
Carbox: an end-to-end differentiable astrochemical simulation framework

Gijs Vermariën^{*1,2} Tommaso Grassi^{3,4} Marie Van de Sande¹
Serena Viti^{1,5,6} Stefano Bovino^{7,8,9} Alessandro Lupi¹⁰
Alexander Ruf^{11,4} Lorenzo Branca¹² Catherine Walsh¹³

¹Leiden Observatory ²SURF ³MPE ⁴Origins
⁵University of Bonn ⁶UCL ⁷University of Rome Sapienza
⁸Universidad de Concepción ⁹INAF Arcetri ¹⁰Università degli Studi dell’Insubria
¹¹LMU Munich ¹²Heidelberg University ¹³University of Leeds
* `vermarien@strw.leidenuniv.nl`

Abstract

Since the first observations of interstellar molecules, astrochemical simulations have been employed to model and understand its formation and destruction pathways. With the advent of high-resolution telescopes such as JWST and ALMA, the number of detected molecules has increased significantly, thereby creating a need for increasingly complex chemical reaction networks. To model such complex systems, we have developed Carbox, a new astrochemical simulation code that leverages the modern high-performance transformation framework Jax. With Jax enabling computational efficiency and differentiability, Carbox can easily utilize GPU acceleration, be used to study sensitivity and uncertainty, and interface with advances in Scientific Machine Learning. All of these features are crucial for modeling the molecules observed by current and next-generation telescopes.

1 Astrochemical reaction networks

The practice of simulating molecules in space has been a key tool since their earliest observations. However, with increasing resolution in both space and frequency over the entire electromagnetic spectrum, the number of molecular detections has risen sharply to 340¹. Likewise, the modeling of reaction networks has evolved from relatively simple hydrogen-based primordial (early universe) chemistry to increasingly complex organic molecular species that are the precursors to biological life. Many codes that evolve the corresponding differential equations to model the abundances of chemical species as a function of time and space have been developed in FORTRAN and C, with some exceptions in Python. The most notable are UCLCHEM[Holdship et al., 2021], Nautilus[Ruaud et al., 2016], and Magickal[Garrod et al., 2022], which focus solely on the chemical evolution, and KROME[Grassi et al., 2014], PRIZMO[Grassi et al., 2020], and Grackle[Smith et al., 2017], which also consider time-dependent coupled thermal processes. Although these codes are applied to different environments and chemical conditions, and may contain various chemical networks that differ in terms of size and composition, they are all based on standard implicit numerical ordinary differential equation (ODE) solvers.

The primary numerical integration method is the Backward Differentiation Formula (BDF) [Byrne and Hindmarsh, 1975], which is well-suited for solving (very) stiff problems, i.e., the ODEs that dictate the chemistry. Apart from some outliers, this integration method dominated the astrochemical codes, with some variants, from the initial implementations (GEAR Hindmarsh [1972]) to modern SUNDIALS [Byrne et al., 2024]. These methods have improved by exploiting numerical efficiency,

¹<https://cdms.astro.uni-koeln.de/classic/molecules>

for example, DLSODES, which takes advantage of the Jacobian sparsity of the chemical ODEs, or simply by leveraging the capabilities of newer and improved CPUs.

Conversely, the few examples of GPU-based solvers are subdominant and limited by memory, by well-established computational techniques available to CPU users, or by the limited coding capabilities and human resources within the astrochemical community.

Recently, scientific machine learning (SciML) approaches to evolve astrochemical systems have been explored by several groups [Grassi et al., 2021, Holdship et al., 2021, Branca and Pallottini, 2022, Tang and Turk, 2022, Sulzer and Buck, 2023, Branca and Pallottini, 2024, Maes et al., 2024, Vermariën et al., 2025], however, the stiffness, the aforementioned non-linearity, and the high dimensionality of chemical systems have limited the applicability of these methods to realistic observed environments. In fact, a more effective use of (SciML) has been the coupling of classical time-integration codes (or methods) to statistical inference pipelines [Holdship et al., 2022, Heyl et al., 2022, Behrens et al., 2024], used to analyze the sensitivity of both chemical and physical parameters Heyl et al. [2023], Vermariën et al. [2025], Grassi et al. [2025]. This approach enables us to utilize statistical methods to predict the input parameters of chemical systems, including their physical boundary conditions. These methods try to determine how the output of a chemical model is influenced by any input parameter, either chemical or physical, sometimes trying to minimize the output in relation to the corresponding observable counterpart. Fully differentiable modeling can help both with the forward sensitivity analysis and enable better inverse parameter studies.

Thanks to the Jax framework [Bradbury et al., 2018], Carbox² is a fully differentiable open-source astrochemical simulation framework that can simulate the evolution of chemical species. By utilizing the power of Jax, Carbox has native support for GPU acceleration and numerical automatic differentiation. Carbox is thus interoperable with SciML within the Jax ecosystem, allowing easier development of Machine Learning (ML) methods and even classical-ML hybrid solutions. It also means that Carbox can seamlessly be integrated with other emerging astrophysical codes implemented in Jax, such as radiative transfer and hydrodynamics [Storeks and Buck, 2024].

2 Simulating astrochemical reaction networks

A chemical kinetic reaction network consists of a set of chemical species connected by a set of reactions. In astrochemistry, the reactions, their reaction mechanism, and the reaction rate coefficients (k_j) can be inferred from theoretical studies, observations, or laboratory measurements. The mathematical counterpart of a chemical network is a set of differential equations that describes the evolution in the number density $n_i(t)$ in cm^{-3} of each chemical species. The numerical problem becomes to solve an initial value problem (IVP), where, if we omit reactions with order greater than two, we have

$$\frac{dn_i}{dt} = \pm \sum_j k_j n_{s_1 \in j} \pm \sum_j k_j n_{s_1 \in j} n_{s_2 \in j},$$

where the subscripts $s_1 \in j$ and $s_2 \in j$ indicate respectively the first and the second product or reactant involved in the j th reaction; both the unimolecular and bimolecular reactions lead to the production (+) and destruction (-) of the species i . The ODE system is integrated in time, given some initial conditions $n_i(t=0)$ for each i th species.

There are several classes of rate reaction coefficients, and their exact expression depends on the chemistry involved [Bovino and Grassi, 2023]. Their description is beyond the aims of this paper; nevertheless, to illustrate the nature of our systems, we provide two notable examples. For a two-body reaction in the gas phase, the reaction rate is described in general by a temperature-dependent function, and often by the Kooji-Arrhenius equation,

$$k = \alpha T_{300}^{\beta} \exp(-\gamma/T),$$

with α , β , and γ specific reaction parameters, and T_{300} the temperature in units of 300 K. On the other hand, reaction coefficients can be functions of parameters other than temperature, depending on the specific core process or environmental conditions, as ultraviolet and X-ray photons, and cosmic rays, which ionize or dissociate species with a rate that usually depends on their energy source. For example, for the direct cosmic-rays ionization ($X + \text{CR} \rightarrow X^+ + e^-$) we have

$$k = \alpha \zeta,$$

²<https://github.com/gijsvermarien/carbox>

with α a parameter depending on the given chemical species (X), and ζ the ionization rate per H_2 molecule, a value that depends on the modeled astrophysical environment, and which plays a crucial role in the evolution of some observed chemical species.

This suggests that the ODE system can be generalized as $\dot{n}(t) = f(n(t), k(\varphi), \varphi)$, where n are all the chemical abundances, k the rate coefficients, and φ the environment parameters (e.g., T , ζ or F_{UV}). This highlights the strong dependence of astrochemical systems on both the physical conditions and the reaction networks, as well as their strong nonlinearity when evolved over time.

3 Sensitivity analysis of astrochemical models

Astrochemical models are often used in a deterministic manner, however, in many realistic applications, it is necessary to determine the effect of varying these parameters. In addition, several of these parameters have intrinsic (large) uncertainties that play a key role in shaping the scientific interpretation of the results. This includes the aforementioned physical input parameters, as well as the rate coefficients k . Assessing the effect of uncertainties in the model input on the model predictions has been done for a variety of astrochemical environments, spanning all stages of stellar and galactic evolution [Wakelam et al., 2005, Vasyunin et al., 2008, Dobrijevic et al., 2010, Penteado et al., 2017, Maes et al., 2023].

The uncertainties on the rate coefficients, i.e., the kinetic data that drive the chemical evolution, can lead to large errors in the abundances predicted by the model. Fig. 3 shows the predicted abundances of a set of molecules throughout a stellar outflow, which is a relatively standard astrochemical benchmark. For some species, the uncertainty in the predicted abundances is limited. For others, the uncertainty spans several orders of magnitude. This has important implications for the interpretation of astrophysical observations, as some species are more reliable probes of the physical conditions than others, and relatively minor discrepancies might lead to different physical interpretations. It also has repercussions for astrochemical model development, as adding chemical or physical complexity might not yield significantly different results, hence over-complicating the model [de Sande et al., submitted]. Finally, determining which reactions’ uncertainties have the largest impact on the predicted chemistry is a crucial step in planning costly computational quantum chemistry calculations and laboratory experiments.

A modern and effective approach to these problems is to use differentiable models, which enable us to more directly assess both the forward and inverse problems. One of the main aims of Carbox is to be a differentiable tool that enables uncertainty and sensitivity analysis with respect to any of the input parameters, both physical and chemical.

4 Carbox

Carbox aims to provide a modern open-source implementation of an astrochemical simulation framework, with a core routine that can build a system of differential equations based on the species and reactions in a given arbitrary astrochemical network. The core routine relies on these species and reactions to have a uniform format, also providing the capability to parse three existing simulation frameworks: PRIZMO[Grassi et al., 2020], UMIST[Millar et al., 2024], and UCLCHEM[Holdship et al., 2017]. These parsers can be easily extended to different reaction database formats and custom rate equations. After parsing the species and reactions and providing them to Carbox, the framework constructs the incidence matrix to enable the effective time integration of the corresponding ODE system. The result of the parsing is a *Network* object that provides easy user interaction and a mirrored *JNetwork* object that represents the mathematical evolution of the chemical network and is compiled by and interoperable with Jax. Support is added for the vectorization of rates with identical mechanisms, as well as for sparsity, providing a tradeoff between memory footprint and computational speed, which is needed for large ($\gtrsim 1000$ reactions) chemical networks. These networks can then be integrated using DiffraX[Kidger, 2022], specifically, with the 5th order Kvaerno stiff integrator and custom tolerances of $a_{\text{tol}} = 10^{-9}$ and $r_{\text{tol}} = 10^{-30}$ providing reliable results [Kværnø, 2004].

To verify the robustness of Carbox, we benchmark it against a gas-phase-only setup of UCLCHEM, using two different networks: a relatively small minimal network with only atomic species and a more extensive network that includes molecules such as methanol (CH_3OH). Given the stiffness of the corresponding ODEs, we can generalize the accuracy of the results to more complicated

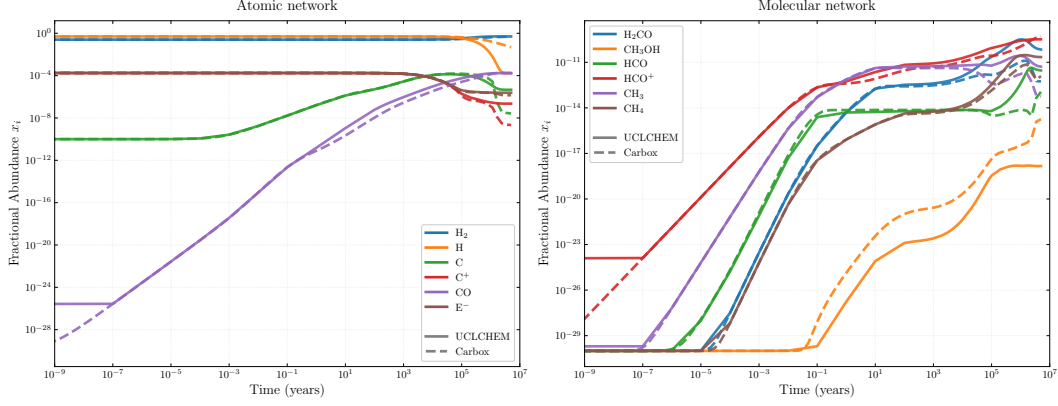


Figure 1: Temporal evolution of the fractional abundances of some selected chemical species for the atomic (left) and molecular (right) reaction networks. The solid line represents the reference integration using the UCLCHEM gas-phase reaction network, while the dashed lines show the results obtained with Carbox. The results are similar for the species not displayed here.

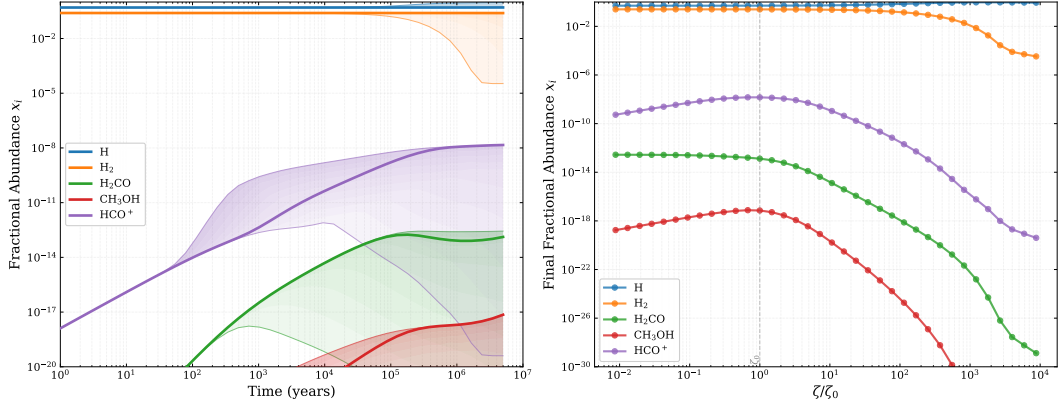


Figure 2: Left: temporal evolution of selected chemical species. The solid line represents the reference evolution ($\zeta/\zeta_0 = 1$), while the shaded area is the maximum extent given by $\zeta/\zeta_0 \in [10^{-2}, 10^4]$. Right: the final ($t = 10^7$ yr) abundances of the same species for different cosmic-ray ionization rate values.

chemical networks. The resulting simulation can be seen in Fig. 1. Showing we can reliably simulate the reaction network. It also highlights the non-linear nature of astrochemical models; small implementation differences can result in order-of-magnitude differences. However, since the ice chemistry is not included, the model provides a feasible approximation of the imperfect gas-phase only abundances.

One of the goals of Carbox is to assist in sensitivity studies. We test the impact of varying the cosmic ray ionization rate ζ , denoted in units of the galactic cosmic ray ionization rate $\zeta_0 = 1.3 \cdot 10^{-17} \text{ s}^{-1}$, on the evolution of different molecular species. This parameter is present in a subset of only 20 out of 2227 reactions involving a subset of chemical species. As expected, varying the ionization efficiency determines a considerable variation in the temporal evolution (Fig. 2, left panel) and in the final ($t = 10^7$ yr) abundances (right panel), both in the directly-affected chemical species (e.g. H_2), but also on the species that are modified indirectly in a non-linear fashion (e.g., H_2CO), being the latter usually difficult to predict even in relatively simple chemical networks.

5 Conclusion

We introduced the open-source Carbox simulation framework, a differentiable Jax-based framework to integrate the ODE system corresponding to astrochemical networks. Our code reproduces the results of well-established astrochemical codes, but it enables differentiability and GPU acceleration; a crucial feature that enables computationally efficiently solving forward and inverse problems.

Acknowledgments

The authors declare no competing interests. This work is financially supported by the advanced ERC grant (ID: 833460, PI: Serena Viti) within the framework of MOleculEs as Probes of the Physics of EXternal (MOPPEX) project. C.W. acknowledges financial support from the Science and Technology Facilities Council and UK Research and Innovation (grant numbers ST/X001016/1 and MR/T040726/1). MVdS thanks the Oort Fellowship at Leiden Observatory. This research was supported by the Munich Institute for Astro-, Particle and BioPhysics (MIAPbP) which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC-2094 – 390783311. SB acknowledges BASAL Centro de Astrofísica y Tecnologías Afines (CATA), project number AFB-17002.

Full affiliations

1. Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands
2. SURF, Amsterdam, The Netherlands
3. Center for Astrochemical Studies, Max Planck Institut für Extraterrestrische Physik, Garching, Germany
4. Exzellenzcluster Origins, Boltzmannstr. 2, 85748 Garching, Germany
5. Transdisciplinary Research Area (TRA) ‘Matter’/Argelander-Institut für Astronomie, University of Bonn, Bonn, Germany
6. Department of Physics and Astronomy, University College London, Gower Street, London, UK
7. Department of Chemistry, University of Rome Sapienza, P.le A. Moro 5, 00185 Rome, Italy
8. Departamento de Astronomía, Facultad de Ciencias Físicas y Matemáticas, Universidad de Concepción, Av. Esteban Iturra s/n Barrio Universitario, Casilla 160, Concepción, Chile
9. INAF, Osservatorio Astrofisico di Arcetri, Largo E. Fermi 5, I-50125 Firenze, Italy
10. Como Lake Center for Astrophysics, DiSAT, Università degli Studi dell’Insubria, via Valleggio 11, 22100 Como, Italy
11. LMU Munich, Faculty of Physics, Schellingstraße 4, 80799 Munich, Germany
12. Interdisciplinary Center for Scientific Computing, Heidelberg University, Heidelberg, Germany
13. School of Physics and Astronomy, University of Leeds, Leeds LS2 9JT, UK

References

- Erica Behrens, Jeffrey G. Mangum, Serena Viti, Jonathan Holdship, Ko-Yun Huang, Mathilde Bouvier, Joshua Butterworth, Cosima Eibensteiner, Nanase Harada, Sergio Martín, Kazushi Sakamoto, Sebastien Muller, Kunihiko Tanaka, Laura Colzi, Christian Henkel, David S. Meier, Víctor M. Rivilla, and Paul P. van der Werf. Neural Network Constraints on the Cosmic-Ray Ionization Rate and Other Physical Conditions in NGC 253 with ALCHEMI Measurements of HCN and HNC. *The Astrophysical Journal*, 977:38, December 2024. ISSN 0004-637X. doi: 10.3847/1538-4357/ad85db. URL <https://ui.adsabs.harvard.edu/abs/2024ApJ...977...38B>. Publisher: IOP ADS Bibcode: 2024ApJ...977...38B.
- Stenafo Bovino and Tommaso Grassi. *ASTROCHEMICAL MODELLING practical aspects of microphysics in numerical*. ELSEVIER - HEALTH SCIENCE, S.I., 2023. ISBN 978-0-323-97257-4. OCLC: 1410622781.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Nectra, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: Composable transformations of Python+NumPy programs, 2018.
- L. Branca and A. Pallottini. Neural networks: Solving the chemistry of the interstellar medium. *Monthly Notices of the Royal Astronomical Society*, 518(4):5718–5733, December 2022. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stac3512.
- Lorenzo Branca and Andrea Pallottini. Emulating the interstellar medium chemistry with neural operators, February 2024.
- Alex N. Byrne, Ci Xue, Troy Van Voorhis, and Brett A. McGuire. Sensitivity Analysis of Aromatic Chemistry to Gas-Phase Kinetics in a Dark Molecular Cloud Model, October 2024. URL <http://arxiv.org/abs/2410.09212>. arXiv:2410.09212 [astro-ph].

- G. D. Byrne and A. C. Hindmarsh. A Polyalgorithm for the Numerical Solution of Ordinary Differential Equations. *ACM Transactions on Mathematical Software*, 1(1):71–96, March 1975. ISSN 0098-3500, 1557-7295. doi: 10.1145/355626.355636.
- M Van de Sande, M Gueguen, T Danilovich, and T J Millar. Sense and Sensitivity - I. Uncertainty analysis of the gas-phase chemistry in AGB outflows. submitted.
- M. Dobrijevic, E. Hébrard, S. Plessis, N. Carrasco, P. Pernot, and M. Bruno-Claeys. Comparison of methods for the determination of key reactions in chemical systems: Application to Titan’s atmosphere. *Advances in Space Research*, 45(1):77–91, January 2010. ISSN 0273-1177. doi: 10.1016/j.asr.2009.06.005. URL <https://www.sciencedirect.com/science/article/pii/S0273117709003834>.
- Robin T. Garrod, Miwha Jin, Kayla A. Matis, Dylan Jones, Eric R. Willis, and Eric Herbst. Formation of Complex Organic Molecules in Hot Molecular Cores through Nondiffusive Grain-surface and Ice-mantle Chemistry. *The Astrophysical Journal Supplement Series*, 259(1):1, February 2022. ISSN 0067-0049. doi: 10.3847/1538-4365/ac3131.
- T. Grassi, S. Bovino, D. R. G. Schleicher, J. Prieto, D. Seifried, E. Simoncini, and F. A. Gianturco. KROME - a package to embed chemistry in astrophysical simulations. *Monthly Notices of the Royal Astronomical Society*, 439(3):2386–2419, March 2014. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stu114.
- T. Grassi, B. Ercolano, L. Szűcs, J. Jennings, and G. Picogna. Modelling thermochemical processes in protoplanetary discs I: Numerical methods. *Monthly Notices of the Royal Astronomical Society*, 494:4471–4491, May 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa971.
- T. Grassi, F. Nauman, J. P. Ramsey, S. Bovino, G. Picogna, and B. Ercolano. Reducing the complexity of chemical networks via interpretable autoencoders, October 2021.
- T. Grassi, M. Padovani, D. Galli, N. Vaytet, S. S. Jensen, E. Redaelli, S. Spezzano, S. Bovino, and P. Caselli. Mapping Synthetic Observations to Prestellar Core Models: An Interpretable Machine Learning Approach, February 2025.
- Johannes Heyl, Elena Sellentin, Jonathan Holdship, and Serena Viti. Identifying the most constraining ice observations to infer molecular binding energies, September 2022. URL <http://arxiv.org/abs/2209.09347>. arXiv:2209.09347 [astro-ph].
- Johannes Heyl, Serena Viti, and Gijs Vermariën. A statistical and machine learning approach to the study of astrochemistry. *Faraday Discussions*, 245(0):569–585, September 2023. ISSN 1364-5498. doi: 10.1039/D3FD00008G. URL <https://pubs.rsc.org/en/content/articlelanding/2023/fd/d3fd00008g>. Publisher: The Royal Society of Chemistry.
- A. C. Hindmarsh. GEAR: Ordinary differential equation system solver. Technical Report UCID–30001(Rev.2), California Univ., Livermore (USA). Lawrence Livermore Lab., August 1972.
- J. Holdship, S. Viti, I. Jiménez-Serra, A. Makrymallis, and F. Priestley. UCLCHEM: A Gas-grain Chemical Code for Clouds, Cores, and C-Shocks. *The Astronomical Journal*, 154(1):38, July 2017. ISSN 1538-3881. doi: 10.3847/1538-3881/aa773f. URL <https://iopscience.iop.org/article/10.3847/1538-3881/aa773f>.
- J. Holdship, S. Viti, T. J. Haworth, and J. D. Ilee. Chemulator: Fast, accurate thermochemistry for dynamical models through emulation. *Astronomy & Astrophysics*, 653:A76, September 2021. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202140357.
- Jonathan Holdship, Jeffrey G. Mangum, Serena Viti, Erica Behrens, Nanase Harada, Sergio Martín, Kazushi Sakamoto, Sebastien Muller, Kunihiko Tanaka, Kouichiro Nakanishi, Rubén Herrero-Illana, Yuki Yoshimura, Rebeca Aladro, Laura Colzi, Kimberly L. Emig, Christian Henkel, Yuri Nishimura, Víctor M. Rivilla, Paul P. van der Werf, and Alma Comprehensive High-Resolution Extragalactic Molecular Inventory (ALCHEMI) Collaboration. Energizing Star Formation: The Cosmic-Ray Ionization Rate in NGC 253 Derived from ALCHEMI Measurements of H₃O⁺ and SO. *The Astrophysical Journal*, 931:89, June 2022. ISSN 0004-637X. doi: 10.3847/1538-4357/ac6753. URL <https://ui.adsabs.harvard.edu/abs/2022ApJ...931...89H>. Publisher: IOP ADS Bibcode: 2022ApJ...931...89H.
- Izakun Jiménez-Serra, Andrés Megías, Joseph Salaris, Herma Cuppen, Angèle Taillard, Miwha Jin, Valentine Wakelam, Anton I. Vasyunin, Paola Caselli, and Yvonne J. Pendleton. Modelling methanol and hydride formation in the JWST Ice Age era. *Astronomy & Astrophysics*, 695:A247, 2025. URL <https://www.aanda.org/articles/aa/abs/2025/03/aa52389-24/aa52389-24.html>. Publisher: EDP Sciences.
- Patrick Kidger. On Neural Differential Equations, February 2022.

- A. Kværnø. Singly Diagonally Implicit Runge–Kutta Methods with an Explicit First Stage. *BIT Numerical Mathematics*, 44(3):489–502, August 2004. ISSN 1572-9125. doi: 10.1023/B:BITN.0000046811.70614.38. URL <https://doi.org/10.1023/B:BITN.0000046811.70614.38>.
- S. Maes, M. Van de Sande, T. Danilovich, F. De Ceuster, and L. Decin. Sensitivity study of chemistry in AGB outflows using chemical kinetics. *Monthly Notices of the Royal Astronomical Society*, 522:4654–4673, July 2023. ISSN 0035-8711. doi: 10.1093/mnras/stad1152. URL <https://ui.adsabs.harvard.edu/abs/2023MNRAS.522.4654M>. Publisher: OUP ADS Bibcode: 2023MNRAS.522.4654M.
- S. Maes, F. De Ceuster, M. Van de Sande, and L. Decin. MACE: A Machine learning Approach to Chemistry Emulation, May 2024.
- T. J. Millar, C. Walsh, M. Van De Sande, and A. J. Markwick. The UMIST Database for Astrochemistry 2022. *Astronomy & Astrophysics*, 682:A109, February 2024. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202346908.
- E. M. Penteado, C. Walsh, and H. M. Cuppen. Sensitivity Analysis of Grain Surface Chemistry to Binding Energies of Ice Species. *The Astrophysical Journal*, 844:71, July 2017. ISSN 0004-637X. doi: 10.3847/1538-4357/aa78f9. URL <https://ui.adsabs.harvard.edu/abs/2017ApJ...844...71P>. Publisher: IOP ADS Bibcode: 2017ApJ...844...71P.
- Maxime Ruaud, Valentine Wakelam, and Franck Hersant. Gas and grain chemical composition in cold cores as predicted by the Nautilus three-phase model. *Monthly Notices of the Royal Astronomical Society*, 459(4): 3756–3767, July 2016. ISSN 0035-8711. doi: 10.1093/mnras/stw887.
- Britton D. Smith, Greg L. Bryan, Simon C. O. Glover, Nathan J. Goldbaum, Matthew J. Turk, John Regan, John H. Wise, Hsi-Yu Schive, Tom Abel, Andrew Emerick, Brian W. O’Shea, Peter Anninos, Cameron B. Hummels, and Sadeh Khochfar. GRACKLE: A chemistry and cooling library for astrophysics. *Monthly Notices of the Royal Astronomical Society*, 466:2217–2234, April 2017. ISSN 0035-8711. doi: 10.1093/mnras/stw3291.
- Leonard Storcks and Tobias Buck. Differentiable Conservative Radially Symmetric Fluid Simulations and Stellar Winds – jfluids, October 2024. URL <http://arxiv.org/abs/2410.23093>. arXiv:2410.23093 [physics] version: 1.
- Immanuel Sulzer and Tobias Buck. Speeding up astrochemical reaction networks with autoencoders and neural ODEs, December 2023.
- Kwok Sun Tang and Matthew Turk. Reduced Order Model for Chemical Kinetics: A case study with Primordial Chemical Network, July 2022.
- A. I. Vasyunin, D. Semenov, Th. Henning, V. Wakelam, Eric Herbst, and A. M. Sobolev. Chemistry in Protoplanetary Disks: A Sensitivity Analysis. *The Astrophysical Journal*, 672:629–641, January 2008. ISSN 0004-637X. doi: 10.1086/523887. URL <https://ui.adsabs.harvard.edu/abs/2008ApJ...672..629V>. Publisher: IOP ADS Bibcode: 2008ApJ...672..629V.
- Gijs Vermariën, Thomas G Bisbas, Serena Viti, Yue Zhao, Xuefei Tang, and Rahul Ravichandran. NeuralPDR: Neural differential equations as surrogate models for photodissociation regions. *Machine Learning: Science and Technology*, 6(2):025069, June 2025. ISSN 2632-2153. doi: 10.1088/2632-2153/ade4ee.
- Gijs Vermariën, Serena Viti, Johannes Heyl, and Francesco Fontani. Understanding molecular ratios in the carbon- and oxygen-poor outer Milky Way with interpretable machine learning. *Astronomy & Astrophysics*, 699:A18, July 2025. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/202553893. URL <https://www.aanda.org/10.1051/0004-6361/202553893>. Publisher: EDP Sciences.
- V. Wakelam, F. Selsis, E. Herbst, and P. Caselli. Estimation and reduction of the uncertainties in chemical models: application to hot core chemistry. *Astronomy and Astrophysics*, 444:883–891, December 2005. ISSN 0004-6361. doi: 10.1051/0004-6361:20053673. URL <https://ui.adsabs.harvard.edu/abs/2005A&A...444..883W>. Publisher: EDP ADS Bibcode: 2005A&A...444..883W.

A Sensitivity analysis of reaction rates

In figure 3, we illustrate a sensitivity analysis of a stellar outflow. We can see that molecules such as CN have a relatively tight peak in abundances, with low sensitivity to the reaction rates, whereas larger carbon molecules such as C₃HCCH have a large sensitivity.

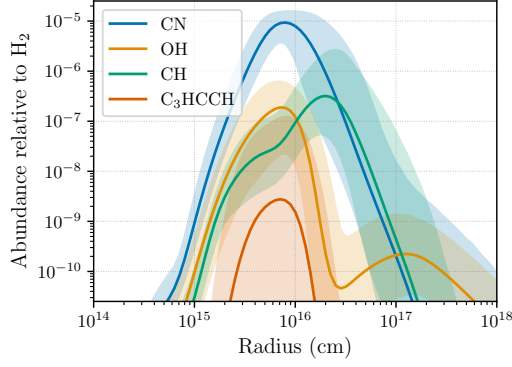


Figure 3: A figure illustrating the sensitivity of a carbon rich stellar outflow with a MCMC sampler varying reaction rates of the entire network.

B Reaction networks

B.1 Atomic Network

The atomic networks describes the most basic of reaction networks containing only 31 species and 380 reactions: C, C⁺, CH, CH⁺, CH₂, CH₂⁺, CH₃, CH₃⁺, CH₄, CH₄⁺, CH₅⁺, CO, CO⁺, e⁻, H, H⁺, H₂, H₂⁺, H₂O, H₂O⁺, H₃⁺, H₃O⁺, HCO, HCO⁺, He, He⁺, Mg, Mg⁺, O, O⁺, O₂, O₂⁺, OH, OH⁺. The reactions contain cosmic ray ionization reactions, photon reactions with rate:

$$k = \alpha F_{UV} \exp(-kA_V)$$

where F_{UV} is the strength of the radiation field often in Habing units, and A_V the visual extinction, which is a measure of the attenuation of the radiation due to the presence of interstellar dust grains, and it depends on the reactant and on the dust optical properties. They contain reactions via secondary photons induced by cosmic rays:

$$\alpha(T_{300})^\beta \frac{E}{1 - \omega} \zeta$$

where E is the cosmic ray efficiency and ω is the dust grain albedo. This is then combined with all the reactions from UMIST that contain all the species in the network. Lastly, we use a special H₂ formation and dissociation, CO dissociation and C⁺ ionization routine that is implemented like-for-like from UCLCHEM.

B.2 Molecular network

The molecular network expands on the atomic network by including the same reaction mechanisms, but now with more species. These species would normally be allowed to freeze out onto dust particles at lower temperatures, introducing ice chemistry that allows larger complex organic molecules such as methanol to form more effectively Jiménez-Serra et al. [2025]. This is turned off however, since the ice-chemistry is not yet implemented in Carbox. This results in a total of 161 species with 2227 reactions.