

---

# Differentiable Conservative Radially Symmetric Fluid Simulations and Stellar Winds $\circ$ `jf1uids`

---

Leonard Storcks and Tobias Buck

Interdisciplinary Center for Scientific Computing, University of Heidelberg  
Im Neuenheimer Feld 205, D-69120 Heidelberg, Germany

leonard.storcks@stud.uni-heidelberg.de & tobias.buck@iwr.uni-heidelberg.de

## Abstract

We present `jf1uids`, a one-dimensional fluid solver that can, by virtue of a *geometric formulation* of the Euler equations, model radially symmetric fluid problems in a conservative manner, i.e., without losing mass or energy. For spherical problems, such as ideal supernova explosions or stellar wind-blown bubble expansions, simulating only along a radial dimension drastically reduces compute and memory demands compared to a full three-dimensional method. This simplification also alleviates constraints on backpropagation through the solver. Written in JAX, `jf1uids` is a GPU-compatible and fully differentiable simulator. We demonstrate the advantages of this differentiable physics simulator by retrieving the wind's parameters for an adiabatic stellar wind expansion from the final fluid state using gradient descent. As part of a larger "stellar winds, cosmic rays and machine learning" research track, `jf1uids` serves as a solid foundation to be extended with additional physics modules, foremost cosmic rays and a neural-net powered gas-cooling surrogate and improved by higher order and more accurate numerical schemes. All code is available under <https://github.com/leo1200/jf1uids/>.

## 1 Motivation and related work

**Our motivation is blowin' in the wind** We are interested in detailed simulations of radial gaseous outflows from a star (stellar winds) [22, 38] and ultimately, inspired by recent observations [24], their interaction with cosmic rays (energetic particles), a challenging pursuit in multiple ways. High-resolution three-dimensional simulations are computationally expensive, and physical effects like (turbulent) cooling of the gas in practice often scale with the resolution of the simulation (dependent on how well the bubble surface is resolved) [17, 18]. We therefore aim for a performant simulation code with radial symmetry,<sup>1</sup> in which detailed physics is modeled by machine learning surrogates in the solver. Finally, once detailed physics is incorporated, it would be desirable to use the simulator to reconstruct stellar wind parameters from observed data (inverse modeling).

**Towards high-performance** Low level classical simulation codes (like [28, 29]) with optimized hand-crafted parallelization schemes are in the process of being extended to High-Performance Computing (HPC) on GPUs (and TPUs) (see [10, 41]) based on vendor specific GPU-programming language extensions like CUDA [21] or a framework like `kokkos` [34]. In another avenue for HPC, which we follow, most of the underlying complexity is abstracted into higher-level numerical libraries

---

<sup>1</sup>This might also be of interest in other areas like studying sonoluminescent bubbles [1].

like `tinygrad` [30], `PyTorch` [23] or `JAX` [4] in the latter of which various simulation codes have already been written (e.g. [2, 5, 15, 19, 33, 39]).<sup>2</sup>

**Differentiability** Another advantage is, that code written in `JAX` is differentiable, the result of a whole simulation can be differentiated with respect to its parameters. This allows us to directly optimize machine learning surrogates for the detailed physics in the simulator, rather than learning an error correction scheme (like `NeuralSim` [12] does in the field of robotics). The same goes for optimizing parts of the numerics like gradient limiters [20], or smoothness indicators [16]. Inverse modeling also benefits from differentiability - in the simplest cases, direct gradient descent on the parameters may be feasible.

**Why not JAX-Fluids?** `JAX-Fluids`<sup>3</sup> is a fluid simulator written in `JAX` featuring high-order<sup>4</sup> numerical schemes. While it supports one-dimensional spherically symmetric simulations, it does so via a geometric source formulation [32, section 1.6.4] which is not conservative. Total energy and mass are not constant over time, an unwanted effect [8].

## 2 The numerics of `jfluids`

### 2.1 Geometric source term vs. geometric formulation of the Euler equations

The geometric difference of a radially symmetric two- or three-dimensional to a Cartesian one-dimensional inviscid fluid can be encompassed by a geometric source term  $\vec{S}(\vec{U})$  in the Euler equations

$$\partial_t \vec{U} + \partial_r \vec{F} = \vec{S} \quad (1)$$

where the state vector  $\vec{U}$ , the flux vector  $\vec{F}$  and geometric source term  $\vec{S}$  are given by

$$\vec{U} = \begin{bmatrix} \rho \\ \rho u \\ E \end{bmatrix}, \quad \vec{F} = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ u(E + p) \end{bmatrix}, \quad \vec{S} = -\frac{\alpha}{r} \begin{bmatrix} \rho u \\ \rho u^2 \\ u(E + p) \end{bmatrix}, \quad \alpha = \begin{cases} 0 & \text{Cartesian} \\ 1 & \text{Cylindrical} \\ 2 & \text{Spherical} \end{cases} \quad (2)$$

with density  $\rho$ , velocity  $u$ , energy  $E$  and pressure  $p$  given via an equation of state. In this formulation, expanding a standard Cartesian one-dimensional solver to radially symmetric two- and three-dimensional simulations only requires implementing the source term  $\vec{S}$ . This was done in `JAX-Fluids` source term solver. Only adapting the solver in terms of a source term, however, leads to numerics inconsistent with the geometry - the proper cell centers to use are the volumetric not the geometric centers and for fluxes to be conservative, the radial increase of area the flux passes through is to be considered (more formal reasoning can be found in [8]).

Based on the geometric form of the Euler equations

$$\partial_t \vec{U} + \frac{1}{r^\alpha} \partial_r (r^\alpha \vec{F}) = \frac{1}{r^\alpha} \vec{N}, \quad \text{with pressure nozzling } \vec{N} = \begin{bmatrix} 0 \\ \alpha r^{\alpha-1} p \\ 0 \end{bmatrix}, \quad (3)$$

Crittenden and Balachandar [8] present the baseline discretization for a radial fluid solver which we implement in Hancock's second order MUSCL-scheme (as described in [35]; we opt for the *minmod gradient limiter*) with a HLL Riemann solver [31].

### 2.2 On conservational properties

To give an example of simulation results obtained using `jfluids` we depict the final density over the radial dimension in a shock test setting with step-functions as initial conditions for the density and

<sup>2</sup>Note, however, that adaptive multiresolution (AMR) in a differentiable and just-in-time compiled domain-specific language like `JAX` has not yet been achieved [3] but is generally possible in a HPC setting (see [11]). We provide a simple prototype for AMR in one dimension in `JAX` under <https://github.com/leo1200/jamr>.

<sup>3</sup><https://github.com/tumaer/JAX-Fluids>

<sup>4</sup>The spatial order of a fluid simulator relates to the speed of numerical convergence with increasing grid resolution (in smooth regimes).

pressure at  $t = 0.2$  on the left side of Fig. 1. For different resolutions and different spatial orders (first order Godunov and second order MUSCL-Hancock scheme (see [32] for details)), as expected, higher resolutions and higher-order spatial schemes lead to sharper solutions.

`jfluids` preserves mass and energy throughout a fluid simulation (for a proof, see [8]). As noted by Crittenden and Balachandar [8], this is a consequence of the discretization of the geometric form of the Euler equations used, not the specific flux-calculation method, the spatial order or the grid size etc. We demonstrate mass and energy conservation for the shock problem in Fig. 1.

While `jfluids` energy and mass residuals fluctuate between  $\pm 4 \cdot 10^{-14}$ ,<sup>5</sup> JAX-Fluids shows increasing residuals with a relative energy error up to 0.1%. Note that the same conservative baseline discretization could - with some implementation effort - also be used with the higher-order methods of JAX-Fluids (see also [37]).

### 2.3 On performance

`jfluids` is written in pure JAX, can be just-in-time compiled and therefore runs on CPUs, GPUs and TPUs alike. Detailed performance-accuracy comparisons between `jfluids` and JAX-Fluids will be provided in the future but are not straightforward. Higher order methods provide better accuracy at lower resolution but increased cost, such that, depending on the setting and exact methods, the accuracy per compute of lower order methods can be superior (as in [9]).

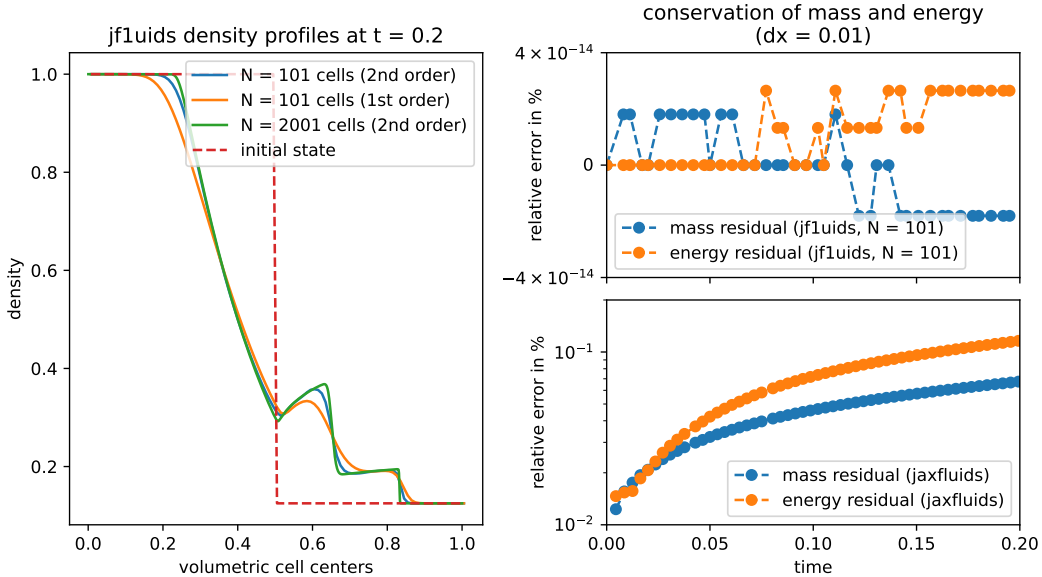


Figure 1: On the left side, the illustrative problem, a radial shock setting is presented, and different solutions compared, on the right the change in mass and energy over time for `jfluids` and JAX-Fluids is shown.

## 3 Differentiation through stellar wind simulations

As a first physics module, we have implemented stellar wind based on the injection of thermal energy.<sup>6</sup> Our simulation successfully replicates the theoretical Weaver solution [38] (see Fig. 2). The main parameters of the wind are the stellar mass  $M$  and outflow velocity of the wind  $v_\infty$ . Fig. 2 shows the simulation of an exemplary wind and the derivative of the final fluid state with respect to

<sup>5</sup>This is due to the limited precision in which the calculations are carried out. Especially with single precision floating point calculations, slight drifts in the *conserved* quantities are possible (just as in symplectic ordinary differential equation solvers, see [27]).

<sup>6</sup>Thermal energy injection is comparable to the situation in a small cluster of stars with thermalization of kinetic energy [17], for a comparison of injection schemes, see [26].

$v_\infty$  - `jf1uids` is fully differentiable. The derivatives are remarkably stable and have simple physical interpretations. For example, nudging up  $v_\infty$  will forward-displace the shock, leading to the spike in the derivatives after 2 parsec. Note that for backpropagation through the solver, all the internal operations (so the whole trajectory over time, or part of it based on checkpointing [3]) are stored [15] which can become a memory bottleneck. This, however, is less pronounced in our one-dimensional compared to a full three-dimensional simulation.<sup>7</sup>




Figure 2: Density, pressure and velocity of a stellar wind simulation in `jf1uids` and their derivatives with respect to the outflow velocity of the stellar wind  $v_\infty$  at  $v_\infty = 2000 \frac{\text{km}}{\text{s}}$ . The resolution is  $N = 401$  cells. The automatic differentiation (autodiff) was done with `jacfwd`, the finite difference plotted is a centered difference quotient with  $\Delta v_\infty = 0.1 \frac{\text{km}}{\text{s}}$

In Fig. 3 we further show, that based on the full final fluid state, simulated using `jf1uids`, and the general simulation setup (background density, pressure, adiabatic index, ...) by gradient descent through the simulation, we can reconstruct stellar wind parameters.

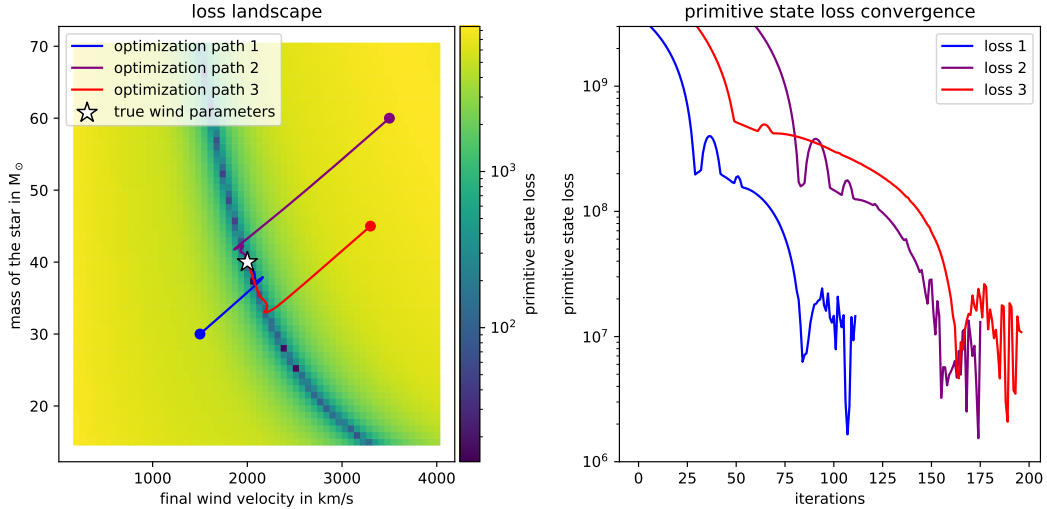


Figure 3: Gradient descent for retrieving stellar wind parameters from the final fluid state of a stellar wind simulation. On the left side, the optimization paths through the two-dimensional parameter space of the wind are shown, on the right, a loss measuring the summed absolute difference to the final state we seek to attain is shown. The resolution is  $N = 401$  cells.

## 4 Conclusion and outlook

We have presented a fully differentiable, conservative and GPU-ready radial fluid solver in JAX, with superior conservative properties compared to JAX-Fluids [3]. Differentiability was showcased on

<sup>7</sup>At high resolutions, which also necessitate smaller time steps [6], this can still be an issue.

stellar wind simulations. Further research as part of the "stellar winds, cosmic rays and machine learning" research track aims at (also see the research overview illustration Fig. 4)

- *numerics*: implementation of higher-order methods and more advanced Riemann solvers (lack thereof is a current limitation), detailed performance profiling
- *machine learning enhanced numerics*: end-to-end training of gradient limiters and smoothness indicators to improve upon the baseline numerics
- *machine learning powered physics*: surrogate modeling for cooling; based on analytical models with free parameters and neural-network cooling terms; aiming for wind evolutions as theorized in [18]
- *cosmic rays*: from summary statistics based on shock energy dissipation [25] over two-fluids models [36] to full cosmic ray energy distributions co-modeled with the fluid flow [14]; the dependence of cosmic ray acceleration on the evolution of the wind will be investigated
- *parameter retrieval*: insights into stellar wind parameters from observable summary statistics via gradient descent, Hamilton-Monte-Carlo [13] and simulation-based inference [7], taking advantage of the differentiability of the simulator [40]

### stellar winds, cosmic rays and machine learning

**sub-track 1:** stellar wind surrogates for the large-scale simulation of Milky Way like galaxies

**problem:** large scale simulations are too coarse to accurately resolve stellar wind

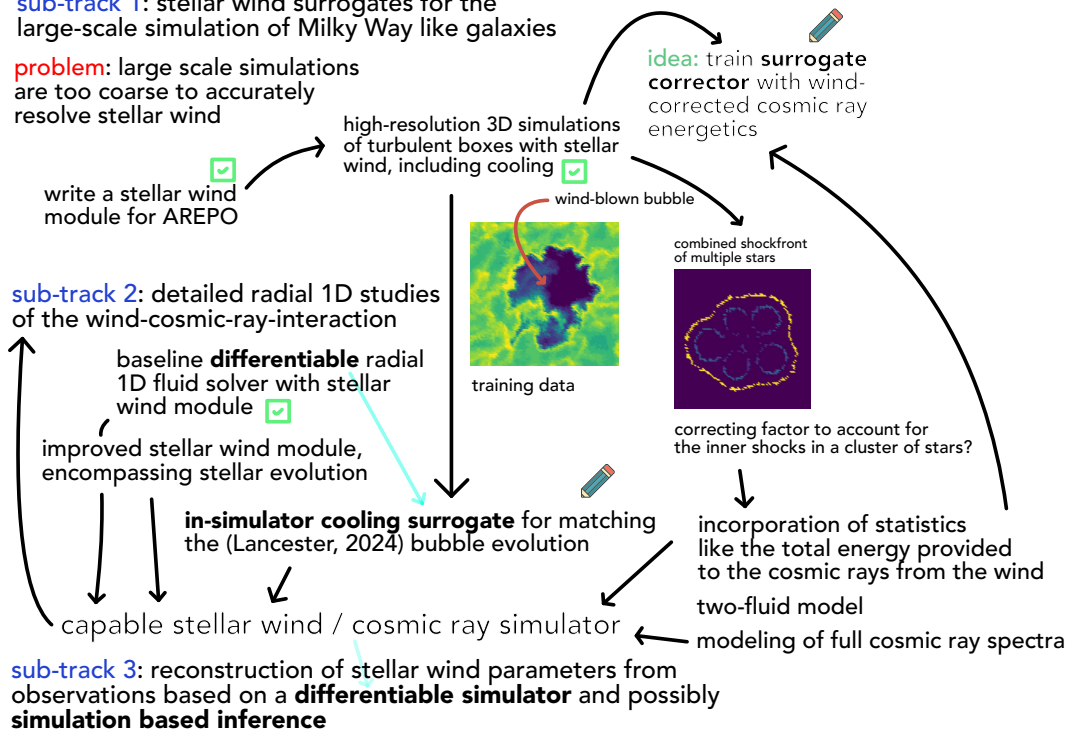


Figure 4: Overview on the broader research context of this work.

## Broader impact statement

The authors are not aware of any immediate ethical or societal implications of this work. This work purely aims to aid scientific research and proposes a method for radial fluid simulations.

## References

- [1] Bradley P. Barber, Robert A. Hiller, Ritva Löfstedt, Seth J. Putterman, and Keith R. Weninger. Defining the unknowns of sonoluminescence. *Physics Reports*, 281(2):65–143, 1997. ISSN 0370-1573. doi: [https://doi.org/10.1016/S0370-1573\(96\)00050-6](https://doi.org/10.1016/S0370-1573(96)00050-6). URL <https://www.sciencedirect.com/science/article/pii/S0370157396000506>.
- [2] Deniz A. Bezgin, Aaron B. Buhendwa, and Nikolaus A. Adams. Jax-fluids: A fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. *Computer Physics Communications*, 282:108527, 2023. ISSN 0010-4655. doi: <https://doi.org/10.1016/j.cpc.2022.108527>. URL <https://www.sciencedirect.com/science/article/pii/S0010465522002466>.
- [3] Deniz A. Bezgin, Aaron B. Buhendwa, and Nikolaus A. Adams. Jax-fluids 2.0: Towards hpc for differentiable cfd of compressible two-phase flows, 2024. URL <https://arxiv.org/abs/2402.05193>.
- [4] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- [5] Jean-Eric Campagne, François Lanusse, Joe Zuntz, Alexandre Boucaud, Santiago Casas, Minas Karamanis, David Kirkby, Denise Lanzieri, Austin Peel, and Yin Li. Jax-cosmo: An end-to-end differentiable and gpu accelerated cosmology library. *The Open Journal of Astrophysics*, 6, April 2023. ISSN 2565-6120. doi: [10.21105/astro.2302.05163](https://doi.org/10.21105/astro.2302.05163). URL <http://dx.doi.org/10.21105/astro.2302.05163>.
- [6] R. Courant, K. Friedrichs, and H. Lewy. Über die partiellen differenzgleichungen der mathematischen physik. *Mathematische Annalen*, 100(1):32–74, December 1928. ISSN 1432-1807. doi: [10.1007/bf01448839](https://doi.org/10.1007/bf01448839). URL <http://dx.doi.org/10.1007/BF01448839>.
- [7] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020. doi: [10.1073/pnas.1912789117](https://doi.org/10.1073/pnas.1912789117). URL <https://www.pnas.org/doi/abs/10.1073/pnas.1912789117>.
- [8] P. E. Crittenden and S. Balachandar. The impact of the form of the euler equations for radial flow in cylindrical and spherical coordinates on numerical conservation and accuracy. *Shock Waves*, 28(4):653–682, March 2018. ISSN 1432-2153. doi: [10.1007/s00193-017-0784-y](https://doi.org/10.1007/s00193-017-0784-y). URL <http://dx.doi.org/10.1007/s00193-017-0784-y>.
- [9] J.A Greenough and W.J Rider. A quantitative comparison of numerical methods for the compressible euler equations: fifth-order weno and piecewise-linear godunov. *Journal of Computational Physics*, 196(1):259–281, 2004. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2003.11.002>. URL <https://www.sciencedirect.com/science/article/pii/S0021999103005965>.
- [10] Philipp Grete, Forrest W. Glines, and Brian W. O’Shea. K-athena: a performance portable structured grid finite volume magnetohydrodynamics code, 2020. URL <https://arxiv.org/abs/1905.04341>.
- [11] Philipp Grete, Joshua C. Dolence, Jonah M. Miller, Joshua Brown, Ben Ryan, Andrew Gaspar, Forrest Glines, Sriram Swaminarayan, Jonas Lippuner, Clell J. Solomon, Galen Shipman, Christoph Junghans, Daniel Holladay, James M. Stone, and Luke F. Roberts. Parthenon – a performance portable block-structured adaptive mesh refinement framework. *arXiv e-prints*, art. arXiv:2202.12309, February 2022. doi: [10.48550/arXiv.2202.12309](https://doi.org/10.48550/arXiv.2202.12309).

- [12] Eric Heiden, David Millard, Erwin Coumans, Yizhou Sheng, and Gaurav S. Sukhatme. Neursim: Augmenting differentiable simulators with neural networks, 2021. URL <https://arxiv.org/abs/2011.04217>.
- [13] Matthew D Hoffman, Andrew Gelman, et al. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, 2014.
- [14] Hyesung Kang and T. W. Jones. Numerical studies of diffusive particle acceleration in supernova remnants. *Monthly Notices of the Royal Astronomical Society*, 249(3):439–451, 04 1991. ISSN 0035-8711. doi: 10.1093/mnras/249.3.439. URL <https://doi.org/10.1093/mnras/249.3.439>.
- [15] Patrick Kidger. *On Neural Differential Equations*. PhD thesis, University of Oxford, 2021.
- [16] Tatiana Kossaczka, Matthias Ehrhardt, and Michael Günther. Enhanced fifth order weno shock-capturing schemes with deep learning. *Results in Applied Mathematics*, 12:100201, 2021. ISSN 2590-0374. doi: <https://doi.org/10.1016/j.rinam.2021.100201>. URL <https://www.sciencedirect.com/science/article/pii/S2590037421000431>.
- [17] Lachlan Lancaster, Eve C. Ostriker, Jeong-Gyu Kim, and Chang-Goo Kim. Efficiently cooled stellar wind bubbles in turbulent clouds. ii. validation of theory with hydrodynamic simulations. *The Astrophysical Journal*, 914(2):90, jun 2021. doi: 10.3847/1538-4357/abf8ac. URL <https://dx.doi.org/10.3847/1538-4357/abf8ac>.
- [18] Lachlan Lancaster, Eve C. Ostriker, Chang-Goo Kim, Jeong-Gyu Kim, and Greg L. Bryan. Geometry, dissipation, cooling, and the dynamical evolution of wind-blown bubbles. *The Astrophysical Journal*, 970(1):18, jul 2024. doi: 10.3847/1538-4357/ad47f6. URL <https://dx.doi.org/10.3847/1538-4357/ad47f6>.
- [19] Yin Li, Libin Lu, Chirag Modi, Drew Jamieson, Yucheng Zhang, Yu Feng, Wenda Zhou, Ngai Pok Kwan, François Lanusse, and Leslie Greengard. pmwd: A differentiable cosmological particle-mesh  $n$ -body library, 2022. URL <https://arxiv.org/abs/2211.09958>.
- [20] Nga Nguyen-Fotiadis, Michael McKerns, and Andrew Sornborger. Machine learning changes the rules for flux limiters. *Physics of Fluids*, 34(8), August 2022. ISSN 1089-7666. doi: 10.1063/5.0102939. URL <http://dx.doi.org/10.1063/5.0102939>.
- [21] NVIDIA, Péter Vingelmann, and Frank H.P. Fitzek. Cuda, release: 10.2.89, 2020. URL <https://developer.nvidia.com/cuda-toolkit>.
- [22] E. N. Parker. Dynamics of the Interplanetary Gas and Magnetic Fields. *Astrophysical Journal*, 128:664, November 1958. doi: 10.1086/146579.
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [24] Giada Peron, Sabrina Casanova, Stefano Gabici, Vardan Baghmanyanyan, and Felix Aharonian. The contribution of winds from star clusters to the galactic cosmic-ray population. *Nature Astronomy*, 8(4):530–537, January 2024. ISSN 2397-3366. doi: 10.1038/s41550-023-02168-6. URL <http://dx.doi.org/10.1038/s41550-023-02168-6>.
- [25] C. Pfrommer, R. Pakmor, K. Schaal, C. M. Simpson, and V. Springel. Simulating cosmic ray physics on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 465(4):4500–4529, November 2016. ISSN 1365-2966. doi: 10.1093/mnras/stw2941. URL <http://dx.doi.org/10.1093/mnras/stw2941>.

- [26] J M Pittard, C J Wareing, and M M Kupilas. How to inflate a wind-blown bubble. *Monthly Notices of the Royal Astronomical Society*, 508(2):1768–1776, September 2021. ISSN 1365-2966. doi: 10.1093/mnras/stab2712. URL <http://dx.doi.org/10.1093/mnras/stab2712>.
- [27] Hanno Rein and David S. Spiegel. ias15: a fast, adaptive, high-order integrator for gravitational dynamics, accurate to machine precision over a billion orbits. *Monthly Notices of the Royal Astronomical Society*, 446(2):1424–1437, November 2014. ISSN 0035-8711. doi: 10.1093/mnras/stu2164. URL <http://dx.doi.org/10.1093/mnras/stu2164>.
- [28] Volker Springel. E pur si muove: galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 401(2):791–851, January 2010. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2009.15715.x. URL <http://dx.doi.org/10.1111/j.1365-2966.2009.15715.x>.
- [29] James M. Stone, Thomas A. Gardiner, Peter Teuben, John F. Hawley, and Jacob B. Simon. Athena: A new code for astrophysical mhd. *The Astrophysical Journal Supplement Series*, 178(1):137–177, September 2008. ISSN 1538-4365. doi: 10.1086/588755. URL <http://dx.doi.org/10.1086/588755>.
- [30] tinygrad. tinygrad: A simple and powerful neural network framework. <https://tinygrad.org>, 2024. Accessed: 2024-09-08.
- [31] Eleuterio F. Toro. *The HLL and HLLC Riemann Solvers*, page 293–311. Springer Berlin Heidelberg, 1997. ISBN 9783662034903. doi: 10.1007/978-3-662-03490-3\_10. URL [http://dx.doi.org/10.1007/978-3-662-03490-3\\_10](http://dx.doi.org/10.1007/978-3-662-03490-3_10).
- [32] Eleuterio F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Springer Berlin Heidelberg, 2009. ISBN 9783540498346. doi: 10.1007/b79761. URL <http://dx.doi.org/10.1007/b79761>.
- [33] Artur P. Toshev, Harish Ramachandran, Jonas A. Erbesdobler, Gianluca Galletti, Johannes Brandstetter, and Nikolaus A. Adams. Jax-sph: A differentiable smoothed particle hydrodynamics framework, 2024. URL <https://arxiv.org/abs/2403.04750>.
- [34] Christian R. Trott, Damien Lebrun-Grandié, Daniel Arndt, Jan Ciesko, Vinh Dang, Nathan Ellingwood, Rahul Kumar Gayatri, Evan Harvey, Daisy S. Hollman, Dan Ibanez, Nevin Liber, Jonathan Madsen, Jeff Miles, David Poliakoff, Amy Powell, Sivasankaran Rajamanickam, Mikael Simberg, Dan Sunderland, Bruno Turcksin, and Jeremiah Wilke. Kokkos 3: Programming model extensions for the exascale era. *IEEE Transactions on Parallel and Distributed Systems*, 33(4):805–817, 2022. doi: 10.1109/TPDS.2021.3097283.
- [35] Bram van Leer. Upwind and high-resolution methods for compressible flow: From donor cell to residual-distribution schemes. In *16th AIAA Computational Fluid Dynamics Conference*. American Institute of Aeronautics and Astronautics, June 2003. doi: 10.2514/6.2003-3559. URL <http://dx.doi.org/10.2514/6.2003-3559>.
- [36] Wagner, A. Y., Falle, S. A. E. G., Hartquist, T. W., and Pittard, J. M. Two-fluid models of cosmic ray modified radiative shocks. *Astronomy and Astrophysics*, 452(3):763–771, 2006. doi: 10.1051/0004-6361:20064885. URL <https://doi.org/10.1051/0004-6361:20064885>.
- [37] Sheng Wang and Eric Johnsen. High-order schemes for the euler equations in cylindrical/spherical coordinates, 2017. URL <https://arxiv.org/abs/1701.04834>.
- [38] R. Weaver, R. McCray, J. Castor, P. Shapiro, and R. Moore. Interstellar bubbles. II. Structure and evolution. *Astrophysical Journal*, 218:377–395, December 1977. doi: 10.1086/155692.
- [39] Tianju Xue, Shuheng Liao, Zhengtao Gan, Chanwook Park, Xiaoyu Xie, Wing Kam Liu, and Jian Cao. Jax-fem: A differentiable gpu-accelerated 3d finite element solver for automatic inverse design and mechanistic data science. *Computer Physics Communications*, 291:108802, October 2023. ISSN 0010-4655. doi: 10.1016/j.cpc.2023.108802. URL <http://dx.doi.org/10.1016/j.cpc.2023.108802>.



- [40] Justine Zeghal, François Lanusse, Alexandre Boucaud, Benjamin Remy, and Eric Aubourg. Neural posterior estimation with differentiable simulators, 2022. URL <https://arxiv.org/abs/2207.05636>.
- [41] Oliver Zier, Rahul Kannan, Aaron Smith, Mark Vogelsberger, and Erkin Verbeek. Adapting arepo-rt for exascale computing: Gpu acceleration and efficient communication, 2024. URL <https://arxiv.org/abs/2404.17630>.