

Neutral hydrogen gas in supernova-driven galactic outflows

Chong-Chong He

chongchong.he@anu.edu.au

Australian National University <https://orcid.org/0000-0002-2332-8178>

Mark Krumholz

Australian National University <https://orcid.org/0000-0003-3893-854X>

Benjamin Wibking

Michigan State University

Aditi Vijayan

Australian National University

Article

Keywords: galaxies: evolution, galaxies: ISM, galaxies: starburst, hydrodynamics, supernovae: general

Posted Date: September 29th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11070170/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

Neutral hydrogen gas in supernova-driven galactic outflows

Chong-Chong He^{1*}, Mark R. Krumholz¹, Benjamin D. Wibking²,
Aditi Vijayan³

^{1*}Research School of Astronomy and Astrophysics, Australian National University,
Canberra, 2601, ACT, Australia.

²Department of Physics and Astronomy, Michigan State University, East Lansing, 48824,
MI, USA.

³School of Earth and Space Exploration, Arizona State University, Tempe, 85281, AZ, USA.

*Corresponding author(s). E-mail(s): chongchong.he@anu.edu.au;

Abstract

Winds that eject mass from galactic discs regulate the growth of galaxies, and observations show that, at least within a few kpc of the disc, most of the mass flux is carried by neutral rather than ionized material. However, it has proven difficult for simulations of supernova-driven winds to reproduce this observation, leading to speculations that the bulk of the mass must be driven by radiation pressure or cosmic rays rather than supernovae. Here we report simulations of supernova-driven winds using the GPU-accelerated code QUOKKA that reach significantly higher resolution and simulation volume than previous work, and that produce winds dominated by neutral mass, in good agreement with observations. The neutral material is not ballistic, but is instead driven by momentum transfer associated with drag and condensation of hot gas onto the neutral phase. By contrast, when we reduce our resolution to a value comparable to that used in earlier work, we also find a lack of neutral material. We show that the change in qualitative outcome with resolution is the result of lower-resolution simulations failing to capture the mixing layers within which condensation occurs, and suggest a new resolution requirement for simulations to capture the phase structure of galactic winds.

Keywords: galaxies: evolution, galaxies: ISM, galaxies: starburst, hydrodynamics, supernovae: general

Galactic winds that remove mass and metals from star-forming discs play a decisive role in galaxy evolution, and in gas-rich regions much of the mass flux appears to be carried by the neutral and molecular phases (Leroy et al. 2015; Martini et al. 2018; Krieger et al. 2019; Veilleux et al. 2020; Yuan et al. 2023). The neutral material in these outflows extends to distances of at least several kpc above the disc. Observations of the Milky Way’s own nuclear wind, where they can achieve the highest resolution, reveal that the neutral material consists of a mist of clouds with typical radii of $\sim 10 - 30$ pc and masses from 10 to $10^5 M_\odot$ (Di Teodoro et al. 2018; Lockman et al. 2020; Noon et al. 2023). Such clouds have been detected up to 2.5–3.5 kpc off the plane, with velocities of $\approx 150 - 300 \text{ km s}^{-1}$ that suggest lifetimes of $\gtrsim 5 - 10$ Myr, and in chemical states that suggest the material must have been entrained from the Galactic plane rather than condensing *in situ* (Noon et al. 2026).

The survival of neutral clouds to such distances and times has proven difficult to explain theoretically: an adiabatic, non-magnetised cold cloud embedded in a hot wind will be disrupted on a cloud-crushing timescale that is far shorter than the inferred cloud lifetimes, and is never accelerated to velocities as high as those seen (Klein et al. 1994). Radiative cooling (Gronke and Oh 2018), thermal conduction (Brüggen and Scannapieco 2016) and magnetic fields (McCourt et al. 2015; Cottle et al. 2020) may mitigate these effects, but have thus far only been demonstrated in idealised “wind-tunnel” simulations that follow a single cloud, not for a full galactic wind. By contrast, stratified-box simulations, which follow a vertically extended patch of a star-forming disc and are the standard tool for studying how galactic winds are launched, have had great difficulty reproducing the observed abundance of neutral gas in winds in simulations where the winds are driven by supernovae (Walch et al. 2015; Girichidis et al. 2016; Kim et al. 2026, 2020; Peschken et al. 2026). In such simulations, while neutral material dominates very close to the disc, almost all the gas

is ionized by the time it reaches distances $\gtrsim 1$ kpc from the galactic plane. Their typical neutral gas mass loading factors – defined as the ratio of the wind mass flux to the star formation rate – at distances of 1 – 2 kpc are $\eta_{\text{H I}} \lesssim 10^{-4}$ (e.g., Kim et al. 2020; Peschken et al. 2026). Significant amounts of neutral material are found only in simulations where the outflow is assumed to be driven primarily by non-thermal pressure sources such as cosmic rays (e.g., Armillotta et al. 2024; Thomas et al. 2025; Sike et al. 2025).

However, typical resolutions in the stratified-box simulations carried out to date are coarse compared with the observed ~ 10 – 30 pc sizes of neutral clouds. For example the SILCC simulation suite reaches only ≈ 4 pc resolution (Girichidis et al. 2016), while the TIGRESS suite reaches 2 pc (Kim and Ostriker 2017), but only within $|z| = 2$ kpc from the disc (Peschken et al. 2026), and only for small disc patches of well under 1 kpc in size. Moreover, due to the numerical difficulty of simulating flows with such extreme density contrasts, many of these simulations resort to relatively diffusive numerical methods, for example piecewise-linear rather than piecewise-parabolic interpolation, further lowering the effective resolution. The highest-resolution simulations of this type carried out to date that include neutral gas are the QED simulation suite (Vijayan et al. 2024, 2025; Vijayan et al. 2026), and there is preliminary evidence from both the QED simulations and other studies (Kim and Ostriker 2017; Peschken et al. 2026) that the properties of cool material in outflows are sensitive to resolution. However, no study has yet demonstrated convergence in the neutral content of a supernova-driven outflow.

Here we report a series of simulations of a supernova-driven outflow from a GPU-accelerated code QUOKKA (Wibking and Krumholz 2022; He et al. 2024a,b; He et al. 2026) in which we reach resolutions of 1 pc at distances of ± 4 kpc from a galactic plane. Our scheme employs piecewise-parabolic reconstruction, which provides a more accurate representation of smooth flows than the piecewise-linear methods commonly used in earlier stratified-box studies, and thus yields sharper shocks and higher effective resolution for advection at fixed grid spacing (e.g. Leidi et al. 2024). These simulations therefore, for the first time, achieve spatial resolutions sufficient to resolve the observed sizes of neutral clouds at their detected distances. We carry out a resolution study in which we hold the physics fixed and vary only the grid spacing. We show that at our highest resolution our simulations produce a neutral mass-dominated wind whose properties are in good agreement with observations, while our lower resolution cases, similar to earlier work, fail to produce a neutral-dominated wind. We analyse the simulations to understand what drives neutral winds and why they have been missed in earlier works.

Results

All our simulations take place in a computational domain $\approx 1 \times 1 \times 8$ kpc in size with periodic boundary conditions in the x and y directions and diode boundaries on the two faces in the z direction. Our highest resolution simulations use a uniform resolution of $1024^2 \times 8192$ cells throughout the domain, giving a cell spacing $\Delta x \approx 1$ pc, and for comparison we also carry out simulations at resolutions with half and one quarter as many cells per linear dimension, with cell sizes $\Delta x \approx 2$ and 4 pc. All simulations begin with a gas surface density $\Sigma_{\text{gas}} = 50 \text{ M}_{\odot} \text{ pc}^{-2}$ and a stellar surface density $\Sigma_{*} = 208 \text{ M}_{\odot} \text{ pc}^{-2}$, comparable to moderately starbursting regions such as the Milky Way centre. Our simulations include hydrodynamics, cooling, self-gravity, and self-consistent star formation and feedback in the form of both FUV radiation and supernovae. See **Methods** for full details on our numerical method and initial conditions, and **Extended Data Table 1** for key summary statistics of the simulations.

Full resolution results

Figure 1 shows the flagship 1 pc simulation at $t = 60$ Myr, close to the time of maximum mass outflow rate. It is clear from the figure that neutral gas reaches the top of the domain 4 kpc above the disc, and that it is structured as clouds and filaments tens of parsecs in size, embedded in a volume-filling wind moving faster than 500 km s^{-1} . Neutral gas is the dominant component of the outflow at every height we measure, carrying 69% of the mass and $\approx 40\%$ of the mass flux above $|z| = 1$ kpc, compared to ≈ 20 – 25% by mass or mass flux for the warm ionised medium and 10% by mass and $\approx 35\%$ by flux for the hot phase; see **Methods** for how we define each phase, and **Extended Data Table 2** for exact numerical values. Crucially, neither the neutral mass loading nor the fraction of outflowing mass carried in the neutral phase decay with height: $\eta_{\text{H I}} \approx 0.3$ and $\eta_{\text{H I}}/\eta_{\text{total}} \approx 0.4$ independent of height from 1 – 4 kpc.

Segmenting the neutral gas into connected three-dimensional structures (see **Methods** and **Extended Data Table 3**) reveals that the neutral material is arranged into a population of discrete structures with a broad spectrum of sizes, effective radii $R_{\text{eff}} \sim 10$ – 70 pc, and masses, $M \sim 10^2$ – $10^5 \text{ M}_{\odot}$; typical column densities are $N_{\text{H}} \sim 10^{20}$ – 10^{21} cm^{-2} . These figures are in broad agreement with the properties observed for clouds in the Milky Way’s nuclear wind (Di Teodoro et al. 2018; Lockman et al. 2020; Noon et al. 2023). Cloud densities, and temperatures are typical of warm neutral ISM, $n_{\text{H}} \sim \text{few cm}^{-3}$ and $T \sim 5000$ – $10,000$ K. The mass distribution includes many small clouds, which tend to be somewhat less dense and hotter,

but much of the mass is in large clouds, and the dominance of larger clouds grows as the clouds move away from the plane. We quantify this trend by measuring the fraction of the mass in structures with effective radii > 30 pc, which grows from $\approx 50\%$ at $|z| \approx 1.5$ kpc to $\approx 70\%$ at $|z| \approx 3.5$ kpc; the largest clouds have filamentary morphologies, with the cloud shown in the inset of Figure 1 being a typical example.

Effects of resolution

Coarsening the grid destroys the neutral outflow rather than scaling it down. Figure 2 shows the same environment and evolutionary stage at $\Delta x = 4, 2$ and 1 pc. We see that at 4 pc neutral gas in the wind is absent or barely resolved, while at 2 pc it is present but is significantly less prominent, and has a qualitatively different spatial distribution that at our maximum resolution. The neutral mass loading factors confirm this visual impression: at the box edge we find η_{HI} and $\eta_{\text{HI}}/\eta_{\text{total}}$ are both $\lesssim 10^{-3}$ with $\Delta x = 4$ pc, $\eta_{\text{HI}} \approx 0.17$ and $\eta_{\text{HI}}/\eta_{\text{total}} \approx 0.3$ with $\Delta x = 2$ pc, compared to $\eta_{\text{HI}} \approx 0.3$ and $\eta_{\text{HI}}/\eta_{\text{total}} \approx 0.4$ and the full 1 pc resolution; see **Extended Data Table 2** for precise figures. Thus a factor of two reduction in resolution decreases the neutral mass loading by a factor of ≈ 2 while a factor of four reduction makes it disappear almost entirely. Examining how η_{HI} changes with height makes it clear that this effect is less a matter of launching than of survival: neutral gas *is* launched at $\Delta x = 4$ pc, with $\eta_{\text{HI}} = 0.008$ at 0.98 kpc, but decays by a factor of 20 with height and reaches zero above 3 kpc, whereas at $\Delta x = 2$ and 1 pc the same quantity is flat to within 11 and 7 per cent over the identical range. By contrast, the total mass loading declines by only 20 – 25 per cent over the same heights at every resolution. Coarsening the grid therefore does not suppress the outflow overall, it simply guarantees that the neutral phase no longer survives the climb.

We can also confirm that the difference in the behaviour of the neutral outflow is not due to a difference in feedback. We illustrate this in Fig 3, which shows the star formation rate (top panel) and both the total and neutral mass loading as a function of time. The figure shows that the total star formation rate in our simulations only varies by $\approx 30\%$ between the different resolutions, and the total mass loading factor η_{total} only differs by tens of percent between 1 and 2 pc resolution, and by a factor of a few from 2 pc to 4 pc, compared to the many order of magnitude difference we find in neutral mass loading.

Acceleration of the neutral gas

In order to better understand why the neutral gas survives and how it is accelerated into the wind in our high-resolution simulation, and why this does not occur at lower resolution, we must investigate its structure and momentum budget. In Fig. 4 we show the vertical profile of gas density, mass-weighted mean vertical velocity, and mass loading for the $\Delta x = 1$ pc simulation at $t = 60$ Myr. We see that neutral gas dominates the mass and mass loading at every height above the disc, tracking the total to within a factor of two, with the hot phase about an order of magnitude less dense and the warm ionised medium between them. The hot gas velocity rises steeply near the disc, reaching ≈ 400 km s $^{-1}$ within a few hundred pc of the midplane and $600 - 700$ km s $^{-1}$ by the time it reaches the edge of the simulation domain, while the neutral velocity rises more slowly, not passing 100 km s $^{-1}$ until heights of ≈ 1 kpc, and reaching ~ 200 km s $^{-1}$ only at the top of the box.

A velocity that rises with height does not by itself mean the gas is accelerating, since the same sort of pattern could be produced by a Hubble-type flow where the highest velocity material is at the largest distances simply because it was launched fastest. To establish that the neutral gas is in fact accelerating, we estimate the Lagrangian derivative of the neutral velocity as described in **Methods**, and plot the result in Fig. 5a. We see that the neutral gas is in fact accelerating upwards at ~ 5 km s $^{-1}$ Myr $^{-1}$ out to distances of ≈ 2.5 kpc from the disc, after which point it coasts at near constant velocity. By contrast, gas travelling ballistically under the influence of gravity would be expected to *decelerate* at $g \approx 5 - 10$ km s $^{-1}$ Myr $^{-1}$ over the same range of heights (dashed green line in Fig. 5a); indeed, a parcel launched ballistically anywhere below 1.0 kpc at the mean neutral gas speed we measure there would stall and turn around before reaching the top of the simulation domain.

The reservoir of pressure, and thus of momentum flux, that can counteract gravity is the hot phase, which carries $\approx 99\%$ of the volume-filling pressure $f_V \langle P \rangle$ (Fig. 5b); here we define $f_V \langle P \rangle$ as the product of a phase's volume filling fraction and its mean pressure, so that the phases sum exactly to the total. The hot phase pressure falls monotonically from $f_V \langle P \rangle / k_B \approx 2 \times 10^5$ to 5.0×10^4 K cm $^{-3}$ between 0.5 and 3 kpc, sustaining an outward gradient across the whole box.

This gradient is much larger than would be needed to support the hot gas against gravity, or even to accelerate it, implying that the decrease in pressure of the hot phase is caused by a transfer of momentum to other phases. To demonstrate this, note that the rate of change of momentum per unit volume, or equivalently the force density, applied to the gas in phase p is $\rho_p(Dv_p/Dt + g)$; when summed over all three phases, thus must equal the decrease in the available momentum flux in the hot gas pressure, $-dP_{\text{hot}}/dz$. We evaluate these various terms as described in **Methods**, and show the result in Fig. 5c. The figure shows

that, integrated over 0.5–3 kpc, the neutral phase takes up 2.5 times as much momentum as the hot phase does, and the two together fully account for the decrease in hot gas pressure to within $\sim 30\%$ – very good agreement given the approximations that are necessary to compute the various terms. The neutral phase dominates the uptake of momentum because, while it moves at a third of the hot velocity, it carries seven times the mass.

The remaining question then becomes how the hot gas transfers momentum to the neutral phase, for which there are two plausible channels. One is ram pressure combined with drag: as the hot gas flows past the cold phase, hot material impacting the surface of cold clouds and the boundary layers around them exerts a force that transmits momentum. The other is condensation: if hot gas cools into the neutral phase, it will carry its higher momentum with it, thereby increasing the mean momentum of the cool gas. Including both sources of acceleration, and the effects of gravity, we can express the total acceleration experienced by a neutral cloud of mass M and cross-sectional area A as

$$\frac{Dv}{Dt} = -|g| + \frac{1}{M} \left(\frac{1}{2} C_D A \rho_h \Delta v^2 + \Gamma_{h \rightarrow H I} \Delta v \right), \quad (1)$$

where g is the gravitational acceleration, C_D is a coefficient of drag of order unity, ρ_h is the density of the hot gas impacting the cloud, Δv is the velocity difference between the hot wind and the neutral cloud, and $\Gamma_{h \rightarrow H I}$ is the net rate at which hot gas condenses onto the cloud; the first term in parentheses represents ram pressure and drag, and the second represents condensation. The main challenge in evaluating these two terms is the choice of ρ_h and Δv , since these vary between the mean properties of the hot gas and the properties in a boundary layer around each neutral cloud. We consider both possibilities, and thereby estimate the mean value of both terms on the right hand side over all neutral clouds, as described in **Methods**. We show the range of estimates at which we arrive as the orange and purple bands in Fig. 5a.

The Figure shows that the two contribute comparably at 1.5 kpc, with condensation increasingly dominant higher up, and that neither suffices alone: even at $C_D = 2$ ram pressure reaches only $6.5 \text{ km s}^{-1} \text{ Myr}^{-1}$ against $|g| = 8.6 \text{ km s}^{-1} \text{ Myr}^{-1}$ at 1.5 kpc, so drag can never lift the gas by itself at any height. By contrast, the sum of condensation and ram pressure is stronger than gravity; we show this sum as the black bars in Fig. 5a, demonstrating that the forces we have computed fully explain the acceleration of H I clouds that we observe in our high-resolution simulation.

On the survival of neutral gas, and the importance of resolution

Condensation also explains why H I structures exist at all. Comparing the mean cloud density to the mean hot gas density the gas in the shell immediately upstream of it (see **Methods**) we find a typical density contrast $\chi = \rho_{cl}/\rho_h \approx 400$ and a typical relative velocity is $\Delta v \approx 150\text{--}180 \text{ km s}^{-1}$. The cloud-crushing time for a cloud of radius R_{eff} is $t_{cc} = \chi^{1/2} R_{\text{eff}}/\Delta v$, and for the median cloud radius identified by our structural analysis we find $t_{cc} \approx 0.7 \text{ Myr}$ at every height, compared to a flow time $z/v_z \sim 15 \text{ Myr}$. Thus the neutral clouds have been in the wind for roughly 20 cloud-crushing times.

An adiabatic cloud would certainly have been destroyed by this time, but condensation can compensate for the mass loss caused by stripping and allow clouds to survive; [Gronke and Oh \(2018\)](#) find that cold clouds survive if their boundary layers cool fast enough, with a quantitative requirement that the cooling time in the mixing layer satisfy $t_{\text{cool,mix}} < t_{cc}$, where the cooling time t_{cool} is defined as the ratio of the gas internal energy to the radiative cooling rate, and $t_{\text{cool,mix}}$ is evaluated at temperature $T_{\text{mix}} = \sqrt{T_h T_{cl}}$ and density $\rho_{\text{mix}} = \sqrt{\rho_h \rho_{cl}}$, i.e., at the geometric mean temperature and density between conditions in the cloud and in the surrounding hot gas.

We evaluate ρ_{mix} , T_{mix} , and $t_{\text{cool,mix}}$ from our simulations as described in **Methods**, and report the median values we obtain for our H I clouds in **Extended Data Table 3**. In each of our three sample slabs we find cooling times $t_{\text{cool,mix}} < 0.08 \text{ Myr}$, roughly an order of magnitude smaller than the cloud crushing time. Equivalently, the critical radius for cloud survival $r_{\text{crit}} = \Delta v t_{\text{cool,mix}}/\chi^{1/2}$ is $\sim 0.5 \text{ pc}$, an order of magnitude below the median cloud effective radius R_{eff} (**Extended Data Table 3**). This explains why we observe that our H I clouds grow rather than being destroyed as the wind moves upward.

This analysis also explains why our lower-resolution simulations, like earlier simulations with similar resolutions, failed to find H I clouds in supernova-driven winds. The characteristic size of the cooling layer is $\ell_{\text{cool}} = c_s t_{\text{cool,mix}}$, where c_s is the sound speed in the mixing layer (i.e., the sound speed evaluated at a temperature T_{mix} ; [Gronke and Oh 2018](#)). Evaluating this for our values of T_{mix} for our sample slabs gives $\ell_{\text{cool}} \approx 3\text{--}6 \text{ pc}$ at heights of 1.5 to 3.5 kpc (**Extended Data Table 3**). In our $\Delta x = 1 \text{ pc}$ resolution simulation, these layers are at least marginally resolved by $\approx 3\text{--}6$ cells, but at a resolution $\Delta x = 2 \text{ pc}$ this becomes even more marginal, $\approx 1.5\text{--}3$ cells, and at $\Delta x = 4 \text{ pc}$ the mixing layer is reduced to a single

cell, and even the typical neutral cloud ($R_{\text{eff}} \sim 6$ pc) is at most two cells in size. Consequently, our lower-resolution simulations fail to capture the mixing layers that are responsible for both half the momentum transfer from the hot phase to the H I, and for all of the condensation that re-supplies the H I.

Discussion and conclusions

One strong implication of our work is that the steep vertical decline of the neutral phase reported by every previous stratified-box and Galactic-centre calculation of supernova-driven outflows is, at least in part, numerical. Our own 4 pc run reproduces that decline; the same physics at 1 pc does not. Our analysis of the boundary layers around the H I clouds produced in our 1 pc resolution simulation provides a quantitative explanation for this phenomenon: in order to produce H I accelerated into winds, simulations must resolve the boundary layers responsible for both re-supplying their mass and transferring hot gas momentum to them. The characteristic sizes of these structures are a few pc, so resolution $\Delta x = 1$ pc is sufficient and $\Delta x = 4$ pc is clearly not; $\Delta x = 2$ pc is marginal, and the outcome likely depends on details of the adopted cooling curve and on the level of numerical dissipation in the simulation method.

We also caution that our results for H I survival are conservative in the sense that our calculations are unmagnetised and include no explicit thermal conduction, meaning that of the three mechanisms that have been proposed to explain the survival of cool, neutral material in winds – radiative cooling, conduction, and magnetic fields – our simulations contain only one. This demonstrates that radiative cooling of the mixing layer is sufficient on its own, but it is possible that adding conduction and magnetic fields might increase the dominance of neutral gas in the mass and mass flux budget even further. It is unclear how including these processes would modify the resolution requirement we have found, since both magnetic fields and conduction conceivably alter the thickness of the mixing layer.

In addition, although we have demonstrated a qualitative change in behaviour with resolution – from $\eta_{\text{HI}} \sim 0$ to $\eta_{\text{HI}} \sim 0.5$ as we go from $\Delta x = 4$ pc to $\Delta x = 1$ pc, we have not yet established full numerical convergence: η_{HI} is noticeably larger for $\Delta x = 1$ pc than for $\Delta x = 2$ pc, and it is entirely possible that it would become larger still if we could simulate at $\Delta x = 1/2$ pc. Our reported neutral mass loading should therefore be treated as a lower bound. Moreover, while we are able to analyse the statistical properties of the H I clouds that form in our simulations, we are limited in our ability to follow the detailed evolution of individual objects by the cadence of our outputs, which is a factor of a few longer than the cloud-crushing time and an order of magnitude longer than the cooling time.

Despite these limitations, our simulations’ demonstration that supernova-driven winds can be dominated by neutral material even at distances many kpc above the galactic plane has important consequences that reach beyond stratified box simulations. Subgrid wind prescriptions in cosmological simulations are calibrated on resolved calculations of exactly this kind (e.g., [Smith et al. 2024](#)), so if those calculations lose their neutral phase numerically due to inadequate resolution, the prescriptions derived from them inherit the loss. The error is not a uniform normalisation but grows with height, since at 4 pc resolution the neutral loading falls by a factor of 20 between 1 kpc and the box edge while at 1 pc resolution it is flat. A second consequence is that, because the neutral phase is fed by condensation as it travels, the phase composition of a wind at the halo boundary is not the composition at launch. Winds cannot be treated as a fixed mixture advected outward. Instead, the neutral fraction delivered to the circumgalactic medium is set en route rather than at the disc.

Our finding that neutral gas can be carried to several kiloparsecs by supernova feedback alone also means that some of the motivation for simulations emphasising cosmic rays or radiation pressure as wind drivers is reduced. Of course this does not mean that these processes do not occur, simply that there is no need to tune uncertain parameters in models of these processes (e.g., cosmic ray diffusion coefficients) in an attempt to reproduce observations of neutral material entrained high into winds. Our findings may therefore lead to a revision of models for cosmic ray- and radiation pressure-driven winds as well.

Methods

Simulations: initial conditions

We evolve a stratified patch of a galactic disc, using periodic boundary conditions in the horizontal (x and y) directions, and “diode” boundary conditions that permit outflow but suppress inflow in the z direction – see [Vijayan et al. \(2024\)](#) for details. The domain is 0.98×0.98 kpc horizontally and spans $|z| \leq 3.92$ kpc, or $0.98 \times 0.98 \times \pm 3.92$ kpc, and we use three cell counts: $1024^2 \times 8192$, $512^2 \times 4096$, and $256^2 \times 2048$, giving spacing of $\Delta x = 0.96$ pc, 1.91 pc, and 3.83 pc, respectively; for brevity we will refer to these cases as $\Delta x = 1$, 2 , and 4 pc. For comparison with previous work, the TIGRESS-NCR LGR4 model reaches $\Delta x = 2$ pc in a $512^2 \times 3072$ pc domain ([Kim et al. 2023](#)), so our 1 pc run resolves roughly nine times the volume at half the cell size.

All runs begin with a gas surface density $\Sigma_{\text{gas}} = 50 \text{ M}_{\odot} \text{ pc}^{-2}$, with the gas initially uniform in x and y and stratified in z with a vertical profile chosen to be in hydrostatic equilibrium. The gas initially has a uniform temperature $T = 10^4$ K, and a metallicity $Z = 1 Z_{\odot}$. In addition to the gravitational potential generated by solving the Poisson equation for the gas and stellar particles (see next section), our simulations also include an “external” potential that combines a stellar disc and a dark-matter halo,

$$\Phi_{\text{*dm}} = 2\pi G \Sigma_{\text{*}} z_{\text{*}} \left[\left(1 + \frac{z^2}{z_{\text{*}}^2} \right)^{1/2} - 1 \right] + 2\pi G \rho_{\text{dm}} R_0^2 \ln \left(1 + \frac{z^2}{R_0^2} \right), \quad (2)$$

where we adopt $\Sigma_{\text{*}} = 208 \text{ M}_{\odot} \text{ pc}^{-2}$, $\rho_{\text{dm}} = 2.4 \times 10^{-2} \text{ M}_{\odot} \text{ pc}^{-3}$, $z_{\text{*}} = 245$ pc, and $R_0 = 4$ kpc; these parameters are identical to those in the R4 model of ([Kim et al. 2020](#)). See [Vijayan et al. \(2024\)](#) for further details on the model setup.

We run each simulation for the times indicated in **Extended Data Table 1**, writing the full 3D structure at a cadence of 2 Myr, and reduced 2D projections of select quantities at a cadence of 0.5 Myr.

Simulations: numerical method

We carry out all simulations with QUOKKA, a GPU-native, block-structured adaptive-mesh radiation-hydrodynamics code built on AMReX ([Wibking and Krumholz 2022](#); [He et al. 2024a,b](#); [He et al. 2026](#)). We use a limited subset of its capabilities: the radiation-transport module is switched off, and mesh refinement is not used, so each run is a single uniform Cartesian grid. The code solves the Euler equations of hydrodynamics coupled to the Poisson equation for self-gravity (along with an external gravitational potential), optically thin radiative cooling, photoelectric heating and star-particle feedback. The hydrodynamics uses a piecewise-parabolic reconstruction of the primitive variables ([Colella and Woodward 1984](#)) with the HLLC Riemann solver, a method-of-lines formulation and second-order strong-stability-preserving Runge–Kutta time-stepping.

Star formation and supernova injection are self-consistent rather than prescribed, and use the “particle–mesh–particle” scheme of [He et al. \(2026\)](#). In cells that satisfy a local Jeans-instability criterion, the code stochastically forms a collisionless star particle at a rate set by the local free-fall efficiency, drawing stellar masses from a Chabrier initial mass function ([Chabrier 2003](#)); stars massive enough to explode as supernovae ($> 9 \text{ M}_{\odot}$ following the models of [Sukhbold et al. 2016](#)) are following individually, while lower-mass stars are lumped together into a composite particle. The resulting star particles are given an initial velocity distribution that includes a tail to high velocities corresponding to runaway stars (see [He et al. 2026](#)), and move freely through the domain under the effects of gravity. Once a particle in the range of initial masses predicted by [Sukhbold et al.](#) to explode as supernovae reaches an age exceeds its progenitor’s main-sequence lifetime (as determined from the MIST stellar evolution tracks – [Choi et al. 2016](#)) it injects $\Delta E_{\text{SN}} = 10^{51}$ erg, $\Delta M_{\text{SN}} = 10 \text{ M}_{\odot}$ and $\Delta M_{Z,\text{SN}} = 1 \text{ M}_{\odot}$ of energy, mass, and metals into its host cell; we treat metals as a passive scalar. Supernova feedback uses a momentum-energy switch, as discussed in [He et al. \(2026\)](#), but due to the very high resolution of our simulations the momentum branch is almost never used, and so in most cases supernova feedback is handled as pure energy deposition.

While we do not include radiative transport, we do include radiative heating and cooling using the GRACKLE library ([Smith et al. 2017](#)). Specifically, we use the GRACKLE tables computed for a uniform far-ultraviolet (FUV) background of [Haardt and Madau \(2012\)](#) with self-shielding; we recompute these tables as functions of (ρ, E_{int}) rather than the default (ρ, T) for computational efficiency, but the underlying physics is identical. The tables include primordial (H and He) cooling and heating, metal-line cooling and heating, and inverse Compton heating by cosmic microwave background photons; as in [Vijayan et al. \(2026\)](#), we hold the metal line cooling table at a fixed metallicity $Z = Z_{\odot}$ rather than evolving it in response to metal deposition by supernovae.

In addition, we include a photoelectric heating rate that tracks the instantaneous, spatially uniform FUV field produced by the simulated stellar population. The instantaneous photoelectric heating rate per H nucleon depends on star formation history in the simulation box as

$$\Gamma_{\text{PE}}(t) = \Gamma_{\text{PE},0} \frac{\int_0^t \psi(t-t') \dot{\Sigma}_*(t') dt'}{\int_0^{1 \text{ Gyr}} \psi(t-t') \dot{\Sigma}_{*,0} dt'} \quad (3)$$

where $\Gamma_{\text{PE},0} = 4.0 \times 10^{-26} \text{ erg s}^{-1} \text{ H}^{-1}$ is the Solar neighborhood photoelectric heating rate (Draine 2011), $\dot{\Sigma}_*(t)$ is the star formation rate per unit area in the simulation box as a function of time, $\psi(t)$ is the FUV luminosity per unit mass produced by a stellar population of age t , and $\dot{\Sigma}_{*,0} = 2.5 \times 10^{-3} \text{ M}_\odot \text{ yr}^{-1} \text{ kpc}^{-2}$ is the star formation rate in the Solar neighbourhood (Fuchs et al. 2009). We compute $\psi(t)$ by using the SLUG stellar population synthesis code (da Silva et al. 2012; Krumholz et al. 2015) to compute the total luminosity in photons with energies $E_\gamma = 8 - 13.6 \text{ eV}$ (the approximate energy range responsible for photoelectric heating) for a continuously-sampled stellar population as a function of age t , using MIST stellar tracks (Choi et al. 2016), the same Chabrier IMF used in the simulations, and SLUG’s “starburst99” spectral synthesis mode (Leitherer et al. 1999; Vázquez and Leitherer 2005). Note that the exact choice of 1 Gyr as the upper limit in the integral in the denominator has almost no effect on the value of the integral, since $\psi(t)$ goes to near-zero values for $t \gtrsim 200 \text{ Myr}$; thus any value greater than a few hundred Myr yields the same result to within a few percent.

Phase-decomposition

In our analysis of the simulation outcomes we split the gas into three phases that partition it exactly: *neutral*, meaning gas with electron abundance $x_e < 0.05$; *hot*, meaning $T > 10^6 \text{ K}$; and the *warm ionised medium* (WIM), the remainder. Our simulations do not follow a non-equilibrium chemical network, so we reconstruct x_e from the tabulated equation of state from GRACKLE, which provides the mean molecular weight μ as well as the cooling rate as a function of gas density and internal energy. For GRACKLE’s adopted hydrogen mass fraction $X = 10/14$ and helium mass fraction $Y = 4/14$, the electron abundance per H nucleon $x_e \equiv n_e/n_{\text{H}}$ is then related to the mean molecular weight as $x_e = (\mu^{-1} - Y/4)/X - 1$. We deliberately define the neutral phase in terms of the electron abundance rather than simply using a temperature cut because the temperature at which the gas transitions from predominantly ionized to predominantly neutral is density-dependent, ranging from $10^{3.8} - 10^{4.2} \text{ K}$, so using a simple temperature cut would not differentiate cleanly between ionized and neutral material at similar temperatures. For the purposes of comparing to observations in the 21 cm line, our primary source of information on the detailed structure of the neutral material entrained into outflows, the ionization fraction is the important quantity.

Outflow rates

For each phase p the mass outflow rate at height h is the net flux through the pair of planes at $z = \pm h$. Formally, we evaluate this as

$$\dot{M}_p(h, t) = \int_{z=+h} \rho_p v_z dA - \int_{z=-h} \rho_p v_z dA, \quad (4)$$

where ρ_p is the gas density for cells classified as belonging to phase p and 0 otherwise, and v_z is the vertical velocity. We evaluate this quantity at the quarter, half, three-quarter and full box half-height, $h = 0.98, 1.96, 2.93$ and 3.91 kpc , which in the text we round to 1, 2, 3 and 4 kpc for brevity. We define the corresponding mass loading factor of each phase as $\eta_p = \dot{M}_p / \langle \text{SFR} \rangle$, where $\langle \text{SFR} \rangle$ is the time-averaged star formation rate. Note that our outflow rates are net mass fluxes rather than the outflow-only values reported in some works (e.g., Rathjen et al. 2023, who replace v_z with $\max(v_z, 0)$ in the first integral and $\min(v_z, 0)$ in the second), so comparison must be made with care.

We measure the outflow rates summed over all phases and for the hot phase at 0.5 Myr intervals, and the outflow rate in the neutral phase at a coarser cadence of 2 Myr, since it is more expensive to compute because x_e must be obtained in post-process rather than on-the-fly as the simulation runs. The WIM rate is the remainder, $\dot{M}_{\text{WIM}} = \dot{M}_{\text{total}} - \dot{M}_{\text{HI}} - \dot{M}_{\text{hot}}$, so wherever the three are compared they carry the neutral measurement’s cadence.

Throughout this work, all time-averaged quantities – including the star formation rate used to normalise η_p – are computed over a common interval of 29–92 Myr. The starting time is chosen to exclude an initial transient phase, which lasts for roughly one free-fall time at the box half-height. We find that star formation in the simulation is bursty and episodic once it has begun, with period of $\approx 120 \text{ Myr}$, and so we choose the final time of our averaging interval to roughly span one of these cycles.

Neutral cloud structural analysis

To quantify the structure of the neutral clouds that form in the full resolution simulation, we select five snapshots that are roughly uniformly-spaced in time from 49 – 70.6 Myr, tracking the period of maximum neutral outflow rate. In each snapshot, we select for analysis three slabs of thickness 600 pc, chosen to provide sufficient statistical sampling, covering heights 1.2 – 1.8 kpc, 2.2 – 2.8 kpc, and 3.2 – 3.8 kpc. Within each slab, we define a neutral cloud as a connected set of cells satisfying $x_e < 0.05$, where connection is evaluated keeping in mind that the domain is periodic in the horizontal direction. We discard objects containing fewer than 500 cells (an effective radius of 4.7 pc) as unresolved, and use the remainder for our analysis.

For each identified cloud, we compute the following statistics: total neutral mass M , effective radius R_{eff} (defined as the radius of a sphere whose volume is equal to the cloud volume), atomic hydrogen number density n_{H} (defined as a neutral mass-weighted mean), temperature T (also a neutral mass-weighted mean), and equivalent column density N_{H} (defined as the line of sight hydrogen column density through the centre of a spherical cloud of equal mass and volume). Quantitatively, for a set of N_c cells belonging to a given cloud, we define

$$M = \Delta x^3 \sum_{i=1}^{N_c} (1 - x_{e,i}) \rho_i \quad (5)$$

$$R_{\text{eff}} = \left(\frac{3}{4\pi} N_c \Delta x^3 \right)^{1/3} \quad (6)$$

$$n_{\text{H}} = \frac{1}{M \mu_{\text{H}}} \Delta x^3 \sum_{i=1}^{N_c} (1 - x_{e,i})^2 \rho_i^2 \quad (7)$$

$$T = \frac{1}{M} \Delta x^3 \sum_{i=1}^{N_c} (1 - x_{e,i}) \rho_i T_i \quad (8)$$

$$N_{\text{H}} = 2R_{\text{eff}} \left(\frac{3M}{4\pi R_{\text{eff}}^3 \mu_{\text{H}}} \right), \quad (9)$$

where ρ_i , T_i , and $x_{e,i}$ are the total density, temperature, and electron fraction in cell i and $\mu_{\text{H}} = 2.3 \times 10^{-24}$ g is the mean mass per H nucleon for our adopted chemical composition.

Finally, we compute the following statistics for the cloud population as a whole in each slab: the number density of clouds n_{cl} (averaged over the five time samples), fraction of cloud mass $f_{M,30}$ held in clouds with $R_{\text{eff}} > 30$ pc (again averaged over the five samples), and the median and inter-quartile range of each of the per-cloud statistics, computed both by number and by mass.

Lagrangian acceleration and the momentum budget

We estimate the Lagrangian vertical acceleration of the neutral gas as a function of height off the plane z ,

$$\frac{Dv_z}{Dt} = \left. \frac{\partial v_z}{\partial t} \right|_z + v_z \left. \frac{\partial v_z}{\partial z} \right|_t, \quad (10)$$

as follows. The advective term $v_z(\partial v_z/\partial z)_t$ is a spatial derivative at one instant, which we evaluate from the single snapshot at $t = 60$ Myr. For this snapshot, we first compute the mass-weighted mean velocity of the neutral gas at height z by averaging the neutral material in slabs $\Delta z = 0.6$ kpc thick, and averaging over the two sides of the midplane; we apply this slab-averaging procedure to reduce small-scale noise, which would otherwise dominate our estimate of the derivative. We then compute the derivative of the vertical velocity with respect to z using centred differences on these slab averages. To estimate the Eulerian term $(\partial v_z/\partial t)_z$, we compute the same $\Delta z = 0.6$ kpc-slab averages on the $t = 55.4$ and 64 Myr snapshots, and estimate $(\partial v_z/\partial t)_z$ at $t = 60$ Myr from a centred difference between these two times.

We obtain estimates of the uncertainty the spatial derivative at fixed time by varying the slab width from 0.5 to 1.5 times our fiducial value, and on the time derivative by computing two one-sided differences from the same three time slices (i.e., from 55.4 to 60 Myr and from 60 to 64 Myr) and taking half the difference between the two one-sided differences as our uncertainty estimate. We obtain a total uncertainty on the final Lagrangian derivative Dv_z/Dt by adding the spatial and temporal uncertainties in quadrature. We show the measurements of Dv_z/Dt produced by this procedure as the blue line and band in Fig. 5a. We also use the same method to compute the Dv_z/Dt term for the hot gas that is required for the red line in Fig. 55c.

Gas acceleration by ram pressure and condensation

Evaluating the terms describing gas acceleration by ram pressure / drag and condensation in Equation 1 requires that we compute or adopt estimates for the cloud mass M , cross-sectional area A , the density of the hot gas incident on the cloud ρ_h , the hot gas-H I velocity difference Δv , the condensation rate $\Gamma_{h \rightarrow \text{H I}}$, and coefficient of drag C_D . We estimate these terms using the sample population of neutral clouds produced by our segmentation procedure, as described above, which provides a catalogue of H I clouds at five times from 49 – 70.6 Myr and at three heights, centred on 1.5, 2.5, and 3.5 kpc.

We have already defined the mass M for each of these clouds, as describe above. To compute the cross-sectional area A , we flatten the set of cells defining the cloud onto the xy plane and computing the area of the resulting set of pixels; mathematically, our definition is

$$A = \Delta x^2 \sum_{ij} \max_k (I_{\text{cloud},ijk}) \quad (11)$$

where $I_{\text{cloud},ijk}$ is an indicator function on cell ijk which is unity if a cell is identified as belonging to a cloud and 0 otherwise. We take the ratio M/A that appears in the drag term in Equation 1 to be simply the mean cloud mass divided by the mean cloud area in each slab.

As noted in the main text, we consider two possible estimates for the hot gas density ρ_h and hot-neutral velocity difference Δv , one based on mean properties and one based on properties in cloud boundary layers. For our mean property-based estimate, we simply evaluate ρ_h and $\Delta v = v_{h,z} - v_{\text{H I},z}$ using the slab-averaged means shown in Fig. 4 at $t = 60$ Myr, and from the analogous quantities computed at our other sample times. For our boundary-layer estimate, we use the same value of $v_{\text{H I},z}$ for the mean H I velocity, and we estimate mean values of the hot gas density and velocity as follows:

1. Define an initial boundary shell around each H I cloud by starting with the cells of the cloud itself and applying $N_d = \max[3, \text{round}(R_{\text{eff}}/\Delta x)]$ iterations of binary dilation with a structuring element with a connectivity of 1. This defines a shell around each cloud with a characteristic radius comparable to the cloud size, but with a minimum size of three cells, and whose shape is contoured to the cloud shape.
2. Next, remove any neutral or non-hot cells, those with $x_e < 0.05$ or $T < 10^6$ K, leaving only cells classified as hot in our phase decomposition.
3. Finally, discard any cells for which the distance from the midplane to their centre, $|z|$, is larger than the cloud's distance to the midplane, defined as the median $|z|$ of cells that belong to the cloud. This step means that we keep only the parts of the shell that are closer to the midplane than cloud centre, and are therefore upstream of it.

Once we have defined the cells belonging to the boundary layer around each cloud, we compute the mean density ρ_h and mass-weighted mean z velocity $v_{h,z}$ over all the cells tagged as belonging to a boundary layer in each slab, and use these for our boundary-layer estimate of the drag / ram pressure force.

To estimate the condensation rate $\Gamma_{h \rightarrow \text{H I}}$ in each of our sample slabs, we note that, for any phase p , the rate of change in the total mass M_p of that phase within a slab from z_0 to z_1 obeys

$$\frac{dM_p}{dt} = \dot{M}_p(z_0, t) - \dot{M}_p(z_1, t) + \Gamma_p, \quad (12)$$

where $\dot{M}_p(h, t)$ is the outward mass fluxes through the surface at height h (Eq. 4) and Γ_p is the net rate at which phase p gains mass from the other two. We estimate dM_p/dt for each phase over our five sample snapshots by computing the mass M_p of each phase inside each slab at each time, and then carrying out a simple least-squares fit to find the slope. We directly measure the fluxes from the simulation at each time; using these fluxes together with our estimate of dM_p/dt yields an estimate for Γ_p at each time, which is the quantity of interest. Rather than use $\Gamma_{\text{H I}}$ estimated in this way directly, we instead evaluate it as $\Gamma_{\text{H I}} = -(\Gamma_{\text{hot}} + \Gamma_{\text{WIM}})$; we do so because the time series of H I mass and instantaneous mass fluxes are rather noisy due to the clumpy H I structure. The WIM and hot gas masses vary much more smoothly in time, allowing significantly better estimates.

Finally, for C_D plausible values are in the range $C_D \approx 1 - 2$, with values smaller for smoothly tapered structures such as spheres and larger values for smooth or blunt bodies. We compute estimates for both $C_D = 1$ and 2 to give a plausible uncertainty range.

Cloud crushing times and mixing layer cooling times

Evaluating the Gronke and Oh (2018) criterion for cloud survival requires that we compute the cloud crushing time t_{cc} and the mixing-layer cooling time $t_{\text{cool,mix}}$; these in turn depend on the density contrast

χ between the H I clouds and the hot material around them, the temperatures T_h and T_{cl} in the hot gas and the cloud, and the densities ρ_{cl} and ρ_h of the hot gas and cloud.

We compute χ in each of our sample slabs by dividing the number-weighted median densities of H I clouds found by our structural analysis by the number-weighted median densities of their boundary layers, with the boundary layers defined in the same way as above. We use the number-weighted medians rather than the mass-weighted medians because the former emphasise values typical of the small clouds that are most vulnerable to disruption by the wind. We report the values of χ and the corresponding cloud crushing times that we obtain via this procedure in **Extended Data Table 3**.

We similarly evaluate T_h , T_{cl} , ρ_h , and ρ_{cl} using the number-weighted medians in the hot regions around the clouds and in the clouds themselves. Once these values are determined, we compute cooling times at $\rho_{mix} = \sqrt{\rho_h \rho_{cl}}$ and $T_{mix} = \sqrt{T_h T_{cl}}$ by reading the cooling rate at that combination of density and temperature from our GRACKLE cooling tables, and dividing this into the internal energy at the corresponding density and temperature. We report the values of $t_{cool,mix}$ that we obtain by this procedure for each of our sample slab regions in **Extended Data Table 3**.




Fig. 1 Edge-on view of the 1 pc resolution simulation at $t = 60$ Myr. Columns show, from left to right, the projected neutral gas column N_{HI} , total density ρ in a slice, gas temperature T in a slice, and gas vertical velocity v_z in a slice, coloured so that velocities away from the midplane are red. Insets zoom in on one example neutral cloud ≈ 3.5 kpc above the plane, with dashed black contours marking the locus $x_e = 0.05$ that we use to define the transition from ionized to neutral.




Fig. 2 Neutral gas column density projections at different simulation resolutions. The three panels show the projected neutral column at $t = 88$ Myr at our three resolutions of $\Delta x = 1$ pc (left column), 2 pc (middle column), and 4 pc (right column). The insets show the same region as the insets in Fig. 1. For clarity we show only a portion of the full simulation domain in these figures.




Fig. 3 Star formation rate (top) and mass loading factor (bottom) versus time in the three simulations. The mass-loading factor is measured at the top of the simulation domain, $z = 4 \text{ kpc}$, and we show both the total mass-loading including all phases (solid) and the component carried by the neutral phase alone (dashed).




Fig. 4 Vertical structure of the outflow at $t = 60$ Myr. The panels show, from top to bottom, the mean density of the gas expressed as H nucleons per unit volume n_H , the mass-weighted mean vertical velocity v_z , and the mass loading factor η as a function of height z ; all quantities are averages over slabs of thickness $\Delta z = \dots$. Black shows the sum over all phases, while blue, green, and red show the neutral, WIM, and hot phases, respectively.




Fig. 5 Neutral gas acceleration and momentum budget. (a) Measured Lagrangian acceleration of the neutral gas (blue, with the shaded band its uncertainty), the gravitational acceleration (green dashed), the momentum transferred from the hot phase to the neutral phase by ram pressure and condensation at $z = 1.5, 2.5,$ and 3.5 kpc (orange and purple bars, respectively), and the sum of ram pressure, condensation, and gravity (black bars) at heights $z \approx 1.5, 2.5,$ and 3.5 kpc. The ram pressure bar spans $C_D = 1$ and 2 , and the condensation bar spans an arrival velocity from the boundary-layer contrast to the hot-phase mean. (b) Volume-filling thermal pressure: total (black), hot (red dashed), and the rest of the gas, H I plus WIM (blue). (c) Force density per phase, H I (blue), hot (red) and their sum (green); solid lines give the total force required to explain the measured acceleration, while dashed lines show the force required to counter gravity. The solid black line is $-dP_{\text{hot}}/dz$, the force provided by the hot phase. Accelerations, pressures, and forces are all measured at $t = 60$ Myr, while the ram pressure and condensation points average over five snapshots spanning $49 - 70.6$ Myr; see **Methods**.

Acknowledgements. CCH, MRK, and AV acknowledge support from the Australian Research Council through Laureate Fellowship FL220100020. This work was supported by computing resources provided by the INCITE program at the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725, and by the Pawsey Supercomputing Research Centre’s Setonix Supercomputer (<https://doi.org/10.48569/18sb-8s43>), with funding from the Australian Government and the Government of Western Australia.

Declarations

- Funding: Australian Research Council Laureate Fellowship FL220100020
- Competing interests: the authors declare no competing interests.
- Ethics approval and consent to participate: not applicable
- Consent for publication: not applicable
- Data availability: QUOKKA is publicly available at <https://github.com/quokka-astro/quokka>. Raw simulation outputs are not publicly available due to their large size, but will be provided on request to the corresponding author.
- Materials availability: not applicable
- Author contributions: CCH designed the simulation setup and analysis pipeline, ran the simulations, carried out the analysis, wrote the majority of the manuscript, and is a core developer of the QUOKKA code. MRK helped design the simulation setup and analysis pipeline, and assisted in writing the manuscript. BDW is the lead developer of the QUOKKA code and contributed to the simulation design. AV contributed to the simulation setup design. All authors edited and approved the manuscript.

Extended Data

Table 1 Summary of all simulations. Here we report the total run duration t_{end} , time-averaged star formation rate $\langle \text{SFR} \rangle$, and time-averaged total, H I, hot, and WIM mass-loading factors η_{total} , $\eta_{\text{H I}}$, η_{hot} , and η_{WIM} , computed at the outer edge of the simulation box. All time-averaged quantities are averaged over the interval from 28 Myr, one free-fall time at the box half-height, to 92 Myr, the end of the first star formation burst.

Δx [pc]	t_{end} [Myr]	$\langle \text{SFR} \rangle$ [$M_{\odot} \text{ yr}^{-1}$]	η_{total}	$\eta_{\text{H I}}$	η_{hot}	η_{WIM}
1	114	0.256	0.661	0.288	0.216	0.157
2	176	0.218	0.560	0.167	0.249	0.145
4	256	0.195	0.135	0.000	0.134	0.000

Table 2 Neutral mass loading as a function height and resolution. Here h is the height, $\eta_{\text{H I}}$ is the neutral mass loading factor, and $\eta_{\text{H I}}/\eta_{\text{total}}$ is the fraction of the total mass flux carried by neutral material. All quantities are time averages over 29–92 Myr.

h [kpc]	$\Delta x =$	4 pc	$\eta_{\text{H I}}$ 2 pc	1 pc	$\Delta x =$	4 pc	$\eta_{\text{H I}}/\eta_{\text{total}}$ 2 pc	1 pc
0.98		0.008	0.193	0.309		0.044	0.288	0.393
1.96		0.003	0.177	0.287		0.020	0.275	0.385
2.93		0.001	0.177	0.304		0.005	0.288	0.426
3.91		0.000	0.171	0.289		0.001	0.301	0.431

Table 3 Structural properties of neutral clouds. Statistics of neutral gas clouds – defined as connected sets of cells with electron fraction $x_e < 0.05$, in three 600 pc-thick slabs centred at the z value, from five snapshots spanning 49–70.6 Myr. *Population* quantities characterise the cloud population as a whole, *Count-weighted* and *Mass-weighted* properties give the median 25th to 75th percentile range (quantities in brackets) by number and mass for each cloud properties, and *Survival criterion* lists quantities used to evaluate the [Gronke and Oh \(2018\)](#) cloud survival criterion, all computed from the number-weighted medians. See **Methods** for exact definitions of each quantity.

	1.5	z [kpc] 2.5	3.5
<i>Population</i>			
n_{cl} [kpc $^{-3}$]	1541	1053	965
$f_M (R_{\text{eff}} \geq 30 \text{ pc})$	52%	70%	73%
<i>Count-weighted</i>			
R_{eff} [pc]	6.0 [5.3–7.4]	6.2 [5.3–7.7]	6.2 [5.3–7.9]
M [$M_{\odot}$]	96 [54–215]	58 [34–134]	45 [26–104]
n_{H} [cm $^{-3}$]	2.81 [2.24–3.80]	1.71 [1.43–2.12]	1.29 [1.03–1.65]
T [K]	9432 [8257–9910]	10014 [9389–10380]	10250 [9700–10619]
N_{H} [10^{20} cm^{-2}]	1.09 [0.80–1.66]	0.66 [0.51–0.96]	0.51 [0.38–0.73]
<i>Mass-weighted</i>			
R_{eff} [pc]	21.6 [8.6–43.4]	42.3 [19.8–58.5]	49.5 [24.6–74.0]
M [$10^4 M_{\odot}$]	1.99 [0.05–8.41]	8.14 [0.41–16.55]	9.36 [1.26–18.42]
n_{H} [cm $^{-3}$]	4.10 [3.96–5.08]	3.00 [3.00–3.00]	1.95 [1.95–2.13]
T [K]	6290 [5549–6497]	5144 [5096–5192]	5213 [5180–5246]
N_{H} [10^{20} cm^{-2}]	8.64 [2.76–14.03]	9.84 [4.73–13.13]	7.75 [5.12–9.75]
<i>Survival criterion</i>			
χ	444	374	344
$t_{\text{cool, mix}}$ [Myr]	0.042	0.062	0.077
t_{cc} [Myr]	0.69	0.71	0.73
r_{crit} [pc]	0.37	0.54	0.66
$R_{\text{eff}}/r_{\text{crit}}$	16.4	11.5	9.5
ℓ_{cool} [pc]	3.4	5.6	6.4

References

- Armillaotta, L., Ostriker, E.C., Kim, C.-G., Jiang, Y.-F.: Cosmic-Ray Acceleration of Galactic Outflows in Multiphase Gas. *ApJ* **964**(1), 99 (2024) <https://doi.org/10.3847/1538-4357/ad1e5c> arXiv:2401.04169 [astro-ph.GA]
- Brüggen, M., Scannapieco, E.: The Launching of Cold Clouds by Galaxy Outflows. II. The Role of Thermal Conduction. *ApJ* **822**(1), 31 (2016) <https://doi.org/10.3847/0004-637X/822/1/31> arXiv:1602.01843 [astro-ph.GA]
- Choi, J., Dotter, A., Conroy, C., Cantiello, M., Paxton, B., Johnson, B.D.: Mesa Isochrones and Stellar Tracks (MIST). I. Solar-scaled Models. *ApJ* **823**, 102 (2016) <https://doi.org/10.3847/0004-637X/823/2/102> arXiv:1604.08592 [astro-ph.SR]
- Chabrier, G.: Galactic Stellar and Substellar Initial Mass Function. *PASP* **115**(809), 763–795 (2003) <https://doi.org/10.1086/376392> arXiv:astro-ph/0304382 [astro-ph]
- Cottle, J., Scannapieco, E., Brüggen, M., Banda-Barragán, W., Federrath, C.: The Launching of Cold Clouds by Galaxy Outflows. III. The Influence of Magnetic Fields. *ApJ* **892**(1), 59 (2020) <https://doi.org/10.3847/1538-4357/ab76d1> arXiv:2002.07804 [astro-ph.GA]
- Colella, P., Woodward, P.R.: The Piecewise Parabolic Method (PPM) for Gas-Dynamical Simulations. *Journal of Computational Physics* **54**, 174–201 (1984) [https://doi.org/10.1016/0021-9991\(84\)90143-8](https://doi.org/10.1016/0021-9991(84)90143-8)
- da Silva, R.L., Fumagalli, M., Krumholz, M.: SLUG – Stochastically Lighting Up Galaxies. I. Methods and Validating Tests. *ApJ* **745**, 145 (2012) <https://doi.org/10.1088/0004-637X/745/2/145> arXiv:1106.3072 [astro-ph.IM]
- Di Teodoro, E.M., McClure-Griffiths, N.M., Lockman, F.J., Denbo, S.R., Endsley, R., Ford, H.A., Harrington, K.: Blowing in the Milky Way Wind: Neutral Hydrogen Clouds Tracing the Galactic Nuclear Outflow. *ApJ* **855**(1), 33 (2018) <https://doi.org/10.3847/1538-4357/aaad6a> arXiv:1802.02152 [astro-ph.GA]
- Draine, B.T.: Physics of the Interstellar and Intergalactic Medium. Princeton University Press: Princeton, NJ, ??? (2011)
- Fuchs, B., Jahreiß, H., Flynn, C.: A Schmidt-Kennicutt Law for Star Formation in the Milky Way Disk. *AJ* **137**, 266–271 (2009) <https://doi.org/10.1088/0004-6256/137/1/266> arXiv:0810.1656
- Gronke, M., Oh, S.P.: The growth and entrainment of cold gas in a hot wind. *MNRAS* **480**(1), 111–115 (2018) <https://doi.org/10.1093/mnrasl/sly131> arXiv:1806.02728 [astro-ph.GA]
- Girichidis, P., Walch, S., Naab, T., Gatto, A., Wünsch, R., Glover, S.C.O., Klessen, R.S., Clark, P.C., Peters, T., Derigs, D., Baczynski, C.: The SILCC (SIMulating the LifeCycle of molecular Clouds) project - II. Dynamical evolution of the supernova-driven ISM and the launching of outflows. *MNRAS* **456**(4), 3432–3455 (2016) <https://doi.org/10.1093/mnras/stv2742> arXiv:1508.06646 [astro-ph.GA]
- Haardt, F., Madau, P.: Radiative Transfer in a Clumpy Universe. IV. New Synthesis Models of the Cosmic UV/X-Ray Background. *ApJ* **746**(2), 125 (2012) <https://doi.org/10.1088/0004-637X/746/2/125>
- He, C.-C., Wibking, B.D., Krumholz, M.R.: An asymptotically correct implicit-explicit time integration scheme for finite volume radiation-hydrodynamics. *Monthly Notices of the Royal Astronomical Society* **531**, 1228–1242 (2024) <https://doi.org/10.1093/mnras/stae1244>
- He, C.-C., Wibking, B.D., Krumholz, M.R.: A novel numerical method for mixed-frame multigroup radiation-hydrodynamics with GPU acceleration implemented in the QUOKKA code. *Monthly Notices of the Royal Astronomical Society* **535**, 3059–3076 (2024) <https://doi.org/10.1093/mnras/stae2580>
- He, C.-C., Wibking, B.D., Vijayan, A., Krumholz, M.R., Li, P.S.: A novel algorithm for GPU-accelerated particle-mesh interactions implemented in the QUOKKA code. *The Open Journal of Astrophysics* **9**, 59235 (2026) <https://doi.org/10.33232/001c.159235> arXiv:2509.18261 [astro-ph.IM]
- Kim, C.-G., Armillaotta, L., Ostriker, E.C., Moon, S., Lancaster, L., Kim, J.-G., Linzer, N.B., Hix, R.N.:

- Cosmic-Ray Feedback in Galactic Disks: Star Formation, Cosmic-Ray Transport, and Multiphase Outflows in TIGRESS++ Simulations. *ApJ* **1006**(2), 123 (2026) <https://doi.org/10.3847/1538-4357/ae76d6> [arXiv:2605.26238](#) [astro-ph.GA]
- Krieger, N., Bolatto, A.D., Walter, F., Leroy, A.K., Zschaechner, L.K., Meier, D.S., Ott, J., Weiss, A., Mills, E.A.C., Levy, R.C., Veilleux, S., Gorski, M.: The Molecular Outflow in NGC 253 at a Resolution of Two Parsecs. *ApJ* **881**(1), 43 (2019) <https://doi.org/10.3847/1538-4357/ab2d9c> [arXiv:1907.00731](#) [astro-ph.GA]
- Krumholz, M.R., Fumagalli, M., da Silva, R.L., Rendahl, T., Parra, J.: SLUG - stochastically lighting up galaxies - III. A suite of tools for simulated photometry, spectroscopy, and Bayesian inference with stochastic stellar populations. *MNRAS* **452**, 1447–1467 (2015) <https://doi.org/10.1093/mnras/stv1374> [arXiv:1502.05408](#)
- Kim, C.-G., Kim, J.-G., Gong, M., Ostriker, E.C.: Introducing TIGRESS-NCR. I. Coregulation of the Multiphase Interstellar Medium and Star Formation Rates. *ApJ* **946**(1), 3 (2023) <https://doi.org/10.3847/1538-4357/acbd3a> [arXiv:2211.13293](#) [astro-ph.GA]
- Klein, R.I., McKee, C.F., Colella, P.: On the Hydrodynamic Interaction of Shock Waves with Interstellar Clouds. I. Nonradiative Shocks in Small Clouds. *ApJ* **420**, 213 (1994) <https://doi.org/10.1086/173554>
- Kim, C.-G., Ostriker, E.C.: Three-phase Interstellar Medium in Galaxies Resolving Evolution with Star Formation and Supernova Feedback (TIGRESS): Algorithms, Fiducial Model, and Convergence. *The Astrophysical Journal* **846**, 133 (2017) <https://doi.org/10.3847/1538-4357/aa8599>
- Kim, C.-G., Ostriker, E.C., Somerville, R.S., Bryan, G.L., Fielding, D.B., Forbes, J.C., Hayward, C.C., Hernquist, L., Pandya, V.: First Results from SMAUG: Characterization of Multiphase Galactic Outflows from a Suite of Local Star-forming Galactic Disk Simulations. *ApJ* **900**(1), 61 (2020) <https://doi.org/10.3847/1538-4357/aba962> [arXiv:2006.16315](#) [astro-ph.GA]
- Leidi, G., Andrassy, R., Barsukow, W., Higl, J., Edelmann, P.V.F., Röpke, F.K.: Performance of high-order Godunov-type methods in simulations of astrophysical low Mach number flows. *A&A* **686**, 34 (2024) <https://doi.org/10.1051/0004-6361/202348882> [arXiv:2402.16706](#) [astro-ph.SR]
- Lockman, F.J., Di Teodoro, E.M., McClure-Griffiths, N.M.: Observation of Acceleration of H I Clouds within the Fermi Bubbles. *ApJ* **888**(1), 51 (2020) <https://doi.org/10.3847/1538-4357/ab55d8> [arXiv:1911.06864](#) [astro-ph.GA]
- Leitherer, C., Schaerer, D., Goldader, J.D., Delgado, R.M.G., Robert, C., Kune, D.F., de Mello, D.F., Devost, D., Heckman, T.M.: Starburst99: Synthesis Models for Galaxies with Active Star Formation. *ApJS* **123**, 3–40 (1999) <https://doi.org/10.1086/313233>
- Leroy, A.K., Walter, F., Martini, P., Roussel, H., Sandstrom, K., Ott, J., Weiss, A., Bolatto, A.D., Schuster, K., Dessauges-Zavadsky, M.: The Multi-phase Cold Fountain in M82 Revealed by a Wide, Sensitive Map of the Molecular Interstellar Medium. *ApJ* **814**, 83 (2015) <https://doi.org/10.1088/0004-637X/814/2/83> [arXiv:1509.02932](#)
- Martini, P., Leroy, A.K., Mangum, J.G., Bolatto, A., Keating, K.M., Sandstrom, K., Walter, F.: H I Kinematics along the Minor Axis of M82. *ApJ* **856**(1), 61 (2018) <https://doi.org/10.3847/1538-4357/aab08e> [arXiv:1802.04359](#) [astro-ph.GA]
- McCourt, M., O’Leary, R.M., Madigan, A.-M., Quataert, E.: Magnetized gas clouds can survive acceleration by a hot wind. *MNRAS* **449**(1), 2–7 (2015) <https://doi.org/10.1093/mnras/stv355> [arXiv:1409.6719](#) [astro-ph.GA]
- Noon, K.A., Krumholz, M.R., McClure-Griffiths, N.M., Di Teodoro, E.M., Armillotta, L.: Modelling the non-equilibrium chemistry of the Milky Way’s cold nuclear wind. *MNRAS* **548**(3), 715 (2026) <https://doi.org/10.1093/mnras/stag715> [arXiv:2602.13580](#) [astro-ph.GA]
- Noon, K.A., Krumholz, M.R., Teodoro, E.M.D., McClure-Griffiths, N.M., Lockman, F.J., Armillotta, L.: Direct observations of the atomic-molecular phase transition in the Milky Way’s nuclear wind. Volume 524, Issue 1, 2023, Pages 1258 - 1268 (2023) [arXiv:2304.06356](#)

- Peschken, N., Di Teodoro, E.M., Hanasz, M., Armillotta, L., Tress, R., Naab, T., Wóltanski, D., Gawryszczak, A.: Simulating winds in the Galactic centre: I. Supernova-driven multiphase outflows and HI cloud acceleration. *A&A* **712**, 148 (2026) <https://doi.org/10.1051/0004-6361/202659639> arXiv:2606.16553 [astro-ph.GA]
- Rathjen, T.-E., Naab, T., Walch, S., Seifried, D., Girichidis, P., Wünsch, R.: SILCC - VII. Gas kinematics and multiphase outflows of the simulated ISM at high gas surface densities. *MNRAS* **522**(2), 1843–1862 (2023) <https://doi.org/10.1093/mnras/stad1104> arXiv:2211.15419 [astro-ph.GA]
- Smith, B.D., Bryan, G.L., Glover, S.C.O., Goldbaum, N.J., Turk, M.J., Regan, J., Wise, J.H., Schive, H.-Y., Abel, T., Emerick, A., O’Shea, B.W., Anninos, P., Hummels, C.B., Khochfar, S.: GRACKLE: a chemistry and cooling library for astrophysics. *MNRAS* **466**(2), 2217–2234 (2017) <https://doi.org/10.1093/mnras/stw3291> arXiv:1610.09591 [astro-ph.CO]
- Sukhbold, T., Ertl, T., Woosley, S.E., Brown, J.M., Janka, H.-T.: Core-collapse Supernovae from 9 to 120 Solar Masses Based on Neutrino-powered Explosions. *ApJ* **821**, 38 (2016) <https://doi.org/10.3847/0004-637X/821/1/38> arXiv:1510.04643 [astro-ph.HE]
- Smith, M.C., Fielding, D.B., Bryan, G.L., Kim, C.-G., Ostriker, E.C., Somerville, R.S., Stern, J., Su, K.-Y., Weinberger, R., Hu, C.-Y., Forbes, J.C., Hernquist, L., Burkhardt, B., Li, Y.: ARKENSTONE - I. A novel method for robustly capturing high specific energy outflows in cosmological simulations. *MNRAS* **527**(1), 1216–1243 (2024) <https://doi.org/10.1093/mnras/stad3168> arXiv:2301.07116 [astro-ph.GA]
- Sike, B., Thomas, T., Ruszkowski, M., Pfrommer, C., Weber, M.: Cosmic-Ray-driven Galactic Winds with Resolved Interstellar Medium and Ion-neutral Damping. *ApJ* **987**(2), 204 (2025) <https://doi.org/10.3847/1538-4357/adda3d> arXiv:2410.06988 [astro-ph.GA]
- Thomas, T., Pfrommer, C., Pakmor, R.: Why are thermally and cosmic ray-driven galactic winds fundamentally different? *A&A* **698**, 104 (2025) <https://doi.org/10.1051/0004-6361/202450817> arXiv:2405.13121 [astro-ph.GA]
- Vijayan, A., Krumholz, M.R., Wibking, B.D.: QUOKKA-based understanding of outflows (QED) - I. Metal loading, phase structure, and convergence testing for solar neighbourhood conditions. *Monthly Notices of the Royal Astronomical Society* **527**, 10095–10110 (2024) <https://doi.org/10.1093/mnras/stad3816>
- Vijayan, A., Krumholz, M.R., Wibking, B.D.: QUOKKA-based understanding of outflows (QED) – III. Outflow loading and phase structure as a function of galactic environment. *Monthly Notices of the Royal Astronomical Society* **539**, 1706–1722 (2025) <https://doi.org/10.1093/mnras/staf136>
- Vijayan, A., Krumholz, M.R., Wibking, B.D.: QUOKKA-based understanding of outflows (QED) – V. Variations in metal loading of galactic winds with element nucleosynthetic origin. *MNRAS* **549**(2), 913 (2026) <https://doi.org/10.1093/mnras/stag913> arXiv:2601.08063 [astro-ph.GA]
- Vázquez, G.A., Leitherer, C.: Optimization of Starburst99 for Intermediate-Age and Old Stellar Populations. *ApJ* **621**, 695–717 (2005) <https://doi.org/10.1086/427866>
- Veilleux, S., Maiolino, R., Bolatto, A.D., Aalto, S.: Cool outflows in galaxies and their implications. *Astron. Astrophys. Rev.* **28**(1), 2 (2020) <https://doi.org/10.1007/s00159-019-0121-9> arXiv:2002.07765 [astro-ph.GA]
- Walch, S., Girichidis, P., Naab, T., Gatto, A., Glover, S.C.O., Wünsch, R., Klessen, R.S., Clark, P.C., Peters, T., Derigs, D., Baczynski, C.: The SILCC (SIMulating the LifeCycle of molecular Clouds) project - I. Chemical evolution of the supernova-driven ISM. *MNRAS* **454**(1), 238–268 (2015) <https://doi.org/10.1093/mnras/stv1975> arXiv:1412.2749 [astro-ph.GA]
- Wibking, B.D., Krumholz, M.R.: QUOKKA: A code for two-moment AMR radiation hydrodynamics on GPUs. *Monthly Notices of the Royal Astronomical Society* **512**, 1430–1449 (2022) <https://doi.org/10.1093/mnras/stac439>
- Yuan, Y., Krumholz, M.R., Martin, C.L.: The observable properties of cool winds from galaxies, AGN, and star clusters - II. 3D models for the multiphase wind of M82. *MNRAS* **518**(3), 4084–4105 (2023) <https://doi.org/10.1093/mnras/stac3241> arXiv:2204.05006 [astro-ph.GA]