂ ($j \geq 1$).

The newest and most complete sets of collisional excitation rates for HCN and HNC involving collisions with para- and ortho-H₂ have been published by M. Hernández Vera et al. (2017). These studies extended the F. Dumouchel et al. (2011) HCN and M. H. Vera et al. (2014) HNC excitation rates to the first 26 rotational energy levels and to kinetic temperatures of 500 K for both molecules. M. Hernández Vera et al. (2017) compared their HCN and HNC para- and ortho-H₂ excitation rates with the He-scaled rates published by F. Dumouchel et al. (2010), noting differences of ~ 2 for HCN–H₂($J = 0$)/HCN–He.²⁹ M. Hernández Vera et al. (2017) also noted that collisions for HCN and HNC involving ortho-H₂ ($j = 1$) are ~ 2 –3 times larger than those involving para-H₂ ($j = 0$). Propensity rules derived by M. Hernández Vera et al. (2017) for HCN and HNC collisions with para- and ortho-H₂ are the same as those derived for collisions with He.

All of these studies have noted that collisional excitation of HCN is very different than that due to HNC, which should have a strong influence on the calculated intensities of these two isomers. We should also note the collisional excitation rates for the C, N, and H isotopologues of HCN and HNC published by D. Navarro-Almáida et al. (2023).

To provide comparison of the commonly used He excitation rates (F. Dumouchel et al. 2010) to the H₂-based excitation rates (M. Hernández Vera et al. 2017) used in the present analysis, we provide in the following a series of comparison diagrams meant to

illuminate the differences between these two sets of collisional excitation rates. These modeled intensities were calculated using both the “old” (He-based) and “new” (ortho+para-H₂-based) rates for the kinetic temperature T_K , H₂ volume density n_{H_2} , and H₂ column density N_{H_2} ranges provided in Table 2. Note that the CRIR was not involved in these calculations, and no beam-filling factor was applied. We combine N_{H_2} with X_{HCN} and X_{HNC} values in the range of $10^{-12} \leq X_{mol} \leq 10^{-6}$ to obtain molecular column densities N_{HCN} and N_{HNC} using $N_{mol} = 2 \times N_{H_2} \times X_{mol}$. This molecular column density calculation yields values in the range of $2 \times 10^{12} \leq N_{mol} \leq 2 \times 10^{19} \text{ cm}^{-2}$. Figures C1, C2, and C3 show a series of comparisons as a function of T_K , n_{H_2} , and N_{mol} , respectively, where the other two parameters in each figure are held constant at “low” (panel (a) in each figure) and “high” (panel (b) in each figure) values. A summary of these comparisons concludes that:

1. Use of the new ortho+para-H₂ excitation rates result in higher (by up to a factor of 2) predicted HCN and HNC integrated transition intensities.
2. Differences between predicted line intensities derived from the old (He-scaled) versus new (ortho+para-H₂) excitation rates are highest (\sim factor of 2) at low temperatures, volume densities, and column densities.
3. Differences between line intensities for each J transition diminish at higher volume and column densities for both sets of collisional excitation rates (see Figure C1).
4. At very high HCN and HNC column densities ($\gtrsim 10^{18} \text{ cm}^{-2}$), the predicted HCN and HNC transition intensities and opacities become so large that the SpectralRadex models do a poor job of predicting integrated intensities. Note that these HCN and HNC column densities are unrealistically large; more than a factor of 1000 larger than measured in the ISM.

We also test the differences between the F. Dumouchel et al. (2010) and M. Hernández Vera et al. (2017) rates in the context of our neural network + Bayesian nested sampling algorithm. The UltraNest inference results using the old and new rates are shown in Figure C4 for regions 26 and 53. Here, we include CRIR as a free parameter using the prior distribution provided in Table 2. Note that no beam-filling factor is applied in these tests. Overall, only minor differences result from the use of the new versus old excitation rates. Though panels (a) and (b) in Figure C4 for region 26 show a bimodality in the posterior distributions for H₂ volume density, CRIR, and H₂ column density, the median predicted values for these parameters are nearly identical using the new and old rates. The minimal differences between parameter values derived from the two sets of rates likely also result from higher predicted kinetic temperatures in both regions and higher H₂ volume and column density estimates in region 53. As noted above, the differences in intensities derived from the two sets of collisional excitation rates diminish at larger values of T_K , n_{H_2} , and N_{H_2} .

²⁹ In this comparison, we have applied the standard scaling factor of 1.363 to convert He excitation rates to para-H₂, which M. Hernández Vera et al. (2017) did not appear to apply.

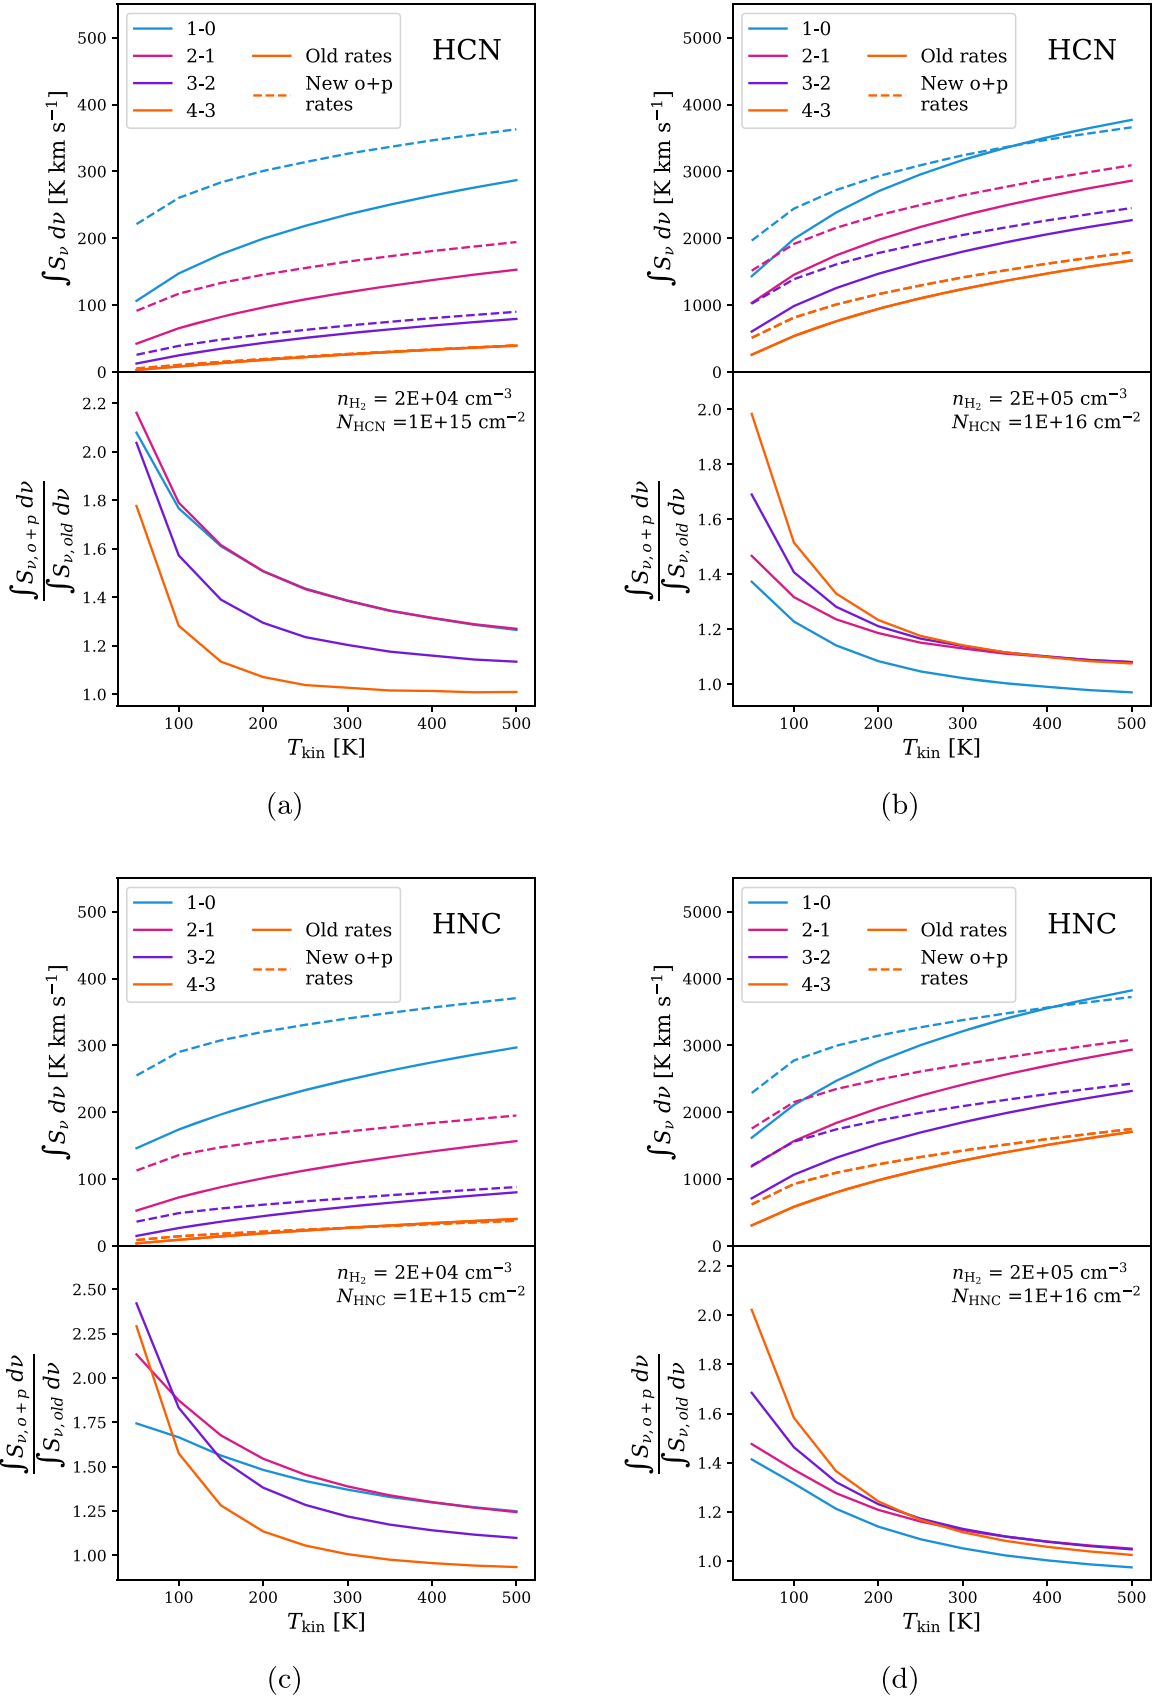


Figure C1. HCN ((a) and (b)) and HNC ((c) and (d)) integrated intensities (top panel in each pair) and integrated intensity ratios (new rates divided by old rates; bottom panels in each pair) as a function of kinetic temperature. Volume and column density are held constant at the listed values for each comparison.

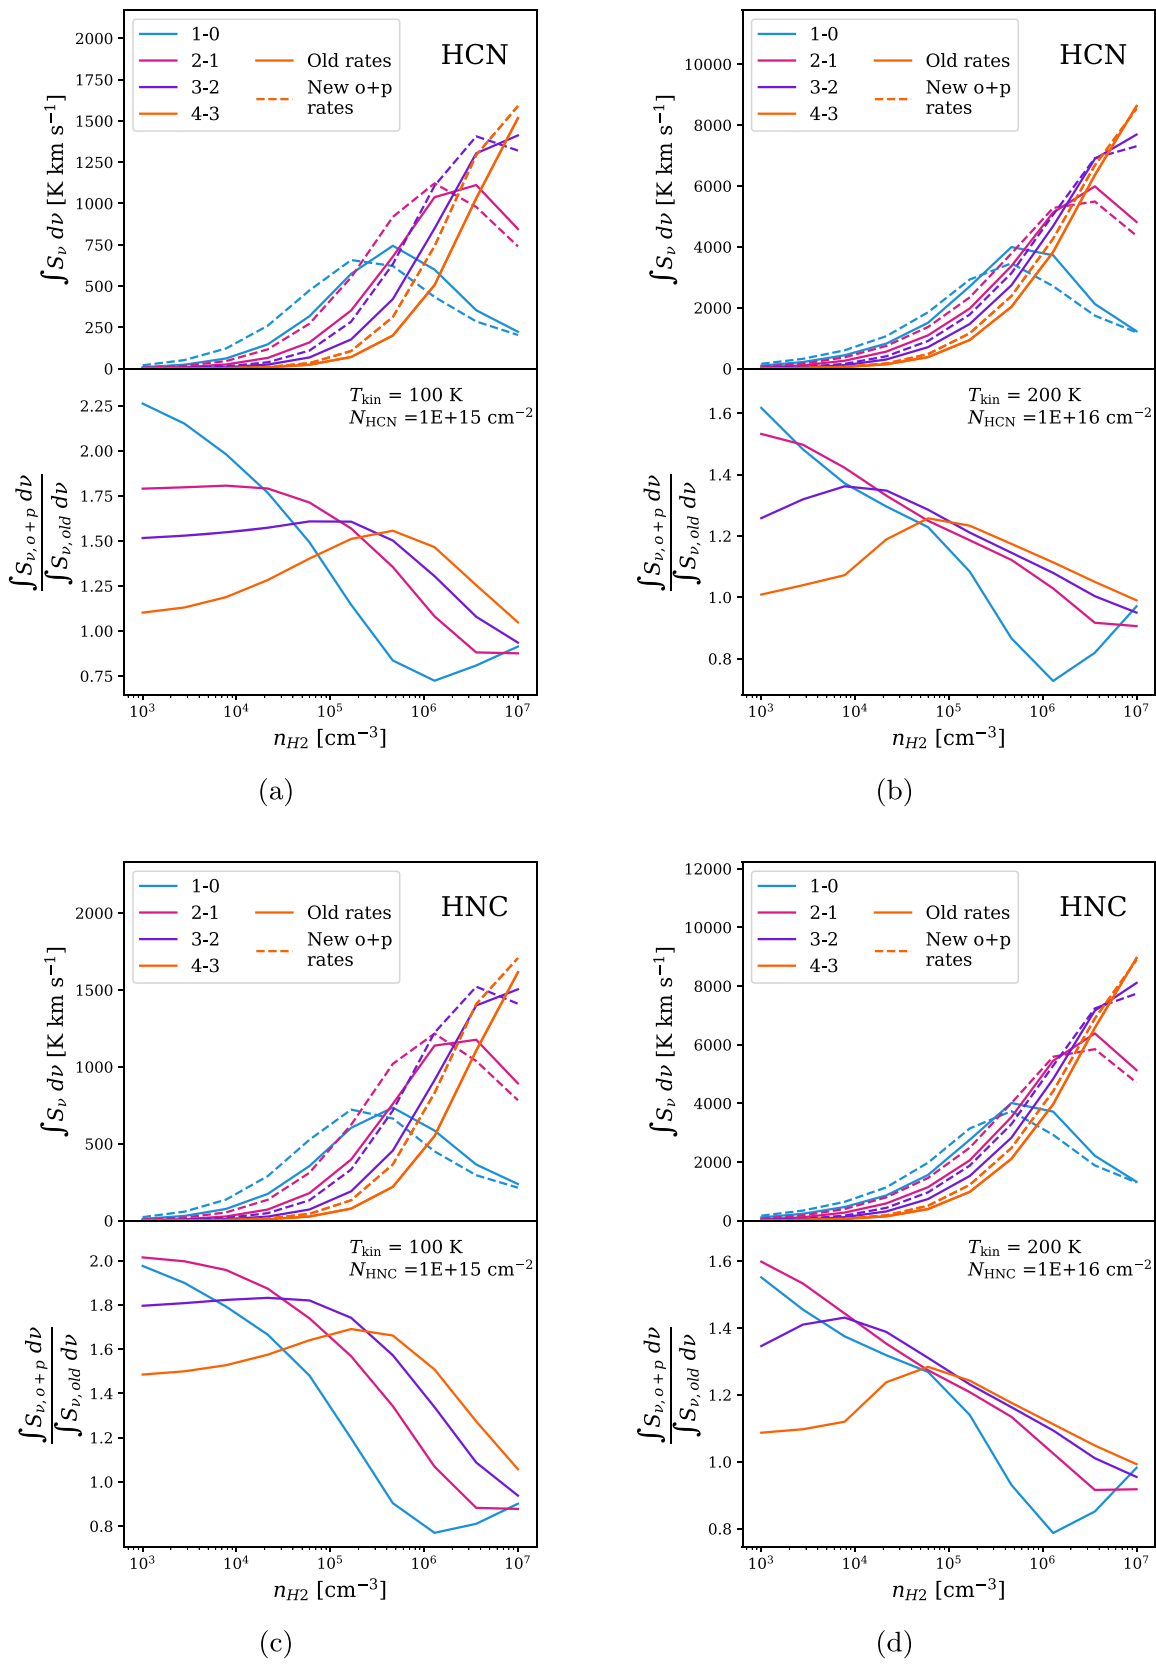


Figure C2. Same as in Figure C1 but plotted as a function of volume density, where kinetic temperature and column density are held constant.

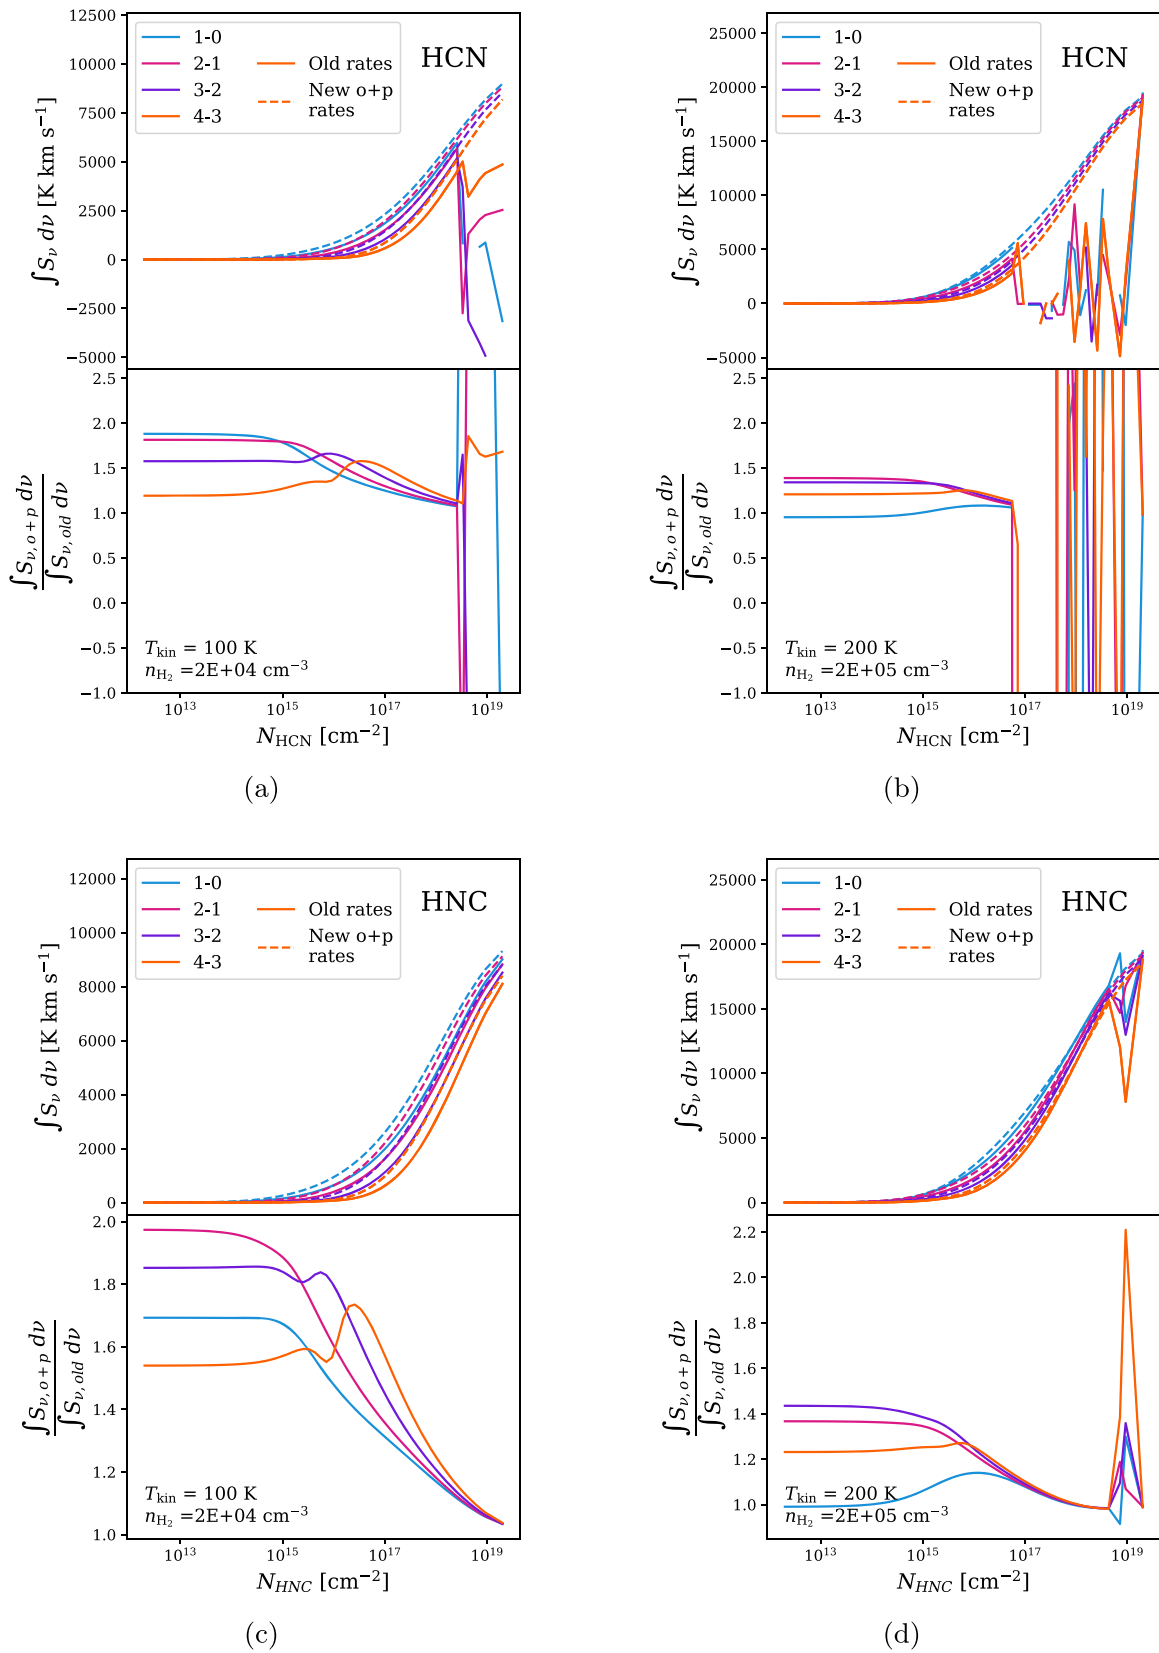


Figure C3. Same as in Figures C1 and C2 but plotted as a function of column density, where kinetic temperature and volume density are held constant.

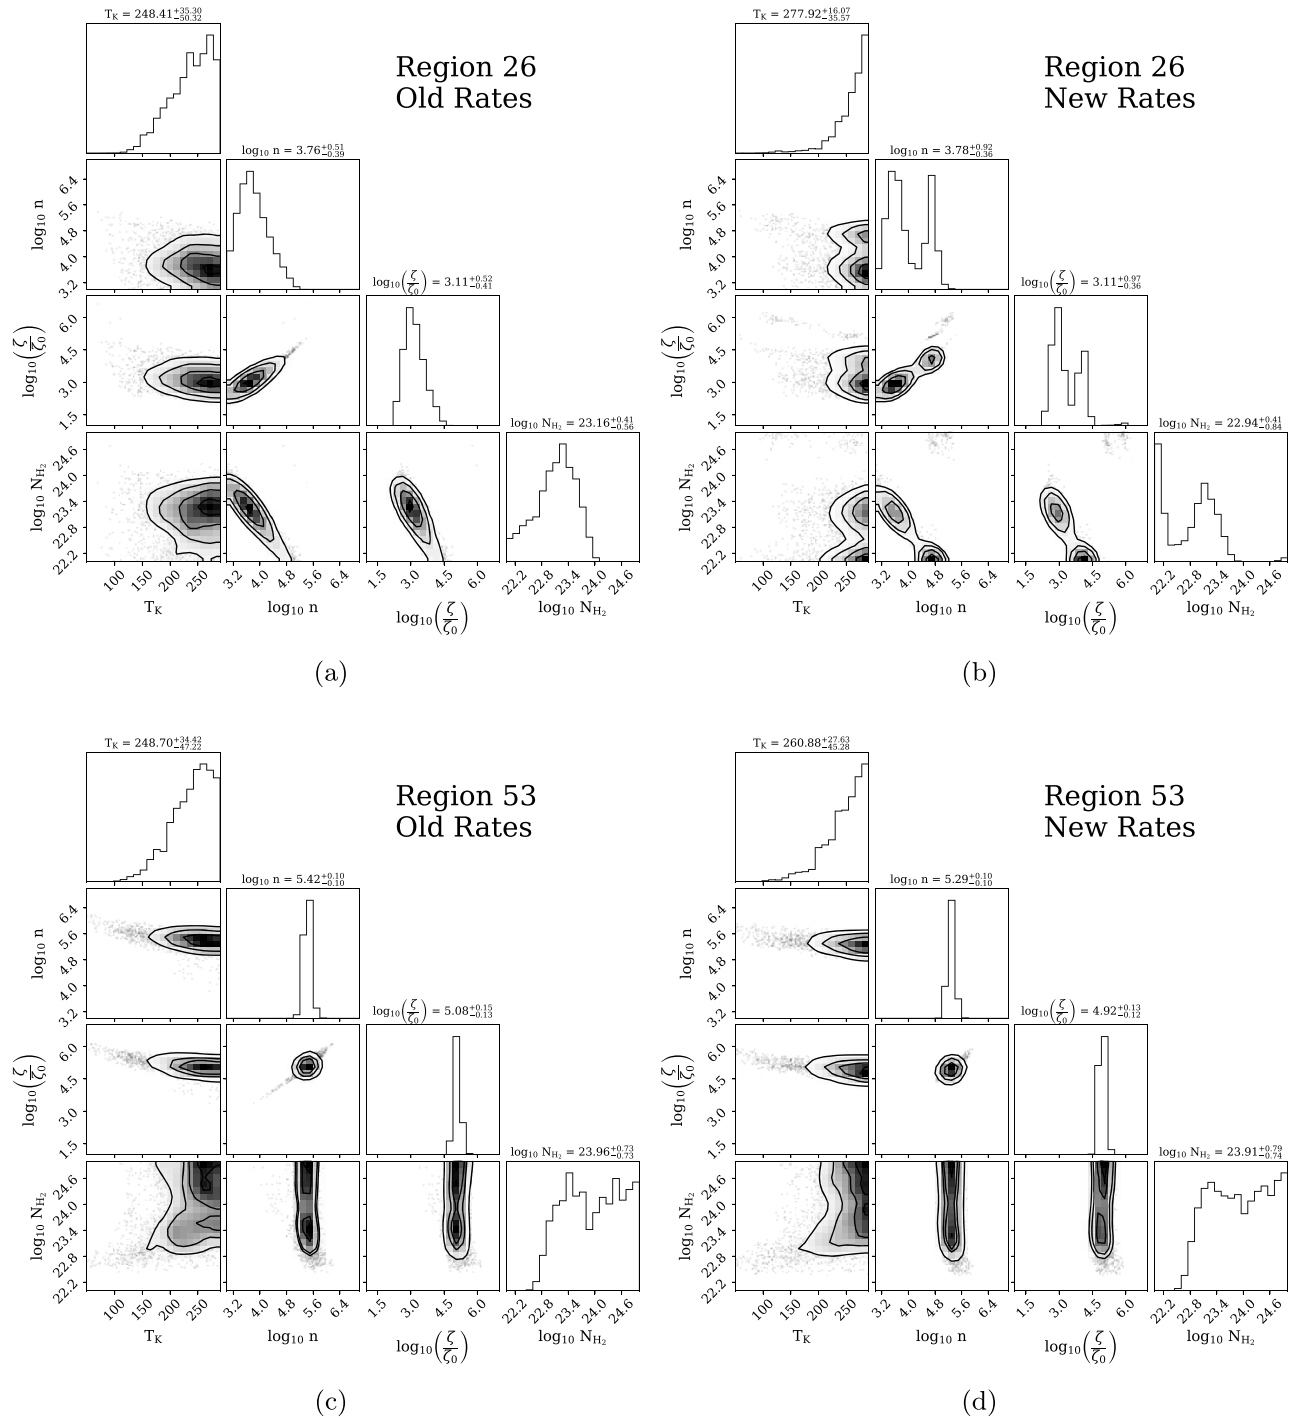


Figure C4. UltraNest inference results derived from HCN and HNC data for regions 26 (top) and 53 (bottom) using old (left) vs. new (right) HCN and HNC collisional excitation rates.

ORCID iDs

Erica Behrens <https://orcid.org/0000-0002-2333-5474>
 Jeffrey G. Mangum <https://orcid.org/0000-0003-1183-9293>
 Serena Viti <https://orcid.org/0000-0001-8504-8844>
 Jonathan Holdship <https://orcid.org/0000-0003-4025-1552>
 Ko-Yun Huang <https://orcid.org/0000-0002-1227-8435>
 Mathilde Bouvier <https://orcid.org/0000-0002-0370-8034>
 Joshua Butterworth <https://orcid.org/0000-0002-5353-1775>
 Cosima Eibensteiner <https://orcid.org/0000-0002-1185-2810>

Nanase Harada <https://orcid.org/0000-0002-6824-6627>
 Sergio Martín <https://orcid.org/0000-0001-9281-2919>
 Kazushi Sakamoto <https://orcid.org/0000-0001-5187-2288>
 Sebastien Muller <https://orcid.org/0000-0002-9931-1313>
 Kunihiko Tanaka <https://orcid.org/0000-0001-8153-1986>
 Laura Colzi <https://orcid.org/0000-0001-8064-6394>
 Christian Henkel <https://orcid.org/0000-0002-7495-4005>
 David S. Meier <https://orcid.org/0000-0001-9436-9471>
 Víctor M. Rivilla <https://orcid.org/0000-0002-2887-5859>
 Paul P. van der Werf <https://orcid.org/0000-0001-5434-5942>

References

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJL*, 709, L152
- Acero, F., Aharonian, F., Akhperjanian, A. G., et al. 2009, *Sci*, 326, 1080
- Ajello, M., Allafort, A., Baldini, L., et al. 2012, *ApJ*, 744, 80
- Armus, L., Heckman, T., & Miley, G. 1987, *AJ*, 94, 831
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167
- 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
- Bamba, A., & Williams, B. J. 2022, in *Handbook of X-Ray and Gamma-Ray Astrophysics*, ed. C. Bambi & A. Santangelo (Berlin: Springer), 77
- Bao, M., Harada, N., Kohno, K., et al. 2024, *A&A*, 687, A43
- Bayet, E., Williams, D. A., Hartquist, T. W., & Viti, S. 2011, *MNRAS*, 414, 1583
- Behrens, E. 2024, ebehrens97/HERA: ApJ Release, v1.0, Zenodo, doi:10.5281/zenodo.13839853
- Behrens, E., Mangum, J. G., Holdship, J., et al. 2022, *ApJ*, 939, 119
- Blasi, P., Amato, E., & Caprioli, D. 2007, *MNRAS*, 375, 1471
- Bolatto, A. D., Warren, S. R., Leroy, A. K., et al. 2013, *Natur*, 499, 450
- Bouvier, M., Viti, S., Behrens, E., et al. 2024, *A&A*, 689, A64
- Brunetti, N., Wilson, C. D., Sliwa, K., et al. 2021, *MNRAS*, 500, 4730
- Buchner, J. 2016, *S&C*, 26, 383
- Buchner, J. 2019, *PASP*, 131, 108005
- Buchner, J. 2021, *JOSS*, 6, 3001
- Butterworth, J., Viti, S., Van der Werf, P. P., et al. 2024, *A&A*, 687, A43
- Callanan, D., Longmore, S. N., Kruijssen, J. M. D., et al. 2021, *MNRAS*, 505, 4310
- CASA Team, Bean, B., Bhatnagar, S., et al. 2022, *PASP*, 134, 114501
- Cenci, E., Feldmann, R., Gensior, J., et al. 2024, *MNRAS*, 527, 7871
- Chevance, M., Kruijssen, J. M. D., Vazquez-Semadeni, E., et al. 2020, *SSRv*, 216, 50
- Chomiuk, L., & Povich, M. S. 2011, *AJ*, 142, 197
- Colzi, L., Fontani, F., Caselli, P., et al. 2018, *A&A*, 609, A129
- Corbelli, E., Braine, J., Bandiera, R., et al. 2017, *A&A*, 601, A146
- Downes, D., & Solomon, P. M. 1998, *ApJ*, 507, 615
- Dumouchel, F., Faure, A., & Lique, F. 2010, *MNRAS*, 406, 2488
- Dumouchel, F., Klos, J., & Lique, F. 2011, *PCCP*, 13, 8204
- Eibensteiner, C., Barnes, A. T., Bigiel, F., et al. 2022, *A&A*, 659, A173
- Feldmann, R., Quataert, E., Faucher-Giguère, C.-A., et al. 2023, *MNRAS*, 522, 3831
- Gaches, B. A. L., & Offner, S. S. R. 2018, *ApJ*, 861, 87
- Gong, M., Ostriker, E. C., Kim, C.-G., & Kim, J.-G. 2020, *ApJ*, 903, 142
- Gorski, M., Ott, J., Rand, R., et al. 2017, *ApJ*, 842, 124
- Haasler, D., Rivilla, V. M., Martín, S., et al. 2022, *A&A*, 659, A158
- Hacar, A., Bosman, A. D., & van Dishoeck, E. F. 2020, *A&A*, 635, A4
- Hanabata, Y., Katagiri, H., Hewitt, J. W., et al. 2014, *ApJ*, 786, 145
- Harada, N., Martín, S., Mangum, J. G., et al. 2021, *ApJ*, 923, 24
- Harada, N., Martín, S., Mangum, J. G., et al. 2022, *ApJ*, 938, 80
- Harada, N., Meier, D. S., Martín, S., et al. 2024, *ApJS*, 271, 38
- Hernández Vera, M., Lique, F., Dumouchel, F., Hily-Blant, P., & Faure, A. 2017, *MNRAS*, 468, 1084
- H.E.S.S. Collaboration, Abdalla, H., Aharonian, F., et al. 2018, *A&A*, 617, A73
- Heyer, M. H., & Brunt, C. M. 2004, *ApJL*, 615, L45
- Heyl, J., Butterworth, J., & Viti, S. 2023, *MNRAS*, 526, 404
- Holdship, J., Mangum, J. G., Viti, S., et al. 2022, *ApJ*, 931, 89
- Holdship, J., Viti, S., Jiménez-Serra, I., Makrymallis, A., & Priestley, F. 2017, *AJ*, 154, 38
- Holdship, J., Viti, S., Martín, S., et al. 2021, *A&A*, 654, A55
- Hollenbach, D. J., & Tielens, A. G. G. M. 1999, *RvMP*, 71, 173
- Huang, K. Y., Viti, S., Holdship, J., et al. 2023, *A&A*, 675, A151
- Humire, P. K., Henkel, C., Hernández-Gómez, A., et al. 2022, *A&A*, 663, A33
- Kim, J., Chevance, M., Kruijssen, J. M. D., et al. 2021, *MNRAS*, 504, 487
- Krieger, N., Bolatto, A. D., Leroy, A. K., et al. 2020, *ApJ*, 897, 176
- Lepp, S., & Dalgarno, A. 1996, *A&A*, 306, L21
- Leroy, A. K., Bolatto, A. D., Ostriker, E. C., et al. 2015, *ApJ*, 801, 25
- Leroy, A. K., Bolatto, A. D., Ostriker, E. C., et al. 2018, *ApJ*, 869, 126
- Leroy, A. K., Walter, F., Sandstrom, K., et al. 2013, *AJ*, 146, 19
- Levy, R. C., Bolatto, A. D., Leroy, A. K., et al. 2022, *ApJ*, 935, 19
- Li, H., & Chen, Y. 2010, *MNRAS*, 409, L35
- Licquia, T. C., & Newman, J. A. 2015, *ApJ*, 806, 96
- Loenen, A. F., Spaans, M., Baan, W. A., & Meijerink, R. 2008, *A&A*, 488, L5
- Mangum, J. G., Ginsburg, A. G., Henkel, C., et al. 2019, *ApJ*, 871, 170
- Martín, S., Mangum, J. G., Harada, N., et al. 2021, *A&A*, 656, A46
- Meijerink, R., Spaans, M., Loenen, A. F., & van der Werf, P. P. 2011, *A&A*, 525, A119
- Meng, F., Sánchez-Monge, Á., Schilke, P., et al. 2019, *A&A*, 630, A73
- Mills, E. A. C., Gorski, M., Emig, K. L., et al. 2021, *ApJ*, 919, 105
- Murphy, E. J., Dong, D., Leroy, A. K., et al. 2015, *ApJ*, 813, 118
- Navarro-Almaida, D., Bop, C. T., Lique, F., et al. 2023, *A&A*, 670, A110
- Padoan, P., & Nordlund, Å. 2002, *ApJ*, 576, 870
- Padovani, M., Bialy, S., Galli, D., et al. 2022, *A&A*, 658, A189
- Padovani, M., Hennebelle, P., Marcowith, A., & Ferrière, K. 2015, *A&A*, 582, L13
- Padovani, M., Marcowith, A., Hennebelle, P., & Ferrière, K. 2016, *A&A*, 590, A8
- Padovani, M., Marcowith, A., Sánchez-Monge, Á., Meng, F., & Schilke, P. 2019, *A&A*, 630, A72
- Papadopoulos, P. P. 2010, *ApJ*, 720, 226
- Phan, V. H. M., Peretti, E., Cristofari, P., Gusdorf, A., & Mertsch, P. 2024, *MNRAS*, 531, 2930
- Ravikularaman, S., Recchia, S., Phan, V. H. M., & Gabici, S. 2024, arXiv:2406.15260
- Rekola, R., Richer, M. G., McCall, M. L., et al. 2005, *MNRAS*, 361, 330
- Renaud, F., Segovia Otero, Á., & Agertz, O. 2022, *MNRAS*, 516, 4922
- Rico-Villas, F., Martín-Pintado, J., González-Alfonso, E., Martín, S., & Rivilla, V. M. 2020, *MNRAS*, 491, 4573
- Rosenberg, M. J. F., van der Werf, P. P., & Israel, F. P. 2013, *A&A*, 550, A12
- Sakamoto, K., Mao, R.-Q., Matsushita, S., et al. 2011, *ApJ*, 735, 19
- Santa-Maria, M. G., Goicoechea, J. R., Pety, J., et al. 2023, *A&A*, 679, A4
- Sarbadhicary, S. K., Badenes, C., Chomiuk, L., Caprioli, D., & Huizenga, D. 2017, *MNRAS*, 464, 2326
- Scherer, K., Fichtner, H., Ferreira, S. E. S., Büsching, I., & Potgieter, M. S. 2008, *ApJL*, 680, L105
- Schinnerer, E., Böker, T., Emsellem, E., & Downes, D. 2007, *A&A*, 462, L27
- Schinnerer, E., & Leroy, A. K. 2024, *ARA&A*, 62, 369
- Schöier, F. L., van der Tak, F. F. S., van Dishoeck, E. F., & Black, J. H. 2005, *A&A*, 432, 369
- Semenov, V. A., Kravtsov, A. V., & Caprioli, D. 2021, *ApJ*, 910, 126
- Tanaka, K., Mangum, J. G., Viti, S., et al. 2024, *ApJ*, 961, 18
- Turner, J. L., & Ho, P. T. P. 1985, *ApJL*, 299, L77
- Ulvestad, J. S., & Antonucci, R. R. J. 1997, *ApJ*, 488, 621
- Vera, M. H., Kalugina, Y., Denis-Alpizar, O., Stoecklin, T., & Lique, F. 2014, *JChPh*, 140, 224302
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, 17, 261